Analysis of Information Leakage in Phenotype and Genotype Datasets

Arif Harmanci[1,2], Mark Gerstein[1,2,3]

1 Program in Computational Biology and Bioinformatics, Yale University, New Haven, CT, USA
2 Department of Molecular Biophysics and Biochemistry, Yale University, New Haven, CT, USA
3 Department of Computer Science, Yale University, New Haven, CT, USA
Corresponding author: Mark Gerstein mark.gerstein@gersteinlab.org

[[Fix \cite's]]

[[Fix Figure/Section References]]

[[Use past tense, all around]]

[[At most 10 equations in the online methods?]]

[[Methods only references]]?

[[Number of references? 26 now, too low?]]

[[Figure Captions]]

[[Read one more time]]

[[TIFF figures: http://www.jimmunol.org/site/pdf/PPTtoTIFF.pdf , RGB, 300 DPI]]

[[1000 => 1,000]]

ABSTRACT

Privacy is receiving much attention with the increase in the breadth and depth of personalized biomedical datasets. Studies on genomic privacy are focused on protection of variants. Molecular phenotype datasets can also contain substantial amount of sensitive information. Although there is no explicit genotypic information in them, subtle genotype-phenotype correlations can be used to statistically link the phenotype and genotype datasets. The links can then be used to characterize individuals' sensitive phenotypes. Here, we first develop a formalism for the quantification of individual characterizing information leakage in a linking attack. We analyze the tradeoff between the predictability of the genotypes and the amount of leaked information that can be used for individual characterization. Then we present a general three step procedure that can be used to practically instantiate a highly accurate attack. We develop a particular realization of the attack for outlier cases and study different aspects of the attack.

1 INTRODUCTION

Genomics has recently emerged as one of the major foci of studies on privacy. This can be attributed to high throughput biomedical data acquisition that bring about a surge of datasets[1,2]. Among these, molecular phenotype datasets, like functional genomics measurements, substantially grow the list of the quasi-identifiers[3] which may lead to re-identification and characterization[3–5]. In general, statistical analysis methods are used to discover genotype-phenotype correlations, which can be utilized by an

Deleted: .*
Deleted: 260 Whitney Avenue,
Deleted: 06520
Deleted: 260 Whitney Avenue,
Deleted: 06520
Deleted: 260 Whitney Avenue,
Deleted: 06520
Deleted: *
Deleted: authors
Deleted: pi@gersteinlab.org
Deleted: ¶
¶
[[HIGHLIGHTED PARTS TO BE MOVED TO SUPPLEMENTARY]]¶
Deleted: unprecedented
Deleted: , particularly personal genomics datasets. Moreover, considering the legislative plans for encouraging public data sharing in biomedical research fields, privacy will be the key consideration in designing data sharing mechanisms. Most studies
Deleted: in personal genomes.
Deleted: , however,
Deleted: individuals by identifying their
Deleted: and breaching privacy.
Deleted: and analysis
Deleted: in linking and
Deleted: show how one could
Deleted: an
Deleted: focusing on the most commonly available data sets, those of RNA-seq and eQTL.
Deleted: three step procedure showing how an attacker would select eQTLs, statistically predict the genotypes, and perform linking based on the predicted genotypes. The linking can be very accurate considering the high dimensionality of phenotypes. The linking attack becomes particularly easy to perform when one deals with outlier gene expression levels. To study this, we developed a
Deleted: this
Deleted: the
Deleted: quantified the amount of information leaked
Deleted: ¶
¶
Deleted: mainly
Deleted: the advancement of technologies for
Deleted: high throughput
Deleted: genomic and metabolomic
Formatted: Font: Not Italic
Deleted: (such as birth date, ZIP code, gender
Deleted:) for participating individuals, which can be used by a

adversary for linking the entries in genotype and phenotype datasets, and revealing sensitive information. The availability of a large number of correlations increases the possibility of linking[6,7].

With the initial genotype-phenotype studies, the protection of privacy of participating individuals emerged as an important issue. Several studies addressed "detection of a genome in a mixture" attacks, which can potentially detect participating indivuduals[8,9] (**Supplementary Fig. 1**). With large sets of seemingly independent datasets, the attacker can also link two or more datasets to pinpoint individuals in datasets and reveal sensitive information. One well-known example of these "linking attacks" is the attack that matched the entries in Netflix Prize Database and the Internet Movie Database (IMDB)[10], and revealed sensitive information. As the size and number of the genotype and phenotype datasets increase, number of potentially linkable datasets will increase, which can make similar scenarios a reality in genomic privacy (Supplementary Note). Several studies addressed multiple scenarios for individual re-identification and genome. In addition, different formalisms are proposed for protection of sensitive information (Supplementary Note). Different aspects of genomic privacy, pertaining linkability of high dimensional phenotype datasets to genotypes, are yet to be explored.

In this paper, we focus on characterizability of the individuals' sensitive information in the context of linking attacks, where the adversary exploits the genotype-phenotype correlations to link datasets and reveal sensitive information. In general, the high dimensional phenotype datasets harbor a number of phenotypes that contain sensitive information, like disease status, and other phenotypes, while not sensitive, may be used for linking to genotypes. Many public quantitative trait loci (QTL) datasets will inevitably help these linkings. Although each QTL has [11–13]. First we study quantification of characterizing information leakage versus risk of characterization. Then, we present a three step analysis framework for analysis of linking attacks. Then we show that the linking can be performed in a much simplified setting by just utilizing the outliers in the data. Compared to a previous publication[14] that relates to our study, we aim at providing quantification measures and privacy analysis frameworks and also demonstrate the accuracy of simplified linking attack under different scenarios. We finally discuss different strategies for risk management against the linking attacks.

2 RESULTS

2.1 Individual Characterization by Linking Attacks

There are three datasets in the context of the breach by linking attacks (**Fig. 1**). First dataset contains the phenotype information for a set of individuals. The phenotypes can include sensitive information such as disease status in addition to several molecular phenotypes such as gene expression levels. The second dataset contains the genotypes and the identities for another set of individuals. The third dataset contains correlations between one or more of the phenotypes in the phenotype dataset and the genotypes. Each entry contains a phenotype, a variant, and the degree to which these values are correlated. We will focus on the gene expression datasets as the representative phenotype dataset. As explained earlier, the abundance of gene expression-genotype correlation (eQTL) datasets makes these datasets most suitable for linking attacks.

Deleted: Many consortia, like GTex[10], ENCODE[11], 1000 Genomes[12], and TCGA[13], are generating large amount of personalized biomedical datasets. Coupled with the generated data, sophisticated analysis methods are being developed to discover correlations between genotypes and phenotypes, some of which can contain sensitive information like disease status. Although these correlations are useful for discovering how genotypes and phenotypes interact, they could also be utilized by an adversary in a linking attack for matching the entries in genotype and phenotype datasets. For example, when a phenotype dataset is available, the adversary can utilize the genotype-phenotype correlations to statistically predict the genotypes, compare the predicted genotypes with the entries in another dataset that contains genotypes. For the entries that are correctly matching, he/she can reveal sensitive phenotypes of the individuals and characterize them. Even when the strength of each genotype-phenotype correlation is not high, the availability of a large number of genotype-phenotype correlations increases the scale of linking. In fact, an adversary can perform correct linking with relatively small number of genotypes[14,15].¶

Different aspects of privacy have been intensely studied. Recently, genomic privacy is receiving much attention as a result of the deluge of personalized genomics datasets that are being generated[16,17]. With the increase in the number of large scale genotyping and phenotyping studies, the protection of privacy of participating individuals emerged as an important issue. Homer et al[18] proposed a statistical testing procedure that enables testing whether a genotyped individual is in a pool of samples, for which only the allele frequencies are known. Im et al[19] showed that given the genotypes of a large set of markers for an individual, an attacker can reliably predict whether the individual participated in a QTL study or not. These attacks, which we refer to as "detection of a genome in a mixture", are one type of attacks on privacy (Fig S6). There is yet another important attack where the attacker links two or more datasets to pinpoint individuals in datasets and reveal sensitive information. One well-known and illustrative example of these "linking attacks", although not in a genomic context, is the linking attack that matched the entries in Netflix Prize Database and the Internet Movie Database (IMDB)[20]. For research purposes, Netflix released an anonymized dataset of movie ratings of thousands of viewers, which they thought was secure as the viewers' names were removed. However, Narayanan et al[20] used IMDB database, a seemingly unrelated and very large database of movie viewers, linked the two databases, and revealed identities and personal information (movie history and choices) of many viewers in the Netflix database. The fact that Netflix and IMDB host millions of individuals in their databases renders the question of detection of an individual in these database irrelevant since any random individual is very likely to be in one or both of these databases but the focus of attacks turns to matching individuals in the databases. Consequently, as the databases grow, the attacks for detection of an individual in a database become unimportant and the linking attacks become more admissible in order to characterize individuals' sensitive information. In the genomic privacy context, as the size and number of the genotype and phenotype datasets ...

Deleted: Overview of the

Deleted: Scenario

Deleted: Figure 1a illustrates the general privacy breaching scenario that is considered.

Deleted: .

Deleted: In this dataset, each

Deleted: In order to formulate and demonstrate the results, we

In an eQTL dataset, genes and variants and a correlation coefficient, denoted by ρ, between the expression levels and genotypes are reported. The absolute value of ρ indicates the strength of association between the eQTL genotype and the eQTL expression level. The sign of ρ represents the direction of association, i.e., which homozygous genotype corresponds to higher expression levels. We assume that the attacker captures this relation with a prediction model and builds the *a posteriori* distribution of the eQTL genotypes given the eQTL expression levels (**Fig. 2a**, **Supplementary Fig. 2**). As a representative dataset for reporting results, we utilize the eQTLs, and gene expression levels from the GEUVADIS project[15] and the genotypes from the 1000 Genomes Project[16] Genotypes Predictability and Information Leakage

We assume that the attacker will behave in a way that maximizes his/her chances of correctly characterizing the most number of individuals. Thus, he or she will try and predict the genotypes, using the phenotype measurements, for the largest set of variants that he or she believes he or she can predict correctly. The most obvious way that the attacker does this is by first sorting the genotype-phenotype pairs with respect to decreasing strength of correlation then predicting the genotypes for each variant (**Supplementary Fig. 3**). The attacker will encounter a tradeoff: As he or she goes down the list, more individuals can be characterized (more genotypes can characterize more individuals) but it also becomes more likely that he or she makes an error in the prediction since the correlation decreases going down the list. This tradeoff can also be viewed as the tradeoff between precision (fraction of the linkings that are correct) and recall (fraction of individuals that are correctly linked). In this section we will propose two measures to quantify this tradeoff.

The attacker aims to correctly characterize n_e individuals in the expression dataset among n_v individuals in the genotype dataset. In order to correctly characterize an individual, given the individual's expression levels, the attacker should predict the genotypes for a set of selected eQTLs correctly such that the predicted set of genotypes are not shared by more than one individual, i.e., the predicted genotypes can be matched to the correct individual. In other words, the joint frequency of the set of predicted genotypes should be $\frac{1}{n_v}$. Equivalently, if the genotypes contain at least $\log_2(n_v)$ bits of information, the individual is vulnerable to characterization of his/her phenotypes. Following this idea, we quantify the leakage of individual characterizing information (ICI) from a set of correctly predicted variant genotypes using the frequencies of the genotypes in the population (**Fig. 3**). *ICI* can be interpreted as a quantification of how rare the predicted genotypes are. The attacker aims to predict as many eQTLs as possible such that *ICI* for the predicted genotypes is at least $\log(n_v)$. *ICI* can also be interpreted as the number of rare SNP genotypes that an individual harbors.

In order to quantify the (correct) predictability of the eQTL genotypes from expression levels, it is necessary to uniformly estimate predictability for a list of eQTLs. Although ρ is a measure of predictability, it is computed differently in different studies and there is no easy way to combine these correlation values when we would like to estimate the joint predictability of multiple eQTL genotypes. We will utilize the exponential of negative conditional entropy of genotypes given gene expression levels for measuring predictability

Moved down [1]: The eQTL dataset is composed of a list of

Deleted: Figure 1b illustrates the eQTL, expression, and genoty

Moved down [2]: represents the genotype ($v_{k,j} \in \{0,1,2\}$) of k

Formatted

Deleted: The absolute value of $\rho(E_k, V_k)$ indicates the strength

Formatted

Formatted

Formatted

Formatted

Deleted: j^{th} individual. The coding of the genotypes from

Formatted

Moved down [3]: eQTL entry are distributed randomly over the

Deleted: The k^{th} row of e, $\boldsymbol{e_k}$, contains the gene expression

Formatted

Formatted

Formatted

Formatted

Formatted

Formatted

Deleted: /…or she will try and predict the genotypes, using the

Deleted: In the context of the linking attack, the

Deleted: he/she should select a set of eQTLs that he/she believ

Deleted: We can rephrase this condition as following in

Formatted

Deleted: statement

Formatted

Deleted: can

Formatted

Deleted: eQTL

Formatted

Deleted: as

Formatted

Deleted: logarithm of their joint frequency. Assuming that

Formatted

Deleted: of different eQTLs (See Section 5) are independent fr

Formatted

Deleted: , and *ICI* denotes the total individual characterizing

Formatted

Deleted: increases as the frequency of the variant's genotype

Formatted

Deleted: maximize the amount of *ICI*, the attacker will aim at

Deleted: whose genotypes are the most correctly predictable

Deleted: (E_k, V_k) is a measure of predictability, …t is computed

Formatted

Deleted: entropy of the

Formatted

Deleted: genotype distribution

Formatted

Deleted: . Given the expression levels

Formatted

Deleted: j^{th} individual, we compute the

Formatted

, which we denote with π (**Fig. 3**). In the case of high predictability, the conditional entropy is close to 0 (very small randomness left in the genotype distribution). Taking the exponential of negative of the entropy converts the entropy to average probability of correct prediction of the genotype, which is close 1, indicating high predictability (Online Methods).

We first considered each eQTL and evaluated the genotype predictability versus the characterizing information leakage. We use the GEUVADIS dataset as a representative dataset for this computation. We computed, for each eQTL, average π and average *ICI* over all the individuals (**Fig. 4a**). Most of the data points are spread along the diagonal; The eQTL variants with high major allele frequencies have high predictability and low *ICI* and vice versa for eQTL variants with lower major allele frequencies (**Fig. 4b**). This is expected because the genotypes of the high frequency variants can be predicted, on average, easily (most individuals will harbor one dominant genotype) and consequently does not deliver much characterizing information and vice versa for the eQTLs with smaller major frequency alleles. The eQTLs with high correlation (**Fig. 4b**) deviate from the diagonal, compared to shuffled data (**Fig. 4c**), with high ICI and high predictability, which will be targeted by an attacker.

To estimate *ICI* leakage with multiple genotypes, we first sorted the eQTLs with respect to the reported correlation, $|\rho|$. Then for top n=1,2,3,…,20 eQTLs, we estimated mean π and mean *ICI* over all the samples as illustrated in Fig S2a. Inspection of mean π versus mean *ICI* for each n (**Fig. 4d**) enables us to estimate the number of vulnerable individuals at different predictability levels. For example, at 20% predictability, there is approximately 8 bits of ICI leakage. At this level of leakage, the adversary can correctly link all individuals, on average with 20% chance, in a sample of $2^8 = 256$ individuals. At 5% predictability, the leakage is 11 bits and the characterizable sample size is $2^{11} = 2048$ individuals, which can be interpreted as a higher risk of characterizability. In addition, auxiliary information can be added into leaking information. For example, gender information, which can be predicted with high accuracy from many molecular phenotype datasets brings 1 bits of auxiliary information.

2.2 Analysis Framework for Individual Characterization

We present a three step framework for individual characterization in the context of linking attacks (**Fig. 2c**). The input is the phenotype measurements for an individual. The aim of the attacker is to link the disease state of the individual to the correct individual in the genotype dataset. In the first step, the attacker selects the QTLs, which will be used in linking. The selection of QTLs can be based on different criteria. As described in the previous section, the most accessible criterion is selection based on the absolute strength of association between the phenotypes and genotypes. In the case of eQTLs, this is the reported correlation coefficient, $|\rho|$. The second step is genotype prediction for the selected QTLs using a prediction model. We assume that the attacker performs maximum *a posteriori* genotype prediction using an estimate of the joint genotype-phenotype distribution. The third and final step of individual characterization is comparison of the predicted genotypes to the genotypes of the n_v individuals in genotype dataset to identify the individual that matches best to the predicted genotypes. In this step, the attacker links the predicted genotypes to the individual in the genotype dataset (Online Methods).

Deleted: , and there is little…(very small randomness left in the genotype distribution….. Taking the exponential of negative of the entropy converts the entropy to average probability of correct prediction of the genotype. In the most predictable case (conditional entropy close to 0), π …

Formatted: Font color: Text 1

Deleted: to

Formatted: Font color: Text 1

Deleted: very …igh predictability (Refer to…nline Methods Section 4.1 for more details …

Deleted: (Refer for Section 5).

Deleted: , which is plotted in

Formatted: Font: Bold, Font color: Text 1

Deleted: 2b…. **4a**). Most of the data points are spread along the diagonal, which indicate that there is a natural tradeoff between correct predictability and *ICI* leakage…. The eQTL variants with high frequency…major allele frequencies have high predictability and …

Formatted …

Deleted: 2b, left

Formatted: Font color: Text 1

Deleted: . The genotypes…and vice versa for the eQTLs with …

Formatted …

Deleted: 2b, right

Formatted: Font color: Text 1

Deleted: . In principle, the adversary… which will aim at …

Deleted: [[FOLLOWING CAN BE MOVED AND REPHRASED]]¶ …

Formatted: Font: Not Italic

Deleted: We then plotted

Deleted: which is shown in

Formatted: Font: Bold, Font color: Text 1

Deleted: 2c. From the plot, we can first

Deleted: These estimates are useful when releasing QTL datas…

Deleted: A General

Deleted: Analysis of

Deleted: [[REPHRASE FOLLOWING]]¶ …

Formatted: Font color: Auto

Deleted: correctly …ink the disease state of the individual to th…

Deleted: gradient or the absolute

Deleted: $\rho(E_k, V_k)$|. In our analysis, we evaluate the effect of …

Deleted: For general applicability of our analysis we are assum…

Formatted: Font: Italic

Deleted: (MAP) …enotype. In this …

Formatted: Font: Italic

Deleted: ,…using an estimate of the attacker assigns the…oint …

Deleted: with the smallest number of mismatches compared t…

2.3 Individual Characterization by Linking Attacks

Using the three step approach, we first evaluated the accuracy of linking when the attacker utilizes MAP genotype prediction where the attacker builds the posterior distribution of genotypes (given expression levels). The attacker utilizes the absolute value of correlation coefficient to select eQTLs (from the GEUVADIS Project[15]) for linking. We assume that the attacker can perfectly recover the joint genotype-expression distribution and uses it to build posterior distribution of genotypes (**Supplementary Fig. 4**). The linking accuracy is above 95% and gets close to 100% when auxiliary information is available (**Fig. 5a**).

In general, correct reconstruction of the *a posteriori* distribution of the genotypes given expression levels may not be possible because the knowledge of only the genotype-phenotype correlation coefficient via eQTLs is not enough to regenerate the *a posteriori* distribution of genotypes given the expression levels. The attacker can, however, utilize a priori knowledge about the relation between gene expression levels and genotypes and build the joint genotype-expression distributions using models with varying complexities and parameters (Online Methods, Supplementary Note, **Supplementary Fig. 5**). We focus on a highly simplified model where the attacker exploits the knowledge that the eQTL genotypes and expression levels are correlated such that the extremes of the gene expression levels (highest and smallest expression levels) are observed with extremes of the genotypes (homozygous genotypes). We used a measure, termed extremity, to quantify the outlierness of expression levels (Online Methods, Supplementary Note, **Supplementary Fig. 6**). Based on the extremity of expression level and the gradient of association, the attacker can coarsely estimate the genotype (**Fig. 2b**). The prediction methodology assigns zero probability to heterozygous genotype and therefore assigns only homozygous genotypes to variants. With this approach, the genotype prediction accuracy increases, as expected, with increasing absolute correlation threshold (**Fig. 5b**). To perform linking with this model, we utilized the correlation based eQTL selection in step 1, then extremity based genotype prediction in step 2. In the step 3, we evaluated two distance measures for linking the predicted genotypes to the individuals in genotype dataset. First is based on comparison of the predicted genotypes to all the genotypes in genotype dataset. Second is based on comparison of the predicted genotypes to only the homozygous genotypes in the genotype dataset, which is motivated by the fact that the attacker only predicts homozygous genotypes in the genotype prediction step (Online Methods). For each measure, the attacker links the predicted genotypes to the individual whose genotypes minimize the selected distance measure (**Supplementary Fig. 8**). More than 95% of the individuals (**Fig. 5c,d**) are vulnerable for most of the parameter selections. When the auxiliary information is present, the fraction of vulnerable individuals increases to 100% for most of the eQTL selections. These results show that linking attack with extremity based genotype prediction, although technically simple, can be extremely effective in characterizing individuals. We will focus on homozygous genotype matching based distance computation in the rest of the paper for simplicity of presentation.

To test reproducibility of linking accuracy, we generated an eQTL training set (210 individuals) and identified eQTLs. The linking of expression and genotype datasets for remaining 211 individuals (testing set) is around 95% (**Supplementary Fig. 9a**), which shows that the linking attack is still effective when testing and training sets do not match. In addition, extremity based linking requires much less

Deleted: <#>Fraction of Vulnerable Individuals with MAP Genotype Prediction¶

Deleted: ¶

Formatted: Font color: Text 1

Deleted: See

Formatted: Font color: Text 1

Deleted: Section 4.8). Even though the genotype prediction wi

Deleted: here. We will assume

Deleted: allelic effects on expression are additive and

Deleted: Therefore, given

Deleted: coarsely the joint distribution of the genotypes and

Formatted: Not Highlight

Deleted: is observed at medium levels of expression. The

Formatted: Not Highlight

Deleted: value 2, and otherwise, genotype value 0. Finally, the

Deleted: is zero. This is expected since we are focusing on the

Formatted: Not Highlight

Deleted: method, i.e., the *a posteriori*

Formatted: Not Highlight

Formatted: Not Highlight

Deleted: can be interpreted as a rank correlation between the

Formatted: Not Highlight

Deleted: . Figure 5b shows the accuracy of genotype prediction

Formatted: Not Highlight

Formatted: Not Highlight

Deleted: correlation threshold. The slight decrease of genotyp

Formatted: Font: Bold

Deleted: 3) and evaluated the fraction of characterizable

Deleted: order to demonstrate

Deleted: utility of the 3-

Deleted: analysis framework;

Deleted: different

Formatted: Not Highlight

Formatted: Not Highlight

Deleted: in the 3rd step of the attack

Formatted: Not Highlight

Deleted: (See Methods Section 4.5 for details). The motivation

Deleted: . Figure 5c and 5d show the fraction of vulnerable

Deleted: for both distance measures. The homozygous genoty

Deleted: gender and/or population information is present as

Deleted: (red and green plots),

Deleted: The previous results show that extremity based linkin

information compared to previously proposed method[14] (Supplementary Note, **Supplementary Table 1c**). To study these further, we evaluated whether the attacker can estimate the reliability of the linkings so as to focus on highly reliable linkings. We observed that the measure we termed, *first distance gap*, denoted by $d_{1,2}$ (Online Methods), serves as a good reliability estimate for each linking. We computed the positive predictive value (PPV) versus sensitivity of the linkings in the testing set with changing $d_{1,2}$ threshold (Online Methods). Compared to random sortings, the attacker can link a large fraction (79%) of the individuals at a PPV higher than 95% (**Supplementary Fig. 9b**).

In order to study the effect of genotype dataset size, we simulated the genotypes of the eQTLs (from the training set) for 100,000 individuals using genotype frequencies from 1000 Genomes Project (Supplementary Note). We then merged simulated dataset with testing dataset (of 100,211 individuals) and observed that linking of testing expression samples to large genotype dataset is around 9% accurate (**Supplementary Fig. 9c**). In addition, the attacker can correctly link 55% of individuals with more than 95% PPV (**Supplementary Fig. 9d**).

We next separated the GEUVADIS samples to 5 populations and identified eQTLs for each population and performed linking using population specific eQTLs **(Supplementary Table S1a)**. The eQTLs from close European populations (CEU, GBR, TSI, FIN) enable high linking accuracy (higher than 95%) for European populations. When the eQTLs are identified from an African population (YRI), the linking accuracies in European populations are smaller. To also study the mismatch between the tissues where linking and eQTL discovery is performed, we downloaded the eQTLs for 6 tissues from the GTex Project[13] and linked the GEUVADIS expression samples to genotypes samples. The accuracy is generally high (>75%) and is highest for blood eQTLs, as expected (**Supplementary Table 1b**). We also observed that when close relatives of individuals in the expression dataset (30 CEU trio dataset from the HAPMAP project[17,18]) are in the genotype dataset, the linking attack assigns much lower ranks to relatives compared to the random individuals (Supplementary Note, **Supplementary Fig. 10**), which may cause privacy concerns for the family of individuals.

3 DISCUSSION

In genomic privacy, it is necessary to consider the basic premise of sharing any type of personal information: There is always an amount of leakage in the sensitive information[19]. In addition, as shown by previous studies, we often cannot propose black-and-white solutions to problems in privacy which mainly roots from the multifaceted nature of privacy. We believe these make it necessary for the genomic data sharing and publishing mechanisms to incorporate statistical quantification methods to objectively quantify risk estimates before the datasets are released. The quantification methodology and the analysis frameworks presented here and in future studies can be applied for analysis of the information leakage in the datasets where the correlative relations can be exploited for performing linking attacks (Supplementary Note, **Supplementary Fig. 11**).

Our study focuses on the individual privacy breaches in the context of linking attacks, where an individual's existence in two seemingly independent databases (e.g., phenotype and the genotype) can cause a privacy concern when an attacker links statistically the databases using the a priori information

Moved down [4]: . In addition, as shown by previous studies,
Deleted: Increasing pace of data generation and the policies to ...
Moved down [5]: Our study focuses on the individual privacy
Moved down [6]: . The obvious risk management strategy
Moved down [7]: . We believe new studies should address
Moved down [20]: Finally, we utilize compute the Shannon
Moved down [22]: We propose using the difference between
Deleted: The analysis of tradeoff between predictability and ...
Moved down [8]: Characterizing Information and
Moved down [9]: and ICI denotes the individual characterizing
Moved down [10]: measure the predictability of eQTL
Moved down [11]: where π denotes the predictability of V_k
Moved down [12]: individual, given the expression levels $e_{k,j}$
Moved down [13]: In addition, this measure is guaranteed to
Moved down [19]: where $I(.)$ is an indicator function for
Moved down [15]: These probabilities are estimated from the
Moved down [21]: Following the previous section, the attacker
Moved down [14]: eQTL variant genotypes and $p(V_k = v)$
Moved down [16]: where $p(V_k = v|E_k = e_{k,j})$ represents the
Moved down [17]: The binning is done for each gene by first
Moved down [18]: where $\min(\boldsymbol{e_k})$ and $\max(\boldsymbol{e_k})$ represents the
Moved down [23]: We first generated the testing and training
Moved down [24]: Modeling of Genotype-Phenotype
Moved down [25]: In the second step of the linking attack, the
Moved down [26]: The main aim of attacker is to maximize the
Moved down [27]: Assuming that mean and variance are
Moved down [28]: This decreases the number of parameters
Moved down [29]: This model requires 4 parameters to be
Moved down [30]: In this model, the prediction only assigns
Deleted: before the datasets are released. This is recently ...
Deleted: The fact that the available molecular phenotypes are ...
Deleted: [24,30,31,57]
Deleted: [59] of the event that RV takes the value $g_{k,j}$:¶ ...
Deleted: ¶ ...
Deleted: ¶ ...
Deleted: For the j^{th}
Deleted: $\pi(\{V_k\},\{E_k = e_{k,j}\}) = \exp(-H(\{V_k\} \mid \{E_k = e_{k,j}\}))$¶ ...
Deleted: [59]:¶ ...
Deleted: This probability can be also interpreted as the ...
Deleted: [59] of the genotypes given the gene expression levels. ...
Deleted: [60] to choose the number of bins. This rule states that ...
Deleted: ¶ ...
Deleted: ¶ ...
Deleted: ¶ ...
Deleted: Maximum *a posteriori* (MAP) Genotype Predicti ...
Deleted: j^{th} individual.
Deleted: ¶ ...
Deleted: On
Deleted: for Genotype Prediction in Linking Attacks¶
...
...
...
...
...
Moved (insertion) [4]
Moved (insertion) [5]

about correlation of different entries in the databases. We showed that the attacks can be effective even when the eQTLs are identified in tissues or populations that are different than the samples being linked, and also the attacks can target not only the linked individual but their families, which increases the impact of breach. The obvious risk management strategy against these attacks is restricting access to the phenotype datasets. The statistical techniques like k-anonymization and differential privacy can also be utilized. These, however, have associated drawbacks about loss of biological utility, and high computational complexity. Moreover, some studies also demonstrated that there are still risks associated with linkability of the anonymized data[20–23]. We believe new studies should address protection and risk management strategies for serving utility maximized and privacy aware high dimensional phenotype datasets.

Moved (insertion) [6]

Moved (insertion) [7]

Deleted: Equations 3 and 4-6). ¶

Accuracy, Positive Predictive Value, and Sensitivity of Linkings¶

The accuracy of the linkings are evaluated with respect to 3 measures. First is the overall accuracy of linkings, which measures the fraction of correctly linked individuals. This is simply among the all the linkings that we performed, the fraction of correct linkings

...

4 DATASETS [[ACCESSION NUMBERS??]]

The normalized gene expression levels for 462 individuals and the eQTL dataset are obtained from gEUVADIS mRNA sequencing project[15]. The eQTL dataset contains all the significant (Identified at most 5% false discovery rate) gene-variant pairs with high genotype-expression correlation. To ensure that there are no dependencies between the variant genotypes and expression levels, we used the eQTL entries where gene and variants are unique. In other words, each variant and gene are found exactly once in the final eQTL dataset (Section S4). The genotype, gender, and population information datasets for 1092 individuals are obtained from 1000 Genomes Project[16]. For 421 individuals, both the genotype data and gene expression levels are available. For tissue analysis, the publicly available significant eQTLs for 6 tissues that are computed by the GTex project are downloaded from the GTex Portal.

[[Accession numbers]]

Deleted: The normalized gene expression levels for 462 individuals and the eQTL dataset are obtained from gEUVADIS mRNA sequencing project[61]. The eQTL dataset contains all the significant (Identified at most 5% false discovery rate) gene-variant pairs with high genotype-expression correlation. To ensure that there are no dependencies between the variant genotypes and expression levels, we used the eQTL entries where gene and variants are unique. In other words, each variant and gene are found exactly once in the final eQTL dataset (See Section S4 for a detailed discussion). The genotype, gender, and population information datasets for 1092 individuals are obtained from 1000 Genomes Project[12]. For 421 individuals, both the genotype data and gene expression levels are available. For tissue analysis, the publicly available significant eQTLs for 6 tissues that are computed by the GTex project are downloaded from the GTex Portal.¶

5 ACKNOWLEDGEMENTS

Authors would like to thank Akdes Serin Harmanci for constructive comments and discussions on study design and running of external tools. Authors also would like to thank the anonymous reviewers for constructive criticism that made the manuscript complete and comprehensive.

6 AUTHOR CONTRIBUTIONS

A.H. designed the study, gathered datasets, performed experiments, and drafted the manuscript. M.G. conceived the study, oversaw the experiments and wrote the manuscript. Both authors approved final manuscript.

Deleted: AH

Deleted: MG

Authors declare no conflict of financial interest.

7 REFERENCES

Moved (insertion) [31]

1. Sboner, A., Mu, X., Greenbaum, D., Auerbach, R. K. & Gerstein, M. B. The real cost of sequencing: higher than you think! *Genome Biol.* **12,** 125 (2011).

2. Rodriguez, L. L., Brooks, L. D., Greenberg, J. H. & Green, E. D. The Complexities of Genomic Identifi ability. *Science (80-.).* **339,** 275–276 (2013).

3. Sweeney, L., Abu, A. & Winn, J. Identifying Participants in the Personal Genome Project by Name. *SSRN Electron. J.* 1–4 (2013). doi:10.2139/ssrn.2257732

4. Sweeney, L. *Uniqueness of Simple Demographics in the U.S. Population, LIDAP-WP4. Forthcom. B. entitled, Identifiability Data.* (2000).

5. Golle, P. Revisiting the uniqueness of simple demographics in the US population. in *Proc. 5th ACM Work. Priv. Electron. Soc.* 77–80 (2006). doi:http://doi.acm.org/10.1145/1179601.1179615

6. Pakstis, A. J. *et al.* SNPs for a universal individual identification panel. *Hum. Genet.* **127,** 315–324 (2010).

7. Wei, Y. L., Li, C. X., Jia, J., Hu, L. & Liu, Y. Forensic Identification Using a Multiplex Assay of 47 SNPs. *J. Forensic Sci.* **57,** 1448–1456 (2012).

8. Homer, N. *et al.* Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays. *PLoS Genet.* **4,** (2008).

9. Im, H. K., Gamazon, E. R., Nicolae, D. L. & Cox, N. J. On sharing quantitative trait GWAS results in an era of multiple-omics data and the limits of genomic privacy. *Am. J. Hum. Genet.* **90,** 591–598 (2012).

10. Narayanan, A. & Shmatikov, V. Robust de-anonymization of large sparse datasets. in *Proc. - IEEE Symp. Secur. Priv.* 111–125 (2008). doi:10.1109/SP.2008.33

11. Xia, K. *et al.* SeeQTL: A searchable database for human eQTLs. *Bioinformatics* **28,** 451–452 (2012).

12. Consortium, T. G. The Genotype-Tissue Expression (GTEx) project. *Nat. Genet.* **45,** 580–5 (2013).

13. Ardlie, K. G. *et al.* The Genotype-Tissue Expression (GTEx) pilot analysis: Multitissue gene regulation in humans. *Science (80-.).* **348,** 648–660 (2015).

14. Schadt, E. E., Woo, S. & Hao, K. Bayesian method to predict individual SNP genotypes from gene expression data. *Nat. Genet.* **44,** 603–608 (2012).

15. Lappalainen, T. *et al.* Transcriptome and genome sequencing uncovers functional variation in humans. *Nature* **501,** 506–11 (2013).

16. The 1000 Genomes Project Consortium. An integrated map of genetic variation. *Nature* **135,** 0–9 (2012).

17. Stranger, B. E. *et al.* Relative impact of nucleotide and copy number variation on gene expression phenotypes. *Science* **315,** 848–853 (2007).

18. The International HapMap 3 Consortium. Integrating common and rare genetic variation in diverse human populations. *Nature* **467,** 52–8 (2010).

19. Narayanan, A. *et al. Redefining Genomic Privacy: Trust and Empowerment*. *bioRxiv* (2014). doi:10.1101/006601

20. Machanavajjhala, A., Kifer, D., Gehrke, J. & Venkitasubramaniam, M. L -diversity. *ACM Trans. Knowl. Discov. Data* **1,** 3–es (2007).

21. Ninghui, L., Tiancheng, L. & Venkatasubramanian, S. t-Closeness: Privacy beyond k-anonymity and ℓ-diversity. in *Proc. - Int. Conf. Data Eng.* 106–115 (2007). doi:10.1109/ICDE.2007.367856

22. Wong, R. C.-W. W., Fu, A. W.-C. C., Wang, K. & Pei, J. Minimality attack in privacy preserving data publishing. in *Proc. 33rd Int. Conf. Very large data bases* 543–554 (2007). at <http://dl.acm.org/citation.cfm?id=1325851.1325914>

23. Fredrikson, M., Lantz, E., Jha, S. & Lin, S. Privacy in Pharmacogenetics: An End-to-End Case Study of Personalized Warfarin Dosing. in *23rd USENIX Secur. Symp.* (2014). at <http://www.biostat.wisc.edu/~page/WarfarinUsenix2014.pdf>

24. Cover, T. M. & Thomas, J. A. *Elements of Information Theory*. *Elem. Inf. Theory* (2005). doi:10.1002/047174882X

25. Herbert A. Sturges. The Choice of a Class Interval. *J. Am. Stat. Assoc.* **21,** 65–66 (1926).

26. Shabalin, A. A. Matrix eQTL: Ultra fast eQTL analysis via large matrix operations. *Bioinformatics* **28,** 1353–1358 (2012).

8 FIGURE LEGENDS

Figure 1: Illustration of the linking attack. (a) Phenotype dataset contains q different phenotype measurements and the HIV Status for a list of individuals. Genotype dataset contains the variants genotypes for m individuals. Phenotype-Genotype correlation datasets contains q phenotypes, variants, and their correlations. The attacker does genotype prediction for all the variants. The attacker then links the phenotype dataset to the genotype dataset by matching the genotypes. The linking potentially reveals the HIV status for the subjects in the genotypes dataset. The IDs and HIV Status are colored to illustrate how the linking combines the entries in the two datasets. The non-shaded columns are used for linking.

Figure 2: Illustration of genotype-expression associations and linking attacks (a) Schematic representation of genotype and expression associations. Genotype (y-axis) and expression (x-axis) are correlated, indicated by line fit and ρ. The rectangles represent conditional distribution of expression given genotype values. (b) Illustration of extremity based genotype prediction. Expression range is divided into two equal ranges (separated by e_{mid}). The blue rectangles represent the distribution used for prediction. Given the distribution of expression (tri-model distribution on right), the positive extremity is assigned genotype 2, and negative extremity is assigned genotype 0. (c) Three step linking process. First step is selection of phenotypes and genotypes to be used in linking. Second step is prediction of genotypes. Last step is linking of predicted genotypes to the genotype dataset.

Figure 3: Illustration of individual characterizing information (ICI) and correct predictability of genotypes. (a) Graphical representation of ICI formulation. ICI for a set of n variant genotypes is computed in terms of population genotype frequencies. Each genotype contributes to ICI additively with the logarithm of reciprocal of the genotype frequency (illustrated by the genotype distributions). (b) Graphical representation of π. Given the joint distribution of genotype and expression (shown below), the conditional distribution of genotypes given expression level e is computed. The exponential of the conditional distribution entropy is used for computing the predictability.

Figure 4: ICI versus π for each eQTL. Plots show, for each eQTL, the information leakage (x-axis) versus correct genotype predictability (y-axis). For each eQTL, the estimated ICI leakage and genotype predictability are plotted. The dots are colored with respect to the major allele frequency (a) and with respect to absolute correlation of the eQTL (b). ICI versus π for shuffled data (red) is compared to the real dataset (blue) in (c).

Figure 5: Accuracy measures for linking attacks. (a) Linking accuracy with MAP genotype predictions. Absolute correlation threshold (x-axis) versus fraction of vulnerable individuals (y-axis). The yellow arrow indicates the maximized position of linking accuracy. Red, green, and cyan plots show linking accuracy with gender, population, and gender + population as auxiliary information. (b) The genotype prediction accuracy. The genotype prediction accuracy (y-axis) of with changing absolute correlation

Deleted: CAPTIONS

Formatted: Highlight

Deleted: (b) Illustration of the expression and genotype datasets. Variant genotype dataset contains the genotypes for q eQTL variants for n_v individuals. j^{th} entry for k^{th} eQTL is denoted by $v_{k,j}$. Similarly, the expression dataset contains the expression levels for q genes. The k^{th}expression level for j^{th} individual is denoted by $e_{k,j}$. The variant genotypes for k^{th} variant is distributed over samples with distribution specified by the random variable V_k. Likewise, the expression levels for k^{th} gene is distributed per random variable E_k. These random variables are correlated with each other with correlation coefficient, denoted by $\rho(E_k, V_k)$ (right).

Formatted: Highlight

Deleted: ***Figure 2:*** Quantification of ICI and correct genotype predictability (a) Adversary's genotype prediction strategy. The phenotype-genotype correlations ρ_1, ρ_2,... are sorted with respect to decreasing absolute value, as shown on each line. For a selected set of n variants, the genotypes are predicted using the phenotypes. The green and red individuals on the right represent the vulnerable and non-vulnerable individuals, respectively. (b)

Formatted: Highlight

Deleted: left) and with respect to absolute correlation of the eQTL (right). (c) Average predictability versus average individual characterizing information leakage. For the top 20 eQTLs, the plot shows the distribution of average predictability and average ICI leakage for the top eQTLs. The number of eQTLs that are used for computing the values at each point are shown next the the point ...

Formatted: Highlight

Deleted: ***Figure 3:*** The figure illustrates three steps of the linki ...

Formatted: Highlight

Deleted: and vulnerable fraction. (a) The number of selected ...

Formatted: Highlight

Deleted: correctly predicted eQTL (red) are shown. At each ...

Formatted: Highlight

Deleted: with MAP genotype prediction. X-axis shows the ...

Formatted: Highlight

Deleted: shows the fraction of vulnerable individuals. At ...

Formatted: Highlight

Deleted: . The red

Formatted: Highlight

Deleted: lines

Formatted: Highlight

Deleted: the fraction of vulnerable individuals when

Formatted: Highlight

Deleted: and

Formatted: Highlight

Deleted: information, respectively, are available

Formatted: Highlight

Deleted: ¶ ...

threshold (x-axis). (c) Linking accuracy with extremity based linking with all genotypes. (d) Linking accuracy with extremity based linking with homozygous genotypes.

[[FOLLOWING ARE NOT UPDATED, YET]]

Table S1: Linking accuracy of extremity based linking attack using the eQTLs are identified in different populations and different tissues. (a) The table shows the linking accuracies (for populations shown in the rows) when the eQTLs that are identified using data (indicated in each column) from different populations. (b) The linking accuracy of individuals in GEUVADIS project when eQTLs identified from different tissues are used in linking.

Table S2: Linking attack accuracy comparison. The table shows linking accuracy for Schadt et al and extremity based linking attack methods. Each row corresponds (for Schadt et al Method) to a different number of data points in the training datasets that is input to Schadt et al method.

Supplementary Figure 1: Schematic comparison of linking attacks (Left) and detection of a genome in a mixture attacks (Right). Each box in the figure represents a dataset in the form of a matrix. Multiple boxes next to each other correspond to concatenation of matrices. Linking attacks aim at linking genotype and phenotype datasets. The phenotype datasets contain both "predicting" phenotypes and other phenotypes, some of which can be sensitive. The attacker first predict genotypes for each of the predicting phenotype. The predicted genotypes are then compared with the genotypes in the genotype dataset. After the linking, all the datasets are concatenated where the identifiers can be matched to the sensitive phenotypes. Different colors indicate how the linking merges different information. The detection of a genome in a mixture attacks start with a genotype dataset. The attacker gets access to the statistics of a GWAS or genotyping dataset (for example, regression coefficients or allele frequencies). Then the attacker generates a statistic and tests it against that of a reference population. The testing result can be converted into the study membership indicator (attended/not attended) which shows whether the tested individual was in the study cohort or not.

Supplementary Figure 2: Representation of the eQTLs. (a) The average ICI leakage versus the genotype predictability is shown for real (red) and shuffled (blue) eQTL dataset is shown. (b) The absolute correlation versus predictability is shown.

Figure shows the attacker's presumed strategy for linking attack. (a) The phenotype and variant pairs are sorted with respect to decreasing absolute correlations values. For the top *n* pairs, joint predictability and ICI are computed. (b) Illustration of prior, joint, and posterior distributions of genotypes and expression levels. Leftmost figure shows the distribution of genotypes over the sample set, which is labelled as the prior distribution. Middle figure shows the joint distribution of genotypes and expression levels. Notice that there is a significant negative correlation between genotype values and the

Deleted: attack (a) Figure illustrates the extremity based genotype prediction. The joint distribution of expression levels and genotypes is shown on left. Given the relation between expression and genotypes, the lower expression levels (Labelled with "Negative Extremity" shown in red ellipse on left) are assigned the genotype "TT" and higher expression levels (Labelled with "Posit ...

Formatted: Highlight

Deleted: only matching

Formatted: Highlight

Formatted: Highlight

Deleted: ***Figure 6:*** The accuracy of linking on the testing datase ...

Formatted: Highlight

Moved down [32]: (a) The average ICI leakage versus the

Deleted: ***Figure S4***

Moved down [35]: : The median absolute gene expression

Formatted: Highlight

Moved down [36]: The attacker first predicts the genotypes

Deleted: ***Figure S1:*** Different properties of the eQTLs.

Deleted: ***Figure S2:***

Formatted: Highlight

Moved down [33]: Figure shows the attacker's presumed

Formatted: Highlight

Deleted: ***Figure S3***

Deleted: ***Figure S5:*** Illustration of linking for j^{th} individual.

Moved down [34]: : The distribution of ranks of the individuals

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Formatted: Highlight

Deleted: ***Figure S6***

Formatted: Highlight

Formatted: Font: Not Bold, Not Italic, Highlight

Formatted: Highlight

Deleted: ***Figure S7***

Moved (insertion) [32]

Formatted: Highlight

Moved (insertion) [33]

Formatted: Highlight

expression levels. Rightmost figure shows the posterior distribution of genotypes given that the gene expression level is 10. The posterior distribution has a maximum (MAP prediction) at genotype 2, which is indicated by a star.

Supplementary Figure 3: The distribution of ranks of the individuals in the linking step. At each gradient threshold, the box plots show, for each individual, their ranks in the genotype comparison in the 3rd step of linking attack with MAP genotype prediction. Notice that at around 0.35 correlation threshold, the assigned ranks are minimized, i.e., most of the individual are linked correctly.

Supplementary Figure 4: The median absolute gene expression extremity statistics over 462 individuals in GEUVADIS dataset. (a) For each individual, the extremity is computed over all the genes (23,662 genes) reported in the expression dataset. The median of the absolute value of the extremity is plotted. X-axis shows the sample index and y-axis shows the extremity. The absolute median extremity fluctuates around 0.25, which is exactly the midpoint between minimum and maximum values of absolute extremity. (b) For each individual, we count the number of genes above the extremity threshold. The plot shows the extremity threshold versus the median number of genes (over 462 individuals) above the extremity threshold. Around half of the genes (indicated by dashed yellow lines) have higher than almost 0.3 extremity on average over all the individuals. Also, around median number of 1000 genes over the samples have higher than 0.45 extremity (indicated by dashed red lines).

Supplementary Figure 5: Illustration of linking for jth individual. The attacker first predicts the genotypes ($\tilde{\boldsymbol{v}}_{\cdot,j}$) which are then used to compute the distance to all the individuals in the genotype dataset. The computed distances are then sorted in decreasing. The top matching individual (in the example, individual *a*) is assigned as the linked individual. The first distance gap, $d_{1,2}$, is computed as the difference between the second ($d_{j,(2)}$) and the first ($d_{j,(1)}$) distances in the sorted list.

Supplementary Figure 6:

Supplementary Figure 7: A representative example of extremity based linking. The phenotype dataset (Consisting of gene expression levels for 6 genes) is shown above. Each phenotype measurement is represented by green (negative extreme), red (positive extreme), or grey (non-extreme) dots. Based on the extremity of phenotypes, the attacker performs prediction of genotypes, which are shown below in (2). He or she uses the eQTL dataset (with genes and SNPs) for prediction. Blue and brown triangles correspond to the correct genotype predictions. The grey crosses correspond to the incorrect or unavailable genotype predictions. The attacker compares the predicted genotypes to the genotype dataset in (3), where triangles show the genotypes, and performs linking. The attacker links the predicted genotypes to the genotype dataset. 3 individuals (Bob, Alice, and John) are highlighted. The attacker can link Bob and John by matching them to their genotypes. The correct prediction of rs7274244 (in yellow dashed rectangle) enables the attacker to distinguish between correct entries and reveal both of their disease status as positive. For Alice, the predicted genotypes are equally matching at two entries both of which match at 2 genotypes; PID-b and PID-k (with negative and positive disease status) thus the attacker cannot exactly reveal Alice's disease status.

Moved (insertion) [34]
Formatted: Highlight
Moved (insertion) [35]
Formatted: Highlight
Moved (insertion) [36]
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Deleted: /
Formatted: Highlight

Supplementary ***Figure 8:*** Illustration of risk assessment procedure for joint genotyping/phenotyping data generation. There are two paths of risk assessment to be performed. The first path evaluates the risks associated with release of the QTL datasets. The genotype and phenotype data (on the left) is first used for quantitative trait loci identification (QTL identification box). This generates the significant QTLs. These are then utilized, in addition to the list of external QTL databases, in quantification of leakage versus predictability, as presented in Section 2.2. These results are then relayed to the risk assessment procedures. The second risk assessment procedure evaluates the release of genotype and phenotype datasets. For this, the datasets are input to application of a list of linking attacks (Presented in Sections 2.3, and 2.4, and other linking attacks in the literature) for evaluation of characterization risks. The results are then relayed to risk assessment procedures.

Supplementary ***Figure 9:*** Models of joint genotype-expression distribution with varying numbers of parameters for a positively correlated eQTL. (a) shows the true distribution where grey boxes represent the expression distributions given different genotypes. Red line show the gradient of correlation between genotype and expression. First simplification of the model is shown in (b). The expression distribution can be modeled with Gaussians with different means and variances with total of 6 parameters. The variances can be assumed same for different genotypes (c), where 4 parameters are required. (d) illustrates a representation of the uniform expression distribution given genotypes, where 4 parameters are required. The conditional distribution of expression is uniform (cross shaded rectangles) over the ranges (e_1, e_2), (e_2, e_3), and (e_3, e_4) given genotypes 0, 1, and 2, respectively. The transparent grey rectangles shows the original distributions. (e) is a simplification of (d) where no conditional probability of expression is assigned given genotype is 1. In this model, only one parameter (e_{mid}) is necessary. The conditional probability of expression given genotypes 0 and 2 are uniform for expression levels below e_{mid} and above e_{mid}, respectively (shown with cross shaded rectangles). The original distribution is included with grey rectangles for comparison. Extremity based prediction is an instantiation of the model in (e).

9 ONLINE METHODS

9.1 Genotype, Expression, and eQTL Datasets

The eQTL, expression, and genotype datasets contain the information for linking attack (**Supplementary Fig. 2**). The eQTL dataset is composed of a list of gene-variant pairs such that the gene expression levels and variant genotypes are significantly correlated. We will denote the number of eQTL entries with q. The eQTL (gene) expression levels and eQTL (variant) genotypes are stored in $q \times n_e$ and $q \times n_v$ matrices e and v, respectively, where n_e and n_v denotes the number of individuals in gene expression dataset and individuals in genotype dataset. The kth row of e, $\boldsymbol{e}_k$, contains the gene expression values for kth eQTL entry and $e_{k,j}$ represents the expression of the kth gene for jth individual. Similarly, kth row of v, $\boldsymbol{v}_k$, contains the genotypes for kth eQTL variant and $v_{k,j}$ represents the genotype ($v_{k,j} \in \{0,1,2\}$) of k variant for jth individual. The coding of the genotypes from homozygous or heterozygous genotype categories to the numeric values are done according to the correlation dataset (Online Methods). We assume that the variant genotypes and gene expression levels for the kth eQTL entry are

Deleted: *S8*
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Deleted: *S9*
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Deleted: 1
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Formatted: Highlight
Moved (insertion) [1]
Formatted: Not Highlight
Formatted: Not Highlight
Formatted: Not Highlight
Formatted: Not Highlight
Formatted: Not Highlight
Formatted: Not Highlight
Formatted: Not Highlight
Formatted: Not Highlight
Formatted: Not Highlight
Moved (insertion) [2]
Formatted: Not Highlight
Formatted: Not Highlight
Formatted: Not Highlight
Moved (insertion) [3]
Formatted: Not Highlight

distributed randomly over the samples in accordance with random variables (RVs) which we denote with V_k and E_k, respectively. We denote the correlation between the RVs with $\rho(E_k, V_k)$. In most of the eQTL studies, the value of the correlation is reported in terms of a gradient (or the regression coefficient) in addition to the significance of association (p-value) between genotypes and expression levels.

9.2 Quantification of Characterizing Information and Predictability

The genotype RV V_k takes 3 different values, {0,1,2}, where the genotype coding is done per counting the number of alternate alleles in the genotype. Given that the genotype is $g_{k,j}$, we quantify the individual characterizing information in terms of *self-information*[24] of the event that RV takes the value $g_{k,j}$:

$$ICI(V_k = g_{k,j}) = I(V_k = g_{k,j}) = -\log(p(V_k = g_{k,j})) \quad (1)$$

where V_k is the RV that represents the kth eQTL genotype, $p(V_k = g_{k,j})$ is the probability (frequency) of that V_k takes the value $g_{k,j}$, and ICI denotes the individual characterizing information. Given multiple eQTL genotypes, assuming that they are independent, the total individual characterizing information is simply summation of those:

$$ICI(\{V_1 = v_{1,j}, V_2 = v_{2,j}, \dots, V_N = v_{N,j}\}) = -\sum_{k=1}^{N} \log\left(p(V_k = v_{k,j})\right). \quad (2)$$

The genotype probabilities are estimated by the frequency of genotypes in the genotype dataset. We measure the predictability of eQTL genotypes using an entropy based measure. Given the genotype RV, V_k, and the correlated gene expression RV, E_k,

$$\pi(V_k|E_k = e) = \exp(-H(V_k|E_k = e)) \quad (3)$$

where π denotes the predictability of V_k given the gene expression level e, and H denotes the entropy of V_k given gene expression level e for E_k. The extension to multiple eQTLs is straightforward. For the kth individual, given the expression levels $e_{k,j}$ for all the eQTLs, the total predictability is computed as

$$\pi(\{V_k\}, \{E_k = e_{k,j}\}) = \exp(-H(\{V_k\} \mid \{E_k = e_{k,j}\})) = \exp\left(-\sum_k H(V_k|E_k = e_{k,j})\right) \quad (4)$$

In addition, this measure is guaranteed to be between 0 and 1 such that 0 represents no predictability and 1 representing perfect predictability. The measure can be thought as mapping the prediction process to a uniform random guessing where the average correct prediction probability is measured by π.

Formatted: Not Highlight
Formatted: Not Highlight
Formatted: Not Highlight
Formatted: Not Highlight
Formatted: Not Highlight
Formatted: Not Highlight
Formatted: Not Highlight
Moved (insertion) [8]
Moved (insertion) [9]
Moved (insertion) [10]
Moved (insertion) [11]
Moved (insertion) [12]
Moved (insertion) [13]

9.3 Estimation of Genotype Entropy

We estimate the genotype entropy using the Shannon's entropy[24]:

$$H(V_k) = -\sum_{v\in\{0,1,2\}} p(V_k = v) \times \log(p(V_k = v)) \quad (5)$$

where V_k represents the RV for kth eQTL variant genotypes and $p(V_k = v)$ represents the probability that V_k takes the value v. This probability can be also interpreted as the population frequency of the genotype v at the kth eQTL's variant locus. These probabilities are estimated from the distribution of genotypes over all the samples. As the genotypes are discrete valued, the above formula can be computed in a straightforward way by the summation after the probabilities are estimated.

Moved (insertion) [14]

Moved (insertion) [15]

In the formulation for conditional predictability of genotypes given expression levels, we also use the conditional specific entropies[24] of the genotypes given the gene expression levels. For this, we use the following formulation:

$$H(V_k|E_k = e_{k,j}) = -\sum_{v\in\{0,1,2\}} p(V_k = v|E_k = e_{k,j}) \times \log(p(V_k = v\,|E_k = e_{k,j})) \quad (6)$$

where $p(V_k = v|E_k = e_{k,j})$ represents the conditional probability that V_k takes the value v under the condition that the RV representing gene expression level for k^{th} eQTLs (E_k) is $e_{k,j}$. Since the gene expression levels are continuous, to estimate the conditional probabilities of genotypes given expression levels; we start with the joint distribution of E_k and V_k, then bin the gene expression levels. For this, we use Sturges' rule[25] to choose the number of bins. This rule states that the number of bins should be selected as $n_b = \lceil \log_2(n_e) \rceil + 1 = \lceil \log_2(421) \rceil + 1 = 10$. The binning is done for each gene by first sorting the expression levels for all the individuals, then the range of gene expression levels are divided into $n_b = 10$ bins of equal size and each expression level is mapped to a value between in $[0, n_b - 1]$. The expression level of k^{th} gene in j^{th} individual, $e_{k,j}$, is mapped to

Moved (insertion) [16]

Moved (insertion) [17]

$$\tilde{e}_{k,j} = \left\lceil \frac{(e_{k,j} - \min(\boldsymbol{e_k})) \times n_b}{\max(\boldsymbol{e_k}) - \min(\boldsymbol{e_k})} \right\rceil \quad (7)$$

where $\min(\boldsymbol{e_k})$ and $\max(\boldsymbol{e_k})$ represents the minimum and maximum values, respectively, for the k^{th} expression level over all the samples and $\tilde{e}_{k,j}$ represents the binned expression level. After the gene expression levels are binned, we use the binned expression levels and compute the conditional distribution of the variant genotypes at each binned gene expression level using the histograms:

Moved (insertion) [18]

$$p(V_k = v|\tilde{E}_k = \tilde{e}_{k,j}) = \frac{\sum_i I(\tilde{e}_{k,i} = \tilde{e}_{k,j}, V_{k,i} = v)}{\sum_i I(\tilde{e}_{k,i} = \tilde{e}_{k,j})} \quad (8)$$

where $I(.)$ is an indicator function for counting the number of matching mapped expression and genotype values:

Moved (insertion) [19]

$$I\left(\tilde{e}_{k,i} = \tilde{e}_{k,j}, V_{k,i} = v\right) = \begin{cases} 1; & \text{if } \tilde{e}_{k,i} = \tilde{e}_{k,j}, V_{k,i} = v \\ 0; & \text{otherwise} \end{cases} \quad (9)$$

Finally, we utilize compute the Shannon entropy of the estimated conditional distribution as the condition specific entropies.

Moved (insertion) [20]

9.4 First Distance Gap Statistic Computation

Following the previous section, the attacker computes, for each individual, the distance to all the genotypes in genotype dataset, then identifies the individual with smallest distance. Let $d_{j,(1)}$ and $d_{j,(2)}$ denote the minimum and second minimum genotype distances (among $d^H\left(\tilde{\boldsymbol{v}}_{\cdot,\boldsymbol{j}}, \boldsymbol{v}_{\cdot,\boldsymbol{a}}\right)$ for all $\boldsymbol{a}$) for jth individual. We propose using the difference between these distances, termed *first distance gap statistic*, as a measure of reliability of linking. For this, the attacker computes following difference:

Moved (insertion) [21]

Moved (insertion) [22]

$$d_{1,2}(j) = d_{j,(2)} - d_{j,(1)} \quad (10)$$

First distance gap can be computed without the knowledge of the true genotypes, and is immediately accessible by the attacker with no need for auxiliary information (**Supplementary Fig. 8**). The basic motivation for this statistic comes from the observation that the first distance gap for correctly linked individuals are much higher compared to the incorrectly linked individuals.

9.5 eQTL Identification with Matrix eQTL

For identification of eQTLs, we used Matrix eQTL[26] method. We first generated the testing and training sample lists by randomly picking 210 and 211 individuals, respectively, for testing and training sets. We then separated the genotype and expression matrices into training and testing sets. Matrix eQTL is run to identify the eQTLs using the training dataset. In order to decrease the run time, Matrix eQTL is run in cis-eQTL identification mode. After the eQTLs are generated, we filtered out the eQTLs whose FDR (as reported by Matrix eQTL) was larger than 5%. We finally removed the redundancy by ensuring that each gene and each SNP is used only once in the eQTL final list. To accomplish this, we selected the eQTL that is correlated with highest association with each gene. The association statistic reported by Matrix eQTL was used as the measure of strength of association between expression levels and genotypes. Similar procedure is applied when eQTLs for 30 trios are identified.

Moved (insertion) [23]

9.6 Modeling of Genotype-Phenotype Distribution

In the second step of the linking attack, the genotype predictions are performed. The genotype predictions are used, as an intermediate information, as input to the third step (**Fig. 2c**), where linking is performed. The main aim of attacker is to maximize the linking accuracy (not the genotype prediction accuracy), which depends jointly on the genotype prediction accuracy and the accuracy of the genotype matching in the 3rd step. Other than the accuracy of linking, another important consideration, for risk management purposes, is the amount of auxiliary input data (like training data for prediction model) that the genotype prediction takes. The prediction methods that require high amount of auxiliary data would decrease the applicability of the linking attack as the attacker would need to gather extra information before performing the attack. On the other hand, the prediction methods that require little or no auxiliary data makes the linking attack much more realistic and prevalent. It is therefore useful, in

Moved (insertion) [24]

Moved (insertion) [25]

Moved (insertion) [26]

the risk management strategies, to study complexities of genotype prediction methods and evaluate how these translate into assessing the accuracy and applicability of the linking attack. We study different simplifications of genotype prediction, and illustrate different levels of complexity for genotype prediction.

In MAP based genotype prediction and linking attack, we assume that the attacker estimates the posterior distribution of genotypes and utilizes the maximum *a posteriori* estimate of the genotype as the general prediction method. For this, attacker must first model the joint genotype-phenotype distribution and then build the posterior genotype distribution (**Supplementary Fig. 5a**). The first level level of model can be built by decomposing the conditional distribution of expression with independent variances and means (**Supplementary Fig. 5b**). Assuming that mean and variance are sufficient statistics for the conditional distributions (e.g., normally distributed), the joint distributions can be modeled when the 6 parameters (3 means and 3 variances) are trained. The training can be performed using unsupervised methods like expectation maximization or can be performed using training data. This would, however, increase the required auxiliary data and decrease the applicability of the linking attack. A simplification of the model by assuming the variances of the conditional expression distributions are same for each genotype (**Supplementary Fig. 5c**). This decreases the number of parameters to be trained to 4 (3 means and 1 variance). An equally complex model with 4 parameters can be built assuming the conditional distributions are uniform at non-overlapping ranges of expression for each genotype (**Supplementary Fig. 5d**). This model requires 4 parameters to be trained corresponding to the expression range limits. Another simplification of the genotype prediction can be performed (**Supplementary Fig. 5e**), which requires only one parameter to be trained. In this model, the prediction only assigns uniform probability for homozygous genotypes when expression levels higher or lower than e_{mid} and assigns 0 conditional probability to the heterozygous genotypes, which brings up an important point: This simplified model is exactly the distribution that is utilized in the extremity based genotype prediction. In the extremity based prediction, we estimate e_{mid} simply as the mid-point of the range of gene expression levels within the expression dataset (Supplementary Note).

9.7 Code Availability

All the analysis code that is used to generate results can be obtained from http://privaseq.gersteinlab.org

10 METHODS ONLY REFERENCES

Moved (insertion) [27]

Moved (insertion) [28]

Moved (insertion) [29]

Moved (insertion) [30]

Moved up [31]: <#>REFERENCES¶

Deleted: <#>1. Sboner, A., Mu, X., Greenbaum, D., Auerbach, R. K. & Gerstein, M. B. The real cost of sequencing: higher than you think! *Genome Biol.* **12,** 125 (2011).¶
2. . Rodriguez, L. L., Brooks, L. D., Greenberg, J. H. & Green, E. D. The Complexities of Genomic Identifi ability. *Science (80-.).* **339,** 275–276 (2013).¶
3. Sweeney, L., Abu, A. & Winn, J. Identifying Participants in the Personal Genome Project by Name. *SSRN Electron. J.* 1–4 (2013). doi:10.2139/ssrn.2257732¶
4. . infographic-printable.pdf. at <http://www.nih.gov/precisionmedicine/infographic-printable.pdf>¶
5. Collins, F. S. A New Initiative on Precision Medicine. *N. Engl. J. Med.* **372,** 793–795 (2015).¶
6. . Plan for Increasing Access to Scientific Publications - NIH-Public-Access-Plan.pdf. at <https://grants.nih.gov/grants/NIH-Public-Access-Plan.pdf>¶
7. GENOMIC DATA SHARING (GDS) Home. at <http://gds.nih.gov/index.html>¶
8. . Sweeney, L. *Uniqueness of Simple Demographics in the U.S. Population, LIDAP-WP4. Forthcom. B. entitled, Identifiability Data.* (2000).¶
9. Golle, P. Revisiting the uniqueness of simple demographics in the US population. in *Proc. 5th ACM Work. Priv. Electron. Soc.* 77–80 (2006). doi:http://doi.acm.org/10.1145/1179601.1179615¶
10. . Consortium, T. G. The Genotype-Tissue Expression (GTEx) project. *Nat. Genet.* **45,** 580–5 (2013).¶
11. Bernstein, B. E. *et al.* An integrated encyclopedia of DNA elements in the human genome. *Nature* **489,** 57–74 (2012).¶
12. . The 1000 Genomes Project Consortium. An integrated map of genetic variation. *Nature* **135,** 0–9 (2012).¶
13. Collins, F. S. The Cancer Genome Atlas (TCGA). *Online* 1–17 (2007).¶
14. . Pakstis, A. J. *et al.* SNPs for a universal individual identification panel. *Hum. Genet.* **127,** 315–324 (2010).¶
15. Wei, Y. L., Li, C. X., Jia, J., Hu, L. & Liu, Y. Forensic Identification Using a Multiplex Assay of 47 SNPs. *J. Forensic Sci.* **57,** 1448–1456 (2012).¶
16. . Church, G. *et al.* Public access to genome-wide data: Five views on balancing research with privacy and protection. *PLoS Genet.* **5,** (2009).¶
17. Lunshof, J. E., Chadwick, R., Vorhaus, D. B. & Church, G. M. From genetic privacy to open consent. *Nat. Rev. Genet.* **9,** 406–411 (2008).¶
18. . Homer, N. *et al.* Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays. *PLoS Genet.* **4,** (2008).¶
19. Im, H. K., Gamazon, E. R., Nicolae, D. L. & Cox, N. J. On sharing quantitative trait GWAS results in an era of multiple-omics data and the limits of genomic privacy. *Am. J. Hum. Genet.* **90,** 591–598 (2012).¶
20. . Narayanan, A. & Shmatikov, V. Robust de-anonymization of large sparse datasets. in *Proc. - IEEE Symp. Secur. Priv.* 111–125 (2008). doi:10.1109/SP.2008.33¶
...

Formatted: Font: 11 pt