

Multi-intent-aware Session-based Recommendation

Minjin Choi Sungkyunkwan University Suwon, Republic of Korea zxcvxd@skku.edu

Hyunsouk Cho Ajou University Suwon, Republic of Korea hyunsouk@ajou.ac.kr

ABSTRACT

Session-based recommendation (SBR) aims to predict the following item a user will interact with during an ongoing session. Most existing SBR models focus on designing sophisticated neural-based encoders to learn a session representation, capturing the relationship among session items. However, they tend to focus on the last item, neglecting diverse user intents that may exist within a session. This limitation leads to significant performance drops, especially for longer sessions. To address this issue, we propose a novel SBR model, called Multi-intent-aware Session-based Recommendation Model (MiaSRec). It adopts frequency embedding vectors indicating the item frequency in session to enhance the information about repeated items. MiaSRec represents various user intents by deriving multiple session representations centered on each item and dynamically selecting the important ones. Extensive experimental results show that MiaSRec outperforms existing state-of-the-art SBR models on six datasets, particularly those with longer average session length, achieving up to 6.27% and 24.56% gains for MRR@20 and Recall@20. Our code is available at https://github.com/jin530/MiaSRec.

CCS CONCEPTS

• Information systems \rightarrow Recommender systems.

KEYWORDS

session-based recommendation; multiple intents

ACM Reference Format:

Minjin Choi, Hye-young Kim, Hyunsouk Cho, and Jongwuk Lee. 2024. Multi-intent-aware Session-based 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, 5 pages. https://doi.org/10.1145/3626772.3657928

1 INTRODUCTION

Session-based recommendation (SBR) [16, 30] aims to learn hidden user preferences in a session and provide personalized items for each user. A session refers to a sequence of user-item interactions

 $^{\ast} \text{Corresponding author}$



This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

SIGIR '24, July 14–18, 2024, Washington, DC, USA © 2024 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-0431-4/24/07. https://doi.org/10.1145/3626772.3657928 Hye-young Kim Sungkyunkwan University Suwon, Republic of Korea khyaa3966@skku.edu

Jongwuk Lee* Sungkyunkwan University Suwon, Republic of Korea jongwuklee@skku.edu



Figure 1: A session example with multiple user intents, such as *travel*, *fashion*, *sun protection*, and *photo*. Dotted rectangles represent items related to each user intent, and solid rectangles represent recommendations for each user intent.

over time, *e.g.*, consecutive clicks on multiple products during a transaction. It is particularly effective for anonymous or first-time users in web applications like e-commerce and streaming services, *e.g.*, Amazon, YouTube, Netflix, and Spotify [1, 5, 6, 21]. SBR inherently suffers from extreme data sparsity because it only deals with user actions during an ongoing session, making it challenging to capture dynamic and intricate item correlations.

Existing SBR models [11, 12, 17, 20, 25, 32, 33, 35] have primarily focused on extracting a single representation from a session to capture and express user preferences. They mainly aimed to model a session consisting of multiple items using various neural-based session encoders, including recurrent neural networks (RNNs) [11, 12, 20], graph neural networks (GNNs) [25, 32, 33], or Transformers [17, 35]. However, despite their advanced encoder designs, modeling only a single representation cannot express multiple user intents.

Figure 1 illustrates the importance of using multiple user intents. While the user may be interested in *photo* when focusing on the last item "camera", looking at the entire session suggests that the user will click items about *travel*. Considering the other items in the session, *fashion* or *sun protection* also align with different user intents. In this scenario, it is more appropriate to recommend a top-*N* item list with appropriate multiple intents, *e.g.*, ("travel bag", "sneakers", and "photo frame"). On the other hand, in some sessions, not all items are important. For example, in Figure 1, "speaker" is less relevant to the other items for capturing user intents.

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



Figure 2: The model architecture of MiaSRec.

Recently, some SBR models, such as MSGIFSR [7] and Atten-Mixer [36], have attempted to capture multiple user intents, focusing primarily on the last few consecutive items to represent diverse user intents. However, they cannot accurately capture user intent if the last item is less important or noisy. Besides, some studies [2, 28, 31, 37] have attempted to identify multiple user interests over a long user-item history. They employ a fixed number of user interests and extract the same number of interests for all users. It may miss some interests or include unnecessary ones since the number of user interests varies by user. We claim the challenges for modeling multiple user intents in the session: (i) how to fully capture multiple user intents inherent in each session and (ii) how to filter out unimportant ones among multiple intents.

To address these issues, we propose a novel SBR model, called *Multi-intent-aware Session-based Recommendation Model (MiaS-Rec)*, as shown in Figure 2. First, MiaSRec encodes the session items with position and frequency embeddings, reflecting sequential information and repeat patterns. It then employs a self-attention mechanism and a high-way network to derive different user intents in the session. Then, it adaptively extracts diverse user intents in the session. Lastly, MiaSRec decodes multiple session representations into item distributions and aggregates them using pooling functions. Despite its simplicity, extensive experiments demonstrate that MiaSRec outperforms existing SBR models on six benchmark datasets. Notably, MiaSRec achieves significant gains for longer sessions (≥ 10) with multiple user intents, up to 13.51% in Recall@20.

2 PROPOSED MODEL

2.1 Session-based Recommendation

Let $\mathcal{V} = \{v_1, \ldots, v_n\}$ denote a set of *n* unique items, *e.g.*, products and songs. An arbitrary session $s = (v_{t_1}, \ldots, v_{t_{|s|}})$ represents a sequence of |s| items that a user interacts with, *e.g.*, clicks, views,

Minjin Choi, Hye-young Kim, Hyunsouk Cho, and Jongwuk Lee

and purchases. Here, t_i indicates the index of the *i*-th item in the session. Given a session, the goal of SBR is to predict the next item $v_{t_{|s|+1}}$ that the user is likely to consume. The SBR model takes a session as an input and returns a top-*N* item list to recommend.

2.2 Model Architecture

In this section, we present a novel SBR model called MiaSRec, which aims to address (i) how to represent multiple user intents in the session and (ii) how to prune out unnecessary user intents.

2.2.1 Embedding Layer. We first embed each session item v_{t_i} into item embedding vector $\mathbf{v}_{t_i} \in \mathbb{R}^d$, and generate the mean item embedding $\mathbf{v}_m = \frac{1}{|s|} \sum_{i=1}^{|s|} \mathbf{v}_{t_i}$, capturing the global session information. To better capture the importance of each item in a session, we incorporate the absolute position embedding vector $\mathbf{a}_i \in \mathbb{R}^d$ to distinguish sequential order and the frequency embedding vector $\mathbf{r}_{f_i} \in \mathbb{R}^d$ to express the importance of repeated items in a session. Here, f_i denotes the frequency of the *i*-th item in a session. In Figure 2, given session $s = (v_1, v_3, v_2, v_3)$, the frequency is (1, 2, 1, 2). Finally, the item, position, and frequency embedding vectors are combined as the model input x_i .

$$\mathbf{x}_i = \mathbf{v}_{t_i} + \mathbf{a}_i + \mathbf{r}_{f_i} \text{ for } i \in \{1, \dots, |s|, m\}.$$

$$\tag{1}$$

Note that we sort session items in reverse order as in [32], so the last item $v_{t_{|s|}}$ always corresponds to the first positional embedding \mathbf{a}_1 . And, we randomly initialize the learnable parameter of position and frequency embeddings.

2.2.2 Multi-intent Representation. We employ the self-attention mechanism [29] to capture the complex relationships among session items. Using the bi-directional self-attention layer Self-attention(\cdot), we generate multiple contextualized representations of the session information associated with each item as follows:

$$\mathbf{c}_1, \dots, \mathbf{c}_{|s|}, \mathbf{c}_m = \text{Self-attention}([\mathbf{x}_1, \dots, \mathbf{x}_{|s|}, \mathbf{x}_m])). \tag{2}$$

To ensure that multiple contextualized representations do not become similar and to better reflect different user intent, we leverage the high-way network [25] by emphasizing item embeddings. Specifically, we combine contextualized representation \mathbf{c}_i and item embedding \mathbf{v}_i to derive user intent representation $\mathbf{o}_i \in \mathbb{R}^d$ that takes into account both the overall information of the session and the information of each item.

$$\mathbf{o}_{i} = \mathbf{g} \odot \mathbf{v}_{i} + (1 - \mathbf{g}) \odot \mathbf{c}_{i} \text{ for } i \in \{1, \dots, |s|, m\},$$

where $\mathbf{g} = \sigma(\mathbf{W}_{\mathbf{g}}[\mathbf{v}_{i}; \mathbf{c}_{i}]^{\top}),$
(3)

where $\mathbf{W}_{\mathbf{g}} \in \mathbb{R}^{d \times 2d}$ is a learnable weight matrix, $\mathbf{g} \in \mathbb{R}^{d}$ is a gating vector, and $\sigma(\cdot)$ is the sigmoid function.

2.2.3 Intent Selection. We employ multiple session representations to fully exploit the potential of each session item. However, not all session items may be necessary, and some may be noisy. To extract essential user intents in a session, we calculate the importance of multiple representations and remove unimportant ones. We utilize a sparse transformation α -entmax [26, 35], assigning a zero probability to unimportant representations.

$$\alpha \text{-entmax}(\mathbf{z}) = \operatorname*{argmax}_{\mathbf{p} \in \Delta^{l}} \mathbf{p}^{\top} \mathbf{z} + H_{\alpha}^{T}(\mathbf{p}), \tag{4}$$

Multi-intent-aware Session-based Recommendation

where $H_{\alpha}^{T}(\mathbf{p}) = \frac{1}{\alpha(\alpha-1)} \sum_{j} (\mathbf{p}_{j} - \mathbf{p}_{j}^{\alpha})$ if $\alpha \neq 1$, else $H_{1}^{T}(p) = -\sum_{j} \mathbf{p}_{j} \log \mathbf{p}_{j}^{-1}$. Depending on α , 1-entmax and 2-entmax are equivalent to softmax and sparsemax [24], respectively. A larger α value generates a more sparse probability distribution, and we set α as 1.5 empirically. We extract session representation $\mathbf{h}_{i} \in \mathbb{R}^{d}$ by masking unnecessary user intents.

$$\gamma = \alpha \operatorname{entmax}(\mathbf{w} \cdot [\mathbf{o}_1; \dots; \mathbf{o}_{|s|}; \mathbf{o}_m]^{\top}),$$

$$\{\mathbf{h}_1, \dots, \mathbf{h}_k\} = \{\gamma_i \mathbf{o}_i | \gamma_i > 0, i \in \{1, \dots, |s|, m\}\},$$

(5)

where $\mathbf{w} \in \mathbb{R}^d$ is a learnable parameter and $\gamma \in \mathbb{R}^{|s|+1}$ is the importance weight vector for each item. k is the number of non-zero elements in γ . Unlike the previous multiple representation models [2, 7, 36, 37] utilize a fixed number of representations regardless of the session, MiaSRec dynamically selects multiple session representations for a given session, up to the number of session items.

2.2.4 *Multi-intent Aggregation.* The multi-intent aggregation process of MiaSRec is divided into two parts: (i) decoding item distributions from multiple session representations and (ii) aggregating the distributions for the final recommendation.

To decode each session representation into the item distribution, we employ cosine similarity with item embedding matrix, *i.e.*, dot product with L2-normalization. For simplicity, we reuse the item embedding look-up table $\mathbf{V} \in \mathbb{R}^{n \times d}$, so the number of model parameters does not increase.

$$\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_k = [\tilde{\mathbf{h}}_1 \tilde{\mathbf{V}}^\top, \dots, \tilde{\mathbf{h}}_k \tilde{\mathbf{V}}^\top], \tag{6}$$

where $\tilde{\mathbf{h}}_i$ and $\tilde{\mathbf{V}}$ are the normalized session vector and the normalized item embedding matrix, respectively. $\hat{\mathbf{y}}_i \in \mathbb{R}^n$ represents an item distribution decoded by the session vector \mathbf{h}_i .

To aggregate multiple item distributions, we adopt max-pooling and average-pooling functions. Max-pooling maintains the principal features for multiple intents, and average-pooling captures consistent intent over the session.

$$\hat{\mathbf{y}} = \beta \text{MaxPool}(\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_k) + (1 - \beta) \text{AvgPool}(\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_k), \quad (7)$$

where $\hat{\mathbf{y}} \in \mathbb{R}^n$ is the final aggregated item distribution, and β is a combination hyperparameter. MaxPool(·) and AvgPool(·) represent the max- and average-pooling functions that fetch the maximum (or average) value for each dimension from multiple vectors.

Lastly, we formulate a cross-entropy loss function for training.

$$L(\mathbf{y}, \hat{\mathbf{y}}) = -\sum_{j=1}^{n} \mathbf{y}(j) \log(\frac{\exp(\hat{\mathbf{y}}(j)/\tau)}{\sum_{i} \exp(\hat{\mathbf{y}}(i)/\tau)}),$$
(8)

where $\mathbf{y} \in \mathbb{R}^n$ is a one-hot vector of the target item, *i.e.*, $\mathbf{y}(j) = 1$ if the *j*-th item is the target item; otherwise, $\mathbf{y}(j) = 0$. Here, τ is the hyperparameter to control the temperature [14] for better convergence [9].

3 EXPERIMENTS

3.1 Experimental Setup

Datasets. We conduct extensive experiments on six real-world datasets collected from e-commerce and music streaming services:

Table 1: Statistics of the various benchmark datasets. AvgLen indicates the average length of entire sessions.

Dataset	# Interacts	# Sessions	# Items	AvgLen
Diginetica	786,582	204,532	42,862	4.12
Retailrocket	871,637	321,032	51,428	6.40
Yoochoose	1,434,349	470,477	19,690	4.64
Tmall	427,797	66,909	37,367	10.62
Dressipi	4,305,641	943,658	18,059	6.47
LastFM	3,510,163	325,543	38,616	8.16

Diginetica, Retailrocket, Yoochoose, Tmall², Dressipi, and LastFM. For data pre-processing, we follow the conventional procedure [20, 22, 23, 33]. We discard the sessions with a single item and the items that occur less than five times in entire sessions. We split training, validation, and test sets chronologically as the 8:1:1 ratio. Table 1 summarizes detailed statistics on all benchmark datasets.

Baselines. We compare MiaSRec with the following SBR models: **SASRec** [17], **SR-GNN** [33], **NISER+** [9], **SGNN-HN** [25], **DSAN** [35], **LESSR** [4], **CORE** [15]. We also compare with the subsequent multiple representations models. **SINE** [28], **ComiRec** [2], **Re4** [37], **Atten-mixer (A-mixer)** [36], **MSGIFSR** [7]. We do not consider SBR models which use additional information, *e.g.*, temporal information [8, 27] or content-based features [3, 13, 19, 40].

Evaluation protocol and metrics. As the common protocol to evaluate SBR models [20, 33], we adopt the *iterative revealing scheme*, which iteratively exposes an item from a session to the model. We adopt Recall (R@20) and Mean Reciprocal Rank (M@20) to quantify the prediction accuracy of the next single item. All experimental results are averaged over three runs with different seeds, and we conduct the significance test using a paired t-test.

Implementation details. For reproducibility, we implement MiaSRec and the other baseline models on an open-source recommendation system library RecBole³ [38, 39]. We optimize all baselines using Adam optimizer [18] with a learning rate of 0.001. We set the embedding dimension to 100 and the max session length to 50. We stop the training if the validation MRR@20 decreases for three consecutive epochs ⁴. For all methods, we set the batch size to 1024. For MiaSRec, we set α as 1.5 for α -entmax and tune the temperature τ among {0.01, 0.05, 0.07, 0.1, 0.5, 1}, dropout ratio δ among {0, 0.1, 0.2, 0.3, 0.4, 0.5}. We search β from 0 to 1 in 0.1 increments. We follow the original papers' settings for other hyperparameters of baseline models, but if not available, we thoroughly tune them.

3.2 Experimental Results

Overall comparison. Table 2 reports the performance comparison between MiaSRec and other baseline models. (i) MiaSRec shows the best performance on all datasets. Note that MiaSRec demonstrates performance improvements of up to 24.56% for R@20 compared to the best baseline. In particular, MiaSRec shows substantial improvements in Recall on datasets with longer average session lengths (*e.g.*, Tmall and LastFM). (ii) Multiple representation models, such

 $[\]overline{{}^{1}\Delta^{l} := \{\mathbf{p} \in \mathbb{R}^{l} : \mathbf{p} \ge 0, \|\mathbf{p}\|_{1} = 1\}}$ denotes the *l*-probability simplex.

²Since it has been widely used in previous studies [10, 15, 34], we adopt Tmall even though it consists of timestamps in units of days, not in minutes or seconds. ³https://github.com/RUCAIBox/RecBole

⁴We report the performance on the test set using the models that show the highest performance on the validation set.

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

Table 2: Performance comparison for MiaSRec and baseline models. Imp. indicates how much better MiaSRec is than the best baseline model. The best model is marked **bold** and the second best model is <u>underlined</u>. Significant differences ($\rho < 0.01$) between the best baseline model and MiaSRec are reported with \dagger .

Dataset	Metric	SASRec	SR-GNN	NISER+	SGNN-HN	DSAN	LESSR	CORE	SINE	ComiRec	Re4	A-mixer	MSGIFSR	MiaSRec	Imp.
Diginetica	R@20	49.86	48.01	51.11	50.60	52.06	48.70	52.89	46.45	51.22	51.59	49.84	53.20	53.54	0.65%
	M@20	17.20	16.60	18.21	17.28	18.25	16.96	<u>18.53</u>	16.10	18.35	18.47	17.07	18.37	19.47 [†]	5.04%
Retailrocket R@	R@20	59.70	58.01	60.70	57.43	61.13	56.56	61.77	55.11	61.56	61.65	59.49	<u>63.04</u>	63.37	0.26%
	M@20	35.71	36.01	38.18	35.39	<u>38.68</u>	36.82	38.49	34.15	38.16	38.10	36.25	38.42	39.23 †	1.41%
Yoochoose Ro M	R@20	63.64	62.28	63.50	61.60	63.73	62.78	64.64	57.50	63.13	62.67	63.73	65.20	65.37	0.26%
	M@20	28.66	28.36	29.06	27.97	29.23	28.84	28.25	25.07	28.29	28.00	29.32	<u>30.02</u>	30.74 [†]	2.39%
Tmall	R@20	35.80	33.47	40.39	39.71	42.82	32.59	<u>44.91</u>	35.66	42.40	41.56	38.76	35.39	55.94 [†]	24.56%
	M@20	25.08	24.75	29.48	24.16	30.85	24.19	<u>31.59</u>	22.41	28.43	28.56	28.52	22.19	33.57 [†]	6.27%
Dressipi R(R@20	37.18	36.10	38.19	38.35	37.77	37.71	38.14	38.18	<u>39.60</u>	39.15	37.75	38.43	42.26	6.73%
	M@20	14.31	14.51	15.34	15.05	15.13	14.73	15.54	15.46	16.07	15.92	15.24	15.90	16.70 [†]	3.92%
LastFM	R@20	20.53	21.80	22.50	22.72	22.47	22.31	22.75	22.17	22.13	23.02	22.93	22.73	25.85 [†]	12.32%
	M@20	6.22	8.70	8.79	7.66	7.93	<u>8.80</u>	7.83	7.57	7.83	8.50	8.74	8.20	9.95 [†]	13.06%



Figure 3: Performance comparison of SBR models over varying session lengths. Sessions are divided into six groups depending on session length.

as MSGIFSR [7], ComiRec [2] and Re4 [37], tend to show higher accuracy than single representation models. This implies that various intents can exist in session, and it is necessary to capture them.

Effect of session length. Figure 3 illustrates the accuracy of SBR models as the session length varies. (i) The accuracy of all models decreases as the session length increases since user intent can vary. For example, in Diginetica, CORE [15] shows 13.2% performance drop for long sessions ($|s| \ge 10$) compared to short sessions (|s| < 5). (ii) MiaSRec shows the highest performance in most cases, regardless of session length, and a comparatively modest performance drop, which indicates that MiaSRec effectively captures various user intents. Particularly, for long sessions ($|s| \ge 10$), it shows significant improvements over CORE, *e.g.*, 6.03% for R@20 in Diginetica.

Ablation study. Table 3 shows the ablation study of MiaSRec for additional embeddings and multi-intent selection. (i) Both frequency and position embeddings have a significant impact on performance. This indicates that reflecting the importance of each item through sequential and occurrence information is effective in improving performance. (ii) It is always better to use multiple representations than a single representation. In particular, MiaSRec shows up to 2.71% improvements in R@20 compared to single representation variants using mean vector (\mathbf{o}_m) and last item vector ($\mathbf{o}_{|s|}$). (iii) The intent selection method of MiaSRec is more effective than heuristic

Table 3: Ablation study for MiaSRec. "PE" and "FE" mean position and frequency embedding. "mean" indicates only using the mean vector (o_m) , and "last k" indicates selecting the last k representations as session representations.

Model	Diginetica		Retail	rocket	Yoochoose					
	R@20	M@20	R@20	M@20	R@20	M@20				
MiaSRec	53.54	19.47	63.37	39.23	65.37	30.74				
Variants for embedding layers										
w/o PE (a _i)	51.36	18.15	61.06	37.51	61.13	26.51				
w/o FE (r _{fi})	53.48	19.23	63.28	38.92	65.15	29.91				
Variants for intent selection										
mean (\mathbf{o}_m)	52.73	18.66	61.70	38.25	64.71	29.84				
last 1 ($\mathbf{o}_{ s }$)	52.34	18.37	61.90	37.79	64.10	30.05				
last 3 ($\mathbf{o}_{ s -2: s }$)	53.08	19.20	62.07	38.68	64.77	30.34				
last 5 ($\mathbf{o}_{ s -4: s }$)	53.38	19.29	63.01	38.90	65.13	30.51				

multiple representation variants. It outperforms other methods that adopt multiple representations of the last few item vectors, suggesting the importance of extracting the representation dynamically over the session.

4 CONCLUSION

This paper proposed a novel SBR model, **MiaSRec**, which exploits various user intents in a session. Unlike previous SBR models that only use a single session representation, MiaSRec fully captures a variety of intents utilizing each session item using multiple representations and dynamically selects more important ones using the intent selection layer. It then effectively decodes and aggregates the multiple representations and provides recommendations that reflect the various intents. Extensive experiments showed that MiaSRec outperformed twelve baseline models on six benchmark datasets.

ACKNOWLEDGEMENT

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant and National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2019-0-00421, 2022-0-00680-003, IITP-2024-2020-0-01821, and NRF-2018R1A5A1060031). Multi-intent-aware Session-based Recommendation

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

REFERENCES

- Ching-Wei Chen, Paul Lamere, Markus Schedl, and Hamed Zamani. 2018. Recsys challenge 2018: automatic music playlist continuation. In *RecSys.* 527–528.
- [2] Gaode Chen, Xinghua Zhang, Yanyan Zhao, Cong Xue, and Ji Xiang. 2021. Exploring Periodicity and Interactivity in Multi-Interest Framework for Sequential Recommendation. In IJCAI. 1426–1433.
- [3] Jinpeng Chen, Yuan Cao, Fan Zhang, Pengfei Sun, and Kaimin Wei. 2022. Sequential Intention-aware Recommender based on User Interaction Graph. In *ICMR*. 118-126.
- [4] Tianwen Chen and Raymond Chi-Wing Wong. 2020. Handling Information Loss of Graph Neural Networks for Session-based Recommendation. In KDD. 1172–1180.
- [5] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep Neural Networks for YouTube Recommendations. In *RecSys.* 191–198.
- [6] Carlos Alberto Gomez-Uribe and Neil Hunt. 2016. The Netflix Recommender System: Algorithms, Business Value, and Innovation. ACM Trans. Manag. Inf. Syst. 6, 4 (2016), 13:1-13:19.
- [7] Jiayan Guo, Yaming Yang, Xiangchen Song, Yuan Zhang, Yujing Wang, Jing Bai, and Yan Zhang. 2022. Learning Multi-granularity Consecutive User Intent Unit for Session-based Recommendation. In WSDM. 343–352.
- [8] Jiayan Guo, Peiyan Zhang, Chaozhuo Li, Xing Xie, Yan Zhang, and Sunghun Kim. 2022. Evolutionary Preference Learning via Graph Nested GRU ODE for Session-based Recommendation. In CIKM. 624–634.
- [9] Priyanka Gupta, Diksha Garg, Pankaj Malhotra, Lovekesh Vig, and Gautam M. Shroff. 2019. NISER: Normalized Item and Session Representations with Graph Neural Networks. *CoRR* (2019).
- [10] Qilong Han, Chi Zhang, Rui Chen, Riwei Lai, Hongtao Song, and Li Li. 2022. Multi-Faceted Global Item Relation Learning for Session-Based Recommendation. In SIGIR. 1705–1715.
- [11] Balázs Hidasi and Alexandros Karatzoglou. 2018. Recurrent Neural Networks with Top-k Gains for Session-based Recommendations. In CIKM. 843–852.
- [12] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. Session-based Recommendations with Recurrent Neural Networks. In *ICLR*.
- [13] Balázs Hidasi, Massimo Quadrana, Alexandros Karatzoglou, and Domonkos Tikk. 2016. Parallel Recurrent Neural Network Architectures for Feature-rich Sessionbased Recommendations. In *RecSys.* 241–248.
- [14] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the Knowledge in a Neural Network. CoRR (2015).
- [15] Yupeng Hou, Binbin Hu, Zhiqiang Zhang, and Wayne Xin Zhao. 2022. CORE: Simple and Effective Session-based Recommendation within Consistent Representation Space. In SIGIR. 1796–1801.
- [16] Dietmar Jannach, Malte Ludewig, and Lukas Lerche. 2017. Session-based item recommendation in e-commerce: on short-term intents, reminders, trends and discounts. User Model. User Adapt. Interact. 27, 3-5 (2017), 351–392.
- [17] Wang-Cheng Kang and Julian J. McAuley. 2018. Self-Attentive Sequential Recommendation. In *ICDM*. 197–206.
- [18] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- [19] Haoyang Li, Xin Wang, Ziwei Zhang, Jianxin Ma, Peng Cui, and Wenwu Zhu. 2022. Intention-Aware Sequential Recommendation With Structured Intent Transition. *IEEE Trans. Knowl. Data Eng.* 34, 11 (2022), 5403–5414.
- [20] Jing Li, Pengjie Ren, Zhumin Chen, Zhaochun Ren, Tao Lian, and Jun Ma. 2017. Neural Attentive Session-based Recommendation. In CIKM. 1419–1428.
- [21] Greg Linden, Brent Smith, and Jeremy York. 2003. Amazon.com Recommendations: Item-to-Item Collaborative Filtering. *IEEE Internet Comput.* 7, 1 (2003), 76–80.
- [22] Malte Ludewig and Dietmar Jannach. 2018. Evaluation of session-based recommendation algorithms. User Model. User Adapt. Interact. 28, 4-5 (2018), 331–390.

- [23] Malte Ludewig, Noemi Mauro, Sara Latifi, and Dietmar Jannach. 2021. Empirical analysis of session-based recommendation algorithms. User Model. User Adapt. Interact. 31, 1 (2021), 149–181.
- [24] André F. T. Martins and Ramón Fernandez Astudillo. 2016. From Softmax to Sparsemax: A Sparse Model of Attention and Multi-Label Classification. In ICML (JMLR Workshop and Conference Proceedings, Vol. 48). 1614–1623.
- [25] Zhiqiang Pan, Fei Cai, Wanyu Chen, Honghui Chen, and Maarten de Rijke. 2020. Star Graph Neural Networks for Session-based Recommendation. In CIKM. 1195– 1204.
- [26] Ben Peters, Vlad Niculae, and André F. T. Martins. 2019. Sparse Sequence-to-Sequence Models. In ACL. 1504–1519.
- [27] Qi Shen, Shixuan Zhu, Yitong Pang, Yiming Zhang, and Zhihua Wei. 2021. Temporal aware Multi-Interest Graph Neural Network For Session-based Recommendation. CoRR (2021).
- [28] Qiaoyu Tan, Jianwei Zhang, Jiangchao Yao, Ninghao Liu, Jingren Zhou, Hongxia Yang, and Xia Hu. 2021. Sparse-Interest Network for Sequential Recommendation. In WSDM. 598–606.
- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NeurIPS*. 5998–6008.
- [30] Shoujin Wang, Longbing Cao, Yan Wang, Quan Z. Sheng, Mehmet A. Orgun, and Defu Lian. 2022. A Survey on Session-based Recommender Systems. ACM Comput. Surv. 54, 7 (2022), 154:1–154:38.
- [31] Shoujin Wang, Liang Hu, Yan Wang, Quan Z. Sheng, Mehmet A. Orgun, and Longbing Cao. 2019. Modeling Multi-Purpose Sessions for Next-Item Recommendations via Mixture-Channel Purpose Routing Networks. In IJCAI. 3771–3777.
- [32] Ziyang Wang, Wei Wei, Gao Cong, Xiao-Li Li, Xianling Mao, and Minghui Qiu. 2020. Global Context Enhanced Graph Neural Networks for Session-based Recommendation. In SIGIR. 169–178.
- [33] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. 2019. Session-Based Recommendation with Graph Neural Networks. In AAAI. 346–353.
- [34] Xin Xia, Hongzhi Yin, Junliang Yu, Yingxia Shao, and Lizhen Cui. 2021. Self-Supervised Graph Co-Training for Session-based Recommendation. In CIKM. 2180–2190.
- [35] Jiahao Yuan, Zihan Song, Mingyou Sun, Xiaoling Wang, and Wayne Xin Zhao. 2021. Dual Sparse Attention Network For Session-based Recommendation. In AAAI. 4635–4643.
- [36] Peiyan Zhang, Jiayan Guo, Chaozhuo Li, Yueqi Xie, Jaeboum Kim, Yan Zhang, Xing Xie, Haohan Wang, and Sunghun Kim. 2023. Efficiently Leveraging Multilevel User Intent for Session-based Recommendation via Atten-Mixer Network. In WSDM. 168–176.
- [37] Shengyu Zhang, Lingxiao Yang, Dong Yao, Yujie Lu, Fuli Feng, Zhou Zhao, Tat-Seng Chua, and Fei Wu. 2022. Re4: Learning to Re-contrast, Re-attend, Re-construct for Multi-interest Recommendation. In WWW. 2216–2226.
- [38] Wayne Xin Zhao, Yupeng Hou, Xingyu Pan, Chen Yang, Zeyu Zhang, Zihan Lin, Jingsen Zhang, Shuqing Bian, Jiakai Tang, Wenqi Sun, Yushuo Chen, Lanling Xu, Gaowei Zhang, Zhen Tian, Changxin Tian, Shanlei Mu, Xinyan Fan, Xu Chen, and Ji-Rong Wen. 2022. RecBole 2.0: Towards a More Up-to-Date Recommendation Library. In CIKM. 4722–4726.
- [39] Wayne Xin Zhao, Shanlei Mu, Yupeng Hou, Zihan Lin, Yushuo Chen, Xingyu Pan, Kaiyuan Li, Yujie Lu, Hui Wang, Changxin Tian, Yingqian Min, Zhichao Feng, Xinyan Fan, Xu Chen, Pengfei Wang, Wendi Ji, Yaliang Li, Xiaoling Wang, and Ji-Rong Wen. 2021. RecBole: Towards a Unified, Comprehensive and Efficient Framework for Recommendation Algorithms. In CIKM. 4653–4664.
- [40] Nengjun Zhu, Jian Cao, Yanchi Liu, Yang Yang, Haochao Ying, and Hui Xiong. 2020. Sequential Modeling of Hierarchical User Intention and Preference for Next-item Recommendation. In WSDM. 807–815.