Private Movie Recommendations for Children

Anh Pham, Mohammad Samragh, Sameer Wagh, Emily Willson

Microsoft Private AI Bootcamp 2019

1 Introduction

Data-driven business models such as recommender systems (Netflix, Pandora) and targeted advertising (Facebook, Google) rely strictly on consumer data and the information they contain on individuals' behavioral patterns and preferences. This reliance effectively opens door to the longstanding conflict of privacy versus convenience: as customers expect the rendered goods to be content-relevant to their specific needs, a certain degree of user data exploitation by service providers is mandatory. At the same time, the mounting number of data collection and data use breaches has prompted the public to grow hostile towards the tech sector; a recent case is the $170 million fine imposed in September 2019 on Google and YouTube Kid for violating federal requirements for child privacy protection [3]. Homomorphic encryption offers a solution to this pressing problem: the fully homomorphic encryption (FHE) scheme can be used to construct a private recommender system with which user data is not exposed to service providers in their raw form, and only data "masked" by encryption are sent to providers for recommendation inferencing. In this project, we construct a general framework for a FHE-backed recommender that can be extended to different applications, with the particular use case of YouTube Kid as a proof-of-concept.

To this end, we propose a *private video recommendation system*, appropriate for a platform like YouTube Kids. This system, built on fully homomorphic encryption, would allow the platform to make tailored recommendations to children without exposing their private information. In this writeup, we briefly describe the motivation and construction of our system. We then discuss the novelty, soundness, feasibility, and impact of such a system.

2 Background

The Child Online Protection Privacy Act (COPPA) was passed in the US in 1999 and significantly revised in 2011 [1]. It provides legal protection for children's activity and data shared online. Among its many requirements, it mandates that online provider "...establish and maintain reasonable procedures to protect the confidentiality, security, and integrity of the personal information collected from children under age 13". The recent legal action against Google is a sober reminder that such privacy-protection procedures are often not in place. Nonetheless, the fact remains that a recommender system requires collecting of personal data.

FHE implementation within the recommender pipeline can resolve the tension between privacy preservation and pattern mining over personal data. FHE allows the ability to compute on encrypted data and yields results that, when decrypted, would match the computation as if it has been done on plaintexts. This capacity enables recommenders to execute their algorithms while respecting the privacy of users.

Prior work on private recommendation algorithms informs our prototype of a private video recommendation system for children [2].

3 Proposed Implementation

The abstract design of our recommender system is shown in Figure 1. In this setting, the server wishes to utilize the client's confidential feature vectors $x \in \mathbb{R}^n$ and provide a proper recommendation. The protocol involves the following steps: (1) The client encrypts its private message $x \rightarrow \mathsf{Enc}_{\mathsf{pk}}(x)$. (2) The server receives the encrypted message and computes $f(\mathsf{Enc}_{\mathsf{pk}}(x))$ homomorphically. During this process, no information about input x or the analysis result is revealed to the server since only the client has the decryption keys. (3) The client decrypts $f(\mathsf{Enc}_{\mathsf{pk}}(x)) \rightarrow f(x)$ using her secret key and retrieves the recommendation.

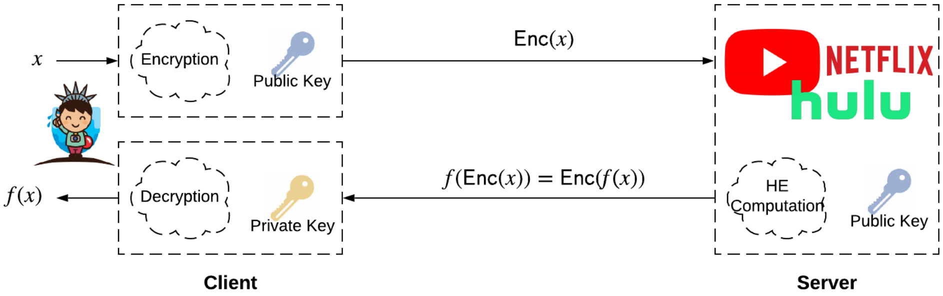

Fig. 1. Our proposed design for a private recommender system based on homomorphic encryption.

In our prototype, the training of the recommender model is done on public ratings from n clients, i.e., $x_1, \dots, x_n$. The data matrix $X = [x_1|x_2 \dots |x_n]$ is formed by stacking the data from the users. The element located at the i-th row and j-th column, $X_{i,j}$, is the rating that the j-th user provides for the i-th movie. For recommendation, we utilize Content Based Filtering (CBF) and Collaborative Filtering (CF) which are widely adopted in recommendation systems. The key steps are as follows:

Offline/training Phase. The server computes a similarity matrix $S \in \mathbb{R}^{m\times m}$ which will later be used for making recommendation. Formally, the similarity matrix is computed as $S = \frac{X \cdot X^T}{W \cdot W^T}$, where $X \in \mathbb{R}^{m\times n}$ is the data matrix and $W \in \mathbb{R}^m$ is the norm of ratings computed as $W = \sqrt{\sum_{i=1}^{n}(x_i^2)}$. For this phase of the prototype, we perform the computation of S in plain text. However, this step can be done homomorphically to enhance privacy.

Online/inference phase. Depending on the underlying data, the server either uses CBF or CF for recommendation making. For a feature vector x, the server computes $f(x)$ as follows:

$$f(x) = \begin{cases} \frac{S\cdot x}{Y} & CBF \\ A - \frac{S\cdot A}{Y} + \frac{S\cdot x}{Y} & CF \end{cases} \tag{1}$$

where $Y \in \mathbb{R}^m$ is achieved by summing the columns of S and $A \in \mathbb{R}^m$ is the average rating for each movie. Note that the term $A - \frac{S\cdot A}{Y}$ in the CF method is a constant and can be computed offline. In both CF and CBF methods, the core computation involving user's data is $\frac{S\cdot x}{Y}$. To ensure user privacy, this step should be done homomorphically. Figure 2 summarizes the required HE-based operations. In the next section, we elaborate on the design of the homomorphic encryption portion of our application.

4 HE Technical Details

We choose the non-interactive model to be fairly optimal in terms of communication. The total communication is the size of two ciphertexts which is $2 \cdot \log q \cdot N \approx 364$ KB each direction, where q is the ciphertext modulus and N is the degree of the cyclotoimic polynomial.

In terms of computation, the server performs matrix-vector multiplication with the vector encrypted and the matrix in plaintext. This requires Galois keys from the user and consumes

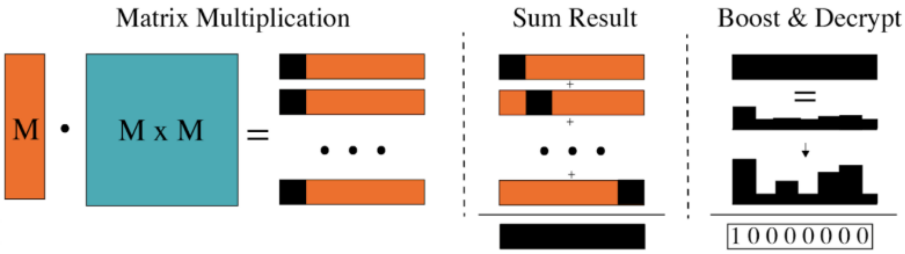

Fig. 2. Our proposed design for a private recommender system based on homomorphic encryption. The underlying algorithm is described in Section 4. We use boosting to increase the weight of the "high" recommendations and reduce that of "low" recommendations.

one depth of computation. Further, to compress all the products into one ciphertext, we use a sequential masking (all slots encode the dot product so on further rotation required) and then add these up into one ciphertext. This consumes one depth of computation.

The boosting is a simple observation that for a given vector $\{a_1, a_2, \ldots, a_n\}$, computing the vectors $\{a_1^k, \ldots a_n^k\}$ for a given k will increase the separation between the top values of a_i from the rest. In our work, we set $k = 2$ and hence require 2 depths of HE computations.

The performance bottleneck is the matrix-vector multiplication which we parallelize using 8 cores. Overall, we require parameters to be set that support at least depth 4 computation. This puts a lower bound on the degree of the cyclotomic polynomail N. We need to use at least $N = 8192$. This gives us about 218-bits of space for the ciphertext modulus. For performance reasons, we use these values and enable computations over a movie corpus of upto 2048 movies. The keys are stored locally on the end devices and the Galois keys are communicated in an offline/set-up phase.

5 Discussion

In this section, we evaluate our solution under the the proposed comparison metrics.

Novelty While previous work has designed a private recommendation system (see [2]), current recommender systems do not incorporate privacy-preserving computations on user data. our work designs and tests a *practical* and low-overhead implementation of such a system. Furthermore, the implementation we propose can extend to domains beyond video recommendation. This opens the door to a wide variety of new applications at the nexus of homomorphic encryption and machine learning.

Soundness The security of our system is inherent due to the underlying homomorphic encryption scheme. The polynomial degree and ciphertext modulus are set to prevent information leakage, as discussed in Section 4. In this short article we only discuss non-interactive HE-based scenarios where the functionality of the recommender system is solely rendered by linear operations. Designing interactive protocols that support more nonlinear operations is a promising future direction. In fact, it may be reasonable to assume an interactive user for video recommendation systems since the user is present during the process. Such interactive protocols will allow for private evaluation of more complex machine learning models, e.g., deep neural networks, which can provide a higher quality of service.

Feasibility Section 4 describes the technical details of our implementation. The size of the recommendation matrix scales linearly with the number of users. The size of the matrix may become unwieldy with very large numbers of users, so we recommend optimization of this matrix computation as future work.

Impact Our work creates alternative ways for companies like YouTube to recommend content to users without violating their privacy. As data-sharing scandals continue to surface, privacy-centric systems become increasingly important.

References

1. Children's Online Privacy Protection Act. https://en.wikipedia.org/wiki/Children's_Online_Privacy_Protection_Act, 2019.
2. Badsha, S., Yi, X., and Khalil, I. A practical privacy-preserving recommender system. *Data Science and Engineering 1*, 3 (2016), 161–177.
3. FTC. Google and YouTube Will Pay Record $170 Million for Alleged Violations of Children's Privacy Law. https://www.ftc.gov/news-events/press-releases/2019/09/google-youtube-will-pay-record-170-million-alleged-violations, 2019.