

Article

Not peer-reviewed version

Enhancing User Intent for Recommendation Systems via Large Language Models

[Xiaochuan Xu](#)^{*}, Zeqiu Xu, [Peiyang Yu](#), Jiani Wang

Posted Date: 8 January 2025

doi: 10.20944/preprints202501.0627.v1

Keywords: LLMs; recommendation systems; user intent prediction



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Article

Enhancing User Intent for Recommendation Systems via Large Language Models

Xiaochuan Xu ^{1,*}, Zeqiu Xu ¹, Peiyang Yu ¹ and Jiani Wang ²

¹ Carnegie Mellon University, Pittsburgh, USA

² Stanford University, Stanford, USA

* Correspondence: xiaochux@alumni.cmu.edu

Abstract: Recommendation systems play a critical role in enhancing user experience and engagement in various online platforms. Traditional methods, such as Collaborative Filtering (CF) and Content-Based Filtering (CBF), rely heavily on past user interactions or item features. However, these models often fail to capture the dynamic and evolving nature of user preferences. To address these limitations, we propose DUIP (Dynamic User Intent Prediction), a novel framework that combines LSTM networks with Large Language Models (LLMs) to dynamically capture user intent and generate personalized item recommendations. The LSTM component models the sequential and temporal dependencies of user behavior, while the LLM utilizes the LSTM-generated prompts to predict the next item of interest. Experimental results on three diverse datasets—ML-1M, Games, and Bundle—show that DUIP outperforms a wide range of baseline models, demonstrating its ability to handle the cold-start problem and real-time intent adaptation. The integration of dynamic prompts based on recent user interactions allows DUIP to provide more accurate, context-aware, and personalized recommendations. Our findings suggest that DUIP is a promising approach for next-generation recommendation systems, with potential for further improvements in cross-modal recommendations and scalability.

Keywords: LLMs; recommendation systems; user intent prediction

1. Introduction

Recommendation systems are critical in various online services, from e-commerce and social media platforms to content streaming services. These systems aim to predict and suggest items that align with user preferences, improving user experience and engagement. Traditional recommendation algorithms typically rely on Collaborative Filtering (CF) and Content-Based Filtering (CBF) methods. CF uses historical interaction data, such as user-item interactions, to predict user preferences, while CBF relies on item features to match users with relevant content [1,2]. Despite their widespread adoption, these approaches often fall short in capturing the dynamic and context-dependent nature of user preferences, especially in real-time context.

In recent years, there has been significant interest in improving recommendation systems by incorporating more sophisticated models capable of understanding user intent and context. Traditional methods often assume that user preferences are static, based on past behavior, but emerging techniques have begun leveraging contextual information, such as user sentiment[3], emotional state, event information extraction[4,5] or temporal factors, to dynamically adjust recommendations. Large Language Models (LLMs) like GPT and BERT have shown exceptional promise in understanding complex user input and context[6], enabling them to extract deeper insights from user-generated content such as reviews, comments, and queries [7]. Recent studies have also explored the integration of LLMs with recommendation systems to improve personalized recommendations by leveraging natural language understanding [8].

Despite the progress, several challenges persist in the development of next-generation recommendation systems. One major issue is that traditional models often fail to accurately capture dynamic user intent, which can change over time and across contexts. While context-aware recommendation

systems have made strides in integrating real-time factors like sentiment or social context, they still lack a robust mechanism for predicting and adapting to user intent shifts in a timely and reliable manner [9].

Furthermore, the cold-start problem remains a significant hurdle, especially for new users or items where limited interaction data is available, making it difficult to infer user preferences effectively [10]. There is a need for systems that can dynamically infer user intent from sparse data, adapt recommendations in real-time, and seamlessly integrate cross-modal information (e.g., text, images, social context) [11].

To address these challenges, this paper proposes a novel framework that integrates LSTM networks and LLMs for dynamic user intent prediction, named **DUIP**. In our approach, an LSTM model is first employed to capture the dynamic intent of the user based on their recent interactions. The output from the LSTM, which serves as a prompt template, is then fed into a Large Language Model (such as GPT-2) to predict the next item the user may be interested in. By using this approach, we combine the LSTM's ability to capture user intent over time with the LLM's powerful language understanding to make more contextually aware and accurate predictions. This model not only addresses the cold-start problem by using dynamic prompts but also enhances personalization by continuously adapting to shifting user needs and preferences.

The rest of the paper is organized as follows: Section 2 provides an overview of the related work in the area of recommendation systems, focusing on approaches for dynamic intent modeling, context-aware recommendations, and the application of Large Language Models in these domains. Section 3 presents the detailed methodology of the proposed approach, explaining the integration of LSTM networks with LLMs for dynamic user intent prediction and item recommendation. Section 4 describes the experimental setup, including datasets, evaluation metrics, and comparison with baseline methods. Finally, Section 5 concludes the paper by summarizing the findings, discussing the implications of the results, and suggesting directions for future research.

2. Related Work

2.1. Intent Modeling in Recommendation Systems

Intent modeling has become a crucial research direction in recommendation systems, aiming to better understand and predict user needs. Traditional recommendation methods, such as Collaborative Filtering and Content-Based Filtering, rely heavily on user historical data [1]. However, these methods often fail to capture real-time user needs and emotional fluctuations, which has prompted interest in dynamic intent modeling approaches [12]. Sentiment analysis has been widely used to understand user intent, as it helps identify users' emotional states and adjusts recommendations accordingly [13]. Several studies have focused on integrating contextual information such as emotional states, location, and social context to enhance recommendation systems [8].

2.2. Application of Large Language Models in Recommendation Systems

The application of Large Language Models (LLMs), such as GPT and BERT, in recommendation systems has been gaining attention due to their ability to process and understand large amounts of text data. These models can generate more accurate and personalized recommendations by understanding the context and intent behind user-generated content [8]. The main research directions in this area include,

Sentiment-Aware Recommendation: By using LLMs to perform sentiment analysis on user reviews and social media data, systems can infer user emotional states and adjust recommendations accordingly [9].

Conversational Recommendation Systems: LLMs are applied in dialog-based recommendation systems, allowing the system to generate personalized recommendations based on real-time feedback from users [13].

Generative Recommendation: Some studies explore using LLMs for generative recommendation, where the model generates personalized recommendations directly from user inputs such as search queries or reviews, instead of matching items from historical data [7].

However, most existing research is primarily focused on sentiment analysis and conversational recommendation, with little emphasis on how LLMs can be used to enhance real-time user intent modeling and dynamically adjust recommendation strategies.

2.3. Intent Inference and Context-Aware Recommendations

Intent inference and context-aware recommendation have become core topics in modern recommendation system research. Unlike traditional static methods, dynamic intent inference captures real-time user needs and provides personalized recommendations accordingly [2]. Incorporating multi-modal information (e.g., text, images, social network data) allows recommendation systems to better understand users' current intent and emotional state [8].

Context-Aware Systems: Context-aware recommendation systems integrate user sentiment, social context, and emotional states to provide more accurate and personalized recommendations [9]. For example, combining sentiment analysis with historical data can improve recommendation responsiveness and user satisfaction.

Cross-Modal Recommendation Systems: There has been growing interest in cross-modal recommendation, which combines text, image, and video data to provide more diverse and personalized recommendations [1].

3. Methodology

This section introduces a novel methodology combining Long Short-Term Memory (LSTM) networks and Large Language Models (LLMs) to model dynamic user intent and generate personalized item recommendations. The key idea is to leverage the LSTM's hidden state to create a learnable soft prompt template that will guide the LLM in predicting the next item the user is likely to interact with, based on the user's evolving preferences.

3.1. Dynamic User Intent Modeling Using LSTM

The primary objective of this step is to capture dynamic user intent—the preferences and interests that change over time—as users interact with the system. To achieve this, we utilize LSTM networks, which excel at modeling sequential data and capturing temporal dependencies in user behavior.

Consider the sequence of user interactions as a time-series of events:

$$X = \{x_1, x_2, \dots, x_t\} \quad (1)$$

where each x_t represents the feature vector associated with an item interaction at time t . The feature vector x_t contains various pieces of information, such as: item-specific features (e.g., item ID, category, price) and user context (e.g., time of interaction, session length, device used). At each timestep, the LSTM network processes this sequence of interactions and produces a hidden state h_t , which encapsulates the user's evolving intent:

$$\begin{aligned} \mathbf{I}_t &= \sigma(\mathbf{X}_t \mathbf{W}_{xi} + \mathbf{H}_{t-1} \mathbf{W}_{hi} + \mathbf{b}_i), \\ \mathbf{F}_t &= \sigma(\mathbf{X}_t \mathbf{W}_{xf} + \mathbf{H}_{t-1} \mathbf{W}_{hf} + \mathbf{b}_f), \\ \mathbf{O}_t &= \sigma(\mathbf{X}_t \mathbf{W}_{xo} + \mathbf{H}_{t-1} \mathbf{W}_{ho} + \mathbf{b}_o), \\ \tilde{\mathbf{C}}_t &= \tanh(\mathbf{X}_t \mathbf{W}_{xc} + \mathbf{H}_{t-1} \mathbf{W}_{hc} + \mathbf{b}_c) \\ \mathbf{C}_t &= \mathbf{F}_t \odot \mathbf{C}_{t-1} + \mathbf{I}_t \odot \tilde{\mathbf{C}}_t \\ \mathbf{H}_t &= \mathbf{O}_t \odot \tanh \mathbf{C}_t \end{aligned} \quad (2)$$

where X_t is input, I_t , F_t , O_t , \tilde{C}_t and C_t are input gate, forgetting gate, output gate, candidate memory cell and memory cell respectively. W_{*i} are weight parameter and b_* are biases. σ and \tanh is activation function.

The hidden state h_t encodes both short-term preferences (e.g., recently viewed items) and long-term patterns (e.g., sustained interest in certain categories), allowing the model to adapt to changing user behaviors.

3.2. Soft Prompt Construction with LSTM's Hidden State

Once the LSTM has generated the hidden state h_t , the next step is to transform this hidden state into a soft prompt that will guide the LLM in making predictions. The soft prompt is a learnable representation that adapts to the user's evolving preferences and provides contextual information for the LLM to use when predicting the next item.

To form the soft prompt P , we first utilize the LSTM's hidden state h_t , which encodes the user's current intent. The hidden state is passed through a learnable transformation function $f(\cdot)$, which converts it into a structured vector suitable for input to the LLM.

$$P = f(h_t) \quad (3)$$

where P is the soft prompt derived from the LSTM's hidden state h_t , and $f(\cdot)$ is a transformation function (such as a linear layer or multi-layer perceptron). This function maps the LSTM hidden state to a more interpretable format that encapsulates the user's preferences and intent in a contextual vector.

In addition to the soft prompt, we also include hard prompts, which are fixed or predefined information about the user or system. Hard prompts can include static information such as:

- User interaction history: previous items the user interacted with.
- Item metadata: item categories, tags, etc.

The final soft prompt P is formed by combining both the dynamic (soft) prompt from the LSTM and the static (hard) prompts representing user-item interactions:

$$P = \langle \text{user } i \rangle \text{ has interacted with } \{z\} \text{ soft} \\ + \text{hard prompt } \langle r, p_i \rangle \quad (4)$$

where $\langle \text{user } i \rangle$ represents the user identity, $\{z\}$ represents the dynamic soft prompt constructed from the LSTM's hidden state h_t , and $\langle r, p_i \rangle$ represents hard prompts that encode historical information about the user's interactions with specific items.

By combining the soft and hard prompts, the model is able to adapt dynamically to the user's current preferences (via the LSTM) while also using historical context to inform the predictions.

Consider the case where the user has recently interacted with laptops and accessories. The LSTM's hidden state h_t could encode this intent, and the transformation function $f(h_t)$ would produce a soft prompt that reflects the user's interest in laptops. For example, the generated soft prompt might look like:

$$P = \langle \text{user}_i \rangle \text{ has interacted with } \{\text{laptop, accessories}\} \\ \text{soft} + \text{hard prompt} \langle \text{item}_j, \text{item}_k \rangle \quad (5)$$

This prompt now includes both the dynamic intent from the LSTM and historical interaction data, which can be passed to the LLM.

3.3. LLM for Item Prediction

Once the soft prompt P has been constructed, it is input into the Large Language Model (LLM), such as GPT-2, to predict the next item the user might be interested in. The LLM uses the contextual information from the prompt to generate predictions.

The LLM processes the soft prompt P and generates a ranked list of predicted items that are relevant to the user's current intent. By conditioning the LLM on the dynamic soft prompt, the model can generate context-aware recommendations.

For example, if the prompt indicates that the user is interested in laptops and accessories, the LLM might generate recommendations like:

$$\begin{aligned} &\text{Recommendations} \\ &= \{ \text{"Laptop Stand"}, \text{"Wireless Mouse"}, \text{"Laptop Bag"} \} \end{aligned} \quad (6)$$

The LLM generates the next item \hat{y} by selecting the item with the highest conditional probability:

$$\hat{y} = \arg \max_{y \in Y} P(y | P) \quad (7)$$

where Y is the set of candidate items, $P(y | P)$ is the probability of item y being the next relevant item, conditioned on the prompt P . This probabilistic framework ensures that the LLM can select the most contextually relevant item, guided by the user's current preferences and historical interaction data.

In this methodology, we integrate LSTM-based dynamic user intent modeling with LLM-based item prediction. The LSTM captures the user's evolving preferences over time, while the soft prompt generated from the LSTM's hidden state guides the LLM in predicting the next relevant item. By combining both soft and hard prompts, our system ensures that the recommendations are not only personalized but also contextually relevant based on the user's interaction history and current intent.

4. Experiments

In our experiments, we use three real-world datasets from diverse domains to evaluate the performance of our recommendation framework.

4.1. Datasets

The MovieLens-1M (ML-1M) dataset contains 3,416 items and 784,860 sessions, with an average session length of 6.85 and a density indicator of 1573.86. This dataset consists of user ratings for movies and is commonly used for evaluating collaborative filtering approaches.

The Amazon Games dataset is a subcategory of the larger Amazon dataset, containing 17,389 items and 100,018 sessions. The average session length is 4.18, and the density indicator is 24.04. This dataset includes user ratings for video games, which provides a suitable test for item recommendations in the entertainment domain.

The Amazon Bundle dataset contains 14,240 items and 2,376 sessions across three subcategories: Electronics, Clothing, and Food. The average session length is 6.73, with a density indicator of 1.12. This dataset includes session data with explicitly annotated user intents, making it particularly valuable for testing intent-based recommendation models.

For all datasets, we preprocessed the data by ordering the user interactions chronologically and dividing them into sessions. For ML-1M and Amazon Games, the interactions are grouped into sessions based on daily interactions. The Amazon Bundle dataset is already sessionized and annotated, so it is used directly. Each dataset is split into training, validation, and test sets using a chronological approach. Specifically, the first 80% of sessions are used for training, the subsequent 10% for validation, and the final 10% for testing. This split ensures that the model is trained on earlier data and tested on more recent interactions, mimicking real-world recommendation systems.

4.2. Baselines

To evaluate the performance of DUIP, we compare it with 11 baseline models categorized into three types: conventional methods, deep learning-based methods, and LLM-based methods. These baselines represent a wide spectrum of approaches from traditional techniques to advanced neural

network-based models, each capturing user intent and session information in different ways. Below is a detailed list of these baselines:

- Mostpop: Recommends the most popular items based on user interactions.
- SKNN [14]: Recommends session-level similar items using session-based nearest neighbor search.
- FPMC [15]: A matrix factorization method that incorporates the first-order Markov chain to model user-item interactions.
- NARM [16]: An RNN-based model with attention mechanisms to capture the main user intent from hidden states.
- STAMP [17]: Learns the user’s primary intent by focusing on the impact of the last item in the context.
- GCE-GNN [18]: Uses both local and global graphs to learn item representations and identify the main intent of a session.
- MCPRN [19]: Models users’ multiple purposes to derive a final session representation.
- HIDE [20]: Splits item embeddings into multiple chunks, each representing a specific intention to learn diverse user intents.
- Atten-Mixer [21]: Learns multi-granularity consecutive user intents for more accurate session representations.
- UniSRec [22]: A cross-domain model that uses item descriptions to learn transferable representations across different domains.
- NIR [23]: Adopts zero-shot prompting to recommend the next item.

4.3. Results and Analysis

The experimental results demonstrate that DUIP significantly outperforms a range of baseline models across multiple datasets, highlighting its effectiveness in addressing the dynamic nature of user preferences. The combination of LSTM-based dynamic user intent modeling and LLM-based next-item prediction allows DUIP to capture and adapt to evolving user behaviors, which is a major advantage over traditional methods and even some deep learning-based models.

As shown in Table 1, on the ML-1M dataset, DUIP delivered impressive performance, especially in HR@1 and HR@5, showing a clear improvement over models such as SKNN, NARM, and STAMP, which rely on static representations of user preferences or sequential patterns without the ability to adapt dynamically to shifts in intent. The LSTM component of DUIP processes user-item interaction sequences, capturing long-term dependencies in user behavior, while the LLM provides a powerful mechanism for generating context-aware predictions. This synergy enables DUIP to not only recommend relevant items but also rank them appropriately, as evidenced by its superior NDCG scores.

Table 1. Performance comparison on ML-1M

Model	HR@1	HR@5	NDCG@1	NDCG@5
MostPop	0.0004	0.0070	0.0004	0.0053
SKNN	0.1270	0.3600	0.1270	0.2530
FPMC	0.1132	0.3748	0.1132	0.2464
NARM	0.1692	0.5230	0.1692	0.3501
STAMP	0.1584	0.5078	0.1584	0.3367
GCE-GNN	0.1312	0.4748	0.1312	0.3044
MCPRN	0.1434	0.4788	0.1434	0.3157
HIDE	0.1498	0.4998	0.1498	0.3256
Atten-Mixer	0.1490	0.4932	0.1490	0.3216
UniSRec	0.0508	0.2508	0.0508	0.1459
NIR	0.0572	0.2326	0.0572	0.1436
DUIP	0.1883	0.5348	0.1883	0.3674

As shown in Table 2, the Games dataset, which presents a challenge due to sparse data and the variety of user interests, further showcases DUIP’s strength. Here, DUIP outperformed the baselines in both HR@1 and NDCG@5, with a notable improvement in NDCG@1, indicating its ability to provide both relevant and well-ranked recommendations. The LSTM’s ability to continuously update the user’s intent based on recent interactions gives DUIP an edge in adapting to the rapidly changing interests of users, especially in domains like entertainment, where preferences can vary significantly over short periods of time.

Table 2. Performance comparison on Games

Model	HR@1	HR@5	NDCG@1	NDCG@5
SKNN	0.0020	0.0020	0.0020	0.0020
FPMC	0.0498	0.2564	0.0498	0.1508
NARM	0.0572	0.2574	0.0572	0.1534
STAMP	0.0556	0.2586	0.0556	0.1555
GCE-GNN	0.0692	0.2744	0.0692	0.1701
MCPRN	0.0522	0.2416	0.0522	0.1432
HIDE	0.0696	0.2694	0.0696	0.1662
Atten-Mixer	0.0530	0.2472	0.0530	0.1475
UniSRec	0.0544	0.2512	0.0544	0.1482
NIR	0.1168	0.3406	0.1168	0.2310
DUIP	0.1732	0.3729	0.1732	0.2561

As shown in Table 3, The performance on the Bundle dataset, which is more session-based and has annotated user intents, also reflects DUIP’s capability in handling session-specific dynamics. Although the performance gap was narrower compared to ML-1M and Games, DUIP still outperformed the baselines in terms of HR@1 and HR@5, demonstrating that the model can effectively work with session-level data and adapt its recommendations to the immediate needs of users, even in sparse interaction environments.

Table 3. Performance comparison on Bundle

Model	HR@1	HR@5	NDCG@1	NDCG@5
MostPop	–	0.0042	–	0.0021
FPMC	0.0398	0.2475	0.0398	0.1395
NARM	0.0322	0.2322	0.0322	0.1303
STAMP	0.0365	0.2352	0.0365	0.1339
GCE-GNN	0.0360	0.2237	0.0360	0.1267
MCPRN	0.0360	0.2352	0.0360	0.1490
HIDE	0.0458	0.2585	0.0458	0.1495
Atten-Mixer	0.0525	0.2644	0.0525	0.1549
UniSRec	0.0496	0.2402	0.0496	0.1430
NIR	0.0975	0.2832	0.0975	0.1939
DUIP	0.1233	0.3001	0.1233	0.2217

5. Conclusion

In this paper, we proposed DUIP, a novel recommendation framework that integrates LSTM-based dynamic user intent modeling with LLM-based item prediction. The results of our extensive experiments across three diverse datasets—ML-1M, Games, and Bundle—demonstrate that DUIP significantly outperforms existing baseline models, including traditional methods, deep learning-based approaches, and LLM-based models. By capturing the dynamic nature of user intent and leveraging LSTM to model sequential interactions, DUIP is able to generate highly personalized, context-aware recommendations. The integration of LLMs further enhances the model’s ability to generate accurate

and relevant predictions, ensuring that recommendations are not only personalized but also reflect real-time shifts in user preferences.

The superior performance of DUIP across HR@1, HR@5, NDCG@1, and NDCG@5 metrics highlights its ability to improve both top-k recommendation accuracy and ranking quality, which are crucial in modern recommendation systems. This makes DUIP a promising solution to long-standing challenges such as cold-start and dynamic user intent prediction. Despite its impressive results, further research can focus on optimizing the model for handling sparse datasets and improving computational efficiency, especially for large-scale real-world applications.

Looking ahead, there are several potential directions for future research and improvement of DUIP. First, integrating additional multi-modal data (such as images, text, and social context) could further enhance the model's adaptability and understanding of user intent. Cross-domain recommendations could also be explored by expanding the model's capability to transfer knowledge across different types of recommendation tasks, such as between movies, products, and music. Furthermore, addressing the scalability of the model in handling large datasets and real-time user interactions could make DUIP even more effective in production environments. Lastly, enhancing the real-time adaptation mechanism of DUIP through online learning techniques could allow it to update its recommendations instantaneously as new user data becomes available.

References

1. Bondevik, J.N.; Bennin, K.E.; Önder Babur.; Ersch, C. A systematic review on food recommender systems. *Expert Systems with Applications* **2024**, *238*, 122166. <https://doi.org/https://doi.org/10.1016/j.eswa.2023.122166>.
2. Wang, S.; Zhang, X.; Wang, Y.; Ricci, F. Trustworthy recommender systems. *ACM Transactions on Intelligent Systems and Technology* **2024**, *15*, 1–20.
3. Yu, P.; Cui, V.Y.; Guan, J. Text classification by using natural language processing. In Proceedings of the Journal of Physics: Conference Series. IOP Publishing, 2021, Vol. 1802, p. 042010.
4. Liu, W.; Zhou, L.; Zeng, D.; Xiao, Y.; Cheng, S.; Zhang, C.; Lee, G.; Zhang, M.; Chen, W. Beyond Single-Event Extraction: Towards Efficient Document-Level Multi-Event Argument Extraction. *arXiv preprint arXiv:2405.01884* **2024**.
5. Liu, W.; Cheng, S.; Zeng, D.; Qu, H. Enhancing document-level event argument extraction with contextual clues and role relevance. *arXiv preprint arXiv:2310.05991* **2023**.
6. Yu, P.; Xu, X.; Wang, J. Applications of Large Language Models in Multimodal Learning. *Journal of Computer Technology and Applied Mathematics* **2024**, *1*, 108–116.
7. Zhu, Y.; Wu, L.; Guo, Q.; Hong, L.; Li, J. Collaborative large language model for recommender systems. In Proceedings of the Proceedings of the ACM on Web Conference 2024, 2024, pp. 3162–3172.
8. Ren, X.; Wei, W.; Xia, L.; Su, L.; Cheng, S.; Wang, J.; Yin, D.; Huang, C. Representation learning with large language models for recommendation. In Proceedings of the Proceedings of the ACM on Web Conference 2024, 2024, pp. 3464–3475.
9. Bakhshizadeh, M. Supporting Knowledge Workers through Personal Information Assistance with Context-aware Recommender Systems. In Proceedings of the Proceedings of the 18th ACM Conference on Recommender Systems, 2024, pp. 1296–1301.
10. Chaimalas, I.; Walker, D.M.; Gruppi, E.; Clark, B.R.; Toni, L. Bootstrapped personalized popularity for cold start recommender systems. In Proceedings of the Proceedings of the 17th ACM conference on recommender systems, 2023, pp. 715–722.
11. Li, X.; Sun, A.; Zhao, M.; Yu, J.; Zhu, K.; Jin, D.; Yu, M.; Yu, R. Multi-intention oriented contrastive learning for sequential recommendation. In Proceedings of the Proceedings of the sixteenth ACM international conference on web search and data mining, 2023, pp. 411–419.
12. Chen, Y.; Liu, Z.; Li, J.; McAuley, J.; Xiong, C. Intent contrastive learning for sequential recommendation. In Proceedings of the Proceedings of the ACM Web Conference 2022, 2022, pp. 2172–2182.
13. Sun, Z.; Liu, H.; Qu, X.; Feng, K.; Wang, Y.; Ong, Y.S. Large language models for intent-driven session recommendations. In Proceedings of the Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2024, pp. 324–334.

14. Jannach, D.; Ludewig, M. When recurrent neural networks meet the neighborhood for session-based recommendation. In Proceedings of the Proceedings of the eleventh ACM conference on recommender systems, 2017, pp. 306–310.
15. Rendle, S.; Freudenthaler, C.; Schmidt-Thieme, L. Factorizing personalized markov chains for next-basket recommendation. In Proceedings of the Proceedings of the 19th international conference on World wide web, 2010, pp. 811–820.
16. Li, J.; Ren, P.; Chen, Z.; Ren, Z.; Lian, T.; Ma, J. Neural attentive session-based recommendation. In Proceedings of the Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, 2017, pp. 1419–1428.
17. Liu, Q.; Zeng, Y.; Mokhosi, R.; Zhang, H. STAMP: short-term attention/memory priority model for session-based recommendation. In Proceedings of the Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining, 2018, pp. 1831–1839.
18. Wang, Z.; Wei, W.; Cong, G.; Li, X.L.; Mao, X.L.; Qiu, M. Global context enhanced graph neural networks for session-based recommendation. In Proceedings of the Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval, 2020, pp. 169–178.
19. Wang, S.; Hu, L.; Wang, Y.; Sheng, Q.Z.; Orgun, M.; Cao, L. Modeling multi-purpose sessions for next-item recommendations via mixture-channel purpose routing networks. In Proceedings of the International Joint Conference on Artificial Intelligence. International Joint Conferences on Artificial Intelligence, 2019.
20. Li, Y.; Gao, C.; Luo, H.; Jin, D.; Li, Y. Enhancing hypergraph neural networks with intent disentanglement for session-based recommendation. In Proceedings of the Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval, 2022, pp. 1997–2002.
21. Zhang, P.; Guo, J.; Li, C.; Xie, Y.; Kim, J.B.; Zhang, Y.; Xie, X.; Wang, H.; Kim, S. Efficiently leveraging multi-level user intent for session-based recommendation via atten-mixer network. In Proceedings of the Proceedings of the sixteenth ACM international conference on web search and data mining, 2023, pp. 168–176.
22. Hou, Y.; Mu, S.; Zhao, W.X.; Li, Y.; Ding, B.; Wen, J.R. Towards universal sequence representation learning for recommender systems. In Proceedings of the Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2022, pp. 585–593.
23. Wang, L.; Lim, E.P. Zero-shot next-item recommendation using large pretrained language models. *arXiv preprint arXiv:2304.03153* **2023**.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.