

Fast Adaptively Weighted Matrix Factorization for Recommendation with Implicit Feedback

Jiawei Chen,^{1,2} Can Wang,^{1,2,*} Sheng Zhou,¹ Qihao Shi,¹
Jingbang Chen,¹ Yan Feng,^{1,2} Chun Chen¹

¹College of Computer Science, Zhejiang University, China

²Zhejiang University-LianlianPay Joint Research Center

{sleepyhunt, wcan, zhousheng_zju, shiqihao321, chenjb, fengyan, chenc}@zju.edu.cn

Abstract

Recommendation from implicit feedback is a highly challenging task due to the lack of the reliable observed negative data. A popular and effective approach for implicit recommendation is to treat unobserved data as negative but downweight their confidence. Naturally, how to assign confidence weights and how to handle the large number of the unobserved data are two key problems for implicit recommendation models. However, existing methods either pursuit fast learning by manually assigning simple confidence weights, which lacks flexibility and may create empirical bias in evaluating user's preference; or adaptively infer personalized confidence weights but suffer from low efficiency.

To achieve both adaptive weights assignment and efficient model learning, we propose a fast adaptively weighted matrix factorization (FAWMF) based on variational auto-encoder. The personalized data confidence weights are adaptively assigned with a parameterized neural network (function) and the network can be inferred from the data. Further, to support fast and stable learning of FAWMF, a new specific batch-based learning algorithm fBGD has been developed, which trains on all feedback data but its complexity is linear to the number of observed data. Extensive experiments on real-world datasets demonstrate the superiority of the proposed FAWMF and its learning algorithm fBGD.

Introduction

Recommender systems play an important role in many internet services. Since in practise most of recommender systems only have the implicit feedback (e.g. items consumed), recently research attention is increasingly shifted from explicit feedback (e.g. rating prediction) to implicit feedback. However, learning a recommender system from implicit feedback is more challenging. In implicit feedback scenarios, only positive feedback are observed, and the unobserved user-item feedback data (e.g. a user has not bought an item yet) are a mixture of real negative feedback (i.e. a user does not like it) and missing values (i.e. a user just does not know it). Existing methods address this problem by treating all

the un-observed data as negative (dislike) but downweighting their confidence. Although it is reasonable, it poses two important questions: (1) How to assign confidence weights for each data? (2) How to handle the massive volume of the unobserved data efficiently?

While there is a rich literature on recommendation from implicit feedback, to our knowledge, all existing methods lack one or more desiderata. Most of these methods rely on the meticulous assignment of confidence weights to the data. For example, WMF (Hu, Koren, and Volinsky 2008) and eALS (He et al. 2016) give uniform or popularity-based confidence weights; BPR (Rendle et al. 2009), CDAE (Wu et al. 2016), or other sophisticated models downweight the contributions of the negative data in their heuristical (e.g. uniformly) negative sampling strategy. However, choosing these weights usually involves heuristic alterations to the data and needs expensive exhaustive grid search via cross-validation. Furthermore, it is unrealistic for researchers to manually set flexible and diverse weights for millions of data. In practical scenarios, the data confidence weights may change for various user-item combinations. Some unobserved data can be attributed to user's preference while others are the results of users' limited scopes. To this end, exposure-based matrix factorization (EXMF) (Liang et al. 2016) has been proposed to downweight negative data automatically by predicting how a user knows a item. However, EXMF involves iterative inference of the exposure for each data, which is computationally expensive and also potentially suffers from overfitting. More recently, SamWalker (Chen et al. 2019b) adaptively assign confidence weights from users' social contexts. However, the social network information is hard obtained in many recommender systems.

To address these problems, we propose a fast adaptively weighted matrix factorization (FAWMF) based on variational auto-encoder (Kingma and Welling 2013) for both adaptive weights assignment and efficient learning. We first analyze EXMF from variational perspective and find that the variational posterior of user's exposure acts as data confidence weights in learning users' preference with MF. It is consistent with our intuitions. Only if the user is exposed to the item, he can decide whether to consume the item based on his preference. The data with larger expo-

*Corresponding author: wcan@zju.edu.cn

sure are more reliable in deriving user’s preference. Further, we propose to replace confidence weights (variational posterior of user’s exposure) with a parameterized inference neural network (function), which reduces the number of the inferred parameters and is capable of capturing latent correlations between users’ exposure. In fact, the independent assumption of the exposure is not practical in real world. Typically, recent literatures (Palla et al. 2005; Zhou 2011) in social science claim that each of us belongs to some information-sharing communities. Naturally, users are exposed to these communities and thus their exposure exhibit correlations when they belong to common communities. Motivated by this point, a specific community-based inference neural network has been designed, which explores latent communities among users and infer users’ exposure from the communities that they belong to. It will capture more precise exposure and can mitigate the overfitting problem. By optimizing the variational lower bound, both the inference neural network and MF can be learned from the data.

Efficiently learning a recommendation model from implicit feedback is also challenging since it requires to account for the large-scale unobserved data. Although stochastic gradient descent optimizer (SGD) with sampling strategy can be employed to speed up the inference, it usually suffers from low convergence and high gradient instability. Especially in the implicit recommendation task, the optimizer usually samples the uninformative data, which have low confidence and make limited contributions on gradient update (Chen et al. 2019b). The final recommendation performance will suffer. Instead, we turn to employ batch gradient descent optimizer (BGD), which computes stable gradient from all feedback data and usually converges to a better optimum. Unfortunately, the naive implement of BGD will suffer from low efficiency caused by the expensive full-batch gradient computation over all data. To address this problem, we develop a fast BGD-based learning algorithm fBGD for our FAWMF. With rigorous mathematical reasoning, massive repeated computations of original BGD can be avoided and the learning process can be accelerated. Notably, despite fBGD computes gradient over all data, its actual complexity is linear with the number of the unobserved positive data. Due to the sparsity of the implicit feedback data, fBGD achieves a significant acceleration.

We summarize our key contributions as follows:

- We propose a FAWMF model based on variational auto-encoder to achieve both adaptive weights assignment and efficient model learning.
- A batch-based learning algorithm fBGD has been developed to learn FAWMF efficiently and stably.
- Extensive experimental evaluations on three well-known benchmark datasets demonstrate the superiority of our FAWMF over the existing implicit MF methods.

Related work

Recommendation from implicit feedback data. Most of the existing methods manually assign coarse-grained confidence weights. For example, the classic weighted factoriza-

tion matrix model (WMF) (Hu, Koren, and Volinsky 2008), eALS (He et al. 2016), ICD-MF (Bayer et al. 2017), used a simple heuristic where the confidence weights are assigned uniformly or based on item popularity; Logistical matrix factorization (Johnson 2014), BPR (Rendle et al. 2009), or neural-based recommendation models (e.g. CDAE (Wu et al. 2016), NCF(He et al. 2017)) downweight the contribution of negative data implicitly in their heuristic negative sampling strategy (e.g. uniformly). More recently, a new probabilistic model EXMF(Liang et al. 2016) was proposed to incorporate user’s exposure to items into the CF methods. When inferring user’s preference, user’s exposure can be translated as data confidence. However, this method suffers from the efficiency problem. Also, some sophisticated models have been proposed to learn confidence weights from users’ social contexts (Wang et al. 2018; Jiawei et al. 2018; Chen et al. 2019b). However, the social information is not available in many cases.

Efficient learning algorithm for implicit recommendation models. To handle the large-scale unobserved data, two types of strategies have been proposed for efficient learning: sample-based learning and whole-data based learning. The first type achieves fast learning with stochastic gradient descent (SGD) and negative sampling. The most popular sampling strategy is to draw un-observed feedback data uniformly, which is widely adopted in LMF(Johnson 2014), BPR(Rendle et al. 2009), CDAE (Wu et al. 2016), NCF (He et al. 2017), Mult-DAE (Liang et al. 2018) etc. Also, (Yu, Bilenko, and Lin 2017), (Chen et al. 2017) and (Hernández-Lobato, Houlshby, and Ghahramani 2014) further propose item popularity-based and item-user co-bias sampling strategy. However, in recommendation task SGD usually suffers from slow convergence and gradient instability when the number of items is large (Chen et al. 2019b).

Thus, some other dynamic sampling strategies are proposed to improve convergence and accuracy. Some literatures (Rendle and Freudenthaler 2014; Yu et al. 2018; Park and Chang 2019; Wang et al. 2017) propose to oversample the “difficult” negative instances which are hard to be discriminated by the models. However, they will expose efficiency issues in sampling. Also, the stochastic gradient estimator is biased and may amplify the natural noise in users feedback data (Li et al. 2018). More recently, SamWalker (Chen et al. 2019b) conduct the random walk along the social network to select informative instances. Although effectively, it has two weakness: (1) SamWalker requires additional social information but it is not available in many recommender systems. (2) SamWalker still employs uniformly sampling strategy to update its dynamic sampler, leading to insufficient training of it.

A more effective and stable way is updating the model from the whole-data, but it face efficiency challenge due to the large number of the negative data. Thus, memorization strategies (e.g. ALS, eALS) has been proposed (Hu, Koren, and Volinsky 2008; He et al. 2016; Bayer et al. 2017; Chen et al. 2019a) to speed up learning. However, these algorithms are just suitable for the MF with simple manual confidence weights, which lacks flexible and may create empirical bias. The recommendation performance will suffer.

Adaptive weights assignment and efficient learning are both important in recommendation from implicit feedback. Existing methods are not able to provide both of these, which motivates the approach described in this paper.

Preliminaries

In this section, we first give the problem definition of implicit recommendation. Then, we introduce exposure-based matrix factorization (EXMF) (Liang et al. 2016) framework from variational perspective to provide usual insight about the relation between user’s exposure and data confidence.

Problem definition

Suppose we have a recommender system with user set U (including n users) and item set I (including m items). The implicit feedback data is represented as $n \times m$ matrix \mathbf{X} with entries x_{ij} denoting whether or not the user i has consumed the item j . The task of a recommender system can be stated as follow: recommending items for each user that are most likely to be consumed by him.

Exposure-based matrix factorization (EXMF)

Note that unobserved feedback data $\{x_{ij} \in \mathbf{X} : x_{ij} = 0\}$ contain the real negative data (dislike) and the missing values (unknown). EXMF introduces a Bernoulli variable a_{ij} to model users’ exposure: $a_{ij} = 1$ denotes that the user i knows the item j and $a_{ij} = 0$ denotes not. Then, EXMF models user’s consumption x_{ij} based on a_{ij} as follow:

$$a_{ij} \sim \text{Bernoulli}(\eta_{ij}) \quad (1)$$

$$(x_{ij}|a_{ij} = 1) \sim N(\mathbf{u}_i^\top \mathbf{v}_j, \lambda_x) \quad (2)$$

$$(x_{ij}|a_{ij} = 0) \sim \delta_0 \approx N(\varepsilon, \lambda_x) \quad (3)$$

where δ_0 denotes $p(x_{ij} = 0|a_{ij} = 0) = 1$; η_{ij} is the prior probability of exposure. Here we relax function δ_0 as $N(\varepsilon, \lambda_x)$ to make model more robust, where ε is a small constant (e.g. $\varepsilon=1e-5$). When $a_{ij} = 0$, we have $x_{ij} \approx 0$, since the user does not know the item and he can not consume it. When $a_{ij} = 1$, when the user has learned the item, he will decide whether or not to consume the item based on his preference. Thus, x_{ij} is modeled as the classic matrix factorization model, where \mathbf{u}_i denotes the K -dimensional latent factor (preference) of user i , and \mathbf{v}_j denotes the K -dimensional latent factor (attribute) of item j .

Analyses of EXMF from variational perspective

The marginal likelihood of EXMF is composed of a sum over each datapoint $\log p(\mathbf{X}) = \sum_{i \in U, j \in I} \log p(x_{ij})$, which can be rewritten as:

$$\begin{aligned} \log p(x_{ij}) &= E_q[\log p(x_{ij}, a_{ij}) - \log q(a_{ij}|x_{ij})] \\ &\quad + E_q[\log p(a_{ij}|x_{ij}) - \log q(a_{ij}|x_{ij})] \\ &= L(u, v, q; x_{ij}) + D_{KL}(q(a_{ij}|x_{ij})||p(a_{ij}|x_{ij})) \end{aligned} \quad (4)$$

where $q(a_{ij}|x_{ij})$ is defined as an approximated variational posterior of a_{ij} . Since the second term – KL-divergence – is non-negative, the first term $L(\theta, q; x_{ij})$ is the evidence

lower bound on (ELBO) the margin likelihood. Classic variational methods (Hoffman et al. 2013) usually employ conjugate variational distribution and individual variational parameters¹, i.e. $q(a_{ij}|x_{ij}) = \text{Bernoulli}(\gamma_{ij})$. Then, optimizing EXMF can be transferred to minimize the following objective function w.r.t $\mathbf{u}_i, \mathbf{v}_j, \gamma_{ij}$:

$$\begin{aligned} J(u, v, \gamma; \mathbf{X}) &= -\frac{2}{\lambda_x} L(u, v, q; x_{ij}) \\ &= \sum_{i \in U, j \in I} \gamma_{ij} (\mathbf{u}_i^\top \mathbf{v}_j - x_{ij})^2 + \sum_{i \in U, j \in I} (1 - \gamma_{ij}) (\varepsilon - x_{ij})^2 \\ &\quad - \frac{2}{\lambda_x} \sum_{i \in U, j \in I} D_{KL}(q(a_{ij}|x_{ij})||p(a_{ij})) \\ &= \sum_{i \in U, j \in I} \gamma_{ij} (\mathbf{u}_i^\top \mathbf{v}_j - x_{ij})^2 + \sum_{i \in U, j \in I} l(\gamma_{ij}) \end{aligned} \quad (5)$$

Exposure as data confidence. The objective function consists of three parts: (1) Weighted matrix factorization to learn user’s preference; (2) The loss when the data are predicted as unknown; (3) The regularization term – KL divergency between the prior and the variational posterior. A good property is observed that the parameters γ_{ij} , which characterize the probability that user i is exposed to item j , act as the confidence of the corresponding data in learning user’s preference. It is consistent with our intuitions. Only if the user is exposed to the item, he can decide whether to consume the item based on his preference. The data with larger exposure are more reliable in deriving user’s preference.

Inefficiency problem. EXMF suffers efficiency problems. The number of inferred variational parameters (confidence weights) γ_{ij} grows quickly with the number of users and items ($n \times m$), which easily scale to billion or even larger level. It will potentially suffer from overfitting and become the time and space bottleneck for practise recommender systems. Further, the inference of EXMF requires to sum over the terms for each data, which is time-consuming.

Fast adaptively weighted matrix factorization

To address these problems, we propose a fast adaptively weighted matrix factorization (FAWMF) based on variational auto-encoder to achieve both adaptive weight assignment and efficient model learning. To reduce the number of the inferred parameters, FAWMF replaces individual parameters γ_{ij} with a parameterized function $\gamma_{ij} = g_\Phi(i, j, \mathbf{X})$ that maps users’ consumption on items into their exposure with parameters Φ . It is more reasonable since there exists interactions between users. Independent assumption of exposure is not practical in real world. Typically, recent literatures (Palla et al. 2005) in social science claim that each of us belongs to some information-sharing communities. Users are exposed to these communities and thus their exposure exhibit commonality when they belong to similar communities. Motivated by this point, we summarize communities as collections of similar users and further infer users’ exposure

¹Note that the EM algorithm presented in (Liang et al. 2016) is a special case of the classic variational inference.




Figure 1: Illustration of the proposed FAWMF

from the communities that they belong to. Concretely, the inference function can be modeled as follow:

$$g_{\Phi}(i, j, \mathbf{X}) = \theta_i^T \sigma(w_j \sum_{k \in U} \theta_k \alpha_k x_{kj} + b_j) \quad (6)$$

where vector θ_i denotes the membership that allocates each user i to a fixed D number of communities and meet $\theta_i \succeq 0$, $|\theta_i|_1 = 1$. We cumulate the consumptions of the users in the community to infer how items are exposed to the community, where α_k captures the heterogenous roles of users. Different users may have different influence strength on the communities (Shi et al. 2019a; 2019b). Then, a logistic linear function with parameters w_j, b_j has been employed to map cumulated consumptions into the exposure. Intuitively, the more influence users in the common community have consumed the item, the user are more likely to know it. Finally, user’s exposure can be depicted by how the user belongs to the communities and how items are exposed to these communities. Overall, the inference of the data confidence can be translated into learning the parameters of inference function $g_{\Phi}(i, j, \mathbf{X})$ ($\Phi \equiv \{\theta, \alpha, w, b\}$), which captures latent correlations between γ_{ij} for better estimation and reduces the number of inferred parameters from $O(nm)$ to $O(D(n + m))$. Different from existing clustering-based methods (e.g. (Chen et al. 2015; Lee et al. 2016)) which cluster users into communities/submatrices based on users’ preference, our FAWMF deduces implicit communities based on users’ exposure.

Also, FAWMF can be well understood from neural network perspective. As shown in Figure 1, at first an outer product between each user’s consumption and his membership has been employed for interactions. The obtained tensor \mathbf{Z} can be regarded as the influence of user’s consumptions along different dimensions (communities). Then, we use a specific striped CNN (Krizhevsky, Sutskever, and Hinton 2012) layer and a linear layer to encode the influences (consumptions) into the “kernel maps” of users’ exposure. Here we choose striped CNN since the adjacent elements of \mathbf{Z} are just caused by the adjacent users(items) id and do not suggest they will have share more commonality. Further, the exposure for different user-item pairs can be depicted as the combination of the “kernel maps”. Finally, the exposure-based MF acts as a probabilistic decoder to predict users’ consumption. Overall, FAWMF forms an auto-encoder to learn both data confidence weights and latent factors of MF.




Figure 2: Illustration of how FAWMF mitigates over-fitting

Discussion

How does FAWMF mitigate overfitting? To better understand this effect, let us draw an analogy with the floating balls in the water, as illustrated in Figure 2. Learning EXMF model by optimizing equation (5) will give a force to push up these positive balls (instances) and push down these unobserved balls (instances). Thus, without strong priors, the confidence weights of the data easily achieve extremely values ($\gamma_{ij} \approx 1$ for the positive instances and $\gamma_{ij} \approx 0$ for the unobserved instances), where unobserved data makes few contributions on training. In this situation, all the instances will be predicted as positive by MF and the model will suffer from over-fitting. In our FAWMF, this over-fitting effect will be mitigated by introducing inference network of the γ_{ij} . There exists correlations between users’ exposure, which can be analogies with elastic links lines between the balls. Typically, users tend to have similar exposure (γ_{ij}), if they share more common communities. Naturally, with the training of the model, the unobserved data which have correlations with positive instances, will be pulled up due to the force from the lines. This way, when the model has well fitted the data, the positive and the unobserved instances will get stable at different depth in water. The unobserved instances which have stronger correlations with positive instances, usually reaches higher position than other unobserved instances.

Fast learning algorithm fBGD

Learning a recommendation model from implicit feedback is time-cost expensively due to the massive volume of the unobserved data. The popular solution addressing this problem is stochastic gradient descent (SGD) with negative sampling. However, SGD usually suffers from low convergency and gradient instability, which affects the performance of the model. Especially, in the recommendation task, the sampler usually select uninformative instances that have small confidence weights γ_{ij} and make limited contributions on update. Thus, we turn to use batch gradient descent optimizer (BGD), which computes stable gradient from all feedback data. Unfortunately, the original BGD will suffer from low efficiency due to the full-batch computation of the gradient over instances. To address this problem, we develop new specific learning algorithm fBGD for our FAWMF. We speed up learning of BGD by caching some specific intermediate variables to avoid the massive repeated computations. Here we detail the derivation of the gradient w.r.t α_k ; where

the counterpart of others $(\theta_i, w_j, b_j, \mathbf{u}_i, \mathbf{v}_j)$ is achieved likewise and presented in supplemental materials.

For better derivation, we let $\mathbf{q}_j = \sigma(w_j \sum_{k \in U} \theta_k \alpha_k x_{kj} + b_j)$. Also, we drop out the regularization term in the objective function (Eq. (5)) since the regularization term usually makes no contribution on recommendation accuracy of FAWMF and hinders fast training. It also can be regarded that we set a large value of parameter λ_x . Then, the gradient of objective function (Eq. (5)) w.r.t α_k for each user k can be derived as follow:

$$\frac{\partial J}{\partial \alpha_k} = \sum_{j \in I} \frac{\partial J}{\partial \mathbf{q}_j} \cdot \frac{\partial \mathbf{q}_j}{\partial \alpha_k} \quad (7)$$

$$\frac{\partial J}{\partial \mathbf{q}_j} = \sum_{i \in U} (\mathbf{u}_i^T \mathbf{v}_j \mathbf{u}_i^T \mathbf{v}_j - 2x_{ij}(\mathbf{u}_i^T \mathbf{v}_j - \varepsilon) - \varepsilon^2) \theta_i \quad (8)$$

$$\frac{\partial \mathbf{q}_j}{\partial \alpha_k} = \mathbf{q}_j \cdot (1 - \mathbf{q}_j) \cdot \theta_k w_j x_{kj} \quad (9)$$

Clearly, the computational cost lies in two parts: (1) the calculation of the gradients of J w.r.t \mathbf{q}_j for each item $j \in I$, which requires the summation over the terms of each user i ; (2) the summation of $\frac{\partial J}{\partial \mathbf{q}_j} \cdot \frac{\partial \mathbf{q}_j}{\partial \alpha_k}$ over each item $j \in I$ to get final result. In fact, the first part produces the main cost since it requires a traversal of all users and repeat this operation over all items; while the second part can be accelerated by just iterating over these positive instances ($x_{kj} = 1$). The overall computational complexity to update $\{\alpha_k : k \in U\}$ is $O(nmKD)$, which is generally infeasible since nm can easily reach billion level or even higher in practise.

To speed up the learning, we rewrite the computational bottlenecks – Eq. (8) – by isolating item-independent terms:

$$\begin{aligned} \frac{\partial J}{\partial \mathbf{q}_j} &= \sum_{k=1}^K \sum_{l=1}^K v_{jk} v_{jl} \sum_{i \in U} u_{ik} u_{il} \theta_i - \varepsilon^2 \sum_{i \in U} \theta_i \\ &\quad - 2 \sum_{i \in U} x_{ij} (\mathbf{u}_i^T \mathbf{v}_j - \varepsilon) \theta_i \end{aligned} \quad (10)$$

By this reformulation, we can see that the major computation – the $\sum_{i \in U} u_{ik} u_{il} \theta_i$ and $\sum_{i \in U} \theta_i$ over all users – is independent of item j . Thus, we could achieve a significant speed-up by caching these terms. That is, we cache $\mathbf{M}_{kl*}^{(q)} = \sum_{i \in U} u_{ik} u_{il} \theta_i$ for each $1 \leq k \leq K, 1 \leq l \leq K$ and $\mathbf{S}^{(q)} = \sum_{i \in U} \theta_i$, where $\mathbf{M}^{(q)}$ is a $K \times K \times D$ -dimensional tensor and $\mathbf{S}^{(q)}$ is a D -dimensional vector. Then, $\frac{\partial J}{\partial \mathbf{q}_j}$ can be calculated as follow:

$$\begin{aligned} \frac{\partial J}{\partial \mathbf{q}_j} &= \sum_{k=1}^K \sum_{l=1}^K v_{jk} v_{jl} \mathbf{M}_{kl*}^{(q)} - \varepsilon^2 \mathbf{S}^{(q)} \\ &\quad - 2 \sum_{i \in U} x_{ij} (\mathbf{u}_i^T \mathbf{v}_j - \varepsilon) \theta_i \end{aligned} \quad (11)$$

The rearrangement of nested sums is the key transformation that allows the fast optimization. The time computation can be reduced from $O(nmKD)$ to $O(K^2D(n+m))$

Table 1: Statistics of three datasets.

Datasets	#Users	#Items	#Observed positive feedback
Movielens	6,040	3,952	1,000,209
Amazon	10,619	37,762	256,287
Douban	123,480	20,029	16,624,937

Table 2: characteristics of the compared methods.

Methods	Adaptive weights?	Without sampling?	Complexity
WMF(ALS)	\	\	$O((n+m)K^3 + \mathbf{X}^+ K^2)$
eALS	\	\	$O((n+m)K^2 + \mathbf{X}^+ K)$
BPR	\	\	$O((n+m + \mathbf{X}^+)K)$
CDAE	\	\	$O((n+m + \mathbf{X}^+)K)$
EXMF	\	\	$O(nmK^3)$
FAWMF	\	\	$O((n+m)K^2D + \mathbf{X}^+ (K+D))$

$|\mathbf{X}^+|(K+D))$, where $|\mathbf{X}^+|$ denotes the number of the observed positive data. Here we spend $O(K^2Dn)$ time on calculating $\mathbf{M}^{(q)}$, $\mathbf{S}^{(q)}$ and $O(K^2Dm)$ time on calculating the first and second term in equation (11) for each item $j \in I$. For the third term, we just sum over the terms with $x_{ij} = 1$ with complexity $O(|\mathbf{X}^+|(K+D))$. Similar strategy can be used to accelerate the update of other parameters. With this algorithm, learning FAWMF model is efficient and stable, which trains on the all data but its complexity is linear to the number of the observed data.

Experiments and analyses

In this section, we conduct experiments to evaluate the performance of FAWMF. Our experiments are intended to address the following major questions:

- (Q1) Does FAWMF outperform state-of-the-art implicit MF methods?
- (Q2) How does the proposed batch-based learning algorithm fBGD perform?

Experimental protocol

Datasets. Three benchmark datasets Moivlens², Amazon (food-reviews)³, Douban⁴ are used in our experiments. These datasets contain users’ feedback on items. The dataset statistics are presented in Table 1. Similar to (Chen et al. 2019b), we preprocess the datasets so that all items have at least three interactions and “binarize” user’s feedback into implicit feedback. That is, as long as there exists some user-item interactions (ratings or reviews), the corresponding implicit feedback is assigned a value of 1. Grid search and 5-fold cross validation are used to find the best parameters. In our FAWMF, we set $D=K=20$, $\varepsilon=1e-5$ and learning rate=0.1 across all datasets.

Compared methods. We compare FAWMF with following baselines. Table 2 concludes their characteristics.

- Item-pop: This is a simple baseline which recommends items based on global item popularity.

²<https://grouplens.org/datasets/movielens/>

³<https://www.kaggle.com/snap/amazon-fine-food-reviews>

⁴<https://www.cse.cuhk.edu.hk/irwin.king.new/pub/data/douban>

Table 3: The performance metrics of the compared methods. The boldface font denotes the winner in that column. The row ‘Impv’ indicates the relative performance gain of our FAWMF compared to the best results among baselines. ‘†’ indicates that the improvement is significant with t-test at $p < 0.05$.

Methods	Movielens				Amazon				Douban			
	Pre@5	Rec@5	NDCG@5	MRR	Pre@5	Rec@5	NDCG@5	MRR	Pre@5	Rec@5	NDCG@5	MRR
Item-pop	0.2092	0.0400	0.2201	0.8958	0.0027	0.0048	0.0055	0.0191	0.1409	0.0332	0.1582	0.6308
WMF(ALS)	0.3841	0.0924	0.4059	1.5751	0.0789	0.0406	0.0858	0.3269	0.2400	0.0656	0.2598	1.0113
eALS	0.3955	0.0917	0.4175	1.5998	0.0984	0.0348	0.1051	0.3951	0.2329	0.0646	0.2520	0.9880
BPR	0.3613	0.0798	0.3794	1.5023	0.0988	0.0469	0.1060	0.3969	0.2371	0.0582	0.2570	1.0093
CDAE	0.3786	0.0860	0.3950	1.5454	0.0948	0.0472	0.0947	0.3994	0.2377	0.0589	0.2573	1.0162
EXMF	0.3871	0.0936	0.4071	1.5720	0.0847	0.0418	0.0928	0.3683	0.2353	0.0666	0.2588	1.0016
FAWMF	0.4054	0.0949	0.4275	1.6279	0.1129	0.0441	0.1285	0.4470	0.2661	0.0680	0.2915	1.0984
Impv%	2.49%†	1.41%†	2.40%†	1.76%†	14.34%†	-6.42%	21.18%†	11.93%†	10.88%†	2.11%†	12.19%†	8.09%†

- WMF(ALS) (Hu, Koren, and Volinsky 2008; Pan et al. 2008): The classic weighted matrix factorization model for implicit feedback data. The corresponding ALS-based (Hu, Koren, and Volinsky 2008) algorithm can reduce inference complexity.
- eALS (He et al. 2016): The improved weighted matrix factorization model, where the data confidence weights are assigned heuristically based on item-popularity.
- BPR (Rendle et al. 2009): The classic pair-wise method for recommendation, coupled with matrix factorization. BPR implicitly downweights the unobserved data in their uniformly negative strategy.
- CDAE (Wu et al. 2016): The advanced recommendation method based on Auto-Encoders, which is a generalization of WMF with more flexible components. CDAE also employs uniformly negative strategy to learn their model.
- EXMF (Liang et al. 2016): A probabilistic model that directly incorporates user’s exposure to items into traditional matrix factorization.

Evaluation Metrics. We adopt four well-known metrics Precision@K (Pre@K), Recall@K (Rec@K), Normalized Discounted Cumulative Gain (NDCG@K) and Mean Reciprocal Rank (MRR) to evaluate recommendation performance: Recall@K (Rec@K) quantifies the fraction of consumed items that are in the top-K ranking list; Precision@K (Pre@K) measures the fraction of the top-K items that are indeed consumed by the user; NDCG@K and MRR evaluate ranking performance of the methods. Refer to our supplemental material for more details about these metrics.

Performance comparison (Q1)

Table 3 presents the performance of the compared methods in terms of three evaluation metrics. The boldface font denotes the winner in that column. For the sake of clarity, the last row of Table 3 also show the relative improvements achieved by our FAWMF over the baselines. Generally speaking, with one exception, FAWMF outperforms all compared methods on all datasets for all metrics.

The improvement of FAWMF over these baselines can be attributed to three aspects: (1) In the real world, users have personalized communities and thus are exposed to diverse information. Correspondingly, the data confidence (exposure) will vary with different user-item combinations. That

is, some unobserved feedback are more likely attributed to user’s preference while others are the results of users’ awareness. By adaptively learning fine-grained data confidence weights from the data, FAWMF achieves better performance than those baselines with manually coarse-grained confidence weights. (2) FAWMF models confidence weight with a community-based inference network, which is capable of capturing latent interactions between users and mitigating the over-fitting problems. It can be seen from the better performance of FAWMF over EXMF. (3) Instead of employing sampling-based stochastic gradient descent optimizer (SGD), FAWMF employs a specific fast batch-based learning algorithm, which has better convergence and more stable results. We also conduct specific experiments in the next subsection to validate this point.

Running time comparisons. Figure 3 depicts running time of the six compared recommendation methods. As we can see, the speed up of our FAWMF over EXMF is significant. Especially in the largest dataset Douban, EXMF requires 56 hours for training, while FAWMF only takes 1.8 hours. The acceleration of FAWMF over EXMF can be attributed two aspects: (1) Confidence weights for each data have been modeled with a parameterized function, which reduces the number of learned parameters. (2) fBGD speed up gradients calculations. This way, FAWMF achieves similar analytical time complexity as other compared methods, which aim at fast learning but sacrifice the flexibility of the confidence weights. Their actual running time are also in the same magnitude. Also, we observe that these BGD-based methods (WMF, eALS, FAWMF) are relatively more efficient than SGD-based methods (BPR, CDAE), although they have similar analytical time complexity. It is caused by the poor convergence of SGD, which usually requires more iterations.

Effect of batch-based learning algorithm (Q2)

In this subsection, we compare our fBGD with two popular SGD-based learning strategies: (1) Uniform-1X (Uniform-5X, Uniform-25X), which uniformly sample a part of unobserved instances to update the model. Note that the training time and prediction accuracy are largely determined by the size of negative samples. Here we test this learning strategy with three different sampling-size for comparisons, where Uniform-25X (Uniform-5X, Uniform-1X) denotes the number of the sampled instances is 25 (5, 1) times as large as the




Figure 3: Running time comparisons.




Figure 4: Recommendation accuracy for different learning strategy versus the number of iterations and running time.

number of positive instances. (2) Itempop-1X (Itempop-5X, Itempop-25X) whose probability of sampling a un-observed data is proportional to the item popularity. We also present the performance of original BGD algorithm for running time comparisons. Here we do not compare with existing dynamic sampling strategies, since they either suffer from efficiency problems or require other side information.

Figure 4 presents pre@5 of FAWMF on dataset MovieLens with different learning strategies versus the number of iteration and running time. There exists the efficiency and effectiveness trade-off of SGD-based learning algorithm with varying sampling-size. The smaller sampling-size will cause insufficient learning and gradient instability. It can be seen from the poor performance final and heavy fluctuation of Uniform-1X (Itempop-1X). On the contrary, when the sampling-size become larger, SGD performs better but spend much more time. However, even if the stochastic optimizer uses a large sampling-size, it can not achieve good performance as BGD-based learning algorithm. Also, we can observe the inefficiency of original BGD. Our proposed fBGD accelerates original BGD for 16 times and even performs more efficient than Itempop-1X. Overall, our proposed fBGD performs better than others in all convergence, speed and recommended performance.

Effect of the parameter K

Figure 5 shows the performance of FAWMF with varying latent dimension K in MF. We also present the results of two close MF methods for comparisons. First, with few exceptions, FAWMF consistently outperforms eALS and WMF across K , demonstrating the effectiveness of adaptively assigning personalized confidence weights. Second, all methods can be improved with a larger K . But a large K might have the risk of overfitting. It can be seen from the worse




Figure 5: Impact of the parameter K .

Community 1		Community 2	
Movie Name	Genres	Movie Name	Genres
1.Gattaca	Drama,Sci-Fi,Thriller	1.My Favorite Year	Comedy
2.Thinner	Horror,Thriller	2.Annie Hall	Comedy,Romance
3.Bride of Frankenstein	Horror	3.Citizen Kane	Drama
4.The New Age	Drama	4.Unforgiven	Western
5.Halloween 4	Horror	5.Father of the Bride	Comedy
Ratio of male/female	Average age	Ratio of male/female	Average age
82.3%/17.7%	33.8	66.70%/33.3%	40.5

Figure 6: Case study of two communities: the top rows show top-5 items that have highest exposure to the users in the two communities; the bottom rows show the ratio of male/female and the average age of the users in each communities. Note that these side information is not used in model training.

results when $K = 60$ on the dataset MovieLens.

Case study

We also conduct a case study to show the learned communities. We do statistical analyses on the communities, two of which are presented in Figure 6. We can find that the genres of the exposed items in the community 1 are mainly about “Horror” and “Thriller”, but the genres in the community 2 are mainly about “Comedy”. These results can be explained by the different kinds of users in the communities. Most of users in the community 1 are young men, who are more likely to enjoy these exciting movies and share these movies with each other. Meanwhile, comedy is a hot topic for the elderly people, who are the main constituents of the community 2. These results validate the effectiveness of our FAWMF. Although the genre information of users and items is not used in model training, the latent correlations between users/items can be captured. Also, we can find different users will belong to different communities and thus have different exposure. Personalized confidence weights are necessary for the implicit recommendation models.

Conclusion

In this paper, we present a novel recommendation method FAWMF based on variational auto-encoder to achieve both adaptive weights assignment and efficient model learning. On the one hand, FAWMF models data confidence weights with a community-based inference neural network, which reduces the number of inferred parameters and is capable of capturing latent interactions between users. On the other hand, a specific batch-based learning algorithm fBGD has been developed to learn FAWMF fast and stably. The exper-

imental results on three real-world datasets demonstrate the superiority of FAWMF over existing implicit MF methods.

ACKNOWLEDGMENTS

This work is supported by National Natural Science Foundation of China (Grant No: U1866602) and National Key Research and Development Project (Grant No: 2018AAA0101503).

References

- Bayer, I.; He, X.; Kanagal, B.; and Rendle, S. 2017. A generic coordinate descent framework for learning from implicit feedback. In *International Conference on World Wide Web*, 1341–1350. IW3C2.
- Chen, C.; Li, D.; Zhao, Y.; Lv, Q.; and Shang, L. 2015. Wemarec: Accurate and scalable recommendation through weighted and ensemble matrix approximation. In *SIGIR*, 303–312. ACM.
- Chen, T.; Sun, Y.; Shi, Y.; and Hong, L. 2017. On sampling strategies for neural network-based collaborative filtering. In *International Conference on Knowledge Discovery and Data Mining*, 767–776. ACM.
- Chen, J.; Wang, C.; Shi, Q.; Feng, Y.; and Chen, C. 2019a. Social recommendation based on users’ attention and preference. *Neurocomputing* 341:1–9.
- Chen, J.; Wang, C.; Zhou, S.; Shi, Q.; Feng, Y.; and Chen, C. 2019b. Samwalker: Social recommendation with informative sampling strategy. In *The World Wide Web Conference*, 228–239. ACM.
- He, X.; Zhang, H.; Kan, M.-Y.; and Chua, T.-S. 2016. Fast matrix factorization for online recommendation with implicit feedback. In *SIGIR*, 549–558. ACM.
- He, X.; Liao, L.; Zhang, H.; Nie, L.; Hu, X.; and Chua, T.-S. 2017. Neural collaborative filtering. In *Proceedings of the 26th International Conference on World Wide Web*, 173–182. ACM.
- Hernández-Lobato, J. M.; Houlsby, N.; and Ghahramani, Z. 2014. Stochastic inference for scalable probabilistic modeling of binary matrices. In *International Conference on Machine Learning*, 379–387.
- Hoffman, M. D.; Blei, D. M.; Wang, C.; and Paisley, J. 2013. Stochastic variational inference. *The Journal of Machine Learning Research* 14(1):1303–1347.
- Hu, Y.; Koren, Y.; and Volinsky, C. 2008. Collaborative filtering for implicit feedback datasets. In *International Conference on Data Mining*, 263–272. Ieee.
- Jiawei, C.; Yan, F.; Martin, E.; Sheng, Z.; Chun, C.; and Wang, C. 2018. Modeling users’ exposure with social knowledge influence and consumption influence for recommendation. In *International Conference on Information and Knowledge Management*, 953–962. ACM.
- Johnson, C. C. 2014. Logistic matrix factorization for implicit feedback data. *Advances in Neural Information Processing Systems* 27.
- Kingma, D. P., and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105.
- Lee, J.; Kim, S.; Lebanon, G.; Singer, Y.; and Bengio, S. 2016. Llorma: Local low-rank matrix approximation. *The Journal of Machine Learning Research* 17(1):442–465.
- Li, D.; Chen, C.; Lv, Q.; Gu, H.; Lu, T.; Shang, L.; Gu, N.; and Chu, S. M. 2018. Adaerror: An adaptive learning rate method for matrix approximation-based collaborative filtering. In *The World Wide Web Conference*, 741–751. IW3C2.
- Liang, D.; Charlin, L.; McInerney, J.; and Blei, D. M. 2016. Modeling user exposure in recommendation. In *International Conference on World Wide Web*, 951–961. ACM.
- Liang, D.; Krishnan, R. G.; Hoffman, M. D.; and Jebara, T. 2018. Variational autoencoders for collaborative filtering. In *World Wide Web Conference*, 689–698. IW3C2.
- Palla, G.; Derényi, I.; Farkas, I.; and Vicsek, T. 2005. Uncovering the overlapping community structure of complex networks in nature and society. *Nature* 435(7043):814–818.
- Pan, R.; Zhou, Y.; Cao, B.; Liu, N. N.; Lukose, R.; Scholz, M.; and Yang, Q. 2008. One-class collaborative filtering. *ICDM* 502–511.
- Park, D. H., and Chang, Y. 2019. Adversarial sampling and training for semi-supervised information retrieval. In *The World Wide Web Conference*, 1443–1453. ACM.
- Rendle, S., and Freudenthaler, C. 2014. Improving pairwise learning for item recommendation from implicit feedback. In *International conference on Web search and data mining*, 273–282. ACM.
- Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2009. Bpr: Bayesian personalized ranking from implicit feedback. In *conference on uncertainty in artificial intelligence*, 452–461. AUAI Press.
- Shi, Q.; Wang, C.; Chen, J.; Feng, Y.; and Chen, C. 2019a. Location driven influence maximization: Online spread via offline deployment. *Knowledge-Based Systems* 166:30–41.
- Shi, Q.; Wang, C.; Chen, J.; Feng, Y.; and Chen, C. 2019b. Post and repost: A holistic view of budgeted influence maximization. *Neurocomputing* 338:92–100.
- Wang, J.; Yu, L.; Zhang, W.; Gong, Y.; Xu, Y.; Wang, B.; Zhang, P.; and Zhang, D. 2017. Irgan: A minimax game for unifying generative and discriminative information retrieval models. In *SIGIR*, 515–524. ACM.
- Wang, M.; Zheng, X.; Yang, Y.; and Zhang, K. 2018. Collaborative filtering with social exposure: A modular approach to social recommendation. In *AAAI, New Orleans, Louisiana, USA, February 2-7, 2018*.
- Wu, Y.; DuBois, C.; Zheng, A. X.; and Ester, M. 2016. Collaborative denoising auto-encoders for top-n recommender systems. In *International Conference on Web Search and Data Mining*, 153–162. ACM.
- Yu, L.; Zhang, C.; Pei, S.; Sun, G.; and Zhang, X. 2018. Walkranker: A unified pairwise ranking model with multiple relations for item recommendation. In *Conference on Artificial Intelligence*. AAAI.
- Yu, H.-F.; Bilenko, M.; and Lin, C.-J. 2017. Selection of negative samples for one-class matrix factorization. In *Proceedings of the 2017 SIAM International Conference on Data Mining*, 363–371. SIAM.
- Zhou, T. 2011. Understanding online community user participation: a social influence perspective. *Internet research* 21(1):67–81.