



UniSAR: Modeling User Transition Behaviors between Search and Recommendation

Teng Shi

Zihua Si

Jun Xu*

Xiao Zhang

Gaoling School of Artificial Intelligence

Renmin University of China

Beijing, China

{shiteng,zihua_si,junxu,zhangx89}@ruc.edu.cn

Xiaoxue Zang

Kai Zheng

Dewei Leng

Yanan Niu

Yang Song

Kuaishou Technology Co., Ltd.

Beijing, China

{zangxiaoxue,zhengkai,lengdewei}@kuaishou.com

{niuyanan,yangsong}@kuaishou.com

ABSTRACT

Nowadays, many platforms provide users with both search and recommendation services as important tools for accessing information. The phenomenon has led to a correlation between user search and recommendation behaviors, providing an opportunity to model user interests in a fine-grained way. Existing approaches either model user search and recommendation behaviors separately or overlook the different transitions between user search and recommendation behaviors. In this paper, we propose a framework named **UniSAR** that effectively models the different types of fine-grained behavior transitions for providing users a **Unified Search And Recommendation** service. Specifically, UniSAR models the user transition behaviors between search and recommendation through three steps: extraction, alignment, and fusion, which are respectively implemented by transformers equipped with pre-defined masks, contrastive learning that aligns the extracted fine-grained user transitions, and cross-attentions that fuse different transitions. To provide users with a unified service, the learned representations are fed into the downstream search and recommendation models. Joint learning on both search and recommendation data is employed to utilize the knowledge and enhance each other. Experimental results on two public datasets demonstrated the effectiveness of UniSAR in terms of enhancing both search and recommendation simultaneously. The experimental analysis further validates that UniSAR enhances the results by successfully modeling the user transition behaviors between search and recommendation.

CCS CONCEPTS

• **Information systems** → **Recommender systems.**

*Corresponding author. Work partially done at Engineering Research Center of Next-Generation Intelligent Search and Recommendation, Ministry of Education. Work done when Teng Shi and Zihua Si were interns at Kuaishou.

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

SIGIR '24, July 14–18, 2024, Washington, DC, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0431-4/24/07

<https://doi.org/10.1145/3626772.3657811>

KEYWORDS

Recommendation; Search; Contrastive Learning

ACM Reference Format:

Teng Shi, Zihua Si, Jun Xu, Xiao Zhang, Xiaoxue Zang, Kai Zheng, Dewei Leng, Yanan Niu, and Yang Song. 2024. UniSAR: Modeling User Transition Behaviors between Search and Recommendation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*, July 14–18, 2024, Washington, DC, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3626772.3657811>

1 INTRODUCTION

Search engines and recommender systems have become important tools for users to access information. Many commercial platforms, such as YouTube, TikTok, and Kwai, have provided both search and recommendation (**S&R**) services in one app, blending the boundaries between users' S&R behaviors. In these apps, users may further search for related items after browsing their preferred items, or they might explore similar items following a search for a specific item. These phenomena indicate that understanding the user behavior transitions between S&R services is highly beneficial for modeling user interests.

Early studies like JSR [41, 42] optimize S&R models by a joint training loss, where both models share item representations. Recently, researchers have paid attention to modeling user behaviors in two scenarios. SESRec [27] and UnifiedSSR [38] proposed to capture user interests from S&R behaviors separately with distinct encoders. Meanwhile, USER [40] proposed to mix the S&R histories into one sequence according to timestamp and leverage a single encoder.

Despite their effectiveness, existing studies have no mechanism to handle the fine-grained user transitions between S&R explicitly. The mixture of search and recommendation scenarios in one app makes user behaviors much more complex than those of a single search engine or recommender system. For example, users may switch to search/recommendation services if the current recommendation/search results cannot satisfy their needs well, leading to a large number of transition behaviors. Specifically, as illustrated in Figure 1, a typical user may exhibit four types of behaviors: (1) staying in the search scenario and issuing new queries (denoted as **s2s**); (2) keeping browsing new items in the recommendation scenario

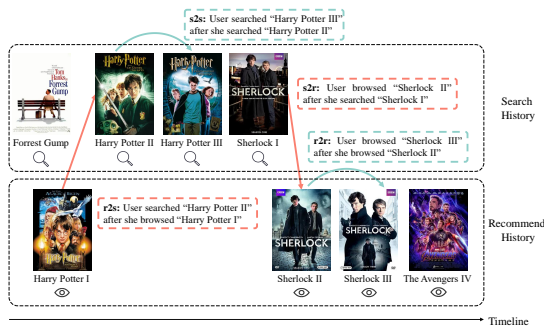


Figure 1: An example of the user search and recommendation history in an App with both search and recommendation services. The user exhibits four types of transition behaviors, as shown by four arrows.

(denoted as $r2r$); (3) searching for certain items after browsing recommended ones, driven by new information needs (transits from recommendation to search, denoted as $r2s$); and (4) browsing the recommended items after searching, as they align well with user interests (transits from search to recommendation, denoted as $s2r$).

We conducted an analysis based on KuaiSAR [29] to show whether the information needs of a user also changed when she transits between search and recommendation. For example, given a randomly sampled set of recommended and clicked items (current item), we calculated the percentage that the immediate preceding clicked item is correlated¹, grouped by the preceding scenario (search or recommendation). From the results of the left two columns of Figure 2, we can see that if the user’s preceding scenario is also the recommendation (no scenario transition), the correlated percentage is 7.99%. However, if the user’s preceding scenario is search (transit from search to recommendation), the correlated percentage drops to 4.86%. Similarly, given a clicked item in search, we also calculated the aforementioned two percentages, shown in the right two columns of Figure 2. We can see that if the user preceding scenario is the same (i.e., search), the correlated percentage is 17.14%. The number drops to 3.67% if the preceding scenario is different (i.e., recommendation). From the analysis, we conclude that it has much more possibility that a user has some new information needs if she transits between S&R. The phenomenon motivates the necessity of fine-grained modeling of user transition behaviors.

This paper proposes an approach named **UniSAR** which explicitly models different types of fine-grained user transition behaviors for both S&R. UniSAR primarily models user transitions through three components: *extraction*, *alignment*, and *fusion*. (1) In the extraction, the chronologically sorted S&R history is fed into a transformer equipped with the mask mechanism. The mask mechanism enables attention calculations between different behaviors, aiding in the extraction of $r2s$ and $s2r$ transitions. Additionally, the S&R histories are separately inputted into two other transformers to extract $s2s$ and $r2r$ transitions. (2) The alignment aims to align transitions from the same scenarios with those from different scenarios (more likely indicating new information needs) to enable the model to learn relationships between different transitions.

¹Two items are considered correlated if they are within the same category.




Figure 2: The percentage of the immediately preceding clicked item correlates with the current item.

Specifically, contrastive learning is employed to align $r2s$ with $s2s$ and $s2r$ with $r2r$, enabling the model to comprehend relationships between them. (3) The fusion utilizes the cross-attention to fuse $r2s$ and $s2s$, resulting in the representation of search. At the same time, another cross-attention is used to fuse $s2r$ and $r2r$, achieving the representation of recommendation. The two representations constitute the overall user representation.

These representations, along with candidate items and queries, can be used to predict user preferences, as have been used in traditional search and recommendation models. As an example, we jointly train UniSAR on S&R data, enabling its application in both scenarios. MMoE [21, 31] is employed to alleviate the seesaw phenomenon for multi-task training.

The major contributions of the paper are summarized as follows:

- UniSAR proposes three components of extraction, alignment, and fusion, to effectively capture the fine-grained transitions between users’ S&R behaviors. We leverage different transformers to extract four transitions and contrastive learning techniques to align them. Cross-attention mechanisms are further used to fuse different transitions.
- UniSAR is trained jointly on search and recommendation tasks and can be applied to both scenarios, effectively utilizing knowledge from each to enhance the other.
- Experiment results on two public datasets demonstrate the effectiveness of UniSAR. UniSAR has outperformed not just the traditional models for a single scenario but also surpassed the existing joint search and recommendation models, achieving state-of-the-art (SOTA) performance.

2 RELATED WORK

Sequential Recommendation. Sequential recommendation aims to model user historical behaviors to capture users’ dynamic preferences. Early works [24] adopt the Markov Chain to learn item transition probability. With the advancement of deep learning, various works have been proposed to model user historical behaviors using neural networks, including methods based on Recurrent Neural Network (RNN) [15, 18], Convolutional Neural Network (CNN) [32], Transformer [16, 28], Multilayer Perceptron (MLP) [47], and Graph Neural Networks (GNN) [5, 37], etc. In addition to employing better neural network architectures, some studies have introduced additional self-supervised signals through contrastive learning [7, 12] to help the training of recommendation models [8, 23, 39]. Unlike these methods, we leverage search history to enhance S&R tasks.

Personalized Search. The goal of personalized search is to provide relevant items based on the user’s searched queries. Early

works [1, 3] only considered the similarity between queries and items, such as QEM [1] and DREM [3]. To better understand users' search intent, the personalized search begins to incorporate user profiles and their search histories [1, 2, 4, 9, 20]. HEM [2] introduces user embeddings. It learns representations of users and items based on reviews through a generative language model. AEM [1] further incorporates user search histories. It employs an attention mechanism to aggregate the user history. ZAM [1] builds upon AEM by incorporating a zero vector to control the degree of personalization. TEM [4] replaces the attention mechanism in AEM with a transformer. Unlike existing approaches, we introduce recommendation history to perform both S&R tasks.

Joint Search and Recommendation. In recent years, there has been a trend toward integrating S&R. These works primarily fall into two categories: (a) Search enhanced recommendation [14, 25, 27, 30, 36]. This type of work utilizes search data as supplementary information to enhance the recommendation performance. IV4Rec [25, 26] utilizes causal learning, treating user-searched queries as instrumental variables to reconstruct user and item embeddings. Query-SeqRec [14] is a query-aware model which incorporates user queries to model users' intent. SESRec [27] uses contrastive learning to disentangle similar and dissimilar interests between user S&R behaviors. (b) Unified S&R [11, 38, 40, 41, 43]. This kind of work performs joint learning of S&R to enhance the model performance in both scenarios. JSR [41, 42] simultaneously trains two models for S&R using a joint loss function. USER [40] integrates user S&R behaviors and feeds them into a transformer encoder. UnifiedSSR [38] proposes a dual-branch network to encode the product history and query history in parallel. In this paper, we develop a framework that utilizes a unified set of parameters to handle both S&R tasks.

3 PROBLEM FORMULATION

Let $\mathcal{U}, \mathcal{I}, \mathcal{Q}$ denote the sets of users, items, and queries, respectively. Each user $u \in \mathcal{U}$ has a chronologically ordered interaction history S_u that encapsulates all her past S&R behaviors. $S_u = \{(x_1, b_1), (x_2, b_2), \dots, (x_N, b_N)\}$, where N stands for the number of behaviors exhibited by user u . The variable $b_t \in \{1, 0\}$ signifies the type of the t -th behavior, with 1 indicating a recommendation and 0 indicating a search. Here, x_t denotes the t -th behavior:

$$x_t = \begin{cases} i_t, & \text{if } b_t = 1 \text{ (recommendation),} \\ \langle q_t, C_{q_t} \rangle, & \text{if } b_t = 0 \text{ (search),} \end{cases}$$

where $i_t \in \mathcal{I}$ represents the t -th interacted item, and $q_t \in \mathcal{Q}$ is the t -th searched query. $C_{q_t} = \{i_1, i_2, \dots, i_{N_{q_t}}\}$ are the N_{q_t} items that u clicked after issuing query q_t . Please note that C_{q_t} could be an empty set, indicating that u may not click any items ($N_{q_t} = 0$) after issuing q_t .

The unified S&R interactions, denoted as $\mathcal{D} = \mathcal{D}_R \cup \mathcal{D}_S$, contain the recommendation dataset $\mathcal{D}_R = \{(u, i_{N+1}, q, S_u, y_{u, i_{N+1}, q}^R)_k\}_{k=1}^{|\mathcal{D}_R|}$ and the search dataset $\mathcal{D}_S = \{(u, i_{N+1}, q, S_u, y_{u, i_{N+1}, q}^S)_k\}_{k=1}^{|\mathcal{D}_S|}$. Here, $y_{u, i_{N+1}, q}^R$ and $y_{u, i_{N+1}, q}^S$ are user u 's preference score for the next item $i_{N+1} \in \mathcal{I}$ after issuing query $q \in \mathcal{Q}$ in recommendation and search scenarios, respectively. It's worth noting that q is empty for the recommendation scenario. The objective is to train on dataset \mathcal{D}

to find the optimal function f_Θ that predicts $y_{u, i_{N+1}, q}^R$ and $y_{u, i_{N+1}, q}^S$ based on the history S_u . The function can be formulated as follows:

$$\begin{bmatrix} \hat{y}_{u, i_{N+1}, q}^R \\ \hat{y}_{u, i_{N+1}, q}^S \end{bmatrix} = f_\Theta(u, i_{N+1}, q, S_u),$$

where f_Θ denotes our unified model with parameters Θ . It's important to note that f_Θ have two outputs: $\hat{y}_{u, i_{N+1}, q}^R$ for recommendation and $\hat{y}_{u, i_{N+1}, q}^S$ for search.

4 OUR APPROACH

This section introduces the proposed method UniSAR. UniSAR's overall architecture is shown in Figure 3. UniSAR first converts the input data into dense representations. Then, it employs three distinct transformers using the mask mechanism to separately extract different transitions between user S&R behaviors. Subsequently, contrastive learning is used to align different transitions to learn the relationships between them. Cross-attentions are also used to fuse these different transitions to obtain the final representations of S&R histories. Following this, an attention mechanism is used to aggregate the historical representations. Finally, UniSAR is jointly trained on S&R data, enabling its application in both scenarios. MMoE is employed to address the seesaw phenomenon for multi-task training.

4.1 Embedding Module

4.1.1 Embedding Layer. We have three embedding tables: $\mathbf{M}_U \in \mathbb{R}^{|\mathcal{U}| \times d}$, $\mathbf{M}_I \in \mathbb{R}^{|\mathcal{I}| \times d}$ and $\mathbf{M}_W \in \mathbb{R}^{|\mathcal{W}| \times d}$ for users, items, and the words of all queries. Here, \mathcal{W} is the set composed of all the words included in the queries. d is the embedding dimension. Given a user and an item, through the lookup operation, we can get their embeddings: $\mathbf{e}_u \in \mathbb{R}^d$ and $\mathbf{e}_i \in \mathbb{R}^d$. For each query q , it contains several words $\{w_1, w_2, \dots, w_{|q|}\} \subseteq \mathcal{W}$. Due to the infrequency of repeated queries in the search data and the typically brief nature of each query, lacking discernible sequential patterns as noted in [6, 38], we derive the query embedding directly by performing mean pooling over the word embeddings: $\mathbf{e}_q = \text{Mean}(\mathbf{e}_{w_1}, \mathbf{e}_{w_2}, \dots, \mathbf{e}_{w_{|q|}}) \in \mathbb{R}^d$, where $\mathbf{e}_{w_t} \in \mathbb{R}^d$ refers to the ID embedding of the t -th word.

Regarding the embeddings of behaviors within the user history S_u , we directly use the item embeddings as the embeddings for recommendation behaviors. For search behaviors, we combine the query embeddings with the mean pooling of clicked item embeddings to form the embeddings. The process is detailed below:

$$\mathbf{e}_{x_t} = \begin{cases} \mathbf{e}_{i_t}, & \text{if } b_t = 1 \text{ (recommendation),} \\ \mathbf{e}_{q_t} + \text{Mean}(C_{q_t}), & \text{if } b_t = 0 \text{ (search),} \end{cases} \quad (1)$$

where $\text{Mean}(C_{q_t}) = \text{Mean}(\mathbf{e}_{i_1}, \mathbf{e}_{i_2}, \dots, \mathbf{e}_{i_{N_{q_t}}})$. Following this, we obtain the representation $\mathbf{E}_u = [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_N]^\top \in \mathbb{R}^{N \times d}$ for the user history S_u .

4.1.2 Query-Item Contrastive Learning. Search and recommendation behaviors exhibit distinct characteristics. Most notably, search behavior includes queries, while recommendation does not. Merging these behaviors directly poses a challenge for the model to capture the transitions between S&R. Hence, we employ contrastive learning to align the representations of queries and items into the same space. This enables the model to capture the relationship

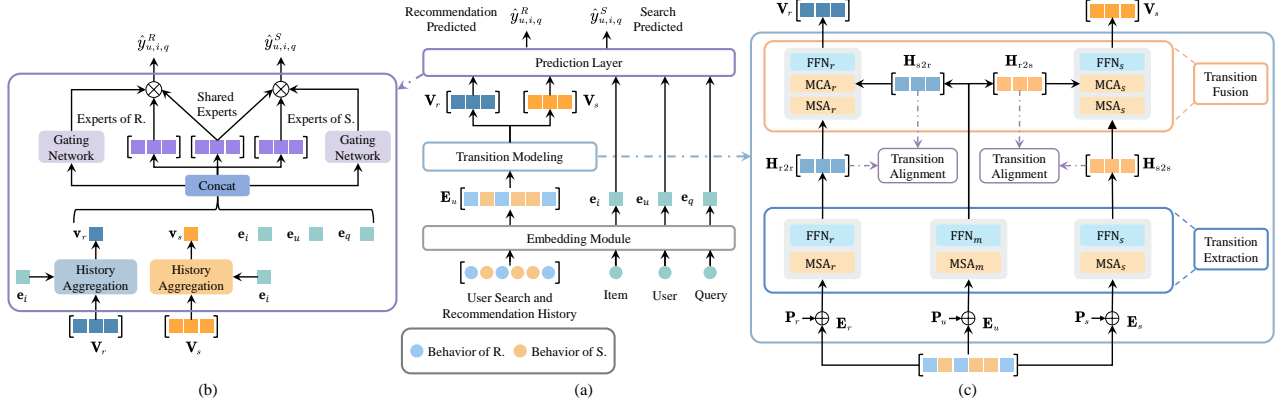


Figure 3: The architecture of UniSAR. (a) Model workflow; (b) Prediction layer; (c) User transition modeling.

between S&R behaviors better. Moreover, this approach helps the model comprehend the relevance between queries and items, which is crucial for search performance.

For the $\langle q, C_q \rangle$ pairs specifically, we treat the query q and items in C_q as positive samples, while randomly sampling other queries and items as negative samples. The contrastive loss is:

$$\mathcal{L}_{\text{Rel}} = - \left[\sum_{i \in C_q} \log \frac{\exp(\text{sim}(\mathbf{e}_q, \mathbf{e}_i) / \tau_1)}{\sum_{i^- \in \mathcal{I}_{\text{neg}}} \exp(\text{sim}(\mathbf{e}_q, \mathbf{e}_{i^-}) / \tau_1)} + \sum_{i \in C_q} \log \frac{\exp(\text{sim}(\mathbf{e}_q, \mathbf{e}_i) / \tau_1)}{\sum_{q^- \in \mathcal{Q}_{\text{neg}}} \exp(\text{sim}(\mathbf{e}_{q^-}, \mathbf{e}_i) / \tau_1)} \right], \quad (2)$$

where τ_1 is a learnable temperature coefficient, \mathcal{I}_{neg} and \mathcal{Q}_{neg} are the set of randomly sampled negative items and queries, respectively. Here, $\text{sim}(\mathbf{a}, \mathbf{b}) = \tanh(\mathbf{a}\mathbf{W}\mathbf{b}^T)$ computes the similarity between $\mathbf{a} \in \mathbb{R}^d$ and $\mathbf{b} \in \mathbb{R}^d$. The \tanh is the activate function and $\mathbf{W} \in \mathbb{R}^{d \times d}$ maps \mathbf{a} and \mathbf{b} to the representation space where the contrastive loss is calculated.

4.2 Transition Modeling

In this section, we introduce the transition modeling module in UniSAR, which explicitly models different types of fine-grained user transition behaviors between S&R. It consists of three components: (1) The *extraction* utilizes transformers equipped with mask mechanism to extract different transitions; (2) The *alignment* employs contrastive learning to align different transitions to learn the relationships between them; (3) The *fusion* uses cross-attention to fuse different transitions to get the overall user representations.

4.2.1 Transition Extraction. Firstly, in order to model the transitions of **s2s** and **r2r**, we extract the S&R sub-histories from S_u , denoted as $S_s = \{(x_1, x_2, \dots, x_{N_s})\}$ and $S_r = \{(x_1, x_2, \dots, x_{N_r})\}$. These sub-histories, S_s or S_r , contain all chronologically ordered search or recommendation behaviors within S_u , respectively. Here, N_s and N_r represent the number of search and recommendation behaviors within the user history S_u , and they satisfy the condition $N_s + N_r = N$. Then, according to Eq. (1), we derive their corresponding representations: $\mathbf{E}_s = [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_{N_s}]^T \in \mathbb{R}^{N_s \times d}$ and $\mathbf{E}_r = [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_{N_r}]^T \in \mathbb{R}^{N_r \times d}$.

To capture the sequential relationships, we further introduce positional embeddings: $\mathbf{P}_s \in \mathbb{R}^{N_s \times d}$, $\mathbf{P}_r \in \mathbb{R}^{N_r \times d}$, and $\mathbf{P}_u \in \mathbb{R}^{N \times d}$ for the sequences S_s , S_r and S_u respectively. The final representations of the three sequences are computed as follows:

$$\hat{\mathbf{E}}_s = \mathbf{E}_s + \mathbf{P}_s, \quad \hat{\mathbf{E}}_r = \mathbf{E}_r + \mathbf{P}_r, \quad \hat{\mathbf{E}}_u = \mathbf{E}_u + \mathbf{P}_u.$$

For the transitions of **s2s** and **r2r**, they respectively model the transition within the S&R sub-histories S_s and S_r . So we feed representations $\hat{\mathbf{E}}_s$ and $\hat{\mathbf{E}}_r$ of S_s and S_r into two transformer [34] encoders for their modeling. Each encoder consists of a Multi-head Self-Attention (MSA) layer and a Feed-Forward (FFN). For **s2s**, MSA takes $\hat{\mathbf{E}}_s$ as the input, where $\mathbf{Q} = \mathbf{K} = \mathbf{V} = \hat{\mathbf{E}}_s$. For **r2r**, MSA takes $\hat{\mathbf{E}}_r$ as its input. The details are as follows:

$$\mathbf{H}_{s2s} = \text{FFN}_s(\text{MSA}_s(\hat{\mathbf{E}}_s, \hat{\mathbf{E}}_s, \hat{\mathbf{E}}_s)), \quad \mathbf{H}_{r2r} = \text{FFN}_r(\text{MSA}_r(\hat{\mathbf{E}}_r, \hat{\mathbf{E}}_r, \hat{\mathbf{E}}_r)), \quad (3)$$

where $\mathbf{H}_{s2s} \in \mathbb{R}^{N_s \times d}$ and $\mathbf{H}_{r2r} \in \mathbb{R}^{N_r \times d}$ capture the transitions for **s2s** and **r2r**, respectively.

For the transitions of **r2s** and **s2r**, they capture the cross-behavior transitions between search and recommendation. Directly inputting the mixed sequence S_u into a transformer makes it challenging to directly model these transitions because its output would not only include **r2s** and **s2r** but also **s2s** and **r2r**, causing redundancy with \mathbf{H}_{s2s} and \mathbf{H}_{r2r} . To disentangle different transitions and enable the model to better capture different transitions separately, we introduce a mask matrix $\mathbf{M} \in \{0, 1\}^{N \times N}$. The values of \mathbf{M} are as follows:

$$\mathbf{M}_{ij} = \begin{cases} 0, & \text{if } b_i = b_j, \\ 1, & \text{if } b_i \neq b_j. \end{cases}$$

We feed the representations $\hat{\mathbf{E}}_u$ of S_u and \mathbf{M} into a single transformer encoder. Within this encoder, the MSA takes $\hat{\mathbf{E}}_u$ and \mathbf{M} as inputs, where $\mathbf{Q} = \mathbf{K} = \mathbf{V} = \hat{\mathbf{E}}_u$:

$$\mathbf{H}_m = \text{FFN}_m(\text{MSA}_m(\hat{\mathbf{E}}_u, \hat{\mathbf{E}}_u, \hat{\mathbf{E}}_u, \mathbf{M})). \quad (4)$$

With the incorporation of the mask matrix \mathbf{M} in the input of MSA_m , the computation of the attention in MSA_m is as follows:

$$\text{Attention}(\hat{\mathbf{E}}_u, \hat{\mathbf{E}}_u, \hat{\mathbf{E}}_u, \mathbf{M}) = \text{Softmax}(\hat{\mathbf{E}}_u \hat{\mathbf{E}}_u^T / \sqrt{d/h} \odot \mathbf{M}) \hat{\mathbf{E}}_u,$$

where \odot denotes the Hadamard product, and h denotes the number of heads in the MSA_m . Figure 4 further demonstrates the varied attention computation methods in different MSA after incorporating

masks. In Eq. (4), $\mathbf{H}_m \in \mathbb{R}^{N \times d}$ captures transitions specific to $\mathbf{r}2\mathbf{s}$ and $\mathbf{s}2\mathbf{r}$, respectively. The mask matrix \mathbf{M} ensures that attention is computed exclusively among different behaviors within MSA_m , thereby effectively capturing the transitions for $\mathbf{r}2\mathbf{s}$ and $\mathbf{s}2\mathbf{r}$. Finally, we extract the search and recommendation sub-sequences from \mathbf{H}_m , yielding $\mathbf{H}_{r2s} \in \mathbb{R}^{N_s \times d}$ and $\mathbf{H}_{s2r} \in \mathbb{R}^{N_r \times d}$. These matrices encapsulate the transitions for $\mathbf{r}2\mathbf{s}$ and $\mathbf{s}2\mathbf{r}$.

4.2.2 Transition Alignment. Based on the representations of these four different types of transitions extracted, we further align the varying representations of the same behavior types among users. For example, \mathbf{H}_{s2s} and \mathbf{H}_{r2s} represent the transitions from search and recommendation behaviors in the context of search history, respectively. Considering that both of them contain information related to the search history, we guide the model to bring their semantics closer. We utilize contrastive learning to align transitions from the same scenarios with those from different scenarios, enabling the model to learn the correlations between them. This can facilitate a better fusion of various transitions, which is discussed in Section 4.2.3.

Specifically, for \mathbf{H}_{s2s} and \mathbf{H}_{r2s} , they respectively involve transitions from the same and different scenarios in terms of search history. We treat them as positive samples for each other and align them. Initially, we perform mean pooling on \mathbf{H}_{s2s} and \mathbf{H}_{r2s} to obtain $\mathbf{h}_{s2s} \in \mathbb{R}^d$ and $\mathbf{h}_{r2s} \in \mathbb{R}^d$, i.e., $\mathbf{h}_{s2s} = \text{Mean}(\mathbf{H}_{s2s})$ and $\mathbf{h}_{r2s} = \text{Mean}(\mathbf{H}_{r2s})$. Then, we treat \mathbf{h}_{s2s} and \mathbf{h}_{r2s} as positive samples, while \mathbf{h}_{s2s}^- and \mathbf{h}_{r2s}^- from other users' histories within the same batch as negative samples. Then, the contrastive loss is:

$$\mathcal{L}_{\text{Align}}^S = - \left[\log \frac{\exp(\text{sim}(\mathbf{h}_{s2s}, \mathbf{h}_{r2s})/\tau_2)}{\sum_{\mathbf{h}_{r2s}^- \in \mathcal{H}_{\text{neg}}^{r2s}} \exp(\text{sim}(\mathbf{h}_{s2s}, \mathbf{h}_{r2s}^-)/\tau_2)} + \log \frac{\exp(\text{sim}(\mathbf{h}_{s2s}, \mathbf{h}_{r2s})/\tau_2)}{\sum_{\mathbf{h}_{s2s}^- \in \mathcal{H}_{\text{neg}}^{s2s}} \exp(\text{sim}(\mathbf{h}_{s2s}^-, \mathbf{h}_{r2s})/\tau_2)} \right], \quad (5)$$

where τ_2 is a learnable temperature coefficient, $\mathcal{H}_{\text{neg}}^{r2s}$ and $\mathcal{H}_{\text{neg}}^{s2s}$ are the set of in-batch negative samples, and the function $\text{sim}(\mathbf{a}, \mathbf{b})$ is defined the same as in Eq. (2). Similarly, we obtain the contrastive loss $\mathcal{L}_{\text{Align}}^R$ for \mathbf{H}_{r2r} and \mathbf{H}_{s2r} . The total contrastive loss for transition alignment is:

$$\mathcal{L}_{\text{Align}} = \mathcal{L}_{\text{Align}}^S + \mathcal{L}_{\text{Align}}^R. \quad (6)$$

4.2.3 Transition Fusion. After extracting and learning the relationships between different transitions, we fuse them to get the overall user representation. Specifically, we fuse \mathbf{H}_{s2s} with \mathbf{H}_{r2s} as well as \mathbf{H}_{r2r} with \mathbf{H}_{s2r} to derive the final representations for S&R histories. Drawing inspiration from prior work on multi-modal fusion [10, 19], we utilize Multi-head Cross-Attention (MCA) for this fusion. For search history, MCA takes \mathbf{H}_{s2s} with \mathbf{H}_{r2s} as the inputs for fusion, where $\mathbf{Q} = \mathbf{H}_{r2s}$, $\mathbf{K} = \mathbf{V} = \mathbf{H}_{s2s}$. While, for recommendation history, MCA takes \mathbf{H}_{r2r} with \mathbf{H}_{s2r} as the inputs. The specific process is:

$$\begin{aligned} \mathbf{F}_s &= \text{MSA}_s(\mathbf{H}_{s2s}, \mathbf{H}_{s2s}, \mathbf{H}_{s2s}), \quad \mathbf{V}_s = \text{FFN}_s(\text{MCA}_s(\mathbf{H}_{r2s}, \mathbf{F}_s, \mathbf{F}_s)), \\ \mathbf{F}_r &= \text{MSA}_r(\mathbf{H}_{r2r}, \mathbf{H}_{r2r}, \mathbf{H}_{r2r}), \quad \mathbf{V}_r = \text{FFN}_r(\text{MCA}_r(\mathbf{H}_{s2r}, \mathbf{F}_r, \mathbf{F}_r)), \end{aligned} \quad (7)$$

where $\mathbf{V}_s \in \mathbb{R}^{N_s \times d}$ and $\mathbf{V}_r \in \mathbb{R}^{N_r \times d}$ are the final representations for S&R histories, respectively. The two representations constitute the overall user representation.




Figure 4: The attention computation methods differ across various MSA (Multi-head Self-Attention) modules within UniSAR. In MSA_r or MSA_s , attention is specifically computed within recommendation or search behaviors to extract $\mathbf{r}2\mathbf{r}$ or $\mathbf{s}2\mathbf{s}$, respectively, as detailed in Eq. (3); MSA_m computes attention exclusively between different behaviors, outlined in Eq. (4), thereby extracting $\mathbf{r}2\mathbf{s}$ and $\mathbf{s}2\mathbf{r}$.

4.3 Model Prediction and Training

4.3.1 History Aggregation. After obtaining representations \mathbf{V}_s and \mathbf{V}_r in Eq. (7) for S&R histories, when predicting whether the next item i will be clicked, we aggregate the history using attention mechanism to extract user interest from the transitions concerning the target item i . The calculation of the attention is as follows:

$$\mathbf{v}_s = \mathbf{V}_s^T \text{Softmax}(\mathbf{V}_s \mathbf{W}_s \mathbf{e}_i), \quad \mathbf{v}_r = \mathbf{V}_r^T \text{Softmax}(\mathbf{V}_r \mathbf{W}_r \mathbf{e}_i) \quad (8)$$

where $\mathbf{W}_s \in \mathbb{R}^{d \times d}$ and $\mathbf{W}_r \in \mathbb{R}^{d \times d}$ are learnable parameters, $\mathbf{v}_s \in \mathbb{R}^d$ and $\mathbf{v}_r \in \mathbb{R}^d$ are the aggregated representations for S&R histories. The attention mechanism in Eq. (8) computes the similarities between the target item and historical behaviors; the higher the similarities, the greater the influence of historical behaviors on the transition to the target item.

4.3.2 Prediction. UniSAR adopts a unified model for S&R tasks. Building on the previous modules' modeling of user transition behaviors, we can predict two tasks separately. In the field of information retrieval, a multi-task framework [21, 31, 35] is commonly used to predict different tasks. Considering that direct making predictions faces challenges due to the differences in data distributions between S&R, we employ a multi-task approach to separately predict the two tasks, incorporating behavior-shared and behavior-specific modules.

UniSAR employs a single model to predict both S&R tasks. We introduce a learnable query $q_\phi \in \mathbb{R}^d$ for the recommendation scenario to unify the input for predictions between S&R. Firstly, we concatenate representations of users, items, queries, and the S&R histories to obtain a shared bottom representation, denoted as $\mathbf{X}_b = \text{Concat}(\mathbf{e}_u, \mathbf{e}_i, \mathbf{e}_q, \mathbf{v}_s, \mathbf{v}_r)$, where $\mathbf{X}_b \in \mathbb{R}^{d_b}$. Here $d_b = 5 \times d$. Next, to balance the relationship between different tasks, we exploit the Multi-gated Mixture of Experts (MMoE) [21, 31] technique. We denote the shared and specific expert sets for S&R as \mathcal{O}_m , \mathcal{O}_s , and \mathcal{O}_r , where the number of experts for each is represented by n_m , n_s , and n_r , respectively. The output of the MMoE module consists of two scores, corresponding to S&R, calculated as follows:

$$\begin{aligned} \hat{y}_{u,i,q}^S &= \sum_{o \in \mathcal{O}_s} g_o^s(\mathbf{X}_b) o(\mathbf{X}_b) + \sum_{o \in \mathcal{O}_m} g_o^s(\mathbf{X}_b) o(\mathbf{X}_b), \\ \hat{y}_{u,i,q}^R &= \sum_{o \in \mathcal{O}_r} g_o^r(\mathbf{X}_b) o(\mathbf{X}_b) + \sum_{o \in \mathcal{O}_m} g_o^r(\mathbf{X}_b) o(\mathbf{X}_b), \end{aligned} \quad (9)$$

where $g^s(\mathbf{X}_b) = \text{Softmax}(\mathbf{W}^s \mathbf{X}_b)$ and $g^r(\mathbf{X}_b) = \text{Softmax}(\mathbf{W}^r \mathbf{X}_b)$ are the gating networks for S&R, $\mathbf{W}^s \in \mathbb{R}^{(n_s+n_m) \times d_b}$ and $\mathbf{W}^r \in \mathbb{R}^{(n_r+n_m) \times d_b}$

Table 1: Comparison of different joint S&R methods. "Explicit" denotes whether the model explicitly considers the differences between different transitions. "One Model" refers to whether a unified set of parameters is used to simultaneously serve both search and recommendation scenarios.

Methods	Transition Modeling					Task		One Model
	r2r	s2s	r2s	s2r	Explicit	R.	S.	
JSR [41]	✓	✗	✗	✗	✗	✓	✓	✗
SESRec [27]	✓	✓	✗	✗	✗	✓	✗	✓
UnifiedSSR [38]	✓	✓	✗	✗	✗	✓	✓	✗
USER [40]	✓	✓	✓	✓	✗	✓	✓	✗
UniSAR (ours)	✓	✓	✓	✓	✓	✓	✓	✓

$\mathbb{R}^{(n_r+n_m) \times d_b}$ are used for linear transformation. Here $o(\cdot)$ denotes the expert networks, Multi-Layer Perceptrons (MLPs) in our implementation. g_o^s and g_o^r are the weights of expert $o(\cdot)$ generated by the gating network for S&R.

4.3.3 Training. Following previous works [27, 45, 46], we utilize binary cross-entropy loss to optimize our model:

$$\mathcal{L}_{\text{Click}}^S = -\frac{1}{|\mathcal{D}_S|} \sum_{(u,i,q) \in \mathcal{D}_S} y_{u,i,q}^S \log(\hat{y}_{u,i,q}^S) + (1 - y_{u,i,q}^S) \log(1 - \hat{y}_{u,i,q}^S),$$

$$\mathcal{L}_{\text{Click}}^R = -\frac{1}{|\mathcal{D}_R|} \sum_{(u,i,q) \in \mathcal{D}_R} y_{u,i,q}^R \log(\hat{y}_{u,i,q}^R) + (1 - y_{u,i,q}^R) \log(1 - \hat{y}_{u,i,q}^R),$$
(10)

where $\mathcal{L}_{\text{Click}}^S$ and $\mathcal{L}_{\text{Click}}^R$ are the click loss for S&R respectively. The total loss for S&R includes the click loss in Eq. (10) and the contrastive loss in Eq. (2) and Eq. (6):

$$\mathcal{L}_S = \mathcal{L}_{\text{Click}}^S + \alpha \mathcal{L}_{\text{Rel}} + \beta \mathcal{L}_{\text{Align}},$$

$$\mathcal{L}_R = \mathcal{L}_{\text{Click}}^R + \alpha \mathcal{L}_{\text{Rel}} + \beta \mathcal{L}_{\text{Align}},$$
(11)

where α and β are the hyper-parameters for contrastive loss. Finally, the overall loss for joint training of S&R is:

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_R + \gamma \mathcal{L}_S + \lambda \|\Theta\|_2, \quad (12)$$

where γ controls the trade-off between S&R, Θ denotes the parameters of UniSAR, and λ controls the L_2 regularization.

Although S&R tasks are related, they have fundamental distinctions. Training a single model for two tasks is not trivial. Joint training may lead to a seesaw effect, where an improvement in one task's performance may lead to a decline in the other (see (4) in Section 5.3). Our introduction of the MMoE module is crucial for model convergence during training.

4.4 Discussion

Table 1 summarized the ability of UniSAR and existing joint S&R models in terms of modeling user S&R histories in user modeling. JSR [41, 42] neglects the transitions between users' different behaviors. SESRec [27] and UnifiedSSR [38] separately input the S&R histories into two encoders, modeling the transitions of **s2s** and **r2r**. USER [40] mixes the user S&R histories into one sequence and implicitly captures the transitions through a single encoder. In USER, different transitions are modeled in the same manner. UniSAR explicitly models fine-grained user transitions in different ways, which considers the differences between various transitions.

Table 2: Statistics of the datasets used in this paper.

Dataset	#Users	#Items	#Queries	#Action-S	#Action-R
KuaiSAR	25,877	6,890,707	453,667	5,059,169	14,605,716
Amazon	68,223	61,934	4,298	934,664	989,618

Furthermore, JSR employs two sets of parameters that are jointly trained by sharing the item set. SESRec utilizes search as supplementary information to enhance recommendations. USER and UnifiedSSR derive two parameter sets for S&R tasks via pre-training and fine-tuning. In this paper, we leverage MMoE to achieve optimal performance in both S&R tasks using a single unified set of parameters, reducing parameter redundancy. In summary, UniSAR is more efficient and effective.

5 EXPERIMENTS

We conducted experiments on two public datasets to evaluate the performance of UniSAR. The source code is available at <https://github.com/TengShi-RUC/UniSAR>.

5.1 Experimental Setup

5.1.1 Dataset. The experiments were conducted on the following publicly available datasets. The statistics of these datasets are shown in Table 2.

KuaiSAR² [29]: The dataset offers authentic user S&R behaviors. As for pre-processing, we only kept users who exhibited both S&R behaviors and filtered data of the last two days due to their sparse search behaviors. Within the remaining data, following [16, 47], we filtered items and users with fewer than five interaction records. Following [44], we used the data from the last day for testing, the second last day for validation, and the rest for training.

Amazon³ [13, 22]: We also adopted a widely accepted semi-synthetic dataset. Following previous works [1, 2, 27, 41], we generated synthetic search behaviors for this recommendation dataset. We adopt the "Kindle Store" subset of the five-core Amazon dataset that covers data in which all users and items have at least five interactions. We strictly adhered to the established practices, as mentioned in [27], for data construction and the leave-one-out strategy to split the dataset.

5.1.2 Evaluation Metrics. Following previous works [27, 28, 47], we utilize the ranking metrics, *Hit Ratio* (HR) and *Normalized Discounted Cumulative Gain* (NDCG). We report HR at positions {1, 5, 10} and NDCG at positions {5, 10}. Following the common strategy [16, 27, 47], we pair the ground-truth item with 99 randomly sampled negative items that the user has not interacted with.

5.1.3 Baselines. In this work, we compare UniSAR with the following state-of-the-art methods.

Firstly, we compare our methods with the following **sequential recommendation models** that do not utilize search data, including (1) **DIN** [46] uses a local activation unit to adaptively learn the representation of user interests; (2) **GRU4Rec** [15] applies GRUs to model users' interaction history. (3) **SASRec** [16] is a unidirectional Transformer-based model; (4) **BERT4Rec** [28]

²<https://kuaisar.github.io/>

³<http://jmcauley.ucsd.edu/data/amazon/>

Table 3: Comparisons of the overall recommendation performance of different methods on both datasets. The best and the second-best methods are highlighted in bold and underlined fonts, respectively. “*” indicates that improvements over the second-best methods are statistically significant (t -test, p -value < 0.01).

Datasets	Metric	Sequential Recommendation					Joint Search and Recommendation						
		DIN	GRU4Rec	SASRec	BERT4Rec	FMLP-Rec	NRHUB	Query-SeqRec	SESRec	JSR	USER	UnifiedSSR	UniSAR
KuaiSAR	HR@1	0.1629	0.1097	0.1249	0.1061	0.1370	0.1243	0.1166	<u>0.1827</u>	0.1754	0.1489	0.1225	0.1990*
	HR@5	0.4509	0.3764	0.4065	0.3699	0.4292	0.3862	0.3920	<u>0.4956</u>	0.4791	0.4086	0.3981	0.5169*
	HR@10	0.6179	0.5788	0.6007	0.5885	0.6159	0.5610	0.5890	<u>0.6643</u>	0.6453	0.5627	0.5939	0.6792*
	NDCG@5	0.3104	0.2435	0.2671	0.2381	0.2851	0.2572	0.2552	<u>0.3432</u>	0.3315	0.2820	0.2617	0.3632*
	NDCG@10	0.3643	0.3087	0.3298	0.3083	0.3453	0.3136	0.3186	<u>0.3978</u>	0.3853	0.3318	0.3249	0.4158*
Amazon (Kindle Store)	HR@1	0.2159	0.1725	0.2059	0.2481	0.1991	0.1889	0.2186	<u>0.2726</u>	0.2346	0.2361	0.2013	0.3010*
	HR@5	0.5170	0.4949	0.5295	0.5311	0.5356	0.4988	0.5401	<u>0.5623</u>	0.5467	0.5441	0.5196	0.5874*
	HR@10	0.6525	0.6548	0.6772	0.6658	<u>0.6879</u>	0.6503	0.6860	0.6864	0.6779	0.6854	0.6707	0.7020*
	NDCG@5	0.3726	0.3388	0.3747	0.3954	0.3739	0.3487	0.3859	<u>0.4245</u>	0.3970	0.3964	0.3662	0.4513*
	NDCG@10	0.4165	0.3907	0.4225	0.4390	0.4232	0.3977	0.4333	<u>0.4648</u>	0.4396	0.4422	0.4151	0.4885*

uses a cloze objective loss for sequential recommendation by the bidirectional Transformer; and (5) **FMLPRec** [47] is an all-MLP model with learnable filters;

Secondly, we compare our methods with the following **personalized search models** that do not utilize recommendation data, including (6) **QEM** [1] only considers the matching scores between items and queries; (7) **HEM** [2] is a latent vector-based personalized model; (8) **AEM** [1] is an attention-based personalized model that aggregates the user’s historical interacted items with the current query; (9) **ZAM** [1] improves AEM by concatenating a zero vector to the item list. (10) **TEM** [4] upgrades the attention layer in AEM with transformer encoder; and (11) **CoPPS** [9] uses contrastive learning techniques;

Finally, we compare our methods with the following two classes of **joint S&R models**: (a) **search enhanced recommendation**: (12) **NRHUB** [36] is a news recommendation model leveraging heterogeneous user behaviors; (13) **Query-SeqRec** [14] is a query-aware sequential model which incorporates queries into user behaviors using transformers; (14) **SESRec** [27] learns disentangled search representation for recommendation using contrastive learning; (b) **unified S&R**: (15) **JSR** [41] is a general framework which optimizes a joint loss. Following [25], the recommendation model in JSR is implemented using DIN in our experiment; (16) **USER** [40] integrates the user’s behaviors in S&R into a heterogeneous behavior sequence. (17) **UnifiedSSR** [38] jointly learns the user behavior history in both S&R scenarios.

5.1.4 Implementation Details. All hyper-parameters for the baselines are searched according to the settings in the original papers. The embedding dimension d is set to 64 for both the KuaiSAR and Amazon datasets. For both datasets, the maximum length for S&R histories is set to 30. The numbers of shared and specific experts, n_m , n_s , and n_r in Section 4.3.2, are both set to 4. The temperature parameters, τ_1 in Eq. (2) and τ_2 in Eq. (6), are tuned among $[0.1 : +0.1 : 1.0]$. The weight α and β in Eq. (11) are tuned among $\{1e-1, 1e-2, 1e-3, 1e-4, 1e-5\}$. The weight γ in Eq. (12) is tuned among $[0.01, 1.0]$. The batch size is set as 1024. We train all the models with 100 epochs and adopt early-stopping to avoid over-fitting. Adam [17] is used to conduct the optimization. The learning rate is tuned among $\{1e-3, 1e-4, 1e-5\}$, and λ in Eq. (12) is searched from $\{1e-4, 1e-5, 1e-6\}$.

5.2 Experimental Results

Tables 3 and 4 report the results on two datasets of S&R tasks, respectively. From the experimental results, we can observe that:

- Firstly, compared to existing joint S&R models, UniSAR has achieved optimal results on all datasets. Overall, UniSAR outperforms the SOTA models, passing the significance test (p -value < 0.01) in most cases. This demonstrates the effectiveness of modeling various user transitions and introducing MMoE to facilitate joint training.
- Secondly, compared to existing models for sequential recommendation or personalized search, UniSAR has consistently yielded superior results. This demonstrates that the integration of S&R in UniSAR benefits both tasks. Furthermore, we can observe that not all models combining S&R yield better results than typical sequential recommendation or personalized search models. For instance, NRHUB and Query-SeqRec don’t exhibit this trend. This indicates that simply blending S&R data doesn’t suffice. Instead, a fine-grained consideration of transitions between user S&R behaviors is necessary.
- Finally, on KuaiSAR data, the model jointly trained for S&R significantly outperforms the personalized search model that only considers search history. This is because, as indicated by the statistics in Table 2, the KuaiSAR search data is extremely sparse. Joint training allows the utilization of recommendation information to alleviate the sparsity of search data, thus leading to a substantial improvement in search performance.

5.3 Ablation Study

We conducted an ablation study on KuaiSAR datasets to investigate the effectiveness of different modules in UniSAR.

(1) Transition Extraction: UniSAR extracts four types of transitions from S&R behaviors, and we removed each transition individually to verify its role. After removing a transition, we remove its corresponding alignment loss and modify some module inputs to keep the architecture’s consistency. For instance, the variant ‘w/o **r2r**’ removes \mathcal{L}_{Align}^R in Eq. (6) and replaces \mathbf{H}_{r2r} with \mathbf{H}_{s2r} in Eq. (7). The first four rows in Table 5 show that removing any transition leads to decreased performances, verifying that each transition is beneficial for both tasks.

To further validate the effectiveness of our design in extracting two types of cross-behavior transitions (i.e., **r2s** & **s2r**), we removed

Table 4: Comparisons of the overall search performance of different methods on both datasets. The best and the second-best methods are highlighted in bold and underlined fonts, respectively. “*” indicates that improvements over the second-best methods are statistically significant (t -test, p -value < 0.01).

Datasets	Metric	Personalized Search					Joint Search and Recommendation				
		QEM	HEM	AEM	ZAM	TEM	CoPPS	JSR	USER	UnifiedSSR	UniSAR
KuaiSAR	HR@1	0.2944	0.3337	0.2703	0.2815	0.3045	0.3117	0.4543	<u>0.4628</u>	0.4389	0.5282*
	HR@5	0.6020	0.6505	0.5956	0.6117	0.6502	0.6616	0.7162	0.7304	<u>0.7377</u>	0.7476*
	HR@10	0.7182	0.7653	0.7182	0.7344	0.7632	0.7707	0.7961	0.8149	<u>0.8320</u>	0.8369*
	NDCG@5	0.4575	0.5029	0.4415	0.4560	0.4887	0.4977	0.5962	<u>0.6069</u>	0.5991	0.6417*
	NDCG@10	0.4953	0.5400	0.4812	0.4959	0.5254	0.5331	0.6221	<u>0.6342</u>	0.6297	0.6708*
Amazon (Kindle Store)	HR@1	0.2772	0.2497	0.2916	0.2954	0.4090	0.4052	0.3176	<u>0.4123</u>	0.3663	0.5343*
	HR@5	0.7100	0.6778	0.7095	0.7109	<u>0.8185</u>	0.8169	0.7038	0.7631	0.7744	0.8190
	HR@10	0.8186	0.8267	0.8443	0.8468	0.9051	<u>0.9051</u>	0.8225	0.8697	0.8812	0.8977
	NDCG@5	0.5066	0.4736	0.5114	0.5147	<u>0.6303</u>	0.6281	0.5173	0.6000	0.5847	0.6875*
	NDCG@10	0.5422	0.5221	0.5554	0.5590	<u>0.6587</u>	0.6570	0.5563	0.6348	0.6196	0.7132*

Table 5: Ablation study on the KuaiSAR dataset. “w/o” indicates the corresponding module in UniSAR is removed.

Model	Recommendation		Search	
	NDCG@5	NDCG@10	NDCG@5	NDCG@10
UniSAR	0.3632	0.4158	0.6417	0.6708
w/o r2r	0.3258	0.3798	0.6007	0.6319
w/o r2s	0.3606	0.4136	0.6376	0.6690
w/o s2r	0.3512	0.4036	0.6292	0.6584
w/o s2s	0.3399	0.3941	0.6404	0.6706
w/o mask M	0.3503	0.4024	0.6217	0.6507
w/o \mathcal{L}_{Align}	0.3458	0.3956	0.6069	0.6381
w/o \mathcal{L}_{Rel}	0.3472	0.3988	0.6177	0.6465
w/o MCA_r	0.3384	0.3924	0.6094	0.6397
w/o MCA_s	0.3561	0.4091	0.6257	0.6545
w/o MMoE	0.3132	0.3692	0.6143	0.6438
w/o Joint Training	0.3569	0.4105	0.5850	0.6167

the mask matrix \mathbf{M} in Eq. (4). The fifth row’s results in Table 5 confirm the effectiveness of separate and fine-grained modeling cross-behavior transitions and the other two types of transitions.

(2) **Transition Alignment:** We investigated the impact of transition alignment in Eq. (6) by removing the alignment loss \mathcal{L}_{Align} . We observed that removing \mathcal{L}_{Align} would lead to a decrease in the performance of both S&R. This indicates that \mathcal{L}_{Align} can assist the model in learning the correlations between different transitions, leading to an effective fusion of transitions.

Additionally, we investigated the impact of \mathcal{L}_{Rel} in Eq. (2). Removing \mathcal{L}_{Rel} also leads to a decrease in S&R performance. This is because \mathcal{L}_{Rel} helps the model learn the relevance between different queries and items, which forms the basis for the model’s semantic alignment of different transitions.

(3) **Transition Fusion:** We explore the effects of cross-attention in Section 4.2.3, which are used for transition fusion to get the overall user representations. The results are shown in the lower part of Table 5. ‘w/o MCA_r ’ indicates that we replace this module with an element-wise addition of \mathbf{H}_{s2s} with \mathbf{H}_{r2s} . Similar operations are done for ‘w/o MCA_s ’. From the results, we observed that the model’s performance declined on both tasks, demonstrating the rationality of using cross-attention for fusion.

(4) **Multi-task Training:** We explore the impact of MMoE and the joint training of two tasks. The results are presented in the bottom part of Table 5. ‘w/o MMoE’ means that we directly remove the MMoE module, and employ two separate MLPs to predict S&R tasks, respectively. Two tasks share most model parameters except these two MLPs. ‘w/o Joint Training’ denotes the separate training of S&R, resulting in two sets of model parameters dedicated to each task. Since each set of parameters handles a single task, MMoE is no longer necessary and hence removed. The performance of these two variants was significantly lower than UniSAR, indicating that removing MMoE or removing joint training adversely affects the model performance.

Further, by comparing these two variants, we observed a seesaw phenomenon in model performance when training on both tasks. ‘w/o MMoE’ leads to a significant improvement in search and a large decline in recommendation compared to ‘w/o Joint Training’. This means that joint training improves the performance of one task but reduces that of the other. This is because, in KuaiSAR, the volume of interactions in recommendation data is much greater than that in search data. There is a trade-off between two tasks. Incorporating MMoE helps us overcome this issue.

5.4 Experimental Analysis

5.4.1 **Impact of \mathcal{L}_{Align} on the Learned Transitions.** To further explore the impact of the contrastive learning in Section 4.2.2 on the learned representations of different transitions, we computed the cosine similarity between \mathbf{h}_{r2s} and \mathbf{h}_{s2s} , as well as between \mathbf{h}_{s2r} and \mathbf{h}_{r2r} , with and without introducing the contrastive loss. Figure 5 depicts the distributions of the cosine similarities. Without the loss \mathcal{L}_{Align} , the similarity was primarily distributed around zero, indicating the lack of correlation between different transitions. However, the similarity increased after integrating the contrastive loss \mathcal{L}_{Align} . This phenomenon indicates that \mathcal{L}_{Align} enables the model to capture similarities between different transitions.

We additionally employed t-SNE [33] to visualize the representations of different transitions, as shown in Figure 6. Without introducing the contrastive loss, we can observe that different transitions were relatively scattered, with transitions like $\mathbf{s2r}$ and $\mathbf{s2s}$ unable to cluster effectively. However, after incorporating the contrastive loss, various transitions were clustered more densely and compactly.




Figure 5: The histogram of the cosine similarities between the hidden representations, with and without introducing \mathcal{L}_{Align} .

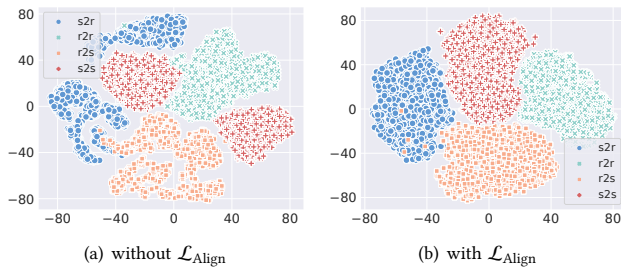


Figure 6: The t-SNE visualization of different transitions' hidden representations.

Specifically, representations of $s2r$ and $r2r$ became closer. Since $s2r$ and $r2s$ originate from the same transformer, their representations tend to be closer. Overall, the inclusion of \mathcal{L}_{Align} facilitates a better understanding and learning of different transitions.

5.4.2 Impact of Hyper-parameters. We analyzed the effects of varying α and β on the S&R performance, as depicted in Figure 7. While adjusting one parameter, the other remained constant at $1e-3$ for α and $1e-1$ for β . Figure 7(a) illustrates that an increase in the weight α of \mathcal{L}_{Rel} corresponds to improved search performance but decreased recommendation performance. This is because \mathcal{L}_{Rel} tends to bring the query closer to its corresponding item, enabling the model to better learn their relevance. Search performance heavily relies on the relevance between queries and items. Regarding recommendation, \mathcal{L}_{Rel} can augment performance to some extent. However, an excessively large α may hinder the optimization of click losses as defined in Eq. (11), consequently leading to a decline in overall performance. Concerning the weight β of \mathcal{L}_{Align} , optimal S&R performance is achieved at a value of $1e-1$. When β is too small, the model cannot effectively learn the similarity of different transitions, failing to efficiently fuse them.

We further investigated the impact of γ in Eq. (12) on the S&R performance, as illustrated in Figure 8. It can be observed that without MMoE for joint training, an increase in γ enhances search performance while diminishing recommendation effectiveness. Additionally, in this scenario, the recommendation performance falls below that of a model trained solely on recommendation data. This signifies the seesaw phenomenon between S&R, highlighting the difficulty of effectively handling both tasks simultaneously through direct joint training without using MMoE. With the introduction




Figure 7: The recommendation and search performance with different values of α and β in Eq. (11)

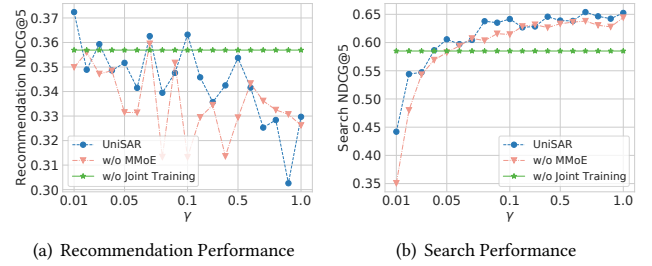


Figure 8: The recommendation and search performance with different values of γ in Eq. (12)

of MMoE, UniSAR manages to achieve better results in most cases compared to models trained separately on S&R data. This further demonstrates the significance of MMoE in learning shared and specific information between S&R, thereby simultaneously enhancing the performance of both tasks.

6 CONCLUSION

In this paper, we propose UniSAR which models fine-grained transitions between user S&R behaviors. UniSAR employs different transformers equipped with the mask mechanism to extract four transitions between S&R and utilizes contrastive learning techniques to align them to learn the relationships between different transitions. Additionally, cross-attentions are employed to fuse different transitions to get the user representations. UniSAR is jointly trained on S&R tasks using MMoE and can be applied to both scenarios, effectively utilizing knowledge from each to enhance the other. Results on two public datasets validate the effectiveness of UniSAR. The experimental analysis further illustrates the importance of each module within UniSAR.

ACKNOWLEDGMENTS

This work was funded by the National Key R&D Program of China (2023YFA1008704), the National Natural Science Foundation of China (No. 62376275, 62377044), Beijing Key Laboratory of Big Data Management and Analysis Methods, Major Innovation & Planning Interdisciplinary Platform for the “Double-First Class” Initiative, PCC@RUC, funds for building world-class universities (disciplines) of Renmin University of China, Kuaishou Technology.

REFERENCES

- [1] Qingyao Ai, Daniel N Hill, SVN Vishwanathan, and W Bruce Croft. 2019. A zero attention model for personalized product search. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 379–388.
- [2] Qingyao Ai, Yongfeng Zhang, Keping Bi, Xu Chen, and W Bruce Croft. 2017. Learning a hierarchical embedding model for personalized product search. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 645–654.
- [3] Qingyao Ai, Yongfeng Zhang, Keping Bi, and W Bruce Croft. 2019. Explainable product search with a dynamic relation embedding model. *ACM Transactions on Information Systems (TOIS)* 38, 1 (2019), 1–29.
- [4] Keping Bi, Qingyao Ai, and W Bruce Croft. 2020. A transformer-based embedding model for personalized product search. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1521–1524.
- [5] Jianxin Chang, Chen Gao, Yu Zheng, Yiqun Hui, Yanan Niu, Yang Song, Depeng Jin, and Yong Li. 2021. Sequential recommendation with graph neural networks. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*. 378–387.
- [6] Jia Chen, Jiabin Mao, Yiqun Liu, Fan Zhang, Min Zhang, and Shaoping Ma. 2021. Towards a better understanding of query reformulation behavior in web search. In *Proceedings of the web conference 2021*. 743–755.
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR, 1597–1607.
- [8] Yongjun Chen, Zhiwei Liu, Jia Li, Julian McAuley, and Caiming Xiong. 2022. Intent contrastive learning for sequential recommendation. In *Proceedings of the ACM Web Conference 2022*. 2172–2182.
- [9] Shitong Dai, Jiongnan Liu, Zhicheng Dou, Haonan Wang, Lin Liu, Bo Long, and Ji-Rong Wen. 2023. Contrastive Learning for User Sequence Representation in Personalized Product Search. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2023, Long Beach, CA, USA, August 6-10, 2023*. ACM, 380–389.
- [10] Zi-Yi Dou, Yichong Xu, Zhe Gan, Jianfeng Wang, Shuohang Wang, Lijuan Wang, Chengguang Zhu, Pengchuan Zhang, Lu Yuan, Nanyun Peng, et al. 2022. An empirical study of training end-to-end vision-and-language transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 18166–18176.
- [11] Yuqi Gong, Xichen Ding, Yehui Su, Kaiping Shen, Zhongyi Liu, and Guannan Zhang. 2023. An Unified Search and Recommendation Foundation Model for Cold-Start Scenario. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 4595–4601.
- [12] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 9729–9738.
- [13] Ruining He and Julian McAuley. 2016. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *proceedings of the 25th international conference on world wide web*. 507–517.
- [14] Zhankui He, Handong Zhao, Zhaowen Wang, Zhe Lin, Ajinkya Kale, and Julian McAuley. 2022. Query-Aware Sequential Recommendation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management (Atlanta, GA, USA) (CIKM '22)*. Association for Computing Machinery, New York, NY, USA, 4019–4023.
- [15] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based Recommendations with Recurrent Neural Networks. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.).
- [16] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 197–206.
- [17] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [18] Jing Li, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Tao Lian, and Jun Ma. 2017. Neural attentive session-based recommendation. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 1419–1428.
- [19] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems* 34 (2021), 9694–9705.
- [20] Jiongnan Liu, Zhicheng Dou, Qiannan Zhu, and Ji-Rong Wen. 2022. A Category-aware Multi-interest Model for Personalized Product Search. In *Proceedings of the ACM Web Conference 2022*. 360–368.
- [21] Jiaqi Ma, Zhe Zhao, Xinyang Yi, Jilin Chen, Lichan Hong, and Ed H Chi. 2018. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 1930–1939.
- [22] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*. 43–52.
- [23] Ruihong Qiu, Zi Huang, Hongzhi Yin, and Zijian Wang. 2022. Contrastive learning for representation degeneration problem in sequential recommendation. In *Proceedings of the fifteenth ACM international conference on web search and data mining*. 813–823.
- [24] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing personalized markov chains for next-basket recommendation. In *Proceedings of the 19th international conference on World wide web*. 811–820.
- [25] Zihua Si, Xueran Han, Xiao Zhang, Jun Xu, Yue Yin, Yang Song, and Ji-Rong Wen. 2022. A Model-Agnostic Causal Learning Framework for Recommendation Using Search Data. In *Proceedings of the ACM Web Conference 2022 (Virtual Event, Lyon, France) (WWW '22)*. Association for Computing Machinery, New York, NY, USA, 224–233.
- [26] Zihua Si, Zhongxiang Sun, Xiao Zhang, Jun Xu, Yang Song, Xiaoxue Zang, and Ji-Rong Wen. 2023. Enhancing Recommendation with Search Data in a Causal Learning Manner. *ACM Transactions on Information Systems* 41, 4 (2023), 1–31.
- [27] Zihua Si, Zhongxiang Sun, Xiao Zhang, Jun Xu, Xiaoxue Zang, Yang Song, Kun Gai, and Ji-Rong Wen. 2023. When Search Meets Recommendation: Learning Disentangled Search Representation for Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, July 23-27, 2023*. ACM, 1313–1323.
- [28] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (Beijing, China) (CIKM '19)*. ACM, New York, NY, USA, 1441–1450.
- [29] Zhongxiang Sun, Zihua Si, Xiaoxue Zang, Dewei Leng, Yanan Niu, Yang Song, Xiao Zhang, and Jun Xu. 2023. KuaiSAR: A Unified Search And Recommendation Dataset. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 5407–5411.
- [30] Zhongxiang Sun, Zihua Si, Xiao Zhang, Xiaoxue Zang, Yang Song, Hongteng Xu, and Jun Xu. 2024. To Search or to Recommend: Predicting Open-App Motivation with Neural Hawkes Process. *arXiv preprint arXiv:2404.03267* (2024).
- [31] Hongyan Tang, Junling Liu, Ming Zhao, and Xudong Gong. 2020. Progressive layered extraction (ple): A novel multi-task learning (mtl) model for personalized recommendations. In *Proceedings of the 14th ACM Conference on Recommender Systems*. 269–278.
- [32] Jiayi Tang and Ke Wang. 2018. Personalized top-n sequential recommendation via convolutional sequence embedding. In *Proceedings of the eleventh ACM international conference on web search and data mining*. 565–573.
- [33] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).
- [34] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [35] Yuhao Wang, Ha Tsz Lam, Yi Wong, Zirui Liu, Xiangyu Zhao, Yichao Wang, Bo Chen, Huifeng Guo, and Ruiming Tang. 2023. Multi-Task Deep Recommender Systems: A Survey. *arXiv preprint arXiv:2302.03525* (2023).
- [36] Chuhan Wu, Fangzhao Wu, Mingxiao An, Tao Qi, Jianqiang Huang, Yongfeng Huang, and Xing Xie. 2019. Neural News Recommendation with Heterogeneous User Behavior. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 4874–4883.
- [37] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. 2019. Session-based recommendation with graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 346–353.
- [38] Jiayi Xie, Shang Liu, Gao Cong, and Zhenzhong Chen. 2023. UnifiedSSR: A Unified Framework of Sequential Search and Recommendation. *arXiv preprint arXiv:2310.13921* (2023).
- [39] Xu Xie, Fei Sun, Zhaoyang Liu, Shiwen Wu, Jinyang Gao, Jiandong Zhang, Bolin Ding, and Bin Cui. 2022. Contrastive learning for sequential recommendation. In *2022 IEEE 38th international conference on data engineering (ICDE)*. IEEE, 1259–1273.
- [40] Jing Yao, Zhicheng Dou, Ruobing Xie, Yanxiong Lu, Zhiping Wang, and Ji-Rong Wen. 2021. USER: A Unified Information Search and Recommendation Model Based on Integrated Behavior Sequence. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management (Virtual Event, Queensland, Australia) (CIKM '21)*. Association for Computing Machinery, New York, NY, USA, 2373–2382.
- [41] Hamed Zamani and W. Bruce Croft. 2018. Joint Modeling and Optimization of Search and Recommendation. In *Proceedings of the First Biennial Conference on Design of Experimental Search & Information Retrieval Systems, Bertinoro, Italy, August 28-31, 2018 (CEUR Workshop Proceedings, Vol. 2167)*. CEUR-WS.org, 36–41.
- [42] Hamed Zamani and W. Bruce Croft. 2020. Learning a Joint Search and Recommendation Model from User-Item Interactions. In *Proceedings of the 13th International*

- Conference on Web Search and Data Mining (Houston, TX, USA) (WSDM '20)*. Association for Computing Machinery, New York, NY, USA, 717–725.
- [43] Kai Zhao, Yukun Zheng, Tao Zhuang, Xiang Li, and Xiaoyi Zeng. 2022. Joint Learning of E-Commerce Search and Recommendation with a Unified Graph Neural Network. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining (Virtual Event, AZ, USA) (WSDM '22)*. Association for Computing Machinery, New York, NY, USA, 1461–1469.
- [44] Wayne Xin Zhao, Junhua Chen, Pengfei Wang, Qi Gu, and Ji-Rong Wen. 2020. Revisiting alternative experimental settings for evaluating top-n item recommendation algorithms. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 2329–2332.
- [45] Guorui Zhou, Na Mou, Ying Fan, Qi Pi, Weijie Bian, Chang Zhou, Xiaoqiang Zhu, and Kun Gai. 2019. Deep interest evolution network for click-through rate prediction. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 5941–5948.
- [46] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep Interest Network for Click-Through Rate Prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '18)*. 1059–1068.
- [47] Kun Zhou, Hui Yu, Wayne Xin Zhao, and Ji-Rong Wen. 2022. Filter-Enhanced MLP is All You Need for Sequential Recommendation. In *Proceedings of the ACM Web Conference 2022 (Virtual Event, Lyon, France) (WWW '22)*. Association for Computing Machinery, New York, NY, USA, 2388–2399.