bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

RESEARCH

Privacy-preserving search for chemical compound databases

Kana Shimizu[1*], Koji Nuida[1,2], Hiromi Arai[3], Shigeo Mitsunari[4], Nuttapong Attrapadung[2], Michiaki Hamada[5], Koji Tsuda[6], Takatsugu Hirokawa[7], Jun Sakuma[8], Goichiro Hanaoka[2] and Kiyoshi Asai[6]

*Correspondence: shimizu-kana@aist.go.jp
[1]Computational Biology Research Center, National Institute of Advanced Industrial Science and Technology (AIST), 2-4-7 Aomi Koto-ku, Tokyo, Japan
Full list of author information is available at the end of the article

Abstract

Background: Searching for similar compounds in a database is the most important process for in-silico drug screening. Since a query compound is an important starting point for the new drug, a query holder, who is afraid of the query being monitored by the database server, usually downloads all the records in the database and uses them in a closed network. However, a serious dilemma arises when the database holder also wants to output no information except for the search results, and such a dilemma prevents the use of many important data resources.

Results: In order to overcome this dilemma, we developed a novel cryptographic protocol that enables database searching while keeping both the query holder's privacy and database holder's privacy. Generally, the application of cryptographic techniques to practical problems is difficult because versatile techniques are computationally expensive while computationally inexpensive techniques can perform only trivial computation tasks. In this study, our protocol is successfully built only from an additive-homomorphic cryptosystem, which allows only addition performed on encrypted values but is computationally efficient compared with versatile techniques such as general purpose multi-party computation. In an experiment searching ChEMBL, which consists of more than 1,200,000 compounds, the proposed method was 36,900 times faster in CPU time and 12,000 times as efficient in communication size compared with general purpose multi-party computation.

Conclusion: We proposed a novel privacy-preserving protocol for searching chemical compound databases. The proposed method, easily scaling for large-scale databases, may help to accelerate drug discovery research by making full use of unused but valuable data that includes sensitive information.

Keywords: Chemical Compound; Similarity Search; Privacy Preserving Data Mining; Tversky Index; Additive Homomorphic Cryptosystem

Introduction

In recent years, the increasing cost of drug development and decreasing number of new chemical entities have become growing concerns [1]. One of the most popular approaches for overcoming these problems is searching for similar compounds in databases [2]. In order to improve the efficiency of this task, it is important to utilize as many data resources as possible. However, the following dilemma prevents the use of many existing data resources. Unpublished experimental results have been accumulated at many research sites, and such data has scientific value [3]. Since data holders are usually afraid of sensitive information leaking from the data resources,

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

they do not want to release the full data, but they might allow authorized users to search the data as long as the users obtain only search results from which they cannot infer sensitive information. Likewise, private databases of industrial research might be made available if the sensitive information were sufficiently protected. On the other hand, query compounds are also sensitive information for the users, and thus the users usually avoid sending queries and want to download all of the data in order to conduct search tasks on their local computers. In short, we cannot utilize important data resources because both the data holder and the data user insist on their privacy. Therefore, an emerging issue is to develop novel technology that enables privacy-preserving similarity searches. We show several use cases in the next section.

Let us start by clarifying privacy problems in database searches. In a database search, two types of privacy are of concern: *"user privacy"* (also known as input privacy) and *"database privacy"* (also known as output privacy). The first is equal to protecting the user's query from being leaked to others including the database holder. The second is equal to protecting the database contents from being leaked to others including the database user, except for the search results held by the user. Here we firstly consider the case of using no privacy-preserving techniques; namely, the user sends a plain query to the server and the server sends the search result. In this case, the user's query is fully obtained by the server. On the database side, the server's data is not directly leaked to the user. However, there is a potential risk that the user may infer the database contents from the search results. To protect user privacy, a scheme called single-database private information retrieval (PIR) has been proposed [4]. The simplest method for achieving PIR is that the user downloads all the contents of the database and searches on his/her local computer. Since this naive approach needs a huge communication size, several cryptographic techniques have been developed, in which the query is safely encrypted/randomized in the user's computer and the database conducts the search without seeing the query. Although PIR is useful for searching public databases, it does not suit the purpose of searching private databases because of the lack of database privacy. Likewise, similarity evaluation protocols keep user privacy [5–7] but they do not sufficiently protect database privacy because the server directly outputs similarity scores that become important hints for inferring database contents.

Generally speaking, it is very difficult to keep both user privacy and database privacy, because the database side must prevent various attacks without seeing the user's query. Among them, the following two attacks are major concerns.

- Regression attack
  Given one data point, the similarity between a target and the data point becomes a strong hint for detecting the target. The accuracy of the detection increases as the number of given data points becomes larger. In fact, a protocol that is not suitably designed may lead to even a small number of queries enabling the database user to detect the target. For example, when the server returns the exact distance between a query and a database entry, the range of the entry is rapidly narrowed as the number of queries increases, and the entry is finally detected uniquely by only almost the same number of queries as the dimension of the entry. For example, in the case of using the MACCS

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

keys, which are 166 bit structural key descriptors and often used for representing chemical compounds, a database entry is detected by sending only 166 queries. Therefore, it is necessary for the server to return the minimum information that is sufficient for the purpose of the search (see Fig. 1 for a detailed explanation).

- Illegal query attack
  Searching with an illegal query often causes unexpected server behaviour. In such a case, the server might return unexpected results that include important server information. To prevent this, the server should ensure the correctness of the user's query.

A schematic view of the privacy-preserving database search problems discussed here is shown in Fig. 2.

In the field of cryptography, there have been studies of versatile techniques such as general purpose multi-party computation (GP-MPC) [8] and fully homomorphic encryption (FHE) [9], which enable the design of systems that maintain both user privacy and database privacy. However, these techniques require huge computational costs as well as intensive communications between the parties, so they are scarcely used in practical applications. In order to avoid using such techniques, a similarity search protocol using a trusted third party [10] and a privacy preserving SQL database using a trusted proxy server [11] have been proposed, but those methods assure privacy only when the third party does not collude with the user or the server, which is not convenient for many real problems. As far as we know, no practical method has been proposed despite the great importance of privacy-preserving similarity searching. To overcome this lack, we propose a novel privacy-preserving similarity search method that can strongly protect database privacy as well as user privacy while keeping a significantly low computational cost and small communication size.

The rest of this paper is organized as follows. In the next section, we summarize our achievements in this study. This is followed by the Cryptographic background section and the Method section, where we define the problem and introduce details of the proposed protocol. In the Security analyses section, both the user privacy and database privacy of the proposed protocol are discussed in detail. In the Performance evaluation section, the central processing unit (CPU) time and communication size of the proposed protocol are evaluated for two datasets extracted from ChEMBL. Finally, we present our conclusions for this study in the Conclusion section.

Our Achievements

Here we focus on similarity search with the Tversky index of fingerprints, which is the most popular approach for chemical compound searches [12] and is used for various search problems in bioinformatics. To provide a concrete application, we address the problem of counting the number of similar compounds in a database, which solves various problems appearing in chemical compound searches. The following model describes the proposed method.

Model 1 *The user is a private chemical compound holder, and the server is a private database holder. The user learns nothing but the number of similar compounds*

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

in the server's database, and the server learns nothing about the user's query compound.

Here we introduce only a small fraction of the many scientific or industrial problems solved by Model 1.

1. Secure pre-purchase inspection service for chemical compound.
   When a client considers the purchase of a commercial database such as a focused library [13], he/she wants to check whether the database includes a sufficient number of similar compounds, without sending his/her private query, but the server does not allow downloading of the database.
2. Secure patent compound filtering.
   When a client finds a new compound, he/she usually wants to know whether it infringes on competitors' patents by searching the database of patent-protected compounds maintained by third parties. The same problem occurs when the client wants to check whether or not the compound is undesirable.
3. Secure negative results check.
   It is a common perception that current scientific publication is strongly biased against negative results [3], although a recent study showed statistically that negative results brought meaningful benefit [14]. Since researchers are reluctant to provide negative results, which often include sensitive information, a privacy-preserving system for sharing those results would greatly contribute to reducing redundant efforts for similar research topics. For example, it would be useful to have a system that allows a user to check whether the query is similar to failed compounds that have previously been examined in other laboratories.

In this study, we propose a novel protocol called the **secure similar compounds counter** (SSCC) which achieves Model 1. The first main achievement of this study is that SSCC is remarkably tolerant against regression attacks compared with existing protocols which directly output the similarity score. Moreover, we propose an efficient method for protecting the database from illegal query attacks. These points are discussed in the Security analyses section.

The second main achievement is that SSCC is significantly efficient both in computational cost and communication size. We carefully designed the protocol such that it uses only an additive-homomorphic cryptosystem, which is computationally efficient, and does not rely on any time-consuming cryptographic methods such as GP-MPC or FHE. Hence the performance of the protocol is sufficiently high for a large-scale database such as ChEMBL [15], as is shown in the Performance evaluation section.

Cryptographic background

Additively homomorphic encryption scheme

In this paper, we use an additive-homomorphic cryptosystem to design our protocol. The key feature of the additive-homomorphic cryptosystem is that it enables to perform additive operations on *encrypted* values. Therefore, intuitively, any standard computation algorithm can be converted into the privacy-preserving computation algorithm, if operations used in the standard algorithm can be replaced by additions.

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

More formally, we use a public-key encryption scheme ($\mathsf{KeyGen}; \mathsf{Enc}; \mathsf{Dec}$), which is semantically secure; that is, an encryption result (ciphertext) leaks no information about the original message (plaintext) [16]. Here, KeyGen is a key generation algorithm for selecting a pair $(\mathsf{pk}, \mathsf{sk})$ of a public key pk and a secret key sk; $\mathsf{Enc}(m)$ denotes a ciphertext obtained by encrypting message m under the given pk; and $\mathsf{Dec}(c)$ denotes the decryption result of ciphertext c under the given sk. We also require the following additive-homomorphic properties:

- Given two ciphertexts $\mathsf{Enc}(m_1)$ and $\mathsf{Enc}(m_2)$ of messages m_1 and m_2, $\mathsf{Enc}(m_1 + m_2)$ can be computed without knowing m_1, m_2 and the secret key (denoted by $\mathsf{Enc}(m_1) \oplus \mathsf{Enc}(m_2)$).
- Given a ciphertext $\mathsf{Enc}(m)$ of a message m and an integer e, $\mathsf{Enc}(e \cdot m)$ can be computed without knowing m and the secret key (denoted by $e \otimes \mathsf{Enc}(m)$).

For example, we can use either the Paillier cryptosystem [17] or the "lifted" version of the ElGamal cryptosystem [18] as such an encryption scheme; now the second operation $\otimes$ can be achieved by repeating the first operation $\oplus$. We notice that the range of plaintexts for those cryptosystems can be naturally set as an integer interval $[-N_1, N_2]$ for some sufficiently large $N_1, N_2 > 0$; therefore, the plaintexts are divided into positive ones, negative ones, and zero.

Non-interactive zero-knowledge proof

Below, we discuss the following situation: A user (a prover) wants to make a server (a verifier) convinced that a ciphertext c generated by the user corresponds to a message m in $\{0, 1\}$, but does not want to reveal any information about which of 0 and 1 is m. This can be achieved by using a cryptographic tool called *non-interactive zero-knowledge* (*NIZK*) *proof*. In the present case, it enables the user to generate a "proof" associated with c, so that:

- If m is indeed in $\{0, 1\}$, then the server can verify this fact by testing the proof (without knowing m itself).
- If $m \notin \{0, 1\}$, then the user cannot generate a proof that passes the server's test.
- The server cannot obtain any information about m from the proof, except for the fact that $m \in \{0, 1\}$.

(See [19] for a general formulation.) Besides the existing general-purpose NIZK proofs, Sakai et al. [20] proposed an efficient scheme specific to the "lifted" ElGamal cryptosystem, which we use below.

Method

The goal of this study is to design a protocol between a user and a server that enables the user to obtain the number of compounds in the server's database that are similar to the user's target compound. Here, a fingerprint of compound is modeled as $\vec{p} \in \{0, 1\}^\ell$ (*i.e.,* a bit string of length ℓ). An equivalent way to refer to $\vec{p}$ is the set of all indices i where $p_i = 1$. We denote such a set by $\boldsymbol{p}$. The similarity of two compounds $\boldsymbol{p}, \boldsymbol{q}$ is then measured by *Tversky index* which is parameterized by $\alpha, \beta > 0$ and is defined as:

$$\mathsf{TI}_{\alpha,\beta}(\boldsymbol{p}, \boldsymbol{q}) = \frac{|\boldsymbol{p} \cap \boldsymbol{q}|}{|\boldsymbol{p} \cap \boldsymbol{q}| + \alpha|\boldsymbol{p} \setminus \boldsymbol{q}| + \beta|\boldsymbol{q} \setminus \boldsymbol{p}|} \quad .$$

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Tversky index is useful since it includes several important similarity measurements such as Jaccard Index (JI, which is exactly $\mathsf{TI}_{1,1}$ and also known as Tanimoto Index) and Dice index (which is exactly $\mathsf{TI}_{1/2,1/2}$) [21]. First, we introduce the basic idea and two efficient techniques for improving database privacy. Then, we describe our full proposed protocol.

Basic idea

We firstly consider the simplest case that the user has (the fingerprint of) a target compound $\boldsymbol{q}$ as a query and the server's database consists of only a single fingerprint $\boldsymbol{p}$. The case of a larger database is discussed later. The goal here is to detect whether or not the Tversky index of $\boldsymbol{p}$ and $\boldsymbol{q}$ is larger than a given threshold $1 \geq \theta > 0$. The main idea of our approach is to calculate the score

$$\theta^{-1}(|\boldsymbol{p} \cap \boldsymbol{q}|) - (|\boldsymbol{p} \cap \boldsymbol{q}| + \alpha|\boldsymbol{p} \setminus \boldsymbol{q}| + \beta|\boldsymbol{q} \setminus \boldsymbol{p}|) \tag{1}$$

from *encrypted* fingerprints $\boldsymbol{p}$ and $\boldsymbol{q}$ by an additive-homomorphic cryptosystem. The score is non-negative if and only if the Tversky index of $\boldsymbol{p}$ and $\boldsymbol{q}$ is at least θ. Now since $|\boldsymbol{p} \setminus \boldsymbol{q}| = |\boldsymbol{p}| - |\boldsymbol{p} \cap \boldsymbol{q}|$ and a similar relation holds for $|\boldsymbol{q} \setminus \boldsymbol{p}|$, the score (1) is positively proportional to

$$\lambda_1|\boldsymbol{p} \cap \boldsymbol{q}| - \lambda_2|\boldsymbol{p}| - \lambda_3|\boldsymbol{q}| \ ,$$

where $\lambda_1 = c(\theta^{-1} - 1 + \alpha + \beta)$, $\lambda_2 = c\alpha$, $\lambda_3 = c\beta$ and any positive value c. We assume that the parameters and the threshold for the Tversky index are rational numbers denoted by $\alpha = \mu_a/\gamma$, $\beta = \mu_b/\gamma$ and $\theta = \theta_n/\theta_d$, where μ_a, μ_b, γ, θ_n and θ_d are non-negative integers. By using $c = \gamma\theta_n g^{-1}$ under this assumption, λ_1, λ_2 and λ_3 become non-negative integers where g is the greatest common divisor of $\gamma(\theta_d - \theta_n) + \theta_n(\mu_a + \mu_b)$, $\theta_n\mu_a$ and $\theta_n\mu_b$.

Motivated by this observation, we define the following modified score, called the **threshold Tversky index**:

Definition 1 *Given parameters α and β and a threshold θ for the Tversky index which are rational numbers denoted by $\alpha = \mu_a/\gamma$, $\beta = \mu_b/\gamma$ and $\theta = \theta_n/\theta_d$ where μ_a, μ_b, γ, θ_n and θ_d are non-negative integers, then the* threshold Tversky index $\overline{\mathsf{TI}}_{\alpha,\beta,\theta} = \overline{\mathsf{TI}}_{\alpha,\beta,\theta}(\boldsymbol{p}, \boldsymbol{q})$ *for fingerprints $\boldsymbol{p}$ and $\boldsymbol{q}$ is defined by*

$$\overline{\mathsf{TI}}_{\alpha,\beta,\theta} := \lambda_1|\boldsymbol{p} \cap \boldsymbol{q}| - \lambda_2|\boldsymbol{p}| - \lambda_3|\boldsymbol{q}| \ ,$$

and non-negative integer parameters λ_1, λ_2 and λ_3 are defined by

$$\begin{aligned} \lambda_1 &= \gamma\theta_n g^{-1}(\theta^{-1} - 1 + \alpha + \beta) \ , \\ \lambda_2 &= \gamma\theta_n g^{-1}\alpha \ , \\ \lambda_3 &= \gamma\theta_n g^{-1}\beta \ , \end{aligned}$$

where g is the greatest common divisor of $\gamma(\theta_d - \theta_n) + \theta_n(\mu_a + \mu_b)$, $\theta_n\mu_a$ and $\theta_n\mu_b$.

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

By the above argument, we have $\mathsf{TI}_{\alpha,\beta}(\boldsymbol{p},\boldsymbol{q}) \geq \theta$ if and only if $\overline{\mathsf{TI}}_{\alpha,\beta,\theta}(\boldsymbol{p},\boldsymbol{q}) \geq 0$. Therefore, the user can know whether or not his/her target compound $\boldsymbol{q}$ is similar (i.e., $\mathsf{TI}_{\alpha,\beta}(\boldsymbol{p},\boldsymbol{q}) \geq \theta$) to the fingerprint $\boldsymbol{p}$ in the database, by obtaining only the value $\overline{\mathsf{TI}}_{\alpha,\beta,\theta}(\boldsymbol{p},\boldsymbol{q})$.

In the protocol, the bits of the user's target fingerprint $\boldsymbol{q}$ and the value $|\boldsymbol{p}|$ held by the server are both encrypted using the user's public key. Since $\overline{\mathsf{TI}}_{\alpha,\beta,\theta}(\boldsymbol{p},\boldsymbol{q})$ can be computed by the addition of these values and multiplication by integers, the protocol can calculate (without the secret key) a ciphertext of $\overline{\mathsf{TI}}_{\alpha,\beta,\theta}(\boldsymbol{p},\boldsymbol{q})$, which is then decrypted by the user. For simplicity, we will abuse the notation and write $\mathsf{TI}(\boldsymbol{p},\boldsymbol{q})$, $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ without subscripts α,β,θ when the context is clear.

We emphasize that our protocol does not use time-consuming cryptographic methods such as GP-MPC and FHE, and data transfer occurs only twice during an execution of the protocol. Hence, our protocol is efficient enough to scale to large databases.

Database security enhancement techniques against regression attack

The ideal situation for the server is that the user learns only the similarity/non-similarity property of fingerprints $\boldsymbol{p}$ and $\boldsymbol{q}$, without knowing any other information about the secret fingerprint $\boldsymbol{p}$. This means that only the sign of $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ should be known by the user. However, in our basic protocol, the value of $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ is fully obtained by the user; Database privacy is not protected from regression attacks. (See the Security analyses section for details.) In order to send only the sign of $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$, we firstly considered using a bit-wise decomposition protocol [22] for extracting and sending only the sign bit of $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$. Although this approach is ideal in terms of security, the protocol requires more than 30 rounds of communications, which is much more efficient than using GP-MPC or FHE, but rather time-consuming for large-scale databases. Therefore, here we propose the novel technique of using dummy replies, which requires only one round of communication while sufficiently minimizing information leakage of $\boldsymbol{p}$. In the proposed technique, besides its original reply $t = \mathsf{Enc}(\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q}))$, the server also chooses random integers $\phi_1,\ldots,\phi_n$ from a suitable interval and encrypts those values under the user's public key pk. Then the server sends the user a collection of ciphertexts $t, \mathsf{Enc}(\phi_1),\ldots,\mathsf{Enc}(\phi_n)$ that are shuffled to conceal the true ciphertext t, as well as the number s_{d} of dummy values ϕ_k with $\phi_k \geq 0$. The user decrypts the received $n+1$ ciphertexts, counts the number s_{c} of non-negative values among the decryption results, and compares s_{c} to s_{d}. Now we have $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q}) \geq 0$ if and only if $s_{\mathrm{c}} - s_{\mathrm{d}} = 1$; therefore, the user can still learn the sign of $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$, while the actual value of $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ is concealed by the dummies. We have confirmed that the information leakage of $\boldsymbol{p}$ approaches zero as the number of dummies becomes large; see the Security analyses for pudding dummies section for more detailed discussion. (We have also developed another security enhancement technique using sign-preserving randomization of $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$; see the Further security enhancement technique by using sign-preserving randomization section in Supporting Information for details.)

Database security enhancement technique against illegal query attack

Illegal query attacks can be prevented if the server can detect whether or not the user's query is valid. To keep user privacy, the server must conduct this task without obtaining more information than the validity/invalidity of the query. In fact,

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

this functionality can be implemented by using the NIZK proof by Sakai et al. [20] mentioned in the Non-interactive zero-knowledge proof section. The improved protocol requires the user to send the server a proof associated with the encrypted fingerprint bits q_i, from which the server can check whether $\boldsymbol{q}$ is indeed a valid fingerprint (without obtaining any other information about $\boldsymbol{q}$); the server aborts the protocol if $\boldsymbol{q}$ is invalid. Here we use the "lifted" ElGamal cryptosystem as our basic encryption scheme to apply Sakai's scheme. (We note that if we require the user to send $\mathsf{Enc}(-|\boldsymbol{q}|)$ used by server's computation, then another NIZK proof is necessary to guarantee the validity of the additional ciphertext, which decreases the communication efficiency of our protocol. Hence our protocol requires the server to calculate $\mathsf{Enc}(-|\boldsymbol{q}|)$ by itself.)

Secure similar compounds counter

For the general case that the database consists of more than one fingerprint $\boldsymbol{p}$, we propose the protocol shown in Algorithm 1 to count the number of fingerprints $\boldsymbol{p}$ similar to the target fingerprint $\boldsymbol{q}$. In the protocol, the server simply calculates the encryption of the threshold Tversky indices for all database entries and, as discussed above, replies with a shuffled collection of these true ciphertexts and dummy ciphertexts, as well as the number s_{d} of non-negative dummy values. Then the value $s_{\mathrm{c}} - s_{\mathrm{d}}$ finally obtained by the user is equal to the number of similar fingerprints $\boldsymbol{p}$ in the database.

Parameter settings of the protocol

Decrypting an encrytion of too large value needs huge computation cost if the lifted-ElGamal cryptosystem is used. Therefore, in order to keep the consistency and efficiency of the protocol, the range of $\overline{\mathsf{TI}}(\boldsymbol{p}, \boldsymbol{q})$ should not be too large. i.e., the integer parameters λ_1, λ_2 and λ_3 in the threshold Tversky index should not be too large. In fact, this will not cause a problem in practice; For example, the parameters become $\lambda_1 = 9$, $\lambda_2 = \lambda_3 = 4$ for computing $\overline{\mathsf{TI}}_{1,1,0.8}$ which is a typical setting of a chemical compound search. In this case, a minimum value and a maximum value of $\overline{\mathsf{TI}}(\boldsymbol{p}, \boldsymbol{q})$ is -664 and 166 for 166 MACCS keys, which is a sufficiently small range. (See Ranges of $\overline{\mathsf{TI}}$ for typical parameter settings section in Supporting Information for details.)

Security analyses

In the area of cryptology, the following two standard security models for two-party computation have been considered:

- *Semi-honest model*: Both parties follow the protocol, but an adversarial one attempts to infer additional information about the other party's secret input from the legally obtained information.
- *Malicious model*: An adversarial party cheats even in the protocol (e.g., by inputting maliciously chosen invalid values) in order to illegally obtaining additional information about the secret.

In this section, we evaluate security of SSCC in both the semi-honest and malicious models.

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Algorithm 1 The secure similar compounds counter (SSCC)

- Public input: Length of fingerprints ℓ and parameters for the Tversky index $\theta = \theta_n/\theta_d, \alpha = \mu_a/\gamma, \beta = \mu_b/\gamma$
- Private input of a user: Target fingerprint $\boldsymbol{q}$
- Private input of a server: Set of fingerprints $P = \{\boldsymbol{p}^{(1)}, \ldots, \boldsymbol{p}^{(M)}\}$

1 (Key setup of cryptosystem) The user generates a key pair (pk, sk) by the key generation algorithm KeyGen for the additive-homomorphic cryptosystem and sends public key pk to the server (the user and the server share public key pk and only the user knows secret key sk).

2 (Initialization) The user encrypts his/her fingerprint $\boldsymbol{q}$ as a vector of ciphertexts: $\vec{\mathsf{Enc}}(q_k) := (\mathsf{Enc}(q_1), \ldots, \mathsf{Enc}(q_\ell))$. He/she also generates $\boldsymbol{v}$ as a vector of proofs. Each proof v_i is associated with $\mathsf{Enc}(q_i)$.

3 (Query of entry) The user sends the vector of ciphertexts $\vec{\mathsf{Enc}}(q_k)$ and the vector of proofs $\boldsymbol{v}$ to the server as a query.

4 (Query validity verification) The server verifies the validity of $\vec{\mathsf{Enc}}(q_k)$ by testing the vector of proof $\boldsymbol{v}$. If $\boldsymbol{v}$ does not pass the server's test, the user cannot move on to the next step.

5 (Calculation of threshold Tversky index)

(a) The server calculates the greatest common divisor of $\gamma(\theta_d - \theta_n) + \theta_n(\mu_a + \mu_b)$, $\theta_n\mu_a$ and $\theta_n\mu_b$ as g, and calculates $\lambda_1 = \gamma\theta_n g^{-1}(\theta^{-1} - 1 + \alpha + \beta)$, $\lambda_2 = \gamma\theta_n g^{-1}\alpha$, and $\lambda_3 = \gamma\theta_n g^{-1}\beta$.

(b) The server calculates $\mathsf{Enc}(-|\boldsymbol{q}|) = \mathsf{Enc}(-\sum_{i=1}^{\ell} q_i)$ from $\vec{\mathsf{Enc}}(q_k)$: $\mathsf{Enc}(-|\boldsymbol{q}|) = -1 \otimes \bigoplus_{i=1}^{\ell} \mathsf{Enc}(q_i)$.

(c) **for** $j = 1$ to M **do**

i. The server calculates $-|\boldsymbol{p}^{(j)}| = -\sum_{i=1}^{\ell} p_i^{(j)}$ and encrypts it to obtain a ciphertext $\mathsf{Enc}(-|\boldsymbol{p}^{(j)}|)$.

ii. The server calculates a ciphertext t_j of threshold Tversky index $\overline{\mathsf{TI}}(\boldsymbol{p}^{(j)}, \boldsymbol{q})$.

```
c ← Enc(0)
for k = 1 to ℓ do
if p_k^(j) = 1
    c ← c ⊕ Enc(q_k)          ▷ Computing Enc(|p^(j) ∩ q|)
end if
end for
t_j ← λ1 ⊗ c ⊕ λ2 ⊗ Enc(−|p^(j)|) ⊕ λ3 ⊗ Enc(−|q|)
end for
```

6 (Padding of dummies)

(a) The server generates a set of dummy values $\{\phi_1, \ldots, \phi_n\}$ and counts the number s_{d} of non-negative dummies $\phi_i \geq 0$.

(b) The server encrypts ϕ_i to obtain a ciphertext $\mathsf{Enc}(\phi_i)$ for $i = 1, \ldots, n$.

(c) The server shuffles the contents of the set $T = \{t_1, \ldots, t_M, \mathsf{Enc}(\phi_1), \ldots, \mathsf{Enc}(\phi_n)\}$.

7 (Return of matching results) The server sends T and s_{d} to the user.

8 (Decryption and counting) The user decrypts the contents of T and counts the number s_{c} of non-negative values.

9 (Evaluation) The user obtains $s_{\mathrm{c}} - s_{\mathrm{d}}$ as the number of similar fingerprints in the database.

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

User privacy

The semantic security of the encryption scheme used in the protocol (see the Additively homomorphic encryption scheme section) implies immediately that the server cannot infer any information about the user's target fingerprint $\boldsymbol{q}$ during the protocol. This holds in both the semi-honest and malicious models.

Thresholding largely improves database privacy

By the analysis described below, we show the large difference between the "ideal" case, in which the user learns only the sign of $\overline{\mathsf{TI}}(\boldsymbol{p}, \boldsymbol{q})$ during the protocol, and the "plain" case, in which the user fully learns the value $\mathsf{TI}(\boldsymbol{p}, \boldsymbol{q})$. Note that SSCC aims for the "ideal" case while the methods proposed in previous studies aim for the "plain" case. Here we consider the general case in which the user is allowed to send more than one query and those queries are searched by Jaccard Index. We also suppose that the database consists of a single fingerprint $\boldsymbol{p}$ in order to clarify the effect of thresholding.

The goal of an attacker is to reveal $\boldsymbol{p}$ by analysing the results returned from the server. It is generally effective for the attacker to exploit the difference between the two outputs obtained by sending two different queries. In fact, when the server returns $\overline{\mathsf{TI}}$, $\overline{\mathsf{TI}}(\boldsymbol{p}, \boldsymbol{q})$ - $\overline{\mathsf{TI}}(\boldsymbol{p}, \boldsymbol{0})$ becomes positive if and only if $p_i = 1$, where $\boldsymbol{q} = (0, \ldots, q_i = 1, \ldots, 0)$ and $\boldsymbol{0} = (0, \ldots, 0)$. This means that the attacker can reveal any bit in $\boldsymbol{p}$ by sending the single query after sending the first query $\boldsymbol{0}$. Therefore, $\boldsymbol{p}$ can be fully revealed by sending only $\ell + 1$ queries. On the other hand, there is no deterministic attack for revealing $\boldsymbol{p}$ from only the sign of $\overline{\mathsf{TI}}$, because two different inputs do not always lead to different outputs. Since we know of a linear algorithm that fully reveals $\boldsymbol{p}$ in response to at most 2ℓ queries after making a "hit" query $\boldsymbol{q}$ such that $\overline{\mathsf{TI}}(\boldsymbol{p}, \boldsymbol{q}) > 0$, here we evaluate database privacy by the probability of making at least one hit query when the user is allowed to send x queries. (See the The attack algorithm by using a hit query section in Supporting Information for details.) This probability is denoted as

$$\sum_{\boldsymbol{p}} \Pr\Big(X = \boldsymbol{p}\Big) \cdot \Big(1 - (1 - f_{\boldsymbol{p}})^x\Big) \ , \tag{2}$$

where $f_{\boldsymbol{p}}$, defined as follows, is the probability that the user makes one hit query with a single trial when $\boldsymbol{p}$ is given.

$$f_{\boldsymbol{p}} := \sum_{\boldsymbol{q}} \Pr\Big(Y = \boldsymbol{q}\Big) \cdot \Pr\Big(\ \overline{\mathsf{TI}}(\boldsymbol{p}, Y) > 0 \ \Big| \ Y = \boldsymbol{q} \ \Big) \ .$$

For ease of calculation, we computed the upper bound of equation (2) for $x = 1, 10, 10^2, \ldots, 10^6$ and $\theta = 0.7, 0.8, 0.9, 1.0$. (See the Derivation and calculation of upper bound of the probability for making at least one hit query section in Supporting Information for details.) Since publicly available 166 MACCS keys are the most popular fingerprint for chemical compound searches, we set ℓ to 166. From the results shown in Fig. 3, we can see that the probability of making a hit query is sufficiently small even though the user is allowed to send a large number of queries. Considering that the user learns $\boldsymbol{p}$ by using no more than $\ell + 1$ queries when he/she

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

learns $\overline{\mathsf{TI}}$, we can conclude that database privacy is largely improved by thresholding. In other words, the proposed protocol, which aims to output only the sign of the similarity score, has stronger security than other previous methods, which directly output similarity scores.

Security analyses for padding dummies

We showed that the output privacy in the "ideal" case is significantly improved from the "plain" case. Here we experimentally evaluate how the actual situation of our proposed protocol is close to the "ideal" case.

Before going into detail analyses, let us discuss how to generate dummies. It is ideal for the server privacy to generate a dummy according to the same distribution where $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ is generated from. However, this is not realistic because $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ is determined by both $\boldsymbol{p}$ and $\boldsymbol{q}$ which is user's private information. Therefore, in our analyses, we assume that a dummy is generated from uniform distribution over possible values of $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$. For example, if possible values of $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ is $\{1,2,3,4,5\}$, dummies are randomly selected from any one of them. The purpose of pudding dummies is to mitigate the risk of leaking $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$. In order to clarify the effect of the use of dummy values, we concentrate on the basic case; the database contains a single $\boldsymbol{p}$, and there exist k possible values of $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$. i-th value of the k possible values arises as the true $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ according to the probability w_i. Namely, true $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ is generated from the multinomial distribution with k different probabilities $\boldsymbol{w} = w_1,\ldots,w_k$, while dummies are generated from the multinomial distribution with equal probability $1/k$. To conduct stringent analyses, we assume that the user knows $\boldsymbol{w}$, and he/she also knows that dummies are uniformly distributed over k possible $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$.

In this experiment, we evaluate the security of our protocol by comparing the probabilities that the user correctly guesses the value $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ in two cases: The case in which the user makes a guess based only on a prior knowledge $\boldsymbol{w}$, and the other case in which the user makes a guess based on the observation of the search result under the condition that he/she knows $\boldsymbol{w}$.

For the first case, the user's best strategy for guessing $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ is to choose the i_0-th possible value, where

$$i_0 = \arg\max_{1\le i\le k} w_i \ . \tag{3}$$

In this case, the success probability of the guess is w_{i_0}.

Let us consider the best strategy for the second case. As described above, we consider an practical case that n dummy values $\phi_1,\ldots,\phi_n$ chosen from the k possible values uniformly at random, and the user makes a guess from the received $n+1$ shuffled values $\phi_1,\ldots,\phi_n,\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$. Now suppose that the user received the i-th possible value a_i times for each $1\le i\le k$ (hence $\sum_{i=1}^k a_i = n+1$). Since the choices of $\phi_1,\ldots,\phi_n$ are independent of $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$, the probability that the user received i-th possible value a_i times for each $1\le i\le k$ and that $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ is i_0-th possible value is

$$\binom{n}{a_1,\ldots,a_{i_0}-1,\ldots,a_k}\left(\frac{1}{k}\right)^n \cdot w_{i_0} = a_{i_0}\cdot w_{i_0}\ \frac{n!}{a_1!\cdots a_k!k^n}\ .$$

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Therefore, the conditional probability that $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ is the i_0-th possible value, conditioned on the set of the user's received values, is

$$\frac{\binom{n}{a_1,\ldots,a_{i_0}-1,\ldots,a_k}\left(\frac{1}{k}\right)^n \cdot w_{i_0}}{\sum_{i=1}^{k}\binom{n}{a_1,\ldots,a_i-1,\ldots,a_k}\left(\frac{1}{k}\right)^n \cdot w_i} = \frac{a_{i_0}\cdot w_{i_0}}{\sum_{i=1}^{k} a_i \cdot w_i} \ .$$

This implies that the user's best strategy is to guess that $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ is the i_0-th possible value, where

$$i_0 = \arg\max_{1\le i\le k} a_i \cdot w_i \ . \tag{4}$$

We estimated success probabilities of user's guess for the both cases by simulation experiments. Here we assumed typical case when $\mathsf{TI}_{1,1,0.8}$ and 166 MACCS keys are used. In this case, $k = 831$ and we performed the experiments for $n = 831 \times 10^0, 831 \times 10^1, \ldots, 831 \times 10^4$ on three different distributions of $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ which were obtained by the following schemes:

1. We randomly selected one fingerprint $\boldsymbol{q}$ from ChEMBL and calculated $\overline{\mathsf{TI}}(\boldsymbol{q},\boldsymbol{p})$ for all the entries in ChEMBL and used the observed distribution as $\boldsymbol{w}$. In our experiment, 177159-th fingerprint was selected as $\boldsymbol{q}$ (referred as $\boldsymbol{w}^{\mathrm{ChEMBL-177159}}$).
2. The same scheme as 1) was used when $\boldsymbol{q}$ was 265935-th fingerprint (referred as $\boldsymbol{w}^{\mathrm{ChEMBL-265935}}$).
3. We randomly selected a value from $1,\ldots,k$ for m times and count frequency of i as h_i and set $w_i = h_i/m$ (referred as $\boldsymbol{w}^{\mathrm{random}}$). We used $k \times 5$ as m.

All the distributions used here are shown in the Distribution of $\boldsymbol{w}$ used in the experiments section in Supporting Information.

We performed $100,000$ trials for each experiment. Each trial consisted of choosing $\phi_1,\ldots,\phi_n$ uniformly at random; choosing $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ according to $\boldsymbol{w}$; deciding the user's guess i_0 by formula (3) and formula (4) respectively (we adopted a uniformly random choice if there were more than one such i_0); and checking whether or not $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ was the i_0-th possible value for both rules (i.e., the user's guess succeeded). The results of the experiment are given in Table 1; they show that the user's attack success probability became significantly close to the ideal case when a sufficiently large number of dummies were used; therefore, our technique of using dummies indeed improves the output privacy.

One might suspect that the attacker can detect the true $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ by sending the same query twice and finding the value which is appeared in both results. However, this attack does not easily succeed if n is sufficiently larger than k (i.e., ideally, all possible values of $\overline{\mathsf{TI}}$ are covered by sufficient number of dummies), and we consider that k is not too large in practice as we discussed in Parameters settings of the protocol section. In order to evaluate the security of the case when the user is allowed to submit L queries, we performed following analyses. For the efficiency of the attack, we assumed that the attacker keeps sending the same query L times. For this case, the probability that $\overline{\mathsf{TI}}(\boldsymbol{p},\boldsymbol{q})$ is the i_0-th possible value after sending

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

L queries on condition that frequency of i-th possible value of j-th query $a_i^{(j)}$ for $j = 1, \ldots, L$ is

$$\prod_{j=1}^{L} \frac{a_i^{(j)} \cdot w_i}{\sum_{h=1}^{k} a_h^{(j)} \cdot w_h} \quad . \tag{5}$$

This implies that the user's best strategy is to choose i-th possible value which maximizes equation (5). In our experiment, we compared the success ratio of the attack based on the above strategy and the ideal success ratio when the user makes the guess only from the given distribution $\boldsymbol{w}$. We also assumed more realistic case that user did not know the exact distribution of dummy but knew the distribution that was similar to the actual distribution the server used. For the evaluation of this case, we generated dummies from the distribution $\boldsymbol{u}$, which was slightly different from uniform distribution, while the user assumed that dummies were generated from uniform distribution. $\boldsymbol{u}$ was generated as follows:

$$u_i = r \cdot 1/k, \quad \text{where} \quad r \sim N(1, \delta^2).$$

We performed the experiment for $L = 1, 10, 10^2, \ldots, 10^5$, $n = 831 \times 10, 831 \times 50, 831 \times 10^2$ and $\delta = 0, 0.05, 0.1, 0.15, 0.2$ based on the same approach used in the evaluation of single query security. i.e., for each trial, n dummies were randomly chosen according to $\boldsymbol{u}$ (note that $\boldsymbol{u}$ was equal to uniform distribution when $\delta = 0$), true value $\overline{\mathsf{TI}}(\boldsymbol{p}, \boldsymbol{q})$ was selected according to $\boldsymbol{w}$ and the attacker's guess was made based on the equation (5). We performed $10,000$ trials for each triplet of L, n and δ. Those experiments were conducted for the same three distributions: $\boldsymbol{w}^{\mathrm{ChEMBL-177159}}$, $\boldsymbol{w}^{\mathrm{ChEMBL-265935}}$ and $\boldsymbol{w}^{\mathrm{ChEMBL-random}}$. We compared the success ratio of the attack and the ideal success ratio when the user made the guess without seeing search results. The results are shown in Fig. 4. The success ratio of user's attack decreased as the number of dummies increased and became closer to the ideal value when the sufficient number of dummies are given, even for the case that a large number of queries were sent. Although an efficient method for dummy generation remains as a future task, the results also show that database privacy is largely improved by hiding the distribution of dummy and the user has to know it with high accuracy in order to attack the server successfully.

Database privacy in malicious model

For our protocol, the difference between the malicious and semi-honest models is that in the malicious model the user may use an invalid input $\boldsymbol{q}$ whose components q_i are not necessarily in $\{0, 1\}$. If the user chooses $\boldsymbol{q}$ in such a way that some component q_i is extremely large and the remaining $\ell - 1$ components are all zero, then $\overline{\mathsf{TI}}(\boldsymbol{p}, \boldsymbol{q})$ will also be an extreme value (distinguishable from the dummy values) and depend dominantly on the bit p_i; therefore, the user can almost surely guess the secret bit p_i. Since our protocol detects whether or not q_i is a bit value without invading user privacy, it can safely reject illegal queries and prevent any illegal query attacks, including above case.

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Performance evaluation

In this section, we evaluate the performance of the proposed method on two datasets created from ChEMBL. We implemented the proposed protocol based on lifted ElGamal encryption. For the implementation, we use elliptic curve parameters called secp192k1, as recommended by SECG (The Standards for Efficient Cryptography Group). These parameters are considered to be more secure than 1024-bit RSA encryption, which is the most commonly used public-key cryptosystem. Owing to the limitation of the range of plaintext, the implementation here does not include sign-preserving randomization. For the purpose of comparison, we also implemented a GP-MPC protocol by using Fairplay [23] whose input and output are a fingerprint and the signs of the Tversky indices, respectively. The Jaccard index along with the threshold $\theta = 0.8$ were used for both protocols. For SSCC, we used 10,000 dummies. These two implementations were tested on two datasets: one, referred to as ChEMBL_1000, was the first 1000 fingerprints stored in ChEMBL, and the other, referred to as ChEMBL_Full, was 1,292,344 fingerprints in the latest version of ChEMBL. All the programs were run on a single core of an Intel Xeon 2.9 GHz on the same machine. The results are shown in Table 2. To avoid environmental effects, we repeated the same experiment five times and calculated average values. Since both CPU time and communication size are exactly linear to the size of database for the GP-MPC protocol, results of ChEMBL_Full for GP-MPC were estimated from the results of ChEMBL_1000 because of the limitation of computational resources.

Despite the proposed method including elaborate calculation like the NIZK proof, we can see from the results that both the CPU time and communication size of the proposed method are significantly smaller than those of the GP-MPC protocol. Furthermore, it is clear that SSCC provides industrial-strength performance, considering that it works, even on a huge database like ChEMBL_Full, taking no more than 167 s and 173 s for the server and client respectively. By using simple data parallelization, the computational speed will be improved linearly with the number of CPUs. Since all the programs were run on the same machine there was almost no latency for the communication between the two parties in these experiments. Therefore, GP-MPC, whose communication size is huge, is expected to require far more time when it runs on an actual network that is not always in a good condition. The other important point is that SSCC requires only two data transfers, which enables data transfer after off-line calculation. On the other hand, GP-MPC must keep online during the search because of the high communication frequency. We also note that it took less than 100 MB to compile SSCC, while GP-MPC required more than 16 GB. Considering these observations, SSCC is efficient for practical use. It is known that several techniques improve the performance of GP-MPC and the previous work by Pinkas et al. [24] reported that Free XOR [25] and Garbled Row Reduction [24], which are commonly used in state-of-the-art GP-MPC methods [26] [27] [28] [29], reduced running time and communication size by factors of 1.8 and 6.3 respectively when a circuit computing an encryption of AES was evaluated. Though these techniques are not implemented in Fairplay, we consider that GP-MPC is yet far less practical for the large-scale chemical compound search problem compared to our method which improved running time and communication size by factors of $36,900$ and $12,000$.

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Conclusion

In this study, we proposed a novel privacy-preserving protocol for searching chemical compound databases. To our knowledge, this is the first practical study for privacy-preserving (for both user and database sides) similarity searching in the fields of bioinformatics and chemoinformatics. Moreover, the proposed method could be applied to a wide range of life science problems such as searching for similar single-nucleotide polymorphism (SNP) patterns in a personal genome database. While the protocol proposed here focuses on searching for a number of similar compounds, we are examining further improvements of the protocol such as the client being able to download similar compounds; we expect this on-going study to further contribute to the drug screening process. In recent years, open innovation has been attracting attention as a promising approach for speeding up the process of new drug discovery [30]. For example, research on neglected tropical diseases including malaria has been promoted by the recent attempt to share chemical compound libraries in the research community. In spite of high expectations, such an approach is still limited to economically less important problems on account of privacy problems [31]. Therefore, privacy-preserving data mining technology is expected to be the breakthrough promoting open innovation and we believe that our study will play an important role.

Competing interests

The authors declare that they have no competing interests.

Author's contributions

KS, HA, KN and KA designed the protocol inspired by the discussion with TH. KS, KN, NA and GH conducted security analyses. KS and SM implemented the protocol. KS evaluated performance of the protocol. All authors wrote the manuscript. All authors read, commented and approved the final manuscript.

Acknowledgement

KS thanks Yusuke Sakai and Takahiro Matsuda for fruitful discussions.

Author details

[1]Computational Biology Research Center, National Institute of Advanced Industrial Science and Technology (AIST), 2-4-7 Aomi Koto-ku, Tokyo, Japan. [2]Research Institute of Secure Systems, National Institute of Advanced Industrial Science and Technology (AIST), 1-1-1 Umezono Tsukuba, Ibaraki, Japan. [3]Information Technology Center, The University of Tokyo 7-3-1 Hongo Bunkyo-ku, Tokyo, Japan. [4]Cybozu Labs 12F, Koraku Mori Bldg. 1-4-14, Tokyo, Japan. [5]Faculty of Science and Engineering, Waseda University 3-4-1 Okubo Shinjuku-ku, Tokyo, Japan. [6]Graduate School of Frontier Sciences, The University of Tokyo 5-1-5, Chiba, Japan. [7]Molecular Profiling Research Center for Drug Discovery, National Institute of Advanced Industrial Science and Technology (AIST), 2-4-7 Aomi Koto-ku, Tokyo, Japan. [8]Graduate School of SIE, University of Tsukuba 1-1-1 Tennodai Tsukuba, Ibaraki, Japan.

References

1. Subbaraman, N.: Flawed arithmetic on drug development costs. Nature Biotechnology **29**(5), 381–381 (2011)
2. Miller, M.a.: Chemical database techniques in drug discovery. Nature Reviews Drug Discovery **1**(3), 220–7 (2002)
3. Schooler, J.: Unpublished results hide the decline effect. Nature **470**, 437 (2011)
4. Ostrovsky, R., Skeith, W.E. III.: A survey of single-database private information retrieval: techniques and applications. In: Proceedings of the 10th International Conference on Practice and Theory in Public-key Cryptography. PKC'07, pp. 393–411 (2007)
5. Goethals, B., Laur, S., Lipmaa, H., Mielikäinen, T.: On private scalar product computation for privacy-preserving data mining. In: Proceedings of the 7th Annual International Conference on Information Security and Cryptology. ICISC 2004, pp. 104–120 (2004)
6. Blundo, C., Cristofaro, E.D., Gasti, P.: EsPRESSo : Efficient Privacy-Preserving Evaluation of Sample Set Similarity. In: Proceedings of Data Privacy Management and Autonomous Spontaneous Security: 7th International Workshop, DPM 2009 and 5th International Workshop, SETOP 2012. DMP/SETOP 2012, pp. 89–103 (2012)
7. Murugesan, M., Jiang, W., Clifton, C., Si, L., Vaidya, J.: Efficient privacy-preserving similar document detection. The VLDB Journal **19**(4), 457–475 (2010)
8. Yao, A.C.-C.: How to generate and exchange secrets. In: Proceedings of the 27th Annual Symposium on Foundations of Computer Science. SFCS '86, pp. 162–167 (1986)
9. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing. STOC '09, pp. 169–178 (2009)

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

10. Laur, S., Lipmaa, H.: On private similarity search protocols. In: Proceedings of the 9th Nordic Workshop on Secure IT Systems. NordSec 2004, pp. 73–77 (2004)
11. Popa, R.A., Redfield, C.M.S., Zeldovich, N., Balakrishnan, H.: In: Proceedings of the 23rd ACM Symposium on Operating Systems Principles. SOSP 11, pp. 85–100
12. Martin, Y.C., Kofron, J.L., Traphagen, L.M.: Do structurally similar molecules have similar biological activity? Journal of Medicinal Chemistry **45**(19), 4350–4358 (2002)
13. Miller, J.L.: Recent developments in focused library design: targeting gene-families. Current Topics in Medicinal Chemistry **6**(1), 19–29 (2006)
14. Curty, R., Tang, J.: Someone's loss might be your gain: A case of negative results publications in science. In: Proceedings of the American Society for Information Science and Technology. ASISTS, vol. 49 (2012)
15. Gaulton, A., Bellis, L.J., Bento, A.P., Chambers, J., Davies, M., Hersey, A., Light, Y., McGlinchey, S., Michalovich, D., Al-Lazikani, B., Overington, J.P.: ChEMBL: a large-scale bioactivity database for drug discovery. Nucleic Acids Research **40**(Database issue), 1100–1107 (2012)
16. Goldwasser, S., Micali, S.: Probabilistic encryption. J. Comput. Syst. Sci. **28**(2), 270–299 (1984)
17. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Proceedings of the 17th International Conference on Theory and Application of Cryptographic Techniques. EUROCRYPT'99, pp. 223–238 (1999)
18. ElGamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. IEEE Transactions on Information Theory **31**(4), 469–472 (1985)
19. Goldreich, O.: Foundations of Cryptography: Volume 1. Cambridge University Press, ??? (2001)
20. Sakai, Y., Emura, K., Hanaoka, G., Kawai, Y., Omote, K.: Methods for restricting message space in public-key encryption. IEICE Transactions **96-A**(6), 1156–1168 (2013)
21. Tversky, A.: Features of similarity. Psychological Review **84**(4), 327–352 (1977)
22. Damgård, I., Fitzi, M., Kiltz, E., Nielsen, J.B., Toft, T.: Unconditionally secure constant-rounds multi-party computation for equality, comparison, bits and exponentiation. In: Proceedings of the 3rd Theory of Cryptography Conference. TCC 2006, pp. 285–304 (2006)
23. Ben-david, A., Nisan, N., Pinkas, B.: Fairplaymp: A system for secure multi-party computation. In: Proceedings of ACM Conference on Computer and Communications Security. CCS 2008, pp. 17–21 (2008)
24. Pinkas, B., Schneider, T., Smart, N.P., Williams, S.C.: Secure two-party computation is practical. In: Proceedings of the 15th International Conference on the Theory and Application of Cryptology and Information Security. ASIACRYPT 2009, pp. 250–267 (2009)
25. Kolesnikov, V., Schneider, T.: Improved garbled circuit: Free XOR gates and applications. In: Proceedings of the 35th International Colloquium on Automata, Languages and Programming. ICALP 2008, pp. 486–498 (2008)
26. Henecka, W., Kögl, S., Sadeghi, A., Schneider, T., Wehrenberg, I.: TASTY: tool for automating secure two-party computations. In: Proceedings of the 17th ACM Conference on Computer and Communications Security. CCS 2010, pp. 451–462 (2010)
27. Huang, Y., Evans, D., Katz, J., Malka, L.: Faster secure two-party computation using garbled circuits. In: Proceedings of the 20th USENIX Security Symposium. USENIX 2011 (2011)
28. Huang, Y., Shen, C.-H., Evans, D., Katz, J., Shelat, A.: Efficient secure computation with garbled circuits. In: Proceedings of the 7th International Conference on Information Systems Security. ICISS, pp. 28–48 (2011)
29. Kreuter, B., Shelat, A., Shen, C.: Billion-gate secure computation with malicious adversaries. In: Proceedings of the 21th USENIX Security Symposium. USENIX Security 2012, pp. 285–300 (2012)
30. Williams, A.J., Harland, L., Groth, P., Pettifer, S., Chichester, C., Willighagen, E.L., Evelo, C.T., Blomberg, N., Ecker, G., Goble, C., Mons, B.: Open PHACTS: semantic interoperability for drug discovery. Drug Discovery Today **17**(21-22), 1188–1198 (2012)
31. Hunter, J., Stephens, S.: Is open innovation the way forward for big pharma? Nature Reviews Drug Discovery **9**(2), 87–88 (2010)

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Figures

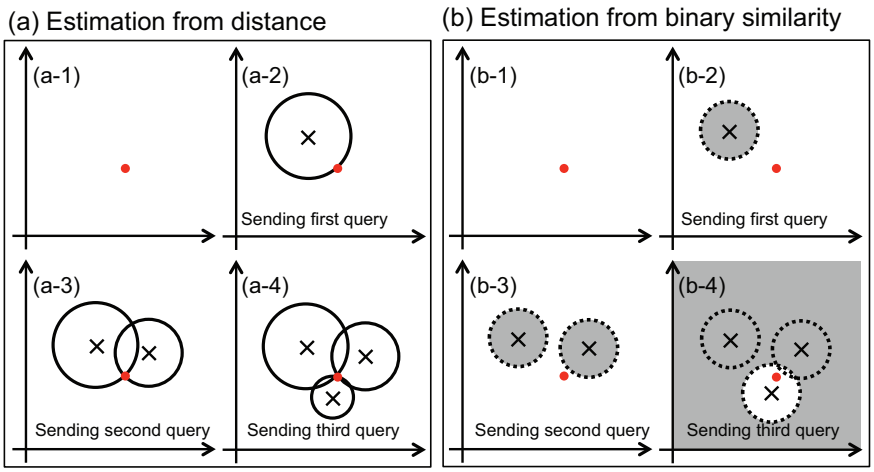

Figure 1 Schematic view showing a large difference in tolerance against the regression attack between two cases: (a) The server's reply is the distance between the attacker's query and the server's data, (b) The server's reply is the binary sign that shows whether or not the distance between the attacker's query and the server's data is larger than the given threshold. The red point represents the server's data and x represents the attacker's query. Prior to the query, the search spaces (white areas) in (a-1) and (b-1) are equal. After the first query has been sent, the search space in (a-2) is limited to the circle whose radius is the distance between the attacker's query and the server's data. On the other hand in (b-2), only the small area of the dashed circle whose radius is the given threshold (gray area) is excluded from the search space. By sending the second query, the attacker knows that one of the two intersections of the two circles in (a-3) is equal to the server's data, while the search space is large in (b-3). Finally, the server's data is detected by sending the third query in (a-4), however in (b-4), the search space is still large, even though the third query is within the given threshold.

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

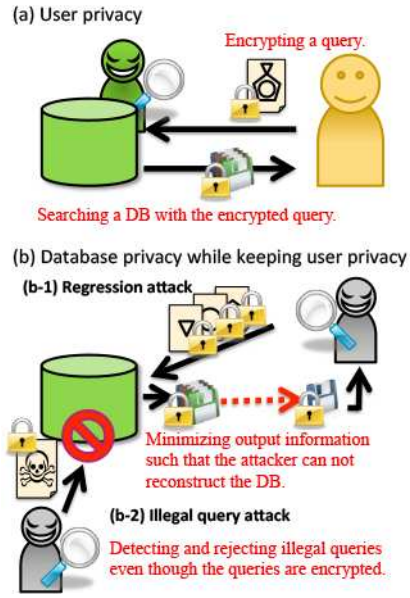

Figure 2 Schematic view of protection of (a) user privacy and (b) database privacy while keeping user privacy. For user privacy, the user's query and the search result which includes the query information must be invisible to the database side during the search task. For database privacy, the server minimizes output information for preventing regression attacks (b-1), and also detects and rejects illegal queries that might cause unexpected information leakage (b-2). These server's tasks must be carried out with the encrypted queries in order to keep user privacy.

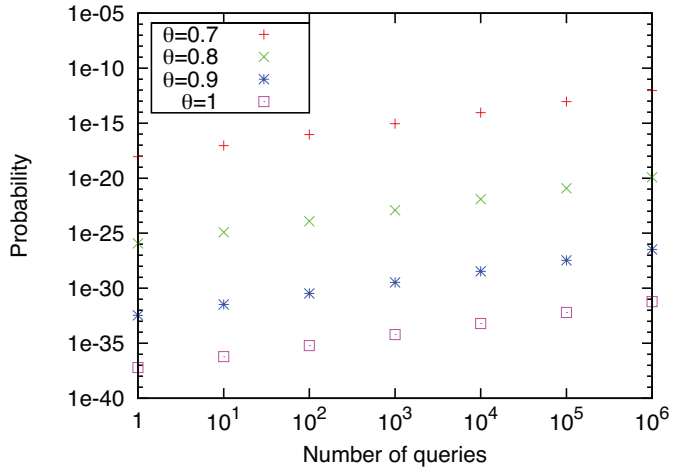

Figure 3 Upper bounds of the probabilities that the user has at least one hit query out of making $1, 10, \ldots, 10^6$ queries. Note that the hit query becomes the critical hint for revealing database information. Each line shows the results with one of the four different thresholds.

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

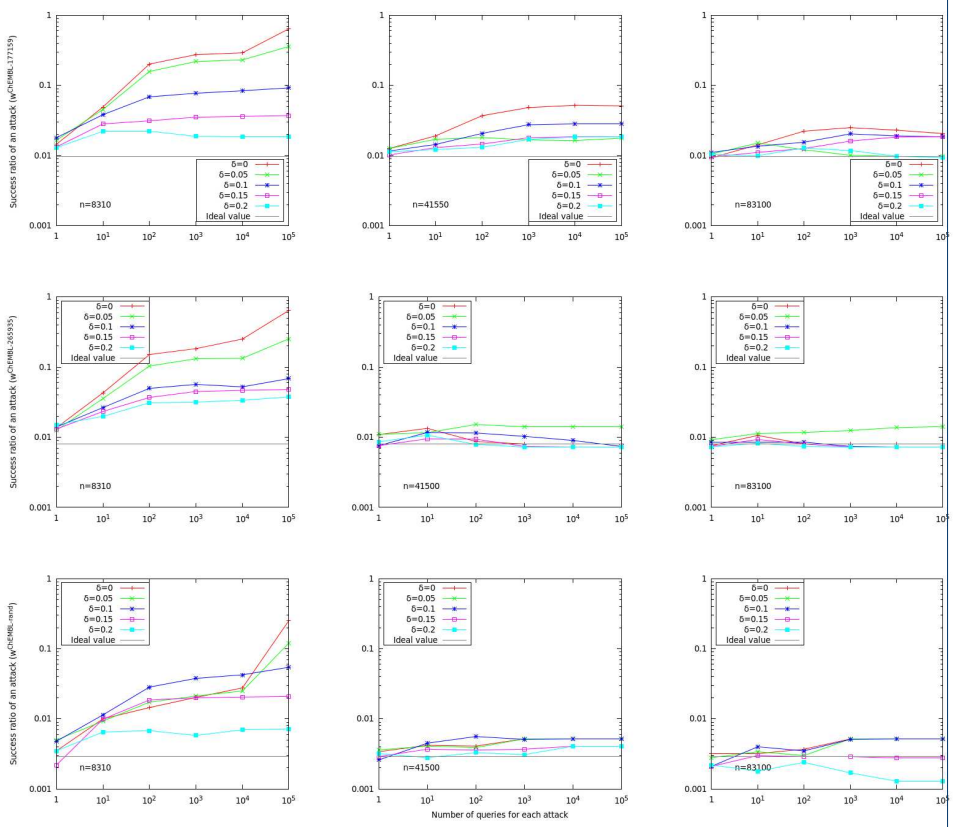

Figure 4 The comparison of the experimental success ratios of the user's guess based on the server's return as well as the prior distribution of true value when the user sends many queries, and success probability based only on a guess from the prior distribution. $\overline{\mathrm{TI}}_{1,1,0.8}$ $(k = 831)$ is assumed and results are calculated for three different numbers of dummies ($n = 831 \times 10$, 831×50, 831×10^2) when the user sends $L = 1, 10, \ldots, 10^5$ queries and three different distributions: $\boldsymbol{w}^{\mathrm{ChEMBL-177159}}$ and $\boldsymbol{w}^{\mathrm{ChEMBL-265935}}$ are actual distributions of $\overline{\mathrm{TI}}_{1,1,0.8}$ on ChEMBL obtained by querying two randomly selected fingerprints from ChEMBL, $\boldsymbol{w}^{\mathrm{rand}}$ is obtained by randomly selecting a value from $1, \ldots, k$ for $m = 5 \times 831$ times and dividing each observed frequency by m.

bioRxiv preprint doi: https://doi.org/10.1101/013995; this version posted January 19, 2015. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under aCC-BY-NC-ND 4.0 International license.

Tables

Table 1 The experimental success ratios of the user's guess based on the server's return and the prior distribution of true value, and success probability based only on a guess from the prior distribution. $\overline{\mathrm{TI}}_{1,1,0.8}$ ($k = 831$) is assumed and results are calculated for five different numbers of dummies ($n = 831, 831 \times 10^1$, 831×10^2, 831×10^3, 831×10^4) are used for three different distributions: $\boldsymbol{w}^{\mathrm{ChEMBL-177159}}$ and $\boldsymbol{w}^{\mathrm{ChEMBL-265935}}$ are actual distributions of $\overline{\mathrm{TI}}_{1,1,0.8}$ on ChEMBL obtained by querying two randomly selected fingerprints from ChEMBL, $\boldsymbol{w}^{\mathrm{rand}}$ is obtained by randomly selecting a value from $1, \ldots, k$ for $m = 5 \times 831$ times and dividing each observed frequency by m.

	$n = 831$	$n = 831 \times 10^1$	$n = 831 \times 10^2$	$n = 831 \times 10^3$	$n = 831 \times 10^4$	Ideal value
$\boldsymbol{w}^{\mathrm{ChEMBL-177159}}$	0.03552	0.01738	0.01101	0.01009	0.00977	0.00981
$\boldsymbol{w}^{\mathrm{ChEMBL-265935}}$	0.02991	0.01337	0.00903	0.00798	0.00784	0.00807
$\boldsymbol{w}^{\mathrm{rand}}$	0.00914	0.0041	0.00309	0.00279	0.00305	0.00289

Table 2 CPU time and communication size of secure similar compounds counter (SSCC) and those of general-purpose multi-party computation (GP-MPC).

	ChEMBL_1000	ChEMBL_Full
CPU time (s)		
SSCC (server)	0.69	167.19
SSCC (client)	1.53	172.37
GP-MPC (server)	4,075.15	6,081,252.70
GP-MPC (client)	4,366.18	6,472,194.42
Communication size (MB)		
SSCC (server → client)	2.24	265.33
SSCC (client → server)	0.03	0.03
GP-MPC (server → client)	42.50	63,195.68
GP-MPC (client → server)	2,128.00	3,165,738.41