Adversarial Tensor Factorization for Context-aware Recommendation

Huiyuan Chen and Jing Li

Electrical Engineering and Computer Science, Case Western Reserve University

Introduction

Contextual factors such as time, location, or tag, can affect user preferences. Context-aware recommendations are thus critical to improve both quality and explainability of recommender systems, compared to traditional recommendations that are solely based on user-item interactions. Tensor factorization machines have achieved the state-of-the-art performance due to their capability of integrating users, items, and contextual factors in one unify way. However, few work has focused on the robustness of a context-aware recommender system. Improving the robustness of a tensor-based model is challenging due to the sparsity of the observed tensor and the multi-linear nature of tensor factorization. In this paper, we propose *ATF*, a model that combines tensor factorization and adversarial learning for context-aware recommendations, which enhancing the robustness of a recommender model. Empirical studies on two real-world datasets show that the proposed method outperforms standard tensor-based methods.

Problem Definition

We use the notation in tag recommender [4], in which the contextual factor is the *tag*. Let $\mathcal{U}$, $\mathcal{I}$, and $\mathcal{T}$ denote the set of users, items, and tags, respectively. The historical user-item-tag events can be represented by the set $\mathcal{D}$, *i.e.*, $(p,q,r) \in \mathcal{D}$ means that user p has tagged an item q with the tag r. We can use a tensor $\mathcal{X} \in \mathbb{R}^{P\times Q\times R}$ to indicate the interactions among the users, items and tags, where P,Q,R are the number of users, items and tags, respectively. if $(p,q,r) \in \mathcal{D}$, $\mathcal{X}_{pqr}=1$, otherwise, $\mathcal{X}_{pqr}=0$. In tag-aware recommendations, for a given user-item pair (p,q), the goal is to provide a list of tags that user p is likely to label item q. This can be achieved by predicting the missing entries in the tensor $\mathcal{X}$.

Tensor Factorization

The PITF [4] decomposes tensor $\mathcal{X}$ via three pairs of inner products among the latent vectors of the user $(\mathbf{u}_p^{(v)}, \mathbf{u}_p^{(t)})$, item $(\mathbf{v}_q^{(u)}, \mathbf{v}_q^{(t)})$ and tag $(\mathbf{t}_r^{(u)}, \mathbf{t}_r^{(v)})$:

$$\hat{\mathcal{X}}_{pqr}(\Theta) = \langle \mathbf{u}_p^{(v)}, \mathbf{v}_q^{(u)}\rangle + \langle \mathbf{u}_p^{(t)}, \mathbf{t}_r^{(u)}\rangle + \langle \mathbf{v}_q^{(t)}, \mathbf{t}_r^{(v)}\rangle, \tag{1}$$

where $\mathbf{u}_p^{(v)}, \mathbf{u}_p^{(t)} \in \mathbb{R}^K$ denote the latent factors of the user p interacting with the item q like $\mathbf{v}_q^{(u)}$ and the tag r like $\mathbf{t}_r^{(u)}$, respectively.
The PITF model is optimized with Bayesian Personalized Ranking (BPR) criterion from implicit feedback [5]. The core idea is to optimize rankings by considering $(p,q,r,r') \in \mathcal{D}_{pq}$, where

$$\mathcal{D}_{pq} = \{(p,q,r,r') | (p,q,r) \in \mathcal{D} \wedge (p,q,r') \notin \mathcal{D}\}.$$

The BPR tries to minimize the following objective function:

$$\mathcal{L}_{\mathsf{BPR}}(\Theta) = \sum_{(p,q,r,r')\in\mathcal{D}_{pq}} -\ln\sigma(\hat{\mathbf{A}}_{pqrr'}(\Theta)) + \lambda\|\Theta\|_F^2, \tag{2}$$

where $\sigma(\cdot)$ is the sigmoid function, $\|\cdot\|_F$ is the Frobenius norm, λ is the regularization parameter, and $\hat{\mathbf{A}}_{pqrr'}(\Theta) = \hat{\mathcal{X}}_{pqr}(\Theta) - \hat{\mathcal{X}}_{pqr'}(\Theta)$ for short.

Our Proposed ATF

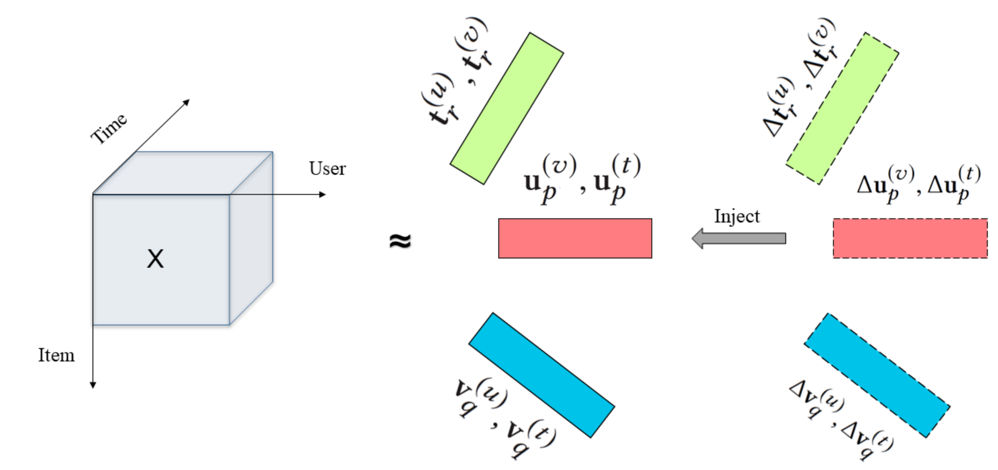

Fig. 1: Adversarial Tensor Factorization.

Inspired by [2, 3], we inject adversarial perturbations Δ on latent factors to quantify the loss of a tensor model under perturbations on its parameters:

$$\hat{\mathcal{X}}_{pqr}(\Theta+\Delta) = \langle \mathbf{u}_p^{(v)} + \Delta\mathbf{u}_p^{(v)}, \mathbf{v}_q^{(u)} + \Delta\mathbf{v}_q^{(u)}\rangle + \langle \mathbf{u}_p^{(t)} + \Delta\mathbf{u}_p^{(t)}, \mathbf{t}_r^{(u)} + \Delta\mathbf{t}_r^{(u)}\rangle + \langle \mathbf{v}_q^{(t)} + \Delta\mathbf{v}_q^{(t)}, \mathbf{t}_r^{(v)} + \Delta\mathbf{t}_r^{(v)}\rangle, \tag{3}$$

where the perturbation vectors Δ are coupled with their corresponding latent factors, *i.e.*, $\Delta\mathbf{u}_p^{(v)} \in \mathbb{R}^K$ denotes the perturbation vector for latent vector $\mathbf{u}_p^{(v)}$. Moreover, the goal of adversarial perturbations is to cause largest influence on the model, which are also known as the worse-case perturbations [1]. Therefore, we find the optimal adversarial perturbations by maximize the BPR loss:

$$\Delta_{adv} = \arg\max_{\Delta} \mathcal{L}_{\mathsf{BPR}}(\hat{\Theta}+\Delta), \qquad s.t. \quad \|\Delta\| \le \epsilon, \tag{4}$$

where ϵ controls the magnitude of adversarial perturbations; $\hat{\Theta}$ is the intermediate parameters.
We minimize the adversarial BPR loss by combining Eq.(2) and Eq.(4) as follow:

$$\begin{aligned} \mathcal{L}_{\mathsf{ATF}}(\Theta) &= \mathcal{L}_{\mathsf{BPR}}(\Theta) + \alpha\mathcal{L}_{\mathsf{BPR}}(\Theta+\Delta_{adv}), \\ \text{where} \quad \Delta_{adv} &= \arg\max_{\Delta,\|\Delta\|\le\epsilon} \mathcal{L}_{\mathsf{BPR}}(\hat{\Theta}+\Delta), \end{aligned} \tag{5}$$

where α controls the impact of the adversarial permutations on the model optimization.

Learning Algorithm

As the intermediate variable Δ maximizes the objective function that is minimized by Θ, the optimization in Eq. (5) can be formulated as a minimax objective function:

$$\Theta^*, \Delta^* = \arg\min_{\Theta} \max_{\Delta,\|\Delta\|\le\epsilon} \mathcal{L}_{\mathsf{BPR}}(\Theta) + \alpha\mathcal{L}_{\mathsf{BPR}}(\Theta+\Delta), \tag{6}$$

where the optimization of model parameters Θ is the minimizing player and adversarial perturbations Δ is the maximizing player. The two players alternately play the minimax game until convergence.
Updating Δ: Given a training instance (p,q,r,r'), the adversarial perturbations Δ can be updated by maximizing:

$$\max_{\Delta,\|\Delta\|\le\epsilon} l_{adv}(\Delta) = -\alpha\ln\sigma(\hat{\mathbf{A}}_{pqrr'}(\hat{\Theta}+\Delta)), \tag{7}$$

Updating Θ: The model parameters Θ can be obtained by minimizing:

$$\min_{\Theta} l_{ATF}(\Theta) = -\ln\sigma(\hat{\mathbf{A}}_{pqrr'}(\Theta)) - \alpha\ln\sigma(\hat{\mathbf{A}}_{pqrr'}(\Theta+\Delta_{adv})) + \lambda\|\Theta\|_F^2, \tag{8}$$

where $\hat{\mathbf{A}}_{pqrr'}(\hat{\Theta}+\Delta) = \hat{\mathcal{X}}_{pqr}(\hat{\Theta}+\Delta) - \hat{\mathcal{X}}_{pqr'}(\hat{\Theta}+\Delta)$.

Experiments

We consider two public datasets from HetRec 2011: MovieLens and Last.fm. We compare with three popular tensor factorization methods: CANDECOMP/PARAFAC (CP), HOSVD and PITF. We use the common evaluation scheme of F1-measure for top-N recommendations.

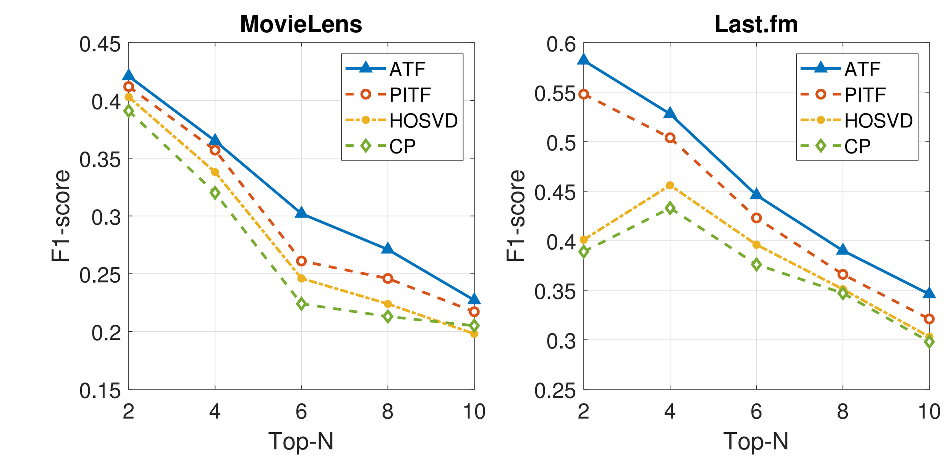

Fig. 2: Performance comparison of different tensor models in top-N recommendations.

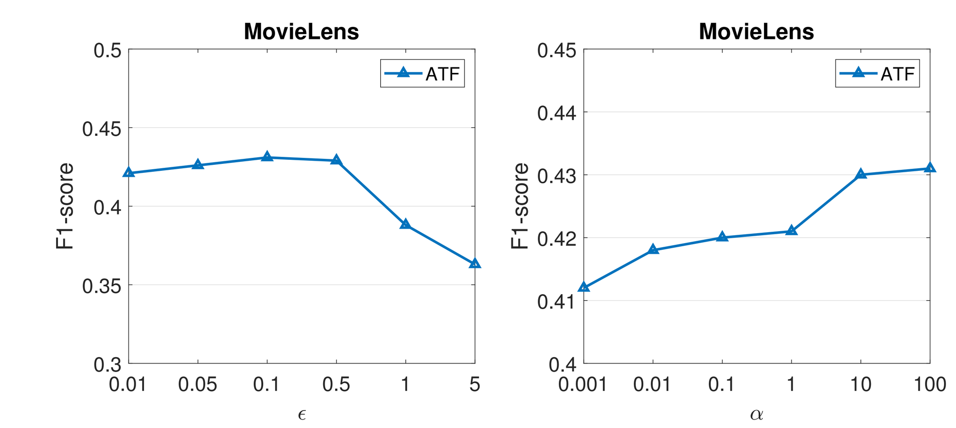

Fig. 3: The impacts of the ϵ and α for top-2 recommendations for the MovieLens dataset.

Conclusion

In this work, we consider the robustness of tensor-based models for context-aware recommendations, and design a new method ATF, which combines tensor factorization and adversarial learning. We also develop a learning algorithm to solve our minimax optimization. In the future, we plan to further investigate the problem of incorporating more contextual factors to better understand the users' behaviors and interests, such as time, location, and users' social networks.

References

References

[1] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. "Explaining and harnessing adversarial examples". In: *ICLR*. 2015.

[2] Xiangnan He et al. "Adversarial personalized ranking for recommendation". In: *SIGIR*. 2018.

[3] Seyed-Mohsen Moosavi-Dezfooli et al. "Universal adversarial perturbations". In: *CVPR*. 2017.

[4] Steffen Rendle and Lars Schmidt-Thieme. "Pairwise interaction tensor factorization for personalized tag recommendation". In: *WSDM*. 2010.

[5] Steffen Rendle et al. "BPR: Bayesian personalized ranking from implicit feedback". In: *UAI*. 2009.