PPGenCDR: A Stable and Robust Framework for Privacy-Preserving Cross-Domain Recommendation

Xinting Liao[1], Weiming Liu[1], Xiaolin Zheng[1], Binhui Yao[2], Chaochao Chen[1*]

[1]College of Computer Science and Technology, Zhejiang University, China
[2]Midea Group, Foshan, China
{xintingliao, 21831010, xlzheng, zjuccc}@zju.edu.cn, tony.yao@midea.com

Abstract

Privacy-preserving cross-domain recommendation (PPCDR) refers to preserving the privacy of users when transferring the knowledge from source domain to target domain for better performance, which is vital for the long-term development of recommender systems. Existing work on cross-domain recommendation (CDR) reaches advanced and satisfying recommendation performance, but mostly neglects preserving privacy. To fill this gap, we propose a privacy-preserving generative cross-domain recommendation (`PPGenCDR`) framework for PPCDR. `PPGenCDR` includes two main modules, i.e., *stable privacy-preserving generator* module, and *robust cross-domain recommendation* module. Specifically, the former isolates data from different domains with a generative adversarial network (GAN) based model, which stably estimates the distribution of private data in the source domain with Ŕenyi differential privacy (RDP) technique. Then the latter aims to robustly leverage the perturbed but effective knowledge from the source domain with the raw data in target domain to improve recommendation performance. Three key modules, i.e., *(1) selective privacy preserver*, *(2) GAN stabilizer*, and *(3) robustness conductor*, guarantee the cost-effective trade-off between utility and privacy, the stability of GAN when using RDP, and the robustness of leveraging transferable knowledge accordingly. The extensive empirical studies on Douban and Amazon datasets demonstrate that `PPGenCDR` significantly outperforms the state-of-the-art recommendation models while preserving privacy.

Introduction

In recent years, cross-domain recommendation (CDR) is boosting for its effectiveness in alleviating the sparsity problem by transferring informative knowledge across related domains (Zhu et al. 2021a). On the one hand, the user-item interaction data plays a vital role in enlightening the recommendation performance in CDR (Zhu et al. 2021b; Chen et al. 2020; Li and Tuzhilin 2020). On the other hand, the interaction data is closely related to user privacy.

Although existing work on CDR achieves advanced performance, most of them assume that it is accessible to transfer interaction data across domains in plaintext (Chen et al. 2020; Li and Tuzhilin 2020). In reality, different domains maintain different types of user-sensitive data, e.g., Amazon has rich user-commodity interaction data, while IMDb is rich in user-movie interaction data. As shown in Fig. 1, it assumes that directly transferring the knowledge without privacy-preserving of a source domain (movie) to boost the performance of a target domain (book) is accessible. However, constrained by legal regulations, a domain may not be able to utilize the interaction data of other domains in plaintext. Therefore, privacy-preserving CDR (PPCDR) is vital for the long-term development of recommender systems.

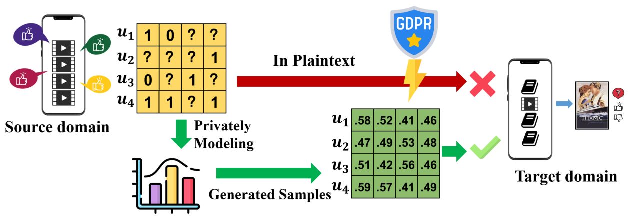

Figure 1: An example of the necessity of `PPGenCDR`.

Recently, several work was proposed to preserve the privacy of recommender systems in two data storage locations, i.e., (1) *the end-user device* that generates data, and (2) *the central database* that legally collects data uploaded by the end-user devices. The *former* avoids publishing user data to the server by modeling the data locally, while unifying a model globally for maintaining performance (Ammad-Ud-Din et al. 2019). However, the efficiency of the former is limited by the frequent communications among local devices and central server (if any). The *latter* prevents privacy leakage from modeling data (Liu et al. 2022a), or transferring information among different central platforms (Gao et al. 2019; Cui et al. 2021). Nevertheless, these methods do not devote to protecting the interaction data of users in CDR. PriCDR (Chen et al. 2022) directly applies differential privacy (DP) to interaction data in a source domain and publishes the perturbed data to a target domain. However, PriCDR *causes looser privacy budget bound compared with using Ŕenyi DP (RDP)* (Mironov 2017) and *can not perform recommendation robustly* (Gao et al. 2022).

To take advantage of interaction data, and bypass its disadvantage of leaking privacy, we devote to using *privacy-preserving data publishing* in PPCDR, which which is a

[*]Chaochao Chen is the corresponding author.

popular technique to release data to public in the perturbed form (Dwork et al. 2009; Fung et al. 2010; Beaulieu-Jones et al. 2019). As Fig. 1 depicts, we use privacy-preserving data publishing to isolate the raw data (the yellow matrix) in the source domain from target domain. To reach this goal, we propose a framework based on generative adversarial network (GAN), namely, privacy-preserving generative model-based CDR (`PPGenCDR`) in this work. Firstly, we adversarially model the distribution of the raw data in the source domain with RDP. Then we generate dense but fake ratings (the green matrix) to target domain.

Nevertheless, three challenges make it non-trivial to implement `PPGenCDR`. Firstly, **CH1**: *How to preserve privacy in the most economical way?* In order to preserve privacy during modeling, DP clips and perturbs the gradients of every layer in the model (Abadi et al. 2016). This approach reduces the utility of the model due to adding a lot of noise, but gains insignificant benefits of privacy-preserving. To address **CH1**, `PPGenCDR` contains a **selective privacy preserver (SPP)**, which only applies RDP on the gradients of model layer that directly accesses the raw data. The overall privacy can be theoretically guaranteed by post-processing (Dwork, Roth et al. 2014) and chain rule.

Secondly, **CH2**: *How to improve the stability of `PPGenCDR` when using RDP to preserve privacy?* GAN-based model optimizes a minmax goal, which is difficult to converge under the game of generator and discriminator (Roth et al. 2017). It worsens the process of modeling data when applying RDP to perturb the gradients of model in source domain. To tackle **CH2**, we design a **GAN stabilizer (GS)** to stabilize the objective of discriminator by an extra regularization derived from control theory.

Lastly, **CH3**: *How to improve the robustness of CDR models in target domain at the lowest cost when leveraging perturbed data?* There are several advanced CDR models in target domain that perform well in modeling source domain data in plaintext. But it requires a costly modification for these models to adapt to the perturbed data. To resolve **CH3**, we leverage **robust conductor (RC)** as an extra and flexible plugin of CDR model. Specifically, RC improves the model performance in target domain by disentangling each dimension of the redundant representations in a batch.

We summarize our contributions as follows: (1) We are the first to privately model the distribution of interaction data, and transfer it as the knowledge of source domain to improve the recommendation performance in target domain. (2) We propose a `PPGenCDR` framework with three special modules, i.e., *SPP*, *GS*, and *RC*, to stably model the raw data of source domain and robustly enhance the performance of target domain in a privacy-preserving way. (3) We conduct extensive experiments to empirically prove the utility of our proposed `PPGenCDR` for PPCDR.

Related Work

Deep Learning-based CDR. CDR has been proposed to address the data sparsity problem in recommender systems by leveraging the denser information from related domains (Zhu et al. 2021a). Existing CDR models can not only model user-item relationships more accurately, but also bridge domain knowledge more effectively. In the former, PPGN (Zhao, Li, and Fu 2019) uses graph neural network to model the high-order user-item interactions. In the latter, various techniques are utilized for bridging the information across domains, such as cross connection (Hu, Zhang, and Yang 2018), dual learning (Li and Tuzhilin 2020), and domain adaption (Liu et al. 2022b). However, they overlook the importance of preserving the privacy of users, e.g., the interaction behaviors of users.

Privacy-preserving CDR. User data in recommender systems are generally stored in two locations, i.e., (1) the *end-user device* that generates data, and (2) the *central database* that legally collects data. The former first models data in decentralized devices, and then unifies the model parameters. FCF (Ammad-Ud-Din et al. 2019) protects the behaviors of users by perturbing the local gradients. MetaMF (Lin et al. 2020) uses meta-learning to the federated recommender systems. But these solutions are inefficient in frequently communicating among different devices. The later focuses on privately modeling data, and transferring the centralized data across domains. UPC-SDG (Liu et al. 2022a) preserves privacy by synthesizing the less contributive items in single-domain recommendation. Several work focuses on transferring different kinds of knowledge, e.g., locations (Gao et al. 2019) and social relationships (Cui et al. 2021), without leaking privacy. Nevertheless, these work overlooks protecting the interaction data for recommendation in CDR. PriCDR is the first work to protect user rating information for transferring. However, it causes looser privacy budget bound and unreliable recommendation performance.

GAN-based Recommendation. IRGAN (Wang et al. 2017) is the first trial to apply GAN-based model to recommender systems, which generates user-item pairs and discriminates their relevance. CFGAN (Chae et al. 2018) proposes a model to generate the continuous purchase vectors, which avoids to deteriorating the performance of discriminator caused by contradicting labels. Several work devotes to augment the rating data, e.g., (1) using the generated negative samples (Chae, Shin, and Kim 2019), and (2) using auxiliary side information (Wang et al. 2019). GANMF (Dervishaj and Cremonesi 2022) notices that a single-valued feedback from discriminator is deficient in generating high-dimensional data, and models the difference of the real and fake profiles by a autoencoder like EBGAN (Zhao, Mathieu, and LeCun 2016). Most of these GAN-based models cannot be directly applied to `PPGenCDR`, due to the several reasons, (1) overlooking privacy-preserving, (2) not stabilizing the modeling process, and (3) requiring raw data as input.

Preliminaries of Differential Privacy

Definition 1 (Ŕenyi Differential Privacy (RDP) (Mironov 2017))**.** A mechanism $\mathcal{M}$ is $(\alpha, \varepsilon)-$RDP with order $\alpha \in (1, \infty)$ if for all neighboring datasets $\mathbf{x}$ and $\mathbf{x}'$, the Ŕenyi divergence satisfies: $D_\alpha \left(\mathcal{M}(\mathbf{x}) \| \mathcal{M}\left(\mathbf{x}'\right)\right) \leq \varepsilon$, and further equals to $\left(\varepsilon + \frac{\log 1/\delta}{\alpha - 1}, \delta\right) -$DP, for any $\delta \in (0, 1)$.

Proposition 1 (RDP with Gaussian Mechanism (Dwork, Roth et al. 2014; Mironov 2017))**.** For a d-dimensional function $f : \mathbf{x} \to \mathbb{R}^d$ with sensitivity $\Delta_f = \max_{\mathbf{x},\mathbf{x}'} \|f(\mathbf{x}) - f(\mathbf{x}')\|_2$, the Gaussian mechanism $\mathcal{M}_{\sigma_p,B} = f(\mathbf{x}) + \mathcal{N}(0, \sigma_p^2 B^2 \mathbf{I})$, parameterized by the variance of Gaussian noise σ_p, is $(\alpha, {\alpha \Delta_f}^2 / 2\sigma_p^2 B^2)-$RDP.

Methodology

Framework of `PPGenCDR`

Firstly, we describe notations. Without loss of generality, we assume there are two domains (i.e., $\mathcal{S}$ and $\mathcal{T}$) in CDR, which fully share the user set $\mathcal{U}$ with $N_{\mathcal{U}}$ users, but differ in rating matrices $\mathbf{R}_{\mathcal{S}} \in \mathbb{R}^{N_{\mathcal{U}} \times N_{\mathcal{S}}}$ and $\mathbf{R}_{\mathcal{T}} \in \mathbb{R}^{N_{\mathcal{U}} \times N_{\mathcal{T}}}$ for $N_{\mathcal{S}}$ items in $\mathcal{S}$ and $N_{\mathcal{T}}$ items in $\mathcal{T}$, respectively. Let $\boldsymbol{r}_i$, the i-th row of $\mathbf{R}$, be the user preference of user i. The goal of PPCDR is to transfer the information of $\mathbf{R}_{\mathcal{S}}$ in source domain to target domain **in a privacy-preserving way**, and further improve the recommendation performance with $\mathbf{R}_{\mathcal{T}}$.

To resolve PPCDR, we propose a framework named as `PPGenCDR` in this paper. As Fig. 2 depicts, `PPGenCDR` includes two modules, i.e., (1) *stable privacy-preserving generator* (*SPPG*) module in the source domain, and (2) *robust cross-domain recommendation* (*RCDR*) module in the target domain. Specifically, *SPPG* aims at stably modeling a privacy-preserving distribution of $\mathbf{R}_{\mathcal{S}}$ in source domain by selectively using RDP. To leverage the information of user identifications ($\mathbf{u}_{id}$), we formulate *SPPG* as a conditional GAN (CGAN) model and take $\mathbf{u}_{id}$ as conditions. Then *SPPG* publishes its *generator* to target domain, which is a flexible *private data publisher* to generate perturbed $\widetilde{\mathbf{R}}_{\mathcal{S}}$ while allowing frequent queries in downstream tasks. Thirdly, *RCDR* robustly enhances the recommendation performance in target domain by leveraging the perturbed knowledge $\widetilde{\mathbf{R}}_{\mathcal{S}}$ in source domain and the sparse information $\mathbf{R}_{\mathcal{T}}$ in target domain.

Stable Privacy-Preserving Generator (SPPG)

Motivation. In source domain, we aim to protect the private data by a *private data publisher*, which isolates the source domain who provides data, and the target domain who requires data. Specifically, we train *SPPG* using RDP, to obtain a privacy-preserving distribution of data, and publish the generator of *SPPG* as a private data publisher for two benefits. Firstly, the target domain gets access to the knowledge of source domain by sampling fake user preferences generated from the generator in *SPPG*, *avoiding the risk of privacy exposure*. After well trained, the privacy budget will be a fixed value, since the generator in *SPPG* does not need to visit the private data any more. Secondly, *SPPG augments the sparse rating data* in an adversarial way to better capture the user preferences in source domain.

Backbone of *SPPG*. To start with, we introduce *SPPG* in details. As the left of Fig. 2 illustrates, *SPPG* has a *generator* ($\mathcal{G}$) and a *discriminator* ($\mathcal{D}$), and takes user identifications $\mathbf{u}_{id}$ as conditions for both of them. Specifically, $\mathcal{G}$ first concatenates a random noise vector $\boldsymbol{\sigma}_g \sim \mathcal{N}(0, \boldsymbol{I})$ and conditions $\mathbf{u}_{id}$. Then $\mathcal{G}$ takes the concatenation as the input to construct fake user preferences $\tilde{\boldsymbol{r}}_{\mathcal{S}} = \mathcal{G}(\boldsymbol{\sigma}_g, \mathbf{u}_{id})$. Given the conditions $\mathbf{u}_{id}$, $\mathcal{D}$ is encouraged to correctly classify the real user preference $\boldsymbol{r}_{\mathcal{S}}$ and the fake one $\tilde{\boldsymbol{r}}_{\mathcal{S}}$. We measure the distribution divergence of $\boldsymbol{r}_{\mathcal{S}}$ and $\tilde{\boldsymbol{r}}_{\mathcal{S}}$ by Wasserstein distance (Arjovsky, Chintala, and Bottou 2017), which is good for modeling high-dimensional user preferences (Dervishaj and Cremonesi 2022). To generate realistic user preferences, $\mathcal{G}$ cheats $\mathcal{D}$ by minimizing the Wasserstein distance. On the contrary, $\mathcal{D}$ maximizes the Wasserstein distance to distinguish the distributions of fake user preferences and the real ones. We maximize the loss of $\mathcal{L}_{\mathcal{G}}$ and $\mathcal{L}_{\mathcal{D}}$ iteratively to obtain the distribution of data in source domain:

$$\begin{cases} \mathcal{L}_{\mathcal{G}} = \mathbb{E}_{\tilde{\boldsymbol{r}} \sim p_{\tilde{\boldsymbol{r}}_{\mathcal{S}}})} \left[\mathcal{D}\left(\tilde{\boldsymbol{r}}, \mathbf{u}_{id}\right)\right], \\ \mathcal{L}_{\mathcal{D}} = \mathbb{E}_{\boldsymbol{r} \sim p_{\boldsymbol{r}_{\mathcal{S}}}} \left[\mathcal{D}\left(\boldsymbol{r}, \mathbf{u}_{id}\right)\right] - \mathbb{E}_{\tilde{\boldsymbol{r}} \sim p_{\tilde{\boldsymbol{r}}_{\mathcal{S}}}} \left[\mathcal{D}\left(\tilde{\boldsymbol{r}}, \mathbf{u}_{id}\right)\right]. \end{cases} \tag{1}$$

Enhancing *SPPG* with Privacy and Stability. Traditionally, we can directly perturb the gradients of *SPPG* by adding random noise in the training procedure, e.g., adopting DP-SGD (Abadi et al. 2016), and get the distribution of private data in source domain. However, it causes two challenges inevitably, i.e., *(1) trade-off between privacy and performance*, and *(2) stability of optimizing minmax objective*. The former means that adding too much noise will surely reduce the utility of the model, but may gain insignificant benefits of privacy-preserving. The latter takes place because GAN-based models are deficient in stabilizing the process of modeling (Roth et al. 2017), which will become worse when adding noise on the gradients for preserving privacy. To tackle the two challenges, we devise selective privacy preserver (SPP) and GAN stabilizer (GS) for *SPPG*, which have a united optimization objective of selectively preserving privacy and stably modeling a GAN-based model.

Preserving privacy selectively for SPPG with SPP. To protect the private data in source domain and avoid severe performance degradation in target domain, we devise SPP which uses RDP technique to selectively perturb the gradients of *SPPG*. As shown in the left of Fig. 2, only $\mathcal{D}$ in *SPPG* can directly access to the private real ratings. Then generator obtains the knowledge of the private real ratings from the feedback of $\mathcal{D}$. Thus, SPP only needs to apply RDP on the gradients of $\mathcal{D}$, i.e., $\boldsymbol{g}_{\mathcal{D}}$, in back-propagation, and keeps the remaining gradients of *SPPG* in private, according to the post-processing property of RDP (Dwork, Roth et al. 2014).

Additionally, we expand $\boldsymbol{g}_{\mathcal{D}}$ in *SPPG* by chain rule:

$$\boldsymbol{g}_{\mathcal{D}} = \nabla_{\mathcal{D}} \mathcal{L}_{\mathcal{D}}\left(\boldsymbol{r}_{\mathcal{S}}, \tilde{\boldsymbol{r}}_{\mathcal{S}}; \boldsymbol{\theta}_{\mathcal{D}}, \mathbf{u}_{id}\right) \cdot J_{\boldsymbol{\theta}_{\mathcal{D}}}, \tag{2}$$

where $J_{\boldsymbol{\theta}_{\mathcal{D}}}$ is the Jacobian matrix of $\mathcal{D}\left(\boldsymbol{r}_{\mathcal{S}}, \tilde{\boldsymbol{r}}_{\mathcal{S}}; \boldsymbol{\theta}_{\mathcal{D}}, \mathbf{u}_{id}\right)$. Hence we can simply perturb the gradients of the first layer in $\mathcal{D}$ (Chen, Orekondy, and Fritz 2020) and preserve privacy.

Specifically, SPP implements RDP on $\boldsymbol{g}_{\mathcal{D}}$ to get its sanitized form $\hat{\boldsymbol{g}}_{\mathcal{D}}$ by clipping and applying Gaussian mechanism on $\boldsymbol{g}_{\mathcal{D}}$. Firstly, SPP bounds the sensitivity of optimization on training examples by clipping the gradient, i.e., ${\boldsymbol{g}} / {\max(1, \|\boldsymbol{g}\|/B)}$ to ensure $\|\boldsymbol{g}\| \leq B$, where B is a clipping constant. After that, SPP gets the sanitized form $\hat{\boldsymbol{g}}_{\mathcal{D}}$ by applying Gaussian mechanism to perturb the values of $\nabla_{\mathcal{D}} \mathcal{L}_{\mathcal{D}}$:

$$\hat{\boldsymbol{g}}_{\mathcal{D}} = \mathrm{SPP}(\nabla_{\mathcal{D}} \mathcal{L}_{\mathcal{D}}\left(\boldsymbol{r}_{\mathcal{S}}, \tilde{\boldsymbol{r}}_{\mathcal{S}}; \boldsymbol{\theta}_{\mathcal{D}}, \mathbf{u}_{id}\right)) \cdot J_{\boldsymbol{\theta}_{\mathcal{D}}}, \tag{3}$$

where $\mathrm{SPP}(\cdot) = \mathrm{clip}(\cdot) + \mathcal{N}\left(0, \sigma_p^2 B^2 \boldsymbol{I}\right)$, and σ_p is the variance of noise in Gaussian mechanism. Therefore, SPP

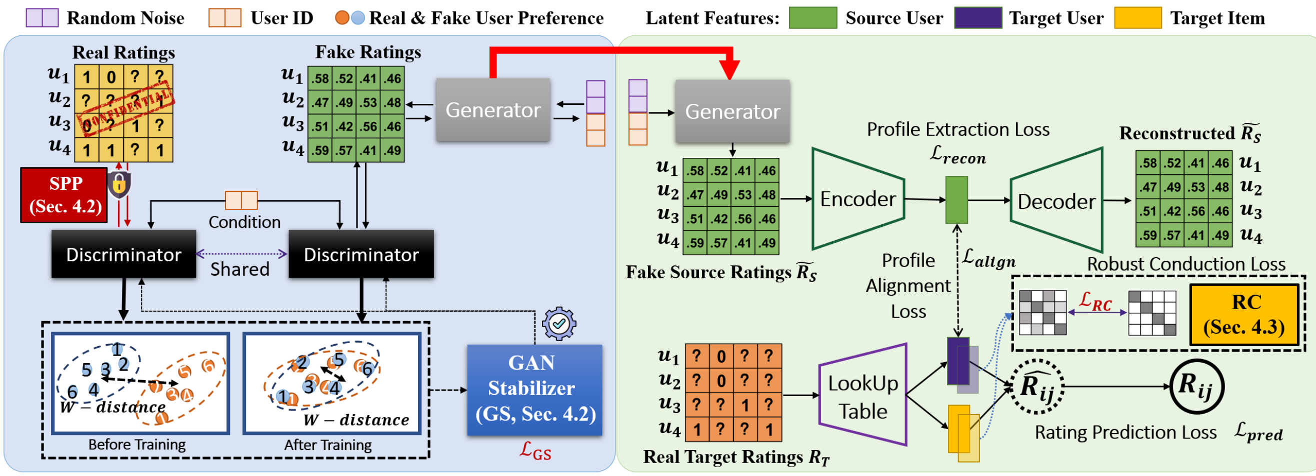

(a) Stable Privacy-Preserving Generator Module in Source Domain (b) Robust Cross-Domain Recommendation Module in Target Domain

Figure 2: Framework of `PPGenCDR`.

replaces $\boldsymbol{g}_{\mathcal{D}}$ (Eq. (2)) with $\hat{\boldsymbol{g}}_{\mathcal{D}}$ (Eq. (3)) in updating, which preserves privacy of the modeling procedure selectively.

Theorem 1 (Privacy Bound of Each *SPPG* Update Step)**.** *For batch size N, clipping constant B, and Gaussian mechanism $\mathcal{M}_{\sigma_p,B}$, each update step for training SPPG satisfies $(\alpha, 2N\alpha/\sigma_p^2)-RDP$ for $B=1$, and $\left(2N\alpha/\sigma_p^2 + \frac{\log 1/\delta}{\alpha-1}, \delta\right)-DP$.*

Proof. The RDP of updating a single user is $(\alpha, 2\alpha/\sigma_p^2)-$RDP with $B = 1$ in Gaussian mechanism, i.e., $\mathcal{M}_{\sigma_p,1} = \text{clip}(\boldsymbol{g}_{\mathcal{D}}) + \mathcal{N}\left(0, \sigma_p^2 \boldsymbol{I}\right)$. Then, the privacy amplification of updating a batch of users could be obtained by using the composition of N Gaussian mechanisms (Dwork 2008). Thus, we obtain the RDP privacy upper bound for a batch of users, i.e., $(\alpha, 2N\alpha/\sigma_p^2)-$RDP, which is measured by $\alpha-$order Rényi divergence as bellow:

$$D_\alpha\left(\hat{\boldsymbol{g}}_{\mathcal{D}}(\boldsymbol{r}_{1:N}), \hat{\boldsymbol{g}}_{\mathcal{D}}\left(\boldsymbol{r}'_{1:N}\right)\right) \leq N \cdot 2\alpha/\sigma_p^2. \tag{4}$$

According to (Dwork, Roth et al. 2014) and (Mironov 2017), the DP privacy upper bound further equals to $\left(2N\alpha/\sigma_p^2 + \frac{\log 1/\delta}{\alpha-1}, \delta\right) -$DP. □

Stabilizing the modeling process of SPPG with GS. Note that SPP applies Gaussian mechanism with noise σ_{gr}, ranging from $[-\sigma_p B, \sigma_p B]$, on the gradients in *SPPG*. In this way, the ground truth distribution of user preferences is not definite during the modeling procedure. To simulate the perturbed ground truth distribution, we first simplify *SPPG* to a Dirac GAN (Mescheder, Geiger, and Nowozin 2018), which is widely studied for analyzing the stability of GAN-based model. In Dirac GAN, the ground truth distribution of $\boldsymbol{r}$ is $p(\boldsymbol{r}) = \text{Dirac}(\boldsymbol{r} - \boldsymbol{c})$, where all elements of the vector $\boldsymbol{c}$ are constant. Then we substitute the vector $\boldsymbol{c}$ in $p(\boldsymbol{r})$ with $\boldsymbol{m}(t) = \boldsymbol{c} - (1 - 2\phi)\boldsymbol{\sigma}_{\text{gr}}$ ($\phi \in [0, 1]$) by interpolating the noise brought from SPP. In terms of Eq. (1), *SPPG* aims to estimate the distribution of perturbed ground truth $p(\boldsymbol{r}) = \text{Dirac}(\boldsymbol{r} - \boldsymbol{m}(t))$ as follows:

$$\frac{d\boldsymbol{\theta}_{\mathcal{D}}}{dt} = \boldsymbol{m}(t) - \boldsymbol{\theta}_{\mathcal{G}}(t), \qquad \frac{d\boldsymbol{\theta}_{\mathcal{G}}}{dt} = \boldsymbol{\theta}_{\mathcal{D}}(t). \tag{5}$$

From the perspective of control theory, we can treat Eq. (5) as a dynamic system, which takes $\boldsymbol{m}(t)$ as input, and $\boldsymbol{y}(t) = (\boldsymbol{\theta}_{\mathcal{D}}(t), \boldsymbol{\theta}_{\mathcal{G}}(t)), \forall t > 0$ as output. And we figure out that the modeling procedure of *SPPG* in Eq. (5) is not a stable dynamic system. Moreover, the stability of modeling $\mathcal{G}$ depends on the stability of modeling $\mathcal{D}$.

Motivated by (Xu et al. 2020), we take GS, as a controller parameterized by τ ($\tau > 0$), to formulate a closed-loop control system (Kailath 1980), which is widely used to improve the stability of a non-linear dynamic system like Eq. (5):

$$\frac{d\boldsymbol{\theta}_{\mathcal{D}}}{dt} = \boldsymbol{m}(t) - \boldsymbol{\theta}_{\mathcal{G}}(t) - \tau\boldsymbol{\theta}_{\mathcal{D}}(t). \tag{6}$$

Comparing Eq. (5) and Eq. (6), we find that GS brings an extra optimization objective to the loss function of $\mathcal{D}$:

$$\mathcal{L}_{\text{GS}} = -\frac{\tau}{2}\left(\mathbb{E}_{\boldsymbol{r}\sim p_{\boldsymbol{r}_{\mathcal{S}}}}\left[\mathcal{D}^2\left(\boldsymbol{r}, \mathbf{u}_{id}\right)\right] + \mathbb{E}_{\tilde{\boldsymbol{r}}\sim p_{\tilde{\boldsymbol{r}}_{\mathcal{S}}}}\left[\mathcal{D}^2\left(\tilde{\boldsymbol{r}}, \mathbf{u}_{id}\right)\right]\right).$$

Intuitively, we can find that the loss of GS $\mathcal{L}_{\text{GS}}$ is actually a regularization term, which penalizes the abrupt outputs from $\mathcal{D}$. Given the feedback of $\mathcal{L}_{\text{GS}}$, *SPPG* adjusts its $\mathcal{D}$ to measure the distance of the real and fake user preferences ($\boldsymbol{r}$ and $\tilde{\boldsymbol{r}}$) in a tighter latent representation space, which stably estimates the distribution of private data in source domain.

Optimization in *SPPG*. Finally, derived from the Eq. (1), we can achieve a stable modeling procedure for *SPPG* by maximizing the loss of *SPPG* with SPP and GS:

$$\mathbf{L}_{\mathcal{G}} = \mathcal{L}_{\mathcal{G}}, \quad \mathbf{L}_{\mathcal{D}} = \mathcal{L}_{\mathcal{D}} + \mathcal{L}_{\text{GS}}. \tag{7}$$

In summary, SPP (red block in Fig. 2) preserves the privacy of data in source domain without decreasing the performance severely, and GS (blue block in Fig. 2) stabilizes the modeling procedure of `PPGenCDR` to better capture the distribution of private data in source domain.

Robust CDR (RCDR)

Motivation. After privately modeling the distribution of data by *SPPG* in source domain, we introduce the way of leveraging it in target domain. Specifically, we devote to leverage the knowledge contained in the generator published from *SPPG*, and improve the recommendation performance with the raw data of target domain. *RCDR* should extract user preferences from the published generator and combine the user preferences of two domains together to enhance the target domain recommendation.

Backbone of *RCDR*. We introduce the general overview of *RCDR* in the right of Fig. 2, which contains (1) a *generator* published from *SPPG*, and (2) a *CDR module* to leverage the perturbed information from source domain and perform recommendation in target domain. Additionally, CDR module consists of three sub-modules, i.e., a *profile extraction* module, a *recommendation prediction* module, and a *profile alignment* module. Firstly, given user identification $\mathbf{u}_{id}$ and a random noise vector $\boldsymbol{\sigma}_g \sim \mathcal{N}(0, \boldsymbol{I})$, the well-trained generator samples perturbed, but effective user preferences $\tilde{\boldsymbol{r}}_{\mathcal{S}} = \mathcal{G}(\boldsymbol{\sigma}_g, \mathbf{u}_{id})$. Then the *profile extraction* module reconstructs $\tilde{\boldsymbol{r}}_{\mathcal{S}}$ to extract latent feature as user profile $\boldsymbol{z}_{\mathcal{S}} = \mathrm{Enc}(\tilde{\boldsymbol{r}}_{\mathcal{S}})$ ($\boldsymbol{z}_{\mathcal{S}} \in \mathbb{R}^{1\times K}$ with K denoting the dimension size of the source domain), which is empirically useful for target model (Chen et al. 2022). To capture a batch of user profiles in source domain $\boldsymbol{Z}_{\mathcal{S}} = [\boldsymbol{z}_{\mathcal{S},1}, \ldots, \boldsymbol{z}_{\mathcal{S},N}] \in \mathbb{R}^{N\times K}$, we minimize the reconstruction loss as follows:

$$\mathcal{L}_{\text{recon}} = \sum_{u=1}^{N} \left\|\mathrm{Dec}\left(\mathrm{Enc}\left(\tilde{\boldsymbol{r}}_{\mathcal{S},u}\right)\right) - \tilde{\boldsymbol{r}}_{\mathcal{S},u}\right\|_F^2, \tag{8}$$

where N is the batch size. Meanwhile, a *recommendation prediction* module, i.e., DMF (Xue et al. 2017), learns a lookup table from $\mathbf{R}_{\mathcal{T}}$ in the target domain, to obtain a batch of user profiles $\boldsymbol{Z}_{\mathcal{T}} = [\boldsymbol{z}_{\mathcal{T},1}, \ldots, \boldsymbol{z}_{\mathcal{T},N}]$ and the item features $\boldsymbol{V}_{\mathcal{T}} = [\boldsymbol{v}_{\mathcal{T},1} \ldots, \boldsymbol{v}_{\mathcal{T},N}]$, where $\boldsymbol{Z}_{\mathcal{T}}, \boldsymbol{V}_{\mathcal{T}} \in \mathbb{R}^{N\times K}$. Additionally, we bridge the knowledge of source domain to the target domain by *profile alignment* module, which minimizes the user profiles in both domains ($\boldsymbol{z}_{\mathcal{S}}$ and $\boldsymbol{z}_{\mathcal{T}}$):

$$\mathcal{L}_{\text{align}} = \sum_{u=1}^{N} \left\|\boldsymbol{z}_{\mathcal{S},u} - \boldsymbol{z}_{\mathcal{T},u}\right\|_F^2. \tag{9}$$

Lastly, *recommendation prediction* module predicts user preference based on user profile $\boldsymbol{z}_{\mathcal{T}}$ and item property $\boldsymbol{v}_{\mathcal{T}}$ by minimizing binary cross entropy $\mathrm{F}_{\text{bce}}(\cdot)$ as below:

$$\mathcal{L}_{\text{pred}} = \sum_{u=1}^{N} \mathrm{F}_{\text{bce}}(\widehat{\boldsymbol{r}}_{\mathcal{T},u}, \boldsymbol{r}_{\mathcal{T},u}), \tag{10}$$

where $\widehat{\boldsymbol{r}}_{\mathcal{T},u} = \boldsymbol{z}_{\mathcal{T},u}{}^{\top}\boldsymbol{v}_{\mathcal{T}}$ is the predicted user preference of user u in target domain.

Enhancing *RCDR* with Robustness. Few existing work in CDR prevents the performance from deteriorating in target domain when utilizing the information from other domains (Gao et al. 2022). In PPCDR, source domain inevitably transfers the perturbed data to target domain for the sake of preserving privacy, which potentially decreases the performance of models in target domain.

Robustifying the recommendation performance in RCDR with RC. We propose a robustness conductor (RC) in *RCDR* to robustly enhance the performance of recommendation model in target domain. Specifically, RC minimizes the impact of noisy data generated by the published generator from *SPPG*, which empirically causes the latent features, i.e., $\boldsymbol{Z}_{\mathcal{S}}$, $\boldsymbol{Z}_{\mathcal{T}}$, and $\boldsymbol{V}_{\mathcal{T}}$, redundant and dependent in a batch of samples.

To enhance robustness, RC disentangles the features of different dimensions in $\boldsymbol{Z}_{\mathcal{T}}$ and $\boldsymbol{V}_{\mathcal{T}}$, which are directly used for predicting user preferences in target domain. The user features $\boldsymbol{Z}_{\mathcal{T}}$ and the item features $\boldsymbol{V}_{\mathcal{T}}$ have the same processing procedure. Take user features as an example, RC first uses Z-score normalization (Lv et al. 2022) to every column (dimension) of the batched user features $\boldsymbol{Z}_{\mathcal{T}}$ to get $\bar{\boldsymbol{Z}}_{\mathcal{T}}$, whose i-th column (dimension) is denoted as $\bar{z}_{\mathcal{T},i}$ ($i \le K$). Then RC measures the cross correlation matrix among different columns of normalized user features, i.e., $\boldsymbol{C}_{\bar{\boldsymbol{Z}}_{\mathcal{T}}} \in \mathbb{R}^{K\times K}$, where $\boldsymbol{C}_{\bar{\boldsymbol{Z}}_{\mathcal{T},ij}} = {\langle \bar{z}_{\mathcal{T},i}, \bar{z}_{\mathcal{T},j}\rangle}/{\|\bar{z}_{\mathcal{T},i}\|\|\bar{z}_{\mathcal{T},j}\|}$ (for $i, j \in 1, \ldots, K$). Intuitively, RC disentangles the redundant representations in a batch from two aspects: (1) keeping the representations of the same dimension consistent, and (2) enforcing the representations of different dimensions mutually independent. To realize the former, RC approximates the diagonal elements of cross-correlation matrix $\boldsymbol{C}_{\bar{\boldsymbol{Z}}_{\mathcal{T},ii}}$ to 1. To reach the latter, RC approximates $\boldsymbol{C}_{\bar{\boldsymbol{Z}}_{\mathcal{T},ij}}$ (for $i \ne j$) to 0. Similarly, RC processes the latent *items* representations $\boldsymbol{V}_{\mathcal{T}}$ to get its cross correlation matrix, i.e., $\boldsymbol{C}_{\bar{\boldsymbol{V}}_{\mathcal{T}}} \in \mathbb{R}^{K\times K}$, and disentangles the features of different dimensions. Finally, we have robust regularization for *RCDR* to minimize:

$$\mathcal{L}_{RC} = \|\boldsymbol{C}_{\bar{\boldsymbol{Z}}_{\mathcal{T}}} - \boldsymbol{I}\|_F^2 + \|\boldsymbol{C}_{\bar{\boldsymbol{V}}_{\mathcal{T}}} - \boldsymbol{I}\|_F^2, \tag{11}$$

where $\boldsymbol{I} \in \mathbb{R}^{K\times K}$ is the identity matrix. Therefore, RC can make the noisy and redundant representations into clean and independent ones, and further enhance recommendation prediction in a robust way. RC is marked in the right of Fig. 2 as a yellow block, which is a flexible plugin to improve the robustness of recommendation models.

Optimization in *RCDR*. Overall, the optimization of *RCDR* is to minimize:

$$\mathbf{L}_{\textit{RCDR}} = \mathcal{L}_{\text{recon}} + \mathcal{L}_{\text{pred}} + \lambda_A \mathcal{L}_{\text{align}} + \lambda_R \mathcal{L}_{\text{RC}}, \tag{12}$$

where λ_A and λ_R are hyper-parameters to balance different types of losses. `PPGenCDR` can stably model the distribution of private data in source domain and robustly enhance the recommendation performance in the target domain.

Experiments and Discussion

In this section, we aim to answer the following questions through empirical studies: **Q1**: Can `PPGenCDR` outperform existing single-domain recommendation models, the state-of-the-art (SOTA) CDR models in plaintext, and the SOTA PPCDR models? **Q2**: How can GS and RC contribute to the *performance* of `PPGenCDR`? **Q3**: How can SPP preserve privacy in `PPGenCDR` *in a cost-effective way*? **Q4**: How do hyper-parameters impact `PPGenCDR`?

Experimental Setup

Datasets. We use two datasets, i.e., **Amazon** (Ni, Li, and McAuley 2019) and **Douban** (Zhu et al. 2021b). Following (Chen et al. 2022), we first select three domains on Amazon dataset, i.e., Movies and TV (Movie), Books (Book), and

	Douban book → music						**Amazon** music → book					
	HR@5	**NDCG@5**	**MRR@5**	**HR@10**	**NDCG@10**	**MRR@10**	**HR@5**	**NDCG@5**	**MRR@5**	**HR@10**	**NDCG@10**	**MRR@10**
BPR	.1284	.0984	.0556	.2364	.1221	.0917	.2556	.1981	.1887	.3711	.2688	.2473
NeuMF	.1317	.0916	.0677	.2416	.1375	.1032	.2660	.2177	.1952	.4023	.2729	.2536
DMF	.1678	.1143	.1029	.2533	.1377	.1014	.2735	.2435	.2171	.4216	.2805	.2703
CFGAN	.1234	.0745	.0723	.2024	.0999	.0769	.1168	.0752	.0617	.2071	.1009	.0722
GANMF	.1385	.0879	.0714	.2219	.1146	.0824	.0632	.0369	.0285	.1180	.0546	.0357
CoNet	.2042	.1325	.1082	.3128	.1611	.1282	.3616	.2824	.2564	.4693	.3172	.2707
DDTCDR	.2100	.1446	.1116	.3315	.1701	.1329	.4091	.3199	.2823	.4864	.3540	.3296
ETL	.2781	.1974	.1691	.4048	.2379	.1857	.4883	.3798	.3439	.6092	.4188	.3600
DARec	.2056	.1346	.1108	.3149	.1677	.1243	.2975	.2119	.1846	.4175	.2503	.1995
PriCDR	.3074	.2107	.1791	.4448	.2549	.1972	.5605	.4527	.4177	.6663	.4875	.4321
`PPGenCDR`-**GS**	.2944	.1958	.1693	.4448	.2434	.1885	.5701	.4639	.4289	.6767	.4986	.4432
`PPGenCDR`-**RC**	.3084	.2123	.1808	.4686	.2639	.2022	.2862	.1953	.1659	.4175	.2375	.1832
`PPGenCDR`	**.3701**	**.2691**	**.2362**	**.5097**	**.3116**	**.2532**	**.6012**	**.4886**	**.4514**	**.7078**	**.5231**	**.4657**
Improvement (↑)	**20.40%**	**27.72%**	**31.88%**	**14.59%**	**22.24%**	**28.40%**	**7.26%**	**7.93%**	**8.07%**	**6.23%**	**7.30%**	**7.78%**
	Douban movie → book						**Amazon** movie → book					
BPR	.1835	.1100	.1041	.2858	.1440	.1285	.2642	.1981	.1887	.3837	.2169	.2247
NeuMF	.1902	.1266	.1093	.3164	.1562	.1302	.2866	.1922	.1920	.3923	.2443	.2581
DMF	.2132	.1220	.1170	.3536	.1625	.1488	.3035	.2144	.2417	.4116	.2568	.2670
CFGAN	.1760	.1113	.0873	.2745	.1363	.0977	.1377	.0885	.0762	.2176	.1143	.0839
GANMF	.1131	.0841	.0746	.1557	.0978	.0802	.1169	.0779	.0653	.1883	.1010	.0747
CoNet	.2516	.1681	.1517	.3432	.2316	.1782	.3611	.2571	.3095	.4366	.3088	.3242
DDTCDR	.2638	.1793	.1603	.3699	.2481	.1837	.4393	.3434	.3231	.4838	.3476	.3427
ETL	.3571	.2432	.2091	.4981	.2922	.2294	.5436	.4288	.3909	.6626	.4673	.4068
DARec	.2700	.1871	.1611	.3958	.2286	.1771	.3665	.2736	.2430	.4950	.3150	.2600
PriCDR	.3577	.2434	.2101	.4987	.2922	.2301	.5862	.4868	.4538	.6889	.5200	.4675
`PPGenCDR`-**GS**	.3647	.2640	.2323	.5114	.3088	.2510	.5976	.4970	.4633	.6923	.5275	.4759
`PPGenCDR`-**RC**	.4123	.2976	.2572	.5654	.3457	.2768	.3614	0.2641	.2322	.5003	.3088	.2505
`PPGenCDR`	**.4238**	**.3050**	**.2655**	**.5680**	**.3524**	**.2849**	**.6253**	**.5218**	**.4874**	**.7181**	**.5518**	**.4999**
Improvement (↑)	**18.48%**	**25.31%**	**26.37%**	**13.90%**	**20.60%**	**23.82%**	**6.67%**	**7.19%**	**7.40%**	**4.24%**	**6.12%**	**6.93%**
	Douban movie → music						**Amazon** movie → music					
BPR	.1535	.0920	.0941	.2958	.1640	.1115	.2443	.1578	.1631	.3793	.2238	.1971
NeuMF	.1625	.1005	.1061	.3027	.1498	.1220	.2664	.2137	.1801	.4052	.2372	.2188
DMF	.1872	.1023	.1039	.3060	.1712	.1267	.3427	.2219	.1958	.4126	.2528	.2375
CFGAN	.1422	.1850	.1845	.2351	.1983	.1895	.0966	.0618	.0484	.1694	.0833	.0571
GANMF	.1213	.0761	.0626	.2114	.0966	.0710	.0795	.0463	.0356	.1444	.0672	.0441
CoNet	.2331	.1347	.1322	.3217	.1924	.1355	.3712	.2824	.2964	.4693	.3172	.2707
DDTCDR	.2456	.1572	.1518	.3877	.2163	.1421	.3819	.3508	.3148	.4877	.3860	.3389
ETL	.3270	.2243	.1908	.4578	.2665	.2082	.5561	.4300	.3884	.6792	.4706	.4052
DARec	.2427	.1626	.1353	.3469	.1949	.1505	.3890	.3859	.3519	.5170	.4280	.3693
PriCDR	.3346	.2305	.1981	.4398	.2689	.2138	.6015	.4872	.4498	.7079	.5216	.4640
`PPGenCDR`-**GS**	.3318	.2268	.1972	.4777	.2740	.2171	.6071	.4905	.4515	.7200	.5267	.4662
`PPGenCDR`-**RC**	.2739	.1911	.1658	.4095	.2345	.1836	.3164	.2159	.1828	.4536	.2602	.2010
`PPGenCDR`	**.3782**	**.2727**	**.2335**	**.5185**	**.3181**	**.2512**	**.6379**	**.5187**	**.4779**	**.7460**	**.5529**	**.4921**
Improvement (↑)	**13.03%**	**18.31%**	**17.87%**	**13.26%**	**18.30%**	**17.49%**	**6.05%**	**6.47%**	**6.25%**	**5.38%**	**6.00%**	**6.06%**

Table 1: Experimental results of target domain on Douban and Amazon datasets. We bold the best result, underline the runner-up, and all of the improvements are significant with *p-value* < 0.01 using the paired sample t-test.

CD Vinyl (Music), and three domains on Douban dataset, i.e., Book, Music, and Movie. We binarize the ratings larger or equal to 3 as positive, and others as negative. We filter the users and items with less than 5 interactions.

Comparison methods. We compare `PPGenCDR` with the following representative models: (1) **BPR** (Rendle et al. 2009) optimizes matrix factorization with implicit feedback by a pairwise ranking loss. (2) **NeuMF** (He et al. 2017) is a neural collaborative filtering method which uses multi-layer perceptron (MLP) to learn the inner product for prediction. (3) **DMF** (Xue et al. 2017) is a deep matrix factorization model that maps both users and items into a low-dimensional space with non-linear projection. (4) **CFGAN** (Chae et al. 2018) generates the user purchase vector based on conditional GAN. (5) **GANMF** enhances the feedback of discriminator derived from energy-based GAN (Zhao, Mathieu, and LeCun 2016). (6) **CoNet** (Hu, Zhang, and Yang 2018) introduces a cross connection unit to enable dual knowledge transfer across domains. (7) **DDTCDR** (Li and Tuzhilin 2020) introduces deep dual learning to transfer knowledge in CDR. (8) **DARec** (Yuan, Yao, and Benatallah 2019) shares user representation across different domains via domain adaption. (9) **ETL** (Chen et al. 2020) adopts the equivalent transformation to selectively share domain-related and domain-specific knowledge. (10) **PriCDR** (Chen et al. 2022) is the SOTA PPCDR model that publishes the perturbed interaction data using DP to target domain.

Implementation details. We tune the hyper-parameters of `PPGenCDR` and baseline models to their best values for a fair comparison. Specifically, we choose RMSprop (Tieleman, Hinton et al. 2012) as the optimizer for *SPPG*, and Adam (Kingma and Ba 2014) for *RCDR*. We set batch size N as 128, learning rate $\eta = 0.01$ for Douban, and $\eta = 0.0005$ for Amazon. We set clipping constant $B = 1$, and the dimension of latent features $K = 200$. We evaluate the top 5 and 10 of ranking results by Hit Ratio (HR), Normalized Discounted Cumulative Gain (NDCG), and Mean Reciprocal Rank (MRR).

Recommendation Performance Comparison

Model Comparison (Q1). We report the comparison results on **Douban** and **Amazon** datasets in the Tab. 1, where $A \to B$ means transferring knowledge of domain A to domain B. We can find that: (1) Compared with the single-domain models, all of the CDR models contribute to alleviating the sparsity issue by transferring the knowledge

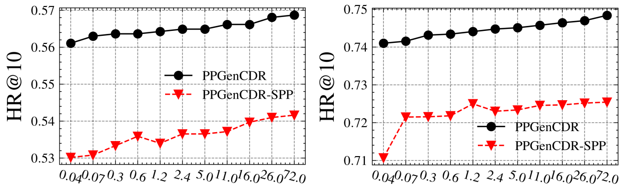

(a) **Douban** movie → book (b) **Amazon** movie → music

Figure 3: Privacy analysis of ε on PPGenCDR.

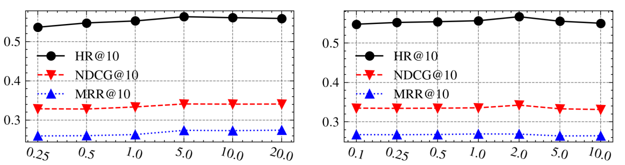

(a) Effect of regularizing λ_A (b) Effect of controller τ

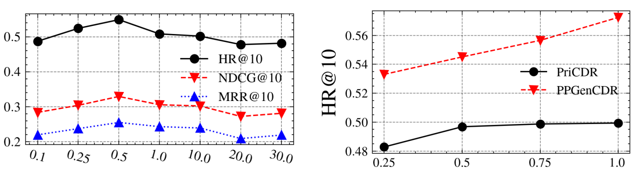

(c) Effect of regularizing λ_R (d) Effect of sparsity degree

Figure 4: Parameters analysis on Douban movie → book.

from other domains. (2) Compared with the SOTA CDR models in plaintext, PPGenCDR and PriCDR also achieve comparable performance, indicating that it is possible to find a trade-off between preserving privacy and maintaining performance. (3) PPGenCDR significantly outperforms the SOTA PPCDR model (i.e., PriCDR) by at least 13.03% and 6.05% on Douban and Amazon in terms of HR@5, respectively, illustrating that PPGenCDR breaks the performance bottleneck by enhancing model stability and robustness.

In-depth Model Analysis

Ablation Study (Q2). We remove GS and RC from PPGenCDR , to get PPGenCDR-GS and PPGenCDR-RC, in order to evaluate the utility of GS and RC, respectively. From Tab. 1, we can discover that: (1) Both PPGenCDR-GS and PPGenCDR-RC degrade their performance compared with PPGenCDR. This implies that the performance of PPCDR depends on not only stably modeling the distribution of private data in source domain, but also robustly using the perturbed information to recommend in the target domain. (2) Compared with PPGenCDR-GS, the performance of PPGenCDR-RC decreases more severely, meaning that it is indispensable to improve the robustness for PPCDR.

Empirical Study of Privacy (Q3). We analyze the effectiveness of SPP by replacing it with PPGenCDR-SPP which exploits RDP to the gradients of *the whole model*. We evaluate the contribution of SPP from two aspects, i.e., (1) preserving the privacy of source domain and (2) maintaining the performance of target domain, with two tasks, i.e., **Douban** movie → book, and **Amazon** movie → music. Specifically, to study the former one, we compare the performance of the generators, published from PPGenCDR and PPGenCDR-SPP, to predict user preferences in the *source* domain, and report the results in Tab. 2. To evaluate the latter one, we depict the change of HR@10 in *target* domain in Fig. 3, for PPGenCDR and PPGenCDR-SPP, by varying privacy budget $\varepsilon \in \{0.04, 0.07, 0.3, 0.6, 1.2, 2.4, 5, 11, 16, 26, 72\}$. We can find that: (1) The generator of PPGenCDR leaks a little more privacy, since it has slightly better results than PPGenCDR-SPP. (2) PPGenCDR significantly outperforms PPGenCDR-SPP with different privacy budgets. The results indicate that SPP is a cost-effective module for balancing preserving privacy and enhancing performance in PPCDR.

	Douban movie			**Amazon** movie		
	HR	**NDCG**	**MRR**	**HR**	**NDCG**	**MRR**
PPGenCDR	.0882	.0258	.0403	.0945	.0401	.0522
PPGenCDR-SPP	.0872	.0201	.0353	.1098	.0324	.0500

Table 2: Performance (top@10) of the generator of PPGenCDR in source domain.

Sensitivity of Hyper-parameters (Q4). We show the effect of hyper-parameters on **Douban** movie → book in Fig. 4. In Fig. 4, we compare the performance by varying the hyper-parameter of alignment $\lambda_A \in \{0.25, 0.5, 1, 5, 10, 20\}$ in (a), the parameter of GS $\tau \in \{0.1, 0.25, 0.5, 1, 2, 5, 10\}$ in (b), and the hyper-parameter of robustness $\lambda_R \in \{0.1, 0.25, 0.5, 1, 10, 20, 30\}$ in (c), respectively. To study the effect of sparsity degree, we compare the performance of PPGenCDR and PriCDR by randomly sampling 25%, 50%, 75%, and 100% of the data in source domain, and report the results in Fig.4 (d). From the results, we can conclude that: (1) PPGenCDR achieves the best performance when $\lambda_A = 1, \tau = 2$, and $\lambda_R = 0.5$. (2) The performance of PPGenCDR changes slightly with λ_A in Fig. 4 (a), meaning that PPGenCDR can better leverage knowledge of the source domain by alignment while maintain the performance with RC. (3) The bell-shaped curves in Fig. 4 (b) and Fig. 4 (c) indicate that it is effective to choose a proper hyper-parameter for better recommendation performance. (4) PPGenCDR outperforms PriCDR under all sparsity degrees in Fig. 4 (d), indicating that PPGenCDR can better address the sparsity issue.

Conclusion

In this paper, we propose PPGenCDR, i.e., a novel framework for privacy-preserving generative cross-domain recommendation. PPGenCDR devotes to stably model the distribution of private data in source domain by stable privacy-preserving generator module, and robustly enhance the performance of target domain by robust cross-domain recommendation module. The extensive empirical studies on two datasets (i.e., Douban and Amazon) demonstrate that PPGenCDR significantly outperforms the state-of-the-art recommendation models while preserving privacy.

Acknowledgements

This work was supported in part by the "Pioneer" and "Leading Goose" R&D Program of Zhejiang (No. 2022C01126), the National Natural Science Foundation of China (No. 62172362), and Leading Expert of "Ten Thousands Talent Program" of Zhejiang Province (No.2021R52001).

References

Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; Talwar, K.; and Zhang, L. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 308–318.

Ammad-Ud-Din, M.; Ivannikova, E.; Khan, S. A.; Oyomno, W.; Fu, Q.; Tan, K. E.; and Flanagan, A. 2019. Federated collaborative filtering for privacy-preserving personalized recommendation system. *arXiv preprint arXiv:1901.09888*.

Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein generative adversarial networks. In *International conference on machine learning*, 214–223. PMLR.

Beaulieu-Jones, B. K.; Wu, Z. S.; Williams, C.; Lee, R.; Bhavnani, S. P.; Byrd, J. B.; and Greene, C. S. 2019. Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes*, 12(7): e005122.

Chae, D.-K.; Kang, J.-S.; Kim, S.-W.; and Lee, J.-T. 2018. Cfgan: A generic collaborative filtering framework based on generative adversarial networks. In *Proceedings of the 27th ACM international conference on information and knowledge management*, 137–146.

Chae, D.-K.; Shin, J. A.; and Kim, S.-W. 2019. Collaborative adversarial autoencoders: An effective collaborative filtering model under the GAN framework. *IEEE Access*, 7: 37650–37663.

Chen, C.; Wu, H.; Su, J.; Lyu, L.; Zheng, X.; and Wang, L. 2022. Differential Private Knowledge Transfer for Privacy-Preserving Cross-Domain Recommendation. In *Proceedings of the ACM Web Conference 2022*, 1455–1465.

Chen, D.; Orekondy, T.; and Fritz, M. 2020. Gs-wgan: A gradient-sanitized approach for learning differentially private generators. *Advances in Neural Information Processing Systems*, 33: 12673–12684.

Chen, X.; Zhang, Y.; Tsang, I. W.; Pan, Y.; and Su, J. 2020. Towards equivalent transformation of user preferences in cross domain recommendation. *ACM Transactions on Information Systems (TOIS)*.

Cui, J.; Chen, C.; Lyu, L.; Yang, C.; and Li, W. 2021. Exploiting Data Sparsity in Secure Cross-Platform Social Recommendation. *Advances in Neural Information Processing Systems*, 34: 10524–10534.

Dervishaj, E.; and Cremonesi, P. 2022. GAN-based matrix factorization for recommender systems. In *Proceedings of the 37th ACM/SIGAPP Symposium on Applied Computing*, 1373–1381.

Dwork, C. 2008. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, 1–19. Springer.

Dwork, C.; Naor, M.; Reingold, O.; Rothblum, G. N.; and Vadhan, S. 2009. On the complexity of differentially private data release: efficient algorithms and hardness results. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, 381–390.

Dwork, C.; Roth, A.; et al. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4): 211–407.

Fung, B. C.; Wang, K.; Chen, R.; and Yu, P. S. 2010. Privacy-preserving data publishing: A survey of recent developments. *ACM Computing Surveys (Csur)*, 42(4): 1–53.

Gao, C.; Huang, C.; Yu, Y.; Wang, H.; Li, Y.; and Jin, D. 2019. Privacy-preserving cross-domain location recommendation. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 3(1): 1–21.

Gao, Y.; Du, Y.; Hu, Y.; Chen, L.; Zhu, X.; Fang, Z.; and Zheng, B. 2022. Self-Guided Learning to Denoise for Robust Recommendation. In Amigó, E.; Castells, P.; Gonzalo, J.; Carterette, B.; Culpepper, J. S.; and Kazai, G., eds., *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*, 1412–1422. 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.

Hu, G.; Zhang, Y.; and Yang, Q. 2018. Conet: Collaborative cross networks for cross-domain recommendation. In *Proceedings of the 27th ACM international conference on information and knowledge management*, 667–676.

Kailath, T. 1980. *Linear systems*, volume 156. Prentice-Hall Englewood Cliffs, NJ.

Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Li, P.; and Tuzhilin, A. 2020. Ddtcdr: Deep dual transfer cross domain recommendation. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, 331–339.

Lin, Y.; Ren, P.; Chen, Z.; Ren, Z.; Yu, D.; Ma, J.; Rijke, M. d.; and Cheng, X. 2020. Meta matrix factorization for federated rating predictions. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 981–990.

Liu, F.; Cheng, Z.; Chen, H.; Wei, Y.; Nie, L.; and Kankanhalli, M. 2022a. Privacy-Preserving Synthetic Data Generation for Recommendation Systems. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '22, 1379–1389. New York, NY, USA: Association for Computing Machinery. ISBN 9781450387323.

Liu, W.; Zheng, X.; Hu, M.; and Chen, C. 2022b. Collaborative Filtering with Attribution Alignment for Review-based Non-overlapped Cross Domain Recommendation. In *Proceedings of the ACM Web Conference 2022*, 1181–1190.

Lv, F.; Liang, J.; Li, S.; Zang, B.; Liu, C. H.; Wang, Z.; and Liu, D. 2022. Causality Inspired Representation Learning for Domain Generalization. In *Proceedings of* *the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8046–8056.

Mescheder, L.; Geiger, A.; and Nowozin, S. 2018. Which training methods for GANs do actually converge? In *International conference on machine learning*, 3481–3490. PMLR.

Mironov, I. 2017. Rényi differential privacy. In *2017 IEEE 30th computer security foundations symposium (CSF)*, 263–275. IEEE.

Ni, J.; Li, J.; and McAuley, J. 2019. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, 188–197.

Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, 452–461.

Roth, K.; Lucchi, A.; Nowozin, S.; and Hofmann, T. 2017. Stabilizing training of generative adversarial networks through regularization. *Advances in neural information processing systems*, 30.

Tieleman, T.; Hinton, G.; et al. 2012. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2): 26–31.

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 *Proceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval*, 515–524.

Wang, Q.; Yin, H.; Wang, H.; Nguyen, Q. V. H.; Huang, Z.; and Cui, L. 2019. Enhancing collaborative filtering with generative augmentation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 548–556.

Xu, K.; Li, C.; Zhu, J.; and Zhang, B. 2020. Understanding and stabilizing GANs' training dynamics using control theory. In *International Conference on Machine Learning*, 10566–10575. PMLR.

Xue, H.-J.; Dai, X.; Zhang, J.; Huang, S.; and Chen, J. 2017. Deep matrix factorization models for recommender systems. In *IJCAI*, volume 17, 3203–3209. Melbourne, Australia.

Yuan, F.; Yao, L.; and Benatallah, B. 2019. DARec: deep domain adaptation for cross-domain recommendation via transferring rating patterns. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 4227–4233.

Zhao, C.; Li, C.; and Fu, C. 2019. Cross-domain recommendation via preference propagation graphnet. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 2165–2168.

Zhao, J.; Mathieu, M.; and LeCun, Y. 2016. Energy-based generative adversarial network. *arXiv preprint arXiv:1609.03126*.

Zhu, F.; Wang, Y.; Chen, C.; Zhou, J.; Li, L.; and Liu, G. 2021a. Cross-Domain Recommendation: Challenges, Progress, and Prospects. In Zhou, Z., ed., *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, 4721–4728. ijcai.org.

Zhu, F.; Wang, Y.; Zhou, J.; Chen, C.; Li, L.; and Liu, G. 2021b. A unified framework for cross-domain and cross-system recommendations. *IEEE Transactions on Knowledge and Data Engineering*.