Integrating Prior Knowledge from Meta-Learning and Large Language Models for Cold-Start Recommendation

Li, Yu
Graduate School of Economics, Kyushu University

Liu, Yixiao
Graduate School of Economics, Kyushu University

Furukawa, Tetsuya
Graduate School of Economics, Kyushu University

https://doi.org/10.5109/7157997

出版情報：Proceedings of International Exchange and Innovation Conference on Engineering & Sciences (IEICES). 9, pp.327-333, 2023-10-19. 九州大学大学院総合理工学府
バージョン：

Integrating Prior Knowledge from Meta-Learning and Large Language Models for Cold-Start Recommendation

Yu Li[1*], Yixiao Liu[1], Tetsuya Furukawa[1]
[1]Graduate School of Economics, Kyushu University, Japan
*Corresponding author email: li.yu.652@s.kyushu-u.ac.jp

Abstract: *Recommender Systems (RSs) often suffer from the cold-start problem, which leads to poor recommendations when dealing with new items or users with limited interaction histories. Recent research has leveraged meta-learning to overcome this issue by extracting informative prior knowledge from the interaction histories of warm users. This prior knowledge helps recommenders identify the preferences of new users with only a few interactions. In this work, we tackle the cold-start problem using two types of prior knowledge: recommendation-specific knowledge derived from meta-learning and general knowledge obtained from Large Language Models (LLMs) pre-trained on vast amounts of internet data. Specifically, our approach involves encoding item and user textual information using LLMs, allowing us to effectively incorporate general prior knowledge. Additionally, we train the recommender using meta-learning techniques to acquire recommendation-specific prior knowledge. By combining these two types of knowledge, our method demonstrates promising performance under various cold-start settings.*

Keywords: Recommender Systems; Large Language Models; Meta-Learning; Prior Knowledge

1. INTRODUCTION

Recommender Systems (RSs) play a vital role in enhancing the online user experience and addressing the challenge of information overload [1,2]. By offering personalized suggestions for candidate items, RS tailors recommendations to match individual user preferences across various online application domains [3,4]. For instance, in movie recommendations, RS utilizes movie content and users' past interaction histories to suggest the latest movies that align with their interests, enabling users to discover new movies they may enjoy. The core principle behind RS involves leveraging user-item interactions and associated side information, particularly textual data such as item titles, descriptions, user profiles, and item reviews, to predict matching scores between users and items [5]. This textual side information contains valuable knowledge that can assist in calculating matching scores, providing excellent opportunities to reveal user preferences and advance RS [6]. However, RSs often suffer from the cold-start problem, i.e., dealing with new items or users with limited interaction histories leads to poor recommendation [7].

To tackle the cold-start issue in recommendation systems, researchers have drawn inspiration from few-shot learning and integrated meta-learning, notably Model-Agnostic Meta-Learning (MAML) [8]. MAML's core concept involves acquiring knowledge from past recommendation tasks and leveraging it to expedite the learning of a new user's preferences through minimal interactions. As an illustration, Lee et al. propose a MAML-based recommendation system which learns a neural network via MAML for quickly adapt its parameter for a new user or a new item with few interaction histories. Moreover, they acquired a universal parameter initialization for the preference estimator as prior knowledge [9]. Dong et al. design two types of memories: i.e., task-specific memory and feature-specific memory which further facilitate the cold-start recommendation [10]. Although the meta-learning based method has shown promise in improving cold-start recommendation quality by utilizing prior knowledge from previous tasks, it may not be sufficient. To enhance the recommendation quality even further, we propose incorporating prior knowledge learned from internet data.

Recently, advanced Natural Language Processing (NLP) techniques, such as Large Language Models (LLMs), have demonstrated remarkable achievements in NLP domain because of their formidable ability to model text [11,12,13]. Unlike conventional models, which are trained directly on labeled data for particular tasks, LLMs undergo pre-training on extensive volumes of textual data from various sources like articles, books, websites, and publicly accessible written materials [12]. Consequently, LLMs provides a prior knowledge which is general for different down-stream tasks [14,15]. Moreover, since LLMs learns from a large amount of text data, LLMs have greater potential in modeling complicated contextual information.

In this paper, we propose LLM-MetaRec to enhance cold-start recommendation quality using two types of prior knowledge: recommendation-specific prior knowledge learned from meta-learning, and general prior knowledge from the LLMs. Our approach involves encoding both item and user information using the LLMs, allowing the downstream recommender neural network to leverage LLMs embeddings and effectively incorporate general prior knowledge. Furthermore, we train the recommendation neural network with meta-learning to enable the recommender to acquire recommendation-specific prior knowledge from previous tasks. By combining these two sources of prior knowledge, the recommender can infer user preferences even with limited item ratings. In experiments, we present evidence of LLM-MetaRec's superiority over baseline techniques on two benchmark datasets under three distinct cold-start scenarios, i.e., recommending existing items for cold users, recommending cold items for existing users, and recommending cold items for cold users.

2. RELATED WORK

2.1 Recommender system

To address the information overload problem, recommender systems have emerged as a crucial tool in various online applications by providing personalized content and services to individual users [16,17]. Typically, most existing recommendation approaches can fall into two main categories: Collaborative Filtering (CF) and Content-based recommendation. As the most common technique, CF-based recommendation methods aim to find similar behaviors patterns of users to predict the likelihood of future interactions [18], which can be achieved by utilizing the historical interaction behaviors between users and items, such as purchase history or rating data. For example, as one of the most popular CF methods, Matrix Factorization (MF) is introduced to modelling the relationship between users and items with user-item interaction histories [19]. In other words, unique identities of users and items, i.e., discrete IDs are encoded to continue embedding vectors so that the matching score can be calculated for recommendations [20,21]. Content-based recommendation methods generally take advantage of additional knowledge about users or items, such as user demographics or item descriptions, to enhance user and item representations for improving recommendation performance [22]. For example, LLAE approach comprises an encoder responsible for transforming the user behavior space into the user attribute space, alongside a decoder tasked with reconstructing the user behavior using the user attribute information [23]. The user and item attribute can build connections beween new and old users by inferring similarities among them. However, only leveraging auxiliary information is not sufficient to predict user preference, i.e., users with similar attributes may have different preferences for items [9].

2.2 Meta-learning for cold-start recommendation

Meta-learning, a promising and innovative approach, aims to amplify the adaptability of models to novel tasks by distilling task-related insights from supplementary data [24]. Researchers have recognized the potential of meta-learning techniques to tackle this cold-start recommendation problem effectively. For example, Vartak et al. model the cold-start recommendation as a meta-learning problem and propose to use Deep Neural Network (DNN) as the recommender [25]. s^2 Meta proposes an novel model updating strategy and a early-stop strategy. Wit the modules, the recommender can quickly converge to a good sub-optimal solution [26]. MeLU , another notable contribution in this research domain, combines cold-start recommendation with a popular meta-learning method, i.e., Model-Agnostic Meta-Learning (MAML). Moreover, MeLU learns a universal parameter initialization for the estimator, which serves as prior knowledge, aiding in its rapid adaptation to new tasks [9]. Similarly, the MAMO method also capitalizes on attribute information of user and item. It propose a task-specific memory for quickly adapt model for a new user or a new item. Additionally, the specific memories are used to facilitate personalized parameter initialization [10].

3. METHOD

In this section, we introduce the details of LLM-empowered meta-learning based recommendation method (LLM-MetaRec). Firstly, we present the general user preference estimator. Following this, we elaborate on integrating LLMs into the estimator, aiming to enhance recommendations by leveraging general prior knowledge from the LLMs. Finally, the estimator is trained with meta-learning to attain recommendation-specific prior knowledge.

3.1 General user preference estimator

The general framework of the user preference estimator, employed in numerous contemporary methods [9], is illustrated in Figure 1. This framework comprises two fundamental components: an embedding module, designed to acquire embeddings of items from their textual information, and a rating prediction module, which is implemented to compute personalized rating scores for item ranking.

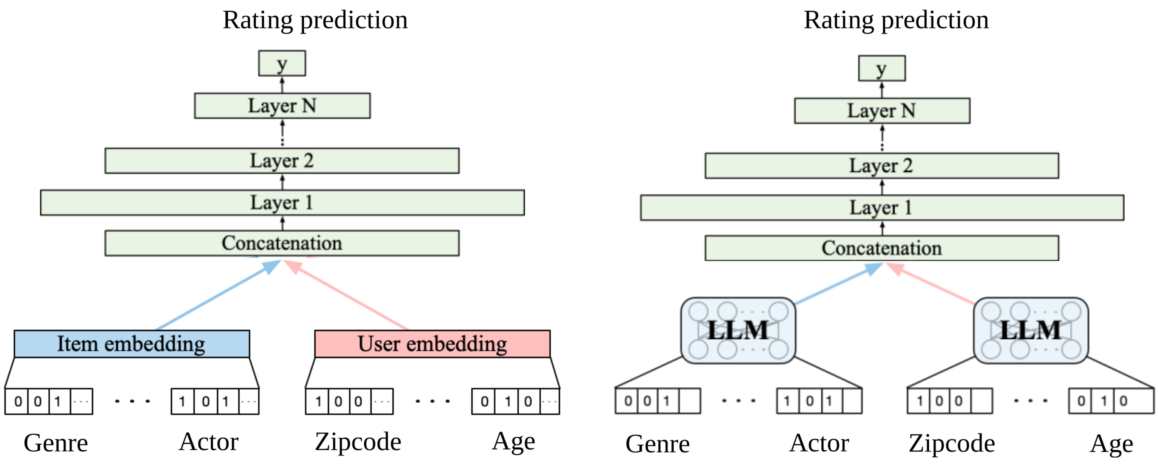

Figure 1. A General user preference estimator (left) and LLM-empowered user preference estimator (right)

Embedding module. The embedding module plays a critical role in learning representations of users and items within the context of recommendation systems. Since the attributes associated with users and items often possess distinct categorical values, it is imperative to employ two distinct strategies for learning embeddings of user and item. To illustrate the process for item embedding, let's consider the specific item v as an example. When the attribute is categorical, e..g, rating, the corresponding embedding is randomly initialized. Subsequently, these attribute-specific embeddings are concatenated together to generate the initial item embedding. Formally, assuming there are *m* attributes $c_v^1, c_v^2, ..., c_v^m$ for each item, and the procedure for item embedding can be described as follows:

$$e_v^1 = [E_1 c_v^1 \oplus E_2 c_v^2 \oplus \ldots, \oplus E_m c_v^m] \quad (1)$$

where $[\cdot \oplus \cdot]$ is the symbolic representation of the concatenation operation, and E_m is the symbolic embedding matrix. When the attributes could have multiple possible categories, such as actor, the attributes are initially transformed into a multi-hot representation, where multiple elements are set to '1' to indicate the presence of various categories and '0' otherwise. Following this, a low-dimensional representation of the multi-hot vector is generated using a feed-forward network with a single layer as below:

$$e_v^2 = (W\boldsymbol{e_i} + b) \quad (2)$$

where W denotes weight matrix, $\boldsymbol{e_i}$ represents the multi-hot representation of c_v^i and b denotes bias vector. Therefore, we concatenate the all the embeddings to get e_v, i.e., $e_v = [e_v^1 \oplus e_v^2]$.

Recommendation module. We concatenate e_u and e_v, and then feed the concatenated embedding to the downstream neural network to generate a lower-dimensional representation. The final layer of the recommendation module is responsible for computing the preference rating $\hat{r}_u^v$, which can be expressed as follows:

$$\hat{r}_u^v = F_{\theta_r}([e_u \oplus e_v]) \quad (3)$$

where the neural network denoted by F is equipped with fully connected components and corresponding parameters θ_r.

3.2 LLM empowered user preference estimator

In this section, we present the LLM-MetaRec framework, illustrated in Figure 1. The embedding module is instantiated using general prior knowledge from pre-trained large language models (LLMs). This allows us to capture contexts within user attributes, such as age and occupation, and item attributes, including genre, writers, directors, actors, and publication year. To prepare the data for the LLMs input, we start by converting all attributes into a textual format. For example, "*The title of the movie is {row['title']}. This movie is released at {row['released']}. The genre are {row['genre']}. The director of this movie is {row['director']}. The writer of this movie is {row['writer']}. The actors of this movie are {row['actors']}, ...*". The LLMs takes the user and item textual descriptions and converts them into embeddings. We concatenate the user and item embeddings to form a unified representation. By leveraging this combined embedding, the downstream neural network, which consists of multiple layers, can effectively incorporate prior knowledge from the internet dataset for recommendation purposes. The reason is that the LLM will generate item embeddings that group items with implicit relationships. For instance, if two movies are directed by two directors who share similar styles, these two movies will likely be grouped closely together in the embedding space.

3.3 Model Training

Following [9], we employ meta-learning to learn a recommendation specific prior knowledge from a labeled training dataset. With the learned prior knowledge, a model can quickly recognize the preference of a new user with a few interactions with this user. Specifically, we utilize the MAML (Model-Agnostic Meta-Learning) method, a popular meta-learning approach. MAML is designed to learn effective initial neural network weights that can be fine-tuned with only a few labeled data points. With MAML, the recommendation specific prior knowledge is stored in the initial neural network weights.

MAML employs an inner- and outer-loop learning strategy. For the training, we divide a labeled training dataset into a support set and a query set. In the inner-loop, the neural network (with its initial weights) is fine-tuned using a small support set containing a few labeled data points. In the outer-loop, the initial weights are updated based on the test loss calculated with the fine-tuned model on a query set. This learning strategy encourages the model to learn effective initial weights, which can then be fine-tuned with a small support set, resulting in minimal loss on the query set—data points without labels during the test phase. Figure 2 visualizes the learning strategy, which includes an example support set and an example query set.

In the context of recommendation tasks, we use an L1-norm loss for prediction since the model is tasked with predicting user ratings (ranging from 1 to 5) for recommended items. The parameters of the LLM are frozen to prevent forgetting prior knowledge gained from internet datasets when exposed to learning signals from a specific task.

4. EXPERIMENT

4.1 Experimental settings

The proposed method undergoes validation across four distinct scenarios involving both non-cold and cold settings: a) recommending warm item to warm user; b) recommending warm item to new user; c) recommending new item to warm user; and d) recommending new item to new user.

Dataset description. The present study conducts experiments on two distinct datasets: MovieLens-1M and Bookcrossing. Both datasets are well-established sources in the field of recommendation systems research, encompassing explicit feedback data, as well as fundamental user and item attribute information. A comprehensive overview of the key statistics pertaining to these two datasets can be found in Table 1.

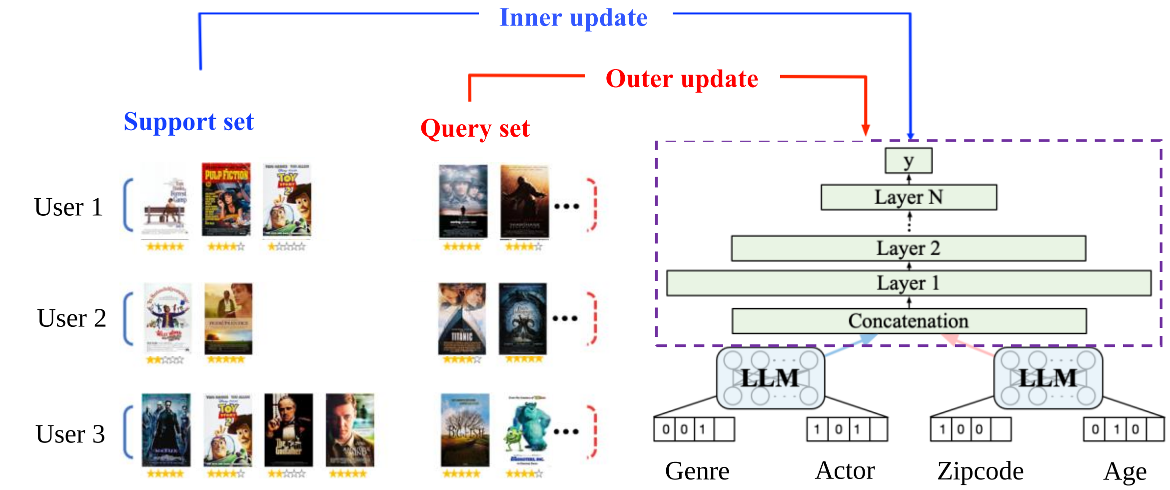

Figure 2. The diagram of MAML and LLM based approach for cold-start recommendation.

Table 1. Summary of the MovieLens-1M and Bookcrossing datasets.

Dataset	MovieLens-1M	Bookcrossing
Number of users	6,040	278,858
Number of items	3,706	271,379
Number of ratings	1,000,209	1,149,780
Density	4.468 %	0.0015 %
User attributes	Gender, Age, Occupation, Zip code	Age, Location
Item attributes	Publication year, Rate, Genre	Publication year, Author, Publisher

Baseline methods. To validate the effectiveness of LLM-MetaRec, we have selected two distinct categories of baseline approaches for comparison purposes. (1) Traditional recommendation methods: Specifically, we include the widely recognized and extensively utilized PPR [27] and Wide & Deep [28]. These conventional methods have been extensively employed in the field of recommendation systems and serve as established benchmarks for performance evaluation. (2) Meta-learning based recommendation methods: This category encompasses innovative approaches such as s^2 Meta [26], MeLU [9], and MAMO [10], which leverage meta-learning algorithms to address the intricate cold-start problem.

Evaluation metrics. In this paper, the evaluation of recommendation methods hinges upon two key performance metrics: prediction accuracy and ranking correctness. To quantify these aspects rigorously, we adopt two metric. Mean Absolute Error (MAE) evaluates the accuracy of rating prediction, and a lower MAE value indicates that the model's predictions are closer to the actual ratings provided by users, signifying better overall performance in predicting user preferences. Normalized Discounted Cumulative Gain @N (nDCG@N) evaluates the preference ordering performance on observed query sets. In simple terms, it assesses how well the model ranks and prioritizes the recommended items based on user preferences. A higher nDCG@N value indicates better performance in suggesting items that align with users' interests, leading to more relevant and satisfying recommendations. In this paper, we utilize the nDCG with N=3 as in MAMO [10].

4.2 Dataset preprocessing

We conduct preprocessing on the MovieLens-1M and Bookcrossing datasets in the subsequent manner:

- MovieLens-1M: This dataset exhibits a well-balanced distribution of interactions, i.e., the dataset have an average number of interactions. To conduct experiments, we adopted the methodology described in [9]. Specifically, we randomly selected 80% of the users from the dataset to serve as training users, representing existing users with prior interactions. The remaining 20% of users were designated as test users, representing cold users with limited or no interaction history. To ensure rigorous evaluation, A 5-fold cross-validation is performed on the partitioned dataset. Additionally, in order to create a clear demarcation between existing and cold items, we ordered the movies based on their release time and identified the top 80% as existing items, while the bottom 20% were categorized as cold items.
- Bookcrossing: The presence of evidently misleading values is evident in this dataset. For instance, the user age field exhibited unrealistic values, ranging from 0 to 237, which clearly don't align with real-world scenarios. To address this issue, we took a practical approach and restricted the user age range to a more reasonable span of 5 to 90 years. Following the guidelines presented in [10], we partitioned the dataset for training and testing purposes. Specifically, the warm/cold user split is 9/1, representing cold users who have limited or no previous interactions in the system. This division allowed us to evaluate the performance of our approach effectively. To assess the robustness and generalization capabilities of our method, we employed 10-fold cross-validation on the refined dataset. The 10-fold cross-validation extract 10 subsets from the original dataset and performs training and test on the 10 subsets. We average the results over the results from the 10 subsets. Regarding the items, we categorized them into warm and cold items based on their popularity. Items with five or

more ratings were considered warm items, indicating that they had received a substantial level of user engagement. On the other hand, items with fewer than five ratings were categorized as cold items, signifying their limited exposure in the system.

Table 2. Comparison of the performance of different recommendation methods.

Type	Method	MovieLens-1M		Bookcrossing	
		MAE	nDCG@3	MAE	nDCG@3
Recommendation on existing items for existing users	PPR	**0.1820**	**0.9831**	3.8092	0.8494
	Wide & Deep	0.9047	<u>0.9117</u>	1.6206	0.9172
	s^2Meta	0.7567	0.8870	1.5147	0.8781
	MeLU-1	0.7661	0.8904	**0.7799**	**0.9572**
	MeLU-5	0.7567	0.8919	<u>0.7955</u>	<u>0.9552</u>
	MAMO	0.8725	0.8866	1.4879	0.8879
	LLM-MetaRec	0.7362	0.8814	1.2535	0.8638
Recommendation on existing items for cold users	PPR	1.0748	0.8468	3.8430	0.8434
	Wide & Deep	1.0694	0.8639	2.0457	0.8515
	s^2Meta	0.8443	0.8401	1.8738	0.8490
	MeLU-1	0.7884	0.8810	1.8701	0.8527
	MeLU-5	<u>0.7854</u>	<u>0.8812</u>	1.8767	0.8532
	MAMO	0.8967	0.8799	<u>1.6217</u>	<u>0.8565</u>
	LLM-MetaRec	**0.7548**	**0.8823**	**1.2394**	**0.8602**
Recommendation on cold items for existing users	PPR	1.2441	0.7632	3.6821	0.8367
	Wide & Deep	1.2655	0.7721	2.2648	0.8437
	s^2Meta	**0.8860**	**0.8323**	1.9767	<u>0.8463</u>
	MeLU-1	0.9361	0.7990	2.1047	0.8441
	MeLU-5	0.9275	0.8005	2.1236	0.8440
	MAMO	0.9306	<u>0.8315</u>	<u>1.7379</u>	0.8402
	LLM-MetaRec	<u>0.8901</u>	0.8248	**1.2584**	**0.8571**
Recommendation on cold items for cold users	PPR	1.2596	0.7634	3.7046	0.8381
	Wide & Deep	1.3114	0.7874	2.3088	0.8405
	s^2Meta	0.9130	0.8148	2.0921	0.8422
	MeLU-1	0.9299	0.8011	2.1475	0.8410
	MeLU-5	0.9235	0.7752	2.1721	<u>0.8422</u>
	MAMO	<u>0.8894</u>	**0.8709**	<u>1.8188</u>	0.8384
	LLM-MetaRec	**0.8832**	<u>0.8311</u>	**1.2703**	**0.8511**

4.3 Result and analysis

In this experiment, we comprehensively evaluate the performance of all methods on two widely used datasets across four distinct recommendation settings. These settings include recommending existing items for existing users, recommending existing items for cold users, recommending cold items for existing users, and recommending cold items for cold users. Table 2 shows the test results of of all the baseline methods on the 2 datasets under various cold-start setting. We emphasize the best results with bold fonts and the second-best results with underlining. This analysis provides a comprehensive comparison of the various approaches, allowing us to gain insights into their effectiveness in addressing different cold-start scenarios. From Table 2, we have following observations:

- In the non-cold-start scenario, the method named PPR achieves the best results, attaining top-notch scores in MAE and nDCG@3. Since this is the non-cold start scenario, the result reflects that PPR may overfit to the warm users and items. In contrast, PPR and Wide & Deep methods exhibited poor performance under the other three cold-start scenarios, indicating their unsuitability for cold-start recommendations. On the other hand, LLM-MetaRec method consistently outperforms other baseline methods in all three cold-start settings while achieving comparable performance levels. This indicates that LLM-MetaRec effectively addresses the cold-start problem and provides promising results even in situations with limited interaction data for new users/items. The findings highlight the significance of incorporating meta-learning and leveraging general knowledge from Large Language Models to enhance recommendation quality, particularly in challenging cold-start scenarios.
- Meta-learning based methods, namely s^2 Meta, MAMO, LLM-MetaRec, and MeLU, consistently demonstrate superior performance in cold scenarios compared to other approaches. The good performance results in the learning strategy of meta-learning. Specifically, in the cold-start recommendation, the model is allowed to access to a few interactions of a user, and meta-learning force the model to learn the user preference from a small support set during training phase. The alignment between the test setting and training setting significantly improve the performance of meta-learning based mehtods.
- In the cold-start setting, the combination of LLM-MetaRec consistently enhances the cold-start recommendation performance of the vanilla meta-learning based recommendation method, MeLU.

This improvement is attributed to LLM-MetaRec's ability to leverage the general prior knowledge provided by LLMs, which offer a wealth of pre-trained information derived from extensive textual data. By effectively incorporating this knowledge, LLM-MetaRec overcomes the limitations of MeLU and achieves superior results in various cold-start scenarios.

5. CONCLUSION

In this work, LLM-MetaRec significantly improve the cold-start recommendation performance of a meta-learning based recommendation based method. By leveraging both recommendation-specific and general prior knowledge, we bridge the gap between new users/items and personalized recommendations, thereby improving the overall user experience in various online application domains. We demonstrate the performance of LLM-MetaRec with two datasets, i.e., MoveiLens-1M and Bookcrossing. The experiment results show that LLM-MetaRec outperforms baseline methods in various cold-start settings.

6. REFERENCES

[1] Y. Li, T. Furukawa, Information gain based dynamic support set construction for cold-start recommendation, Journal of Intelligent Information Systems. (2023) 1-21.

[2] X. Chen, W. Fan, J. Chen, H. Liu, Z. Liu, Z. Zhang, Q. Li, Fairly adaptive negative sampling for recommendations, Proceedings of the ACM Web Conference 2023. (2023) 3723-3733.

[3] Y. Gao, T. Sheng, Y. Xiang, Y. Xiong, H. Wang, J. Zhang, Chat-rec: Towards interactive and explainable llms-augmented recommender system, arXiv preprint arXiv:2303.14524, (2023).

[4] J. Chen, L. Ma, X. Li, N. Thakurdesai, J. Xu, J. H. Cho, K. Nag, E. Korpeoglu, S. Kumar, K. Achan, Knowledge graph completion models are few-shot learners: An empirical study of relation labeling in e-commerce with llms, arXiv preprint arXiv:2305.09858, (2023).

[5] W. Fan, X. Zhao, X. Chen, J. Su, J. Gao, L. Wang, Q. Liu, Y. Wang, H. Xu, L. Chen et al., A comprehensive survey on trustworthy recommender systems, arXiv preprint arXiv:2209.10117, (2022).

[6] L. Zheng, V. Noroozi, P. S. Yu, Joint deep modeling of users and items using reviews for recommendation, Proceedings of the tenth ACM International Conference on Web Search and Data Mining, (2017) 425-434.

[7] D. K. Panda, S. Ray, Approaches and algorithms to mitigate cold start problems in recommender systems: a systematic literature review, Journal of Intelligent Information Systems. 59 (2022) 341-366.

[8] C. Finn, P. Abbeel, S. Levine, Model-agnostic meta-learning for fast adaptation of deep networks, International Conference on Machine Learning. (2017) 1126-1135.

[9] H. Lee, J. Im, S. Jang, H. Cho, S. Chung, Melu: Meta-learned user preference estimator for cold-start recommendation. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. (2019) 1073-1082.

[10] M.Dong, F.Yuan, L.Yao, X.Xu, L.Zhu, Mamo: Memory-augmented meta-optimization for cold-start recommendation, Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. (2020) 688-697.

[11] A.Vaswani, N. Shazeer, N.Parmar, J.Uszkoreit, L. Jones, A. N.Gomez, I. Polosukhin, Attention is all you need. Advances in Neural Information Processing Systems. (2017) 30.

[12] J. D. M. W. C.Kenton, L. K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, Proceedings of NaacL-HLT (2019).

[13] L. Dong, N. Yang, W. Wang, F. Wei, X. Liu, Y. Wang, J. Gao, M. Zhou, H. W. Hon, Unified language model pre-training for natural language understanding and generation, Advances in Neural Information Processing Systems. (2019) 32.

[14] X. Qiu, T. Sun, Y. Xu, Y. Shao, N. Dai, X. Huang, Pre-trained models for natural language processing: A survey, Science China Technological Sciences. 63 (2020) 1872-1897.

[15] Y. Kim, Convolutional neural networks for sentence classification, arXiv preprint arXiv: 1408.5882. (2014).

[16] J. Wu, W. Fan, J. Chen, S. Liu, Q. Li, K. Tang, Disentangled contrastive learning for social recommendation, Proceedings of the 31st ACM International Conference on Information & Knowledge Management. (2022) 4570-4574.

[17] W. Fan, X. Liu, W. Jin, X. Zhao, J. Tang, and Q. Li, "Graph trend filtering networks for recommendation," Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. (2022) 112-121.

[18] W. Fan, Y. Ma, D. Yin, J. Wang, J. Tang, and Q. Li, Deep social collaborative filtering, in Proceedings of the 13th ACM Conference on Recommender Systems. (2019) 305-313.

[19] W. Fan, Q. Li, and M. Cheng, Deep modeling of social relations for recommendation, Proceedings of the AAAI Conference on Artificial Intelligence. (2018).

[20] X. Zhao, H. Liu, W. Fan, H. Liu, J. Tang, and C. Wang, Autoloss: Automated loss function search in recommendations, Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining. (2021) 3959-3967.

[21] X. Zhao, H. Liu, W. Fan, H. Liu, J. Tang, C. Wang, M. Chen, X. Zheng, X. Liu, and X. Yang, Autoemb: Automated embedding dimensionality search in streaming recommendations, IEEE International Conference on Data Mining (ICDM). (2021) 896-905.

[22] F. Vasile, E. Smirnova, and A. Conneau, Meta-prod2vec: Product embeddings using side-information for recommendation, Proceedings of the 10th ACM conference on recommender systems. (2016) 225-232.

[23] J. Li, M. Jing, K. Lu, L. Zhu, Y. Yang, Z. Huang, From zero-shot learning to cold-start

recommendation, Proceedings of the AAAI conference on artificial intelligence. 33 (2019) 4189-4196.
[24] T. Hospedales, A. Antoniou, P. Micaelli, A. Storkey, Meta-learning in neural networks: A survey, IEEE transactions on pattern analysis and machine intelligence. 44(2021) 5149-5169.
[25] M. Vartak, A. Thiagarajan, C. Miranda, J. Bratman, H. Larochelle, A meta-learning perspective on cold-start recommendations for items, Proceedings of the Annual Conference on Neural Information Processing Systems. (2017) 6904-6914.
[26] Z. Du, X. Wang, H. Yang, J. Zhou, J. Tang, Sequential scenariospecific meta learner for online recommendation, Proceedings of the the International Conference on Knowledge Discovery & Data Mining. (2019) 2895-2904 .
[27] S. Park, W. Chu, Pairwise preference regression for cold-start recommendation, Proceedings of the 2009 ACM Conference on Recommender Systems. (2009) 21-28.
[28] H. Cheng, L. Koc, J. Harmsen, T. Shaked, T. Chandra, H. Aradhye, G. Anderson, G. Corrado, W. Chai, M. Ispir, R. Anil, Z. Haque, L. Hong, V. Jain, X. Liu, H. Shah, Wide & deep learning for recommender systems, Proceedings of the 1st Workshop on Deep Learning for Recommender Systems. (2016) 7-10.