

Continual Collaborative Distillation for Recommender System

Gyuseok Lee*
Pohang University of
Science and Technology
Pohang, Gyeongbuk
Republic of Korea
gyuseok.lee@postech.ac.kr

SeongKu Kang*
University of Illinois at
Urbana-Champaign
Champaign, IL, USA
seongku@illinois.edu

Wonbin Kweon
Pohang University of
Science and Technology
Pohang, Gyeongbuk
Republic of Korea
kwb4453@postech.ac.kr

Hwanjo Yu†
Pohang University of
Science and Technology
Pohang, Gyeongbuk
Republic of Korea
hwanjoyu@postech.ac.kr

ABSTRACT

Knowledge distillation (KD) has emerged as a promising technique for addressing the computational challenges associated with deploying large-scale recommender systems. KD transfers the knowledge of a massive teacher system to a compact student model, to reduce the huge computational burdens for inference while retaining high accuracy. The existing KD studies primarily focus on one-time distillation in static environments, leaving a substantial gap in their applicability to real-world scenarios dealing with continuously incoming users, items, and their interactions. In this work, we delve into a systematic approach to operating the teacher-student KD in a non-stationary data stream. Our goal is to enable efficient deployment through a compact student, which preserves the high performance of the massive teacher, while effectively adapting to continuously incoming data. We propose Continual Collaborative Distillation (CCD) framework, where both the teacher and the student continually and collaboratively evolve along the data stream. CCD facilitates the student in effectively adapting to new data, while also enabling the teacher to fully leverage accumulated knowledge. We validate the effectiveness of CCD through extensive quantitative, ablation, and exploratory experiments on two real-world datasets. We expect this research direction to contribute to narrowing the gap between existing KD studies and practical applications, thereby enhancing the applicability of KD in real-world systems.

CCS CONCEPTS

• **Information systems** → **Retrieval models and ranking; Recommender systems; Retrieval efficiency.**

KEYWORDS

Recommender System, Knowledge Distillation, Continual Learning

ACM Reference Format:

Gyuseok Lee, SeongKu Kang, Wonbin Kweon, and Hwanjo Yu. 2024. Continual Collaborative Distillation for Recommender System. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, August 25–29, 2024, Barcelona, Spain. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3637528.3671924>

*Both authors contributed equally to this research.

†Corresponding author.



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

KDD '24, August 25–29, 2024, Barcelona, Spain
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0490-1/24/08.
<https://doi.org/10.1145/3637528.3671924>

1 INTRODUCTION

Recommender systems (RS) have been used in diverse industrial platforms to enhance user experience, foster loyalty, and contribute to business success [15, 18, 23]. In recent years, increasingly large and sophisticated recommendation models, from graph neural networks [38, 46] to large language models [2, 7, 50], have been developed to identify users' intricate preferences. Furthermore, these large-scale models are often employed concurrently through model ensembles to further improve recommendation performance [16, 49]. **However, this improved performance comes at the cost of increased computations, memory resources, and inference latency**, which poses significant challenges for deployment in **real-time** services and resource-constrained environments [4, 16, 43].

To overcome this issue, recent studies have employed knowledge distillation (KD) [10, 33] (Figure 1a). KD serves as a model **compression technique** that transfers knowledge from a massive system (teacher) into a compact model (student), aiming to generate a model that achieves both high effectiveness and efficiency. It first constructs a massive teacher system, which often comprises multiple large-scale models, to attain high performance [16, 49]. Then, during the distillation process, the student is trained to replicate the high-quality recommendations from the teacher system, which are derived from its massive capacity. After the distillation, the student becomes capable of achieving comparable performance to the teacher [16, 36, 43]. Furthermore, the student has significantly reduced inference latency, making it suitable for deployment.

However, existing KD studies have focused on one-time distillation in static environments, overlooking real-world scenarios handling a non-stationary data stream where new users, items, and interactions are continuously incoming. A naive approach to applying KD in the data stream is to repeatedly update the teacher and generate a new student via distillation each time new data arrives. However, this approach raises two critical problems: First, the massive teacher system cannot be updated frequently, as training large-scale models requires significant time and resources [6]. Consequently, the deployed student remains static until the next teacher update cycle, failing to provide recommendations that reflect up-to-date interactions as well as new users and items. Second, conducting the distillation independently each time fails to fully leverage the previously generated knowledge along the data stream. This results in inefficient training and suboptimal performance.

A well-established approach for training a model with a non-stationary data stream is continual learning (CL) [20, 26] (Figure 1b). When trained with the data stream, **an ideal model should seamlessly adapt to newly incoming data without forgetting previous knowledge from historical data**. CL aims to strike a balance

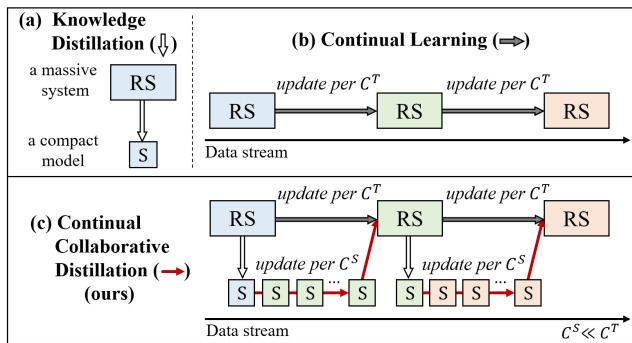


Figure 1: A conceptual comparison of (a) knowledge distillation, (b) continual learning, and (c) the proposed continual collaborative distillation. C^T and C^S denote the update cycle for the massive system (e.g., a weekly update) and the compact model (e.g., a daily update), respectively.

between these two aspects, termed *plasticity* (i.e., adaptation to new data) and *stability* (i.e., preserving previous knowledge) [26]. Although CL has been studied for RS [39, 40, 48], operating the teacher-student KD in a data stream still remains unexplored. Most existing CL methods focus on the continual updates of a *single model* with *sufficient capacity* to learn effectively on its own. However, KD necessitates handling both the teacher and the student which have distinct natures and update cycles. Also, the student has a highly limited learning capability due to its small size, which makes it more challenging to update effectively.

We propose Continual Collaborative Distillation (CCD), a new approach to systematically operate the teacher-student KD in a non-stationary data stream (Figure 1c). Our goal is to enable efficient deployment through a compact student which preserves the high performance of the massive teacher, while effectively adapting to continuously incoming data. In CCD, the student model, which can be updated frequently due to its small size, learns to adapt to new interactions with a short update cycle (e.g., a daily update). Meanwhile, the teacher system uncovers richer knowledge through its larger capacity with a longer update cycle (e.g., a weekly update). As shown in Figure 1c, within each teacher update cycle, we proceed through three consecutive stages: (1) Compact student generation via KD, (2) Continual student updates to provide recommendations reflecting up-to-date interactions, (3) Teacher system update by leveraging knowledge accumulated along the data stream.

We introduce new effective update strategies for both the teacher and the student, considering their distinct natures. To complement the student’s limited learning capability, we propose two techniques: *entity embedding initialization* which facilitates the learning of new users and items based on the recent prominent trends, and *proxy-guided replay learning* which identifies forgotten knowledge and aids in its preservation by replaying past predictions. Furthermore, we propose an effective teacher update strategy that fully leverages previously obtained model knowledge, including both from the teacher-side and the student-side. By repeating the updating along the data stream, CCD allows for collaboratively enhancing both the teacher and the student; the student is enhanced by effectively adapting to new entities and preserving previous knowledge with the replay learning. This enhanced student knowledge is then

subsequently harnessed to improve the teacher system, which in turn allows for generating a more powerful student model. Our contributions are summarized as follows:

- We explore a new problem of operating teacher-student KD in a non-stationary data stream, which has not been studied well. This direction of research can contribute to bridging the gap between the existing KD techniques and real-world applications.
- We introduce CCD framework, where both the teacher and the student collaboratively evolve along the data stream. CCD facilitates the student’s effective adaptation to new data, while also enabling the teacher to fully leverage accumulated knowledge.
- We validate the effectiveness of CCD through extensive experiments on real-world datasets. Furthermore, we provide thorough analyses to verify the validity of each proposed component.

2 RELATED WORK

Knowledge distillation (KD). KD serves as a model compression technique that transfers knowledge from a large teacher model to a lightweight student model [10, 33]. In recent years, the inference costs of recommender systems (RS) have progressively increased, presenting challenges for their practical deployment. Consequently, KD has gained much research attention to reduce inference costs while maintaining high recommendation performance [14, 16, 19, 21, 49]. Earlier work [36] transfers the point-wise importance of top-ranked items by the teacher, and [14] proposes a topology distillation that transfers relational knowledge from the teacher representation space. Recent studies [12, 13, 16, 17, 24] have delved into list-wise distillation to transfer the ranking orders of items directly. They formulate the distillation as a *ranking matching* problem and train the student to preserve the teacher’s permutation. This approach has shown state-of-the-art performance in various ranking-oriented applications such as recommendation [12, 13, 17] and document retrieval [31, 47]. Furthermore, considering the remarkable performance of a massive system consisting of multiple large-scale models, [16, 49] introduce distillation methods tailored for compressing knowledge of an ensembled system. These KD methods have greatly alleviated the huge computational burdens of deploying a large-scale RS, thereby expanding its applicability to various environments.

However, existing KD studies have focused on one-time distillation in static environments, leaving a substantial gap in their application to real-world scenarios with continuously incoming data. In this work, we delve into a systematic approach to operating the teacher-student KD in a non-stationary data stream.

Continual learning (CL). Continual learning [20, 26, 30], also known as lifelong learning or incremental learning, is a concept to train a model with a non-stationary data stream. A major challenge is to strike a balance between plasticity and stability [26], where plasticity refers to the ability to learn new knowledge, and stability focuses on retaining previous knowledge. Naively updating a model on new data often fails to achieve this balance, resulting in catastrophic forgetting of its previous knowledge [20]. Two popular CL approaches are regularization [44, 45] and experience replay [1, 35]. The regularization approach typically imposes regularization constraints on the parameter space, discouraging the model parameters from drastically changing from the previously trained

ones. On the other hand, the replay approach involves reusing historical data, by storing a subset of representative samples [1] or by using a generative model trained on input distribution [35].

Due to its practical importance, CL has been actively studied for RS. [39, 40, 44] have delved into the regularization approach. [39, 44] have focused on structure-aware regularization for graph-neural network-based models. More recently, [40] proposes a new personalized weight adjustment to tailor the impact of regularization for each user, pointing out that preserving the same amount of historical information for all users is sub-optimal. On the other hand, [1, 3, 29, 48] have focused on the replay approach. [1, 29] select representative data based on interaction frequency, while [3, 48] have focused on self-correcting learning by using an error memory to store samples that the model previously failed to predict.

Though effective, these CL methods have mostly been studied for updating a *single* model having *sufficient capacity* to learn effectively on its own, which differs significantly from our problem. Operating the teacher-student KD in a data stream still remains unexplored.

3 PROBLEM FORMULATION

3.1 Concept Definition

Definition 1 (Teacher-Student Knowledge Distillation). A massive teacher recommender system RS^T , which typically comprises several large-scale models, achieves high performance through its large capacity. However, it also incurs high computational costs for inference. KD is employed to compress RS^T into a lightweight student model S [16]. The student model has significantly reduced inference latency, making it well-suited for real-time services and resource-constrained environments.

Definition 2 (Recommendation with Data Stream). A real-world recommender system operates in a non-stationary data stream where new users, items, and interactions are continuously incoming. **In many practical scenarios, a fixed-size time window of recent data is employed to update the system [44].** Naturally, the data stream D is viewed as consecutive blocks of interaction data $[D_1, D_2, \dots, D_K]$ with a certain time interval C . D_k corresponds to set of interactions collected from the timestamp t_{k-1} to t_k , where $t_k = t_{k-1} + C$. Each block serves as training data for updating the system at each time segment. At k -th block, the system optimizes its performance on D_k while leveraging the knowledge obtained from the previous blocks, D_1, \dots, D_{k-1} . Note that the system only uses the most recent block for training without directly accessing the previous blocks [40, 44, 48]. The performance is evaluated across the entire timeline.

Definition 3 (Teacher/Student update cycle). An update cycle for a system signifies the frequency at which the system is updated to adapt newly incoming interaction data [29]. It is often empirically decided considering various factors such as available resources and the amount of collected data. Let C^T and C^S denote update cycles for the teacher and the student, respectively. In practice, C^T is much larger than C^S , as updating the teacher system requires more time and resources compared to the student model. For instance, one might opt for a weekly update cycle for the teacher ($C^T = 7$ days) while allowing the student to update on a daily basis ($C^S = 1$ day). It is important to note that we define the data blocks in terms of the teacher system, thus the time interval is defined as $C = C^T$.

3.2 Problem Definition

We delve into a systematic approach to operate teacher-student KD in a non-stationary data stream. Our goal is to enable efficient deployment through a lightweight S , preserving the high accuracy of RS^T while effectively adapting to continuously incoming data. This problem has two unique desiderata:

- (1) The student should effectively adapt to continuously incoming data by itself during C^T . That is, while the teacher remains static, the student should cope with evolving preferences as well as new users and items. However, due to its limited capacity, naive updates of the student often lead to ineffective adaptation and a substantial loss of previously acquired knowledge.
- (2) The knowledge of teacher and student should be effectively accumulated and leveraged throughout the data stream. The conventional KD approach conducts distillation independently for each time segment. This cannot fully leverage the knowledge generated from previous blocks, leading to inefficient training and suboptimal performance.

Lastly, it is worth noting that we aim for a model-agnostic solution, enabling service providers to choose their preferred model for both the teacher and the student according to their environments.

4 METHODOLOGY

4.1 Overview

We introduce a new **C**ontinual **C**ollaborative **D**istillation (CCD) framework, where both the teacher and the student evolve collaboratively along the non-stationary data stream. Assume that we have RS^T updated using up to $(k-1)$ -th data block. During the subsequent teacher update cycle (i.e., during collecting k -th block for teacher update), CCD follows three consecutive stages:

- (1) **Student model generation (§4.2)**. We generate a compact student model S by compressing RS^T through KD. The student model is used for deployment.
- (2) **Continual update of student (§4.3)**. The student model is continually updated using non-stationary data with the cycle C^S . This process allows recommendations to aptly reflect new users, new items, and up-to-date interactions. We introduce new embedding initialization and proxy-guided replay learning strategies to improve the training while preventing the loss of previous knowledge.
- (3) **Teacher system update (§4.4)**. When the teacher's next update cycle arrives, RS^T is updated using new interaction data D_k accumulated during the cycle. We propose a new strategy for selectively harnessing the knowledge obtained from the student side to further enhance RS^T .

For each data block, CCD iterates these three consecutive stages (§4.2-§4.4). In the subsequent sections, we explain how CCD operates for the k -th data block. Figure 2 provides an overview of CCD.

4.2 Stage 1: Student model generation via KD

Given the massive teacher system $RS^T(\cdot; \Theta_{k-1}^T)$ which was updated up to $(k-1)$ -th block, we generate a lightweight student model $S(\cdot; \theta_k^S)$ which will be deployed for the current k -th block. As KD is the model-agnostic technique, the student model can be any existing recommendation model that predicts the ranking score for

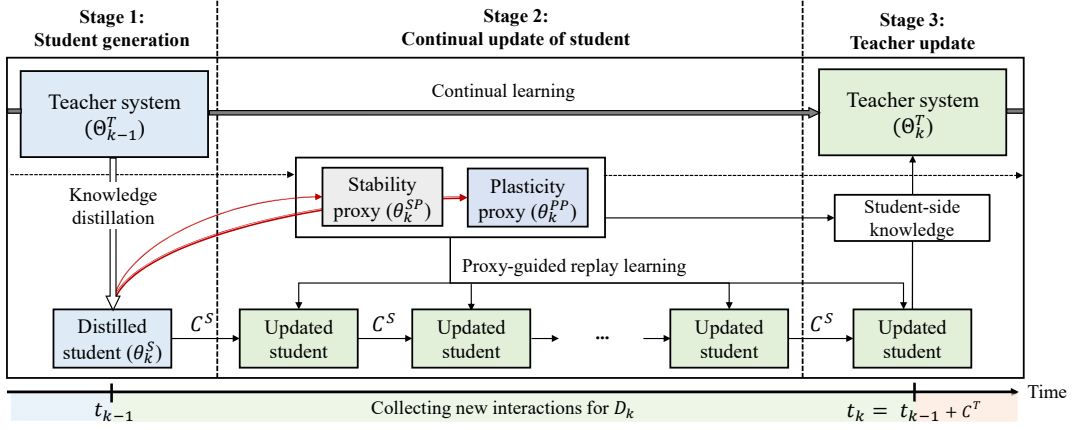


Figure 2: Overview of CCD framework for k -th data block.

each user-item pair, i.e., $S : \mathcal{U} \times \mathcal{I} \rightarrow \mathbb{R}$, where \mathcal{U} and \mathcal{I} denote the set of users and items, respectively.

In CCD framework, any off-the-shelf KD technique can be flexibly employed to generate the student model. In this work, we employ the recent list-wise distillation [12, 16] that trains the student to emulate the item permutation (i.e., ranking orders) predicted by the teacher. Let π_u^T denote the item ranking list for each user u predicted by RS^T . The list-wise distillation loss is defined as the negative log-likelihood of permutation probability [16]¹:

$$\min_{\theta^S} \mathcal{L}_{KD} = - \sum_{u \in \mathcal{U}} \log p(\pi_u^T | \theta^S). \quad (1)$$

After KD, the distilled student can achieve low inference latency due to its small size while preserving knowledge of RS^T .

4.3 Stage 2: Continual update of student with incoming interactions

As RS^T cannot be updated during its cycle C^T , the student model needs to effectively adapt to continuously incoming data by itself. Specifically, the student is updated per its own cycle C^S ($\ll C^T$).

In this section, we explain how we update the student model. We first present how CCD learns with new data to provide recommendations reflecting up-to-date interactions (§4.3.1). Then, we introduce a new proxy-guided replay learning to support effective training by preventing the loss of previous knowledge (§4.3.2). Lastly, we summarize the overall training objective (§4.3.3).

4.3.1 Learning with new interactions. We update the student model using new interactions collected for the current data block. Here, an important challenge arises when dealing with new entities that did not exist in the previous data blocks. In practical applications, new users and new items continuously emerge. However, due to the limited capacity, it is challenging for the student to effectively learn their preference from scratch.

To facilitate learning for these new entities, we introduce a simple yet effective **entity embedding initialization** technique. Let G denote the user-item bipartite graph, where nodes are users and items, and edges represent their interactions observed within the current data

block. For a given node i , we denote \mathcal{N}_i^h as the set of h -hop neighboring nodes in the graph. Recent studies [34, 41] have utilized the 1-hop neighbors to generate embeddings for new entities. Specifically, the embeddings are initialized by aggregating the directly interacted entity embeddings, i.e., $\mathbf{e}_i = \text{AGG}(\{\mathbf{e}_u : u \in \mathcal{N}_i^1\})$.

However, for new entities, 1-hop interactions are typically highly limited. To supplement limited interactions, we leverage the 2-hop relations (e.g., items purchased together, users who bought the same item) based on the frequency information from the current data block. Formally, we identify the set of *prominent* users and items, denoted as \mathcal{P}_U and \mathcal{P}_I , which have top interaction frequency in the current block. These entities reveal the up-to-date prominent trends, such as highly in-demand or trending items, assisting in comprehending newly emerging entities. We initialize embeddings for new users and items as follows:

$$\begin{aligned} \mathbf{e}_u &= \text{AGG}(\{\mathbf{e}_x : x \in \mathcal{N}_u^1 \cup (\mathcal{N}_u^2 \cap \mathcal{P}_U)\}), \\ \mathbf{e}_i &= \text{AGG}(\{\mathbf{e}_x : x \in \mathcal{N}_i^1 \cup (\mathcal{N}_i^2 \cap \mathcal{P}_I)\}). \end{aligned} \quad (2)$$

We use the **simple mean pooling** for the aggregation function $\text{AGG}(\cdot)$. These initialized embeddings are added to the embedding table of the student model.

Then, the student model is updated with new interactions within the current data block. We employ the pair-wise ranking loss [32]:

$$\min_{\theta^S} \mathcal{L}_{BPR} = - \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{N}_u^1} \sum_{j \notin \mathcal{N}_u^1} \log \sigma(\hat{r}_{u,i}^S - \hat{r}_{u,j}^S), \quad (3)$$

where σ denotes the sigmoid function. $\hat{r}_{u,i}^S = S(u, i)$ denotes the ranking score predicted by the student.

4.3.2 Proxy-guided replay learning. Naively updating a model with newly incoming interactions can result in catastrophic forgetting, where the model significantly loses previously acquired knowledge [39, 40]. This largely degrades the overall performance of the model and further hinders the learning of new data. Furthermore, the student model has a highly limited capacity, which makes it more challenging to update effectively.

As a solution, we introduce a new proxy-guided replay learning, which employs an external memory called *proxy* to assist the student's learning. The learning theory in neuroscience [28] posits that humans do effective learning through two complementary systems: a fast learning system for short-term adaptation to specific

¹As this KD technique is not our contribution, we provide its details in Appendix A.1. Note that CCD framework is not dependent on this specific KD technique.

	Student w/ naive update					
After D_4	0.034	0.049	0.061	0.096	0.088	
After D_5	0.030	0.038	0.046	0.071	0.080	0.093
	Stability proxy					
After D_4	0.121	0.094	0.089	0.092	0.083	
After D_5	0.119	0.093	0.092	0.094	0.087	0.069
	Plasticity proxy					
After D_4	0.068	0.074	0.087	0.107	0.113	
After D_5	0.072	0.077	0.084	0.098	0.099	0.104
	D_0	D_1	D_2	D_3	D_4	D_5

Figure 3: Effects of proxy. After D_4 and D_5 (y-axis), we assess Recall@20 on test sets from the previous blocks (x-axis). A naive update of the student results in significant catastrophic forgetting. Two proxies effectively accumulate the previous knowledge from complementary perspectives. (Dataset: Yelp) experiences, and a slow learning system for gradual acquisition of structured knowledge. Drawing from the theory, we employ two complementary memories for the student, referred to as *plasticity proxy* focusing on rapid adaptation, and *stability proxy* emphasizing long-term retention of knowledge. During the student’s update, we utilize these proxies to identify forgotten preference knowledge and assist in its preservation by replaying past predictions.

Stability and plasticity proxies. A proxy serves as a memory that accumulates the model knowledge acquired along the data stream. To efficiently implement the proxy, we employ the *temporal mean* technique [37] which accumulates a model’s knowledge by the exponential moving average of the model parameters. Specifically, we construct the proxy by the temporal mean of the distilled student models generated by KD (§4.2). We introduce two complementary proxies: the *plasticity proxy* S^{PP} and the *stability proxy* S^{SP} . At k -th data block, the proxies are updated as follows:

$$\begin{aligned}\theta_k^{SP} &= (1 - w_{SP}) \cdot \theta_{k-1}^{SP} + w_{SP} \cdot \theta_k^S, \\ \theta_k^{PP} &= (1 - w_{PP}) \cdot \theta_{k-1}^{PP} + w_{PP} \cdot \theta_k^S.\end{aligned}\quad (4)$$

θ_k^S denotes the parameters of the distilled student model from §4.2. θ_k^{SP} and θ_k^{PP} denote the parameters for the stability and plasticity proxies that support training on k -th block, respectively. The scalar weights are hyperparameters following the relationship: $0 < w_{SP} \ll w_{PP} < 1$. Therefore, the stability proxy is slowly updated with the long-term retention of previous knowledge, while the plasticity proxy is rapidly updated with an emphasis on recent trends. These distinct proxies provide complementary views for the previously acquired knowledge (Figure 3). Note that proxies are updated once per block, right after KD from the teacher system.

Replay learning for forgotten knowledge. Using the proxies, we identify forgotten preference knowledge and train the student to recover it by replaying past predictions. A recommendation model produces a ranked list of items for each user. In the list, the top-ranked items correspond to the most probable predictions based on the model’s knowledge of the user preference. In this regard, we identify the forgotten knowledge that was previously captured but is not well-reflected in the current model based on rank disparities in the ranking lists; if items that were previously ranked near the top now hold significantly lower ranks, it may suggest that the model has forgotten knowledge related to them.

Given a model A , $\text{rank}_{u,i}^A$ denotes the rank of item i in the ranking list of user u . A lower value indicates a higher ranking position, i.e., $\text{rank}_{u,*}^A = 0$ is the highest rank. For each item i , the rank disparity with respect to another model B is defined as follows:

$$d_i^u(A, B) = \exp(\epsilon \cdot (\text{rank}_{u,i}^A - \text{rank}_{u,i}^B)), \quad (5)$$

where $\epsilon > 0$ is a hyperparameter to control the sharpness of the distribution. We utilize the exponential function to put a stronger emphasis on items with large rank disparities. A high value of $d_i^u(A, B)$ indicates that item i is ranked significantly higher by model B compared to model A .

We create top- N recommendation lists for each user from the student and proxy models, then construct item sets for replay learning based on the rank disparity. Specifically, we obtain I_{SP} by sampling items from the top- N list of S^{SP} , using a probability distribution $p(i) \propto d_i^u(S, S^{SP})$. Similarly, we acquire I_{PP} by sampling items from the top- N list of S^{PP} , using a distribution $p(i) \propto d_i^u(S, S^{PP})$. Then, we train the student to recover previous knowledge on the identified items by replaying the predictions from the proxies:

$$\min_{\theta^S} \mathcal{L}_{RE} = - \sum_u \left(\sum_{i \in I_{SP}} \ell(\hat{r}_{u,i}^S, \hat{r}_{u,i}^{SP}) + \sum_{i \in I_{PP}} \ell(\hat{r}_{u,i}^S, \hat{r}_{u,i}^{PP}) \right). \quad (6)$$

$\ell(\cdot, \cdot)$ denotes the error function between two predictions. In this work, we employ the simple binary cross-entropy loss. It is worth noting that I_{SP} and I_{PP} are dynamically constructed based on the current state of the student, enabling the effective identification and recovery of forgotten knowledge. This is an important distinction from the existing replay-based CL methods for RS [1, 29, 48], which leverage a pre-constructed set of historical data without considering the current state of the student.

4.3.3 The overall objective for student update. The student model is updated with the new interaction data per C^S , assisted by the stability and plasticity proxies. The overall objective is:

$$\min_{\theta^S} \mathcal{L}_{BPR} + \lambda_{RE} \cdot \mathcal{L}_{RE}. \quad (7)$$

λ_{RE} is a hyperparameter controlling the impact of replay learning.

4.4 Stage 3: Teacher system update

Once the teacher update cycle arrives (i.e., the k -th block has been completely collected), we update RS^T . To update a system on the data stream, two types of losses are typically employed [39, 40]: (1) collaborative filtering loss to learn new user-item interactions, and (2) continual learning loss to mitigate forgetting of previous system knowledge. We also employ these two objectives. Moreover, we introduce an additional objective to fully leverage the knowledge obtained from the student side.

4.4.1 Leveraging the student-side knowledge. We argue that RS^T can be further improved by leveraging knowledge obtained from the student-side in two aspects: First, the updated student model contains up-to-date knowledge of the current data block, unlike RS^T which has not been updated for a while. In particular, using the student model, we can obtain potential interactions that are not directly observed in the current block but are likely to be observed in the future. Second, the proxies, which accumulate knowledge acquired throughout the data stream, serve as valuable knowledge sources to mitigate catastrophic forgetting. While it is

Algorithm 1: CCD algorithm for k -th data block

Input : Data stream D , Teacher system $RS^T(\cdot; \Theta_{k-1}^T)$, Proxies $S^{SP}(\cdot; \theta_{k-1}^{SP}), S^{PP}(\cdot; \theta_{k-1}^{PP})$
Output : Updated teacher system $RS^T(\cdot; \Theta_k^T)$, Updated proxies $S^{SP}(\cdot; \theta_k^{SP}), S^{PP}(\cdot; \theta_k^{PP})$

- 1 Generate a new distilled student S via KD ▷ Eq.1
- 2 Update the stability and plasticity proxies S^{SP}, S^{PP} ▷ Eq.4
- 3 **for every** C^S **do**
- 4 Update the student model S ▷ Eq.7
- 5 Update the teacher system RS^T ▷ Eq.9

possible to consider utilizing separate proxies for RS^T , given the large size of RS^T , employing the student-side proxies through KD is much more cost-effective in terms of both space and time.

During the teacher training, we selectively leverage the most confident predictions from the student-side models (i.e., the student and two proxies) based on the rank disparity. Specifically, we identify items ranked near the top by the student-side models but assigned significantly lower rankings by RS^T . From the top- N list of each student-side model $S^* \in \{S, S^{SP}, S^{PP}\}$, we obtain $I_{S^* \rightarrow T}$ by sampling items using a probability distribution $p(i) \propto d_i^u(RS^T, S^*)$. The student-side knowledge is transferred to the teacher as follows:

$$\mathcal{L}_{S \rightarrow T} = - \sum_u \sum_{S^*} \sum_{i \in I_{S^* \rightarrow T}} \ell(\hat{r}_{u,i}^T, \hat{r}_{u,i}^{S^*}). \quad (8)$$

4.4.2 The overall objective for teacher update. RS^T is updated with new interaction data, assisted by the student-side knowledge.

$$\min_{\Theta^T} \mathcal{L}_{CF} + \lambda_{CL} \cdot \mathcal{L}_{CL} + \lambda_{S \rightarrow T} \cdot \mathcal{L}_{S \rightarrow T}. \quad (9)$$

Θ^T denotes the training parameters of RS^T . \mathcal{L}_{CF} is the original collaborative filtering loss (e.g., binary cross-entropy) used to train RS^T . \mathcal{L}_{CL} is the continual learning loss. CCD does not require a specific CL technique, and various CL methods for RS can be flexibly employed. In this work, we use the recently proposed method [40]. λ_{CL} and $\lambda_{S \rightarrow T}$ are the hyperparameters balancing the losses. As RS^T evolves during its training, we gradually reduce the impact of the student-side knowledge. We use a simple annealing schedule, where its impact at t -th epoch is controlled as $\lambda_{S \rightarrow T} = \lambda_{S \rightarrow T}^0 \cdot e^{-t/\tau}$. Here, $\lambda_{S \rightarrow T}^0$ controls the initial impact, and τ controls the annealing speed. The overall training process is outlined in Algorithm 1.

Remarks. A key aspect of CCD involves the collaborative evolution of the teacher and student. To elaborate, we enhance the student by effectively adapting to new entities and mitigating the forgetting problem using proxies (§4.3). This enhanced student knowledge is then subsequently harnessed to improve the teacher (§4.4), which in turn allows for generating a more powerful student (§4.2).

5 EXPERIMENTS

5.1 Experimental Setup

We provide details of setup in Appendix A.2.

5.1.1 Datasets. We use two real-world datasets: Gowalla and Yelp [21, 44]. To simulate the non-stationary data streams, we split each dataset such that the first 50% forms the base block (D_0), and the remaining 50% is evenly divided into 5 incremental blocks

Table 1: Data block statistics of two datasets.

Data Blocks		D_0	D_1	D_2	D_3	D_4	D_5
Gowalla	# of accumulated users	19,668	21,266	24,107	26,840	29,106	29,858
	# of accumulated items	39,354	39,809	40,381	40,691	40,908	40,988
	# of new users	–	1,598	2,841	2,733	2,266	752
	# of new items	–	455	572	310	217	80
	# of interactions	513,732	62,336	112,167	123,042	126,474	89,713
Yelp	# of accumulated users	12,248	13,367	13,896	14,399	14,761	14,950
	# of accumulated items	10,822	11,345	11,634	11,970	12,197	12,261
	# of new users	–	1,119	529	503	362	189
	# of new items	–	523	289	336	227	64
	# of interactions	151,084	51,280	34,501	33,593	37,970	34,189

(D_1, \dots, D_5), according to the temporal timestamp [29, 44]. For each incremental block, we randomly divide the interactions of each user into train/validation/test sets in 80%/10%/10% split. Table 1 presents the statistics of each block.

5.1.2 Teacher-student KD setup. As our focus is on operating the teacher-student KD in a data stream, we closely follow the setup of the existing KD studies [16, 24, 36]. For **teacher**, we construct a massive system by an ensemble of more than five large models [16]. We increase the capacity of the teacher system until its performance is no longer improved. For **student**, we employ two widely used backbone models: a matrix factorization (MF)-based model [32] and a graph neural network (GNN)-based model [8]. We set a small embedding size for the student (8 for Yelp, and 16 for Gowalla), considering the teacher size for each dataset [16, 24]. The student update cycle is set as one-tenth of the teacher cycle. Note that to ensure consistent evaluation for both the teacher and student, all models are assessed on the same test set after training with all interactions within each data block. Further details on KD setup including teacher configuration, and the number of parameters are provided in Appendix A.2.2.

5.1.3 Evaluation protocol. We closely follow the evaluation protocol of the existing CL methods [5, 40, 44]. All models are first trained on D_0 , then continually updated using the incremental blocks. Specifically, at k -th block, the model is updated using D_k , while access to past blocks D_0, \dots, D_{k-1} is forbidden [5, 40, 44].

The evaluation of CL focuses on assessing how effectively a model achieves a balance between plasticity and stability. We employ two CL metrics proposed in [5]. We construct the matrix $A \in \mathbb{R}^{K \times K}$, where each element a_{ij} denotes the recommendation performance on block j after completing training on block i . After training on each k -th block, we report two metrics:

- **Learning Average (LA)** = $\frac{1}{k} \sum_{i=1}^k a_{i,i}$. It assesses how well a model adapts to new blocks, focusing on the plasticity aspect.
- **Retained Average (RA)** = $\frac{1}{k} \sum_{i=1}^k a_{k,i}$. It assesses how well a model retains past knowledge, focusing on the stability aspect.

A model’s overall capability is summarized by the harmonic mean of LA and RA, denoted by **H-mean** [5]. To evaluate the recommendation performance, we employ Recall@20 and NDCG@20 [5, 22, 40, 44]. We report the results from five independent runs.

5.1.4 Baselines. To evaluate CCD, which trains both the teacher and student in a data stream, we compare various CL approaches.

- **Full-Batch** uses *all the historical data* to train the model from scratch. We report its results solely for reference purposes.
- **Fine-Tune** is a naive baseline that updates the model using its original loss function on new data.

- **LWC-KD-PIW** [40] is the state-of-the-art regularization-based CL method for recommendation. It applies the regularization (i.e., LWC-KD [39]) with personalized weights, adjusting the regularization effects considering the dynamics of each user.
- **ReLoop2** [48] is the state-of-the-art replay-based CL method for recommendation. It introduces self-correcting learning using an error memory that stores samples that the model failed to predict. For the student, we further devise two baselines that represent the best solution that combines the existing KD and CL methods; along the data stream, we update the teacher using the state-of-the-art CL method (i.e., LWC-KD-PIW). Then, for each block, we generate a compact student via KD [16], which is subsequently updated by either Fine-tune or LWC-KD-PIW, referred to as **KD + Fine-tune** and **KD + LWC-KD-PIW**, respectively. Lastly, our approach is:
- **CCD** trains the teacher and student collaboratively along the data stream. We use the same KD technique [16] for stage 1 (\mathcal{L}_{KD} in Eq.1) and LWC-KD-PIW [40] for stage 3 (\mathcal{L}_{CL} in Eq.9).

5.1.5 Implementation Details. We utilize PyTorch with CUDA from RTX A6000 and AMD EPYC 7313 CPU. All hyperparameters are tuned via grid search on the validation set. The learning rate and L_2 regularization for the Adam optimizer are chosen from $\{0.001, 0.005, 0.01\}$ and $\{0.0001, 0.0005, 0.001\}$, respectively. For hyperparameters related to the list-wise KD, we follow the setup from [16]. The impacts of continual learning λ_{CL} is chosen from $\{0.001, 0.01, 0.1\}$. For CCD, w_{SP} and w_{PP} for updating stability and plasticity proxies are chosen from $\{0.0, 0.1\}$ and $\{0.9, 1.0\}$, respectively. The number of item samples $I_{PP}, I_{SP}, I_{S \rightarrow T}$ is chosen from $\{1, 3, 5\}$. The remaining hyperparameters of CCD are set to their default values. The number of prominent users \mathcal{P}_U and items \mathcal{P}_I are set to 20. The rank disparity-based sampling is conducted from the top-50 list of each model (i.e., $N = 50$). We set $\epsilon = 10^{-3}$, $\tau = 5$, $\lambda_{RE} = 0.5$, and $\lambda_{S \rightarrow T}^0 = 0.5$. For baseline-specific hyperparameters, we follow the search ranges provided in the original papers.

5.2 Performance Comparison

5.2.1 Main results. Table 2 and Table 3 show the overall performances of the teacher systems and student models optimized by each compared method. Overall, CCD performs better than all baselines in terms of both adapting to new data (LA) and retaining previous knowledge (RA), achieving a good balance between them (H-mean). Also, CCD consistently improves both the teacher and student, with the improvements often gradually increasing throughout the data stream (Figure 4). We analyze the results from various perspectives:

Teacher-side. Teacher systems trained with CCD consistently achieve higher H-mean across all data blocks compared to state-of-the-art CL methods, including both regularization-based (i.e., LWC-KD-PIW) and replay-based (i.e., ReLoop2) methods. Unlike the CL methods that update the teacher by itself, CCD additionally leverages the student-side knowledge. As discussed in §4.4.1, the student-side models provide up-to-date knowledge from the updated student as well as historical knowledge accumulated via two complementary proxies. This can aid the teacher in adapting to new data more effectively while retaining past knowledge. Interestingly, we observe that the forgetting phenomenon is exacerbated through

the ensemble for the teacher. We conjecture that this amplification could stem from the combined effect of individual model forgetting within the ensemble, leaving further investigation for future study.

Student-side. Student models trained with CCD consistently achieve higher H-mean compared to both state-of-the-art CL methods and the KD-enhanced variants (i.e., KD + Fine-Tune/LWC-KD-PIW). Overall, methods that leverage KD from the teacher systems (i.e., KD + Fine-Tune/LWC-KD-PIW, CCD) achieve higher performance than the remaining methods training the student by itself. This shows the importance of leveraging extensive teacher knowledge. Moreover, we observe that the effectiveness of the CL methods is rather limited for the student models, particularly for the latter blocks (i.e., D4 and D5). The student has a highly limited capacity and learning capability, making it more challenging to update effectively. This aspect has not been explicitly considered in previous CL methods. CCD introduces two new strategies, entity embedding initialization and proxy-guided replay learning, which effectively support the student’s learning process.

It is also noted that CCD even generally outperforms Full-Batch in terms of LA. Full-Batch is certainly the strongest baseline for RA (stability), as it learns from all observed historical data. However, it is not always the strongest method for LA (plasticity), as it treats all data equally without putting emphasis on the latest interactions. CCD effectively learns the latest interactions while preserving previous knowledge, which leads to enhanced performance in both LA and RA.

Collaborative evolution. Figure 4 presents the absolute H-mean gain of CCD over the best CL competitor. We report the results of Recall@20 with the GNN-based student. In CCD, the teacher and the student collaboratively improve each other by mutual knowledge exchange (§4.2, §4.4). These improvements accumulate over time, progressively refining both the teacher and student models. As a result, the performance gap between CCD and the CL competitor generally widens over time, which ascertains the collaborative evolution in CCD.

5.2.2 Accuracy-efficiency comparison. Table 4 compares the accuracy and efficiency aspects of teacher and student, trained by the best CL method (i.e., LWC-KD-PIW) and CCD, respectively. We increase the student size until it achieves *comparable* performance to the teacher. Then, we report the average H-mean (Recall@20) for all blocks, the number of parameters, and the time required for training and inference. We utilize PyTorch with CUDA from RTX A6000 and AMD EPYC 7313 CPU. Compared to the massive teacher system which has significant computational costs for training and consolidating the multiple large-scale models, the compact student model trained with CCD significantly reduces the huge computational burdens while maintaining the high performance of the teacher system.

Lastly, we highlight that the KD process takes negligible time compared to the teacher update. Specifically, on the Yelp dataset, KD takes about 8 minutes, whereas the teacher update takes about 107 minutes. Considering the student update takes about 5 minutes, the KD process, followed by student updates, is more efficient than repetitive updates of the large-scale teacher, making it more suitable for real-time applications.

Table 2: The overall performance (Recall@20) comparison. * denotes $p < 0.05$ for the paired t-test on CCD with the best baseline.

Gowalla		After D_1			After D_2			After D_3			After D_4			After D_5		
		LA	RA	H-mean	LA	RA	H-mean	LA	RA	H-mean	LA	RA	H-mean	LA	RA	H-mean
Teacher	Full-Batch	0.1758	0.1752	0.1755	0.1760	0.1804	0.1782	0.1739	0.1774	0.1756	0.1753	0.1794	0.1773	0.1767	0.1784	0.1775
	Fine-Tune	0.1725	0.1544	0.1629	0.1714	0.0957	0.1228	0.1695	0.0828	0.1113	0.1700	0.0717	0.1009	0.1722	0.0577	0.0864
	LWC-KD-PIW	0.1826	0.1483	0.1637	0.1791	0.1141	0.1394	0.1772	0.1023	0.1297	0.1779	0.0903	0.1198	0.1797	0.0826	0.1132
	ReLoop2	0.1706	0.1570	0.1635	0.1595	0.1454	0.1521	0.1536	0.1243	0.1374	0.1512	0.1087	0.1265	0.1482	0.0980	0.1180
	CCD w/ MF student	0.1823	0.1623*	0.1717*	0.1784	0.1524*	0.1644*	0.1763	0.1429*	0.1579*	0.1764	0.1378*	0.1547*	0.1779	0.1302*	0.1504*
CCD w/ GNN student	0.1809	0.1619	0.1709*	0.1784	0.1513*	0.1637*	0.1764	0.1400	0.1561*	0.1764	0.1333*	0.1519*	0.1787	0.1246*	0.1468*	
Student (MF)	Full-Batch	0.1303	0.1318	0.1310	0.1296	0.1327	0.1311	0.1293	0.1338	0.1315	0.1310	0.1360	0.1335	0.1324	0.1383	0.1353
	LWC-KD-PIW	0.1261	0.1134	0.1194	0.1281	0.0897	0.1055	0.1287	0.0914	0.1069	0.1305	0.0789	0.0983	0.1342	0.0694	0.0915
	ReLoop2	0.1101	0.1119	0.1110	0.1066	0.1055	0.1060	0.1023	0.0998	0.1010	0.1008	0.0892	0.0946	0.1003	0.0826	0.0906
	KD + Fine-Tune	0.1309	0.1457	0.1379	0.1273	0.1153	0.1210	0.1271	0.1023	0.1134	0.1272	0.1031	0.1139	0.1292	0.0835	0.1014
	KD + LWC-KD-PIW	0.1292	0.1442	0.1363	0.1249	0.1235	0.1242	0.1247	0.1047	0.1138	0.1284	0.0888	0.1050	0.1301	0.0842	0.1022
CCD	0.1446*	0.1457	0.1451	0.1333	0.1432*	0.1381*	0.1280	0.1416*	0.1345*	0.1255	0.1363*	0.1307*	0.1249	0.1286*	0.1267*	
Student (GNN)	Full-Batch	0.1390	0.1393	0.1391	0.1408	0.1449	0.1428	0.1411	0.1446	0.1428	0.1431	0.1454	0.1442	0.1456	0.1468	0.1462
	LWC-KD-PIW	0.1380	0.1034	0.1182	0.1369	0.0837	0.1039	0.1369	0.0892	0.1080	0.1370	0.0794	0.1005	0.1398	0.0753	0.0979
	ReLoop2	0.1277	0.1025	0.1137	0.1256	0.0950	0.1082	0.1248	0.0878	0.1031	0.1254	0.0844	0.1009	0.1282	0.0759	0.0953
	KD + Fine-Tune	0.1523	0.1495	0.1507	0.1506	0.1318	0.1406	0.1486	0.1115	0.1274	0.1496	0.0991	0.1192	0.1525	0.0872	0.1110
	KD + LWC-KD-PIW	0.1523	0.1496	0.1509	0.1507	0.1323	0.1409	0.1493	0.1052	0.1234	0.1502	0.0966	0.1176	0.1534	0.0871	0.1111
CCD	0.1681*	0.1421	0.1540	0.1622	0.1428*	0.1519*	0.1590	0.1376*	0.1475*	0.1584	0.1330*	0.1446*	0.1598	0.1202*	0.1372*	
Yelp		After D_1			After D_2			After D_3			After D_4			After D_5		
		LA	RA	H-mean	LA	RA	H-mean	LA	RA	H-mean	LA	RA	H-mean	LA	RA	H-mean
Teacher	Full-Batch	0.1135	0.1145	0.1140	0.1096	0.1164	0.1129	0.1070	0.1159	0.1113	0.1064	0.1141	0.1101	0.1044	0.1151	0.1095
	Fine-Tune	0.1210	0.1001	0.1096	0.1122	0.0716	0.0874	0.1087	0.0562	0.0741	0.1058	0.0488	0.0668	0.1060	0.0548	0.0722
	LWC-KD-PIW	0.1235	0.1022	0.1118	0.1141	0.0675	0.0848	0.1118	0.0552	0.0739	0.1083	0.0491	0.0676	0.1075	0.0499	0.0682
	ReLoop2	0.1051	0.0836	0.0931	0.0949	0.0563	0.0707	0.0927	0.0550	0.0690	0.0898	0.0514	0.0654	0.0894	0.0469	0.0615
	CCD w/ MF student	0.1253	0.1128	0.1187	0.1214*	0.1007*	0.1101*	0.1231*	0.0980*	0.1091*	0.1213*	0.0940*	0.1059*	0.1197*	0.0921*	0.1041*
CCD w/ GNN student	0.1244	0.1079	0.1156*	0.1202	0.0991*	0.1086*	0.1212	0.0971*	0.1078*	0.1209	0.0950*	0.1064*	0.1198*	0.0914*	0.1037*	
Student (MF)	Full-Batch	0.0805	0.0805	0.0805	0.0766	0.0785	0.0775	0.0746	0.0755	0.0750	0.0723	0.0764	0.0743	0.0705	0.0767	0.0735
	LWC-KD-PIW	0.0811	0.0795	0.0803	0.0770	0.0621	0.0688	0.0743	0.0591	0.0658	0.0725	0.0553	0.0627	0.0719	0.0516	0.0601
	ReLoop2	0.0796	0.0796	0.0796	0.0775	0.0778	0.0776	0.0734	0.0748	0.0741	0.0699	0.0690	0.0694	0.0671	0.0643	0.0657
	KD + Fine-Tune	0.0882	0.1058	0.0962	0.0858	0.0962	0.0907	0.0846	0.0879	0.0862	0.0843	0.0822	0.0832	0.0806	0.0757	0.0781
	KD + LWC-KD-PIW	0.0863	0.1037	0.0942	0.0867	0.0995	0.0927	0.0843	0.0885	0.0863	0.0842	0.0824	0.0833	0.0833	0.0735	0.0781
CCD	0.1037*	0.1043	0.1040*	0.0978*	0.1004	0.0991*	0.0923*	0.0902	0.0912*	0.0901*	0.0873	0.0887*	0.0885*	0.0870*	0.0877*	
Student (GNN)	Full-Batch	0.0873	0.0894	0.0883	0.0840	0.0874	0.0857	0.0824	0.0884	0.0853	0.0817	0.0898	0.0856	0.0799	0.0890	0.0842
	LWC-KD-PIW	0.0937	0.0726	0.0818	0.0878	0.0495	0.0633	0.0876	0.0442	0.0588	0.0856	0.0383	0.0529	0.0856	0.0383	0.0529
	ReLoop2	0.0825	0.0726	0.0772	0.0769	0.0550	0.0641	0.0765	0.0494	0.0600	0.0752	0.0528	0.0620	0.0762	0.0510	0.0611
	KD + Fine-Tune	0.0997	0.1022	0.1009	0.0992	0.0940	0.0965	0.0970	0.0869	0.0917	0.0953	0.0845	0.0896	0.0974	0.0808	0.0883
	KD + LWC-KD-PIW	0.0986	0.1007	0.0996	0.0974	0.0928	0.0950	0.0954	0.0874	0.0912	0.0936	0.0843	0.0887	0.0958	0.0809	0.0877
CCD	0.1136*	0.1054*	0.1093*	0.1084*	0.0982	0.1030*	0.1097*	0.0930	0.1007*	0.1090*	0.0913	0.0994*	0.1100*	0.0878	0.0977*	

Table 3: The overall performance (NDCG@20) comparison. * denotes $p < 0.05$ for the paired t-test on CCD with the best baseline.

Gowalla		After D_1			After D_2			After D_3			After D_4			After D_5		
		LA	RA	H-mean	LA	RA	H-mean	LA	RA	H-mean	LA	RA	H-mean	LA	RA	H-mean
Teacher	Full-Batch	0.0897	0.0893	0.0895	0.0879	0.0900	0.0889	0.0864	0.0874	0.0869	0.0863	0.0870	0.0866	0.0864	0.0860	0.0862
	Fine-Tune	0.0858	0.0752	0.0802	0.0848	0.0457	0.0594	0.0837	0.0380	0.0523	0.0838	0.0329	0.0472	0.0852	0.0258	0.0408
	LWC-KD-PIW	0.0919	0.0725	0.0811	0.0889	0.0543	0.0674	0.0875	0.0477	0.0617	0.0876	0.0426	0.0573	0.0885	0.0390	0.0541
	ReLoop2	0.0853	0.0769	0.0809	0.0777	0.0697	0.0735	0.0741	0.0589	0.0656	0.0721	0.0506	0.0595	0.0702	0.0455	0.0552
	CCD w/ MF student	0.0915	0.0804	0.0856*	0.0877	0.0751*	0.0797*	0.0857	0.0681*	0.0759*	0.0851	0.0651*	0.0738*	0.0853	0.0606*	0.0709*
CCD w/ GNN student	0.0903	0.0794	0.0845*	0.0871	0.0718	0.0787*	0.0854	0.0660*	0.0745*	0.0848	0.0628*	0.0722*	0.0852	0.0584*	0.0693*	
Student (MF)	Full-Batch	0.0664	0.0665	0.0664	0.0638	0.0645	0.0641	0.0635	0.0649	0.0642	0.0633	0.0652	0.0642	0.0634	0.0653	0.0643
	LWC-KD-PIW	0.0620	0.0554	0.0585	0.0619	0.0426	0.0505	0.0616	0.0427	0.0504	0.0620	0.0368	0.0462	0.0641	0.0328	0.0434
	ReLoop2	0.0550	0.0554	0.0552	0.0515	0.0506	0.0510	0.0491	0.0473	0.0482	0.0478	0.0418	0.0446	0.0472	0.0382	0.0422
	KD + Fine-Tune	0.0643	0.0731	0.0684	0.0610	0.0549	0.0578	0.0607	0.0481	0.0537	0.0603	0.0485	0.0538	0.0611	0.0389	0.0475
	KD + LWC-KD-PIW	0.0641	0.0726	0.0681	0.0607	0.0601	0.0604	0.0601	0.0486	0.0537	0.0612	0.0419	0.0497	0.0616	0.0392	0.0479
CCD	0.0733*	0.0738	0.0735*	0.0653*	0.0683*	0.0668*	0.0622	0.0674*	0.0647*	0.0602	0.0643*	0.0622*	0.0591	0.0600*	0.0595*	
Student (GNN)	Full-Batch	0.0709	0.0707	0.0708	0.0700	0.0719	0.0709	0.0696	0.0706	0.0701	0.0696	0.0702	0.0699	0.0701	0.0701	0.0701
	LWC-KD-PIW	0.0692	0.0496	0.0578	0.0682	0.0398	0.0503	0.0672	0.0405	0.0505	0.0665	0.0384	0.0487	0.0682	0.0278	0.0395
	ReLoop2	0.0550	0.0554	0.0552	0.0515	0.0506	0.0510	0.0491	0.0473	0.0482	0.0478	0.0418	0.0446	0.0472	0.0382	0.0422
	KD + Fine-Tune	0.0755	0.0697	0.0725	0.0731	0.0591	0.0654	0.0717	0.0477	0.0573	0.0714	0.0448	0.0551	0.0732	0.0385	0.0505
	KD + LWC-KD-PIW	0.0751	0.0695	0.0722	0.0729	0.0589	0.0652	0.0715	0.0491	0.0582	0.0712	0.0454	0.0554	0.0730	0.0381	0.0501
CCD	0.0716	0.0715	0.0715	0.0707	0.0674*	0.0690*	0.0703	0.0650*	0.0675*	0.0704	0.0620*	0.0659*	0.0712	0.0557*	0.0625*	

Table 5: Results of various ablations on teacher system.

Method	After D_4			After D_5		
	LA	RA	H-mean	LA	RA	H-mean
CCD	0.1209	0.0950	0.1064	0.1198	0.0914	0.1037
w/o student-side knowledge	0.1083	0.0491	0.0676	0.1075	0.0499	0.0682
w/o proxies (student only)	0.0933	0.0683	0.0789	0.0956	0.0666	0.0785
w/o annealing	0.1021	0.0645	0.0791	0.1034	0.0649	0.0797

Table 6: Results of various ablations on student model.

Method	After D_4			After D_5			
	LA	RA	H-mean	LA	RA	H-mean	
CCD	0.1090	0.0913	0.0994	0.1100	0.0878	0.0977	
Proxy-guided learning	w/o proxy learning	0.0972	0.0656	0.0783	0.0965	0.0596	0.0737
	w/o S-proxy	0.0959	0.0713	0.0818	0.0965	0.0727	0.0829
	w/o P-proxy	0.0981	0.0694	0.0813	0.0981	0.0713	0.0826
Proxy update	extreme stability	0.1035	0.0904	0.0965	0.1052	0.0869	0.0952
	extreme plasticity	0.1062	0.0910	0.0980	0.1075	0.0871	0.0962

5.3 Study of CCD

We provide comprehensive analyses for an in-depth understanding of CCD. We report the results with the GNN-based student on Yelp.

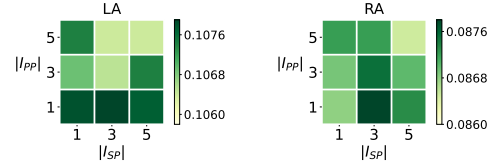
5.3.1 Ablation study. We present the ablation study of CCD to demonstrate the effectiveness of each proposed component.

Teacher side. Table 5 presents the ablation results on the teacher. From §4.4, we exclude three proposed components: (1) ‘w/o student-side knowledge’ excludes $\mathcal{L}_{S \rightarrow T}$, (2) ‘w/o proxies (student only)’ only uses the updated student excluding the proxies in $\mathcal{L}_{S \rightarrow T}$, and (3) ‘w/o annealing’ excludes the loss annealing. First, excluding the student-side knowledge largely degrades the effectiveness of CCD, and the best performance is achieved by leveraging both the updated student and the proxies accumulating the historical knowledge. Also, the simple annealing that gradually reduces the impacts of the student-side knowledge effectively improves the teacher.

Student side. Table 6 presents the ablation results on the student. First, we examine the effects of proxy-guided replay learning (§4.3.2). ‘w/o proxy learning’ excludes \mathcal{L}_{RE} , and ‘w/o S-/P-proxy’ exclude the stability proxy and the plasticity proxy, respectively. We observe that both proxies are indeed beneficial for the student, as they provide complementary views of the previous knowledge, as shown in Figure 3. Second, we investigate the effect of w_{SP} and w_{PP} , used to update the proxies in Eq.4. We assess the extreme case for each proxy. When $w_{SP} = 0$, the stability proxy does not accept any recent knowledge (i.e., extreme stability). Conversely, when $w_{PP} = 1$, the plasticity proxy is entirely replaced by the recent knowledge (i.e., extreme plasticity). The best performance is achieved when both proxies are updated with an appropriate balance. However, as long as the proxies accumulate past knowledge, the choice of update weights has little impact on the final performance. We set $w_{SP} = 0.1$ and $w_{PP} = 0.9$, respectively.

5.3.2 Impacts of the replay size. Figure 5 presents the performance (Recall@20) of the student model trained with varying replay sizes in Eq.6. The best performance is achieved with a small sampling size (≤ 3), suggesting that our proxy-guided replay learning does not notably increase the training complexity.

5.3.3 Analysis on dormant/new users. We further analyze the recommendation quality of the student for specific user groups having distinct characteristics. In Table 7, we present the recommendation performance for *dormant users* who haven’t used the

**Figure 5: Performance with varying replay sizes.****Table 7: Recommendation performance on dormant users.**

Method	Recall@20	NDCG@20
LWC-KD-PIW	0.0081	0.0024
KD + Fine-Tune	0.0645	0.0286
CCD	0.0806	0.0311

Table 8: Effects of new entity initialization on new users.

	Student		Teacher	
	Recall@20	NDCG@20	Recall@20	NDCG@20
CCD	0.1313	0.0559	0.1500	0.0605
w/ random init.	0.1270	0.0546	0.1387	0.0573
w/ one-hop init.	0.1116	0.0436	0.1037	0.0420

system for an extended period. We select users who were active in D_1 , remained dormant from D_2 to D_4 , and became active again in D_5 . We evaluate performance for the users on the test set of D_5 , after updating the student up to D_4 . We compare the baselines that show high effectiveness in previous experiments. We observe that CCD generates more accurate recommendations for the dormant users, indicating that CCD effectively preserves previous knowledge.

In Table 8, we present the average recommendation performance for *new users* for each data block. CCD proposes to use the prominent-entity information to facilitate the understanding of new entities (§4.3.1). To assess its effectiveness, we report the results when replacing the proposed technique with conventional random initialization and the widely used 1-hop initialization [34, 41]. The proposed technique effectively enhances the student’s ability to adapt to new entities. Moreover, as the teacher subsequently leverages the student-side knowledge, it leads to further improvement in teacher performance. These results collectively show that CCD achieves a good balance between stability and plasticity.

6 CONCLUSION

We propose CCD framework for effective and efficient deployment through a compact student model having the high performance of the massive teacher system, while aptly adapting to continuously incoming data. Unlike the existing KD studies have focused on one-time distillation in static environments, CCD updates both the teacher and the student continually and collaboratively along the data stream. CCD facilitates the student’s effective adaptation to new data, while also enabling the teacher to fully leverage accumulated knowledge. Given its great compatibility with existing models, we expect that our CCD framework can provide a better accuracy-efficiency trade-off for practical recommender systems.

ACKNOWLEDGEMENT

This work was supported by the IITP grant funded by the MSIT (South Korea, No.2018-0-00584, RS-2019-II191906), the NRF grant funded by the MSIT (South Korea, No.RS-2023-00217286, No.2020R1A2B5B03097210), the TIP funded by the MOTIE (South Korea, No.20014926), and the DIP grant funded by the MSIT and Daegu Metropolitan City (South Korea, No. DBSD1-07).

REFERENCES

- [1] Kian Ahrabian, Yishi Xu, Yingxue Zhang, Jiapeng Wu, Yuening Wang, and Mark Coates. 2021. Structure aware experience replay for incremental learning in graph-based recommender systems. In *CIKM*. 2832–2836.
- [2] Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. Tallrec: An effective and efficient tuning framework to align large language model with recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 1007–1014.
- [3] Guohao Cai, Jieming Zhu, Quanyu Dai, Zhenhua Dong, Xiuqiang He, Ruiming Tang, and Rui Zhang. 2022. ReLoop: A Self-Correction Continual Learning Loop for Recommender Systems. In *SIGIR*. 2692–2697.
- [4] Yankai Chen, Huifeng Guo, Yingxue Zhang, Chen Ma, Ruiming Tang, Jingjie Li, and Irwin King. 2022. Learning Binarized Graph Representations with Multi-faceted Quantization Reinforcement for Top-K Recommendation. In *KDD*.
- [5] Jaime Hieu Do and Hady W Lauw. 2023. Continual Collaborative Filtering Through Gradient Alignment. In *RecSys*. 1133–1138.
- [6] Xiacong Du, Bhargav Bhushanam, Jiecao Yu, Dhruv Choudhary, Tianxiang Gao, Sherman Wong, Louis Feng, Jongsoo Park, Yu Cao, and Arun Kejariwal. 2021. Alternate model growth and pruning for efficient training of recommendation systems. In *20th IEEE International Conference on Machine Learning and Applications (ICMLA)*. IEEE, 1421–1428.
- [7] Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2022. Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5). In *RecSys*. 299–315.
- [8] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *SIGIR*.
- [9] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *WWW*.
- [10] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the knowledge in a neural network. In *NeurIPS*.
- [11] Cheng-Kang Hsieh, Longqi Yang, Yin Cui, Tsung-Yi Lin, Serge Belongie, and Deborah Estrin. 2017. Collaborative metric learning. In *WWW*.
- [12] SeongKu Kang, Junyoung Hwang, Wonbin Kweon, and Hwanjo Yu. 2020. DE-RRD: A Knowledge Distillation Framework for Recommender System. In *CIKM*.
- [13] SeongKu Kang, Junyoung Hwang, Wonbin Kweon, and Hwanjo Yu. 2021. Item-side ranking regularized distillation for recommender system. *Information Sciences* 580 (2021), 15–34. <https://doi.org/10.1016/j.ins.2021.08.060>
- [14] SeongKu Kang, Junyoung Hwang, Wonbin Kweon, and Hwanjo Yu. 2021. Topology Distillation for Recommender System. In *KDD*.
- [15] SeongKu Kang, Junyoung Hwang, Dongha Lee, and Hwanjo Yu. 2019. Semi-supervised learning for cross-domain recommendation to cold-start users. In *CIKM*.
- [16] SeongKu Kang, Wonbin Kweon, Dongha Lee, Jianxun Lian, Xing Xie, and Hwanjo Yu. 2023. Distillation from Heterogeneous Models for Top-K Recommendation. In *WWW*. 801–811.
- [17] SeongKu Kang, Wonbin Kweon, Dongha Lee, Jianxun Lian, Xing Xie, and Hwanjo Yu. 2024. Unbiased, Effective, and Efficient Distillation from Heterogeneous Models for Recommender Systems. *ACM Trans. Recomm. Syst.* (feb 2024). <https://doi.org/10.1145/3649443>
- [18] SeongKu Kang, Dongha Lee, Wonbin Kweon, Junyoung Hwang, and Hwanjo Yu. 2022. Consensus Learning from Heterogeneous Objectives for One-Class Collaborative Filtering. In *WWW*.
- [19] SeongKu Kang, Dongha Lee, Wonbin Kweon, and Hwanjo Yu. 2022. Personalized Knowledge Distillation for Recommender System. *Knowledge-Based Systems* 239 (2022), 107958. <https://doi.org/10.1016/j.knsys.2021.107958>
- [20] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* 114, 13 (2017), 3521–3526.
- [21] Wonbin Kweon, SeongKu Kang, and Hwanjo Yu. 2021. Bidirectional Distillation for Top-K Recommender System. In *WWW*.
- [22] Dongha Lee, SeongKu Kang, Hyunjun Ju, Chanyoung Park, and Hwanjo Yu. 2021. Bootstrapping User and Item Representations for One-Class Collaborative Filtering. In *SIGIR*.
- [23] Youngjune Lee, Yeongjong Jeong, Keunchan Park, and SeongKu Kang. 2023. MvFS: Multi-view Feature Selection for Recommender System. In *CIKM*. 4048–4052.
- [24] Youngjune Lee and Kee-Eung Kim. 2021. Dual Correction Strategy for Ranking Distillation in Top-N Recommender System. In *CIKM*.
- [25] Dawen Liang, Rahul G. Krishnan, Matthew D. Hoffman, and Tony Jebara. 2018. Variational Autoencoders for Collaborative Filtering. In *WWW*.
- [26] Guoliang Lin, Hanlu Chu, and Hanjiang Lai. 2022. Towards better plasticity-stability trade-off in incremental learning: A simple linear connector. In *CVPR*. 89–98.
- [27] John I Marden. 1996. *Analyzing and modeling rank data*. CRC Press.
- [28] James L McClelland, Bruce L McNaughton, and Randall C O'Reilly. 1995. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review* 102, 3 (1995), 419.
- [29] Fei Mi, Xiaoyu Lin, and Boi Faltings. 2020. Ader: Adaptively distilled exemplar replay towards continual learning for session-based recommendation. In *RecSys*. 408–413.
- [30] Quang Pham, Chenghao Liu, and Steven Hoi. 2021. Dualnet: Continual learning, fast and slow. In *NeurIPS*. 16131–16144.
- [31] Sashank Reddi, Rama Kumar Pasumarthi, Aditya Menon, Ankit Singh Rawat, Felix Yu, Seungyeon Kim, Andreas Veit, and Sanjiv Kumar. 2021. Rankdistil: Knowledge distillation for ranking. In *AISTATS*. PMLR.
- [32] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *UAI*.
- [33] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. 2015. Fitnets: Hints for thin deep nets. In *ICLR*.
- [34] Wei Shen, Chuheng Zhang, Yun Tian, Liang Zeng, Xiaonan He, Wanchun Dou, and Xiaolong Xu. 2021. Inductive Matrix Completion Using Graph Autoencoder. In *CIKM*. 1609–1618.
- [35] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. 2017. Continual learning with deep generative replay. In *NeurIPS*.
- [36] Jiayi Tang and Ke Wang. 2018. Ranking distillation: Learning compact ranking models with high performance for recommender system. In *KDD*.
- [37] Antti Tarvainen and Harri Valpola. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*.
- [38] Menghan Wang, Yujie Lin, Guli Lin, Keping Yang, and Xiao-ming Wu. 2020. M2GRL: A multi-task multi-view graph representation learning framework for web-scale recommender systems. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 2349–2358.
- [39] Yuening Wang, Yingxue Zhang, and Mark Coates. 2021. Graph structure aware contrastive knowledge distillation for incremental learning in recommender systems. In *CIKM*. 3518–3522.
- [40] Yuening Wang, Yingxue Zhang, Antonios Valkanias, Ruiming Tang, Chen Ma, Jianye Hao, and Mark Coates. 2023. Structure aware incremental learning with personalized imitation weights for recommender systems. In *AAAI*. 4711–4719.
- [41] Yunfan Wu, Qi Cao, Huawei Shen, Shuchang Tao, and Xueqi Cheng. 2022. Inmo: A model-agnostic and scalable module for inductive collaborative filtering. In *SIGIR*. 91–101.
- [42] Fen Xia, Tie-Yan Liu, Jue Wang, Wensheng Zhang, and Hang Li. 2008. Listwise approach to learning to rank: theory and algorithm. In *ICML*.
- [43] Xin Xia, Hongzhi Yin, Junliang Yu, Qinyong Wang, Guandong Xu, and Quoc Viet Hung Nguyen. 2022. On-Device Next-Item Recommendation with Self-Supervised Knowledge Distillation. In *SIGIR*.
- [44] Yishi Xu, Yingxue Zhang, Wei Guo, Huifeng Guo, Ruiming Tang, and Mark Coates. 2020. Graphsail: Graph structure aware incremental learning for recommender systems. In *CIKM*. 2861–2868.
- [45] Yang Yang, Da-Wei Zhou, De-Chuan Zhan, Hui Xiong, and Yuan Jiang. 2019. Adaptive deep models for incremental learning: Considering capacity scalability and sustainability. In *KDD*. 74–82.
- [46] Rex Ying, Ruiming He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 974–983.
- [47] Hansi Zeng, Hamed Zamani, and Vishwa Vinay. 2022. Curriculum Learning for Dense Retrieval Distillation. In *SIGIR*. 1979–1983.
- [48] Jieming Zhu, Guohao Cai, Junjie Huang, Zhenhua Dong, Ruiming Tang, and Weinan Zhang. 2023. ReLoop2: Building Self-Adaptive Recommendation Models via Responsive Error Compensation Loop. In *KDD*.
- [49] Jieming Zhu, Jinyang Liu, Weiqi Li, Jincai Lai, Xiuqiang He, Liang Chen, and Zibin Zheng. 2020. Ensembled CTR Prediction via Knowledge Distillation. In *CIKM*.
- [50] Yaochen Zhu, Liang Wu, Qi Guo, Liangjie Hong, and Jundong Li. 2024. Collaborative Large Language Model for Recommender Systems. In *Proceedings of the ACM on Web Conference 2024 (WWW '24)*. ACM. <https://doi.org/10.1145/3589334.3645347>

A APPENDIX

The source code of CCD is publicly available through the author’s GitHub repository.²

A.1 List-wise distillation

The list-wise distillation [12, 16, 24] trains the student to emulate the item permutation (i.e., ranking orders) predicted by the teacher. Specifically, they define a probability for observing each permutation based on the Plackett-Luce model [27], then train the student to maximize the likelihood of the teacher permutations [42]. For each user u , the item ranking list predicted by RS^T is denoted by π_u^T . The list-wise KD loss is defined as the negative log-likelihood of permutation probability [16]:

$$\mathcal{L}_{KD} = - \sum_{u \in \mathcal{U}} \log p(\pi_u^T | S) = - \sum_{u \in \mathcal{U}} \log \prod_{n=1}^N \frac{\exp(\hat{r}_{u, \pi_u^T(n)}^S)}{\sum_{i=n}^{|\mathcal{I}|} \exp(\hat{r}_{u, \pi_u^T(i)}^S)}. \quad (10)$$

$\pi_u^T(n)$ denotes the n -th item in π_u^T . N is a hyperparameter reflecting the length of the recommendation list. In this work, we set $N = 50$ [16]. By minimizing the loss, the student model learns to preserve the detailed ranking orders of top- N items in π_u^T , while penalizing the remaining items below the N -th rank.

This distillation approach has shown remarkable performance in many applications such as recommendation [12, 16, 17, 24] and document retrieval [31, 47]. Although more complicated variants have been studied, we obtained satisfactory results with the default distillation loss. We assessed the effectiveness of this distillation in our experimental setup, and the results are presented in Table 10.

A.2 Experiment details

A.2.1 Dataset. We use two public datasets with real-world user-item interactions: Gowalla³ and Yelp⁴. We filter out users and items with fewer than 10 interactions in both datasets. The total data statistics after preprocessing are summarized in Table 9.

Table 9: Dataset statistics.

Dataset	#Users	#Items	#Interactions	Sparsity(%)
Gowalla	29,858	40,988	1,027,464	99.91
Yelp	14,950	12,261	342,617	99.81

A.2.2 Teacher and student configuration. We provide the detailed configuration of the teacher and the student used in the experiments. Table 10 presents the number of learning parameters as well as the effectiveness of distillation in our setup.

Teacher. Following the teacher configuration of the recent KD work [16], we construct a massive teacher system by ensembling multiple large models. Initially, we evaluate the effectiveness of five different recommendation models for each dataset: matrix factorization (MF) [32], metric learning [11], deep neural network [9], graph neural network (GNN) [8], and variational autoencoder (VAE) [25]. Subsequently, we select up to two models for each dataset that

demonstrate high performance. Lastly, we incrementally increase the capacity of the teacher system by augmenting the number of models or dimension sizes until its performance no longer improves. For Gowalla, we construct the teacher system using five MF-based models [32] and two VAE-based models [25], each initialized with distinct random seeds and set to 64 dimensions. For Yelp, we employ five GNN-based models [8], each initialized with distinct random seeds and set to 128 dimensions. We adopt the ensemble scheme utilized in [16].

Student. For the student, we employ two backbone models: MF-based model [32] and GNN-based model [8]. We set a small embedding size for the student (16 for Gowalla and 8 for Yelp), considering the teacher size for each dataset.

Effectiveness of KD. We assessed the effectiveness of the list-wise KD technique (A.1) on our teacher-student setup. The recommendation performance after the distillation is presented in Table 10.

Table 10: Recall@20 results after distillation for each data block. #Params. denotes the number of parameters.

Dataset		D_1	D_2	D_3	D_4	#Params.
Gowalla	Teacher	0.1849	0.1708	0.1698	0.1768	33.47M
	Student	0.1822	0.1686	0.1652	0.1677	1.13M
Yelp	Teacher	0.1237	0.1117	0.1244	0.1195	18.24M
	Student	0.1161	0.1138	0.1256	0.1147	0.22M

²https://github.com/Gyu-Seok0/CCD_KDD24

³<https://snap.stanford.edu/data/loc-gowalla.html>

⁴<https://www.yelp.com/dataset>