



GraphPro: Graph Pre-training and Prompt Learning for Recommendation

Yuhao Yang
University of Hong Kong
Hong Kong, China
yuhao-yang@outlook.com

Lianghao Xia
University of Hong Kong
Hong Kong, China
aka_xia@foxmail.com

Da Luo
Wechat, Tencent
Guangzhou, China
lodalu@tencent.com

Kangyi Lin
Wechat, Tencent
Guangzhou, China
plancklin@tencent.com

Chao Huang*
University of Hong Kong
Hong Kong, China
chaohuang75@gmail.com

ABSTRACT

GNN-based recommendation systems have been successful in capturing complex user-item interactions using multi-hop message passing. However, most of these methods are designed, trained and evaluated on a single, static graph, thus struggling to handle the dynamic nature of user-item interactions. This limits their scalability and performance in real-world dynamic scenarios. In our study, we propose a framework called GraphPro that combines dynamic graph pre-training with prompt learning in an efficient way. This unique approach allows GNNs to effectively capture both long-term user preferences and short-term behavior changes in a dynamic environment, resulting in accurate and up-to-date recommendations. To address the issue of changing user preferences, we integrate a temporal prompt mechanism and a graph-structural prompt learning mechanism into the pre-trained GNN architecture. The temporal prompt mechanism incorporates time-related information into user-item interactions, enabling the model to naturally incorporate temporal dynamics. The graph-structural prompt learning mechanism allows the model to apply pre-trained insights to new behavior dynamics without the need for continuous retraining. We also introduce a dynamic evaluation framework for recommendations that better reflects real-world scenarios and reduces the offline-online discrepancy. Through comprehensive experiments, including deployment in a large-scale industrial scenario, we demonstrate the seamless scalability of GraphPro with various leading recommenders. Our results highlight the superiority of GraphPro in terms of effectiveness, robustness, and efficiency. We release the model implementation at the link: <https://github.com/HKUDS/GraphPro>.

CCS CONCEPTS

• Information systems → Recommender systems.

*Chao Huang is the corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WWW'24, May 13–17, 2024, Singapore, Singapore

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0171-9/24/05...\$15.00
<https://doi.org/10.1145/3589334.3645546>

KEYWORDS

Graph Pre-Training, Prompt Learning, Recommendation

ACM Reference Format:

Yuhao Yang, Lianghao Xia, Da Luo, Kangyi Lin, and Chao Huang. 2024. GraphPro: Graph Pre-training and Prompt Learning for Recommendation. In *Proceedings of the ACM Web Conference 2024 (WWW'24)*, May 13–17, 2024, Singapore, Singapore. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3589334.3645546>

1 INTRODUCTION

Recommender systems are integral to numerous Web platforms, assisting users in navigating through the overwhelming amount of information by suggesting relevant items. In recent years, graph neural networks (GNNs) have emerged as powerful tools for modeling user-item interactions in recommendation tasks, enabling effective representation learning on graph-structured data. By treating users and items as nodes and their interactions as edges, GNNs can capture intricate multi-hop relationships between users and items, facilitating the generation of personalized recommendation results.

Previous studies in GNN-augmented recommendation have centered on crafting effective message passing strategies to delineate the collaborative relations between users and items [1, 47, 55]. These efforts seek to harness GCN's potential to unravel high-order connections within the user-item interaction graph. Follow-up investigations have delved into streamlining the message-passing approach [5, 15], diminishing GNN model intricacies [30, 38], and advancing sampling method quality [21]. Lately, the field has progressed with the integration of self-supervised learning (SSL) into GNN frameworks for recommendation [35]. These cutting-edge methodologies [2, 49, 58] predominantly utilize the InfoNCE loss function [31] to align contrastive perspectives, thus enhancing the performance of the LightGCN architecture [15].

Although these methods demonstrate impressive performance, existing methods have predominantly focused on static recommendation scenarios (Fig 1 upper), neglecting the dynamic nature of recommendations in real-world settings (Fig 1 lower) [51, 54]. In such dynamic environments, recommenders continuously learn from new data and make predictions in real-time. However, these methods are primarily designed for single-graph training and evaluation, resulting in a decline in recommendation dynamics and widening the gap between offline and online performance [22]. Additionally, the arrival of new data often brings distribution shifts,




Figure 1: Our dynamic recommendation setting compared to the vanilla single-graph training in existing methods.

further complicating accurate recommendations for graph-based models that do not incorporate relevant contextual information for the newly arrived data. These challenges severely limit the scalability of existing models and impede their ability to adapt to evolving user preferences promptly. This becomes crucial for providing timely and accurate recommendations in dynamic environments.

To address these challenges, we present a simple and effective framework called GraphPro, which combines parameter-efficient and dynamic Graph Pre-Training with Prompt Learning for recommendation. Our method begins with pre-training Graph Neural Networks (GNNs) on extensive historical interaction data, enabling them to capture long-term user preferences and item relevance. This pre-training phase assimilates knowledge from a substantial amount of historical interactions. Subsequently, during the fine-tuning phase on more recent target data, our model swiftly adapts to evolving user preferences and captures short-term behavior dynamics. This is achieved through a prompt learning schema, which facilitates effective knowledge transfer.

To ensure the effective handling of evolving user preferences by the pre-trained GNN, our GraphPro framework seamlessly integrates a temporal prompt mechanism and a graph-structural prompt learning mechanism. This integration allows for the injection of time-aware context from new data, enabling the model to adapt to changing user preferences. Inspired by advancements in relative positional encoding techniques [33, 41], we meticulously design a dedicated temporal prompt mechanism that aligns with the message aggregation layer of GNNs. Within this prompt mechanism, we encode time information on interaction edges as part of the normalization term for aggregation, all in a parameter-free manner. This innate capability allows the model to naturally incorporate temporal information without the need for additional fine-tuning. By incorporating temporal awareness into the pre-trained graph model, we empower the model to effectively capture vital signals that are highly relevant to the evolving user preferences.

Moreover, our graph-structural prompt learning mechanism facilitates the seamless transfer of knowledge from the pre-trained model to downstream recent recommendation tasks. This framework eliminates the need for continuous incremental learning of the pre-trained model, enabling the transfer of pre-trained knowledge to any future time period to effectively adapt to behavior dynamics. In this mechanism, we include the newly generated interaction edges between the fine-tuning time and the pre-training time as prompt edges, providing the pre-trained model with essential contextual information for the fine-tuning process. Rather than undergoing extensive training, we perform a single non-training forward pass on the prompt edges, prompting the pre-trained model to adapt to the distribution shift of node representations and effectively adjust its predictions. It’s worth mentioning that our GraphPro is model-agnostic and parameter-efficient, making it easy to integrate

into existing GNN recommenders as a plug-in enhancement. In summary, the main contributions of our work are as follows:

- We emphasize the importance of effectively and scalably pre-training and fine-tuning graph-based recommenders to accommodate time-evolving user preferences, thus facilitating up-to-date and accurate recommendations in dynamic environments.
- We introduce GraphPro, which effectively handles evolving user preferences through the pre-training and fine-tuning of GNNs. The proposed prompt learning paradigm enables the transfer of relevant knowledge from the pre-trained model to downstream recommendation tasks in both temporal and structural ways.
- Furthermore, we introduce a snapshot-based dynamic setting for recommender system evaluation, which offers a more realistic approximation to real-world recommendation scenarios compared to the traditional single-time testing approach.
- We conduct experiments on diverse datasets to showcase the robustness, efficiency, and performance advantages of GraphPro. To further validate the effectiveness of our framework, we also present an industry deployment on a large-scale online platform.

2 PRELIMINARIES

We define the task of pre-training and fine-tuning GNNs for recommendation. We denote the user set as \mathcal{U} and the item set as \mathcal{I} . In the context of collaborative filtering, a typical graph structure, constructed using existing methods [15], can be represented as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \mathcal{U} \cup \mathcal{I}$ represents the set of all nodes in user-item interaction graph \mathcal{G} . The edges in \mathcal{E} correspond to interactions between users and items, with a value of $y_{u,i} = 1$.

In order to provide recommendations at time slot T_1 , we gather historical user-item interactions to construct a graph $\mathcal{G}_1 = (\mathcal{V}_1, \mathcal{E}_1)$, where \mathcal{E}_1 represents the user-item interactions collected before T_1 . Existing stationary graph collaborative filtering models typically train the model from scratch using the complete dataset \mathcal{G}_1 . The objective is to optimize time-specific model parameters Θ_1 by maximizing the likelihood of generating accurate recommendations:

$$\operatorname{argmax}_{\Theta_1} P_{f_{\Theta_1}}(y_1 | \mathcal{G}_1). \quad (1)$$

Dynamic Learning in Recommender Systems. In real-world applications, the evaluation of recommenders goes beyond the simplistic static setting commonly used in existing collaborative filtering method [15, 47]. In practice, researchers assess the long-term performance of models by deploying them in a live-update environment, as discussed in [56], where new user-item interactions are continuously generated over time. The model is specifically designed to make ongoing predictions for future user-item interactions based on this evolving data. Formally, the model should have initial weights Θ_{n-1} corresponding to different time intervals T_n , which are then updated through learning on new interactions \mathcal{G}_n to enhance the accuracy of up-to-date predictions.

$$\operatorname{argmax}_{\Theta_n} P_{f_{\Theta_n}}(y_n | \mathcal{G}_n; \Theta_{n-1}). \quad (2)$$

In this study, we draw inspiration from the work [32] and employ a series of graph snapshots to simulate practical dynamic recommendation scenarios. These graph snapshots are represented as

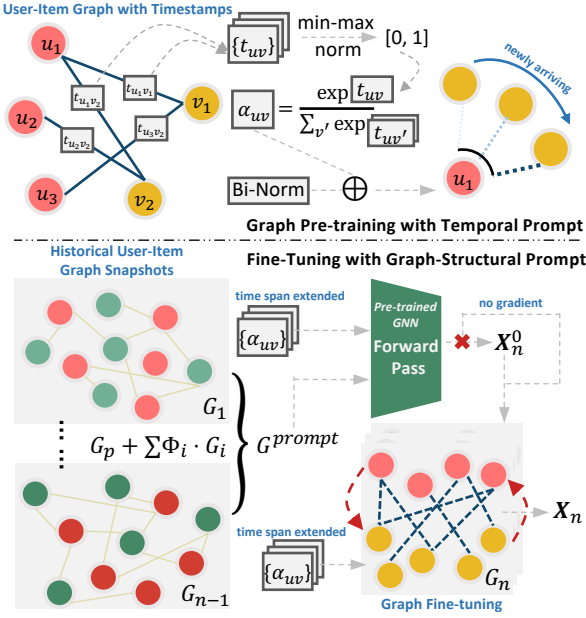


Figure 2: Overall framework of GraphPro.

$[\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_N]$, where we have a total of N snapshots corresponding to different time intervals. Each graph snapshot consists of subsets of users and items from a global set, and the interactions between them evolve over time. Specifically, $\mathcal{G}_n = (\mathcal{V}_n \subset \mathcal{V}, \mathcal{E}_n)$, where \mathcal{V}_n represents the nodes in the snapshot and \mathcal{E}_n denotes the time slot-specific interaction edges. It is important to emphasize that snapshots \mathcal{G}_n are collected within the time period between two consecutive snapshots, namely $[T_{n-1}, T_n]$.

3 METHODOLOGY

In this section, we provide the technical details of our proposed GraphPro framework, depicted in Figure 2, which illustrates its architecture. We introduce two crucial components: graph pre-training with a temporal prompt mechanism and a graph-structural prompt-enhanced fine-tuning mechanism. These components are specifically designed to enhance the performance and scalability of GNN-based recommenders in dynamic recommendation scenarios.

3.1 Graph Pre-training with Temporal Prompt Mechanism

In practical recommendation scenarios, user-item interaction data continues to accumulate over time. Online platforms like Amazon and TikTok constantly receive new user purchases and video watch logs, respectively, on a daily basis. In such dynamic environments, the availability of fresh user-item interactions provides valuable information that can be leveraged to guide pre-trained models in adapting to time-evolving user preferences and providing continuously up-to-date recommendations in dynamic settings.

3.1.1 Temporal Prompt Mechanism. In our framework, we introduce a temporal prompt mechanism to incorporate time-aware contextual information from the latest user preferences and behaviors. This mechanism allows for personalized and timely recommendations by considering the temporal dynamics of user-item

interactions. To capture the temporal sequence of user-item interactions, we propose a relative time encoding scheme, which enables us to incorporate temporal information into graph convolutions. By encoding the relative time between interactions, the model can explicitly capture the temporal dependencies and changes in user preferences that are reflected in the newly arrived data.

Our temporal prompt mechanism offers two significant advantages over existing time encoding techniques when it comes to capturing user behavior dependencies across different time slots.

- **Generalization.** Unlike the use of absolute positional embeddings in models like BERT [9], our mechanism takes inspiration from recent advancements in sequence modeling in NLP and leverages a relative positional encoding. Absolute positional embeddings have limited generalization capabilities across continuous time steps, which is problematic in our dynamic recommendation setting. These embeddings are trained on sequences with varying lengths but struggle to handle sequences beyond the trained lengths during fine-tuning and prediction. This leads to a distribution gap between the pretraining phase, where the model learns from fixed-length data, and the fine-tuning and prediction phase, where longer future time steps are encountered.
- **Scalability.** Our temporal prompt design avoids the need to add a fixed-length positional embedding to node representations. Instead, it generates relative temporal-aware weights that can be seamlessly integrated with message passing. This design allows our pretrained GNN to be easily applied to longer-range graph structures during fine-tuning and testing, greatly improving scalability for dynamic recommendation tasks.

3.1.2 Temporal Prompt-enhanced Graph Convolutions. In order to effectively capture the temporal dynamics of user-item interactions in our model, we have implemented a temporal prompt and incorporated relative time encoding into our graph convolutions. This enables our model to consider the most recent contextual signals from the new data, and adapt to evolving user preferences over time. Within the context of our user-item interaction graph \mathcal{G} , the edge attributes consist of Unix timestamps denoted as t^{unix} . These timestamps represent the exact moments when users u interacted with items v . To prepare these timestamps for encoding in our model, we convert them into relative time steps by dividing them by a fixed time interval τ . This time interval, which is a hyperparameter, can be defined with a resolution of either hour, day, or week. As a result, for any given edge $e_{u,v}$ in the graph, its corresponding timestep attribute can be computed as follows:

$$t_{u,v} = f_{\text{div}}(t_{u,v}^{\text{unix}}, t^{\text{unix}}) = \lfloor \frac{t_{u,v}^{\text{unix}} - \min(t^{\text{unix}})}{\tau} \rfloor, \quad (3)$$

Here, $t_{u,v}^{\text{unix}}$ is the Unix timestamp assigned to the edge $e_{u,v}$, and the $\lfloor * \rfloor$ notation denotes the floor operation. To avoid the influence of specific numerical scales and ensure uniformity, we normalize these time attributes $t = t_{u,v}|e_{u,v} = 1$ to the range of $[0, 1]$.

$$t = \frac{t - \min(t)}{\max(t) - \min(t)}. \quad (4)$$

To consider the temporal information among the neighbors during message aggregation in our GNNs, we apply the softmax function

to the time attributes $t_{u,v}$ of the first-order neighbors on the graph.

$$\alpha_{u,v} = \frac{e^{t_{u,v}}}{\sum_{v' \in \mathcal{N}_u} e^{t_{u,v'}}}, \quad (5)$$

To enable dynamic time-aware graph neural network (GNN) for recommendation pretraining, we introduce an additional normalization term, $\alpha_{u,v}$, into the message passing step of LightGCN. In this case, \mathcal{N}_u represents the neighbors of node u , and $\alpha_{u,v}$ encodes the weight for aggregating information from node v to node u .

$$\mathbf{x}_u^{(l)} = \sum_{v \in \mathcal{N}_u} \left(\frac{1}{2\sqrt{|\mathcal{N}_u||\mathcal{N}_v|}} + \frac{\alpha_{u,v}}{2} \right) \mathbf{x}_v^{(l-1)}, \quad (6)$$

We introduce a normalization term and apply mean-pooling to incorporate time-aware normalization into the original bidirectional graph normalization while preserving embedding magnitude.

Adaptability and Efficiency. The incorporation of the time-aware normalization term $\alpha_{u,v}$ into the message passing of LightGCN enhances the GNN's adaptability to evolving user-item interactions over time. By giving more weight to interactions that are closer in time during neighbor aggregation, the model becomes more attentive to the dynamic nature of user-item interactions and assigns higher importance to recent interactions. This alignment with the objective of recommendation tasks ensures that the model captures timely and relevant user preferences, leading to more accurate and personalized recommendations. Importantly, our time-aware regularization approach does not introduce additional embedding encoding. Instead, it dynamically generates graph regularization terms based on the relative time order, making it a lightweight and efficient solution. This design allows the model to handle varying absolute time lengths effectively, showcasing excellent generalization capabilities and requiring minimal computational overhead.

3.2 Fine-Tuning with Graph Prompt Mechanism

In this section, we will discuss how we effectively transfer knowledge from a pre-trained Graph Neural Network (GNN) model for fine-tuning with future user-item interactions. To begin the fine-tuning process at the target time T_n , which occurs after the pre-training time T_p , the intuitive way is to update the model parameters incrementally by simply fine-tuning. That is, we iteratively provide the model with data from the updated time intervals to fine-tune the node representations that were previously updated in the preceding time intervals. Therefore, the initial embeddings for fine-tuning at T_n are derived from the forward pass after the last fine-tuning step.

$$\mathbf{X}_n^0 = \text{forward}(\mathbf{X}_{n-1}; \mathcal{G}_{n-1}), \quad (7)$$

where $\text{forward}(\ast)$ represents the complete forward pass of the model, utilizing the last fine-tuned embeddings \mathbf{X}_{n-1} and the graph structure \mathcal{G}_{n-1} . This method has the advantage of directly capturing users' continuous interest changes within a specific time span. However, the incremental fine-tuning mechanism has two significant drawbacks. First, iteratively updating model parameters based on small-range interactions may lead the model to converge to a local optimum specific to that time period, limiting the potential for continuous fine-tuning on the updated representations in the future. Secondly, persistently updating the parameters of the pre-trained model can result in a significant computational burden.

3.2.1 Graph-Structural Prompt Mechanism. In our approach, we address the mentioned issues by leveraging the interaction edges between the pre-training time T_p and the current time T_n as prompt edges. This allows the pretrained model to directly fine-tune on future time periods without the need for iterative updates. Inspired by discrete prompt tuning in large language models [37, 39], we treat the edges of the graph during a specific time period as discrete prompts that guide the propagation of pretrained embeddings. This captures the representation shift between the pre-training and fine-tuning time points and provides better temporal-aware initial embeddings for fine-tuning. To generate prompt structures, we concatenate the pre-training graph structure with the sampled future edges between the pre-training and current fine-tuning time. This combination enables the model to capture the temporal dynamics and improve the effectiveness of fine-tuning:

$$\mathcal{G}^{\text{prompt}} = \mathcal{G}_p \oplus \sum_{i=1}^n \Phi_i \odot \mathcal{G}_i; \quad \Phi_i = \begin{cases} 1 - (i-1)\phi, & \phi > 0 \\ 1 + (n-i)\phi, & \phi < 0 \end{cases}, \quad (8)$$

where " \oplus " denotes graph concatenation and " \odot " denotes graph sampling. Here, a hyper-parameter ϕ is introduced as the sampling decay for prompt structures, where a positive ϕ suggests that we include more early structures and less recent ones, and vice versa. After generating the prompt structures, we proceed with a forward pass using the pretrained embeddings \mathbf{X}_p on the prompt graph to generate embeddings for fine-tuning. To mitigate the overfitting effect and improve generalization for more robust fine-tuning in our GraphPro framework, we introduce a random gating [3] mechanism that slightly perturbs the pre-trained embeddings.

$$\tilde{\mathbf{X}}_p = \mathbf{X}_p \odot \text{sigmoid}(\tilde{\mathbf{W}}\mathbf{X}_p + \tilde{\mathbf{b}}), \quad (9)$$

$$\mathbf{X}_n^0 = \text{forward}(\tilde{\mathbf{X}}_p; \mathcal{G}^{\text{prompt}}), \quad (10)$$

The non-learnable random gating weights, $\tilde{\mathbf{W}} \in \mathbb{R}^{d \times d}$ and $\tilde{\mathbf{b}} \in \mathbb{R}^d$, are generated from a Gaussian distribution. It's important to note that the relative time encoding also plays a vital role in facilitating the model's ability to sense relative temporal connections during the prompt propagation process. By propagating the embeddings learned from extensive pretraining over a large time period on the prompt edges, which include interactions from subsequent time periods, we achieve two objectives. Firstly, we enable the obtained embeddings to maintain stable user interests. Secondly, we swiftly capture changes in user interests within the subsequent time span. By refraining from directly training the embeddings on the short-term graph, we mitigate the risk of the model parameters becoming trapped in local optima. This approach grants us greater flexibility for subsequent fine-tuning and enables the model to more effectively adapt to users' evolving interests over time.

3.2.2 Prompt Learning with Adaptive Gating Mechanism.

To address the distribution shift in node representations between the time-aware graph snapshots \mathcal{G}_{n-1} and \mathcal{G}_n , we introduce a learnable gating mechanism that adaptively transforms the input embeddings \mathbf{X}_n^0 . This gating mechanism allows for modeling the changes in user/item representations over time, effectively preserving the informative signals necessary for making accurate future recommendations. We employ gradient truncation on \mathbf{X}_n^0 to prevent direct optimization of the large-scale pre-trained model. Instead,

we fine-tune \mathbf{X}_n^0 using newly interaction structural contexts \mathcal{G}_n to improve the accuracy of predictions at the target time interval T_n :

$$\tilde{\mathbf{X}}_n^0 = \mathbf{X}_n^0 \odot \text{sigmoid}(\mathbf{W}_l \mathbf{X}_n^0 + \mathbf{b}_l), \quad (11)$$

$$\mathbf{X}_n = \text{forward}(\tilde{\mathbf{X}}_n^0; \mathcal{G}_n), \quad (12)$$

At this stage, we have derived the user and item representations \mathbf{X}_n for making predictions starting from time T_n . In this process, the learnable weights \mathbf{W}_l and \mathbf{b}_l , which have the same size as the random gating, are introduced. To estimate the probability of user u interacting with item i , we calculate the dot product between the user and item representations \mathbf{x}_n^u , \mathbf{x}_n^i , denoted as $\hat{y}_{u,i} = \mathbf{x}_n^{u\top} \cdot \mathbf{x}_n^i$.

3.3 Model Learning and Discussion

3.3.1 Optimized Objective. In both the pre-training and fine-tuning stages, we define our training objectives based on optimizing the BPR loss. The BPR loss ensures that the predicted score for an observed interaction is higher than that of its unobserved negative samples. This loss function is commonly used in recommendation systems to model the preference ranking between items for individual users. By optimizing this loss, we aim to improve the model’s ability to accurately rank and predict user-item interactions.

$$\mathcal{L} = - \sum_{(u,i,j) \in D} \log \sigma(\hat{y}_{ui} - \hat{y}_{uj}), \quad (13)$$

In our training strategy, we utilize a dataset D that includes negative items j sampled at each training mini-batch. Our approach follows a two-stage process. In the first stage, we pre-train a GNN-based recommender on a large-time-scale graph until convergence. This involves training the model on a comprehensive set of historical data, allowing it to learn long-term patterns and user preferences. In the second stage, we fine-tune the pre-trained model on small-time-scale graph snapshots that include interactions from a more recent time period. This fine-tuning process helps the model adapt and capture short-term changes in user interests and item dynamics.

3.3.2 Interplotive Parameter Update. To ensure that the model parameters are learned in synchronization with the evolving user and item representations, it is important to update the pre-trained node embeddings over time steps. Inspired by the investigation in [32, 56], we propose an interpolative approach for updating the pre-trained user and item embeddings. Specifically, to estimate the best initial state for training at the next time step T_n , we combine the pre-trained embeddings with the embeddings learned within a sliding window $[T_{n-\omega}, T_{n-1}]$ using interpolation. This allows us to leverage both the long-term knowledge captured during pre-training and the recent changes observed within the sliding window, enabling the model to effectively adapt to the evolving dynamics of user-item interactions:

$$\mathbf{X}_n^{\text{init}} = \text{mean}(\mathbf{X}_p, \sum_{i=1}^{\omega} \frac{i \cdot \mathbf{X}_{n-i}}{\sum_{k=1}^{\omega} k}) \quad (14)$$

\mathbf{X} represents the model parameters, which correspond to the user and item embeddings. The left term of the equation calculates a weighted normalization of the weights $[\mathbf{X}_{n-1}, \dots, \mathbf{X}_{n-\omega}]$, where the more recent fine-tuned representations are given less weight. This weighting helps to mitigate the local optima effect, where the

Table 1: Statistics of the experimental datasets.

Statistics	Taobao	Koubei	Amazon
# Users	117,450	119,962	131,707
# Items	86,718	101,404	107,028
# Interactions	8,795,404	3,986,609	876,237
# Density	8.6e-4	3.3e-4	6.2e-5
Temporal Segmentation			
# Pre-training Span	5 days	4 weeks	4 weeks
# Tuning-Predicting Span	5 days	5 weeks	9 weeks
# Snapshot Granularity	daily	weekly	weekly

model may get stuck in suboptimal solutions based on recent but noisy information. The hyperparameter ω controls the size of the sliding window, which determines the number of previous time steps considered for fine-tuning. As ω becomes smaller, the model updates its evolved representations more frequently, allowing it to capture recent changes. However, this can increase the risk of getting trapped in local optima due to the limited historical information considered. On the other hand, if ω is larger, the model can incorporate longer-term information, but may have reduced sensitivity to recent fine-tuned weights.

4 EVALUATION

In this section, we compare our proposed GraphPro with state-of-the-art methods to address the following research questions.

- **RQ1:** Can GraphPro outperform other time-aware graph learning models and pre-trained GNNs in dynamic recommendations?
- **RQ2:** How does GraphPro perform when integrated as a model-agnostic plug-in component with state-of-the-art recommenders?
- **RQ3:** Can GraphPro perform equally well or even better than the vanilla full-data training paradigm?
- **RQ4:** How does the performance of GraphPro change under different ablation settings of key components and hyper-parameters?
- **RQ5:** How effective is GraphPro in tackling the cold-start issue?
- **RQ6:** How does the potential scalability of GraphPro facilitate efficient model convergence with our prompt learning paradigm?
- **RQ7:** Can GraphPro effectively empower real-world recommendation systems when deployed in industrial applications?

4.1 Experimental Settings

4.1.1 Datasets. We utilize three public datasets that cover diverse real-world scenarios in dynamic recommendation. The **Taobao** dataset captures implicit feedback from Taobao.com, a Chinese e-commerce platform, over a period of 10 days. The **Koubei** dataset, provided for the IJCAI’16 contest, records 9 weeks of user interactions with nearby stores on Koubei in Alipay. The **Amazon** dataset consists of a 13-week collection of product reviews sourced from Amazon. More details about these datasets can be found in Table 1.

4.1.2 Baseline Models. We include the recent dynamic GNNs and graph prompt approaches as our baselines. Specifically, three most relevant research lines are included for comparison:

Dynamic Recommendation Methods. We include DGCN [25], which explores categorizes edges as past and current and designing a GNN to propagate the information from past edges to current ones. However, it does not explicitly consider a dynamic setting

with snapshots and focuses on a single graph.

Graph Prompt Methods. unify pre-training and downstream tasks with a shared prompts for task-specific knowledge retrieval.

- GraphPrompt [28] introduces an approach for pretraining and prompting in graph contexts. It uses a learnable prompt to guide downstream tasks and enables them to access relevant knowledge from pretrained models using a shared template.
- GPF [11] establishes a general approach for tuning prompts in any pre-trained graph neural networks by introducing prompts within the feature space of the graph.

Dynamic Graph Neural Networks. These networks are tailored to dynamic graphs, updating embeddings with time sensitivity to reflect graph changes. We benchmark our approach against notable models EvolveGCN [32] and ROLAND [56]:

- EvolveGCN [32]. It adapts to graph evolution using recurrent neural networks to modify GCN parameters over time, available as hidden state (-H) or recurrent input (-O) variants.
- ROLAND [56]. A dynamic graph method that employs meta-learning to refresh embeddings for subsequent re-initialization, integrating these with GNN’s layer-wise hidden states.

4.1.3 Integration with GNN Recommenders. GraphPro is a versatile architecture that seamlessly integrates as a plug-in component with any GNN-based recommender, highlighting its flexibility. In our evaluation, we implement GraphPro using the efficient LightGCN [15] model. Additionally, we extend GraphPro’s applicability by integrating it with self-supervised learning-based recommenders like SGL [49], MixGCF [21], and SimGCL [58]. This integration provides empirical evidence of GraphPro’s effectiveness in enhancing dynamic and adaptable recommendations.

4.1.4 Evaluation Protocols. In our evaluation, we used graph snapshots at different intervals to model real-world dynamics. We employed a two-step sliding window technique for learning and predicting future changes. Following the Pre-train and Fine-tune framework, we pre-trained the model with a significant dataset fraction and fine-tuned it on subsequent snapshots (Table 1). Baseline models followed the same approach for consistency. Dynamic GNNs started fine-tuning with pre-trained weights. Results were averaged over future snapshots, using standard metrics like Recall@k and nDCG@k at k=20, aligning with existing methodologies [15, 16, 49].

4.2 Performance Comparison (RQ1–RQ3)

4.2.1 Comparison with Baselines. We showcase the performance of GraphPro alongside other methods, such as graph prompt techniques and dynamic GNNs, using LightGCN as the baseline, detailed in Table 2. Key observations from the analysis include:

- Our GraphPro surpasses both graph prompt and dynamic graph learning methods, highlighting the efficacy of our pre-training and prompt learning strategy. This performance can be credited to: 1) Our temporal prompt mechanism that adeptly captures the evolving user-item interactions during pre-training and fine-tuning, and 2) Our graph prompt design that ensures seamless knowledge transfer from the pre-trained model and mitigates distribution shifts across temporal snapshots.
- The varying performance of baseline methods underscores the complexity of dynamic recommendations. EvolveGCN excels

Table 2: When compared to various baselines utilizing different backbone architectures, GraphPro consistently exhibits strong overall performance across different types of datasets. The script * denotes the statistically significant results compared to the second best at $p < 0.01$ level.

Method	Taobao		Koubei		Amazon	
	Recall	nDCG	Recall	nDCG	Recall	nDCG
DGCN	0.0229	0.0228	<u>0.0353</u>	<u>0.0255</u>	0.0158	0.0084
LightGCN+						
GraphPrompt	0.0199	0.0195	0.0342	0.0249	0.0154	0.0075
GPF	0.0223	0.0220	0.0348	0.0251	<u>0.0174</u>	<u>0.0088</u>
EvolveGCN-H	0.0224	0.0221	0.0315	0.0231	0.0138	0.0066
EvolveGCN-O	<u>0.0236</u>	<u>0.0232</u>	0.0334	0.0242	0.0157	0.0084
ROLAND	0.0226	0.0226	0.0301	0.0223	0.0150	0.0069
GraphPro	0.0251*	0.0245*	0.0362*	0.0265*	0.0191*	0.0094*
SGL+						
GraphPrompt	0.0223	0.0220	0.0355	0.0261	0.0161	0.0079
GPF	0.0229	0.0226	0.0363	0.0266	<u>0.0187</u>	<u>0.0096</u>
EvolveGCN-H	0.0235	0.0232	0.0358	0.0263	0.0137	0.0066
EvolveGCN-O	<u>0.0242</u>	<u>0.0238</u>	<u>0.0365</u>	<u>0.0268</u>	0.0173	0.0090
ROLAND	0.0222	0.0222	0.0340	0.0251	0.0161	0.0078
GraphPro	0.0268*	0.0264*	0.0371*	0.0277*	0.0221*	0.0114*
MixGCF+						
GraphPrompt	0.0248	0.0245	0.0377	0.0276	0.0180	0.0089
GPF	0.0251	0.0247	<u>0.0380</u>	<u>0.0278</u>	<u>0.0182</u>	<u>0.0092</u>
EvolveGCN-H	0.0240	0.0237	0.0354	0.0262	0.0129	0.0061
EvolveGCN-O	<u>0.0271</u>	<u>0.0267</u>	0.0375	0.0276	0.0171	0.0085
ROLAND	0.0232	0.0230	0.0349	0.0260	0.0152	0.0072
GraphPro	0.0280*	0.0273*	0.0393*	0.0291*	0.0216*	0.0109*
SimGCL+						
GraphPrompt	0.0239	0.0224	0.0348	0.0258	0.0139	0.0069
GPF	0.0237	0.0220	0.0357	0.0264	<u>0.0182</u>	<u>0.0094</u>
EvolveGCN-H	0.0241	0.0238	<u>0.0356</u>	<u>0.0265</u>	0.0134	0.0067
EvolveGCN-O	<u>0.0241</u>	<u>0.0238</u>	0.0351	0.0258	0.0168	0.0088
ROLAND	0.0228	0.0228	0.0333	0.0246	0.0151	0.0075
GraphPro	0.0280*	0.0276*	0.0368*	0.0276*	0.0205*	0.0108*

with the Taobao dataset but not universally, possibly due to overfitting short-term user behaviors with its complex architecture. In contrast, GraphPro employs streamlined prompt mechanisms that adeptly seize long-term user interests and assimilate fresh behavioral data, offering a more consistent and effective approach.

- Despite being regarded as the meticulously-designed dynamic GNN, ROLAND does not demonstrate superior performance. This limitation may be attributed to its intricate model parameter update schemes, which introduce larger perturbations to embeddings. Consequently, the representation learning for users and items is disrupted, rendering it less effective in capturing the time-evolving user preferences in recommendation tasks.

4.2.2 Integration with SOTA Methods. We also evaluate GraphPro’s versatility with various base recommenders—MixGCF, SGL, and SimGCL—reimplementing all methods under uniform evaluation criteria. The averaged results across time slots can be found in Table 2. Our summarized observations are as follows:

- GraphPro consistently excels alongside advanced recommenders, showcasing our method’s flexibility and capability to elevate performance across scenarios. Baselines, however, show varied performance depending on the base recommender and dataset used, with no single method dominating across all scenarios. While




Figure 3: Key component ablation study for fine-tuning stage. Y-axis denotes performance metrics on the left and epochs (displayed as ∇) for convergence on the right.

EvolveGCN-O often ranks second on Taobao, GPF stands out on Amazon regardless of the base recommender. Such variability in response to dataset specifics and underlying models prevents baselines from achieving consistent, substantial outcomes.

- Superior base model representations typically yield better results. GraphPro sees a 4.5% boost when comparing SGL to SimGCL on Taobao, suggesting it capitalizes well on the base model’s improved representations in dynamic contexts. Conversely, methods like EvolveGCN-O fail to show gains with better base models, hinting at potential generalization weaknesses in their design.

4.2.3 Comparison with Full-Data Training. We place the results and discussion in Appendix A.2 due to space limitation.

4.3 Ablation Study (RQ4)

4.3.1 Key Components in Fine-Tuning. We conduct an ablation study to evaluate the impact of key elements in GraphPro’s fine-tuning process. Three variants of GraphPro are introduced, each omitting a specific component, for comparative analysis.

- (-)PT: This variant removes the **Prompt Tuning** component that leverages prompt edges from past interactions, opting to fine-tune the model directly with new interactions.
- (-)GT: Adaptive **GaTing** mechanism, which facilitates dynamic knowledge transformation in the fine-tuning phase, is omitted.
- (-)IU: The **Interplotive Update** module is excluded, maintaining static pre-trained weights during fine-tuning.

Based on the data shown in Figure 3, we found the following key insights: i) Each of the three components in GraphPro is crucial. Removing any of them not only decreases recommendation precision but also prolongs the convergence time, highlighting their collective effectiveness. ii) The structural prompt and interpolative update features are essential for improving accuracy and preventing local optima during dynamic learning. Their absence leads to faster convergence but significantly reduced accuracy. iii) The adaptive gating feature is vital for faster convergence and improved accuracy by facilitating better gradient flow. Without it, we observe longer convergence times and decreased accuracy, emphasizing its role in handling distribution shifts during fine-tuning across snapshots.

4.3.2 Effect of Pre-trained Model. We assess how different pre-training model architectures influence fine-tuning outcomes by examining four variants, focusing on the use of relative time encoding and the intrinsic representational strength of the models. The baseline model "LGN(+)"TE" incorporates time encoding (TE) with LightGCN. The "LGN(-)"TE" variant excludes TE during pretraining. The other two, "MixGCF(+)"TE" and "SimGCL(+)"TE," employ more robust models for pretraining while maintaining LightGCN




Figure 4: Ablation study for pretrained models.




Figure 5: Evaluation performance for tuned and untuned users on Amazon compared with the best baseline, GPF.

for fine-tuning. Performance contrasts between pretraining and fine-tuning for these variants are depicted in Figure 4.

- The temporal prompt mechanism enhances both the speed of convergence and the accuracy of predictions during pretraining and fine-tuning. It directs the model to capitalize on crucial temporal cues in the message-passing process.
- The use of more robust pretrained models correlates with improved performance, echoing trends seen in pre-trained language models like BERT [9] and vision models like ViT [10]. Our framework’s scalability and adaptability are evident, as it allows the use of powerful pretrained models to excel in recommendation. This underscores the strategy of refining large pretrained models with more streamlined models downstream for enhanced results.

4.4 Learning Impact Analysis (RQ5 & RQ6)

4.4.1 Fine-tuned v.s. Untuned (Cold-start) Nodes. This section highlights the advantages of our fine-tuning approach on node representation learning for recommendations. We categorize users into two groups: those undergoing fine-tuning and cold-start users. Using the Amazon dataset, we evaluate and present the findings for each group independently in Figure 5.

- GraphPro effectively enhances representations for both tuned and cold-start users, consistently outperforming in most scenarios. This success is largely due to the structural prompt design, which leverages past interactions to enrich the representations.
- GraphPro showcases enhanced performance over the baseline with sustained long-term improvement. After an initial lag in the first snapshot, GraphPro overtakes the top baseline from T_2 through T_8 , underscoring its effectiveness at navigating local optima for superior long-term results in dynamic settings.

4.4.2 Efficiency in Learning. This section examines the learning efficiency of our GraphPro. As a parameter-efficient approach, GraphPro minimizes learnable weights during both pretraining and fine-tuning, unlike incremental training methods, thereby reducing training time and computational costs. Our comparisons of training curves, shown in Figure 6, reveal that GraphPro not only




Figure 6: The training curves for GraphPro and baselines on the Taobao and Koubei are presented, with scatters marking performance over various stages (pre-training and fine-tuning) against epochs. A star marker denotes the final convergence, and the right y-axis shows the total time consumed.

Table 3: Online A/B test results spanning 5 days. HPC: highly-personalized content. CC: click count. VCC: video click count.

Model	CTR	HPC CTR	Avg. CC	Avg. VCC
Online Model	10.61%	13.42%	0.6716	0.0188
GraphPro	10.78%	13.89%	0.6831	0.0194
# Improve	1.53%± 0.68%	3.45%± 0.64%	1.71%± 0.84%	3.28%± 1.76%

outperforms EvolveGCN-O and GPF on the Taobao and Koubei datasets but also achieves this with fewer epochs. For example, GraphPro reaches convergence after just four fine-tuning stages, taking roughly half the epochs and time compared to the baselines, underscoring the efficiency improvements GraphPro provides.

4.5 Online Deployment (RQ7)

GraphPro is implemented on Wechat Open Platforms, serving millions of users, to assess its effectiveness in content recommendations. Integrating GraphPro into the core CTR prediction system, it utilizes unsupervised DGI-trained user embeddings as pretraining weights. Historical item-to-user interactions shape prompt edges, and these embeddings are fine-tuned with a 1-layer GNN to inform item representations for the primary model, with updates occurring every ten minutes. During an online A/B test, GraphPro and the current online model each engage approximately 2 million users. Performance is evaluated over a 5-day period using four metrics related to CTR and click count, as presented in Table 3. Results show that GraphPro significantly enhances the real-world recommender system by effectively modeling evolving user and item representations and leveraging deep user interests through pretraining and fine-tuning, all while being easy and cost-effective to deploy.

5 RELATED WORKS

GNNs for Recommendation. GNNs have become key in recommendation systems for capturing multi-hop collaborative signals [12]. Models like NGCF [47] and PinSage [55] recursively enhance user and item embeddings through message-passing. GCCF [6] adds a residual structure, whereas LightGCN [15] streamlines by discarding non-linearities. Further innovations include hypergraph learning [26, 53] and intent disentanglement [36, 48] to parse complex collaborative patterns. To tackle sparsity in graph recommenders, self-supervised techniques such as contrastive [24, 40] and generative approaches [50] have been introduced.

Pre-training and Fine-tuning on GNNs. Inspired by NLP’s pre-training and fine-tuning success, similar techniques have been

applied to GNNs. Contrastive learning and infomax-based pre-training have been developed for advanced representation learning [19, 43, 46, 57]. Innovations in pre-training, including link prediction, feature generation, and prompt-based fine-tuning, have also emerged [11, 17, 20, 28, 44, 45]. However, effectively handling dynamic graph learning remains a challenge. Our GraphPro excels in learning with dynamic graphs during pre-training and fine-tuning.

In the realm of recommendation tasks, [14, 27] suggest pre-training models tailored for user and item modeling. [14] pre-trains a GNN to mimic cold-start conditions, while [27] develops a pre-training strategy using side information. Nevertheless, these approaches concentrate on static recommendation settings and fail to account for the dynamic evolution of user preferences, thus limiting their effectiveness as time-sensitive recommenders.

Dynamic Graph Learning. Research on learning with temporally evolving graphs is gaining traction. Techniques such as EvolveGCN [32], Dyngraph2vec [13], DGNN [29], ROLAND [56], and WinGNN [60] leverage recurrent layers to grasp graph changes. Yet, they miss a dedicated pre-training and fine-tuning strategy specific to dynamic graphs and can introduce noise into user and item representations. DGCN [25] does consider dynamics in graph recommendation learning, but it falls short of handling graphs with evolving snapshots, limiting its assessment to a static graph. [8, 34] introduce a snapshot-based dynamic setting for recommendation, but are designated to predict next items as sequential recommenders.

Sequential Recommendation. Sequential recommendation, a field attuned to time-sensitive recommendation scenarios, showcases works like attention-based approaches: SASRec [23] and BERT4Rec [42], GNN-based methods: SURGE [4] and Retaggn [18], and SSL models such as S3-rec [59], ICL [7], and DCRec [52]. Our approach, while similarly embracing temporal dynamics, diverges from these next-item predictors. Sequential models, with their autoregressive encoders, face limitations in next-item prediction and differ from graph-based methods that excel in top-K item retrieval. Furthermore, they don’t incorporate explicit time intervals like days or weeks, but rather rely on fixed historical sequences.

6 CONCLUSION

This study introduces GraphPro, a new framework combining dynamic graph pre-training with prompt learning, aiming to substantially enhance the adaptability and scalability of recommender systems that are sensitive to temporal dynamics. GraphPro uses a temporal prompt mechanism to transfer knowledge from pre-trained models to recommendation tasks on fresh data. Its graph-structured prompt learning features an adaptive gating mechanism that includes vital contextual information, easing fine-tuning and responsiveness to behavioral shifts. Through an extensive array of experiments plus real-world deployment, we validate that GraphPro delivers exceptionally better performances for recommendation.

ACKNOWLEDGMENTS

This project is partially supported by 2023 Tencent Wechat Rhino-Bird Focused Research Program Research and Weixin Open Platform. This research work is also supported by Department of Computer Science & Musketeers Foundation Institute of Data Science at the University of Hong Kong.

REFERENCES

- [1] Rianne van den Berg, Thomas N Kipf, and Max Welling. 2017. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263* (2017).
- [2] Xuheng Cai, Chao Huang, Lianghao Xia, and Xubin Ren. 2023. LightGCL: Simple Yet Effective Graph Contrastive Learning for Recommendation. *arXiv preprint arXiv:2302.08191* (2023).
- [3] Yi Cai, Xuefei Ning, Huazhong Yang, and Yu Wang. 2023. Ensemble-in-One: Ensemble Learning within Random Gated Networks for Enhanced Adversarial Robustness. In *AAAI*, Vol. 37. 14738–14747.
- [4] Jianxin Chang, Chen Gao, Yu Zheng, Yiqun Hui, Yanan Niu, Yang Song, Depeng Jin, and Yong Li. 2021. Sequential recommendation with graph neural networks. In *SIGIR*. 378–387.
- [5] Lei Chen, Le Wu, Richang Hong, Kun Zhang, and Meng Wang. 2020. Revisiting graph based collaborative filtering: A linear residual graph convolutional network approach. In *AAAI*, Vol. 34. 27–34.
- [6] Lei Chen, Le Wu, Richang Hong, Kun Zhang, and Meng Wang. 2020. Revisiting graph based collaborative filtering: A linear residual graph convolutional network approach. In *AAAI*, Vol. 34. 27–34.
- [7] Yongjun Chen, Zhiwei Liu, Jia Li, Julian McAuley, and Caiming Xiong. 2022. Intent contrastive learning for sequential recommendation. In *WWW*. 2172–2182.
- [8] Zeyuan Chen, Wei Zhang, Junchi Yan, Gang Wang, and Jianyong Wang. 2021. Learning dual dynamic representations on time-sliced user-item interaction graphs for sequential recommendation. In *CIKM*. 231–240.
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xi-aohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [11] Taoran Fang, Yunchao Zhang, Yang Yang, Chunping Wang, and Lei Chen. 2022. Universal Prompt Tuning for Graph Neural Networks. *arXiv preprint arXiv:2209.15240* (2022).
- [12] Chen Gao, Yu Zheng, Nian Li, Yinfeng Li, Yingrong Qin, Jinghua Piao, Yuhan Quan, Jianxin Chang, Depeng Jin, Xiangnan He, et al. 2023. A survey of graph neural networks for recommender systems: Challenges, methods, and directions. *Transactions on Recommender Systems* 1, 1 (2023), 1–51.
- [13] Palash Goyal, Sujit Rokka Chhetri, and Arquimedes Canedo. 2020. dyngraph2vec: Capturing network dynamics using dynamic graph representation learning. *Knowledge-Based Systems* 187 (2020), 104816.
- [14] Bowen Hao, Jing Zhang, Hongzhi Yin, et al. 2021. Pre-training graph neural networks for cold-start users and items representation. In *WSDM*. 265–273.
- [15] Xiangnan He, Kuan Deng, Xiang Wang, et al. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *SIGIR*. 639–648.
- [16] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *WWW*. 173–182.
- [17] Zhenyu Hou, Xiao Liu, Yukuo Cen, Yuxiao Dong, Hongxia Yang, et al. 2022. Graphmae: Self-supervised masked graph autoencoders. In *KDD*. 594–604.
- [18] Cheng Hsu and Cheng-Te Li. 2021. Retaggn: Relational temporal attentive graph neural networks for holistic sequential recommendation. In *WWW*. 2968–2979.
- [19] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. 2019. Strategies for pre-training graph neural networks. *arXiv preprint arXiv:1905.12265* (2019).
- [20] Ziniu Hu, Yuxiao Dong, Kuansan Wang, Kai-Wei Chang, and Yizhou Sun. 2020. Gpt-gnn: Generative pre-training of graph neural networks. In *KDD*. 1857–1867.
- [21] Tinglin Huang, Yuxiao Dong, Ming Ding, Zhen Yang, Wenzheng Feng, Xinyu Wang, and Jie Tang. 2021. Mixgcn: An improved training method for graph neural network-based recommender systems. In *KDD*. 665–674.
- [22] Yitong Ji, Aixin Sun, Jie Zhang, and Chenliang Li. 2023. A critical study on data leakage in recommender system offline evaluation. *Transactions on Information Systems* 41, 3 (2023), 1–27.
- [23] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *ICDM*. IEEE, 197–206.
- [24] Kang Li, Chang-Dong Wang, et al. 2023. Self-Supervised Group Graph Collaborative Filtering for Group Recommendation. In *WSDM*. 69–77.
- [25] Xiaohan Li, Mengqi Zhang, Shu Wu, Zheng Liu, Liang Wang, and S Yu Philip. 2020. Dynamic graph collaborative filtering. In *ICDM*. IEEE, 322–331.
- [26] Yinfeng Li, Chen Gao, Hengliang Luo, Depeng Jin, and Yong Li. 2022. Enhancing hypergraph neural networks with intent disentanglement for session-based recommendation. In *SIGIR*. 1997–2002.
- [27] Siwei Liu, Zaiqiao Meng, Craig Macdonald, and Iadh Ounis. 2023. Graph neural pre-training for recommendation with side information. *TOIS* 41, 3 (2023), 1–28.
- [28] Zemin Liu, Xingtong Yu, Yuan Fang, and Xinming Zhang. 2023. Graphprompt: Unifying pre-training and downstream tasks for graph neural networks. In *WWW*. 417–428.
- [29] Franco Manessi, Alessandro Rozza, and Mario Manzo. 2020. Dynamic graph convolutional networks. *Pattern Recognition* 97 (2020), 107000.
- [30] Kelong Mao, Jieming Zhu, Xi Xiao, Biao Lu, Zhaowei Wang, and Xiuqiang He. 2021. UltraGCN: ultra simplification of graph convolutional networks for recommendation. In *CIKM*. 1253–1262.
- [31] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018).
- [32] Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kanezashi, Tim Kaler, Tao Schardl, et al. 2020. EvolveGCN: Evolving graph convolutional networks for dynamic graphs. In *AAAI*, Vol. 34. 5363–5370.
- [33] Ofir Press, Noah Smith, and Mike Lewis. 2021. Train Short, Test Long: Attention with Linear Biases Enables Input Length Extrapolation. In *ICLR*.
- [34] Jiarui Qin, Kan Ren, Yuchen Fang, Weinan Zhang, and Yong Yu. 2020. Sequential recommendation with dual side neighbor-based collaborative relation modeling. In *WSDM*. 465–473.
- [35] Xubin Ren, Lianghao Xia, Yuhao Yang, Wei Wei, Tianle Wang, Xuheng Cai, and Chao Huang. 2023. SSLRec: A Self-Supervised Learning Library for Recommendation. *arXiv preprint arXiv:2308.05697* (2023).
- [36] Xubin Ren, Lianghao Xia, Jiashu Zhao, Dawei Yin, and Chao Huang. 2023. Disentangled Contrastive Collaborative Filtering. *arXiv preprint arXiv:2305.02759* (2023).
- [37] Timo Schick and Hinrich Schütze. 2020. Exploiting cloze questions for few shot text classification and natural language inference. *arXiv preprint arXiv:2001.07676* (2020).
- [38] Yifei Shen, Yongji Wu, Yao Zhang, Caihua Shan, Jun Zhang, et al. 2021. How powerful is graph convolution for recommendation?. In *CIKM*. 1619–1629.
- [39] Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980* (2020).
- [40] Jie Shuai, Kun Zhang, Le Wu, Peijie Sun, et al. 2022. A review-aware graph contrastive learning framework for recommendation. In *SIGIR*. 1283–1293.
- [41] Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. 2021. RoFormer: Enhanced Transformer with Rotary Position Embedding. *arXiv:2104.09864* [cs.CL]
- [42] Fei Sun, Jun Liu, Jian Wu, et al. 2019. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *CIKM*. 1441–1450.
- [43] Fan-Yun Sun, Jordan Hoffman, Vikas Verma, and Jian Tang. 2019. InfoGraph: Unsupervised and Semi-supervised Graph-Level Representation Learning via Mutual Information Maximization. In *ICLR*.
- [44] Mingchen Sun, Kaixiong Zhou, Xin He, et al. 2022. Gppt: Graph pre-training and prompt tuning to generalize graph neural networks. In *KDD*. 1717–1727.
- [45] Xiangguo Sun, Hong Cheng, Jia Li, Bo Liu, and Jihong Guan. 2023. All in One: Multi-Task Prompting for Graph Neural Networks. (2023).
- [46] Petar Veličković, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, et al. 2018. Deep graph infomax. *arXiv preprint arXiv:1809.10341* (2018).
- [47] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. In *SIGIR*. 165–174.
- [48] Jiahao Wu, Wenqi Fan, Jingfan Chen, Shengcai Liu, et al. 2022. Disentangled contrastive learning for social recommendation. In *CIKM*. 4570–4574.
- [49] Jiancan Wu, Xiang Wang, Fuli Feng, Xiangnan He, Liang Chen, Jianxun Lian, et al. 2021. Self-supervised graph learning for recommendation. In *SIGIR*. 726–735.
- [50] Lianghao Xia, Chao Huang, Chunzhen Huang, Kangyi Lin, Tao Yu, and Ben Kao. 2023. Automated Self-Supervised Learning for Recommendation. In *WWW*. 992–1002.
- [51] Teng Xiao, Zhengyu Chen, and Suhang Wang. 2023. Reconsidering Learning Objectives in Unbiased Recommendation: A Distribution Shift Perspective. In *KDD*. 2764–2775.
- [52] Yuhao Yang, Chao Huang, Lianghao Xia, Chunzhen Huang, Da Luo, and Kangyi Lin. 2023. Debaised Contrastive Learning for Sequential Recommendation. In *WWW*. 1063–1073.
- [53] Yuhao Yang, Chao Huang, Lianghao Xia, Yuxuan Liang, Yanwei Yu, and Chenliang Li. 2022. Multi-behavior hypergraph-enhanced transformer for sequential recommendation. In *KDD*. 2263–2274.
- [54] Zhengyi Yang, Xiangnan He, Jizhi Zhang, et al. 2023. A Generic Learning Framework for Sequential Recommendation with Distribution Shifts. In *SIGIR*.
- [55] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In *KDD*. 974–983.
- [56] Jiaxuan You, Tianyu Du, and Jure Leskovec. 2022. ROLAND: graph learning framework for dynamic graphs. In *KDD*. 2358–2366.
- [57] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, et al. 2020. Graph contrastive learning with augmentations. *NeurIPS* 33 (2020), 5812–5823.
- [58] Junliang Yu, Hongzhi Yin, Xin Xia, Tong Chen, Lizhen Cui, and Quoc Viet Hung Nguyen. 2022. Are graph augmentations necessary? simple graph contrastive learning for recommendation. In *SIGIR*. 1294–1303.
- [59] Kun Zhou, Hui Wang, Wayne Xin Zhao, Yutao Zhu, Sirui Wang, et al. 2020. S3-rec: Self-supervised learning for sequential recommendation with mutual information maximization. In *CIKM*. 1893–1902.
- [60] Yifan Zhu, Fangpeng Cong, Dan Zhang, Wenwen Gong, Qika Lin, Wenzheng Feng, Yuxiao Dong, and Jie Tang. 2023. WinGNN: Dynamic Graph Neural Networks with Random Gradient Aggregation Window. In *KDD*. 3650–3662.

A APPENDIX

A.1 Hyper-parameter Sensitivity

Our research focuses on investigating the response of GraphPro to hyper-parameter adjustments, using three distinct datasets. We conduct a meticulous analysis of the primary hyperparameters that play a crucial role in the operation of GraphPro. These hyperparameters include the temporal prompt interval τ , the updating window ω , which determines the intervals for model updates, and the sampling decay ϕ , which influences the emphasis on past interactions. To ensure a comprehensive evaluation, we carefully select ranges of values for each hyperparameter that are tailored to the unique characteristics of each experimental dataset utilized.

- τ : Taobao-[0.5, 1, 4, 12]; Koubei and Amazon-[24, 48, 72, 96]
- ω : Taobao and Koubei-[1, 2, 3]; Amazon-[2, 4, 6]
- ϕ : [0.05, 0.1, -0.05, -0.1]

Our analysis of hyper-parameter sensitivity in GraphPro is presented in Figure 8, from which we derive several observations:

- 1) GraphPro exhibits a greater degree of sensitivity to hyperparameter changes when applied to the Amazon data in comparison to those of Taobao and Koubei. This heightened sensitivity on Amazon is likely due to its considerably longer fine-tuning and prediction period, which spans an entire 9 weeks. This duration is substantial when contrasted with the other datasets and suggests a correlation between the span of fine-tuning and the impact of hyperparameters on the model’s long-term performance. For enhanced long-term accuracy, we need to be careful when picking these settings.
- 2) When examining the details, it becomes evident that GraphPro is notably more affected by the updating window ω compared to other hyperparameters in its configuration. The size of the updating window requires careful consideration, as it is crucial to strike a delicate balance. If the window is too narrow, the model may fail to incorporate important updates that contribute to the evolving representations of users and items. On the other hand, if the window is too wide, the model may become burdened with irrelevant noise, thus obscuring the current representation learning process.
- 3) Moreover, our findings reveal a reassuring aspect of GraphPro’s design: its resilience. Minor modifications in the hyper-parameter settings do not precipitate substantial declines in the model’s performance. We emphasize the significance of customizing hyperparameter choices to align with the intrinsic characteristics and unique aspects of the dataset being utilized.




Figure 8: Performance change w.r.t. key hyperparameters.

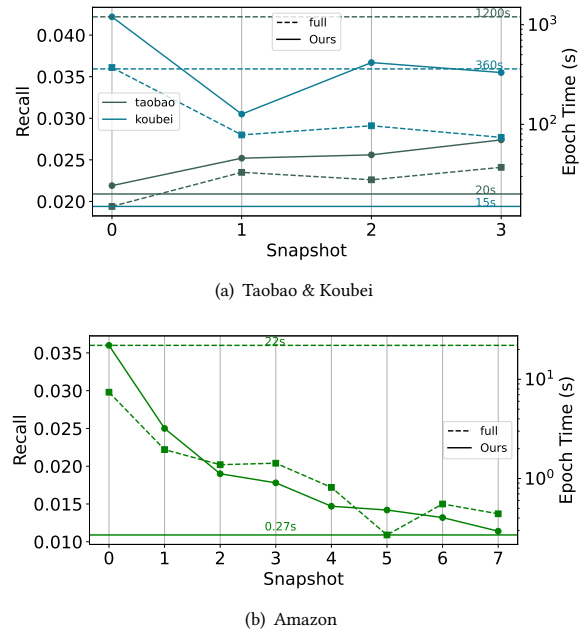


Figure 7: GraphPro v.s. full-data training using LightGCN. The results with the average training time for a single epoch are being depicted by a horizontal line on the right Y-axis.

A.2 Comparison with Full-Data Training

In this section, our primary objective is to conduct a comprehensive performance comparison between our GraphPro and the traditional full-data training approach (FULL). Our aim is to determine if GraphPro can match or exceed the performance of FULL. To achieve this, we conduct a meticulous comparison using three distinct datasets, with a keen focus on Recall and average epoch time. The results of our analysis are visually represented in Figure 7. Our findings can be summarized into two key observations:

- 1) In our study, GraphPro stands out on both the Taobao and Koubei datasets, surpassing FULL’s performance at every testing point. This underscores GraphPro’s robust capability in processing time-sensitive data to yield more precise recommendations. With the Amazon dataset, the two approaches are closely matched, but GraphPro shows a distinct advantage in the initial phases. This suggests that GraphPro’s dynamic method can be immediately beneficial, improving performance from the start without losing ground to FULL as the evaluation progresses over time.
- 2) Furthermore, a significant aspect of GraphPro’s appeal lies in its exceptional efficiency. When compared to the FULL approach, GraphPro cuts down the average training time dramatically, achieving astonishing efficiency gains. Specifically, we observe an impressive 60-fold increase in speed on the Taobao dataset, a 24-fold increase on the Koubei dataset, and an unparalleled 81-fold increase on the Amazon dataset. These improvements in training efficiency can be attributed to the innovative architecture of GraphPro, which seamlessly integrates pre-training and prompt learning techniques. This strategic integration significantly accelerates the learning process, enabling GraphPro to rapidly assimilate and adeptly apply new data to generate accurate and timely recommendations.