Poster: Towards Privacy-Preserving Federated Recommendation via Synthetic Interactions

Thirasara Ariyarathna, Salil S. Kanhere, Hye-Young Paik
School of Computer Science and Engineering, University of New South Wales, Sydney, Australia
t.devanmini, salil.kanhere, h.paik@unsw.edu.au

***Abstract*—We propose a defence against inference attacks in Federated Recommendation Systems (FedRecs) using Generative Adversarial Networks (GAN). Results on real-world datasets show that our method can achieve better privacy utility balance.**

***Index Terms*—Federated Learning, Recommendation, Privacy**

I. Introduction

Service applications use Recommendation Systems to assist users in navigating numerous options by providing personalized suggestions. These systems are trained using users' past interactions, requiring the collection of user activity data (e.g., GPS location data from smartphones). Privacy concerns have led to using Federated Recommenders (FedRecs), which utilize intermediate parameters for training instead of real user data. However, FedRecs can be vulnerable to inference attacks, particularly for user-item interactions. To mitigate these attacks, FedRecs randomly sample pseudo-interacted items. The success rate of the inference attack is directly proportional to the sampling rate of these pseudo-interactions. To effectively minimize the success rate, the set of pseudo-interacted items should be at least equal to the set of actual interacted items. However, randomly sampling an equal number of pseudo interactions significantly decreases recommendation accuracy. To address this issue, we propose a novel method that uses GANs to generate synthetic user-item interactions in FedRecs. This method generates item interactions that closely resemble users' preferences and replaces the actual interactions, thereby protecting privacy-sensitive user-item interactions while maintaining recommendation performance.

II. Methodology

We consider the typical FL setting and server as the adversary trying to infer user-item interactions. We adopt widely used NeuMF and LightGCN as the base recommendation models implemented in the FL setting. FL process begins with the server sharing the global recommendation model and the initial generator, trained on public data. As shown in Fig. 1, the user u's local training procedure involves three steps. **Step-1**: Select the items to replace the original interacted items I_u with generated pseudo-interacted items S_u. We determine the number of items to be replaced using replacement ratio $R = |I_u \cap S_u|/|I_u|$. The more items are replaced, the less data leakage risk. To maximize the utility of the recommendation task, we select the items that are less contributing to the user's preference. For this, we employ the representation of items v_i that u interacted, and then the user's preferences can be

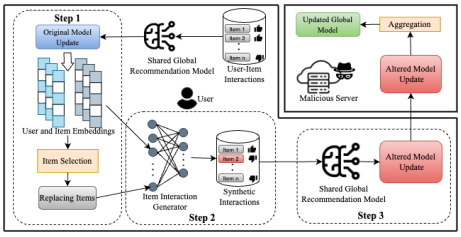

Fig. 1: Overview of the proposed defence

presented as $p_u = 1/|I_u|(\Sigma_{i \in I_u} a_{ui} v_i)$. Here, a_{ui} is a trainable parameter denoting the attention weight of item i for u's preference. **Step-2**: Training the generator to incorporate user preferences so that generated items are likely to be chosen by the user. The generator uses the concatenation of the user embedding vector and the selected item embedding vector as the input to get a latent feature $\mathcal{V}_{u,i} = W(p_u, v_i)+b$ for the output. Then, we calculate the similarity between the latent feature $\mathcal{V}_{u,i}$ and all item embedding D_I as $h_{u,i} = \mathcal{V}_{u,i} D_I^T$ Finally, we estimate the probability distribution over all candidate items as $y_{u,i} = softmax(h_{u,i})$ and then top n candidate items are selected according to replacement ratio. **Step-3**: Sending the local update obtained based on generated items to the server.

III. Experiments and Results

We conducted evaluations using two real-world datasets, MovieLens (ML) and Foursquare, in two cities, New York (NY) and Tokyo (TYO). We compared our method, Synthetic Interaction addition (SI), with Random pseudo-Interaction addition (RI) for three different replacement ratios (R). Fig. 2 shows the performance for the two recommendation methods, NeuMF and LightGCN. We gained up to 56%, 24% and 52% performance improvement for the same privacy protection (same R) for ML, NY and TYO datasets, respectively.

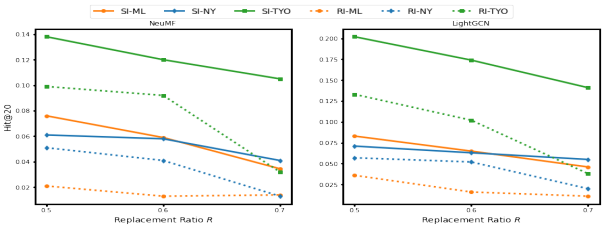

Fig. 2: Hit@20 against R on three datasets for RI and SI

IV. Conclusion

The study concludes that our approach of using GANs in FedRecs can enhance user privacy by effectively combating inference attacks while maintaining recommendation quality.