

NeuSE: A Neural Snapshot Ensemble Method for Collaborative Filtering

DONGSHENG LI*, Fudan University, China

HAODONG LIU*, University of Electronic Science and Technology of China, China

CHAO CHEN, Shanghai Jiao Tong University, China

YINGYING ZHAO[†], Fudan University, China

STEPHEN M. CHU, IBM Research - China, China

BO YANG[†], University of Electronic Science and Technology of China, China

In collaborative filtering (CF) algorithms, the optimal models are usually learned by globally minimizing the empirical risks averaged over all the observed data. However, the global models are often obtained via a performance tradeoff among users/items, *i.e.*, not all users/items are perfectly fitted by the global models due to the hard non-convex optimization problems in CF algorithms. Ensemble learning can address this issue by learning multiple diverse models but usually suffer from efficiency issue on large datasets or complex algorithms. In this paper, we keep the intermediate models obtained during global model learning as the snapshot models, and then adaptively combine the snapshot models for individual user-item pairs using a memory network-based method. Empirical studies on three real-world datasets show that the proposed method can extensively and significantly improve the accuracy (up to 15.9% relatively) when applied to a variety of existing collaborative filtering methods.

Additional Key Words and Phrases: collaborative filtering, snapshot ensemble, neural networks

ACM Reference Format:

Dongsheng Li, Haodong Liu, Chao Chen, Yingying Zhao, Stephen M. Chu, and Bo Yang. 2021. NeuSE: A Neural Snapshot Ensemble Method for Collaborative Filtering. 1, 1 (April 2021), 20 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Collaborative filtering (CF) is the most important class of methods in today's recommender systems [1, 2]. In many existing CF algorithms, *e.g.*, matrix factorization (MF), the models are usually learned via minimizing empirical risks. Since the loss functions usually aggregate the losses over all users and items, convergences are reached when most users and items are appropriately modeled but not all of them [4, 20]. Our case studies

*Dongsheng Li and Haodong Liu contributed equally to this work.

[†]Yingying Zhao and Bo Yang are the corresponding authors.

Authors' addresses: Dongsheng Li, Fudan University, 2005 Songhu Road, Yangpu District, Shanghai, China, dongshengli@fudan.edu.cn; Haodong Liu, University of Electronic Science and Technology of China, 2006 Xiyuan Ave, Chengdu, Sichuan, China, haodong.lhd@gmail.com; Chao Chen, Shanghai Jiao Tong University, 800 Dongchuan Road, Minhang District, Shanghai, China, chao.chen@sjtu.edu.cn; Yingying Zhao, Fudan University, 2005 Songhu Road, Yangpu District, Shanghai, China, yingyingzhao@fudan.edu.cn; Stephen M. Chu, IBM Research - China, 55 Chuanhe Road, Pudong New District, Shanghai, China, schu@cn.ibm.com; Bo Yang, University of Electronic Science and Technology of China, 2006 Xiyuan Ave, Chengdu, Sichuan, China, yangbo@uestc.edu.cn.

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 ACM 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.

© 2021 Association for Computing Machinery.

XXXX-XXXX/2021/4-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

show that the global models can perfectly fit some users/items, but the other users/items may suffer from either overfitting or underfitting. Recently, Beutel et al. [4] also proved that globally optimized models are typically not the best models for every subset of the data. In summary, the globally optimized models often yield suboptimal performances in CF problems [4, 19, 20].

One may imagine that adaptive learning rate methods, *e.g.*, AdaGrad [7], can address this issue by manipulating the learning of individual users/items to achieve a “truly” global optimum. However, the optimization problem in CF algorithms are generally non-convex due to the co-existence of users and items [22, 28], so that adaptive learning rate methods can only help to better converge to one local optimum but cannot explore other local optimums that may achieve better performance as demonstrated in our case studies. Another kind of possible solutions may be Bayesian approaches [26, 29], which perform approximate inference by integrals from sampling. However, Bayesian approaches generally face two main challenges [29]: 1) choosing the prior distributions for the unknown model parameters is non-trivial and 2) these methods are usually computationally demanding and thus limited to small-scale problems. Therefore, ensemble learning-based CF algorithms [19, 20], which learn multiple base models each of which can achieve better performance over a subset of data, are more appropriate to address the suboptimal performance issue of global models [4]. However, the learning of multiple base models in ensemble methods will suffer from efficiency issue on large datasets, *e.g.*, the Yahoo! Music dataset [15], or on very complex models, *e.g.*, deep neural networks [9]. Therefore, it is desirable to design ensemble methods to address the suboptimal performance issue of global models without significantly increasing the computational overhead.

Following recent works that create ensembles from the slices of the learning trajectory [12, 35], we also combine the snapshot models obtained during global model learning, so that ensemble learning can be achieved with little additional training costs. Different from existing works which simply average the snapshot models [12, 35, 41], we propose a **Neural network-based Snapshot Ensemble (NeuSE)** method, in which three memory networks are introduced to weigh different snapshot models: 1) model memory network to capture the global importance of each snapshot model, 2) user memory network to capture the user-wise importance of each snapshot model and 3) item memory network to capture the item-wise importance of each snapshot model. Then, a fully connected layer with three mapping modules are adopted to combine the outputs of memory networks to produce adaptive ensemble weights for different user-item pairs. To achieve better generalization performance, we propose a soft label-based learning method, which can reduce the loss when the predicted snapshot model is close to the optimal snapshot model and increase the loss otherwise. Empirical studies on several real-world datasets demonstrate that the proposed method can extensively improve the performance of existing collaborative filtering methods (up to 15.9% relatively) on almost all kinds of algorithms and recommendation tasks.

2 RELATED WORK

Collaborative filtering (CF) has been extensively studied in recent years and the state-of-the-art performances are achieved via machine learning techniques, *e.g.*, supervised learning [26, 42] and unsupervised learning [13, 16, 37]. However, due to the non-convex nature of CF problems, many existing algorithms cannot learn the “optimal” model [4, 19, 20], *i.e.*, the globally optimized models are typically not optimal for each subset of the data [4]. Therefore, ensemble learning [19, 20], which learn multiple base models each of which can achieve better performance over a subset of data, are more appropriate to address the suboptimal performance issue [4]. However, efficiency issue may arise in ensemble methods due to the learning of multiple models, especially on large datasets [12] or complex neural network methods [9]. In addition, many existing ensemble collaborative filtering methods were proposed for specific CF methods and thus may not be applied to other CF methods, *e.g.*, LLORMA [19] achieved excellent ensemble performance on matrix factorization methods but cannot be applied to UserKNN or ItemKNN methods. Different from existing ensemble CF methods, *the goal of this work*

is to improve the performances of almost all kinds of existing CF algorithms without significantly increasing the training overheads.

Many techniques have been proposed to learn better machine learning models in general tasks, e.g., adaptive learning rate [7, 21], Dropout [33], and Bayesian approaches [26, 29, 40], etc. However, adaptive learning rate methods are not capable of exploring different local optimums to achieve better performance. Dropout can indeed improve the performances of neural network-based CF methods [36] but cannot be applied to non-neural network methods, e.g., matrix factorization [16]. Factorization models learned via Bayesian treatment [26, 29], e.g., MCMC, achieved superior performance over those learned via SGD or ALS. But as mentioned earlier, Bayesian approaches usually suffer from unknown prior distribution and scalability issues. In addition, *the proposed method is orthogonal to the above methods and thus can be combined with them to further improve performance as demonstrated in the experiments.*

This work follows existing works that create ensembles from the slices of learning trajectory. Swann et al. [35] and Xie et al. [41] both suggested to average the outputs of snapshot models during stable range of epochs to improve model performance. Huang et al. [12] proposed to learn multiple neural network models converging to different local minimums during one optimization path via cyclic learning rate schedules. However, simply averaging the outputs of all snapshot models may raise challenges to snapshot model selection. As pointed out by Swann et al. [35], models with poor generalization ability should not be selected, because averaging poor models may affect the ensemble performance. Therefore, the ensemble method should be capable of weighing the importance of different snapshot models so that poor models can be excluded most of the time. Xie et al. [41] also adopted a simple MLP on top of the snapshot models to learn the ensemble weights of different classifiers. However, their method only captures the global importance of snapshot models, which is not enough in collaborative filtering tasks where each user may have varying ensemble weights across different items. Our empirical studies show that a simple MLP can only achieve similar performance with simple averaging the outputs of the snapshot models. Different from the above works, *the proposed method can adaptively change the ensemble weights for different user-item pairs, so that a user may have different ensemble weights on different items to better characterize the non-convexity of collaborative filtering.*

3 PROBLEM FORMULATION

This section first analyzes the non-convex optimization problem in collaborative filtering. Then, we motivate the snapshot ensemble problems by three case studies.

3.1 Preliminaries

For many machine learning-based CF methods, due to the coexistence of users and items, both user parameters and item parameters should be learned via minimizing predefined loss functions over the entire dataset. Then, an interaction function between user parameters and item parameters can be adopted to generate recommendation scores, e.g., matrix factorization methods use dot-product as the interaction function [16] and NeuMF uses multi-layer perceptron as the interaction function [9]. Given the observed rating matrix R and a loss function \mathcal{L} , the optimal user parameters (U^*) and item parameters (V^*) are obtained as follows:

$$U^*, V^* = \arg \min_{U, V} \mathcal{L}(R|U, V). \quad (1)$$

To learn the globally optimal U, V , iterative learning methods, e.g., SGD [16], ALS [11, 26] and MCMC [26], are usually adopted. \mathcal{L} can be convex to U or V alone, e.g., in the mean square loss function, but it is usually non-convex to U, V together. Therefore, many local optimums exist due to the non-convex optimization problem. Although it has been proved that any local optimum is also global optimum in matrix factorization problem under certain assumption [8], i.e., on symmetric positive definite matrices, there are no such results for general CF problems to the best of our knowledge because user-item rating matrices in CF tasks are asymmetric.




Fig. 1. Test RMSE trends on rating prediction task for users with different number of ratings using the RSVD method [25], *i.e.*, “<50” means users with less than 50 ratings and so forth. We train the model using vanilla SGD in figure (a) and using SGD with adagrad [7] in figure (b). Fig. 2. Performance of three strategies: 1) single (the optimal snapshot model), 2) set 1 (epoch = 50, 100, 150, ..., 500), 3) set 2 (epoch = 100, 200, 300).

3.2 Case Studies

This section presents three case studies on the MovieLens 1M dataset to motivate the proposed method, in which we randomly split the data into training and test sets by 9:1. In Figure 1 and Figure 2, we use one of the classic matrix factorization methods – the RSVD method [25] in the rating prediction task.

3.2.1 Global vs. Local Convergence. Figure 1 shows the test errors of different kinds of users with increasing number of epochs in the rating prediction task. As shown in Figure 1 (a), users with different numbers of ratings converge at different epochs and the global optimum is achieved with epoch around 175 which is near the optimum for the majority of users (with 50-500 ratings). Therefore, users with less than 50 ratings and users with more than 500 ratings will suffer from underfitting and overfitting, respectively, if we only use the globally optimal model. Similar phenomenon can also be observed from Figure 1 (b), in which the model is trained with AdaGrad SGD instead of vanilla SGD in Figure 1 (a). Due to the co-existence of users and items, the optimal snapshot models are also varying for different kinds of items. Therefore, the optimal snapshot model for a user could vary across different items (account for about 90% of ratings in the case study). This case study confirms that a globally optimal model is typically not the optimum for all subsets of the data. Therefore, it is necessary to choose different snapshot models for different user-item pairs to achieve better accuracy.

It should be noted that the test errors for different kinds of users in Figure 1 (b) are relatively more stable than those in Figure 1 (a), which seems that the suboptimal performance issue of collaborative filtering can be alleviated using adaptive learning rate methods. However, we can also observe in this study that the optimal (global) model trained via AdaGrad SGD achieved almost the same test RMSE compared with the optimal (global) model trained via vanilla SGD. The same phenomenon has also been confirmed by several recent works [32, 39], in which marginal improvements are observed on adaptive learning rate methods in several machine learning tasks. Therefore, we claim that adaptive learning rate methods can help to easier determine the optimal (global) model, *i.e.*, near-optimal models can be obtained in a relatively large range of epochs, but these methods cannot guarantee superior performance compared to conventional learning methods and thus cannot solve the suboptimal performance issue fundamentally.

3.2.2 Model Selection. In ensemble learning, the performance may degrade if poor base models are selected [43]. Therefore, snapshot model selection is crucial to existing methods [35]. Xie et al. [41] suggested to select snapshots from “a stable range of epoch”, which is not applicable to some CF problems where stable ranges are unavailable (*e.g.*, Figure 1 (a)). Figure 2 compares the accuracy of three selection strategies: 1) single (the globally optimal

snapshot model), 2) set 1 (all snapshot models from every 50 epochs till epoch = 500) and 3) set 2 (snapshot models with epoch = 100, 200, 300). Here, we report the test RMSE of a simple average over all snapshot models with ten random data splits. As shown in the figure, if the snapshot models are not properly selected (set 1), *i.e.*, too noisy, the ensemble performance of averaging the snapshot models is even worse than the optimal single model. On the contrary, if the snapshot models are properly selected (set 2), a simple average can achieve observable accuracy improvement, *e.g.*, around 0.3% relatively in this study. This study indicates that an ensemble learner is needed to assign higher weights to the superior models and lower weights for inferior models for different user-item pairs to overcome the challenge of snapshot model selection.

4 NEUSE: NEURAL SNAPSHOT ENSEMBLE

The proposed method consists of two key components:

- **Snapshot Selection.** Due to the non-convex optimization problems in CF, it is challenging to intentionally learn several locally optimal models. As shown in the case study, each intermediate model in the learning trajectory could be optimal for a subset of users and items, which should be all considered in the ensemble. Thus, we propose to select snapshot models with fixed epoch variations (ΔT), *i.e.*, we keep the snapshot models for every ΔT epochs during training. For non-machine learning-based CF methods, we can also select intermediate models as snapshots, *e.g.*, models with different number of neighbors in memory-based methods.
- **Ensemble Learning.** For each user-item pair, each snapshot model will have a predicted rating, all of which are used as the input of the snapshot ensemble method. The ensemble method adopts a multi-memory network to capture the global, user-specific and item-specific importance of the snapshot models via different memory networks. To ensure robust ensemble learning, we change the hard label (optimal snapshot model as 1 and 0 for all others) to soft labels, because nearby snapshot models of the optimal model could have much smaller error than faraway snapshot models (as shown in Figure 1). After learning the ensemble model, we can use its output as the adaptive ensemble weights.

4.1 Network Architecture

The key step in NeuSE is to adaptively generate the ensemble weights for different ratings. However, the ensemble model could easily overfit if we only consider the limited ratings of individual users/items. To remedy this, we adopt the memory networks [38] in NeuSE, which can maintain long-term memories to facilitate inference. The memory components store the weight vectors of global information and the top k neighbors of each user/item, which can help to address the limited data issue of individual users and items by leveraging the global information and user/item neighborhood information. As illustrated in Figure 3, the proposed ensemble network consists of three layers: 1) input layer, 2) memory layer, and 3) output layer.

4.1.1 Input Layer. We use the identities of a user and an item as the input feature via one-hot encoding, *i.e.*, we represent the u -th user with a sparse vector in which only the u -th element is 1 and all other elements are 0. The sparse representation of a user or an item is then projected to a dense vector $\mathbf{e}_u \in \mathbb{R}^d$ or $\mathbf{e}_i \in \mathbb{R}^d$ using one fully connected layer, which can be seen as the latent vector for the user or item. Given $D = 2d$, we concatenate the two embedding vectors to construct $\mathbf{e}_{ui} \in \mathbb{R}^D$, which is then fed into the subsequent memory layer.

4.1.2 Memory Layer. The memory layer consists of three components: model memory, user memory and item memory.

Model Memory. The model memory consists of an internal memory matrix $\mathbf{M}_m \in \mathbb{R}^{D \times N_m}$ and an external memory matrix $\mathbf{C}_m \in \mathbb{R}^{D \times N_m}$, where N_m represents the number of snapshot models and D denotes the dimensionality of each memory cell. Each snapshot model s is embedded in two memory slots $\mathbf{m}_s \in \mathbf{M}_m$ and $\mathbf{c}_s \in \mathbf{C}_m$, storing the




Fig. 3. Architecture of the proposed multi-memory network for snapshot ensemble with one memory layer.

specific internal and external attributes of it, respectively. For each snapshot model s , the corresponding internal memory slot \mathbf{m}_s is used to measure the relevance between s and user u and item i , *i.e.*, if \mathbf{m}_s is similar to \mathbf{e}_{ui} then we know that model s is important to predict u 's rating on i . In other words, \mathbf{m}_s captures how many user-item pairs can use model s to accurately predict their ratings. Different from internal memory, each slot in the external memory, *e.g.*, \mathbf{c}_s , can be regarded as the representation of each snapshot model.

More specifically, we form a snapshot model preference vector $\mathbf{p}_s^M \in \mathbb{R}^D$ – each dimension p_{uis}^M is the level of compatibility among snapshot model s , user u and item i as follows:

$$p_{uis}^M = \mathbf{m}_s^T \mathbf{e}_{ui}, \quad \forall s \in N(m), \quad (2)$$

where $N(m)$ represents the set of snapshot models. In other words, p_{uis}^M measures the importance of snapshot model s for the rating made by user u on item i .

To adaptively weigh snapshot models, we need to predefine some heuristic functions, *e.g.*, cosine similarity, to distinguish superior snapshot models from inferior ones. Therefore, we introduce the attention mechanism, in which we can learn an adaptive weighting function to focus on only a subset of the influential snapshot models to derive the final score. More formally, we compute the attention weight for a given snapshot model to infer its unique contribution as follows:

$$w_{uis}^M = \frac{\exp(p_{uis}^M)}{\sum_{k \in N(m)} \exp(p_{uik}^M)}, \quad \forall s \in N(m). \quad (3)$$

The attention mechanism allows our model to place higher weights on appropriate snapshot models and lower weights on ones that may be less compatible with user u and item i . Then, we can construct the final representation of model memory by combining the external model memory with the attention weights as follows:

$$\mathbf{q}_m = \sum_{s \in N(m)} w_{uis}^M \mathbf{c}_s. \quad (4)$$

Here, \mathbf{c}_s is the embedding vector for snapshot model s in the external model memory, which stores the long-term information pertaining specifically to each snapshot model. In Equation 4, the attention mechanism can selectively weigh the snapshot models according to the targeted user and item, and the output of model memory \mathbf{q}_m represents a weighted sum over the snapshot models composed of the compatibility among the targeted user, item and all snapshot models.

User Memory and Item Memory. Similar to model memory, the user memory and item memory both consist of an internal memory and an external memory to store the long-term memories of the snapshot models for each user and item, respectively. The memory components work as neighborhood aggregation function to select similar users/items via neural attention mechanism to mitigate the insufficient data issue of individual users/items. For each user v in the neighborhood, we form user u 's preference vector $\mathbf{p}_u^U \in \mathbb{R}^D$ in which each dimension p_{uiv}^U measures the level of similarity between the target user u and its neighbor v on the given item i as follows:

$$p_{uiv}^U = \mathbf{m}_v^T \mathbf{e}_{ui}, \quad \forall v \in N(i), \quad (5)$$

where the neighborhood $N(i)$ represents the set of all users who have interactions with item i , and \mathbf{m}_v is the embedding vector of user v in the internal user memory matrix $\mathbf{M}_u \in \mathbb{R}^{D \times N_i}$, where N_i denotes the number of users in $N(i)$. Then, the weight of each neighbor v on item i can be described as follows:

$$w_{uiv}^U = \frac{\exp(p_{uiv}^U)}{\sum_{k \in N(i)} \exp(p_{uik}^U)}, \quad \forall v \in N(i). \quad (6)$$

w_{uiv}^U produces a distribution over the neighborhood of u on item i . Then, we can construct the final neighborhood-based representation of user u as follows:

$$\mathbf{q}_u = \sum_{v \in N(i)} w_{uiv}^U \mathbf{c}_v, \quad (7)$$

where \mathbf{c}_v is the memory slot for user v in the external user memory matrix $\mathbf{C}_u \in \mathbb{R}^{D \times N_i}$.

The item memory network is similarly designed as the user memory network. Given an item j in the neighborhood $N(u)$ – the set of all items having interactions with user u , we can describe the process in item memory as follows:

$$p_{uij}^I = \mathbf{m}_j^T \mathbf{e}_{ui}, \quad \forall j \in N(u), \quad (8)$$

where \mathbf{m}_j is the embedding vector of item j in the internal item memory matrix $\mathbf{M}_i \in \mathbb{R}^{D \times N_u}$. N_u denotes the number of items in $N(u)$. Then, the weight of each neighboring item j can be described as follows:

$$w_{uij}^I = \frac{\exp(p_{uij}^I)}{\sum_{k \in N(u)} \exp(p_{uik}^I)}, \quad \forall j \in N(u). \quad (9)$$

Finally, we construct the neighborhood-based representation of item i as follows:

$$\mathbf{q}_i = \sum_{j \in N(u)} w_{uij}^I \mathbf{c}_j, \quad (10)$$

where \mathbf{c}_j is the memory slot for item j in the external item memory matrix $\mathbf{C}_i \in \mathbb{R}^{D \times N_u}$.

4.1.3 Output Layer. The output layer takes the predictions of all snapshot models and the outputs of the memory networks to make the final prediction. Similar to Biased MF [16] and FM [26], we introduce global bias, user bias and item bias by a linear projection from their latent representations $(\mathbf{q}_m, \mathbf{q}_u, \mathbf{q}_i)$, which are helpful to capture the unique characteristics of data, user and item, respectively.

We project the final representations of the three memory networks to global/user/item biases as follows:

$$b_m = \mathbf{W}_{model}^T \mathbf{q}_m + b_{model}, \quad (11)$$

$$b_u = \mathbf{W}_{user}^T \mathbf{q}_u + b_{user}, \quad (12)$$

$$b_i = \mathbf{W}_{item}^T \mathbf{q}_i + b_{item}. \quad (13)$$

$\mathbf{W}_{model}, \mathbf{W}_{user}, \mathbf{W}_{item} \in \mathbb{R}^{D \times 1}$ and $b_{model}, b_{user}, b_{item}$ are parameters to be learned.




Fig. 4. Architecture of the proposed deep multi-memory network for snapshot ensemble with multiple memory layers. Each hop denotes one layer of model/user/item memory networks with feature mapping layers. The neighborhood information is the same as that in Figure 3, which is omitted for simplicity.

Then, we concatenate the outputs of all snapshot models (\mathbf{q}_p), the latent representations from the three memory networks ($\mathbf{q}_m, \mathbf{q}_u, \mathbf{q}_i$) and the three bias terms above (b_m, b_u, b_i) together, and feed them into a fully connected layer with softmax activation function as follows:

$$\hat{\mathbf{y}} = \text{softmax}(\mathbf{W}^T [\mathbf{q}_p, \mathbf{q}_m, \mathbf{q}_u, \mathbf{q}_i, b_m, b_u, b_i] + \mathbf{b}). \quad (14)$$

$\mathbf{W} \in \mathbb{R}^{(N_m+3D+3) \times N_m}$ and $\mathbf{b} \in \mathbb{R}^{N_m}$ are the parameters.

Finally, we obtain the prediction $\hat{\mathbf{y}} \in \mathbb{R}^{N_m}$, which represents the ensemble weights for the snapshot models considering not only the predictions from all snapshot models but also global attributes of all snapshot models and local information from the neighborhoods of users and items.

4.1.4 Multi-Layer Memory Networks. We now illustrate how to extend the proposed method to incorporate multiple memory layers in Figure 4. Unlike the single hop memory network in Figure 3, multi-layer memory modules can utilize the neighborhood representation of users and items in previous layers to derive the representation of snapshot models in the current layer and vice versa. In other words, the model has the chance to refine snapshot model representation by looking back and reconsidering the neighborhood information and vice versa. More specifically, multi-layer memory modules are stacked together in which the output of the h^{th} hop is used as the input to the $(h+1)^{\text{th}}$ hop. Each hop h (except for the last hop) queries the internal model memory, user memory and item memory via the attention mechanism and combines the predictions from all snapshot models to derive

the attention weight vector \mathbf{p}_{ui}^h , in which each dimension p_{uis}^h is the importance of snapshot model s in hop h :

$$\mathbf{p}_{ui}^h = \text{softmax}(\mathbf{W}_h^T [\mathbf{q}_m, \mathbf{q}_u, \mathbf{q}_i, \mathbf{q}_p, b_m, b_u, b_i] + \mathbf{b}_h). \quad (15)$$

Here, \mathbf{W}_h is a square weight matrix mapping the concatenated vector to the latent space. Next, we can construct the output of hop h as follows:

$$\mathbf{o}_{ui}^h = \sum_{s \in N(m)} p_{uis}^h c_s. \quad (16)$$

Here, \mathbf{o}_{ui}^h is the representation vector of snapshot models after considering the neighborhood representation propagated from previous hops, which is then transferred to the next hop as follows:

$$\mathbf{z}_{ui}^h = \phi(\mathbf{W}^h \mathbf{z}_{ui}^{h-1} + \mathbf{o}_{ui}^h + \mathbf{b}^h), \quad (17)$$

where the projection matrix $\mathbf{W}^h \in \mathbb{R}^{D \times D}$ and the bias vector $\mathbf{b}^h \in \mathbb{R}^D$ are parameters. ϕ denotes the activation function and \mathbf{z}_{ui}^{h-1} represents the input query of hop h . The initial query $\mathbf{z}_{ui}^0 = \mathbf{e}_{ui}$. In the last hop, the fully connected layer produces the final prediction $\hat{\mathbf{y}} \in \mathbb{R}^{N_m}$ as ensemble weights, which is the same as the single hop memory network.

4.2 Loss Function with Soft Labels

Our goal is to combine snapshot models not to select the optimal one because the optimal snapshot model is risky to infer due to the limited ratings of individual users/items, so that combining snapshot models whose training epochs are close to optimum can help to reduce ensemble variance. Therefore, we formulate the ensemble learning task as a multi-class classification problem. More specifically, we propose a soft labelling method to assign larger weights for snapshot models whose training epochs are close to optimum. The soft label for snapshot model s is defined as follows:

$$y_s = \frac{\exp(x_s)}{\sum_{k \in N(m)} \exp(x_k)}, \quad \forall s \in N(m). \quad (18)$$

Here, $x_s = (|e_s - e_o| + 1)^{-\alpha} \in (0, 1]$ is the transformed epoch distance between the s^{th} snapshot model (e_s) and the optimal one (e_o). And α is a hyperparameter, which is chosen as 1 based on empirical studies.

To measure the difference between the predicted ensemble weights distribution $\hat{\mathbf{y}}$ and the targeted distribution \mathbf{y} defined above, we adopt the KL-divergence described as follows:

$$L = \sum_{s=1}^{N_m} y_s (\log y_s - \log \hat{y}_s). \quad (19)$$

This is also the loss function to train the ensemble network, which can be effectively solved via SGD. After obtaining the ensemble weights for each user-item pair (from the output layer of the proposed ensemble network), we can use a simple weighted average to obtain the recommendation scores as follows:

$$\tilde{r}_{ui} = \hat{\mathbf{y}} \cdot \hat{\mathbf{r}}_{ui}. \quad (20)$$

$\hat{\mathbf{y}}$ is the predicted ensemble weights for the targeted user-item pair, and $\hat{\mathbf{r}}_{ui}$ contains the predicted scores on the targeted user-item pair from all snapshot models.

5 DISCUSSION

5.1 Efficiency Analysis

The computational overhead of training the NeuSE models is small compared to training many collaborative filtering models, *e.g.*, NeuMF [9]. In each epoch, the computational complexity is $O(DHN_R)$, where D is the embedding size of the memory networks, H is the number of hops of the memory networks and N_R is the number of ratings. Due to the adoption of multiple memory network components, the proposed method requires higher memory space in model training compared with NeuMF. However, as shown in the experiments, although NeuSE

Table 1. Statistics of the datasets adopted in the experiments.

Dataset	#users	#items	#ratings	Density
MovieLens 100K ¹	943	1,682	100,000	6.30%
Book Crossing ²	2,946	17,384	272,679	0.53%
Amazon Musical Instruments ³	80,004	13,836	500,176	0.05%

may take more memory space compared with NeuMF, the overall training time of NeuSE is much lower than that of NeuMF due to much faster convergence. More empirical discussions about the efficiency of the proposed method can be found in the experiments.

5.2 Addressing Cold-start User/item

The cold-start problem is a well-known challenge in collaborative filtering [1]. The proposed method can handle cold-start user or cold-start item as follows. We can change the input layer from the simple one-hot vectors to the user/item rating vectors, then the proposed method can handle unseen users who have ratings on seen items (cold-start user problem) or unseen items who have ratings by seen users (cold-start item problem). But for unseen users who only rate unseen items, we may need to either re-train the models or rely on auxiliary information (e.g., user demographics or item genres) as other CF algorithms.

5.3 Application Scope

Existing collaborative filtering algorithms can be divided into two categories: (1) non-learning based methods (e.g., most popular and KNN [10, 30]) and learning based methods (e.g., factorization methods [16, 26, 29] and neural network methods [9, 31]). We now analyze the feasibility of the proposed NeuSE method on top of each kind of methods.

In non-learning based methods, if we cannot create snapshot models, e.g., in most popular-based recommendation, the proposed NeuSE method will fail to improve the performance. However, if we can create snapshot models, e.g., by choosing different number of neighbors in KNN methods, NeuSE can successfully improve the performance as demonstrated in the experiments (Section 6).

In learning based methods, we can analyze the feasibility of NeuSE according to the optimization problems of the collaborative filtering algorithms. In most learning based methods, the model learning problems are formulated as non-convex optimization problems, e.g., as defined in Equation 1, in which iterative methods, e.g., SGD, are popular ways to learn the models. In this case, we can easily obtain snapshot models by choosing the intermediate models during learning. However, if the model learning problems are formulated as convex optimization problems [3, 34], the learned model can be obtained by a closed form solution. In this case, we cannot apply NeuSE to improve their performances because we have no snapshot models. However, convex optimization is rare in collaborative filtering algorithms, so the we can claim that NeuSE can be applied to most of the learning-based collaborative filtering algorithms.

In summary, NeuSE can be applied to improve the model performance as long as we can create snapshot models. As demonstrated in the experiments, snapshot models can be created in most of the popular collaborative filtering algorithms, so that we believe NeuSE is desirable in collaborative filtering tasks.

6 EXPERIMENTS

6.1 Experimental Setup

6.1.1 Dataset Description. We adopt three popular real-world datasets to evaluate the performance of the proposed snapshot ensemble method, the details of which are listed in Table 1. All three datasets can be regarded as both explicit feedback (user ratings on items) and implicit feedback (ratings as positive feedbacks and others as unknown feedbacks). The rating scale is 1-5 in the MovieLens and Amazon datasets and 1-10 in Book Crossing dataset.

6.1.2 Evaluation Protocol. For all experiments, we use the leave-one-out evaluation protocol [9, 27]. We split the datasets into training, validation and test sets by chronological order, in which the most recent rating of each user is used in test, the penultimate rating is used in validation and all the rest ratings are used in training. Then, the best performing models on the validation set are selected for test. For top-N recommendation, we randomly select 99 negative examples for each test example [9]. The negative examples are then fixed in the same experiment for fair comparison.

6.1.3 Compared Methods. We compare the proposed snapshot ensemble method with the following methods:

- **Single:** the single snapshot model with best validation performance is used for prediction;
- **Average:** the horizontal averaging method [41], which averages the outputs of snapshot models for prediction;
- **HSE:** the horizontal stacked ensemble method [41], which takes the outputs of snapshot models as the inputs to another meta model. We use an MLP as the meta learner for HSE in the experiments due to the universal approximation ability of deep neural networks.
- **SE:** the cyclic learning rate-based snapshot ensemble method [12], which is a state-of-the-art snapshot ensemble method but only applicable in models learned via SGD.

6.1.4 Baseline CF Methods. For each snapshot ensemble method, we apply them on the following popular collaborative filtering algorithms: 1) **UserKNN** [10]; 2) **ItemKNN** [30]; 3) **RSVD** [25]; 4) **FM** [26]; 5) **AutoRec** [31]; 6) **NeuMF** [9]. The above baselines cover many kinds of existing CF methods including three algorithm types (memory-based, factorization, and neural network methods) and two tasks (rating prediction and top-N recommendation).

6.1.5 Hyperparameters. In training the proposed NeuSE network, we use the batch size of 128, initial learning rate of 0.01, and Adam [14] as the optimizer to adaptively tune the learning rate. We apply dropout [33] in NeuSE with a dropout rate of 0.5. We use ReLU as the activation function for NeuSE due to better empirical performance in deep models [24]. The user/item embedding size of NeuSE is set to 16 for all the experiments. We randomly initialize the parameters of NeuSE with a Gaussian distribution $\mathcal{N}(0, 0.01)$. We use the two-hop memory network architecture in the experiments unless otherwise specified.

The hyperparameters of all baseline CF methods are kept the same for all snapshot ensemble methods as reported in their original papers, *i.e.*, the performances are only related to the snapshot ensemble methods. For RSVD, FM and AutoRec, we use the epoch variations $\Delta T = 10$ to select snapshot models and $\Delta T = 5$ for NeuMF till 200 epochs. For UserKNN and ItemKNN, we use cosine similarity to find nearest neighbors and use the number of nearest neighbors in $\{10, 20, 30, \dots, 100\}$ to create snapshot models. The embedding sizes or number of hidden factors are 8, 8, 500, 32 for RSVD, FM, AutoRec and NeuMF, respectively. Note that NeuSE can also improve the

¹<https://grouplens.org/datasets/movielens/100k/>

²<http://www2.informatik.uni-freiburg.de/~ctiegl/BX/>

³<http://jmcauley.ucsd.edu/data/amazon/>

performance of baseline CF algorithms with large embedding sizes as demonstrated in the sensitivity analysis. Other hyperparameters of the baseline CF algorithms, e.g., learning rates, weight decay, etc., are adopted from the best-performing ones in the original papers.

6.1.6 Evaluation Metrics. For rating prediction task, we evaluate the performances of different methods using root mean square error (RMSE), in which lower value indicates better performance:

$$\text{RMSE} = \sqrt{\frac{1}{|\Omega|} \sum_{u,i \in \Omega} (R_{u,i} - \hat{R}_{u,i})^2},$$

where $R_{u,i}$ is the rating of user u on item i and $\hat{R}_{u,i}$ is the corresponding predicted rating. Ω contains all the user-item pairs in the test set and $|\Omega|$ denotes the number of test examples.

For top-N recommendation task, we evaluate the performances of different methods using two popular metrics: Hit Ratio@N (HR@N) and NDCG@N, in which higher value indicates better performance:

$$\text{HR@N} = \frac{\#\text{Users w. Hits @N}}{\#\text{Users}}, \quad (21)$$

$$\text{NDCG@N} = \frac{\text{DCG@N}}{\text{IDCG@N}}. \quad (22)$$

Here, $\text{DCG@N} = \sum_{i=1}^N \frac{2^{r_i} - 1}{\log_2(i+1)}$, where r_i is 1 if the i -th item in the top N list is relevant to the targeted user and 0 otherwise. IDCG@N is the value of DCG@N with perfect ranking.

6.2 Recommendation Accuracy

Table 2 compares the accuracy between NeuSE and three snapshot ensemble methods on the MovieLens, Book Crossing and Amazon datasets. As shown in the table, the accuracy gain of NeuSE varies across different baseline CF methods and datasets. The highest improvements are $\sim 7.5\%$ over single model at the NDCG of UserKNN in the Amazon dataset, $\sim 15.4\%$ over Average method at the NDCG of ItemKNN in the Book Crossing dataset, $\sim 15.9\%$ over the HSE method at the NDCG of UserKNN in the Amazon dataset and $\sim 4.1\%$ over the SE method at the NDCG of NeuMF in the Book Crossing dataset. Although the improvements are not so significant everywhere, we can see that the improvements are universal, *i.e.*, over all baseline algorithms on all datasets.

We have the following observations from the results:

- Simply averaging the snapshot models cannot always improve accuracy over the optimal snapshot model, in which simple average increases the accuracy in three out of six baseline algorithms. The main reason is that appropriate subset of models could be better than all of them [43], but simple average takes all snapshot models with equal weights and thus cannot alleviate the effects of poor snapshot models.
- The HSE method, which uses an MLP to combine the snapshot models, achieves better performance than simple average in most of cases. This further confirms that dynamic ensemble (by varying the ensemble weights for different examples) could be more desirable in collaborative filtering tasks. However, the HSE method cannot memorize the user- and item-specific characteristics, e.g., some users/items may converge early as shown in the case studies, and thus cannot achieve optimal accuracy.
- The SE method using cyclic learning rate to obtain snapshot models achieves comparable accuracy to simple average on factorization methods (*i.e.*, RSVD and FM) and slightly better accuracy on neural network models (*i.e.*, AutoRec and NeuMF). The main reason is that neural network models trained with cyclic learning rate may converge to different local optimums [12] but factorization models may converge to the same local optimum if using cyclic learning rate.
- The proposed NeuSE method can achieve higher accuracy than all other three kinds of snapshot ensemble methods, because a) NeuSE adaptively weighs different snapshot models for individual ratings so that the effects of poor models can be alleviated from the ensemble; b) different snapshot models can converge to

Table 2. Accuracy comparison between NeuSE and four snapshot ensemble methods (Single, Average [41], HSE [41] and SE [12]) on the MovieLens 100K, Book Crossing and Amazon Musical Instruments datasets. Note that the SE method can only be applied to models learned via SGD and thus cannot be applied to UserKNN and ItemKNN. For the AutoRec method, we report the results of item-based AutoRec (I-AutoRec), which showed better performance than the U-AutoRec method in their paper. N is set to 20 for HR@ N and NDCG@ N in the results.

	Method	MovieLens 100K			Book Crossing			Amazon		
		RMSE	HR	NDCG	RMSE	HR	NDCG	RMSE	HR	NDCG
UserKNN	Single	1.2727	0.2365	0.0822	2.2056	0.2423	0.0844	1.3126	0.1548	0.0481
	Average	1.2759	0.2354	0.0801	2.2156	0.2410	0.0841	1.3118	0.1538	0.0473
	HSE	1.2704	0.2384	0.0827	2.2124	0.2415	0.0865	1.3142	0.1572	0.0446
	NeuSE	1.2615	0.2398	0.0857	2.1847	0.2562	0.0897	1.3046	0.1597	0.0517
ItemKNN	Single	1.2755	0.2238	0.0811	2.2152	0.2289	0.0791	1.3061	0.2257	0.0851
	Average	1.2753	0.2216	0.0802	2.2156	0.2241	0.0725	1.3085	0.2259	0.0843
	HSE	1.2682	0.2269	0.0812	2.2065	0.2292	0.0738	1.2957	0.2234	0.0845
	NeuSE	1.2621	0.2306	0.0862	2.1828	0.2395	0.0837	1.2879	0.2343	0.0886
RSVD	Single	1.1305	-	-	2.2849	-	-	1.2851	-	-
	Average	1.1171	-	-	2.2762	-	-	1.2831	-	-
	HSE	1.1327	-	-	2.2851	-	-	1.2846	-	-
	SE	1.1130	-	-	2.2758	-	-	1.2818	-	-
	NeuSE	1.0973	-	-	2.2686	-	-	1.2689	-	-
FM (SGD)	Single	1.1449	0.6278	0.2779	2.2351	0.3741	0.1403	1.1136	0.4759	0.1762
	Average	1.0948	0.6215	0.2631	2.1980	0.3688	0.1396	1.1052	0.4758	0.1658
	HSE	1.1047	0.6324	0.2791	2.2031	0.3764	0.1438	1.1037	0.4771	0.1784
	SE	1.0942	0.6307	0.2759	2.1975	0.3798	0.1457	1.1028	0.4789	0.1777
	NeuSE	1.0928	0.6382	0.2893	2.1938	0.3826	0.1497	1.0970	0.4844	0.1813
FM (MCMC)	Single	1.4138	0.5882	0.2631	2.5923	0.3672	0.1285	1.5391	0.4577	0.1525
	Average	1.4157	0.5895	0.2674	2.5892	0.3668	0.1273	1.5348	0.4542	0.1487
	HSE	1.4154	0.5841	0.2609	2.5854	0.3654	0.1274	1.5301	0.4594	0.1549
	NeuSE	1.3812	0.5977	0.2704	2.4647	0.3713	0.1324	1.4954	0.4643	0.1586
I-AutoRec	Single	1.1078	0.6776	0.2983	1.6479	0.3859	0.1427	0.9552	0.5127	0.2359
	Average	1.1164	0.6752	0.2985	1.6482	0.3814	0.1385	0.9517	0.5115	0.2381
	HSE	1.1029	0.6739	0.2901	1.6483	0.3847	0.1417	0.9537	0.5163	0.2379
	SE	1.1084	0.6734	0.2939	1.6375	0.3879	0.1462	0.9586	0.5129	0.2367
	NeuSE	1.0857	0.6845	0.3017	1.5849	0.3924	0.1489	0.9348	0.5231	0.2427
NeuMF	Single	-	0.8134	0.4087	-	0.4423	0.1948	-	0.6750	0.4083
	Average	-	0.8081	0.4097	-	0.4532	0.1943	-	0.6752	0.4085
	HSE	-	0.8092	0.4075	-	0.4527	0.1963	-	0.6784	0.4098
	SE	-	0.8171	0.4121	-	0.4535	0.1952	-	0.6755	0.4083
	NeuSE	-	0.8297	0.4153	-	0.4679	0.2031	-	0.6847	0.4124

different local optimums due to the nature of iterative optimization; c) the proposed ensemble method can fully utilize the global importance of snapshot models, the user and item specific characteristics and the neighborhood information from similar users and items via the memory networks to achieve better performance.

- We can also observe that the proposed NeuSE method can improve the accuracy on models trained with both iterative optimization methods (e.g., SGD) and Bayesian approaches (e.g., MCMC). In addition, the proposed method can also improve the accuracy on memory-based methods (e.g., UserKNN and ItemKNN),




Fig. 5. Sensitivity of NeuSE with the number of memory network layers and labels on three CF algorithms: (a) RSVD, (b) FM and (c) NeuMF. We use ReLU as the activation function and 16 as embedding size in this experiment.




Fig. 6. Sensitivity of NeuSE with CF models of different complexity on three CF algorithms: (a) RSVD, (b) FM and (c) NeuMF. We use ReLU as the activation function in this experiment. The x -axis denotes the complexity (number of latent factors and embedding size) of the CF models in each algorithm.

in which no learning procedure is involved. This indicates that NeuSE can be applied to almost all kinds of existing collaborative filtering algorithms to boost their accuracy.

6.3 Sensitivity Analysis

We first evaluate the sensitivity of NeuSE with different memory layers and labels and then analyze its efficiency. Note that we only report the results on three CF methods: RSVD, FM and NeuMF on the MovieLens 100K dataset because other methods and datasets show similar trends and thus are omitted.

6.3.1 Memory Network Layers. This experiment analyzes how the accuracy of NeuSE changes with the number of memory layers. As shown in Figure 5, for all three CF methods, two-layer memory networks achieve the highest accuracy than one-layer and three-layer memory networks. The main reason is that one-layer memory network cannot leverage the global/neighborhood information in previous layers to infer better representations in current layer and three-layer memory network is prone to overfitting due to more parameters.

6.3.2 Soft vs. Hard Label. This experiment analyzes how the proposed soft labelling technique affects the accuracy of NeuSE. As shown in Figure 5, NeuSE learned with soft labels achieves higher accuracy over hard labels, $\sim 0.5\%$ improvement in RSVD, $\sim 0.2\%$ improvement in FM, and $\sim 1.4\%$ improvement in NeuMF. This confirms that the proposed soft-labelling method can indeed improve the performance of snapshot ensemble due to the ability of adapting the losses on snapshot models with different performances. On the contrary, hard labels are prone to overfitting due to the noisy behaviors of individual users/items.




Fig. 7. Sensitivity of NeuSE with different user/item embedding sizes (16, 32, 64) on three CF algorithms: (a) RSVD, (b) FM and (c) NeuMF. We use ReLU as the activation function in this experiment.




Fig. 8. Sensitivity of NeuSE with different activation functions (Sigmoid, Tanh, ReLU) on three CF algorithms: (a) RSVD, (b) FM and (c) NeuMF. We use the embedding size of 16 in this experiment. On average, ReLU achieves the best performance among all three activation functions.

6.3.3 Model Complexity. Due to the possible concern that NeuSE can only improve the performance of baseline CF algorithms with low capacity, we conduct experiments to analyze the performance of NeuSE in CF models with different model complexity, e.g., different user/item embedding sizes in NeuMF. As shown in Figure 6, NeuSE can consistently improve the performance over the best Single models in all the cases with different complexities of the underlying CF models. This experiment indicates that (1) the performance of CF algorithms will vary with model capacity but the suboptimal performance issue exists in both high capacity models and low capacity models and (2) NeuSE can consistently improve the performance as long as the suboptimal performance issue exists.

6.3.4 Embedding Size. This experiment analyzes how sensitive are NeuSE models to different sizes of user/item embedding. As shown in Figure 7, the accuracy of NeuSE peaks with the embedding size of 16 and the sign of overfitting appears when embedding size grows to 32 and 64. This indicates that NeuSE can achieve decent performance with very light-weight models due to the adoption of multi-layer memory networks. Small embedding size also indicates that training NeuSE models can enjoy fast computation and high memory efficiency, i.e., NeuSE is less likely to suffer from the efficiency bottlenecks compared with many existing ensemble collaborative filtering methods.

6.3.5 Activation Function. This experiment analyzes how sensitive is NeuSE to different activation functions in the hidden layers. As shown in Figure 8, all three activation functions achieve similar performance if we only use one hop of memory network. When we increase the memory networks to 2-hop or 3-hop, Tanh achieves the




Fig. 9. Sensitivity of NeuSE with different number of snapshot models over the RSVD method. The snapshot models are selected from the intermediate models trained every 10 epochs. We can see that selecting 7-9 snapshot models can achieve slightly better performance in this experiment.




Fig. 10. Efficiency comparison between training NeuSE models and training baseline CF models. We use ReLU as the activation function and 16 as embedding size in this experiment. The training time of NeuSE is comparable to training RSVD models and FM models and much lower than training NeuMF models.




Fig. 11. Efficiency comparison between NeuSE and other three snapshot ensemble methods. We use ReLU as the activation function and 16 as embedding size in this experiment. The efficiency of NeuSE is comparable to HSE and SE methods. The simple average method is much more efficient than other methods.

worst performance among all three activation functions. The performance of ReLU is slightly better than Sigmoid and Tanh on average in all three algorithms. Compared with Sigmoid and Tanh, ReLU can reduce the likelihood of gradient vanishing and achieve better convergence in practice [17], which could be the main reasons of the better performance of ReLU.

6.3.6 Snapshot Selection. Figure 9 shows the performance of NeuSE with different number of snapshot models over the RSVD method on three datasets. The snapshot models are selected from the intermediate models trained every 10 epochs. As we can see from the results, the performance variations are not very significant when we change the number of snapshot models from 5 to 10, which indicates that NeuSE can achieve decent performance even without optimal snapshot model selection. In addition, selecting 7-9 snapshot models can achieve slightly better performance than others in this experiment. This indicates that selecting too many snapshot models, *e.g.*, 10, may hurt the accuracy, because some poor performance models may be selected. This conclusion is consistent with the observations from other ensemble learning works, in which researchers found that the ensemble performance may be hurt if the performance of the base models are too poor [5, 18, 43]. However, thanks to the strong ensemble selection capability of NeuSE, the ensemble performance does not degrade much even we select 10 snapshot models.

6.3.7 Efficiency Analysis. Here, we first analyze the efficiency of applying NeuSE on three CF methods: RSVD, FM and NeuMF. As shown in Figure 10, the efficiency of NeuSE is relevant to the baseline CF methods due to different learning trajectories of them but the computation time of NeuSE is comparable (RSVD, FM) or smaller (NeuMF) compared with the three CF methods. Therefore, we can conclude that NeuSE will not significantly increase the training time of CF methods compared with existing ensemble CF methods which train tens of base models [5, 19, 23].

We also compare the additional computation overhead of NeuSE with other snapshot ensemble strategies, *i.e.*, average, HSE and SE. As shown in Figure 11, the simple average method requires very little additional computation overhead due to no model learning steps and all the other snapshot ensemble methods have comparable computation overhead. Although NeuSE takes more additional computation overhead with increasing

Table 3. Accuracy comparison between NeuSE and four snapshot ensemble methods (Single, Average [41], HSE [41] and SE [12]) on the Netflix Prize dataset. N is set to 20 for HR@N and NDCG@N in the results. Average and HSE achieve similar performance with Single in all three algorithms. SE achieves better performance than Single, Average and HSE methods, but NeuSE outperforms all the compared methods in all cases.

	Method	RMSE	HR	NDCG
RSVD	Single	1.1237	-	-
	Average	1.1238	-	-
	HSE	1.1227	-	-
	SE	1.1205	-	-
	NeuSE	1.1193	-	-
FM	Single	1.0737	0.6723	0.4830
	Average	1.0733	0.6718	0.4827
	HSE	1.0741	0.6725	0.4833
	SE	1.0721	0.6798	0.4879
	NeuSE	1.0703	0.6835	0.4918
NeuMF	Single	-	0.8962	0.6254
	Average	-	0.8957	0.6246
	HSE	-	0.8983	0.6271
	SE	-	0.9021	0.6320
	NeuSE	-	0.9073	0.6382

hops of memory networks, the computation time increases are linear. Therefore, we can conclude that NeuSE can achieve the best performance among all the snapshot ensemble methods without significantly increasing the computational overhead compared with HSE and SE.

6.4 Performance on Large Dataset

Next, we evaluate the performance of NeuSE on the Netflix Prize dataset to demonstrate its effectiveness on large datasets. The Netflix prize dataset consists of 100,480,507 ratings from 480,189 users given to 17,770 items. In this study, we compare both the accuracy and efficiency of NeuSE with the other baseline methods when applied on top of RSVD [25], FM [26] and NeuMF [9]. All the evaluation details are consistent with those described in Section 6.1.

Table 3 compares the accuracy of NeuSE with other snapshot ensemble methods (Average, HSE and SE) on the Netflix Prize dataset on top of three popular CF algorithms (RSVD, FM (SGD) and NeuMF). As we can see from the results, NeuSE again outperforms all the other snapshot ensemble methods in all settings, which is consistent with the observations on relatively smaller datasets, e.g., MovieLens and Amazon. This confirms that NeuSE can boost the performance of collaborative filtering algorithms on both small datasets and large datasets.

Figure 12 compares the training time of NeuSE with other snapshot ensemble methods on the Netflix Prize dataset. As we can see from the results, the simple Average method is much more efficient than the other methods because no learning steps are involved. Similar to the comparison on the MovieLens datasets in Figure 11, the training overhead of NeuSE is on par with training a baseline CF model, which is also consistent with the observations on smaller datasets in Figure 11. Therefore, we can conclude that the training time of NeuSE is comparable to popular CF algorithms (e.g., RSVD [25] and NeuMF [9]) on both small datasets and large datasets.

In summary, we can conclude from the experimental results that NeuSE achieves decent performance (accuracy and efficiency) on both small datasets and large datasets.




Fig. 12. Efficiency comparison between NeuSE and other snapshot ensemble methods on the Netflix Prize dataset.

7 CONCLUSION AND FUTURE WORK

This paper proposes NeuSE – a neural snapshot ensemble method for collaborative filtering algorithms. NeuSE integrates the snapshot models during global model learning via multi-memory networks, which only introduces a small amount of extra computational overhead but can extensively and significantly improve the performance of almost all kinds of existing collaborative filtering algorithms. One possible extension of this work is to apply the idea of NeuSE to other domains beyond collaborative filtering, e.g., non convex neural networks [6].

ACKNOWLEDGEMENT

This work is supported by National Natural Science Foundation of China (Project Nos. 61977013, 62090025) and Sichuan Science and Technology Program (Project No. 2019YJ0164).

REFERENCES

- [1] G. Adomavicius and A. Tuzhilin. 2005. Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering* 17, 6 (June 2005), 734–749. <https://doi.org/10.1109/TKDE.2005.99>
- [2] Charu C Aggarwal et al. 2016. *Recommender systems*. Springer.
- [3] Fabio Aioli. 2014. Convex AUC optimization for top-N recommendation with implicit feedback. In *Proceedings of the 8th ACM Conference on Recommender Systems*. 293–296.
- [4] Alex Beutel, Ed H. Chi, Zhiyuan Cheng, Hubert Pham, and John Anderson. 2017. Beyond Globally Optimal: Focused Learning for Improved Recommendations. In *Proceedings of the 26th International Conference on World Wide Web*. 203–212.
- [5] Chao Chen, Dongsheng Li, Yingying Zhao, Qin Lv, and Li Shang. 2015. WEMAREC: Accurate and Scalable Recommendation Through Weighted and Ensemble Matrix Approximation. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '15)*. 303–312.
- [6] Anna Choromanska, Mikael Henaff, Michael Mathieu, Gérard Ben Arous, and Yann LeCun. 2015. The loss surfaces of multilayer networks. In *Artificial Intelligence and Statistics*. 192–204.
- [7] John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *Journal of Machine Learning Research* 12, Jul (2011), 2121–2159.
- [8] Rong Ge, Chi Jin, and Yi Zheng. 2017. No Spurious Local Minima in Nonconvex Low Rank Problems: A Unified Geometric Analysis. In *Proceedings of the 34th International Conference on Machine Learning (ICML '17)*. 1233–1242.
- [9] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *Proceedings of the 26th International Conference on World Wide Web (WWW '17)*. 173–182.
- [10] Jonathan L. Herlocker, Joseph A. Konstan, Al Borchers, and John Riedl. 1999. An Algorithmic Framework for Performing Collaborative Filtering. In *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '99)*. ACM, 230–237.

- [11] Y. Hu, Y. Koren, and C. Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *The 2008 Eighth IEEE International Conference on Data Mining (ICDM '08)*. 263–272.
- [12] Gao Huang, Yixuan Li, Geoff Pleiss, Zhuang Liu, John E Hopcroft, and Kilian Q Weinberger. 2017. Snapshot ensembles: Train 1, get m for free. *arXiv preprint arXiv:1704.00109* (2017).
- [13] Nour El Islem Karabadjji, Samia Beldjoudi, Hassina Seridi, Sabeur Aridhi, and Wajdi Dhifli. 2018. Improving memory-based user collaborative filtering with evolutionary multi-objective optimization. *Expert Systems with Applications* 98 (2018), 153 – 165.
- [14] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [15] Noam Koenigstein, Gideon Dror, and Yehuda Koren. 2011. Yahoo! Music Recommendations: Modeling Music Ratings with Temporal Dynamics and Item Taxonomy. In *Proceedings of the Fifth ACM Conference on Recommender Systems (RecSys '11)*. ACM, 165–172.
- [16] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 42 (2009), 30–37. Issue 8.
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.
- [18] Anders Krogh and Jesper Vedelsby. 1995. Neural network ensembles, cross validation, and active learning. In *Advances in neural information processing systems*. 231–238.
- [19] Joonseok Lee, Seungyeon Kim, Guy Lebanon, and Yoram Singer. 2013. Local low-rank matrix approximation. In *Proceedings of The 30th International Conference on Machine Learning*. 82–90.
- [20] Dongsheng Li, Chao Chen, Wei Liu, Tun Lu, Ning Gu, and Stephen Chu. 2017. Mixture-Rank Matrix Approximation for Collaborative Filtering. In *Advances in Neural Information Processing Systems 30*. 477–485.
- [21] Dongsheng Li, Chao Chen, Qin Lv, Hansu Gu, Tun Lu, Li Shang, Ning Gu, and Stephen M. Chu. 2018. *AdaError: An Adaptive Learning Rate Method for Matrix Approximation-Based Collaborative Filtering*. International World Wide Web Conferences Steering Committee, 741–751.
- [22] Dongsheng Li, Chao Chen, Qin Lv, Li Shang, Stephen M Chu, and Hongyuan Zha. 2017. ERMMA: expected risk minimization for matrix approximation-based recommender systems. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*. 1403–1409.
- [23] Zhiyu Min and Dahua Lin. 2018. Probabilistic Ensemble of Collaborative Filters. In *Thirty-Second AAAI Conference on Artificial Intelligence*. 2371–2378.
- [24] Vinod Nair and Geoffrey E. Hinton. 2010. Rectified Linear Units Improve Restricted Boltzmann Machines. In *Proceedings of the 27th International Conference on International Conference on Machine Learning (ICML '10)*. 807–814.
- [25] Arkadiusz Paterek. 2007. Improving regularized singular value decomposition for collaborative filtering. In *KDD CUP'07*. 5–8.
- [26] Steffen Rendle. 2010. Factorization machines. In *The 2010 IEEE International Conference on Data Mining*. IEEE, 995–1000.
- [27] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (UAI '09)*. 452–461.
- [28] Jasson DM Rennie and Nathan Srebro. 2005. Fast maximum margin matrix factorization for collaborative prediction. In *Proceedings of the 22nd international conference on Machine learning*. ACM, 713–719.
- [29] Ruslan Salakhutdinov and Andriy Mnih. 2008. Bayesian Probabilistic Matrix Factorization Using Markov Chain Monte Carlo. In *Proceedings of the 25th International Conference on Machine Learning (ICML '08)*. 880–887.
- [30] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based Collaborative Filtering Recommendation Algorithms. In *Proceedings of the 10th International Conference on World Wide Web (WWW '01)*. ACM, 285–295.
- [31] Suvash Sedhain, Aditya Krishna Menon, Scott Sanner, and Lexing Xie. 2015. AutoRec: Autoencoders Meet Collaborative Filtering. In *Proceedings of the 24th International Conference on World Wide Web (WWW '15 Companion)*. ACM, 111–112.
- [32] Samuel L Smith, Pieter-Jan Kindermans, Chris Ying, and Quoc V Le. 2017. Don't decay the learning rate, increase the batch size. *arXiv preprint arXiv:1711.00489* (2017).
- [33] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research* 15 (2014), 1929–1958.
- [34] Harald Steck. 2019. Embarrassingly shallow autoencoders for sparse data. In *The World Wide Web Conference*. 3251–3257.
- [35] A. Swann and N. Allinson. 1998. Fast committee learning: Preliminary results. *Electronics Letters* 34 (July 1998), 1408–1410. Issue 14.
- [36] Maksims Volkovs, Guangwei Yu, and Tomi Poutanen. 2017. DropoutNet: Addressing Cold Start in Recommender Systems. In *Advances in Neural Information Processing Systems 30*. 4957–4966.
- [37] Yong Wang, Jiangzhou Deng, Jerry Gao, and Pu Zhang. 2017. A hybrid user similarity model for collaborative filtering. *Information Sciences* 418–419 (2017), 102 – 118.
- [38] Jason Weston, Sumit Chopra, and Antoine Bordes. 2014. Memory networks. *arXiv preprint arXiv:1410.3916* (2014).
- [39] Ashia C Wilson, Rebecca Roelofs, Mitchell Stern, Nati Srebro, and Benjamin Recht. 2017. The marginal value of adaptive gradient methods in machine learning. In *Advances in Neural Information Processing Systems*. 4148–4158.
- [40] Teng Xiao, Shangsong Liang, Weizhou Shen, and Zaiqiao Meng. 2019. Bayesian Deep Collaborative Matrix Factorization. In *Thirty-Third AAAI Conference on Artificial Intelligence*. 5474–5481.

- [41] Jingjing Xie, Bing Xu, and Zhang Chuang. 2013. Horizontal and vertical ensemble with deep representation for classification. *arXiv preprint arXiv:1306.2759* (2013).
- [42] Y. Zhang, J. Wu, and H. Wang. 2019. Neural Binary Representation Learning for Large-Scale Collaborative Filtering. *IEEE Access* 7 (2019), 60752–60763.
- [43] Zhi-Hua Zhou, Jianxin Wu, and Wei Tang. 2002. Ensembling neural networks: Many could be better than all. *Artificial Intelligence* 137, 1 (2002), 239 – 263.