

Active Large Language Model-Based Knowledge Distillation for Session-Based Recommendation

Yingpeng Du¹, Zhu Sun^{2*}, Ziyang Wang¹, Haoyan Chua¹, Jie Zhang¹, Yew-Soon Ong^{1,3}

¹College of Computing and Data Science, Nanyang Technological University, Singapore

²Information Systems Technology and Design, Singapore University of Technology and Design, Singapore

³A*STAR Center for Frontier AI Research, Singapore

dyp1993@pku.edu.cn, zhu_sun@sutd.edu.sg, ziyang.wang1753@e.ntu.edu.sg,

haoyan001@e.ntu.edu.sg, zhangj@ntu.edu.sg, asysong@ntu.edu.sg

Abstract

Large language models (LLMs) provide a promising way for accurate session-based recommendation (SBR), but they demand substantial computational time and memory. Knowledge distillation (KD)-based methods can alleviate these issues by transferring the knowledge to a small student, which trains a student based on the predictions of a cumbersome teacher. However, these methods encounter difficulties for *LLM-based KD in SBR*. 1) It is expensive to make LLMs predict for all instances in KD. 2) LLMs may make ineffective predictions for some instances in KD, e.g., incorrect predictions for hard instances or similar predictions as existing recommenders for easy instances. In this paper, we propose an active LLM-based KD method in SBR, contributing to sustainable AI. To efficiently distill knowledge from LLMs with limited cost, we propose to extract a small proportion of instances predicted by LLMs. Meanwhile, for a more effective distillation, we propose an active learning strategy to extract instances that are as effective as possible for KD from a theoretical view. Specifically, we first formulate gains based on potential effects (e.g., effective, similar, and incorrect predictions by LLMs) and difficulties (e.g., easy or hard to fit) of instances for KD. Then, we propose to maximize the minimal gains of distillation to find the optimal selection policy for active learning, which can largely avoid extracting ineffective instances in KD. Experiments on real-world datasets show that our method significantly outperforms state-of-the-art methods for SBR.

Introduction

Recently, large language models (LLMs) have shown the potential to equip recommender systems (RSs) with their extensive knowledge and powerful reasoning capabilities (Zhao et al. 2023). Most of the existing methods employ LLMs as recommenders based on sophisticated prompt engineering (Gao et al. 2023) and fine-tuning with domain-specific knowledge (Wu et al. 2023). However, employing LLMs as recommenders usually demands substantial computational time and memory, leading to a high latency and computation requirement during the serving time and limiting real-world applications (e.g., on-device session-based recommendation (SBR) (Xia et al. 2022)).

*Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Intuitively, existing knowledge distillation (KD)-based recommendation methods (Tang and Wang 2018; Lee et al. 2019) seem to be potential solutions to these issues, which applied a model-agnostic KD technique to reduce model size while preserving the performance by training a lightweight recommender. Specifically, they train a new compact model (i.e., student recommender) based on the predictions of a well-trained cumbersome model (i.e., teacher recommender). However, it poses significant challenges for *LLM-based KD in SBR*. Firstly, *these methods lead to extreme cost, thus restricting the efficiency of KD*. Specifically, the inference of LLMs is more expensive w.r.t. time and resources than conventional recommenders. Predicting for all instances (e.g., all sessions in SBR) with LLMs is not a sustainable solution. Secondly, *LLMs may make ineffective (e.g., similar and incorrect) predictions for some instances thus hindering the effectiveness of KD*. Specifically, LLMs may make similar predictions as conventional recommenders in some instances, therefore these instances provide limited information to KD and potentially lead to a redundant distillation. In addition, some instances are hard to correctly predict by LLMs (i.e., failing to model users' real preferences), which may mislead the student recommender.

To address the first challenge, we propose only extracting a small proportion of instances to be predicted by the LLM teacher. Compared to most of the existing KD methods, our strategy does not demand LLM predicting for all instances, which can significantly reduce the cost of KD. To address the second challenge, we propose to conduct active learning to elicit effective LLM-based KD. Specifically, we first define the difficulty of instances based on whether their embedded knowledge can be easily learned by the recommender. Typically, an easy instance may be less informative (e.g., low gain) for LLM-based KD, because its knowledge is likely to be already captured by the student recommender. Although a hard instance may contribute informative knowledge, such an instance may also easily lead to an incorrect prediction (e.g., negative gain) by LLMs when it contains much noise. Meanwhile, we also identify the three potential effects of instances when they are effectively, similarly, and incorrectly predicted by the LLM in ranking tasks, abbreviated as effective/similar/incorrect instances. Then, we formulate their gains with consideration of both their diffi-

culties and effects. Finally, we maximize the minimal gains of distillation to collect effective instances by finding the optimal selection policy for KD. We theoretically and empirically prove that the proposed active learning strategy can select effective instances for KD, that is, selecting as informative as possible instances while avoiding incorrect or similar predictions by LLMs. Our method is a tailored and sustainable solution to LLM-based KD for SBR.

Contribution. In summary, we propose an Active LLM-based Knowledge Distillation Recommendation method, named ALKDRec, as a sustainable yet effective solution to SBR. For more efficient LLM-based KD, we propose to elicit student learning from a small proportion of instances. To ensure effective LLM-based KD theoretically, we propose to maximize the minimal gains of distillation by selecting effective instances for KD. To the best of our knowledge, our method is the first trial to theoretically and practically distill knowledge from LLMs to enhance lightweight recommenders with active learning for efficiency and effectiveness purposes. Experiments on real-world datasets show that our method significantly outperforms state-of-the-art KD methods with representative recommendation backbones.

Related Work

LLM-based recommendation. Leveraging LLMs for recommendation has attained popularity recently due to their advanced capabilities (Wu et al. 2023). Typically, two paradigms have been widely applied, i.e., LLM-as-extractor and LLM-as-recommender. The former paradigm leverages the LLM as a feature extractor for downstream recommendation tasks, employing LLMs to capture and infer user profiles (Du et al. 2024a), item descriptions (Ren et al. 2024; Wei et al. 2024a), and other textual data (Liu et al. 2024a; Geng et al. 2022; Du et al. 2024b). However, LLMs in these methods are not directly designed for recommendation, which may lead to a gap between extracted knowledge and recommendation tasks. The latter paradigm adopts the LLM as a predictor to directly provide recommendations for users. They usually employ prompt techniques such as in-context learning (Hou et al. 2024; Sun et al. 2024b; Wang et al. 2024c), prompt tuning (Sun et al. 2024b; Kang et al. 2023b), and parameter fine-tuning (Shi et al. 2023; Luo et al. 2024) to trigger LLMs directly generate recommendations. However, these methods may lead to high latency and increased computation requirements during the serving time.

Knowledge distillation in recommendation. Recently, Knowledge distillation (Kang 2024) has attracted considerable attention in recommendation for low latency and decreased computation. Most of these methods first train a large teacher recommender and then transfer the knowledge from the teacher to the target compact student recommender, which mainly utilizes soft labels of teacher predictions for KD (Lee et al. 2019; Pan, He, and Yu 2019; Kang et al. 2021; Huang et al. 2023). Furthermore, advanced KD technologies are adopted for recommendation such as adversarial distillation (Chen et al. 2018; Wang et al. 2018), bidirectional distillation (Kweon, Kang, and Yu 2021; Lee and Kim 2021), privileged feature distillation (Wang et al. 2021;

Xu et al. 2020; Zhang et al. 2020; Liu et al. 2023a; Deng et al. 2023), debias distillation (Chen et al. 2023), multi-modal/model distillation (Liu et al. 2023b; Wei et al. 2024b; Kang et al. 2023a; Sun et al. 2024a), and sparsity distillation (Wang et al. 2024a), and LLM-based distillation (Wang et al. 2024b; Liu et al. 2024b; Cui et al. 2024). However, these methods are not tailored for efficient LLM-based KD due to (1) extreme costs with LLM inference (e.g., fine-tuning or distilling knowledge based on all instances.). For example, DLLM2Rec (Cui et al. 2024) is trained on all instances predicted by LLM teachers; and (2) ineffective predictions by LLMs for KD in SBR.

The Proposed Method

Problem Definition. Let $\mathcal{V} = \{v_1, v_2, \dots, v_N\}$ represent the item set with N items. Each anonymous session $s \in \mathcal{S}$ records the user’s interaction with l items, which is denoted as $s = \{v_1^s, v_2^s, \dots, v_l^s | v_i^s \in \mathcal{V}\}$. We assume to know the titles of items. Given s , the task of SBR is to rank all items that have not interacted yet and provide a top- K item list for the next item v_{l+1}^s prediction. Our goal is to design a KD method in SBR, enabling the student recommender to effectively and efficiently learn from the LLM teacher with small-scale parameters (e.g., low dimension in latent space) and achieve accurate recommendations.

Model Overview. The overall architecture of the proposed method is shown in Figure 1. Firstly, to make the *LLM teacher* align with domain-specific knowledge, we elicit it to summarize from predictions of the *conventional teacher recommender* that is trained by domain-specific user behaviors (Module 1). Secondly, to distill knowledge from the LLM teacher, we trigger the *student recommender* to learn from its prediction (Module 2). Finally, to efficiently distill knowledge from the LLM teacher, we propose an active learning strategy for KD theoretically and practically (Module 3).

Enhance LLM teacher with conventional teacher recommender

LLMs usually lack domain-specific knowledge (e.g., user behavior patterns) related to the SBR tasks, which may hinder the effective KD of LLMs. To make LLMs align with domain-specific knowledge, we elicit the LLM to make summarization with the help of the conventional teacher recommender $T_{\text{rec}}(\cdot)$ which is trained on all instances \mathcal{S} , and can capture the collaborative signals behind users’ behaviors. Then, we elicit the LLM teacher $T_{\text{llm}}(\cdot)$ to summarize users’ behavior patterns using prompt $\text{prompt}_{\text{smr}}$ as hints by analyzing predictions of the teacher recommender,

$$\text{Hints} = T_{\text{llm}}(\text{prompt}_{\text{smr}}, \{\text{rank}_{s_1}^{\text{rec}}, \dots, \text{rank}_{s_L}^{\text{rec}} | s_i \in \mathcal{S}\}),$$

where $\text{rank}_s^{\text{rec}} = T_{\text{rec}}(s)$ denotes the predictions of teacher recommender on session s . The detailed prompt $\text{prompt}_{\text{smr}}$ for summarization is shown in Appendix A. With summarized hints encapsulated with domain-specific knowledge, we can trigger the LLM teacher to make recommendations as a ranking list $\text{rank}_s^{\text{llm}} = T_{\text{llm}}(\text{Prompt}_{\text{dst}}, s, \text{Hints}, \mathcal{C}_s)$, where $\text{prompt}_{\text{dst}}$ is the prompt that triggers $T_{\text{llm}}(\cdot)$ for SBR shown in Appendix A. $\mathcal{C}_s = \text{rank}_s^{\text{rec}}[1 : \kappa]$ is the candidate set of κ (e.g., $\kappa = 50$) top items suggested by $T_{\text{rec}}(\cdot)$.

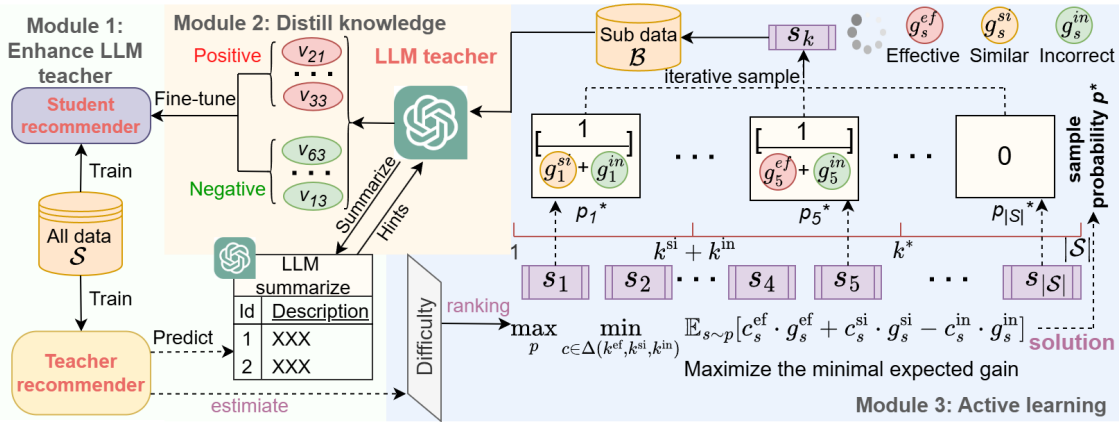


Figure 1: The architecture of the ALKDRec method.

Distill knowledge from the LLM teacher

To distill knowledge from the LLM teacher into the student, we fine-tune the student recommender $S_{\text{rec}}(\cdot)$ that is well-trained based on all instances \mathcal{S} , thus making the student recommender learn to imitate the prediction of the LLM teacher, i.e., $\text{rank}_s^{\text{llm}}$, with pair-wise loss,

$$\max \sum_{s \in \mathcal{S}} \sum_{v \in \text{rank}_s^{\text{llm}}} \alpha_v \cdot \log \sigma(S_{\text{rec}}(s, v) - S_{\text{rec}}(s, v')), \quad (1)$$

where $S_{\text{rec}}(s, v)$ denotes the predicted score for item v in session s by the student recommender, and v' denotes the negative item. α_v denotes the weight of pair-wise loss w.r.t. positive sample v , where we assign the item v to the top position with a higher weight.

Existing KD strategies which learn from teachers regarding all instances \mathcal{S} are extremely costly for LLM-based KD, as the inference of LLM teachers is much more expensive than conventional teachers. To this end, we propose to distill knowledge from a subset of all instances, i.e., $\mathcal{B} \subset \mathcal{S}$ where $|\mathcal{B}| \ll |\mathcal{S}|$. However, how to select an effective instance subset is crucial for LLM-based KD in SBR.

Active learning for LLM-based distillation

Difficulty of instances in distillation. For efficient LLM-based KD, it is intuitive to query the LLM teacher with informative instances, the patterns of which are hardly captured by conventional recommenders. To be specific, if the knowledge of an instance can be easily captured by conventional recommenders, we regard it as an **easy instance** for LLM-based KD. Hence, employing the LLM teacher distill on these easy instances is likely to lead to redundant distillation. Contrarily, the **hard instances** may provide informative knowledge for KD, but it is hard to predict correctly by the LLM teacher. To measure the difficulty of instances, we propose to assess the negative consistency between the session representation and each item within the session in the conventional teacher recommender, i.e., $df_s = -\frac{1}{|\mathcal{S}|} \sum_{v \in \mathcal{S}} \sigma(\text{Encode}(s)^T \cdot e_v)$, where $\text{Encoder}(s) \in \mathbb{R}^d$ and $e_v \in \mathbb{R}^d$ are the embeddings of session s and item v , respectively. Thus, a large negative consistency

(i.e., greater difficulty) indicates that the teacher recommender can hardly capture the knowledge.

Effects of instances in distillation. Instances may show different effects in distillation when distilling knowledge from the LLM teacher, which can be categorized as follows. **Effective instances** indicate that the LLM teacher makes accurate predictions but their patterns are hardly captured by conventional recommenders, thus usually providing informative knowledge to KD. **Similar instances** indicate that the LLM teacher and the student recommender share similar predictions, which usually provide limited informative knowledge to KD. **Incorrect instances** indicate that LLM makes incorrect predictions on users' behaviors, which may mislead the student model training.

Expected gain of distillation. To comprehensively model the gain of instances on KD, we propose considering the difficulty and effect of the instance. Effective instances usually provide informative knowledge to KD. Therefore, we assume the gain of the effective instance g_s^{ef} is proportional to the difficulty of the instance df_s . Similar instances usually provide limited informative knowledge to KD, which can improve the confidence of extracted knowledge thus showing marginal gains to KD, i.e., $g_s^{\text{si}} < g_s^{\text{ef}}$. Incorrect instances may mislead student model training and show the negative gain to KD, i.e., $-g_s^{\text{in}} < 0$ and $|g_s^{\text{in}}| < g_s^{\text{ef}}$. The detailed implementation of these gains can be found in the *Implementation Details* of Section *Experiments*.

To elicit effective LLM-based KD, we propose an active learning strategy to select instances with the maximal expected gain, given by,

$$\max_p \mathbb{E}_{s \sim p} [c_s^{\text{ef}} \cdot g_s^{\text{ef}} + c_s^{\text{si}} \cdot g_s^{\text{si}} - c_s^{\text{in}} \cdot g_s^{\text{in}}], \quad (2)$$

where $p \in \mathbb{R}^{|\mathcal{S}|}$ denotes the probability of selecting one specific instance for active KD. Without loss of generality, we rank instances based on their effective gains, i.e., $g_1^{\text{ef}} > g_2^{\text{ef}} > \dots > g_{|\mathcal{S}|}^{\text{ef}}$. In addition, $c_s^{\text{ef}} = 1$, $c_s^{\text{si}} = 1$, or $c_s^{\text{in}} = 1$ represent the indicator, indicating an effective, similar, or incorrect instance for KD, satisfying $c_s^{\text{ef}} + c_s^{\text{si}} + c_s^{\text{in}} = 1$.

In real-world scenarios, it is hard to identify the indicator value of an instance s before eliciting the LLM teacher to

predict for it, making it intractable to directly optimize the expected gain in Equation (2) for LLM-based KD. To this end, we propose maximizing the expected gain of distillation regardless of the specific cases of all instances, which can be achieved by maximizing the minimal expected gain for the instance selection inspired by (Baykal et al. 2023).

Maximize the minimal gains of distillation. As it is intractable to obtain specific values of indicator of all instances $\{c_1, \dots, c_{|\mathcal{S}|}\}$, we only assume to know the overall number for each type of instances, i.e., $\sum_s c_s^{\text{ef}} = k^{\text{ef}}$ for effective instances, $\sum_s c_s^{\text{si}} = k^{\text{si}}$ for similar instances, and $\sum_s c_s^{\text{in}} = k^{\text{in}}$ for incorrect instances, where $k^{\text{ef}} + k^{\text{si}} + k^{\text{in}} = |\mathcal{S}|$. We formulate to maximize the minimal expected gain,

$$\max_p \min_{c \in \Delta(k^{\text{ef}}, k^{\text{si}}, k^{\text{in}})} \underbrace{\mathbb{E}_{s \sim p} [c_s^{\text{ef}} \cdot g_s^{\text{ef}} + c_s^{\text{si}} \cdot g_s^{\text{si}} - c_s^{\text{in}} \cdot g_s^{\text{in}}]}_{=\sum_s p_s \cdot [c_s^{\text{ef}} \cdot g_s^{\text{ef}} + c_s^{\text{si}} \cdot g_s^{\text{si}} - c_s^{\text{in}} \cdot g_s^{\text{in}}] := z(p, c)}, \quad (3)$$

where $\Delta(k^{\text{ef}}, k^{\text{si}}, k^{\text{in}})$ denotes a set assembling all combinations of instances which satisfies $\sum_s c_s^t = k^t$ for $t \in \{\text{ef}, \text{si}, \text{in}\}$, and probability p is the active strategy for the instance selection.

To solve the max-min problem in Equation (3), we provide two theorems to prove the lower bound and upper bound of the expected gain. Then, we illustrate that the equilibrium point is the exact solution for this problem¹.

Definition 1 (Probability p^*). Suppose we have a vector $p^* = (p_1^*, \dots, p_{|\mathcal{S}|}^*)$ satisfying

$$p_s^* = \begin{cases} 1/(H_{k^*}(g_s^{\text{si}} + g_s^{\text{in}})), & \text{if } 1 \leq s \leq k^{\text{si}} + k^{\text{in}}, \\ 1/(H_{k^*}(g_s^{\text{ef}} + g_s^{\text{in}})), & \text{if } k^{\text{si}} + k^{\text{in}} < s \leq k^*, \\ 0, & \text{if } s > k^*, \end{cases}$$

where

$$H_k = \begin{cases} \sum_{s=1}^k \frac{1}{g_s^{\text{si}} + g_s^{\text{in}}}, & \text{if } 1 \leq k \leq k^{\text{si}} + k^{\text{in}}, \\ H_{k^{\text{si}} + k^{\text{in}}} + \sum_{s=1}^k \frac{1}{g_s^{\text{ef}} + g_s^{\text{in}}}, & \text{if } k^{\text{si}} + k^{\text{in}} < k \leq k^*. \end{cases}$$

We have p^* is a probability distribution as $p_s^* \geq 0$ and $\sum_{s=1}^{|\mathcal{S}|} p_s^* = \frac{1}{H_{k^*}} [\sum_{s=1}^{k^{\text{si}} + k^{\text{in}}} \frac{1}{g_s^{\text{si}} + g_s^{\text{in}}} + \sum_{s=k^{\text{si}} + k^{\text{in}} + 1}^{k^*} \frac{1}{g_s^{\text{ef}} + g_s^{\text{in}}}] = 1$.

To illustrate the property of the probability p^* , we prove its relation to a lower bound of the expected gain $z(p^*, c)$ in Equation (3), which is detailed as follows.

Theorem 1. The expected gain $z(p, c)$ has a constant lower bound with $p = p^*$, i.e.,

$$\min_{c \in \Delta(k^{\text{ef}}, k^{\text{si}}, k^{\text{in}})} \mathbb{E}_{s \sim p^*} [c_s^{\text{ef}} \cdot g_s^{\text{ef}} + c_s^{\text{si}} \cdot g_s^{\text{si}} - c_s^{\text{in}} \cdot g_s^{\text{in}}] = \Gamma(k^*)$$

where $\Gamma(k^*) = [k^{\text{ef}} + k^{\text{si}} - N + G_{k^*}]/H_{k^*}$ and

$$G_k = \begin{cases} \sum_{s=1}^k g_s^{\text{si}} / (g_s^{\text{si}} + g_s^{\text{in}}), & \text{if } 1 \leq k \leq k^{\text{si}} + k^{\text{in}}, \\ G_{k^{\text{si}} + k^{\text{in}}} + \sum_{s=1}^k g_s^{\text{ef}} / (g_s^{\text{ef}} + g_s^{\text{in}}), & \text{if } k^{\text{si}} + k^{\text{in}} < k \leq k^*. \end{cases}$$

¹Due to space limitations, we only provide the main sketch of proofs, while the detailed proofs can be found in Appendix B of our GitHub repository at <https://github.com/kk97111/ALKDRec>.

Proof sketch. Assigning $p = p^*$ in expected gain $z(p, c)$, we can reformulate the expected gain $z(p^*, c)$ as

$$z(p^*, c) = \frac{1}{H_{k^*}} \left[\sum_{s=1}^{k^{\text{si}} + k^{\text{in}}} \left(c_s^{\text{ef}} \frac{g_s^{\text{ef}} + g_s^{\text{in}}}{g_s^{\text{si}} + g_s^{\text{in}}} + c_s^{\text{si}} - \frac{g_s^{\text{in}}}{g_s^{\text{si}} + g_s^{\text{in}}} \right) + \sum_{s=k^{\text{si}} + k^{\text{in}} + 1}^{k^*} \left(c_s^{\text{ef}} + c_s^{\text{si}} \frac{g_s^{\text{si}} + g_s^{\text{in}}}{g_s^{\text{ef}} + g_s^{\text{in}}} - \frac{g_s^{\text{in}}}{g_s^{\text{ef}} + g_s^{\text{in}}} \right) \right].$$

To minimize the gain $z(p^*, c)$, we should assign effective instances ($c_s^{\text{ef}} = 1$) into $[k^{\text{si}} + k^{\text{in}}, |\mathcal{S}|]$, similar instances ($c_s^{\text{si}} = 1$) into $[1, k^{\text{si}} + k^{\text{in}}]$, and incorrect instances into rest place of $[1, |\mathcal{S}|]$. Therefore, we have the minimal gain $z(p^*, c)$ as

$$\begin{aligned} \min_{c \in \Delta(k^{\text{ef}}, k^{\text{si}}, k^{\text{in}})} z(p^*, c) &= \frac{1}{H_{k^*}} \left[0 + k^{\text{si}} + (k^* - k^{\text{si}} - k^{\text{in}}) \right. \\ &+ 0 - \sum_{s=1}^{k^{\text{si}} + k^{\text{in}}} \frac{g_s^{\text{in}}}{g_s^{\text{si}} + g_s^{\text{in}}} - \sum_{s=k^{\text{si}} + k^{\text{in}} + 1}^{k^*} \frac{g_s^{\text{in}}}{g_s^{\text{ef}} + g_s^{\text{in}}} \left. \right] \\ &= \frac{1}{H_{k^*}} [k^{\text{ef}} + k^{\text{si}} - N + G_{k^*}], \end{aligned}$$

where $\Gamma(k^*) = [k^{\text{ef}} + k^{\text{si}} - N + G_{k^*}]/H_{k^*}$. \square

To explore the gain $z(p, c)$ w.r.t. the combination c , we make a specific definition of \hat{c} as follows.

Definition 2 (Combination \hat{c}^{ef} , \hat{c}^{si} , and \hat{c}^{in}). Suppose we have three vectors $\hat{c}^t = (\hat{c}_1^t, \dots, \hat{c}_{|\mathcal{S}|}^t)$ for $t \in \{\text{ef}, \text{si}, \text{in}\}$ as

$$\hat{c}_s^{\text{ef}} = \begin{cases} 0, & \text{if } 1 \leq s \leq k^{\text{si}} + k^{\text{in}}, \\ \frac{\Gamma(k^*) + g_s^{\text{in}}}{g_s^{\text{si}} + g_s^{\text{in}}}, & \text{if } k^{\text{si}} + k^{\text{in}} < s \leq k^*, \\ \left[k^{\text{ef}} - \sum_{s=k^{\text{si}} + k^{\text{in}}}^{k^*} \frac{\Gamma(k^*) + g_s^{\text{in}}}{g_s^{\text{si}} + g_s^{\text{in}}} \right] / (|\mathcal{S}| - k^*), & \text{if } s > k^*, \end{cases}$$

$$\hat{c}_s^{\text{si}} = \begin{cases} \frac{\Gamma(k^*) + g_s^{\text{in}}}{g_s^{\text{ef}} + g_s^{\text{in}}}, & \text{if } 1 \leq s \leq k^{\text{si}} + k^{\text{in}}, \\ 0, & \text{if } k^{\text{si}} + k^{\text{in}} < s \leq k^*, \\ \left[k^{\text{si}} - \sum_{s=1}^{k^{\text{si}} + k^{\text{in}}} \frac{\Gamma(k^*) + g_s^{\text{in}}}{g_s^{\text{ef}} + g_s^{\text{in}}} \right] / (|\mathcal{S}| - k^*), & \text{if } s > k^*, \end{cases}$$

and $\hat{c}_s^{\text{in}} = 1 - \hat{c}_s^{\text{ef}} - \hat{c}_s^{\text{si}}$.

We have $\hat{c} := (\hat{c}^{\text{ef}}, \hat{c}^{\text{si}}, \hat{c}^{\text{in}}) \in \Delta(k^{\text{ef}}, k^{\text{si}}, k^{\text{in}})$ that is a feasible combination as $\sum_{s \leq |\mathcal{S}|} \hat{c}_s^{\text{ef}} = k^{\text{ef}}$, $\sum_{s \leq |\mathcal{S}|} \hat{c}_s^{\text{si}} = k^{\text{si}}$, and $\sum_{s \leq |\mathcal{S}|} \hat{c}_s^{\text{in}} = k^{\text{in}}$ are established.

To illustrate the property of the combination \hat{c} , we prove that its relation to an upper bound of the expected gain $z(p, \hat{c})$ in Equation (3), which is detailed as follows.

Theorem 2. The expected gain $z(p, c)$ has a constant upper bound with $c = \hat{c}$, i.e.,

$$\max_p \mathbb{E}_{s \sim p} [\hat{c}_s^{\text{ef}} \cdot g_s^{\text{ef}} + \hat{c}_s^{\text{si}} \cdot g_s^{\text{si}} - \hat{c}_s^{\text{in}} \cdot g_s^{\text{in}}] = \Gamma(k^*)$$

$$k^* = \operatorname{argmax}_{s=1, \dots, |\mathcal{S}|} \frac{g_s^{\text{ef}} + k^{\text{ef}} + k^{\text{si}} - |\mathcal{S}|}{H_s}$$

where $\Gamma(k^*) = [k^{\text{ef}} + k^{\text{si}} - N + G_{k^*}]/H_{k^*}$ and $g_1 > \dots > g_{|\mathcal{S}|}$.

Proof sketch. Assign $c = \hat{c}$ in gain $z(p, c)$, we have

$$z(p, \hat{c}) = \sum_{s=1}^{k^*} p_s \Gamma(k^*) + \sum_{s=k^*}^{|\mathcal{S}|} p_s \underbrace{[\hat{c}_s^{\text{ef}}(g_s^{\text{ef}} + g_s^{\text{in}}) + \hat{c}_s^{\text{si}}(g_s^{\text{si}} + g_s^{\text{in}}) - g_s^{\text{in}}]}_{:=t(c_s^{\text{ef}}, c_s^{\text{si}}, s)}$$

Furthermore, we reformulate $t(c_s^{\text{ef}}, c_s^{\text{si}}, s)$ when $s \geq k^*$, and prove that $t(c_s^{\text{ef}}, c_s^{\text{si}}, s) \leq g_s^{\text{ef}}$ by inequality scaling.

According to the definition of k^* , we have

$$\begin{aligned} \Gamma(k^*) &\geq \Gamma(k^* + 1) = \frac{k^{\text{ef}} + k^{\text{si}} - N + G_{k^*} + g_{k^*+1}/(g_{k^*+1} + l_{k^*+1})}{H_{k^*} + 1/(g_{k^*+1} + l_{k^*+1})} \\ &\geq \frac{g_{k^*+1}/(g_{k^*+1} + l_{k^*+1})}{1/(g_{k^*+1} + l_{k^*+1})} = g_{k^*+1} \geq g_s^{\text{ef}} \text{ for } (s \geq k^* + 1) \end{aligned}$$

Therefore, we have $\Gamma(k^*) \geq g_s^{\text{ef}} \geq t(c_s^{\text{ef}}, c_s^{\text{si}}, s)$ when $s \geq k^*$.

To maximize $z(p, \hat{c})$, we apparently have $\sum_{s=1}^{k^*} p_s = 1$ and $p_s = 0$ when $s > k^*$. Therefore, we have the conclusion $\max_p \mathbb{E}_{s \sim p} [\hat{c}_s^{\text{ef}} \cdot g_s^{\text{ef}} + \hat{c}_s^{\text{si}} \cdot g_s^{\text{si}} - \hat{c}_s^{\text{in}} \cdot g_s^{\text{in}}] = \Gamma(k^*)$ \square

According to the above theorems, we can achieve the maximum of minimal expected gain according to the following theorem.

Theorem 3. The maximum of the minimal expected gain $z(p, c)$ in Equation (3) can be achieved as $\Gamma(k^*)$ in this max-min game when $p^* = (p_1^*, \dots, p_{|\mathcal{S}|}^*)$ with $k^* = \text{argmax}_{s=1, \dots, |\mathcal{S}|} [g_s^{\text{ef}} + k^{\text{ef}} + k^{\text{si}} - |\mathcal{S}|]/H_s$ as

$$\begin{aligned} \Gamma(k^*) &= \min_{c \in (k^{\text{ef}}, k^{\text{si}}, k^{\text{in}})} \mathbb{E}_{s \sim p^*} [c_s^{\text{ef}} \cdot g_s^{\text{ef}} + c_s^{\text{si}} \cdot g_s^{\text{si}} - c_s^{\text{in}} \cdot g_s^{\text{in}}] \\ &\leq \max_p \min_{c \in (k^{\text{ef}}, k^{\text{si}}, k^{\text{in}})} \mathbb{E}_{s \sim p} [c_s^{\text{ef}} \cdot g_s^{\text{ef}} + c_s^{\text{si}} \cdot g_s^{\text{si}} - c_s^{\text{in}} \cdot g_s^{\text{in}}] \\ &\leq \max_p \mathbb{E}_{s \sim p} [\hat{c}_s^{\text{ef}} \cdot g_s^{\text{ef}} + \hat{c}_s^{\text{si}} \cdot g_s^{\text{si}} - \hat{c}_s^{\text{in}} \cdot g_s^{\text{in}}] = \Gamma(k^*). \end{aligned}$$

To gain the subset $\mathcal{B} \subset \mathcal{S}$ for KD, we select instances one by one where each instance is sampled based on the active learning policy p^* in Definition 1 with k^* shown in Theorem 2, the detailed . In summary, we provide an active learning strategy (i.e., active instance selection policy p^*) that can maximize the minimal expected gain for the instance selection, contributing to the effective LLM-based KD with limited cost. Compared to (Baykal et al. 2023) which focused on binary classification tasks, the proposed active learning strategy in ALKDRec can model the complicated cases for ranking tasks of SBR, i.e., the LLM teacher may make effective, similar, and incorrect predictions for KD in SBR tasks.

Experiments and Analysis

We conduct extensive experiments to evaluate the performance of ALKDRec and answer four research questions. **RQ1:** Whether ALKDRec outperforms state-of-the-art methods for SBR w.r.t. effectiveness and efficiency? **RQ2:** Whether ALKDRec benefits from the knowledge distilled from the LLM teacher for SBR? **RQ3:** Whether ALKDRec benefits from the active learning strategy for SBR? **RQ4:** How do hyper-parameters in ALKDRec affect its performance for SBR?

Datasets. We evaluate ALKDRec and baselines on two real-world datasets, namely Hetrec2011-ML and Amazon-Games. The former dataset contains user ratings of the online movie service MovieLens. The latter dataset is the ‘Video Game’ category of the Amazon dataset, which contains users’ reviews on games. All items in these two datasets are attached with title information. For each user, we extract her interactions within one day as a session sorted by timestamps, and we also filter out sessions with less than 5 records. After data preprocessing, we have 12,323 sessions, 8,475 items, and 112,034 interactions for Hetrec2011-ML, and 20,091 sessions, 26,138 items, and 132,677 interactions for Amazon-Games.

Evaluation Protocol and Metrics. We randomly split sessions into training, validation, and test sets by 6:2:2. For evaluation phase, we adopt the widely used leave-one-out evaluation protocol. Specifically for each session, we hold out the last interacted item for testing. We adopt two widely used evaluation metrics for top- K recommendation: recall and normalized discounted cumulative gain (ndcg) (Sun et al. 2022), where K is set as 5/10 empirically. Experimental results are recorded as the average of five runs.

Backbone models and baselines Inspired by prior setting (Kang et al. 2020), we test KD methods across 3 representative backbone recommenders for SBR. **FPMC** (Rendle, Freudenthaler, and Schmidt-Thieme 2010) adopts the Markov Chain to capture sequential patterns of users’ behaviors. **STAMP** (Liu et al. 2018) proposes to model users’ current interests from the short-term memory of the last clicks. **AttMix** (Zhang et al. 2023) models multi-level reasoning over item transitions from both concept-view and instance-view. For baseline methods, we take the following state-of-the-art *agnostic* KD methods as the baselines for comparison, including **DE** (Kang et al. 2020), **FTD** (Kang et al. 2021), **HTD** (Kang et al. 2021), **unKD** (Chen et al. 2023), **DSL** (Wang et al. 2024a), and **DLLM2Rec** (Cui et al. 2024).

Implementation Details. For an effective instance, the low difficulty implies a less informative contribution to KD in SBR. Therefore, we assign the gain value of g_s^{ef} based on its difficulty with an exponential distribution $g_s^{\text{ef}} = 1/[\text{rank}(df_s)]^\mu$, where $\text{rank}(\cdot)$ denotes the ascending ranking of instance among all instances w.r.t. difficulty, and we set $\mu = 10$ empirically. For similar and incorrect instances, we assign them with lower gain values compared to effective instances, i.e., $g_s^{\text{si}} = g_s^{\text{in}} = g_s^{\text{ef}}/2$. We adopt the GPT-4-turbo-2024-04-09 as the LLM teacher to distill knowledge from 500 instances for all datasets and backbones except 750 instances for FPMC in Amazon-Games. We set the number of effective/similar/incorrect instances as 1:5:4 for all datasets by the grid search. For α_v in Equation (1), we assign 3/2/1 when the item v in ranking position $[1, 5]/(5, 15)/(15, 25]$. Following Lee et al. (2019), we set the latent dimensions for the teacher and student recommenders at 100 and 10 across all KD methods for a fair comparison. We set the learning rate as 1×10^{-3} with Adam optimizer and batch size as 1024 for all methods.

Datasets	Backbone	Metric	Student	Teacher	DE	FTD	HTD	unKD	DSL	DLLM2Rec	ALKDRec	Improve.
HeteRec2011-ML	FPMC	recall@5	0.00730	0.01339	0.00690	<u>0.00933</u>	0.00893	0.00527	0.00527	0.00933	0.01258*	34.78%
		ndcg@5	0.00379	0.00781	0.00358	0.00457	<u>0.00570</u>	0.00361	0.00356	0.00467	0.00747*	31.10%
		recall@10	0.01055	0.02191	0.01258	<u>0.01866</u>	0.01460	0.00893	0.00893	0.01623	0.02150*	15.22%
		ndcg@10	0.00415	0.01002	0.00658	<u>0.00901</u>	0.00723	0.00426	0.00436	0.00731	0.01019*	13.06%
	STAMP	recall@5	0.00609	0.00852	0.00852	0.00811	0.00609	0.00487	<u>0.00852</u>	0.00609	0.01014*	19.05%
		ndcg@5	0.00364	0.00446	<u>0.00529</u>	0.00450	0.00337	0.00309	0.00452	0.00364	0.00670*	26.76%
		recall@10	0.01177	0.01907	<u>0.01501</u>	0.01298	0.01217	0.01014	0.01298	0.01217	0.01623*	8.11%
		ndcg@10	0.00563	0.00937	<u>0.00729</u>	0.00616	0.00517	0.00509	0.00542	0.00575	0.00799*	9.58%
	AttMix	recall@5	0.00122	0.00527	0.00527	0.00649	0.00325	0.00365	0.00527	0.00527	<u>0.00609</u>	-6.25%
		ndcg@5	0.00067	0.00361	0.00306	<u>0.00332</u>	0.00222	0.00235	0.00306	0.00284	0.00370*	11.40%
		recall@10	0.00406	0.00933	0.00933	0.01014	0.00730	0.00568	0.00933	<u>0.01055</u>	0.01095*	3.84%
		ndcg@10	0.00180	0.00419	0.00409	0.00366	0.00351	0.00262	0.00406	<u>0.00448</u>	0.00547*	22.14%
Amazon-Games	FPMC	recall@5	0.03309	0.05773	0.02886	0.03260	<u>0.03459</u>	0.00921	0.03211	0.03409	0.03583*	3.60%
		ndcg@5	0.01857	0.03322	0.01471	0.01869	0.01917	0.00504	0.01802	<u>0.01954</u>	0.02188*	11.98%
		recall@10	0.05101	0.08311	0.04703	0.05126	<u>0.05200</u>	0.01643	0.04331	0.05126	0.05748*	10.53%
		ndcg@10	0.02199	0.03688	0.01938	<u>0.02251</u>	0.02199	0.00696	0.01891	0.02225	0.02610*	15.97%
	STAMP	recall@5	0.03583	0.05375	0.02937	0.03359	0.03085	0.01643	0.03086	<u>0.03608</u>	0.03707	2.76%
		ndcg@5	0.02025	0.03091	0.01624	0.01938	0.01689	0.00850	0.01777	<u>0.02051</u>	0.02103	2.53%
		recall@10	0.04902	0.07066	0.04231	0.05051	0.04877	0.02290	0.04505	<u>0.05051</u>	0.05325*	5.42%
		ndcg@10	0.02119	0.03104	0.01724	<u>0.02246</u>	0.02215	0.00934	0.01905	0.02188	0.02342*	4.26%
	AttMix	recall@5	0.01916	0.04379	0.01095	0.00325	0.01568	0.00406	0.00487	<u>0.01866</u>	0.01916	2.67%
		ndcg@5	0.01144	0.02546	0.00600	0.00206	0.00910	0.00234	0.00280	<u>0.01103</u>	0.01118	1.33%
		recall@10	0.02264	0.05200	0.01717	0.00690	0.02140	0.00649	0.00933	<u>0.02215</u>	0.02936*	32.59%
		ndcg@10	0.00983	0.02315	0.00687	0.00323	0.00977	0.00310	0.00355	<u>0.00987</u>	0.01343*	36.07%

Table 1: Performance of different methods. * indicates statistically significant improvement of the proposed method to baseline models on t-test ($p < 0.05$). Student and Teacher represent the convention student and teacher recommenders, respectively.

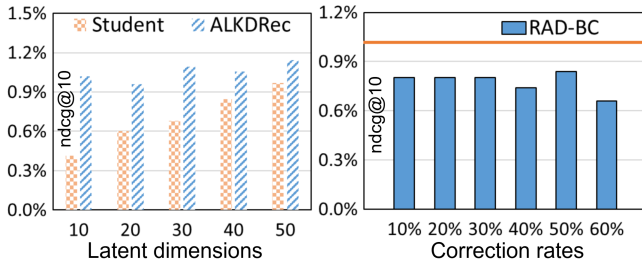


Figure 2: Performance (ndcg@10) with different (a) latent dimensions, and (b) correction rates.

Performance Comparison (RQ1)

Effectiveness. Table 1 presents the performance of various methods, including the student recommenders, teacher recommenders, and KD recommendation methods. To highlight the most notable results, we have bolded the best results and underlined the runner up results for KD-based methods on each dataset. From the experimental results, we can get the following conclusions. Firstly, ALKDRec significantly outperforms all KD-based baselines in most cases, which shows the effectiveness of ALKDRec in KD for SBR. Specifically, ALKDRec achieves average improvements of 17.03%, 9.81%, and 12.97% when compared to the best KD-based baselines on the three backbones FPMC, STAMP, and AttMix, respectively. Second, KD-based baselines outperform student recommender in most cases, which shows the necessity of KD for SBR. However, they may degrade students’ performance sometimes, e.g., STAMP and AttMix perform worse than vanilla students on Amazon-

Games. This may be attributed to their limited model size making it hard to capture the complicated patterns, which demand strong generalization capability with large model size, e.g., multi-level reasoning in AttMix. Therefore, distilling this complicated knowledge is daunting for the student and leads to degradation. This also explains why the complicated AttMix backbone performs worse than other backbones with a limited model size. Third, DLLM2Rec performs well among KD-based baselines but underperforms our ALKDRec, indicating the necessity of our active learning strategy for LLM-based KD.

Efficiency. In the training phase, our ALKDRec only extracts from a subset of instances (e.g., around 44 minutes and 8.6 USD for ChatGPT API in Amazon-Games), which is much faster and computation-saving than training on all samples (e.g., around 1782 minutes and 347.0 USD for ChatGPT API in Amazon-Games). In addition, ALKDRec with a limited model size even outperforms the teacher recommender with a 10X model size, indicating that distilling knowledge from LLMs is an efficient solution for SBR.

Ablation Study (RQ2&3)

To assess the effectiveness of ALKDRec’s module design, we compare with several variants. (1) **TR** distills knowledge from **Teacher Recommender** by mimicking its predictions, which fine-tunes the well-trained student recommender for KD. (2) **Random** distills knowledge from LLM by randomly sampling a subset of instances. (3) **Hardest/Easiest** distill knowledge from LLM for the hardest/easiest instances for their difficulties. (4) **RAD-BC** replaces our active learning strategy in ALKDRec with the **robust active distillation**

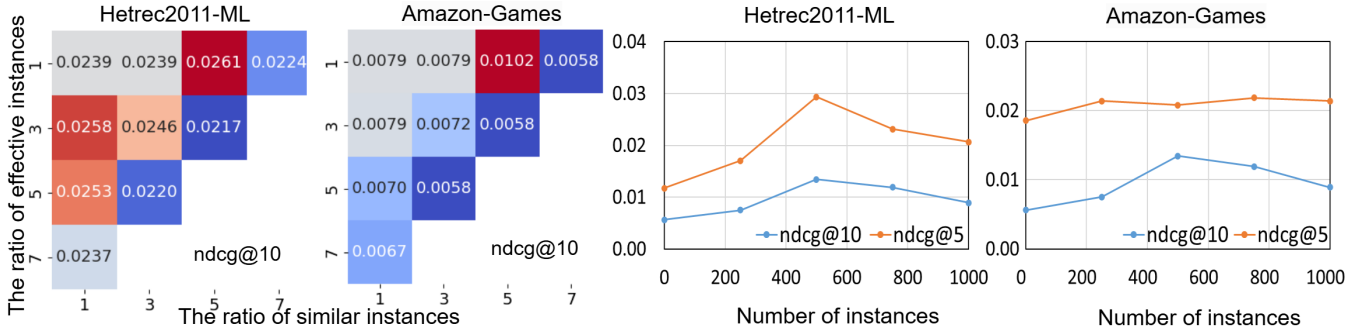


Figure 3: Performance of ALKDRec varying with (a) effective/similar instance ratios; (b) and instance numbers.

Hetrec2011-ML						
Backbone	FPMC		STAMP		AttMix	
Metric	R@10	N@10	R@10	N@10	R@10	N@10
TR	0.01907	0.00897	0.01055	0.00476	0.00852	0.00375
Random	0.01907	0.00856	0.01339	0.00755	0.00974	0.00469
Easiest	0.01623	0.00693	0.01542	0.00754	0.00933	0.00426
Hardest	0.01298	0.00639	0.01582	0.00756	0.00933	0.00439
RAD-BC	0.01420	0.00658	0.01177	0.00601	0.01136	0.00519
ALKDRec	0.02150	0.01019	0.01623	0.00799	0.01095	0.00547

Amazon-Games						
Backbone	FPMC		STAMP		AttMix	
Metric	R@10	N@10	R@10	N@10	R@10	N@10
TR	0.05175	0.02238	0.05200	0.02223	0.02264	0.01002
Random	0.05101	0.02239	0.05225	0.02203	0.02538	0.01132
Easiest	0.05076	0.02238	0.05026	0.02161	0.01966	0.00905
Hardest	0.04877	0.02133	0.05001	0.02173	0.02239	0.00997
RAD-BC	0.05200	0.02321	0.05225	0.02228	0.02712	0.01210
ALKDRec	0.05748	0.02610	0.05325	0.02342	0.02936	0.01343

Table 2: Performance of the variants for ablation studies measured by recall (R@10) and ndcg (N@10).

on binary classification tasks (Baykal et al. 2023) that only models the correct and incorrect instances. We set the number of correct instances in RAD-BC as the total number of effective/similar instances in ALKDRec. Table 2 shows the performance of ALKDRec and its variant, indicating the following conclusions.

RQ2: Firstly, ALKDRec consistently outperforms TR, which shows the necessity of distilling knowledge from the LLM teacher. Secondly, Random shows improvements over TR in most cases, confirming the potential of LLM-based KD. In addition, Figure 2 (a) shows the performance of ALKDRec with FPMC varying with different latent dimensions. It indicates that ALKDRec can consistently improve the student recommenders regardless of their sizes, illustrating its generalizability and robustness for LLM-based KD.

RQ3: ALKDRec outperforms the other instance selection strategies (i.e., Random, Easiest, Hardest, and RAD-BC) in most cases, indicating the effectiveness of LLM-based KD by maximizing the minimal gains of distillation for SBR. Firstly, simply selecting the easiest and hardest samples per-

forms badly among these methods. To be specific, the easiest instance leads to redundant (less informative) distillation as its knowledge is likely to be already captured by the student recommender, contributing less to KD for SBR. The hardest instance makes it hard to be correctly predicted by the LLMs, which may mislead the training of student recommenders. Secondly, ALKDRec outperforms RAD-BC in most cases, which is attributed to ALKDRec can avoid both similar and incorrect predictions of LLMs to KD for SBR. Specifically, Figure 2 (b) compares the performances between ALKDRec and RAD-BC with varying correction rates (Baykal et al. 2023) when adopting backbone FPMC on Hetrec2011-ML. It shows that ALKDRec consistently outperforms RAD-BC, indicating the superiority of the proposed active learning strategy for LLM-based KD.

Hyper-Parameter Study (RQ4)

We now investigate the impacts of key parameters for ALKDRec, including the ratios of effective/similar/incorrect instances and the number of sampled instances for KD $|\mathcal{B}|$. Figure 3 shows the performance of ALKDRec with FPMC across various combinations of effective/similar ratios and different numbers of sampled instances. First, it demonstrates that the optimal performance is achieved with effective, similar, and incorrect instances at a ratio of 1 : 5 : 4. Second, we observe that the performance first improves with more instances, but begins to degrade when exceeding a threshold. We suggest selecting 500 instances for active KD learning for effective and efficient consideration.

Conclusion

In this paper, we propose an Active LLM-based KD Recommendation method for a sustainable and effective solution to SBR. For more efficient LLM-based KD, we propose to elicit student learning from a small proportion of instances. To theoretically ensure effective LLM-based KD, we propose maximizing the minimal gains of distillation by selecting effective instances for KD. Experiments on real-world datasets show that our method significantly outperforms state-of-the-art KD methods with representative recommendation backbones. In future work, we will explore how to efficiently distill embedding-level knowledge from LLMs to the student recommenders.

Acknowledgments

This research is supported by the Agency for Science, Technology and Research (A*STAR) under its RIE 2025 - Industry Alignment Fund - Pre Positioning (IAF-PP) funding scheme (Project No: M23L4a0001) and the Ministry of Education, Singapore, under its MOE AcRF Tier 1, SUTD Kickstarter Initiative (SKI 2021.06.12). This work is also partially supported by the MOE AcRF Tier 1 funding (RG13/23) and Open Foundation of Key Laboratory of Interdisciplinary Research of Computation and Economics (Shanghai University of Finance and Economics), Ministry of Education.

References

- Baykal, C.; Trinh, K.; Iliopoulos, F.; Menghani, G.; and Vee, E. 2023. Robust Active Distillation. In *The Eleventh International Conference on Learning Representations (ICLR)*.
- Chen, G.; Chen, J.; Feng, F.; Zhou, S.; and He, X. 2023. Unbiased Knowledge Distillation for Recommendation. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining (WSDM)*, 976–984.
- Chen, X.; Zhang, Y.; Xu, H.; Qin, Z.; and Zha, H. 2018. Adversarial Distillation for Efficient Recommendation with External Knowledge. *ACM Transactions on Information Systems (TOIS)*, 1–28.
- Cui, Y.; Liu, F.; Wang, P.; Wang, B.; Tang, H.; Wan, Y.; Wang, J.; and Chen, J. 2024. Distillation Matters: Empowering Sequential Recommenders to Match the Performance of Large Language Model. *arXiv preprint arXiv:2405.00338*.
- Deng, Y.; Chen, Y.; Dong, X.; Pan, L.; Li, H.; Cheng, L.; and Mo, L. 2023. BKD: A Bridge-based Knowledge Distillation Method for Click-Through Rate Prediction. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 1859–1863.
- Du, Y.; Luo, D.; Yan, R.; Wang, X.; Liu, H.; Zhu, H.; Song, Y.; and Zhang, J. 2024a. Enhancing Job Recommendation Through LLM-based Generative Adversarial Networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 8363–8371.
- Du, Y.; Wang, Z.; Sun, Z.; Chua, H.; Liu, H.; Wu, Z.; Ma, Y.; Zhang, J.; and Sun, Y. 2024b. Large Language Model with Graph Convolution for Recommendation. *arXiv preprint arXiv:2402.08859*.
- Gao, Y.; Xiong, Y.; Gao, X.; Jia, K.; Pan, J.; Bi, Y.; Dai, Y.; Sun, J.; and Wang, H. 2023. Retrieval-augmented Generation for Large Language Models: A Survey. *arXiv preprint arXiv:2312.10997*.
- Geng, S.; Liu, S.; Fu, Z.; Ge, Y.; and Zhang, Y. 2022. Recommendation as Language Processing (RLP): A Unified Pretrain, Personalized Prompt & Predict Paradigm (p5). In *Proceedings of the 16th ACM Conference on Recommender Systems (RecSys)*, 299–315.
- Hou, Y.; Zhang, J.; Lin, Z.; Lu, H.; Xie, R.; McAuley, J.; and Zhao, W. X. 2024. Large Language Models are Zero-shot Rankers for Recommender Systems. In *European Conference on Information Retrieval (ECIR)*, 364–381.
- Huang, F.; Wang, Z.; Huang, X.; Qian, Y.; Li, Z.; and Chen, H. 2023. Aligning Distillation for Cold-start Item Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 1147–1157.
- Kang, S. 2024. Knowledge Distillation Approaches for Accurate and Efficient Recommender System. *arXiv preprint arXiv:2407.13952*.
- Kang, S.; Hwang, J.; Kweon, W.; and Yu, H. 2020. DERD: A Knowledge Distillation Framework for Recommender System. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (CIKM)*, 605–614.
- Kang, S.; Hwang, J.; Kweon, W.; and Yu, H. 2021. Topology Distillation for Recommender System. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD)*, 829–839.
- Kang, S.; Kweon, W.; Lee, D.; Lian, J.; Xie, X.; and Yu, H. 2023a. Distillation from Heterogeneous Models for Top-K Recommendation. In *Proceedings of the ACM Web Conference 2023*, 801–811.
- Kang, W.-C.; Ni, J.; Mehta, N.; Sathiamoorthy, M.; Hong, L.; Chi, E.; and Cheng, D. Z. 2023b. Do LLMs Understand User Preferences? Evaluating LLMs On User Rating Prediction. *arXiv preprint arXiv:2305.06474*.
- Kweon, W.; Kang, S.; and Yu, H. 2021. Bidirectional Distillation for Top-K Recommender System. In *Proceedings of the Web Conference 2021*, 3861–3871.
- Lee, J.-w.; Choi, M.; Lee, J.; and Shim, H. 2019. Collaborative Distillation for Top-N Recommendation. In *2019 IEEE International Conference on Data Mining (ICDM)*, 369–378.
- Lee, Y.; and Kim, K.-E. 2021. Dual Correction Strategy for Ranking Distillation in Top-N Recommender System. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management (CIKM)*, 3186–3190.
- Liu, C.; Xie, R.; Liu, X.; Wang, P.; Zheng, R.; Zhang, L.; Li, J.; Xia, F.; and Lin, L. 2023a. Future Augmentation with Self-distillation in Recommendation. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML/PKDD)*, 602–618. Springer.
- Liu, F.; Chen, H.; Cheng, Z.; Nie, L.; and Kankanhalli, M. 2023b. Semantic-Guided Feature Distillation for Multimodal Recommendation. In *Proceedings of the 31st ACM International Conference on Multimedia (MM)*, 6567–6575.
- Liu, Q.; Chen, N.; Sakai, T.; and Wu, X.-M. 2024a. Once: Boosting Content-based Recommendation with Both Open and Closed-source Large Language Models. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM)*, 452–461.
- Liu, Q.; Wu, X.; Zhao, X.; Zhu, Y.; Zhang, Z.; Tian, F.; and Zheng, Y. 2024b. Large Language Model Distilling Medication Recommendation Model. *arXiv preprint arXiv:2402.02803*.

- Liu, Q.; Zeng, Y.; Mokhosi, R.; and Zhang, H. 2018. STAMP: Short-Term Attention/Memory Priority Model for Session-based Recommendation. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining (KDD)*, 1831–1839.
- Luo, S.; Yao, Y.; He, B.; Huang, Y.; Zhou, A.; Zhang, X.; Xiao, Y.; Zhan, M.; and Song, L. 2024. Integrating large language models into recommendation via mutual augmentation and adaptive aggregation. *arXiv preprint arXiv:2401.13870*.
- Pan, Y.; He, F.; and Yu, H. 2019. A novel Enhanced Collaborative Autoencoder with Knowledge Distillation for Top-N Recommender Systems. *Neurocomputing*, 332: 137–148.
- Ren, X.; Wei, W.; Xia, L.; Su, L.; Cheng, S.; Wang, J.; Yin, D.; and Huang, C. 2024. RLMRec: Representation Learning with Large Language Models for Recommendation. In *Proceedings of the ACM on Web Conference 2024*, 3464–3475.
- Rendle, S.; Freudenthaler, C.; and Schmidt-Thieme, L. 2010. Factorizing Personalized Markov Chains for Next-Basket Recommendation. In *Proceedings of the 19th international conference on World wide web (WWW)*, 811–820.
- Shi, T.; Zhang, Y.; Xu, Z.; Chen, C.; Feng, F.; He, X.; and Tian, Q. 2023. Preliminary Study on Incremental Learning for Large Language Model-based Recommender Systems. *arXiv preprint arXiv:2312.15599*.
- Sun, W.; Xie, R.; Zhang, J.; Zhao, W. X.; Lin, L.; and Wen, J.-R. 2024a. Distillation is All You Need for Practically Using Different Pre-trained Recommendation Models. *arXiv preprint arXiv:2401.00797*.
- Sun, Z.; Fang, H.; Yang, J.; Qu, X.; Liu, H.; Yu, D.; Ong, Y.-S.; and Zhang, J. 2022. Daisyrec 2.0: Benchmarking Recommendation for Rigorous Evaluation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 8206–8226.
- Sun, Z.; Liu, H.; Qu, X.; Feng, K.; Wang, Y.; and Ong, Y. S. 2024b. Large Language Models for Intent-Driven Session Recommendations. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 324–334.
- Tang, J.; and Wang, K. 2018. Ranking Distillation: Ranking Distillation: Learning Compact Ranking Models With High Performance for Recommender System. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining (KDD)*, 2289–2298.
- Wang, S.; Sui, Y.; Wu, J.; Zheng, Z.; and Xiong, H. 2024a. Dynamic Sparse Learning: A Novel Paradigm for Efficient Recommendation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM)*, 740–749.
- Wang, S.; Zhang, K.; Wu, L.; Ma, H.; Hong, R.; and Wang, M. 2021. Privileged Graph Distillation for Cold Start Recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 1187–1196.
- Wang, X.; Zhang, R.; Sun, Y.; and Qi, J. 2018. KDGAN: Knowledge Distillation with Generative Adversarial Networks. *Advances in neural information processing systems (NeurIPS)*.
- Wang, Y.; Tian, C.; Hu, B.; Yu, Y.; Liu, Z.; Zhang, Z.; Zhou, J.; Pang, L.; and Wang, X. 2024b. Can Small Language Models be Good Reasoners for Sequential Recommendation? In *Proceedings of the ACM on Web Conference 2024*, 3876–3887.
- Wang, Z.; Du, Y.; Sun, Z.; Chua, H.; Feng, K.; Wang, W.; and Zhang, J. 2024c. Re2LLM: Reflective Reinforcement Large Language Model for Session-based Recommendation. *arXiv preprint arXiv:2403.16427*.
- Wei, W.; Ren, X.; Tang, J.; Wang, Q.; Su, L.; Cheng, S.; Wang, J.; Yin, D.; and Huang, C. 2024a. LLMRec: Large Language Models with Graph Augmentation for Recommendation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM)*, 806–815.
- Wei, W.; Tang, J.; Xia, L.; Jiang, Y.; and Huang, C. 2024b. Multi-Modal Knowledge Distillation for Recommendation with Prompt-Tuning. In *The Web Conference*, 3217–3228.
- Wu, L.; Zheng, Z.; Qiu, Z.; Wang, H.; Gu, H.; Shen, T.; Qin, C.; Zhu, C.; Zhu, H.; Liu, Q.; et al. 2023. A Survey on Large Language Models for Recommendation. *arXiv preprint arXiv:2305.19860*.
- Xia, X.; Yin, H.; Yu, J.; Wang, Q.; Xu, G.; and Nguyen, Q. V. H. 2022. On-Device Next-Item Recommendation with Self-Supervised Knowledge Distillation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 546–555.
- Xu, C.; Li, Q.; Ge, J.; Gao, J.; Yang, X.; Pei, C.; Sun, F.; Wu, J.; Sun, H.; and Ou, W. 2020. Privileged Features Distillation at Taobao Recommendations. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, 2590–2598.
- Zhang, P.; Guo, J.; Li, C.; Xie, Y.; Kim, J. B.; Zhang, Y.; Xie, X.; Wang, H.; and Kim, S. 2023. Efficiently Leveraging Multi-level User Intent for Session-based Recommendation via Atten-Mixer Network. In *Proceedings of the sixteenth ACM international conference on web search and data mining (WSDM)*, 168–176.
- Zhang, Y.; Xu, X.; Zhou, H.; and Zhang, Y. 2020. Distilling Structured Knowledge into Embeddings for Explainable and Accurate Recommendation. In *Proceedings of the 13th international conference on web search and data mining (WSDM)*, 735–743.
- Zhao, W. X.; Zhou, K.; Li, J.; Tang, T.; Wang, X.; Hou, Y.; Min, Y.; Zhang, B.; Zhang, J.; Dong, Z.; et al. 2023. A Survey of Large Language Models. *arXiv preprint arXiv:2303.18223*.