The Fortieth AAAI Conference on Artificial Intelligence (AAAI-26)

Dual-Horizon Interest Model for Unified Search and Recommendation

Wenhao Zhu[1], **Yuxin Li**[1], **Shuo Wang**[1], **Hao Wang**[1]*,

[1]School of Computer Science, Wuhan University, China
{zhuwenhao, liyuxuxin, wangshuo.ws, wanghao.cs}@whu.edu.cn

Abstract

Search and recommendation are pivotal for information access and are increasingly unified to exploit shared user-item interactions. Both tasks suffer from data sparsity, which joint modeling can mitigate by integrating behavioral data with or without explicit queries. However, existing unified frameworks rarely distinguish between users' long- and short-term interests, despite their divergent temporal dynamics in search and recommendation. In this work, we propose a novel model, DHIM, which explicitly disentangles and integrates users' long- and short-term interests across both the search and recommendation scenarios. First, long- and short-term interests are independently extracted from search and recommendation using a unified extraction strategy. These interests are then adaptively integrated via a cross-scenario fusion module. A self-supervised contrastive loss supervises the learning of both interest types within and across scenarios. The resulting representations are fed into downstream search and recommendation models for prediction. Extensive experiments on two public benchmarks demonstrate that our approach consistently outperforms single-scenario and state-of-the-art joint models, achieving superior accuracy and generalizability. To our knowledge, this is the first work to incorporate explicit dual-horizon interest modeling into a unified search and recommendation framework with self-supervised contrastive learning.

Introduction

In modern society, search and recommendation (S&R) constitute the primary means of information access. Contemporary platforms increasingly integrate both, exploiting shared user-item interactions to address explicit queries and infer implicit preferences. For example, short video services such as YouTube and TikTok embed S&R within a single application, enabling seamless transitions between active querying and passive browsing. Nonetheless, both paradigms suffer from data sparsity. A unified modeling framework that jointly leverages user behaviors from both scenarios can mitigate this issue by refining item ranking when queries are present and generating personalized recommendations otherwise, thereby enhancing overall performance.

Many studies have explored joint modeling of S&R. Some methods optimize S&R models by a joint training loss (Zamani and Croft 2018, 2020). Some others leverage search data to enhance recommendation performance (Wu et al. 2019; Si et al. 2022, 2023a,b; He et al. 2022). Recent work has developed unified S&R frameworks that jointly model the two tasks to exploit their intrinsic correlations and achieve mutual benefits (Yao et al. 2021; Zhao et al. 2022; Xie et al. 2024; Shi et al. 2024). For example, USER (Yao et al. 2021) unifies S&R behaviors into a heterogeneous sequence and applies a joint encoder to model their interdependencies. UnifiedSSR (Xie et al. 2024) uses a dual-branch encoder for query and interaction histories, augmented by a self-supervised session module to model dynamic user intent. UniSAR (Shi et al. 2024) models search-recommendation transitions via masked transformers, aligns them contrastively, and fuses them with cross-attention.

While prior studies (Guo et al. 2019; Shen et al. 2022; Zheng et al. 2022; Ou et al. 2025) demonstrate that explicitly modeling users' long- and short-term interests enhances performance in standalone recommendation or search tasks, no joint S&R framework has yet effectively modeled these dual-horizon preferences. User interests are inherently dynamic rather than static, which underscores the necessity of modeling short-term preferences to achieve accurate and timely personalization.

Indeed, user behavior in S&R is driven by a combination of relatively stable long-term interests and dynamically changing short-term interests. For instance, a consumer who consistently prefers electronic products may temporarily focus on cars when planning a purchase. Likewise, a user who receives recommendations for a specific travel destination may later initiate related searches, showing how interest cues can propagate between scenarios. These phenomena reveal that user intent in search or recommendation cannot be adequately captured by modeling either long-term or short-term interests in isolation. Instead, a unified modeling framework that adaptively integrates and transfers user interests across both scenarios is essential for enhancing relevance, personalization, and contextual responsiveness.

Accordingly, we present DHIM, a unified S&R model designed to explicitly capture long- and short-term interests. The model infers user interests by encoding behavioral sequences from both search and recommendation contexts,

---

[*] Corresponding author.

and introduces a cross-scenario fusion mechanism to enrich the representation of user preferences. The framework consists of four main modules: (1) a Cross-Scenario Behavioral Fusion Module that employs dual encoders and cross-attention to capture both intra- and inter-scenario behavioral dependencies, following the behavior transition design from UniSAR (Shi et al. 2024); (2) a Dual-Horizon Interest Encoder that separates long- and short-term interests within each scenario; (3) a Cross-Scenario Dual-Horizon Interest Fusion Module that aligns and integrates corresponding interests across domains, enabling mutual reinforcement; and (4) a Self-Supervised Contrastive Learning that imposes within- and cross-scenario constraints to enhance the quality and consistency of interest representations. Finally, prediction is performed using the Progressive Layered Extraction (Tang et al. 2020) (PLE) multi-task learning framework.

The major contributions of the paper are as follows:

- To the best of our knowledge, this is the first work to explicitly model both long- and short-term user interests for S&R within a unified framework.
- We introduce a self-supervised contrastive learning strategy for both intra- and inter-scenario interest modeling. Specifically, we construct proxy signals for long- and short-term interests and formulate contrastive objectives that encourage accurate disentanglement within each scenario and effective alignment across scenarios.
- Experimental results on two public datasets demonstrate that our model outperforms traditional single-scenario and existing joint S&R approaches, and achieves state-of-the-art performance.

Related Work

Sequential Recommendation

Sequential recommendation aims to predict users' next interactions by modeling temporal dependencies in behavior sequences. Compared to static preference models, it better captures evolving interests. Early approaches used Markov chains (Rendle, Freudenthaler, and Schmidt-Thieme 2010), while later methods employed deep neural networks, especially Transformers (Kang and McAuley 2018; Sun et al. 2019; Shin et al. 2024), for improved long-range modeling. To mitigate data sparsity and noise, self-supervised learning (Chen et al. 2022; Zheng et al. 2022; Wang et al. 2023; Ou et al. 2025) and denoising strategies (Yuan et al. 2021) have been introduced. Recent work also explores leveraging large language models (Ye et al. 2025) and auxiliary signals like user search behaviors to enhance sequential modeling.

Personalized Search

Personalized search systems tailor results to individual users by integrating the current query with user-specific signals such as historical behaviors and contextual information, rather than returning identical results for all users. Early methods relied on heuristic rules and hand-crafted features (Dou, Song, and Wen 2007; Wang et al. 2013); supervised models then leveraged labeled relevance data to improve ranking accuracy (White et al. 2013). Recent neural approaches capture high-dimensional representations of implicit intent from queries and clicks (Ge et al. 2018; Lu et al. 2019; Ma et al. 2020; Zhou, Dou, and Wen 2020; Zhou et al. 2021; Bai, Dou, and Wen 2024), with further gains from adversarial learning (Lu et al. 2019), contrastive learning (Zhou et al. 2021; Dai et al. 2023) and large language models (Tang et al. 2025). Despite these advances, personalized search and recommendation remain siloed; our work jointly models S&R tasks within a unified framework.

Search-enhanced Recommendation

Recent research treats user-initiated searches as explicit signals to strengthen recommendation. IV4Rec (Si et al. 2022) and its extension IV4Rec++ (Si et al. 2023a) use search actions as instrumental variables to decompose recommendation embeddings into causal and non-causal components. SESRec (Si et al. 2023b) applies co-attention to align and segment S&R behaviors, then refines recommendation interests via attention over the disentangled segments. Query-SeqRec (He et al. 2022) incorporates query–item co-occurrence into a sequential recommendation framework to better capture user intent. These models leverage search data to improve recommendation performance but overlook the opportunity to integrate S&R for complementary gains across both tasks.

Unified Search and Recommendation

Several recent studies have explored unified modeling of S&R. USER (Yao et al. 2021) merges behaviors into a unified sequence and applies a shared encoder to capture user interests. SRJGraph (Zhao et al. 2022) constructs a user–item graph with query attributes and applies intent-aware aggregation. UnifiedSSR (Xie et al. 2024) adopts dual branches for query and interaction histories with a self-supervised session module. UniSAR (Shi et al. 2024) extracts transition features via masked transformers, aligns them with contrastive learning, and fuses them using cross-attention. However, none of these methods explicitly disentangle dual-horizon interest for S&R within a unified framework. In contrast, our method disentangles dual-horizon interest in both tasks, fuses them across scenarios, and aligns them via contrastive learning.

Preliminaries

Let $\mathcal{U}$, $\mathcal{I}$, and $\mathcal{Q}$ denote the sets of users, items, and queries, respectively. The user subscript u is omitted for clarity when its meaning is evident from context. Each user $u \in \mathcal{U}$ is associated with an ordered interaction history: $S_T = \big[(x_1, b_1), (x_2, b_2), \cdots, (x_T, b_T)\big]$, where T is the total number of interactions and $b_t \in \{0, 1\}$ indicates whether the t-th interaction is a search ($b_t = 0$) or a recommendation ($b_t = 1$). The behavior x_t is defined as:

$$x_t = \begin{cases} i_t, & b_t = 1, \\ \langle q_t, C_t \rangle, & b_t = 0, \end{cases} \tag{1}$$

where $i_t \in \mathcal{I}$ is the item clicked in a recommendation interaction; $q_t \in \mathcal{Q}$ is the user's t-th search query and

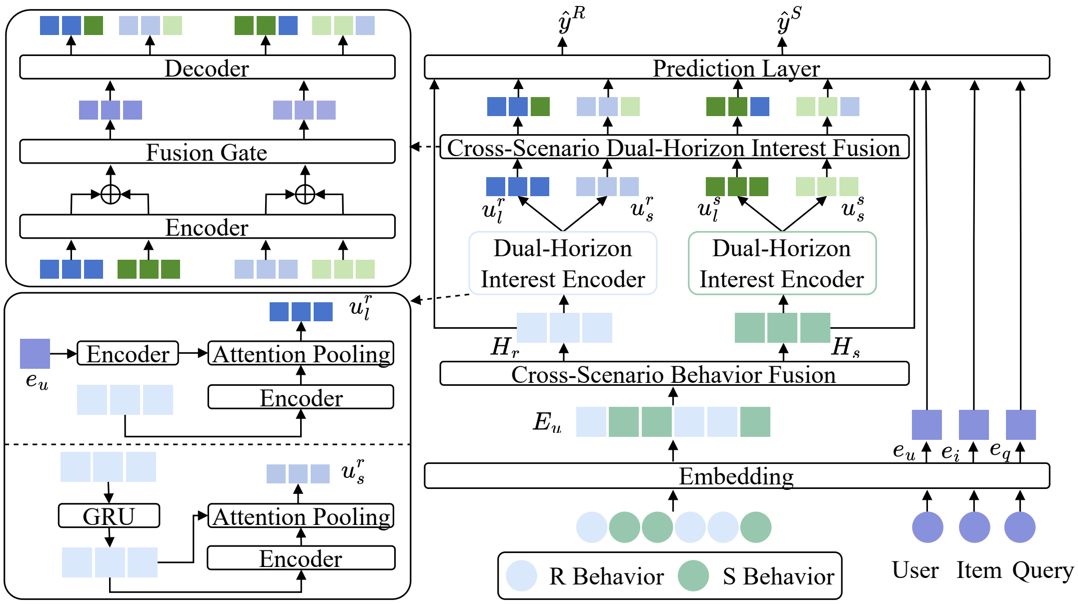

Figure 1: The architecture of the proposed DHIM.

$C_t = \{i_{t,1}, \cdots, i_{t,N_t}\} \subseteq \mathcal{I}$ denotes the set of items clicked in response to q_t.

All interactions are collected into the unified dataset $\mathcal{D} = \mathcal{D}_R \cup \mathcal{D}_S$, where $\mathcal{D}_R = \{(u, S_T, i_{T+1}, y_{T+1})_k\}_{k=1}^{N_R}$ and $\mathcal{D}_S = \{(u, S_T, i_{T+1}, y_{T+1})_k\}_{k=1}^{N_S}$, with y_{T+1} denoting the ground-truth preference scores for the next item i_{T+1} in recommendation and search contexts, respectively. Here, N_S and N_R denote the total number of interactions from recommendation and search scenarios, respectively. To unify the input structure across both tasks, we introduce a learnable query embedding $\boldsymbol{e}_q = \boldsymbol{q}_\phi \in \mathbb{R}^d$ for recommendation. We aim to learn a unified model f_Θ that maps

$$(u, S_T, q, i_{T+1}) \mapsto \hat{y}_{T+1}, \tag{2}$$

by minimizing the prediction error over $\mathcal{D}$, thus enabling accurate next-item preference estimation in both tasks.

Our Method

Overall Framework

Figure 1 depicts the framework. Raw inputs are first embedded into low-dimensional vectors. The Cross-Scenario Behavior Fusion module, designed based on UniSAR (Shi et al. 2024), integrates intra- and inter-scenario features from S&R behaviors. Dual-Horizon Interest Encoders extract dual-horizon interest independently for each scenario. The Cross-Scenario Dual-Horizon Interest Fusion module fuses matching interests from both scenarios into unified horizon-specific embeddings, which are subsequently decoded into scenario-specific representations. Self-supervised learning ensures intra- and inter-scenario consistency of these interests. Adaptive gating and attention fuse and aggregate these interests conditioned on candidate items. Finally, a PLE-based prediction layer produces outputs for both S&R tasks.

Embedding Layer

We define embedding tables $E_\mathcal{U}$, $E_\mathcal{I}$, and $E_\mathcal{W}$ for users, items, and query words, respectively. Here, $\mathcal{W}$ is the set composed of all the words included in the queries.

For a query $q = \{w_1, w_2, \cdots, w_{|q|}\} \subseteq \mathcal{W}$, we represent it by averaging (i.e., MEAN) the embeddings of its constituent words, based on the observation that query terms typically exhibit weak sequential dependencies: $e_q = \text{MEAN}(e_{w_1}, e_{w_2}, \cdots, e_{w_{|q|}})$.

To distinguish search from recommendation, the behavior embedding $e_{x_t} \in \mathbb{R}^d$ is e_{i_t} for recommendation and $e_{q_t} + \text{mean}(C_{q_t})$ for search, where C_{q_t} denotes the set of item embeddings clicked under query q_t.

Query-Item Contrastive Alignment. To enhance semantic relevance in the search scenario, we employ a bidirectional contrastive learning strategy to align queries and their clicked items in a shared latent space. Given each interaction $\langle q, C_q \rangle$, we treat (q, C_q) as positive pairs and randomly sample negative items I_{neg} and negative queries Q_{neg}. The contrastive loss is defined as:

$$\begin{aligned}\mathcal{L}_\text{Q} &= \text{BiCL}\big(e_q, C_q, I_\text{neg}, Q_\text{neg}; \tau_1\big) \\ &= -\Bigg[\sum_{i \in C_q} \log \frac{\exp\big(\text{sim}(e_q, e_i)/\tau_1\big)}{\sum_{i^- \in I_\text{neg}} \exp\big(\text{sim}(e_q, e_{i^-})/\tau_1\big)} \\ &\quad + \sum_{i \in C_q} \log \frac{\exp\big(\text{sim}(e_q, e_i)/\tau_1\big)}{\sum_{q^- \in Q_\text{neg}} \exp\big(\text{sim}(e_{q^-}, e_i)/\tau_1\big)}\Bigg],\end{aligned} \tag{3}$$

where τ_1 is a learnable temperature parameter, and $\text{sim}(a, b) = \tanh(a^\top W b)$ measures similarity in the latent space via a learnable projection matrix W.

Cross-Scenario Behavior Fusion Module

We denote the S&R behavior embeddings as $E_s \in \mathbb{R}^{N_s \times d}$ and $E_r \in \mathbb{R}^{N_r \times d}$, and the global sequence of user behaviors across both scenarios as $E_g \in \mathbb{R}^{(N_r+N_s)\times d}$, each with learnable positional encodings.

Independent transformer (Vaswani et al. 2017) encoders produce intra-scenario representations, $H^t_{\text{intra}} = \text{Transformers}_t\big(E_t\big), t \in \{r, s\}$. Subsequently, a transformer is applied to the global embedding E_g, employing a binary mask $M_{ij} = \mathbb{I}(b_i \neq b_j)$ to restrict attention across different scenarios, yielding the cross-scenario features: $(H^r_{\text{inter}}, H^s_{\text{inter}}) = \text{Transformers}_g\big(E_g, M\big)$.

Scenario-specific multi-head cross-attention (MCA) integrates intra- and inter-scenario representations, followed by a position-wise feedforward network (FFN):

$$H_t = \text{FFN}_t\big(\text{MCA}_t(H^t_{\text{intra}}, H^t_{\text{inter}}, H^t_{\text{inter}})\big), \tag{4}$$

where $t \in \{s, r\}$ and $Q = H^t_{\text{intra}}, K = V = H^t_{\text{inter}}$.

Following (Shi et al. 2024), we align intra- and inter-scenario representations via a bidirectional contrastive loss. Here, h^r_{intra} denotes the mean of recommendation-scenario behavior embeddings, and h^r_{inter} denotes the mean of cross-scenario behavior embeddings.

Positive pairs are constructed from $(h^r_{\text{intra}}, h^r_{\text{inter}})$ of the same user, while negative samples set $\mathcal{H}^{r-}_{\text{inter}}$ and $\mathcal{H}^{r-}_{\text{intra}}$ are randomly sampled from other users in the same batch. The contrastive loss is defined as:

$$\mathcal{L}^r_{\text{Align}} = \text{BiCL}\big(h^r_{\text{intra}}, h^r_{\text{inter}}, \mathcal{H}^{r-}_{\text{inter}}, \mathcal{H}^{r-}_{\text{intra}}; \tau_2\big), \tag{5}$$

where τ_2 is a learnable temperature parameter.

Similarly, the contrastive alignment loss for search is defined, and the overall behavior alignment loss is:

$$\mathcal{L}_{\text{Align}} = \mathcal{L}^s_{\text{Align}} + \mathcal{L}^r_{\text{Align}}. \tag{6}$$

Dual-Horizon Interest Encoder

After obtaining the final behavior representations in the recommendation and search scenarios, we denote them as: $H_t = [h_1, h_2, \cdots, h_{N_t}] \in \mathbb{R}^{N_t \times d}, t \in \{s, r\}$.

As both scenarios share the same architecture for long- and short-term interest encoders, we illustrate the design based on the recommendation scenario.

Long-Term Interest Encoder. The embedding e_u is commonly used to represent the user's relatively stable preferences. To align it with item representations, we project e_u into a latent query vector q_ℓ as the long-term query for long-term interest modeling. Meanwhile, each historical behavior representation h_j is mapped into the same latent space, resulting in the projected vector v_j, enabling semantic alignment with q_ℓ.

We then compute attention weights between the projected behavior features v_j and the long-term query q_ℓ using an interaction-based fusion mechanism:

$$\alpha_j = \text{MLP}(v_j \,||\, q_\ell \,||\, (v_j - q_\ell) \,||\, (v_j \odot q_\ell)), \tag{7}$$

where $||$ denotes concatenation, $\odot$ denotes the element-wise product, and MLP stands for multilayer perceptron.

After softmax normalization, the long-term interest representation is given by:

$$u^r_\ell = \sum_{j=1}^{N_r} a_j h_j \in \mathbb{R}^d, \quad a_j = \text{softmax}(\alpha_j). \tag{8}$$

Short-Term Interest Encoder. Unlike long-term interest extraction, short-term interest focuses on recent behavior patterns. Previous work (Zheng et al. 2022) employed two separate RNNs to extract short-term interests, whereas we simplify this by encoding the user's behavior sequence with a single GRU. The hidden state at each timestep captures the user's transient interest: $\{o_1, o_2, \cdots, o_{N_r}\} = \text{GRU}(\{h_1, h_2, \cdots, h_{N_r}\})$. The final hidden state $o_{N_r} \in \mathbb{R}^d$ representing the user's most recent interest, is directly used as the short-term query vector q_s.

Each hidden state o_j is first projected into the latent space and denoted as $\tilde{o}_j$. We compute attention weights with respect to the query q_s via an interaction-based mechanism:

$$\beta_j = \text{MLP}(\tilde{o}_j \,\|\, q_s \,\|\, (\tilde{o}_j - q_s) \,\|\, (\tilde{o}_j \odot q_s)). \tag{9}$$

After softmax normalization, the short-term interest representation is given by:

$$u^r_s = \sum_{j=1}^{N_r} b_j \tilde{o}_j \in \mathbb{R}^d, \quad b_j = \text{softmax}(\beta_j). \tag{10}$$

Through the above procedures, we obtain the user's long- and short-term interest representations in the recommendation and search scenarios: $u^r_\ell, u^r_s, u^s_\ell$, and u^s_s, respectively.

Cross-Scenario Dual-Horizon Interest Fusion

We assume that search and recommendation yield complementary long- and short-term interest representations. After encoding each scenario's interests into a shared latent space, we fuse them via two lightweight gating networks:

$$\mu_t = \sigma\big(G_t(\, u^r_t,\, u^s_t,\, u^r_t - u^s_t,\, u^r_t \odot u^s_t \,)\big), \tag{11}$$

$$u_t = \mu_t \, u^r_t + (1 - \mu_t) \, u^s_t, \tag{12}$$

where $t \in \{\ell, s\}$, G_ℓ, G_s are MLP-based gates, and σ is the sigmoid function. Decoding to each scenario's original space is performed by separate decoders, yielding the following reconstructed embeddings: $\hat{u}^r_\ell, \hat{u}^r_s, \hat{u}^s_\ell$, and $\hat{u}^s_s$.

Self-Supervised Contrastive Learning

To better capture users' long- and short-term interests, we introduce two complementary self-supervised tasks. Intra-scenario contrast encourages alignment with the corresponding interest proxy and separation from mismatched ones. Inter-scenario alignment encourages consistency between semantically similar interests across scenarios. Together, they provide auxiliary supervision to enhance the generalizability of learned representations.

Intra-Scenario Contrastive Objectives. Inspired by prior work (Zheng et al. 2022; Shen et al. 2022), we formalize two proxy representations per scenario. For recommendation, let p^r_ℓ denote the mean embedding over the complete behavior sequence (long-term) and p^r_s the mean over the most recent K interactions (short-term). Likewise, for search, we define p^s_ℓ and p^s_s analogously.

For the recommendation scenario:

$$\begin{aligned}
\operatorname{sim}(\hat{u}^r_\ell, p^r_\ell) &> \operatorname{sim}(\hat{u}^r_\ell, p^r_s), \\
\operatorname{sim}(p^r_\ell, \hat{u}^r_\ell) &> \operatorname{sim}(p^r_\ell, \hat{u}^r_s), \\
\operatorname{sim}(\hat{u}^r_s, p^r_s) &> \operatorname{sim}(\hat{u}^r_s, p^r_\ell), \\
\operatorname{sim}(p^r_s, \hat{u}^r_s) &> \operatorname{sim}(p^r_s, \hat{u}^r_\ell).
\end{aligned} \tag{13}$$

Inter-Scenario Contrastive Objectives. To promote cross-scenario consistency, we impose alignment constraints that draw embeddings of the same interest category closer while pushing apart those of different categories. Specifically, we project the search-side embeddings into the recommendation space via a shared matrix $W_s \in \mathbb{R}^{d\times d}$, and apply pairwise similarity constraints to align corresponding interests and separate mismatched ones:

$$\begin{aligned}
\operatorname{sim}\big(W_s\hat{u}^s_\ell,\ \hat{u}^r_\ell\big) &> \operatorname{sim}\big(W_s\hat{u}^s_\ell,\ \hat{u}^r_s\big), \\
\operatorname{sim}\big(\hat{u}^r_\ell,\ W_s\hat{u}^s_\ell\big) &> \operatorname{sim}\big(\hat{u}^r_\ell,\ W_s\hat{u}^s_s\big), \\
\operatorname{sim}\big(W_s\hat{u}^s_s,\ \hat{u}^r_s\big) &> \operatorname{sim}\big(W_s\hat{u}^s_s,\ \hat{u}^r_\ell\big), \\
\operatorname{sim}\big(\hat{u}^r_s,\ W_s\hat{u}^s_s\big) &> \operatorname{sim}\big(\hat{u}^r_s,\ W_s\hat{u}^s_\ell\big).
\end{aligned} \tag{14}$$

These constraints enable bidirectional alignment between search and recommendation, ensuring semantic consistency of long- and short-term interests across scenarios and improving cross-task generalization.

Contrastive Loss. We employ a triplet loss to optimize contrastive objectives:

$$\mathcal{L}_{\text{tri}}(a, p, q) = \max\left(d(a, p) - d(a, q) + m, 0\right), \tag{15}$$

where $d(\cdot,\cdot)$ denotes the Euclidean distance and m is a predefined margin.

The total contrastive loss is formulated as:

$$\begin{aligned}
\mathcal{L}^{t_{\text{intra}}}_{\text{con}} &= \mathcal{L}_{\text{tri}}(\hat{u}^t_\ell, p^t_\ell, p^t_s) + \mathcal{L}_{\text{tri}}(p^t_\ell, \hat{u}^t_\ell, \hat{u}^t_s) \\
&\quad + \mathcal{L}_{\text{tri}}(\hat{u}^t_s, p^t_s, p^t_\ell) + \mathcal{L}_{\text{tri}}(p^t_s, \hat{u}^t_s, \hat{u}^t_\ell)
\end{aligned} \tag{16}$$

$$\begin{aligned}
\mathcal{L}^{\text{inter}}_{\text{con}} &= \mathcal{L}_{\text{tri}}(\hat{u}^r_\ell, \hat{u}^s_\ell, \hat{u}^s_s) + \mathcal{L}_{\text{tri}}(\hat{u}^s_\ell, \hat{u}^r_\ell, \hat{u}^r_s) \\
&\quad + \mathcal{L}_{\text{tri}}(\hat{u}^r_s, \hat{u}^s_s, \hat{u}^s_\ell) + \mathcal{L}_{\text{tri}}(\hat{u}^s_s, \hat{u}^r_s, \hat{u}^r_\ell),
\end{aligned} \tag{17}$$

$$\mathcal{L}_{\text{con}} = \mathcal{L}^{r_{\text{intra}}}_{\text{con}} + \mathcal{L}^{s_{\text{intra}}}_{\text{con}} + \gamma \cdot \mathcal{L}^{\text{inter}}_{\text{con}}, \tag{18}$$

where $t \in \{s, r\}$, γ is a weighting factor controlling the influence of inter-scenario alignment.

Prediction and Training

Long-Short Interest Adaptive Fusion. Given the candidate item e_i, to decouple its relevance to the long- and short-term user preferences, we compress the behavior sequence $H_t, t \in \{s, r\}$ using a GRU, and obtain the final hidden state h_{f_t} as a global user representation. A gating network, implemented as an MLP followed by a sigmoid activation, produces the fusion weight:

$$\alpha_t = \sigma(\text{MLP}(h_{f_t} \,\|\, e_i \,\|\, \hat{u}^t_\ell \,\|\, \hat{u}^t_s)), \tag{19}$$

$$u^t_i = \alpha_t \hat{u}^t_\ell + (1 - \alpha_t)\hat{u}^t_s. \tag{20}$$

To capture item-specific global signals, we apply target-conditioned attention to independently aggregate each scenario's historical behaviors, producing $h_r, h_s \in \mathbb{R}^d$. The final fused representations for e_i are then given by $\hat{e}^r_i = h_r \,||\, u^r_i$ and $\hat{e}^s_i = h_s \,||\, u^s_i$.

Prediction Layer via PLE We use the PLE architecture to share low-level features and learn task-specific representations for S&R. The input $\mathbf{x} = \operatorname{concat}\big(e_u,\ e_i,\ e_q,\ \hat{e}^r_i,\ \hat{e}^s_i\big) \in \mathbb{R}^{d_h}$, where for recommendation without a query we set $e_q = q_\phi$ as a learnable embedding. The PLE module comprises shared and task-specific experts, with a gating network routing the input $\mathbf{x}$ to expert outputs and task-specific MLPs generating the final predictions $\hat{y}^R$ and $\hat{y}^S$.

Training Objective. We adopt the binary cross-entropy loss to train for click prediction tasks in both scenarios:

$$\begin{aligned}
\mathcal{L}^t_{\text{Click}} = -\frac{1}{|\mathcal{D}_t|} \sum_{(u,i,q)\in\mathcal{D}_t} \Big[& y^t_{u,i,q} \log \hat{y}^t_{u,i,q} \\
& + (1 - y^t_{u,i,q}) \log(1 - \hat{y}^t_{u,i,q}) \Big],
\end{aligned} \tag{21}$$

where $t \in \{S, R\}$, $\mathcal{D}_S$ and $\mathcal{D}_R$ denote the training data in search and recommendation scenarios, respectively.

Overall Loss. The training objective combines multiple loss components from different tasks and scenarios:

$$\mathcal{L}^{\text{t}} = \mathcal{L}^t_{\text{Click}} + \alpha \mathcal{L}^t_{\text{Align}} + \beta \mathcal{L}^t_{\text{con}} + \kappa \mathcal{L}^t_{\text{Q}}, \tag{22}$$

where $t \in \{S, R\}$, α, β, κ are hyperparameters controlling the contributions of each component.

The overall training objective jointly optimizes the final loss of S&R. Specifically, the total loss is computed as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}^R + \lambda_1 \mathcal{L}^S + \lambda_2 \|\theta\|^2_2, \tag{23}$$

where λ_1 controls the trade-off between S&R, θ denotes the parameters of DHIM, and λ_2 controls the L_2 regularization.

Experiments

Experiment Setup

Datasets. The experiments were conducted on two publicly available datasets. The KuaiSAR (Sun et al. 2023) dataset contains real-world user behaviors from both search and recommendation scenarios and is the only publicly available dataset with authentic cross-scenario S&R interactions. The Amazon (McAuley et al. 2015) dataset is a semisynthetic benchmark based on the "Kindle Store" subset of the five-core Amazon data, with search behaviors generated following prior work (Shi et al. 2024).

Baselines. We benchmark our approach against state-of-the-art methods in three categories: (1) Sequential recommendation models without search data, including DIN (Zhou et al. 2018), FMLPRec (Zhou et al. 2022) and BSARec (Shin et al. 2024); (2) Personalized search models without recommendation data, including AEM (Ai et al. 2019), ZAM (Ai et al. 2019) and MAI (Bai, Dou, and Wen 2024); and (3) Joint S&R models, including SESRec (Si et al. 2023b), UnifiedSSR (Xie et al. 2024) and UniSAR (Shi et al. 2024).

Method	KuaiSAR					Amazon				
	HR@1	HR@5	HR@10	NDCG@5	NDCG@10	HR@1	HR@5	HR@10	NDCG@5	NDCG@10
Recommendation Task										
DIN	0.1663	0.4553	0.6161	0.3145	0.3665	0.2369	0.5114	0.6419	0.3798	0.4221
FMLP-Rec	0.1346	0.4268	0.6125	0.2827	0.3427	0.2594	0.5859	0.7147	0.4308	0.4726
BASRec	0.1053	0.3538	0.5313	0.2307	0.2879	0.2130	0.5444	0.6945	0.3848	0.4334
SESRec†	0.1781	0.4843	0.6505	0.3356	0.3893	0.2975	0.5906	0.7118	0.4516	0.4909
UnifiedSSR†	0.1226	0.3979	0.5849	0.2620	0.3225	0.2121	0.5221	0.6704	0.3730	0.4210
UniSAR†	0.1824	0.4898	0.6568	0.3404	0.3944	0.3056	0.5947	0.7150	0.4574	0.4964
DHIM†	**0.2030***	**0.5217***	**0.6811***	**0.3677***	**0.4193***	**0.3080**	**0.6060***	**0.7270***	**0.4651***	**0.5044***
Personalized Search Task										
AEM	0.2520	0.5987	0.7332	0.4328	0.4764	0.3331	0.7124	0.8312	0.5330	0.5719
ZAM	0.2941	0.6210	0.7399	0.4658	0.5045	0.3324	0.7128	0.8326	0.5331	0.5723
MAI	0.4465	0.7256	0.8180	0.5969	0.6269	0.3676	0.7484	0.8671	0.5697	0.6085
UnifiedSSR†	0.4210	0.7337	0.8321	0.5883	0.6203	0.4276	0.8154	**0.9085***	0.6363	0.6667
UniSAR†	0.4269	0.7394	0.8327	0.5964	0.6267	0.5265	0.8179	0.8957	0.6846	0.7100
DHIM†	**0.4550**	**0.7441**	**0.8420***	**0.6118***	**0.6422***	**0.5282**	**0.8227***	0.9021	**0.6875**	**0.7134***

Table 1: Performance of DHIM and state-of-the-art baselines on KuaiSAR and Amazon datasets. Bold indicates the best result; underline, the second best. †: joint modeling of S&R. *: significant improvement over second best ($p < 0.01$).

Metrics. Consistent with prior studies (Zhou et al. 2022; Si et al. 2023b; Shi et al. 2024), we evaluate performance using Hit Ratio (HR) and Normalized Discounted Cumulative Gain (NDCG). Specifically, we report HR@1,@5, and @10, as well as NDCG@5 and @10. We pair the ground-truth item with 99 randomly sampled negative items from the user's non-interacted pool. In UniSAR's evaluation, the positive example is always fixed in the first position, causing any tied scores to favor it. To eliminate this positional bias, we randomize the ordering of positive and negative candidates at test time, thereby ensuring impartial tie resolution.

Implementation Details. All baseline hyperparameters are tuned following the original papers. The embedding dimension d is 64 and the maximum length for S&R is set to 30 for both datasets. The PLE framework uses 4 shared and 4 task-specific experts. The batch size is set as 1024. The short-term window size K is 2 for search and 4 for recommendation. The coefficient γ is fixed at 0.001, with α and β both set to 0.01. The hyperparameters κ and λ_1 are tuned over $\{0.001, 0.01, 0.1\}$ and $\{0.01, 0.1\}$, respectively.

Main Results

We conduct experiments on two S&R datasets, and the results in Table **??** show the following:

On the recommendation task, our model achieves the best performance across all metrics and datasets. Our proposed model outperforms the strongest unified baseline UniSAR on both datasets, particularly on the KuaiSAR dataset, where it improves HR@1 by 2.06% and NDCG@10 by 2.49%. These improvements stem from enhanced user interest modeling; specifically, our method builds upon UniSAR by introducing dual-horizon interest representations across both search and recommendation scenarios, enabling finer temporal alignment.

On the personalized search task, our model generally outperforms baselines on both datasets. Specifically, our model achieves improvements of 0.85% in HR@1 and 1.53% in NDCG@10 over UniSAR. Notably, our HR@10 on Amazon is marginally lower than UnifiedSSR. However, these differences are minimal and do not diminish the overall advantage of our model across most metrics and datasets.

Furthermore, joint training on search and recommendation consistently outperforms task-specific models, highlighting the benefit of cross-task integration. The performance gain is largely due to the mitigation of data sparsity, as joint modeling enables shared representation learning from complementary signals across tasks.

These results demonstrate that our unified approach, augmented by dual-horizon interest modeling, effectively captures cross-task user preferences and attains state-of-the-art performance.

Ablation Study

Table 2 shows that each component contributes to the joint performance on KuaiSAR, where "w/o" indicates the model without the corresponding component.

Impact of Dual-Horizon Interest Modules We first evaluate the contributions of long- and short-term interest representations in both recommendation and search branches. Removing either long- or short-term interests from either scenario consistently leads to performance degradation. Specifically, excluding long-term interest in the recommendation branch (*w/o rec-long*) results in a significant drop in recommendation performance (NDCG@5: 0.3677 → 0.3437),

Model Variant	Recommendation		Search		Overall	
	NDCG@5	NDCG@10	NDCG@5	NDCG@10	NDCG@5	NDCG@10
Full Model	**0.3677**	**0.4193**	0.6118	0.6422	**0.3820**	**0.4324**
w/o rec-long	0.3437	0.3969	0.6092	0.6399	0.3593	0.4111
w/o rec-short	0.3513	0.4016	**0.6129**	**0.6423**	0.3666	0.4157
w/o src-long	0.3472	0.4001	0.5936	0.6243	0.3616	0.4132
w/o src-short	0.3415	0.3932	0.6094	0.6402	0.3572	0.4077
w/o CSF	0.3376	0.3913	0.6105	0.6378	0.3536	0.4058
w/o $\mathcal{L}_{\text{con}}^{r_{\text{intra}}}$	0.3597	0.4113	0.6106	0.6401	0.3744	0.4247
w/o $\mathcal{L}_{\text{con}}^{s_{\text{intra}}}$	0.3282	0.3818	0.6107	0.6402	0.3448	0.3969
w/o $\mathcal{L}_{\text{con}}^{\text{inter}}$	0.3475	0.4008	0.6060	0.6359	0.3627	0.4146
w/o $\mathcal{L}_{Q}$	0.3534	0.4045	0.5954	0.6268	0.3676	0.4175
w/o PLE	0.3375	0.3907	0.6031	0.6348	0.3531	0.4050
only Rec/Src	0.3062	0.3584	0.5366	0.5719	0.3197	0.3709

Table 2: Ablation study on the KuaiSAR dataset reporting NDCG@5 and NDCG@10 for two tasks and overall performance.

confirming their importance in capturing stable user preferences. Similarly, removing short-term interest (*w/o rec-short*) in recommendation also impairs recommendation accuracy. Notably, excluding short-term interests from the recommendation branch yields a modest gain in search, which may result from reduced interference of short-term noise in the search encoder; however, the overall NDCG@5 still decreases from 0.3820 to 0.3666, indicating that isolated improvements cannot offset the overall performance loss. For the search branch, removing either long-term or short-term components also decreases both search and recommendation performance, demonstrating the mutual influence and necessity of dual-granularity interest modeling across tasks. Furthermore, disabling the Cross-Scene Dual-Horizon Interest Fusion (CSF) module, which integrates these interests across scenarios, consistently degrades both tasks, underscoring the importance of inter-scenario interest interaction.

Effect of Interest Alignment Modules To ensure coherent interest representations both within and across scenarios, we introduce intra-scenario and inter-scenario contrastive losses. First, the intra-scenario losses encourage separation of long- and short-term interests within each task domain. Ablating the search or recommendation intra-scenario loss degrades performance, highlighting their roles in filtering intra-domain noise and stabilizing preference encoding. Beyond within-task alignment, the inter-scenario loss $\mathcal{L}_{\text{con}}^{\text{inter}}$ aligns semantically related interests across recommendation and search. Excluding this loss (*w/o* $\mathcal{L}_{\text{con}}^{\text{inter}}$) decreases search NDCG@5 to 0.6060 and recommendation NDCG@5 to 0.3475, resulting in a decrease in overall NDCG@5 from 0.3820 to 0.3627. This confirms that cross-task contrastive alignment is critical for leveraging shared user signals and achieving robust joint modeling. Together, intra- and inter-scenario contrastive objectives synergize to produce disentangled, yet interoperable, interest representations that drive superior performance.

Effect of Multi-Task Learning Framework The PLE framework is introduced to disentangle shared and task-specific representations, thereby mitigating negative transfer and enabling precise parameter allocation. When PLE is ablated (*w/o PLE*), the expert-based architecture is replaced by two independent MLPs for S&R. Consequently, recommendation NDCG@5 declines from 0.3677 to 0.3375 and search NDCG@5 from 0.6118 to 0.6031, leading to a reduction in overall NDCG@5 from 0.3820 to 0.3531. This decline highlights that simple parameter sharing cannot capture the distinct yet related patterns required by each task. Further, when the model is restricted to a single scenario (*only Rec/Src*), each task is predicted by its own dedicated MLP, and all performance metrics suffer pronounced declines. Such degradation underscores the complementary nature of search and recommendation signals: search benefits from the stable, long-term patterns captured by the recommendation branch, while recommendation leverages the explicit, query-driven preferences revealed in search. These findings validate the imperative of a unified framework that jointly models both tasks, allowing the two to reinforce one another and achieve superior collective performance.

Additionally, the query-item alignment loss ($\mathcal{L}_Q$) aligns query embeddings with clicked item representations. Its removal leads to evident drops in both search and recommendation, indicating its necessity for modeling query–item consistency and ensuring fine-grained supervision.

Conclusion

In this paper, we propose DHIM, a unified S&R model that explicitly captures users' long- and short-term interests across scenarios. Dual-Horizon Interest Encoders extract scenario-specific interests, which are then fused by the Cross-Scenario Interest Fusion module into unified embeddings and decoded into scenario-specific representations. We incorporate self-supervised contrastive learning to enhance long- and short-term interest representations across and within scenarios. The model is jointly trained on search and recommendation tasks via the PLE multi-task framework, effectively leveraging diverse interest signals. Experiments on two public datasets confirm its effectiveness, with ablations highlighting each module's contribution.

Acknowledgments

The authors would like to acknowledge the support from National Natural Science Foundation of China (No. 62441229), National Key R&D Program of China (No. 2024YFE0111800), and State Key Laboratory of Cognitive Intelligence of China (Open Project iED2023-004).

References

Ai, Q.; Hill, D. N.; Vishwanathan, S. V. N.; and Croft, W. B. 2019. A Zero Attention Model for Personalized Product Search. In *CIKM, 2019*, 379–388.

Bai, Y.; Dou, Z.; and Wen, J.-R. 2024. Learning Dynamic Multi-attribute Interest for Personalized Product Search. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *EMNLP-Finding, 2024*, 2984–2993.

Chen, Y.; Liu, Z.; Li, J.; McAuley, J.; and Xiong, C. 2022. Intent Contrastive Learning for Sequential Recommendation. In *WWW, 2022*, 2172–2182.

Dai, S.; Liu, J.; Dou, Z.; Wang, H.; Liu, L.; Long, B.; and Wen, J.-R. 2023. Contrastive Learning for User Sequence Representation in Personalized Product Search. In *KDD, 2023*, 380–389.

Dou, Z.; Song, R.; and Wen, J.-R. 2007. A large-scale evaluation and analysis of personalized search strategies. In *WWW, 2007*, 581–590.

Ge, S.; Dou, Z.; Jiang, Z.; Nie, J.-Y.; and Wen, J.-R. 2018. Personalizing Search Results Using Hierarchical RNN with Query-aware Attention. In *CIKM, 2018*, 347–356.

Guo, Y.; Cheng, Z.; Nie, L.; Wang, Y.; Ma, J.; and Kankanhalli, M. 2019. Attentive Long Short-Term Preference Modeling for Personalized Product Search. *ACM Trans. Inf. Syst.*, 37(2).

He, Z.; Zhao, H.; Wang, Z.; Lin, Z.; Kale, A.; and Mcauley, J. 2022. Query-Aware Sequential Recommendation. In *CIKM, 2022*, 4019–4023.

Kang, W.-C.; and McAuley, J. 2018. Self-attentive sequential recommendation. In *ICDM, 2018*, 197–206. IEEE.

Lu, S.; Dou, Z.; Jun, X.; Nie, J.-Y.; and Wen, J.-R. 2019. PSGAN: A Minimax Game for Personalized Search with Limited and Noisy Click Data. In *SIGIR, 2019*, 555–564.

Ma, Z.; Dou, Z.; Bian, G.; and Wen, J.-R. 2020. PSTIE: Time Information Enhanced Personalized Search. In *CIKM, 2020*, 1075–1084.

McAuley, J.; Targett, C.; Shi, Q.; and van den Hengel, A. 2015. Image-Based Recommendations on Styles and Substitutes. In *SIGIR, 2015*, 43–52.

Ou, Z.; Zhang, X.; Zhu, Y.; Lyu, S.; Liu, J.; and Ao, T. 2025. LS-TGNN: Long and Short-Term Temporal Graph Neural Network for Session-Based Recommendation. 39(12): 12426–12434.

Rendle, S.; Freudenthaler, C.; and Schmidt-Thieme, L. 2010. Factorizing personalized markov chains for next-basket recommendation. In *WWW, 2010*, 811–820.

Shen, Q.; Wen, H.; Zhang, J.; and Rao, Q. 2022. Hierarchically Fusing Long and Short-Term User Interests for Click-Through Rate Prediction in Product Search. In *CIKM, 2022*, 1767–1776.

Shi, T.; Si, Z.; Xu, J.; Zhang, X.; Zang, X.; Zheng, K.; Leng, D.; Niu, Y.; and Song, Y. 2024. UniSAR: Modeling User Transition Behaviors between Search and Recommendation. In *SIGIR, 2024*, 1029–1039.

Shin, Y.; Choi, J.; Wi, H.; and Park, N. 2024. An attentive inductive bias for sequential recommendation beyond the self-attention. In *AAAI, 2024*.

Si, Z.; Han, X.; Zhang, X.; Xu, J.; Yin, Y.; Song, Y.; and Wen, J.-R. 2022. A Model-Agnostic Causal Learning Framework for Recommendation using Search Data. In *WWW, 2022*, 224–233.

Si, Z.; Sun, Z.; Zhang, X.; Xu, J.; Song, Y.; Zang, X.; and Wen, J.-R. 2023a. Enhancing Recommendation with Search Data in a Causal Learning Manner. *ACM Trans. Inf. Syst.*, 41(4).

Si, Z.; Sun, Z.; Zhang, X.; Xu, J.; Zang, X.; Song, Y.; Gai, K.; and Wen, J.-R. 2023b. When Search Meets Recommendation: Learning Disentangled Search Representation for Recommendation. In *SIGIR, 2023*, 1313–1323.

Sun, F.; Liu, J.; Wu, J.; Pei, C.; Lin, X.; Ou, W.; and Jiang, P. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. In *CIKM, 2019*, 1441–1450.

Sun, Z.; Si, Z.; Zang, X.; Leng, D.; Niu, Y.; Song, Y.; Zhang, X.; and Xu, J. 2023. KuaiSAR: A Unified Search And Recommendation Dataset. In *CIKM, 2023*, 5407–5411.

Tang, H.; Liu, J.; Zhao, M.; and Gong, X. 2020. Progressive Layered Extraction (PLE): A Novel Multi-Task Learning (MTL) Model for Personalized Recommendations. In *RecSys, 2020*, 269–278.

Tang, T.; Tian, Z.; Zhu, Z.; Wang, C.; Hu, H.; Tang, G.; Liu, L.; and Xu, S. 2025. LREF: A Novel LLM-based Relevance Framework for E-commerce Search. In *WWW ,2025*, 468–475.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. In *NIPS, 2017*, 6000–6010.

Wang, C.; Ma, W.; Chen, C.; Zhang, M.; Liu, Y.; and Ma, S. 2023. Sequential Recommendation with Multiple Contrast Signals. *ACM Trans. Inf. Syst.*, 41(1).

Wang, H.; He, X.; Chang, M.-W.; Song, Y.; White, R. W.; and Chu, W. 2013. Personalized ranking model adaptation for web search. In *SIGIR, 2013*, 323–332.

White, R. W.; Chu, W.; Hassan, A.; He, X.; Song, Y.; and Wang, H. 2013. Enhancing personalized search by mining and modeling task behavior. In *WWW, 2013*, 1411–1420.

Wu, C.; Wu, F.; An, M.; Qi, T.; Huang, J.; Huang, Y.; and Xie, X. 2019. Neural News Recommendation with Heterogeneous User Behavior. In Inui, K.; Jiang, J.; Ng, V.; and Wan, X., eds., *EMNLP-IJCNLP, 2019*, 4874–4883.

Xie, J.; Liu, S.; Cong, G.; and Chen, Z. 2024. UnifiedSSR: A Unified Framework of Sequential Search and Recommendation. In *WWW, 2024*, 3410–3419.

Yao, J.; Dou, Z.; Xie, R.; Lu, Y.; Wang, Z.; and Wen, J.-R. 2021. USER: A Unified Information Search and Recommendation Model based on Integrated Behavior Sequence. In *CIKM, 2021*, 2373–2382.

Ye, Y.; Zheng, Z.; Shen, Y.; Wang, T.; Zhang, H.; Zhu, P.; Yu, R.; Zhang, K.; and Xiong, H. 2025. Harnessing multimodal large language models for multimodal sequential recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 13069–13077.

Yuan, J.; Song, Z.; Sun, M.; Wang, X.; and Zhao, W. X. 2021. Dual sparse attention network for session-based recommendation. In *AAAI, 2021*, volume 35, 4635–4643.

Zamani, H.; and Croft, W. B. 2018. Joint Modeling and Optimization of Search and Recommendation. In *Biennial Conference on Design of Experimental Search & Information Retrieval Systems*.

Zamani, H.; and Croft, W. B. 2020. Learning a Joint Search and Recommendation Model from User-Item Interactions. In *WSDM, 2020*, 717–725.

Zhao, K.; Zheng, Y.; Zhuang, T.; Li, X.; and Zeng, X. 2022. Joint Learning of E-commerce Search and Recommendation with a Unified Graph Neural Network. In *WSDM, 2022*, 1461–1469.

Zheng, Y.; Gao, C.; Chang, J.; Niu, Y.; Song, Y.; Jin, D.; and Li, Y. 2022. Disentangling long and short-term interests for recommendation. In *WWW, 2022*, 2256–2267.

Zhou, G.; Zhu, X.; Song, C.; Fan, Y.; Zhu, H.; Ma, X.; Yan, Y.; Jin, J.; Li, H.; and Gai, K. 2018. Deep Interest Network for Click-Through Rate Prediction. In *KDD, 2018*, 1059–1068.

Zhou, K.; Yu, H.; Zhao, W. X.; and Wen, J.-R. 2022. Filter-enhanced MLP is All You Need for Sequential Recommendation. In *WWW, 2022*, 2388–2399.

Zhou, Y.; Dou, Z.; and Wen, J.-R. 2020. Encoding History with Context-aware Representation Learning for Personalized Search. In *SIGIR, 2020*, 1111–1120.

Zhou, Y.; Dou, Z.; Zhu, Y.; and Wen, J. 2021. PSSL: Self-supervised Learning for Personalized Search with Contrastive Sampling. In Demartini, G.; Zuccon, G.; Culpepper, J. S.; Huang, Z.; and Tong, H., eds., *CIKM, 2021*, 2749–2758.