A Dynamic Meta-Learning Model for Time-Sensitive Cold-Start Recommendations

Krishna Prasad Neupane, Ervine Zheng, Yu Kong, Qi Yu

Rochester Institute of Technology
{kpn3569, mxz5733, yu.kong, qi.yu}@rit.edu

Abstract

We present a novel dynamic recommendation model that focuses on users who have interactions in the past but turn relatively inactive recently. Making effective recommendations to these *time-sensitive cold-start users* is critical to maintain the user base of a recommender system. Due to the sparse recent interactions, it is challenging to capture these users' current preferences precisely. Solely relying on their historical interactions may also lead to outdated recommendations misaligned with their recent interests. The proposed model leverages historical and current user-item interactions and dynamically factorizes a user's (latent) preference into time-specific and time-evolving representations that jointly affect user behaviors. These latent factors further interact with an optimized item embedding to achieve accurate and timely recommendations. Experiments over real-world data help demonstrate the effectiveness of the proposed time-sensitive cold-start recommendation model.

Introduction

Recommender system has long been used as an effective means to improve user experience and to provide personalized recommendations in diverse fields such as media, entertainment, and e-commerce. (Sun, Parthasarathy, and Varshney 2014; Xie et al. 2018). One effective way of recommendation is via Collaborative Filtering (CF) (Goldberg et al. 1992; Su and Khoshgoftaar 2009), which recommends items based on similar users' preferences. CF assumes that users who had similar interactions with some items in the past are likely to have the same preference on other items, and leverages the observed user-item interactions to make predictions for the missing parts, which indicate the potential items of interest to users. Matrix Factorization (MF) is one commonly used CF technique that exploits user and item latent factors to capture their inherent attributes. Most existing MF methods model user preferences and item attributes as static factors (Koren 2008; Koren, Bell, and Volinsky 2009). Some recent efforts consider the dynamic changes in the user-item interactions by modeling user preferences as shifting latent factors over time, allowing them to provide more timely recommendations (Charlin et al. 2015; Gultekin and Paisley 2014; Sun, Parthasarathy, and Varshney 2014).
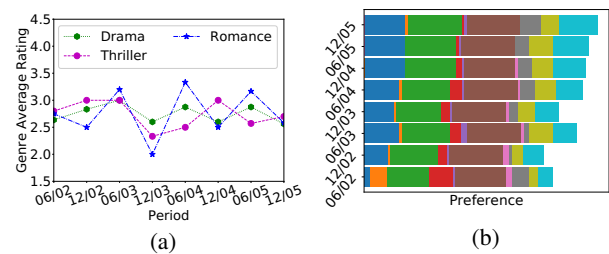
Figure 1: For User (ID:2181970): (a) shows dynamic changes of ratings for three genre types over time (b) demonstrates how user's latent preferences evolve over time.

When user-item interactions are very sparse, making accurate recommendations based on limited information is highly challenging, which is usually referred to as the cold-start problem. A viable solution to handle the cold-start problem is to utilize extra side information such as item descriptions, user profiles (Uyangoda, Ahangama, and Ranasinghe 2018), and social relationships (Yao et al. 2014). While various side information is widely available for items (i.e., movies, songs, and books), the user-related side information is typically very scarce as acquiring users' personalized information may be practically difficult due to privacy issues. Several recent works adopted meta-learning for few-shot learning in the recommender systems to alleviate the cold-start problem (Vartak et al. 2017; Lee et al. 2019; Du et al. 2019). Meta-learning aims to learn global knowledge from the historical information of many similar tasks (users) and then provide quick adaptation for a new task (user) with limited interaction information. Although the meta-learning approaches show promising results (Vartak et al. 2017; Lee et al. 2019), they are primarily designed for static settings and hence not effective in providing *timely recommendations that best reflect users' current interests*.

Figure 1a shows that the average movie ratings for three genre types of an example user from the Netflix dataset, which oscillate significantly from 2002 to 2005. The average rating indicates an overall satisfaction on each type of genre items that the user interacts with for each period, and serves as an indicator for the change of users' preference. Those changes are usually caused by the interplay of two contributing factors. First, a user's preference over different types of items (e.g., movie or music) may change over

time, which we refer to as the *time-evolving factor*. Second, a user's behavior in a specific period may vary significantly from other periods due to the impact of money/time budget or other external causes, which we refer to as the *time-specific factor*. Given sufficient user interactions over time, both factors could be effectively learned to provide accurate and timely recommendations. However, in practice, many users may have interactions in prior periods but become largely inactive with few interactions in recent periods for various reasons. We define these users as the *time-sensitive cold-start users* due to scarce recent interactions. Solely relying on the historical interactions of these users may lead to outdated recommendations that do not match their recent interests. Furthermore, the limited recent interactions pose a user cold-start problem for the current period that makes existing CF-based techniques less effective.

The main challenge with these time-sensitive cold-start users is to simultaneously capture their most recent interests and evolving preferences, which are keys to achieve accurate and timely recommendations. In this paper, we focus on this special *time-sensitive cold-start problem*, which is critical for a recommender system to maintain its user base. We propose to dynamically factorize a user's (latent) preference into time-specific and time-evolving representations in order to capture the time-specific and time-evolving factors from both the historical and current user-item interactions. For example, the Netflix dataset consists of a user's interactions with a large movie set in the form of user ratings. The variation of movie ratings from the same user may be affected by the change of user preference, rating criteria, or other (unknown) factors. In addition, a user's preference for different genres may evolve over multiple periods, as shown in Figure 1b, which corresponds to the proportion of different factors in the time-evolving representation discovered by our proposed model.

The proposed model consists of two distinct modules: a *meta-learning module* and a *recurrent module*. The former aims to capture time-specific latent factors through limited interaction data by leveraging the shared knowledge learned from other users. The latter aims to capture time-evolving latent factors by nesting a recurrent neural network, and it can be jointly optimized with the meta-learning module through the model-agnostic meta-learning approach (Finn, Abbeel, and Levine 2017). Finally, we seamlessly integrate the two modules by merging the time-specific and time-evolving factors to form the user representation. This user representation further interacts with an item embedding (which is also optimized during model training) to provide the final recommendations. Our experimental results clearly show that the proposed model makes timely recommendations that closely resemble the dynamically changed user ratings as a result of effectively integrating the complementary factors capturing the user preferences.

The **main contributions** of this paper are five-fold: (i) the **first work** to formulate the time-sensitive cold-start problem that is critical to maintain the user base of a recommender system; (ii) a **novel integrated recommendation framework** to model sparse dynamic user-item interactions and extract time-evolving and time-specific factors of user preferences simultaneously; (iii) a **time-sensitive meta-learning module** to effectively handle user cold-start problems by leveraging knowledge shared across multiple users from the current recommendation period to adapt to any specific user's case using limited interaction information, (iv) a **time-evolving recurrent module** to effectively capture the gradual shift of users' preferences over time, and (v) an **integrated training process** that combines these two models to simultaneously learn time-specific and time-evolving factors and optimizes the item embedding.

We conduct extensive experiments over multiple real-world datasets and compare with representative state-of-art dynamic and meta-learning-based recommender systems to demonstrate the effectiveness of the proposed model.

Related Work

Matrix Factorization Models. Matrix factorization is a commonly used collaborative filtering approach that characterizes users and items through latent factors inferred from item rating patterns (Koren, Bell, and Volinsky 2009). Using singular value decomposition (SVD) for recommendation is popularized by *Simon Funk* (Funk 2006) in the Netflix prize competition and a probabilistic version is introduced by (Mnih and Salakhutdinov 2008). SVD is extended to SVD++ (Koren 2008) for processing implicit feedback.

The above static models do not include the important temporal information, which should be considered for analyzing user-item interactions in a time-sensitive setting. Dynamic matrix factorization has been developed to address the issue, which allows latent features to change with time. Some works introduce time-specific factors, such as timeSVD++, which uses additive bias to model user-related temporal changes (Koren 2009). There have been works that employ Gaussian state-space models to introduce time-evolving factors with a one-way Kalman filter (Sun, Parthasarathy, and Varshney 2014; Gultekin and Paisley 2014). To handle implicit data, Sahoo et al. propose an extension to the hidden Markov model (Sahoo, Singh, and Mukhopadhyay 2012), where clicks are drawn from a negative binomial distribution and Charlin et al. (Charlin et al. 2015) introduces a Gaussian state-space model with the Poisson emission. However, matrix-factorization-based algorithms may suffer from limited expressive power and may not capture the complex nature of user-item interactions. Instead, our model captures both dynamic and static user preferences.

Deep Learning Models. Recent works in recommender systems (Cheng et al. 2016; Guo et al. 2017; Zhou et al. 2019) utilize deep learning to provide better recommendations. Cheng et al. (Cheng et al. 2016) propose to jointly train wide linear models and deep neural networks to combine the benefits of memorization and generalization. Similarly, DeepFM (Guo et al. 2017) integrates the power of deep learning and factorization machines models to learn low- and high-order feature interactions simultaneously from the input. DIEN (Zhou et al. 2019) formulates interest evolution network as a deep learning model to capture latent temporal interests and evolving interests for better recommendations. These models are very sensitive to the features and might

need important features information and large datasets while training, which is not necessary for our model.

Graph-Based Models. Another popular line of recommendation systems is graph-based models. A graph captures high-order user-item interactions through an iterative process to provide effective recommendations (Guo et al. 2020). Users and items are represented as a bipartite graph in (Berg, Kipf, and Welling 2017) and links are predicted to make recommendations. Similarly, a graph-based framework called Neural Graph Collaborative Filtering (NGCF) (Wang et al. 2019b) explicitly encodes the collaborative signal in the form of high-order connectivities in a user-item bipartite graph via embedding propagation. However, these methods are unable to capture long-term user preferences or deal with cold-start problems.

Sequential Models. Sequential models understand the sequential user behaviors via user-item interactions, and model the evolution of users' preferences and item popularity over time (Wang et al. 2019a; Fang et al. 2020). Tang et al. (Tang and Wang 2018) utilizes convolutional sequence embedding to capture union level and point level contributions of historical items via horizontal and vertical filters and provides top-N sequential recommendations. Similarly, Kang et al. (Kang and McAuley 2018) introduce a self-attentive mechanism to handle both long and short-term user preferences in a sequential setting. Sequential models focus on users' evolving preferences (i.e., recent interactions) but largely neglect long-term users' preferences. In contrast, our model dynamically captures long-term time-evolving preference via RNN and time-specific users preferences through meta-learning to alleviate cold-start problems.

Meta-learning Models. The user-item interaction data is usually sparse because a user may only interact with a few items within the large item pool. In such cases, making recommendations can be viewed as a few-shot learning problem. Meta-learning (Schmidhuber 1987; Bengio et al. 1992) is recently becoming a popular few-shot learning approach that learns from similar tasks and can generalize quickly and efficiently for the few-shot unseen new tasks. Finn et al. (Finn, Abbeel, and Levine 2017) propose a model-agnostic meta-learning model that learns global parameters from a large number of tasks and performs as a good generalization on a new task that has few samples utilizing the few steps of gradients. To address the cold-start problem in item recommendation, Vartak et al. (Vartak et al. 2017) introduce a meta-learning strategy that focuses on the items-cold-start problem to recommend cold-start items considering that items arrive continuously in the Twitter Timeline. Similarly, recent work based on meta-learning is done to estimate user preferences in (Lee et al. 2019) and scenario-specific recommendation in (Du et al. 2019). Both works use the information of users and items to generate user and item embeddings. Meta-learning is specifically utilized to learn heterogeneous information networks to alleviate cold-start problems (Lu, Fang, and Shi 2020). Similarly, meta-learning approachs are utilized in graphs (Wei et al. 2020; Xie et al. 2021) to quickly adapt to new sub-graph that alleviate the cold-start problem. Also, meta-learning is applied in click-through rate (CTR) prediction (Pan et al. 2019) where desirable embeddings for the new ads are generated via meta-learner. These embedding methods are designed for a static setting, making them not applicable to learn user preferences evolving over time. In our work, we simultaneously learn time-evolving and user-specific preferences to provide accurate and timely recommendations for time-specific cold-start users with very limited recent interaction data.

The Dynamic Meta-Learning Recommendation Model

Problem Settings. We propose a dynamic recommendation model, where the input data is represented as $\{\mathcal{U}, \mathcal{I}, \mathcal{H}\}$, $\mathcal{U}$ is the user set, $\mathcal{I}$ is the item set, and $\mathcal{H}$ is the set of time periods. A time period $t \in \mathcal{H}$ defines a particular time interval where the interactions for user u with the item i are aggregated based on timestamps. We perform recommendation in each time period with a recommendation function as

$$f^t_{\theta^t_u, \omega}(i) = \hat{r}^t_{(u,i)} \qquad \forall u \in \mathcal{U}, i \in \mathcal{I}, t \in \mathcal{H} \qquad (1)$$

where $\hat{r}^t_{(u,i)}$ is the recommended score for item i assigned by user u at period t, θ^t_u is user latent factor at time t, and ω is the parameter of the recurrent neural network module. The goal of a recommender system is to predict the recommendation scores that can accurately capture a user's true preference on items over time so that the recommended items are likely to be adopted by the users.

We formulate dynamic recommendations as a few-shot regression problem in the meta-learning setting. Users are *dynamically* partitioned into *meta-train* and *meta-test* sets based on their interactions in the current recommendation period t. In particular, the meta-train user set includes users with sufficient interactions, while the meta-test user set includes time-sensitive cold-start users who have only a few interactions in the current time period. Details for the train-test user splits are discussed in the experiment section. We consider a distribution over tasks $P(\mathcal{T})$, and each user is represented as a few-shot regression task $\mathcal{T}^t_u$ sampled from the given task distribution. In general, a task includes a *support set* $\mathcal{S}^t_u$ and a *query set* $\mathcal{Q}^t_u$. The support set includes a user's interactions at period t where k is interpreted as the number of shots (i.e., interactions). The query set includes the rest interactions of this user at period t.

$$\mathcal{T}^t_u \sim P(\mathcal{T}): \quad \mathcal{S}^t_u = \{(u, i_j), r^t_{(u,i_j)}\}_{j=1:k}, \\ \mathcal{Q}^t_u = \{(u, i_j), r^t_{(u,i_j)}\}_{j=k+1:N_t} \qquad (2)$$

where N_t is the number of items a user interacted with at period t and $r^t_{(u,i_j)}$ represents label (i.e. rating or count) from user u to item i_j. We adopt episodic training (Vinyals et al. 2016), where the training task mimics the test task for efficient meta-learning. The support set $\mathcal{S}$ in each episode works as the labeled training set on which the model is trained to minimize the loss over the query set $\mathcal{Q}$. The training process iterates episode by episode until convergence.
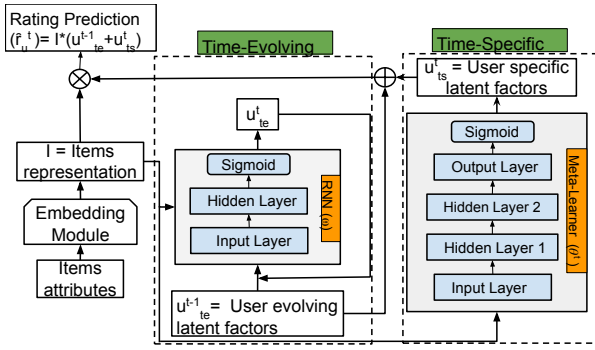
Figure 2: The proposed model captures time-evolving factors via a recurrent network module and time-specific factors through a meta-learning module.

Model Overview

To leverage item information such as text descriptions, our model generates an initial item representation (I) using an embedding matrix $E \in \mathbb{R}^{d \times m}$, where m is the dimension of input item attributes, and d is the dimension of embedding. The embedding is generated from item attributes following (Cheng et al. 2016; Lee et al. 2019). An item i is first encoded as a binary vector $z_i \in \mathcal{R}^m$, where the corresponding index of item attributes is set to 1 and 0 otherwise. The binary vector is then transformed using the embedding matrix: $e_i = E z_i$. The embedding of all items can be stacked as: $I = [e_1, e_2, ..., e_n]$ where n is the total number of items. The embedding matrix E will be optimized along with the model training process after the user latent factor is learned, and details are provided at the end of this section.

Figure 2 shows that the proposed model consists of a time-specific meta-learning module and a time-evolving recurrent neural network module to generate time-specific user latent factors u_{ts}^t and time-evolving latent factors u_{te}^t, both of which contribute to the final prediction $f^t(u, i)$. Details of them are described in following sections. After both modules are trained, the model learns time-specific user factors u_{ts}^t and time-evolving user factors u_{te}^t. These user factors are merged to interact with the item embedding to provide recommendations for the user. The recommendation for a user u at the current period t is denoted as a vector $\hat{r}_u^t$.

Making recommendations can be viewed as a regression problem. By using the mean square error (MSE) function, the loss for a specific user u is formulated as:

$$\mathcal{L}_{\mathcal{T}_u^t}[f_{\theta_u^t, \omega}^t] = \sum_i ||f_{\theta_u^t, \omega}^t(i) - r_{(u,i)}^t||_2^2,$$
$$f_{\theta_u^t, \omega}^t(i) = (u_{ts}^t + u_{te}^{t-1}) \cdot e_i \tag{3}$$

where $r_{(u,i)}^t$ is user-item interaction (rating or count). The user representation is the vector sum of time-specific u_{ts}^t and time-evolving user factors u_{te}^{t-1} (as a compact single representation reducing the number of trainable parameters that helps to avoid overfitting), and the prediction is achieved by the dot product of u and item embedding i. Note that prediction for the current time includes time-specific user factors from current time period i.e. u_{ts}^t and time-evolving user

factors from the previous time period i.e. u_{te}^{t-1}. In general, dynamic recommender systems utilize the information from the previous period to predict for the next period, and we follow this standard setting.

The total loss is formed by aggregating all users in the meta-train set, regularized by the L_2 norm of key model parameters. Let θ_u^t and θ^t denote the local (i.e., user-specific) and global parameters of the time-specific meta-learning module, ω denote the parameters of the time-evolving recurrent neural network module. Training of a dynamic recommendation model can be formulated as:

$$\arg\min_{\theta^t, \omega} \sum_{\mathcal{T}_u^t \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_u^t}[f_{\theta_u^t, \omega}^t] + \frac{\lambda}{2}(||\theta^t||_2^2 + ||\omega||_2^2), \tag{4}$$
$$\theta_u^t = \theta^t - \alpha \nabla_{\theta^t} \mathcal{L}_{\mathcal{T}_u^t}[f_{\theta^t, \omega}^t]$$

where θ_u^t is one gradient step from global parameter θ^t of the meta-learned time-specific module with α being the step size and λ is the regularization parameter.

Time-Specific Meta-Learning Module

This module aims to capture time-specific user factors by only considering the information from the current recommendation period. The meta-learner takes input from the specific period, which is a different setting than the existing meta-learning-based recommender systems (Lee et al. 2019; Du et al. 2019). In this way, the model can capture the latent factors associated with that specific period to provide more accurate and timely recommendations. We consider each user as a learning task. Our goal is to learn a meta parameter θ^t that represents a time-specific global user representation given the meta-training set. We follow the standard setting of few-shot learning (Finn, Abbeel, and Levine 2017), where the distribution over tasks is represented as $p(\mathcal{T})$. The model is trained iteratively by sampling tasks from $p(\mathcal{T})$. The meta-learning module generates time-specific user latent factors (u_{ts}^t) as:

$$u_{ts}^t = f_{meta}^t(\mathcal{T}_u^t; \theta^t) \tag{5}$$

where $\mathcal{T}_u^t$ represents task of a user u at period t. The task $\mathcal{T}_u^t$ includes $\mathcal{S}_u^t$ and $\mathcal{Q}_u^t$. We first pass $\mathcal{S}_u^t$ into f_{meta}^t to adapt user-specific model parameter θ_u^t from the global user model parameter θ^t and then we provide $\mathcal{Q}_u^t$ into the f_{meta}^t to generate time-specific user factors (u_{ts}^t).

We apply an optimization-based meta-learning approach (Finn, Abbeel, and Levine 2017) to learn time-specific user factors, as shown in Figure 2. The meta-learning network consists of one input layer, two fully connected hidden layers, and one output layer. The first and second hidden layers have 128 and 64 hidden units with ReLU activation, while the last layer estimates time-specific user factors with a linear function followed by sigmoid activation. The input to the meta-learning model is the item embedding for the users on a particular period. Algorithm 1 shows the training process that learns the model parameters. For the time-specific module, the local update (line 7) is done for the user specific parameter, which is achieved by one or more gradients from the global parameter:

$$\theta_u^t = \theta^t - \alpha \nabla_{\theta^t} \mathcal{L}_{\mathcal{T}_u^t}[f_{\theta^t, \omega}^t] \tag{6}$$

In this update, the loss function is computed with the support set. Similarly, the global update (line 11) is conducted with the new item interactions of each user from the query set for the meta update:

$$\theta^t = \theta^t - \beta \nabla_{\theta^t} \sum_{\mathcal{T}_u^t \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_u^t}[f_{\theta_u^t, \omega}^t] \qquad (7)$$

This process continues to find a good global parameter shared by all users in each period.

Time-Evolving Module

User preferences usually change dynamically over time. By capturing the time-evolving factors and integrating them with the time-specific factor, the proposed model can recover the user's true preference more accurately. To this end, we formulate time-evolving user factors (u_{te}^t) for each user using a nested recurrent neural network (RNN):

$$u_{te}^t = f_{rnn}^t(u_{te}^{t-1}, D_u^t; \omega) \qquad (8)$$

where D_u^t is the set of items that user u interacted with at time t, u_{te}^{t-1} is the previous time period time-evolving user factors, and ω is the network parameter. Notice that the input and output of the RNN are both latent variables instead of observations. We use SGD to update the parameter of RNN:

$$\omega = \omega - \gamma \nabla_\omega(\mathcal{L}_{\mathcal{T}_u^t}[f_{\theta_u^t, \omega}^t] + \frac{\lambda}{2}||\omega||_2^2) \qquad (9)$$

where γ is the step size. As shown in the time-evolving module of Figure 2, the vector representation of a hidden layer u_{te}^t is a time-evolving factor of user u at period t and helps to propagate influence from the previous period to the next period (Zhang et al. 2014). The updates of time-specific user factors through meta-learning and time-evolving user factors through nested RNN are summarized in Algorithm 1. The recommendation process is summarized in Algorithm 2.

Algorithm 1: Model training

Require: Set of time periods: $\mathcal{H}$
Require: Hyperparameters: α, β, γ
1: **for** $t \in \mathcal{H}$ **do**
2:      Initialize meta learner, θ^t
3:      **while** not converge **do**
4:          Sample tasks $\mathcal{T}_u^t \sim p(\mathcal{T})$
5:          **for** all $\mathcal{T}_u^t$ **do**
6:              Sample support set $\mathcal{S}_u^t$ for local update
7:              Perform local update with $\mathcal{S}_u^t$ for time-specific module using Equation (6)
8:              Sample query set $\mathcal{Q}_u^t$ for meta update
9:              Update time-evolving module with $\mathcal{D}_u^{t-1}$ using Equation (9)
10:          **end for**
11:          Perform meta update with $\mathcal{Q}_u^t$ for time-specific module using Equation (7)
12:      **end while**
13: **end for**

Algorithm 2: Recommendation for time-specific cold-start users

---

Require: Trained meta parameter θ^t, RNN parameter ω, recommendation time period t
1: Identify cold-start user set for t
2: **for** each user u in the set **do**
3:      Form support set $\mathcal{S}_u^t$ from current interactions
4:      Perform local update with $\mathcal{S}_u^t$ for time-specific module using Equation (6)
5:      Compute user factors using Equations (5) and (8)
6:      Make recommendation using Equation (3)
7: **end for**

Joint Item Embedding Optimization. Let $\mathcal{L}_{emb}$ denote a differentiable loss function used to train the embedding matrix E. And let $\mathcal{G}$ denote the decoding module, which is followed by attribute-wise sigmoid transformation:

$$d_i = \mathcal{G}(Ez_i), \quad [\hat{z}_i]_j = \text{sigmoid}(\eta_j^T d_i)$$
$$= \frac{1}{1 + \exp(-\eta_j^T d_i)} \quad \forall j \in \{1, ..., K\} \qquad (10)$$

where z_i is the original item representation in a binary vector, K is the length of z_i, E is the embedding matrix, d_i is the decoded item representation, η is the parameter for attribute-wise Sigmoid transformation, and $\hat{z}_i$ is the recovered item representation. The loss function for learning the embedding matrix is a negative log-likelihood and is represented as:

$$\mathcal{L}_{emb} = -\sum_{i \in I} \sum_j [z_i]_j \log[\hat{z}_i]_j \qquad (11)$$

Other designs for item embedding with a differentiable loss function can also be applied.

To jointly train the embedding matrix, time-specific and time-evolving modules, we combine those losses, and optimize the total loss with respect to θ^t, ω and E.

$$\arg \min_{\theta^t, \omega, E} \sum_{\mathcal{T}_u^t \sim p(\mathcal{T})} \mathcal{L}_{\mathcal{T}_u^t}[f_{\theta_u^t, \omega, E}^t] + \xi \mathcal{L}_{emb} +$$
$$\frac{\lambda}{2}(||\theta^t||_2^2 + ||\omega||_2^2) \qquad (12)$$
$$\theta_u^t = \theta^t - \alpha \nabla_{\theta^t} \mathcal{L}_{\mathcal{T}_u^t}[f_{\theta^t, \omega, E}^t]$$

where ξ is the weight to be tuned. If ξ is set to a very large value, matrix E is determined only by $\mathcal{L}_{emb}$.

Note that the encoded item embedding is used for both modules. By fixing θ_u^t and ω, it is possible to back-propagate and calculate the gradient with respect to E, which is updated in each task as

$$E = E - \gamma \nabla_E(\mathcal{L}_{\mathcal{T}_u^t}[f_{\theta_u^t, \omega, E}^t] + \mathcal{L}_{emb}) \qquad (13)$$

Also notice that $\mathcal{L}_{emb}$ is not dependent on θ^t and ω. Therefore, when embedding matrix is fixed, the loss function reduces to Eq (4), and θ^t and ω are updated without considering $\mathcal{L}_{emb}$.

Experiments

We conduct experiments on two movie datasets: *Netflix* and *MovieLens-1M* that consist of users' ratings of movies as explicit feedback and one music dataset: *Last.fm* that consists of users' play counts of music tracks as implicit feedback. Besides reporting the overall recommendation performance and comparing with state-of-the-art baselines, we also investigate key properties of the model, including: (1) each module's performance when used in isolation, (2) impact of varying time period lengths, (3) recommendation performance with no interactions in the current period, and (4) impact of hyper-parameters in the model. Experimental settings, datasets details, and analysis of the above key properties are presented in the Appendix (Neupane et al. 2021).

Methods for Comparison. For comparison, we include matrix factorization based static and dynamic models, deep learning based models, graph models, sequential models, and meta-learning models:

- *Matrix factorization (MF):* The standard MF model SVD++ (Koren 2008) that also exploits both explicit and implicit feedback is used here as a static baseline.
- *Dynamic models:* We use timeSVD++ (Koren 2009), collaborative Kalman Filter (CKF) (Gultekin and Paisley 2014), and dynamic Poisson factorization (DPF) (Charlin et al. 2015) as the time-evolving models.
- *Deep learning models:* We use Wide and deep (Cheng et al. 2016), DeepFM (Guo et al. 2017) as static, and DIEN (Zhou et al. 2019) as a dynamic models for deep learning-based recommendation. However, most of them are developed for click-through rate prediction in their original forms.
- *Graph-based model:* Most graph-based models are designed for static settings. For comparison, we use graph convolutional matrix completion (GC-MC) (Berg, Kipf, and Welling 2017), which models recommendation as link prediction in the graph, and neural graph collaborative filtering (NGCF) (Wang et al. 2019b) that utilizes embedding propagation over user-item graphs.
- *Sequential model:* We use Sequential Recommendation via Convolutional Sequence Embedding (Caser) (Tang and Wang 2018), which models recommendation as a unified and flexible structure to capture both preferences and sequential patterns, and transformer-based sequential recommendation model (SASRec) (Kang and McAuley 2018) in our comparison.
- *Meta-learning models:* We follow the model-agnostic meta-learning model (MAML) to implement the meta-learning model similar to MeLU (Lee et al. 2019). We also compared with the meta-learning model in (Vartak et al. 2017) that focuses on item cold-start problem (referred to as ML-ICS). The model is also designed for the classification setting, so we have to make adjustments to fit into our context.

Evaluation Metrics. For evaluation, we analyze the experimental results in terms of both the deviation of predicted values from the ground truth and the errors of the ranking sequences. We use Root Mean Squared Error (RMSE) and

Normalized Discounted Cumulative Gain (NDCG) averaged across all test users. RMSE is usually reported for explicit data, while NDCG is usually reported for implicit data:

$$\text{RMSE} = \sqrt{\sum_{r_{u,i} \in O} (\hat{r}_{u,i} - r_{u,i})^2 / |O|},$$

$$\text{NDCG}_u = \sum_n \frac{rel_n^{pred}}{\log_2(1+n)} / \sum_n \frac{rel_n^{ideal}}{\log_2(1+n)} \quad (14)$$

where O is the observation set for the test set and rel_n is the relevancy of n^{th} item in the ranking sequence for user u, which is binary for implicit data or the rating for explicit data. To penalize the negative feedback, we linearly mapped the ratings to a range of [-1,1]. The NDCG is the fraction of Discounted Cumulative Gain (DCG) of recommendation result over the ideal DCG.
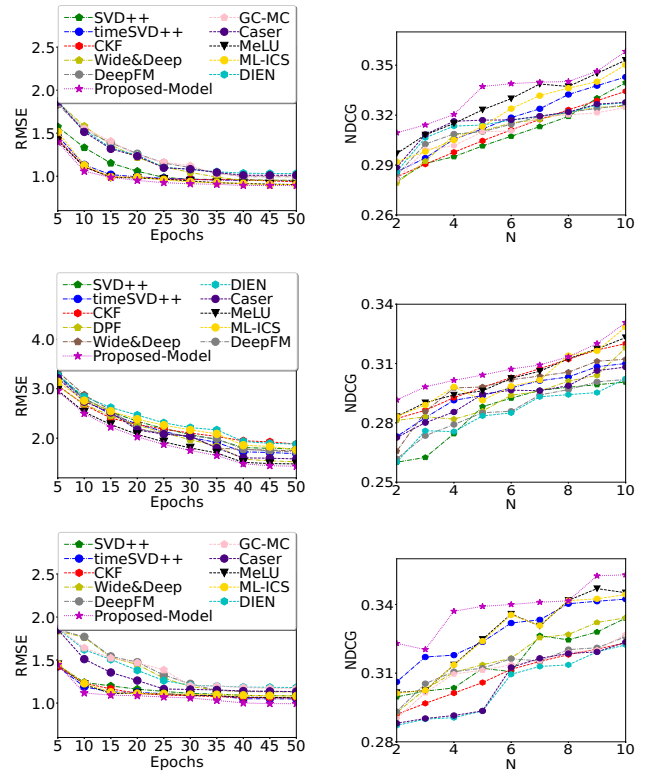
Figure 3: NDCG and RMSE Performance: Netfix (top), Last.fm (middle), and MovieLens-1M (bottom)

Recommendation Performance

The experimental results are shown in Figure 3. We evaluate NDCG based on the top N recommendation list and RMSE based on the training epochs. The RMSE is stable after 30-40 epochs in all datasets.

The average results of NDCG and RMSE considering all periods with the range of deviation are shown in Table 1, for the Netflix, Last.fm, and MovieLens datasets, respectively. The proposed model clearly demonstrates the advantage of combining time-specific and time-evolving user factors that

Category	Model	Netflix		Last.fm		MovieLens-1M	
		RMSE	NDCG	RMSE	NDCG	RMSE	NDCG
MF	SVD++	0.9797±0.03	0.2915	1.7829±0.08	0.2882	1.0825±0.04	0.3023
Dynamic	timeSVD++	0.9538±0.06	0.3115	1.6912±0.11	0.2962	1.0483±0.03	0.3224
	CKF	0.9337±0.04	0.3130	1.5316±0.32	0.3018	1.0652±0.04	0.3151
	DPF	N/A	N/A	1.5227±0.43	0.3085	N/A	N/A
Deep Learning	Wide and Deep	0.9904±0.04	0.2864	1.7253±0.22	0.2727	1.1364±0.06	0.2932
	DeepFM	0.9811±0.03	0.2930	1.6815±0.21	0.2971	1.1723±0.05	0.2882
	DIEN	1.0345±0.04	0.2832	1.9225±0.26	0.2714	1.1872±0.14	0.2843
Graph	GC-MC	1.0760±0.03	0.2901	N/A	N/A	1.1704±0.08	0.2913
	NGCF	1.0321±0.03	0.3026	1.5612±0.23	0.2896	1.1216±0.05	0.3103
Sequential	Caser	1.0124±0.03	0.3101	1.5824±0.31	0.2931	1.1339±0.08	0.3012
	SASRec	N/A	0.3246	N/A	0.3103	N/A	0.3238
Meta-Learning	MeLU	0.9213±0.05	0.3232	1.2580±0.28	0.3122	1.0685±0.08	0.3214
	ML-ICS	0.9332±0.04	0.3173	1.2408±0.24	0.3142	1.0845±0.06	0.3244
Proposed	**Ours**	**0.8925±0.03**	**0.3472**	**1.2203±0.16**	**0.3385**	**0.9945±0.08**	**0.3351**

Table 1: Recommendation Results (RMSE and NDCG)

lead to a superior recommendation accuracy as compared with other competitive models. Both explicit and implicit datasets are highly sparse, and MF models' performance is poor due to the sparse interactions. Also, MF models suffer from cold-start problems, and thus their performances are fairly limited, as shown in Table 1. Similarly, deep learning models require sufficient training data and hence largely suffer in the few-shot recommendation setting. Moreover, these models might need extra side information, like user profile and item details, for better recommendations. For example, DIEN needs cleverly chosen interest features like user behavior, and the absence of those features limits its performance, as shown in Table 1. Similarly, the poor performance of graph-based models in both movie datasets implies that these methods are insufficient to handle cold-start problems. Like other existing models, the performance of a sequential model is less effective for the cold-start users in all three datasets. The reason could be that the model is less effective in capturing long-term user preferences. In contrast, the meta-learning approaches show better results by leveraging shared knowledge across the users. However, in the time-specific cold-start setting, test users have very limited interactions. In particular, the meta-learning model doesn't benefit from time-evolving aspects of the user interests, and thus underperforms the proposed model.

We use an illustrative example to further demonstrate how the proposed model effectively captures the underlying user interest and its evolution in Figure 4. The recommended movie genres are compared to the user's favorite genres based on the provided true ratings. The result shows that the recommendation matches user's changing taste over time well. It is also interesting to see that the proposed model accurately detects some dramatic changes in user's ratings (e.g., from 12/02 to 06/03 and from 06/04 to 12/04), which were likely to be caused by some time-specific factors.

We further present an example to show how the meta-learning module effectively captures time-specific factors in the form of popular trends in a specific period from the global user space and transfers the (meta) knowledge to the cold-start users with very limited interactions. Table 2 demonstrates top-5 time-specific (period 4) popular movies learned by the meta-learning module, which shares


Figure 4: Dynamic trend of movie genres in Netflix: Ground truth (left) and recommended (right)

Users	Movies
Training users	['Best in Show', 'Chicken Run', 'Sommersby', 'Bedrooms and Hallways', 'The Mod Squad']
Test user (ID:5636)	['Mr. Mom', 'Best in Show', 'Shower', "We're No Angels", 'Groundhog Day']
Test user (ID:5539)	['Best in Show', 'An Ideal Husband', 'Life Is Beautiful', 'Breaking Away', 'Kramer vs. Kramer']

Table 2: Time-specific popular movies learned and predicted by the meta-learning module

that knowledge with the test users (i.e., two users are shown in Table 2 and some time-specific popular movies like 'Best in Show' are recommended to them). This example demonstrates how our model provides effective recommendations to users by capturing time-specific factors.

Conclusion

In this paper, we formulate a novel time-sensitive cold-start problem and present a dynamic recommendation framework to address its unique challenges. The framework integrates a time-sensitive meta-learning module with a time-evolving recurrent module. The former handles the user cold-start problem by learning global knowledge among users from their interaction information in the current recommendation period. This module is jointly optimized with the time-evolving recurrent module that captures a user's gradually shifted preferences. A merged user representation is generated using the two modules' outputs and interacts with the item embedding to provide the final recommendations.

Acknowledgements

Ethical Impact

The proposed model is generally applicable to many real-world applications where maintaining existing users' participation is critical, such as product recommendation in e-business and content recommendation in social media. Besides, this model can operate on multiple platforms such as health, education, and e-commerce to handle those sensitive users in the system. Since the proposed model does not rely on using users' profile information, it has the potential to improve the privacy protection of recommender systems.

References

Bengio, S.; Bengio, Y.; Cloutier, J.; and Gecsei, J. 1992. On the optimization of a synaptic learning rule. In *Preprints Conf. Optimality in Artificial and Biological Neural Networks*, volume 2. Univ. of Texas.

Berg, R. v. d.; Kipf, T. N.; and Welling, M. 2017. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263*.

Charlin, L.; Ranganath, R.; McInerney, J.; and Blei, D. M. 2015. Dynamic poisson factorization. In *Proceedings of the 9th ACM Conference on Recommender Systems*, 155–162.

Cheng, H.-T.; Koc, L.; Harmsen, J.; Shaked, T.; Chandra, T.; Aradhye, H.; Anderson, G.; Corrado, G.; Chai, W.; Ispir, M.; et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*, 7–10.

Du, Z.; Wang, X.; Yang, H.; Zhou, J.; and Tang, J. 2019. Sequential Scenario-Specific Meta Learner for Online Recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2895–2904.

Fang, H.; Zhang, D.; Shu, Y.; and Guo, G. 2020. Deep learning for sequential recommendation: Algorithms, influential factors, and evaluations. *ACM Transactions on Information Systems (TOIS)*, 39(1): 1–42.

Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 1126–1135. JMLR. org.

Funk, S. 2006. "Netflix Update: Try This At Home". https://sifter.org/simon/journal/20061211.html.

Goldberg, D.; Nichols, D.; Oki, B. M.; and Terry, D. 1992. Using collaborative filtering to weave an information tapestry. *Communications of the ACM*, 35(12): 61–70.

Gultekin, S.; and Paisley, J. 2014. A collaborative kalman filter for time-evolving dyadic processes. In *2014 IEEE International Conference on Data Mining*, 140–149. IEEE.

Guo, H.; Tang, R.; Ye, Y.; Li, Z.; and He, X. 2017. DeepFM: a factorization-machine based neural network for CTR prediction. *arXiv preprint arXiv:1703.04247*.

Guo, Q.; Zhuang, F.; Qin, C.; Zhu, H.; Xie, X.; Xiong, H.; and He, Q. 2020. A survey on knowledge graph-based recommender systems. *IEEE Transactions on Knowledge and Data Engineering*.

Kang, W.-C.; and McAuley, J. 2018. Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*, 197–206. IEEE.

Koren, Y. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, 426–434.

Koren, Y. 2009. Collaborative filtering with temporal dynamics. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, 447–456.

Koren, Y.; Bell, R.; and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *Computer*, 42(8): 30–37.

Lee, H.; Im, J.; Jang, S.; Cho, H.; and Chung, S. 2019. MeLU: Meta-Learned User Preference Estimator for Cold-Start Recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1073–1082.

Li, B.; Zhu, X.; Li, R.; Zhang, C.; Xue, X.; and Wu, X. 2011. Cross-domain collaborative filtering over time. In *Twenty-Second International Joint Conference on Artificial Intelligence*.

Lu, Y.; Fang, Y.; and Shi, C. 2020. Meta-learning on heterogeneous information networks for cold-start recommendation. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1563–1573.

Mnih, A.; and Salakhutdinov, R. R. 2008. Probabilistic matrix factorization. In *Advances in neural information processing systems*, 1257–1264.

Neupane, K.; Zheng, E.; Kong, Y.; and Yu, Q. 2021. Appendix: A Dynamic Meta-Learning Model for Time-Sensitive Cold-Start Recommendations. https://github.com/ritmininglab/A-Dynamic-Meta-Learning-Model-for-Time-Sensitive-Cold-Start-Recommendations/blob/main/Appendix.pdf.

Pan, F.; Li, S.; Ao, X.; Tang, P.; and He, Q. 2019. Warm up cold-start advertisements: Improving ctr predictions via learning to learn id embeddings. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 695–704.

Sahoo, N.; Singh, P. V.; and Mukhopadhyay, T. 2012. A hidden Markov model for collaborative filtering. *Mis Quarterly*, 1329–1356.

Schmidhuber, J. 1987. *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook*. Ph.D. thesis, Technische Universität München.

Su, X.; and Khoshgoftaar, T. M. 2009. A survey of collaborative filtering techniques. *Advances in artificial intelligence*, 2009.

Sun, J. Z.; Parthasarathy, D.; and Varshney, K. R. 2014. Collaborative kalman filtering for dynamic matrix factorization. *IEEE Transactions on Signal Processing*, 62(14): 3499–3509.

Tang, J.; and Wang, K. 2018. Personalized top-n sequential recommendation via convolutional sequence embedding. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, 565–573.

Uyangoda, L.; Ahangama, S.; and Ranasinghe, T. 2018. User profile feature-based approach to address the cold start problem in collaborative filtering for personalized movie recommendation. In *2018 Thirteenth International Conference on Digital Information Management (ICDIM)*, 24–28. IEEE.

Vartak, M.; Thiagarajan, A.; Miranda, C.; Bratman, J.; and Larochelle, H. 2017. A meta-learning perspective on cold-start recommendations for items. In *Advances in neural information processing systems*, 6904–6914.

Vinyals, O.; Blundell, C.; Lillicrap, T.; Wierstra, D.; et al. 2016. Matching networks for one shot learning. In *Advances in neural information processing systems*, 3630–3638.

Wang, S.; Hu, L.; Wang, Y.; Cao, L.; Sheng, Q. Z.; and Orgun, M. 2019a. Sequential recommender systems: challenges, progress and prospects. *arXiv preprint arXiv:2001.04830*.

Wang, X.; He, X.; Wang, M.; Feng, F.; and Chua, T.-S. 2019b. Neural graph collaborative filtering. In *Proceedings of the 42nd international ACM SIGIR conference on Research and development in Information Retrieval*, 165–174.

Wei, T.; Wu, Z.; Li, R.; Hu, Z.; Feng, F.; He, X.; Sun, Y.; and Wang, W. 2020. Fast Adaptation for Cold-start Collaborative Filtering with Meta-learning. In *2020 IEEE International Conference on Data Mining (ICDM)*, 661–670. IEEE.

Xie, F.; Chen, L.; Ye, Y.; Zheng, Z.; and Lin, X. 2018. Factorization machine based service recommendation on heterogeneous information networks. In *2018 IEEE International Conference on Web Services (ICWS)*, 115–122. IEEE.

Xie, R.; Wang, Y.; Wang, R.; Lu, Y.; Zou, Y.; Xia, F.; and Lin, L. 2021. Long Short-Term Temporal Meta-learning in Online Recommendation. *arXiv preprint arXiv:2105.03686*.

Yao, Y.; Tong, H.; Yan, G.; Xu, F.; Zhang, X.; Szymanski, B. K.; and Lu, J. 2014. Dual-regularized one-class collaborative filtering. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, 759–768.

Zhang, Y.; Dai, H.; Xu, C.; Feng, J.; Wang, T.; Bian, J.; Wang, B.; and Liu, T.-Y. 2014. Sequential click prediction for sponsored search with recurrent neural networks. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*.

Zhou, G.; Mou, N.; Fan, Y.; Pi, Q.; Bian, W.; Zhou, C.; Zhu, X.; and Gai, K. 2019. Deep interest evolution network for click-through rate prediction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 5941–5948.