Engineering 4 (2018) 21–28

Contents lists available at ScienceDirect

Engineering

journal homepage: www.elsevier.com/locate/eng

Research
Cybersecurity—Review

Toward Privacy-Preserving Personalized Recommendation Services

Cong Wang [a,b,*], Yifeng Zheng [a,b], Jinghua Jiang [a,c], Kui Ren [d]

[a] Department of Computer Science, City University of Hong Kong, Hong Kong, China
[b] City University of Hong Kong, Shenzhen Research Institute, Shenzhen, Guangdong 518057, China
[c] Department of Computer Science and Technology, Xi'an Jiaotong University, Xi'an 710049, China
[d] Institute of Cyber Security Research, Zhejiang University, Hangzhou, Zhejiang 310058, China

ARTICLE INFO

Article history:
Received 21 June 2017
Revised 26 September 2017
Accepted 12 February 2018
Available online 16 February 2018

Keywords:
Privacy protection
Personalized recommendation services
Targeted delivery
Collaborative filtering
Machine learning

ABSTRACT

Recommendation systems are crucially important for the delivery of personalized services to users. With personalized recommendation services, users can enjoy a variety of targeted recommendations such as movies, books, ads, restaurants, and more. In addition, personalized recommendation services have become extremely effective revenue drivers for online business. Despite the great benefits, deploying personalized recommendation services typically requires the collection of users' personal data for processing and analytics, which undesirably makes users susceptible to serious privacy violation issues. Therefore, it is of paramount importance to develop practical privacy-preserving techniques to maintain the intelligence of personalized recommendation services while respecting user privacy. In this paper, we provide a comprehensive survey of the literature related to personalized recommendation services with privacy protection. We present the general architecture of personalized recommendation systems, the privacy issues therein, and existing works that focus on privacy-preserving personalized recommendation services. We classify the existing works according to their underlying techniques for personalized recommendation and privacy protection, and thoroughly discuss and compare their merits and demerits, especially in terms of privacy and recommendation accuracy. We also identity some future research directions.

1. Introduction

Nowadays, recommendation systems are increasingly gaining popularity and are widely deployed for online services. The widespread use of recommendation systems allows users to enjoy diverse personalized recommendations for movies, books, ads, restaurants, hotels, and more. Meanwhile, personalized recommendation services have also become extremely effective revenue drivers for online business. For example, recent research suggests that 35% of what consumers purchase on Amazon and 75% of what they watch on Netflix are attributable to personalized recommendations [1]. A study conducted by the research firm Marketing Sherpa showed that 11.5% of the revenue generated in the shopping sessions of involved e-commerce sites comes from purchases of products via personalized recommendation [2].

To support personalized recommendation, the common practice of existing systems usually involves either collaborative filtering-based (CFB) recommendation or content-based (CB) recommendation [3]. CFB recommendation systems usually recommend items based on the similarity between users. For example, a user rating for a movie would be predicted based on the ratings/decisions of other similar users (classified via some metric). CB recommendation systems typically generate recommendations by comparing the properties of items with those of users' personal preference/behavioral data. For example, an ad network may compare the keywords associated with ads with the keywords indicating a user's preference in order to serve personalized ads. To obtain personalized recommendations from these systems, users are typically required to provide their personal data to the recommender for processing and analytics.

Although personalized recommendation is greatly beneficial, directly exposing users' private data to the recommender poses privacy risks for users [4–6]: ① The provided data undesirably discloses the users' personal interests to the recommender; ② the provided data may be abused by the recommender, for example, by a recommender selling user data to third parties for financial incentives without user consent [4]; and ③ the provided data

* Corresponding author.
E-mail address: congwang@cityu.edu.hk (C. Wang).

https://doi.org/10.1016/j.eng.2018.02.005

may be stolen by motivated attackers due to security breaches on the recommender side [5,6]. Therefore, it is of paramount importance to develop privacy-preserving techniques for recommendation systems, so that the intelligence of recommendation systems is preserved while user privacy is respected.

In this paper, we survey the literature related to personalized recommendation services with privacy protection. We first present the general architecture of real-world recommendation systems, and the privacy issues therein. We then provide a comprehensive survey of existing solutions that can support privacy-preserving personalized recommendation services. As mentioned above, the mechanisms adopted in recommendation systems are usually either CFB or CB. Based on this observation, we first classify the existing solutions into two broad categories: ① privacy-preserving CFB recommendation and ② privacy-preserving CB recommendation. In the first category, existing works are further classified into private neighborhood-based approaches and private machine learning-based approaches, according to the concrete plaintext techniques adopted. In the second category, existing works are further classified into private targeted advertising and private targeted coupon delivery, according to the concrete application settings. Therefore, there are four explicit categories in total.

When describing the representative existing works in each category, our key insight is to further classify them based on their underlying security strategies/techniques for privacy protection; for example, some works rely on cryptographic techniques such as homomorphic encryption and garbled circuits (GCs), while others resort to data obfuscation techniques. Among the large volume of works in this trending area, we carefully select highly cited representative works that describe popular techniques, as well as papers that deliver significantly new and emerging techniques. Our goal is to cover each category as comprehensively as possible in order to call for further motivated research activities.

The rest of this paper is organized as follows. Section 2 presents the general architecture and privacy issues of recommendation systems. Section 3 describes existing works on privacy-preserving CFB recommendation. Section 4 describes existing works on privacy-preserving CB recommendation. Section 5 discusses some future research directions, and Section 6 concludes the paper.

2. Recommendation systems

2.1. System model

Recommendation systems aim to provide accurate recommendations for users by collecting and processing their personal data using effective approaches [7]. The system model of a personalized recommendation service is illustrated in Fig. 1. It contains two primary entities: users and recommender. Each user has some personal data on her local device (e.g., a smartphone), which indicates her personal interests/preferences. The recommender collects the users' personal data, processes the collected data, and provides personalized recommendations for the users. The generated recommendations can be provided or shown to users in various ways, such as through messages and pop-up windows.

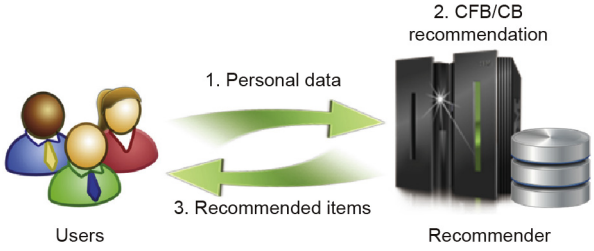

Fig. 1. The system model of a personalized recommendation service. The recommender may adopt either collaborative filtering-based (CFB) techniques or content-based (CB) techniques.

To process the collected user data for recommendation, the recommender may adopt different kinds of techniques. Roughly speaking, according to the adopted techniques, recommendation systems can be classified into two categories: CFB recommendation systems and CB recommendation systems. As mentioned earlier, CFB recommendation systems recommend items based on the similarity between users. That is, items recommended to a particular user are those preferred by other users that share similar preferences [3]. In contrast, CB recommendation systems conduct recommendation based on the properties of items, which may be described by certain explicit features (e.g., attributes and characteristics).

To leverage the similarity between users, CFB recommendation systems usually adopt either neighborhood-based approaches or machine learning-based approaches. Neighborhood-based approaches directly compute the similarity relationship between users [8], and leverage this relationship to generate personalized recommendations. In contrast, machine learning-based approaches first learn a mathematical model from the collected user data, and then apply the model to generate personalized recommendations.

2.2. Privacy issues

The more personal data a recommender collects, the more accurate recommendations users can obtain. The user data collected by the recommender may include information about the users' identity, demographic profile, behavioral data, purchase history, rating history, and more [9]. Such information can be very privacy-sensitive. For example, the demographic profile refers to demographic characteristics of the customer, such as age, gender, weight, and level of education; behavioral data refers to dynamic data of the customer, such as location, activity status, and browsing history; and rating history refers to the votes that the customer has provided on items. Providing such information to the recommender in the clear would pose undesirable privacy risks. For example, user data could be sold to a third party by the recommender without user consent, or could even be stolen by motivated attackers. Therefore, protecting user data in recommendation systems is of critical importance.

3. Privacy-preserving CFB recommendation

CFB recommendation systems typically recommend items based on similarity measures between users [3]. To support the functionality of CFB recommendation while preserving user privacy, a number of works on privacy-preserving CFB recommendation have been undertaken. Because CFB recommendation systems usually adopt either neighborhood-based approaches or machine learning-based approaches, we classify these existing works into two categories: private neighborhood-based approaches and private machine learning-based approaches.

3.1. Private neighborhood-based approaches

Existing solutions under the umbrella of private neighborhood-based approaches usually adopt techniques from two main categories. The first category is cryptographic techniques [4,6,10,11] and the second is randomization techniques [12–14]. Cryptographic technique-based solutions generally require high

computation overhead, which may not be well-suited for large-scale data. However, this type of approach can provide strong protection for user data under semantic security while ensuring recommendation accuracy. Randomization technique-based solutions apply randomized perturbation to users' private data, such as by adding proper noise; this process generally trades recommendation accuracy for user privacy. However, this type of approach has low computation overhead and is much faster than cryptographic technique-based solutions. In what follows, we describe some representative works on cryptographic technique-based solutions [4,6,10,11] and randomization technique-based solutions [12–14].

Erkin et al. [6] proposed an efficient privacy-preserving recommendation system using partially homomorphic encryption (PHE) and secure multiparty computation (SMC) protocols. Their design encrypts users' private data (i.e., user ratings on items) via PHE, so that the recommender cannot access the original ratings while still being able to process the data to generate private recommendations. More specifically, in order to avoid active user interaction with the recommender, their design introduces a semi-trusted third party to assist the recommender to complete the recommendation in the encrypted domain. Neither the recommender nor the extra third party can learn users' private data. With this architecture, users simply upload their encrypted data to the recommender and then stay offline. The recommender then runs a cryptographic protocol with the third party to generate personalized recommendations. Because directly applying PHE would lead to high computation and communication overhead, Erkin et al. [6] further designed packing techniques to pack several numerical values in a compact way prior to encryption, thus significantly boosting the overall system performance.

Later, Badsha et al. [10] proposed a recommendation system based on the ElGamal cryptosystem, which is a kind of PHE. Compared with the work in Ref. [6], their design does not introduce an additional party for assistance. However, all users are required to actively collaborate with the recommender server so as to generate private recommendations for a target user. Very recently, Badsha et al. [11] designed a new recommendation system, which relies on the Boneh–Goh–Nissim (BGN) homomorphic cryptosystem. Similar to the work in Ref. [6], their design adopts an additional server for assistance; however, no interaction is introduced between the recommendation server and the additional server. The additional server only assists the user in decrypting ciphertexts whenever necessary. A user who wants to obtain private recommendations must have active interaction with the recommendation/additional server. That is, a user is not able to simply send her private data to the recommendation server for processing and then later directly obtain personalized recommendations.

Unlike the above works, Li et al. [4] proposed a privacy-preserving recommendation system named YANA (short for "you are not alone") based on the idea of user grouping. In YANA, users are organized into groups with diverse interests, and interact with the recommender via interest-specific pseudo users. In this way, individual users' personal interest information is kept confidential from the server. They also proposed a set of SMC protocols and recommendation strategies to protect individual user privacy against other group members in the recommendation process. The idea adopted in YANA is similar to that of *k*-anonymity and *l*-diversity in private data publishing. However, in *k*-anonymity and *l*-diversity, user data are collected in the clear by a central server, which then anonymizes the collected data so as to prevent third parties from identifying individual users. In the context of private personalized recommendation, the server collecting user data is usually untrusted, so users should be grouped and maintained in a distributed and privacy-preserving manner.

Instead of adopting expensive cryptographic techniques, Polat and Du [12] proposed the application of data randomization techniques to protect user data. In particular, each user perturbs her personal data by adding random noise before sending them to the recommender. As the recommender only receives the noisy data, it is unable to obtain accurate information about users. As long as the number of users is sufficiently large, the recommender can still generate personalized recommendations based on the aggregate information, which can be obtained from the scrambled user data with decent accuracy. However, some studies [15–17] show that this proposed perturbation technique can still leak information about users' private data.

Shokri et al. [13] proposed a new obfuscation mechanism to obfuscate user-item connections from an untrusted server. In their system, they assume that there is a central recommender server storing users' online profiles and generating recommendations for users. A user also stores her own locally named offline profile. The recommender server generates personalized recommendations according to the online profiles only. Each user independently synchronizes her online profile with her offline profile when she updates recently rated items. For privacy consideration, in addition to the users' actual ratings, the user also adds other users' item information to her profile. In particular, the user arbitrarily selects some of her peers and certain information about the selected peers' offline profiles, and adds some of their items to her profile. Therefore, the user's actual profile is hidden from the server since the server does not know which items have been actually rated by the user. The proposed obfuscation mechanism may still reveal user's interests at a higher level. Meanwhile, the obfuscation mechanism incurs a tradeoff in recommendation accuracy.

Targeting practical scalability, Chow et al. [14] later presented a new privacy-preserving collaborative filtering system that has the ability to accommodate a large number of users. In their system, users are first clustered based on mutual similarity, through the proper use of the locality-sensitive hashing (LSH) technique. This is different from the work in Ref. [12] in which similar users are identified via noisy ratings. Next, personalized recommendations are generated based on the aggregate ratings of similar users in the same cluster. This design creates a primitive to cluster similar users in a privacy-preserving manner, by modifying the existing LSH techniques. Chow et al. [14] also injected artificial ratings during the process of privately generating personalized recommendations.

More specifically, their design comprises a clustering step and a recommendation step. In the clustering step, the recommender distributes an LSH algorithm to users. With this hashing function, the similarity between users can be measured by the matches between the LSH values of their ratings. To compute the LSH value, each user does some pre-processing, for example, by standardizing their ratings by subtracting the mean. The LSH value is then uploaded to the recommender. In the recommendation step, the recommender computes the average rating of all the users that share the same LSH value, without seeing them in the clear. To protect user privacy, Chow et al. [14] proposed having users send obfuscated ratings to the recommender, which can still allow the recommender to compute the average rating. In particular, obfuscation is conducted by adding artificial ratings for randomly selected movies. The ratio of the number of artificial ratings against true ratings is a tunable system parameter, which also decides the privacy-utility tradeoff. A higher ratio means stronger privacy, but lower accuracy.

3.2. Private machine learning-based approaches

Other works explore privacy-preserving machine learning-based recommendation. The basic idea in these works is to first

train a machine learning model over collected user data in a privacy-preserving manner, and then apply the model to generate personalized recommendations. The machine learning techniques usually adopted in these works include matrix factorization (MF) and ridge regression (RR). For privacy protection, these works typically rely on cryptographic techniques that include PHE, fully homomorphic encryption (FHE), and GCs. In what follows, we describe some representative works.

A design for a privacy-preserving MF for recommendation systems was proposed by Nikolaenko et al. [18]. MF techniques aim to learn item profiles and user profiles from user ratings. Given an item profile and a user profile (where either of these is a vector), a rating prediction can then be made by the inner product of vectors. In this design, the recommender collects encrypted user ratings and then runs a privacy-preserving MF protocol with a third party called the crypto-service provider (CSP). The security goal of this protocol is mainly to ensure that neither the recommender nor the CSP can learn user ratings. At the end of protocol execution, item profiles are produced, which can then be leveraged to make predictions for users on unrated items. Note that because exposing both item profiles and user profiles can easily violate user privacy, this protocol does not produce the user profiles at the recommender side.

To support privacy-preserving MF, Nikolaenko et al. [18] adopted a hybrid approach that combines PHE and GC. In this approach, user ratings are encrypted by a partially homomorphic cryptosystem under the public key of the CSP. After collecting all encrypted user ratings, the recommender adds random masks in the encrypted domain to the user ratings, and sends the resulting ciphertexts to the CSP. The CSP then decrypts the ciphertexts and obtains masked user ratings. The CSP subsequently prepares a GC that takes as input the garbled values of the masked user ratings and random masks. Inside the circuit, user ratings are first recovered by removing the masks from the masked user ratings; next, MF is performed. After running a protocol based on the built GC, the item profiles are produced.

After the item profiles are obtained, a viable way to produce private predictions is as follows: The recommender sends the item profiles to a user, who first recovers her own profile by solving an optimization problem. Having the item profiles and her own profile, the user can then produce local predictions for ratings on unrated items. In addition to this basic method, Nikolaenko et al. [18] proposed a mechanism to ask the recommender to generate predictions for users in a privacy-preserving manner.

As a follow-up work to improve the efficiency of the work in Ref. [18], Kim et al. [19] proposed an efficient privacy-preserving MF protocol based on FHE. They adopted an architecture similar to that in Ref. [18], in which a CSP is introduced to assist the recommender to perform MF, which is based on the gradient descent method. In this protocol, the CSP holds two key pairs: One is for a partially homomorphic cryptosystem and the other is for a fully homomorphic cryptosystem. Each user encrypts her ratings under the public key of the partially homomorphic cryptosystem and sends the ciphertexts to the recommender, which then runs a protocol with the CSP based on random masking to convert the partially homomorphic ciphertexts into fully homomorphic ciphertexts. Next, the recommender and the CSP jointly perform gradient descent-based MF to produce encrypted item profiles and user profiles, based on FHE and random masking. As gradient descent essentially requires the inner product of vectors, directly applying FHE is inefficient. Therefore, Kim et al. [19] introduced a novel data structure to enable the full use of slots supported by the single-instruction-multiple-data (SIMD) operation in FHE ciphertext. That is, it allows one homomorphic operation for multiple operations over vectors during gradient descent.

Considering that disclosing which items a user rated may leak personal information such as gender, Kim et al. [19] further enhanced their design by injecting fake ratings and then operating in the encrypted domain to remove the effects of fake ratings via the use of input indicators produced by users. Note that while Ref. [18] also mentioned this approach, it is not included in the design therein for the efficiency issue.

Unlike the work in Refs. [18] and [19], Nikolaenko et al. [20] studied another machine learning technique with privacy protection for recommendation: privacy-preserving RR. In RR-based recommendation, the recommender collects the preferences and ratings of many users for different items and runs a learning algorithm on the data. The learning algorithm generates a model that can be used to predict how a new user will rate certain items. Nikolaenko et al. [20] designed a protocol to enable the recommender to learn the model without seeing user data in the clear.

In this system architecture, a third-party CSP is introduced to cooperate with the recommender in order to complete the learning procedure. Similar to the work in Ref. [18], Nikolaenko et al. [20] designed the privacy-preserving RR protocol by combining homomorphic encryption with GC. The proposed protocol contains two phases. In the first phase, each user encrypts her data records under the public key of the CSP, and sends the ciphertexts to the recommender. The recommender then exploits the homomorphic addition property of homomorphic encryption to perform aggregation across different users' data. Such aggregation is viable as they re-formulate the RR problem, leading to significant reduction of the amount of data for further processing. In the second phase, the recommender adds random masks to the encrypted aggregate data in the encrypted domain, and sends the ciphertexts to the CSP. The CSP then performs decryption and obtains masked aggregated data. To find out the model securely, the CSP builds a GC, which takes as input the garbled values of the masked aggregate data and random masks. After the protocol based on the built GC is run, the model is produced and can be further used to make predictions for recommendation. This work does not specifically address how to securely apply the learned model to generate personalized recommendations.

As a follow-up work to improve the efficiency of the work in Ref. [20], Hu et al. [21] proposed a new privacy-preserving RR protocol, which is purely based on PHE and random masking, and which runs between the recommender and a third party. In order to achieve efficiency, they heavily leverage the packing-supported property of PHE, such as the Paillier cryptosystem. In particular, they designed a packed secure multiplication protocol that can compute the products of multiple pairs of private inputs simultaneously in the encrypted domain.

Based on this packed secure multiplication protocol, Hu et al. [21] further reformulated the RR problem into a problem of solving linear equations. They then proposed the adoption of either Gaussian elimination or the Jacobi iterative method to efficiently derive the learned model. In this design, the model is produced in the encrypted domain and the recommender is unable to obtain the model. After the encrypted model is obtained, the recommender sends it to users, who also receive decryption keys from the third party. The users can then decrypt the model and apply the model locally.

4. Privacy-preserving CB recommendation

CB recommendation systems recommend items based on the properties of items. Two common applications of CB recommendation services are targeted advertising and targeted coupon delivery. In targeted advertising, the ad network collects users' personal information and delivers matching ads to targeted users. In targeted coupon delivery, vendors intend to provide targeted coupons to certain users who are likely to become loyal routine

customers, based on the user's behavior profile. Below, we describe how to protect user privacy in each of these applications.

4.1. Private targeted advertising

Numerous works have proposed solutions to achieve targeted ad delivery and protect user's personal information. The privacy-protection mechanisms usually adopted in these works include local targeting [22,23], game theory [23], anonymization [24–26], cryptographic techniques [27], and obfuscation [28,29]. In what follows, we describe some representative works.

Toubiana et al. [22] proposed the adoption of the strategy of local targeting to simultaneously achieve behavioral targeting and privacy, and presented a private targeted advertising architecture named Adnostic. More specifically, Adnostic first pre-fetches a list of ads and stores them locally before the user visits publishers' pages. When the user browses a publisher's page, the ad network sends the list of n ads to the user. The user checks whether the most relevant ad has already existed in the pre-fetched list. If the selected ad is already stored locally, the browser displays it immediately, which leads to the speedup of page display. Otherwise, the user needs to download the selected ad from the ad network for display. Note that regardless of whether the selected ad has already been stored or not, Adnostic always downloads all the non-pre-fetched listed ads, so as to avoid information leakage to the ad network. In Adnostic, n is a configurable parameter; a larger value allows more accurate targeting, but leads to more consumption in network bandwidth, and vice versa. An appropriate value for n is 20, as suggested by Adnostic.

In Adnostic, there is a risk that none of the fetched ads will precisely match the user's preferences, because the ad network only delivers a small number of ads to the user. Therefore, to cover the wide spectrum of user's interests, at least one ad per segment for a given interest-segmentation system should be retrieved from the ad network. Their survey of existing online behavioral advertising systems found that the number of segments used in existing systems ranges from 25 to 100, which gives an upper bound on the number of transmitted ads.

Later, Wang et al. [23] proposed a design that also performs targeting on a user's local device; however, their focus is to provide incentives for users to click on ads that are interesting to them but may cause a potential privacy risk. In particular, they proposed a privacy-aware compensation framework to promote targeted advertising with control of privacy risks. In their framework, Wang et al. [23] considered that users, the ad broker, and advertisers are rational and selfish entities and that each of them only cares about their own interests. In order to encourage users to click on interesting ads, the ad network provides a certain amount of compensation for the users' privacy leakage from ad clicks. This can effectively improve the click-through rate and generate more revenue for both the ad network and advertisers. Wang et al. [23] further modeled the framework as a three-stage Stackelberg game, in which all the entities are considered to be selfish, and all have the goal of maximizing their own utilities by selecting optimal strategies. They obtained the Nash equilibrium by analyzing the cooperation and competition relationship among users, the ad network, and advertisers. In addition, Wang et al. [23] analyzed the competition between advertisers that share the whole market. Furthermore, they modeled the market-sharing scenario as a non-cooperative game and proved that the Nash equilibrium exists. In short, their proposed framework provides strong motivation for the ad network and advertisers to enforce the compensation policy in practice, and for users to embrace the service of targeted advertising.

Instead of performing targeting on a user's local device, Guha et al. [24] proposed a private advertising architecture called Privad, which involves an additional party called the dealer in order to protect user privacy. In Privad, in order to prevent the ad network from learning the identity of the client, the dealer anonymizes the communications between the client and the ad network. In order to prevent the dealer from learning a user's profiles, the communications between the client and the ad network are encrypted, and only the client and the ad network can decrypt the transmitted messages correctly. However, such an architecture has a potential limitation: The dealer needs to stay online all the time. This is not desirable, as a third party should ideally be infrequently contacted in practice. Later, Backes et al. [25] proposed a provably secure architecture called ObliviAd for privacy-preserving online behavioral advertising; this architecture leverages the technology of secure hardware-based private information retrieval (PIR). ObliviAd implements the PIR using oblivious random-access memory (ORAM) running on a secure co-processor (SC) that resides on the ad network side, which allows users to retrieve the ads that best match their profiles without leaking any private personal information. In particular, when a user visits a webpage, ObliviAd first sends the encrypted user profiles to the SC. The SC then securely searches ads and selects a subset of ads that best fit the user profiles based on the specified algorithm of the ad network. The searching and selection schemes are built on the ORAM protocol, which can prevent the ad network from learning any information about the selected ads. In order to support the case of multiple different ads per keyword, Backes et al. [25] further modified the used ORAM scheme. After that, SC delivers selected ads to the user in encrypted form.

Unlike the work in Refs. [24] and [25], Artail and Farhat [26] explored private non-local targeting without introducing an additional third party. Their architecture relies on cooperation among users to request and distribute ads to each other, and implements a shuffling algorithm to hide the interests of individual users from each other and their identity information from the ad network. First, in order to hide the identity information from the ad network when requesting ads, each user aggregates her interests based on a shuffling mechanism in an *ad hoc* network, and then sends them through one of the peer users. The selected peer user acts as a proxy to anonymously contact the ad network with the received aggregated interests. After receiving the ads, the designated user distributes them to the users in the same way they were collected, through an *ad hoc* connection. Second, since users may not trust each other, the aggregated interests should not leak private information to each other. In their design, Artail and Farhat [26] combined asymmetric cryptosystems and shuffling mechanisms to protect individual users' preferences against each other.

Instead of relying on user cooperation, Jiang et al. [27] leveraged the private stream-searching (PSS) technique to design another privacy-preserving targeted advertising system that can offer privacy protection and accurate targeting without introducing an additional third party. In this system, user profiles are inferred from the user's behavioral data and the sensor data on the local mobile device. A user uses the profile to construct an encrypted ad request, which is then sent to the untrusted ad network. The ad network processes the encrypted profiles over all the ads, and returns encrypted matching ads without knowing any underlying content. The user recovers matching ads using her private key. In this way, the ad network can deliver accurate matching ads to users with strong privacy protection. However, directly applying PSS to achieve secure and accurate mobile advertising results in serious practicality issues. That is, it leads to high computation and communication overhead for the resource-constrained mobile devices. Therefore, Jiang et al. [27] further proposed mechanisms to improve the computation and communication performance of the system. In particular, they proposed using a hierarchical structure to represent the user and ad preferences, and to leverage ads auctioning in the ad network. They also proposed encouraging

users to provide broad categories to narrow the search range of matched ads, which can save computation cost for the ad network. They also minimized ad-loading latency by utilizing the prefetching and caching mechanism, which can help to amortize the computation and communication cost.

Other works also exploit obfuscation techniques to enable private targeted advertising. In Ref. [28], Hardt and Nath pointed out that optimizing the following three design goals is hardly feasible in a targeted advertising system: privacy, efficiency, and ad relevance. Hence, they formulated an optimization problem for ad selection in targeted advertising systems. The proposed optimization problem includes three important variables. The first variable is privacy; that is, how much information a user shares with the ad network. The second is bandwidth efficiency; that is, how many ads are sent to the user. The third variable is utility; that is, the relevance of the ads that are delivered to users. More specifically, the proposed framework endows users with the rights to decide the amount of personal information to share with the ad network. According to the information collected from users, the ad network selects a set of ads for users, for which the communication overhead is bounded. Upon receiving the ads from the ad network, the local device of a user selects the most relevant one to display, by analyzing all the private information of that user. Therefore, their framework faces the following challenge: How to choose the set of delivered ads in a proper way such that the variable utility can be maximized while meeting the constraints on the variable privacy and bandwidth efficiency. Hardt and Nath [28] showed that it is a non-deterministic polynomial-time (NP)-hard problem to find the optimal set of ads, and thus employ approximation techniques to solve the optimization problem.

In an independent work, Davidson et al. [29] pointed out that personalization support should be provided by a unified system within the operation system (OS), rather than by individual apps. On the one hand, when collecting user preference information from all apps and OS interactions, it is possible to achieve considerably higher accuracy compared with what any specified app can obtain individually. On the other hand, the user generally trusts the OS, so the trusted computing base would not be expanded by performing personalization there. To advocates for OS support for user-side personalization, Davidson et al. [29] proposed an OS service named MoRePriv. To balance privacy and personalization, MoRePriv generalizes a user's personal profile to a coarse-grained profile, which limits the damage from potential leaks of private information. MoRePriv first captures several streams of (sensitive) information about the user within the OS, such as the user's e-mail, short message service (SMS), and more, which may indicate user preferences. MoRePriv then infers users' interests by parsing and classifying that information. MoRePriv provides a set of application programing interfaces (APIs) that endow third-party applications with limited access to the user profile.

4.2. Private targeted coupon delivery

Although the literature contains a considerable amount of research on private targeted advertising, little work has been done on secure targeted coupon delivery. Compared with private targeted advertising, targeted coupon delivery raises additional security challenges. First, it requires targeted coupons to be delivered to eligible users whose behavioral profiles accurately satisfy the vendors' targeting profile. This is to prevent coupon exploitation attacks. Second, throughout the targeted coupon-delivery procedure, non-eligible users should not learn anything about the vendor's targeting profile, except for their non-eligibility status. Otherwise, there is a risk that non-eligible users may try to exploit the information they learn to get targeted coupons [30–32]. In the literature, the security strategies/techniques for targeted coupon delivery include local targeting [30,31] and cryptographic techniques (i.e., PHE and GC) [32].

Partridge et al. [30] were the first to propose a privacy-preserving targeted coupon delivery framework called PiCoDa. PiCoDa adopts the strategy of local targeting to achieve the privacy protection of users. To verify a user's eligibility for targeted coupons, they resort to the LSH technique to test whether the user's behavioral profile approximately matches the vendor's targeting profile. As a result, if the LSH value of a user's behavioral profile is identical to that of the vendor's targeting profile, the user is able to derive the key for decrypting targeted coupons. The major limitation of this work is that a portion of non-eligible users may be able to obtain targeted coupons, due to the existence of false positives in LSH. This may violate the vendor's interests.

Rane and Uzun [31] later proposed another design for private targeted coupon delivery. They adopted the same security strategy as in Ref. [30] (i.e., local targeting). However, instead of adopting LSH to test a user's eligibility, they proposed the application of error-correcting codes to encode the user's behavioral profile and the vendor's targeting profile. This new mechanism frees the design from the problem of false positives. However, the tradeoff is that some information about the vendor's targeting profile is leaked to users. Jiang et al. [32] recently proposed a new design for private targeted coupon delivery, which combines the techniques of homomorphic encryption and GC. Their design ensures that targeted coupons are accurately delivered to eligible users only, while achieving user privacy and vendor protection.

5. Future research directions

(1) **Robustness against malicious users.** Most of the existing works on private personalized recommendation services assume that users honestly participate in the whole procedure. However, in practice, some users may be malicious and may intentionally provide invalid data to the recommender in order to disrupt the system. This could be a serious threat, and is largely underexplored in existing works on privacy-preserving personalized recommendation services. Defense against such a threat could be challenging, especially when the recommender receives only encrypted user data. To ensure the quality of service, it is imperative to develop verification techniques for privacy-preserving recommendation systems in order to provide robustness against data-faking attacks from malicious users. One possible direction is to leverage some cryptographic techniques such as digital signatures and commitments, thus forcing users to produce committed data.

(2) **Security against malicious recommendation.** An implicit assumption in most of the existing works on private recommendations is that the recommended items from the recommender are benign and will do no harm. However, some media-rich recommendation items (e.g., ads) on mobile devices contain JavaScript, image, or video. It is common for these items to access the external storage—a shared place where different mobile applications store their files. Some recent studies [33,34] have shown that these media-rich recommended items can be exploited to infer privacy-sensitive information such as gender and social circle. To defend against the emerging threats of malicious recommendation, it is critically important to provide an isolated execution environment on mobile devices and to endow it with the capability of satisfying the entire usage requirements of media-rich recommended items. In addition, most existing works that introduce a third party consider that third party to be semi-honest. However, such an assumption may not always hold true in practice; the third party may behave maliciously while assisting the recommender to perform private personalized recommendations. Developing effective

Table 1
A comparison of works on privacy-preserving personalized recommendation.

Category	Techniques	Additional party	Active user participation	Computation overhead	Inter-user interaction	User data protection	Accuracy loss	Refs.
NBCF	PHE, SMC	Yes	Yes	High	No	Strong	No	[6]
NBCF	PHE	No	No	High	No	Strong	No	[10]
NBCF	PHE	Yes	Yes	High	No	Strong	No	[11]
NBCF	SMC	No	Yes	High	Yes	Strong	No	[4]
NBCF	Perturbation	No	No	Low	No	Weak	Yes	[12]
NBCF	Obfuscation	No	Yes	Low	Yes	Weak	Yes	[13]
NBCF	LSH, artificial ratings	No	No	Low	No	Weak	Yes	[14]
MLBCF	MF, PHE, GC	Yes	No	High	No	Strong	No	[18]
MLBCF	MF, FHE	Yes	No	High	No	Strong	No	[19]
MLBCF	RR, PHE, GC	Yes	No	High	No	Strong	No	[20]
MLBCF	RR, PHE	Yes	No	High	No	Strong	No	[21]
CBRTA	Local targeting	No	No	Low	No	Strong	Yes	[22]
CBRTA	Local targeting, game theory	No	No	Low	No	Strong	Yes	[23]
CBRTA	Anonymization, PKE	Yes	No	High	No	Strong	No	[24]
CBRTA	ORAM	Yes	No	High	No	Strong	No	[25]
CBRTA	Anonymization, PKE	No	Yes	High	Yes	Strong	No	[26]
CBRTA	PSS	No	No	High	No	Strong	No	[27]
CBRTA	Obfuscation	Yes	No	Low	No	Weak	Yes	[28]
CBRTA	Obfuscation	No	No	Low	No	Weak	Yes	[29]
CBRTCD	LSH, local targeting	No	No	Low	No	Strong	Yes	[30]
CBRTCD	Fuzzy commitment	No	No	Low	No	Strong	No	[31]
CBRTCD	PHE, GC	Yes	No	High	No	Strong	No	[32]

NBCF: neighborhood-based collaborative filtering; MLBCF: machine learning-based collaborative filtering; CBRTA: content-based recommendation for targeted advertising; CBRTCD: content-based recommendation for targeted coupon delivery; PHE: partially homomorphic encryption; FHE: fully homomorphic encryption; GC: garbled circuit; MF: matrix factorization; RR: ridge regression; SMC: secure multiparty computation; PKE: public-key encryption; LSH: locality-sensitive hashing; ORAM: oblivious random-access memory; PSS: private stream searching.
The indicator "Active user participation" is marked "Yes" if more than one round of user-server interaction is needed upon a request for private recommendations or if inter-user interaction is required. The indicator "User data protection" is marked "Strong" if user data are held locally or protected via cryptographic techniques. The indicator "Computation overhead" is marked "High" if cryptographic operations are involved.

mechanisms to defend against such a malicious adversary is also crucially important.

(3) **Mobile-oriented cost efficiency.** Along with the ubiquity of mobile devices, mobile personalized recommendation services are becoming increasingly prevalent. However, mobile devices are usually resource-constrained, especially in terms of battery and bandwidth. In order to achieve a satisfying user experience in mobile personalized recommendation services while preserving user privacy, a security design should impose lightweight cost on mobile devices. Therefore, it is worthwhile to continue the line of research of developing lightweight privacy-preserving techniques for mobile devices in recommendation services. One possible direction is to integrate techniques such as user network traffic analysis [35] and social networking [36] in order to propose new innovative privacy-preserving solutions. Another direction may be to design lightweight privacy-preserving solutions by leveraging advanced trusted hardware such as Intel® Software Guard Extensions (Intel® SGX), which has the potential to provide secure computations at minimal performance overhead. In addition, when conducting a performance evaluation of a security design for private mobile personalized recommendation services, it is necessary to consider the energy cost on mobile devices, rather than limiting the evaluation to computation and communication costs. This is because battery energy is one of the most precious resources of mobile devices, and is a practical factor that directly affects user experience.

6. Conclusions

In this paper, we surveyed the literature related to privacy-preserving personalized recommendation services. We first presented the system architecture of personalized recommendation services, commonly adopted recommendation techniques, and privacy issues posed by personalized recommendation services. We then described existing privacy-preserving techniques for personalized recommendation services, which are classified into two broad categories: privacy-preserving CFB recommendation and privacy-preserving CB recommendation. We then further classified privacy-preserving CFB recommendation into private neighborhood-based approaches and private machine learning-based approaches, and further classified privacy-preserving CB recommendation into private targeted advertising and private targeted coupon delivery. A comparison of the described existing works on privacy-preserving recommendation is summarized in Table 1 [4,6,10–14,18–32]. Finally, we provided some discussion on future research directions.

Acknowledgements

This work was supported in part by the Research Grants Council of Hong Kong (CityU 11276816, CityU 11212717, and CityU C1008-16G), the Innovation and Technology Commission of Hong Kong (ITS/168/17), and the National Natural Science Foundation of China (61572412 and 61772236).

Compliance with ethics guidelines

Cong Wang, Yifeng Zheng, Jinghua Jiang, and Kui Ren declare that they have no conflict of interest or financial conflicts to disclose.

References

[1] Mackenzie I, Meyer C, Noble S. How retailers can keep up with consumers [Internet]. Chicago: McKinsey & Company; 1996–2018 [cited 2017 Mar 22]. Available from: http://www.mckinsey.com/industries/retail/our-insights/how-retailers-can-keep-up-with-consumers.
[2] Hassler J. The power of personalized product recommendations [Internet]. Carrollton: Intelliverse; 2018 [updated 2017 Aug 22; cited 2017

Mar 22]. Available from: http://www.intelliverse.com/blog/the-power-of-personalized-product-recommendations/.
[3] Leskovec J, Rajaraman A, Ullman J. Mining of massive datasets. 2nd ed. Cambridge: Cambridge University Press; 2014.
[4] Li D, Lv Q, Shang L, Gu N. Efficient privacy-preserving content recommendation for online social communities. Neurocomputing 2017;219:440–54.
[5] Ramakrishnan N, Keller BJ, Mirza BJ, Grama AY, Karypis G. Privacy risks in recommender systems. IEEE Internet Comput 2001;5(6):54–62.
[6] Erkin Z, Veugen T, Toft T, Lagendijk RL. Generating private recommendations efficiently using homomorphic encryption and data packing. IEEE Trans Inf Foren Sec 2012;7(3):1053–66.
[7] Xu K, Yan Z. Privacy protection in mobile recommender systems: A survey. In: Wang G, Ray I, Alcaraz Calero J, Thampi S, editors Proceedings of the 9th International Conference on Security, Privacy and Anonymity in Computation, Communication, and Storage; 2016 Nov 16–18; Zhangjiajie, China. Cham: Springer; 2016. p. 305–18.
[8] Koren Y, Bell R, Volinsky C. Matrix factorization techniques for recommender systems. Computer 2009;42(8):30–7.
[9] Aïmeur E, Brassard G, Fernandez JM, Onana FSM. Alambic: A privacy-preserving recommender system for electronic commerce. Int J Inf Secur 2008;7(5):307–34.
[10] Badsha S, Yi X, Khalil I. A practical privacy-preserving recommender system. Data Sci Eng 2016;1(3):161–77.
[11] Badsha S, Yi X, Khalil I, Bertino E. Privacy preserving user-based recommender system. In: Proceedings of the 37th IEEE International Conference on Distributed Computing Systems; 2017 Jun 5–8; Atlanta, GA, USA. Los Alamitos: IEEE Computer Society Press; 2017. p. 1074–83.
[12] Polat H, Du W. Privacy-preserving collaborative filtering using randomized perturbation techniques. In: Proceedings of the 3rd IEEE International Conference on Data Mining; 2003 Nov 19–22; Melbourne, FL, USA. Los Alamitos: IEEE Computer Society Press; 2003. p. 625–8.
[13] Shokri R, Pedarsani P, Theodorakopoulos G, Hubaux JP. Preserving privacy in collaborative filtering through distributed aggregation of offline profiles. In: Proceedings of the 3rd ACM Conference on Recommender Systems; 2009 Oct 23–25; New York, NY, USA. New York: Association for Computing Machinery, Inc.; 2009. p. 157–64.
[14] Chow R, Pathak MA, Wang C. A practical system for privacy-preserving collaborative filtering. In: Proceedings of the 12th IEEE International Conference on Data Mining Workshops; 2012 Dec 10; Brussels, Belgium. Los Alamitos: IEEE Computer Society Press; 2012. p. 547–54.
[15] Huang Z, Du W, Chen B. Deriving private information from randomized data. In: Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data; 2005 Jun 14–16; Baltimore, MD, USA. New York: Association for Computing Machinery, Inc.; 2005. p. 37–48.
[16] Zhang S, Ford J, Makedon F. Deriving private information from randomly perturbed ratings. In: Ghosh J, Lambert D, Skillicorn D, Srivastava J, editors Proceedings of the 2006 SIAM International Conference on Data Mining; 2006 Apr 20–22; Bethesda, MD, USA. Philadelphia: Society for Industrial and Applied Mathematics; 2006. p. 59–69.
[17] Aggarwal CC. On randomization, public information and the curse of dimensionality. In: Proceedings of the 23rd International Conference on Data Engineering; 2007 Apr 15–20; Istanbul, Turkey. Los Alamitos: IEEE Computer Society Press; 2007. p. 136–45.
[18] Nikolaenko V, Ioannidis S, Weinsberg U, Joye M, Taft N, Boneh D. Privacy-preserving matrix factorization. In: Proceedings of the 2013 ACM SIGSAC Conference on Computer and Communications Security; 2013 Nov 4–8; Berlin, Germany. New York: Association for Computing Machinery, Inc.; 2013. p. 801–12.
[19] Kim S, Kim J, Koo D, Kim Y, Yoon H, Shin J. Efficient privacy-preserving matrix factorization via fully homomorphic encryption: Extended abstract. In: Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security; 2016 May 30–Jun 3; Xi'an, China. New York: Association for Computing Machinery, Inc.; 2016. p. 617–28.
[20] Nikolaenko V, Weinsberg U, Ioannidis S, Joye M, Boneh D, Taft N. Privacy-preserving ridge regression on hundreds of millions of records. In: Proceeding of the 2013 IEEE Symposium on Security and Privacy; 2013 May 19–22; Berkeley, CA, USA. Los Alamitos: IEEE Computer Society Press; 2013. p. 334–48.
[21] Hu S, Wang Q, Wang J, Chow SSM, Zou Q. Securing fast learning! Ridge regression over encrypted big data. In: Proceedings of the 2016 IEEE Trustcom/BigDataSE/ISPA; 2016 Aug 23–26; Tianjin, China. Los Alamitos: IEEE Computer Society Press; 2016. p. 19–26.
[22] Toubiana V, Narayanan A, Boneh D, Nissenbaum H, Barocas S. Adnostic: Privacy preserving targeted advertising. In: Proceedings of the 2010 Network and Distributed System Security Symposium; 2010 Feb 28–Mar 3; San Diego, CA, USA. Reston: Internet Society; 2010. p. 1–16.
[23] Wang W, Yang L, Chen Y, Zhang Q. A privacy-aware framework for targeted advertising. Comput Netw 2015;79:17–29.
[24] Guha S, Cheng B, Francis P. Privad: Practical privacy in online advertising. In: Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation; 2011 Mar 30–Apr 1; Boston, MA, USA. Berkeley: USENIX Association; 2011. p. 169–82.
[25] Backes M, Kate A, Maffei M, Pecina K. ObliviAd: Provably secure and practical online behavioral advertising. In: Proceedings of the 2012 IEEE Symposium on Security and Privacy; 2012 May 20–23; San Francisco, CA, USA. Los Alamitos: IEEE Computer Society Press; 2012. p. 257–71.
[26] Artail H, Farhat R. A privacy-preserving framework for managing mobile ad requests and billing information. IEEE Trans Mobile Comput 2015;14(8):1560–72.
[27] Jiang J, Gui X, Shi Z, Yuan X, Wang C. Towards secure and practical targeted mobile advertising. In: Proceedings of the 11th International Conference on Mobile Ad-hoc and Sensor Networks; 2015 Dec 16–18; Shenzhen, China. Los Alamitos: IEEE Computer Society Press; 2015. p. 79–88.
[28] Hardt M, Nath S. Privacy-aware personalization for mobile advertising. In: Proceedings of the 2012 ACM Conference on Computer and Communications Security; 2012 Oct 16–18; Raleigh, NC, USA. New York: Association for Computing Machinery, Inc.; 2012. p. 662–73.
[29] Davidson D, Fredrikson M, Livshits B. MoRePriv: Mobile OS support for application personalization and privacy. In: Proceedings of the 30th Annual Computer Security Applications Conference; 2014 Dec 8–12; New Orleans, LA, USA. New York: Association for Computing Machinery, Inc.; 2014. p. 236–45.
[30] Partridge K, Pathak MA, Uzun E, Wang C. PiCoDa: Privacy-preserving smart coupon delivery architecture. In: Proceedings of 5th Workshop on Hot Topics in Privacy Enhancing Technologies; 2012 Jul 11–13; Vigo, Spain; 2012. p. 95–108.
[31] Rane S, Uzun E. A fuzzy commitment approach to privacy preserving behavioral targeting. In: Proceedings of the ACM MobiCom Workshop on Security and Privacy in Mobile Environments; 2014 Sep 11; Maui, HI, USA. New York: Association for Computing Machinery, Inc.; 2014. p. 31–6.
[32] Jiang J, Zheng Y, Yuan X, Shi Z, Gui X, Wang C, et al. Towards secure and accurate targeted mobile coupon delivery. IEEE Access 2016;4:8116–26.
[33] Wu D, Chang RKC. Analyzing Android browser apps for file:// vulnerabilities. In: Chow SSM, Camenisch J, Hui LCK, Yiu SM, editors Proceedings of the 17th International Conference on Information Security; 2014 Oct 12–14; Hong Kong, China. Cham: Springer; 2014. p. 345–63.
[34] Son S, Kim D, Shmatikov V. What mobile ads know about mobile users. In: Proceedings of 2016 Network and Distributed System Security Symposium; 2016 Feb 21–24; San Diego, CA, USA. Reston: Internet Society; 2016. p. 1–14.
[35] Su X, Zhang D, Li W, Li W. Android app recommendation approach based on network traffic measurement and analysis. In: Proceedings of the IEEE Symposium on Computers and Communication; 2015 Jul 6–9; Larnaca, Cyprus. Piscataway: Institute of Electrical and Electronic Engineers, Inc.; 2015. p. 988–94.
[36] Li F, He Y, Niu B, Li H, Wang H. Match-MORE: An efficient private matching scheme using friends-of-friends' recommendation. In: Proceedings of the 2016 International Conference on Computing, Networking and Communications; 2016 Feb 15–18; Kauai, HI, USA. Piscataway: Institute of Electrical and Electronic Engineers, Inc.; 2016. p. 1–6.