

Reconciling Utility with Privacy in Genomics

Mathias Humbert Eрман Ayday
Jean-Pierre Hubaux
Laboratory for Communications and Applications
EPFL, Lausanne, Switzerland
firstname.lastname@epfl.ch

Amalio Telenti
Institute of Microbiology
University Hospital of Lausanne
Lausanne, Switzerland
amalio.telenti@chuv.ch

ABSTRACT

Direct-to-consumer genetic testing makes it possible for everyone to learn their genome sequences. In order to contribute to medical research, a growing number of people publish their genomic data on the Web, sometimes under their real identities. However, this is at odds not only with their own privacy but also with the privacy of their relatives. The genomes of relatives being highly correlated, some family members might be opposed to revealing any of the family's genomic data. In this paper, we study the trade-off between utility and privacy in genomics. We focus on the most relevant kind of variants, namely single nucleotide polymorphisms (SNPs). We take into account the fact that the SNPs of an individual contain information about the SNPs of his family members and that SNPs are correlated with each other. Furthermore, we assume that SNPs can have different utilities in medical research and different levels of sensitivity for individuals. We propose an obfuscation mechanism that enables the genomic data to be publicly available for research, while protecting the genomic privacy of the individuals in a family. Our genomic-privacy preserving mechanism relies upon combinatorial optimization and graphical models to optimize utility and meet privacy requirements. We also present an extension of the optimization algorithm to cope with the non-linear constraints induced by the correlations between SNPs. Our results on real data show that our proposed technique maximizes the utility for genomic research and satisfies family members' privacy constraints.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General—Security and protection (e.g., firewalls); J.3 [Life and Medical Sciences]: Biology and genetics; K.4.1 [Computer and Society]: Public Policy Issues—Privacy

Keywords

Genomic Privacy; Obfuscation; Optimization

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
WPES'14, November 3, 2014, Scottsdale, Arizona, USA.
Copyright 2014 ACM 978-1-4503-2948-7/14/11 ...\$15.00.
<http://dx.doi.org/10.1145/2665943.2665945>.

1. INTRODUCTION

Genomic research has revolutionized our understanding of medicine: the “one size fits all” approach has already been replaced by personalized medicine for the treatment of many diseases, for which genetic factors of the individuals play an important role. This also paves the way to the early diagnoses of many serious diseases. The association of genetic factors with diseases and treatments is only possible via large-scale genome-wide association studies (GWAS) that require the availability of a considerably high number of human genomes. Computers are at the core of this endeavor, because (i) high computing power is required to interpret genomic data [11], (ii) hand-held devices are used to visualize this data, and (iii) increasingly more people upload genomic data (and more generally health-related data) on public websites (e.g. OpenSNP.org and personalgenomes.org). Computing systems facilitate data access and processing for legitimate usage, but sometimes also for purposes that were initially unintended, thus raising privacy concerns.

Genomic data carries much sensitive information about its owner. By analyzing the DNA of an individual, it is now possible to learn about his disease predispositions (e.g., for Alzheimer's or Parkinson's), ancestries, and physical attributes [17]. The threat to genomic privacy is magnified by the fact that a person's genome is correlated to his family members' genomes, thus leading to interdependent privacy risks. Kin genomic privacy was popularized by the story of Henrietta Lacks whose cells were sequenced and whose DNA sequence was put online without the consent of her descendants [1]. After complaints from the family, essentially due to privacy concerns, Henrietta's genome was taken offline, and in 2013, the National Institutes of Health (NIH) came to an agreement with the Lacks family, which gave them some control over her genome [2]. Even though this agreement enables the genomic researchers to use Henrietta's genome again, it also draws attention to the lack of techniques for balancing the benefits of genomic research with personal and kin genomic privacy risks. Richard Sharp, the director of biomedical ethics at the Mayo Clinic, warned that the agreement was only a “one-off solution” rather than a broad policy that addresses the tension between research and relatives' privacy, and he added that a “new policy” was absolutely needed [2].

Anonymization was the first countermeasure proposed to protect genomic privacy, but in many different studies it was proven inadequate [20, 21, 33]. Another protection mechanism is to add noise to aggregate statistical results (to satisfy differential privacy) [18, 25], but at the cost of reduced

accuracy. The last option proposed in the literature is to rely on cryptographic techniques [4, 6]. Even though these techniques are proven to be effective for using genomic data in healthcare [4, 12], computational complexity becomes very high when it comes to conducting statistical tests on large numbers of encrypted genomes for genomic research [26].

In this work, we present a genomic-privacy preserving mechanism (GPPM) for reconciling people’s willingness to share their genomes (e.g., to help research¹) with privacy. Our GPPM acts at the individual data level, not at the aggregate data (or statistical) level like in [18, 25]. Focusing on the most relevant type of variants (the SNPs), we study the trade-off between the usefulness of disclosed SNPs (utility) and genomic privacy. We consider an individual who wants to share his genome, yet who is concerned about the subsequent privacy risks for himself and his family. Thus, we design a system that maximizes the disclosure utility but does not exceed a certain level of privacy loss within a family, considering (i) kin genomic privacy, (ii) personal privacy preferences (of the family members), (iii) privacy sensitivities of the SNPs, (iv) correlations between SNPs, and (v) the research utility of the SNPs. Our GPPM can automatically evaluate the privacy risks of all the family members and decide which SNPs to disclose. To achieve this goal, it relies on probabilistic graphical models and combinatorial optimization. Our results indicate that, given the current data model, genomic privacy of an entire family can be protected while an appropriate subset of genomic data can be made available. Our contributions can be summarized as follows:

- We propose a GPPM for enabling genomic research while protecting personal and kin genomic privacy.
- Given the genomic data model, our obfuscation mechanism maximizes the utility and meets all the privacy constraints of a given family.
- Using combinatorial optimization, we first compute the optimal solution without considering correlations between SNPs, and then we extend the algorithm to address non-linear constraints induced by these correlations.

The paper is organized as follows. In Section 2, we present the adversary model and our privacy metrics. In Section 3, we present our GPPM in detail. Next, in Section 4, we evaluate its performance. In Section 5, we summarize the related work, before concluding in Section 6.

2. PRELIMINARIES

In this section, we briefly introduce the adversary model and our quantification framework for measuring genomic privacy. We encourage the reader to look at Section 2.1 of [23] for an introduction to the relevant genomic concepts.

2.1 Adversary Model

The objective of the adversary is to infer the values of hidden SNPs of one or more members of a given family.² To

¹<http://opensnp.wordpress.com/2011/11/17/first-results-of-the-survey-on-sharing-genetic-information/>

²We assume here that the attacker has de-anonymized the genome(s) owner(s) and knows (part of) the familial relationships. For more details, take a look at the Appendix A and Subsection 2.2 of [23].

do so, he relies on some background knowledge, essentially the minor allele frequencies (MAFs) of the SNPs, the linkage disequilibrium (LD) values between the SNPs, and the basic rules of Mendelian inheritance. Note that LD values are only expressed between pairs of SNPs by geneticists, thus only pairwise correlations are available to the adversary. In addition to this background information, the adversary observes a subset of the SNPs of the family members, typically those who have disclosed their genomic data.

The attacker’s ultimate goal is then formally defined as computing posterior marginal probabilities of unknown variables from the global posterior probability distribution $P(\mathbf{X}_H|\mathbf{X}_O, \mathcal{K})$, where \mathbf{X}_H represents the set of hidden SNPs, \mathbf{X}_O the set of observed SNPs, and \mathcal{K} the background knowledge or data model. This inference attack can be efficiently carried out by relying on belief propagation (see Subsection 3.2 of [22] for more details about the algorithm).

2.2 Privacy Metrics

We define \mathbf{F} to be the set of family members in the *targeted family* and \mathbf{S} to be the set of SNP IDs (i.e., positions on the DNA sequence), where $|\mathbf{F}| = n$ and $|\mathbf{S}| = m$. We also let x_j^i be the value of SNP j ($j \in \mathbf{S}$) for individual i ($i \in \mathbf{F}$), where $x_j^i \in \{0, 1, 2\}$. Furthermore, we let $\mathbf{X} = \mathbf{X}_O \cup \mathbf{X}_H$ be the set of SNPs of all family members, hence $|\mathbf{X}| = n \cdot m$.

The inferred marginal probabilities can be expressed as $P(\hat{x}_j^i|\mathbf{X}_O, \mathcal{K})$, for all $i \in \mathbf{F}, j \in \mathbf{S}$. We quantify privacy by measuring the adversary’s *incorrectness*, i.e. the adversary’s error in inferring the targeted SNPs. This quantifies the expected distance between the adversary’s estimate on the value of a SNP, \hat{x}_j^i , and the true value of the corresponding SNP, x_j^i . This expected estimation error can be expressed as follows:

$$E_j^i = \sum_{\hat{x}_j^i \in \{0,1,2\}} P(\hat{x}_j^i|\mathbf{X}_O, \mathcal{K}) ||x_j^i - \hat{x}_j^i||. \quad (1)$$

The above metric is useful for quantifying the genomic privacy of individuals. To quantify individuals’ health privacy, we can focus on the predisposition to different diseases. Let \mathbf{S}_d be the set of IDs of the SNPs that are associated with a disease d . Then, health privacy for an individual i regarding the disease d can be quantified as follows:

$$D_d^i = \frac{1}{\sum_{k \in \mathbf{S}_d} c_k} \sum_{k \in \mathbf{S}_d} c_k E_k^i, \quad (2)$$

where c_k is the contribution of SNP k to disease d .

3. PROPOSED SOLUTION

In order to mitigate attribute-inference attacks and protect genomic and health privacy, the GPPM relies upon an *obfuscation mechanism*. In practice, obfuscation can be implemented by adding noise to the SNP values, by injecting fake SNP values, by reducing precision, or by simply hiding the SNP values. In this paper, we choose SNP hiding, essentially because the genomic research community would not receive other options positively. Indeed, genetic researchers are very reluctant about adding noise or fake data, notably because of the huge investment they make to increase (sequencing) accuracy. We assume one family member, at a given time, who wants to disclose his SNPs and to guarantee a minimum privacy level for him and his family. Fig. 1 provides an overview of our GPPM.

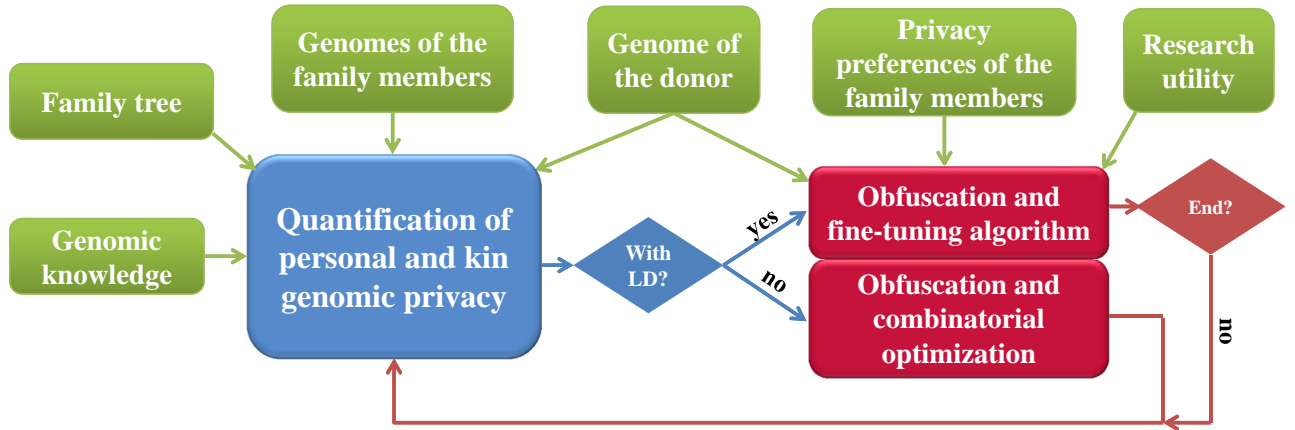


Figure 1: General protection framework. The GPPM takes as inputs (i) the privacy levels of all family members, (ii) the genome of the donor, (iii) the privacy preferences of the family members, and (iv) the research utility. First, LD is not considered in order to use combinatorial optimization (see Subsection 3.2). Note that we go only once through this box. Then, LD is used and a fine-tuning algorithm is used to cope with non-linear constraints. See Subsection 3.3 for details on the end criterion. The algorithm outputs the set of SNPs that the donor can disclose.

3.1 Settings

For clarity of presentation, we focus on one family whose members are defined by the set \mathbf{F} ($|\mathbf{F}| = n$). We assume that there is only one donor D who makes the decision to share his genome at a given time. His relatives might have already publicly shared some of their genomic data on the Internet. D takes this into account when he makes his own disclosure decision. We let \mathbf{S} ($|\mathbf{S}| = m$) be the set of SNP IDs. Its cardinality m can go up to 50 million, as this is currently the approximate number of SNPs in the human population [3]. In practice, however, people put online (e.g., on OpenSNP) up to 1 million of the most significant SNPs. We let $\mathbf{X}^D = \{x_j^D : j \in \mathbf{S}\}$ represent the set of SNPs of D (x_j^D is the value of SNP j of the donor D), that are all initially undisclosed, i.e. $\mathbf{X}^D \subseteq \mathbf{X}_H$ (where \mathbf{X}_H denotes the set of SNPs from \mathbf{X} whose values are hidden, as discussed before). Finally, we let $\mathbf{y}^D = \{y_j^D : j \in \mathbf{S}\}$ represent the decision vector of D , where $y_j^D = 1$ means the corresponding SNP will be disclosed, and $y_j^D = 0$ means x_j^D will remain hidden. Note that the decision to disclose a SNP j could be probabilistic, thus transforming y_j^D into a continuous variable in $[0, 1]$. We leave the study of the continuous case for future work.

We express the privacy constraints of a family member both in terms of genomic and health privacy. Our framework can account for different privacy preferences for different family members, SNPs, and diseases. For all $i \in \mathbf{F}$, $j \in \mathbf{S}$, we define the privacy sensitivity as s_j^i . We can set the s_j^i 's to be equal by default. Then, an individual willing to personalize his privacy preferences may further define his own privacy sensitivities regarding specific SNPs based on his privacy concerns regarding, e.g., certain phenotypes. The most well-known example of such a scenario is the case of James Watson, co-discoverer of DNA, who made his whole DNA sequence publicly available, with the exception of one gene known as Apolipoprotein E (ApoE), one of the strongest

predictors for the development of Alzheimer's disease.³ We let the sets \mathbf{P}_s^i and \mathbf{P}_d^i include the privacy-sensitive SNP IDs and privacy-sensitive diseases of individual i , respectively. We represent the tolerance to the genomic-privacy loss of individual i as $\text{Pri}(i, \mathbf{P}_s^i)$, and the tolerance to the health-privacy loss of individual i regarding disease $d \in \mathbf{P}_d^i$ as $\text{Pri}(i, d)$. These tolerance values represent the maximum privacy loss (after the disclosure of D 's SNPs) that an individual would bear. By considering the privacy losses instead of the absolute privacy levels, we ensure that the donor will more likely reveal a SNP whose value is already well inferred by the attacker before donor's disclosure (e.g., by using SNPs previously shared by the donor's relatives). Note that these tolerance values can always be updated for any new family member willing to disclose his genome. Finally, the utility function is a non-decreasing function of the norm of \mathbf{y}^D , as the knowledge of more SNPs can only help genomic research. In a first step towards enhanced genomic privacy, we assume linear contribution of SNPs on utility.⁴ Formally, we define u_j to be the utility provided by SNP j . Note that, in practice, the utility of the SNPs can be determined by the research authorities and can vary based on the study.

3.2 Linear Optimization

3.2.1 Optimization Problem

The donor faces an optimization problem: How to maximize research utility while protecting his own and his relatives' genomic and health privacy. First, the objective function is formally defined as $\sum_{j \in \mathbf{S}} u_j y_j^D$. Then, privacy constraints are defined, for each individual, as the sum of privacy losses induced by the donor's disclosure over all SNPs. This sum must be capped by the respective privacy loss tolerances of all family members. Formally, for all individuals $i \in \mathbf{F}$ and SNPs $j \in \mathbf{S}$, the privacy loss induced by the dis-

³Later researchers have used correlations in the genome to unveil Watson's predisposition to Alzheimer's [29]. In this work, we also consider such correlations.

⁴We intend to study non-linear utility in future work.

closure of x_j^D is defined as $(E_j^i(y_j^D = 0) - E_j^i(y_j^D = 1))$. Note here that the privacy loss at a given SNP j for any relative is only affected by the donor's decision y_j^D regarding SNP j but no other SNP $k \neq j$, meaning that LD correlations are not taken into account. We make this assumption here in order to define linear constraints. We show how to extend the linear optimization problem to include LD correlations in Subsection 3.3. Finally, note that if an individual i has already revealed his SNP j , i.e. $x_j^i \in \mathbf{X}_O$, the privacy loss at this SNP for i is zero, because $E_j^i(y_j^D = 0) = E_j^i(y_j^D = 1) = 0$.

For all $i \in \mathbf{F}$, $j \in \mathbf{S}$, the privacy weight p_j^i is defined as

$$p_j^i = s_j^i \times (E_j^i(y_j^D = 0) - E_j^i(y_j^D = 1)). \quad (3)$$

Clearly, p_j^i at a given SNP j can be different for each family member, depending on how close he is from the donor in the family tree, on the actual values x_j^i and x_j^D of his and the donor's SNPs, and on his sensitivity. Note that $s_j^i = 0 \forall j \notin \mathbf{P}_s^i$.

We can now define the linear optimization problem as

$$\begin{aligned} & \underset{\mathbf{y}^D}{\text{maximize}} && \sum_{j \in \mathbf{S}} u_j y_j^D \\ & \text{subject to} && \sum_{j \in \mathbf{P}_s^i} p_j^i y_j^D \leq \text{Pri}(i, \mathbf{P}_s^i), \forall i \in \mathbf{F} \\ & && \sum_{k \in \mathbf{S}_d} p_k^i y_k^D \leq \text{Pri}(i, d), \forall d \in \mathbf{P}_d^i, \forall i \in \mathbf{F} \\ & && y_j^D \in \{0, 1\}, \forall j \in \mathbf{S}, \end{aligned} \quad (4)$$

where \mathbf{S}_d is the set of SNPs that are associated with disease d . Note that, for the last inequality, we replace the sensitivity s_k^i in p_k^i by the contribution c_k of SNP k to disease d described in (2), and we embed the normalization factor $\sum_k c_k$ of (2) in $\text{Pri}(i, d)$.

3.2.2 Optimization Algorithm

Our optimization problem is very similar to the multidimensional knapsack problem [19]. We decide to follow the branch-and-bound method proposed by Shih [31], because it finds the optimal solution, represents a good trade-off between time and storage space, and allows for the extension of the algorithm to null and negative (privacy) weights. A branch-and-bound algorithm is a systematic enumeration of all candidate solutions, where large subsets of candidate solutions are pruned by using upper bounds on the quantity being optimized. A branch-and-bound method generally relies on two main rules: (i) the estimation of the upper bound at any node (state of assigned variables) in the search tree, and (ii) a choice criterion for the selection of a branching variable at the node selected for further partitioning.

In order to find (i), Shih suggests treating the C -constraint knapsack problem as C single-constraint knapsack problems with the same objective function, and then computing the value associated to the optimal fractional solution (thus relaxing $y_j^D \in \{0, 1\}$ into $y_j^D \in [0, 1]$) of all of these C problems separately. The fractional optimal solution is easier to solve than the integer solution, as it enables us to sort the items (SNPs), with respect to their ratios between utility and privacy weights $r_j^i = u_j/p_j^i$, from the highest to the lowest ratios, and then to select all the highest ones that can fit under the constraint, with the last SNP being partially included (based on the remaining room). Note that, in our setting,

we can have different orderings of SNPs for different constraints, based on the p_j^i values of the family members. The computation of the fractional optimal solution is repeated C times, for the C different optimization problems, leading to C optimal values. Then, the upper bound at the given node is defined as the minimum among all these C values.

The node selected for the next branching is defined as the one in the search frontier whose upper bound is the highest among all nodes in the frontier, and where the solution associated with this upper bound is infeasible (some variables being different than 0 and 1, or some constraints being not satisfied). The branching variable is the one whose ratio is the smallest among all the non-zero free variables (variables not explicitly assigned to 0 or 1 at a node) in this infeasible solution. If the solution at this node is feasible (all decision variables assigned to 0 or 1 and all constraints satisfied), then it is optimal, and the algorithm stops.

Let us mention that our optimization problem has two main differences with the multidimensional knapsack problem. First, the privacy metrics, hence weights, are expressed in real values, between 0 and 2 for E_j^i , whereas the knapsack problem assumes integer numbers only. In order to obtain integer values, we merely multiply all our privacy weights p_j^i 's and tolerance values $\text{Pri}(\cdot)$ by 10^k , where $k \in \mathbb{N}^+$ depends on the precision we want to attain, and then round the weights to the closest greater integer and the tolerance values to the closest smaller integer. This ensures that all privacy constraints in the space of real numbers are still satisfied. Second, the privacy weight p_j^i can be equal to zero (e.g., if $x_j^i \in \mathbf{X}_O$) or even negative (when the donor reveals a SNP whose value increases the privacy of his relative(s) at the same SNP).⁵ Thus, the ratios r_j^i might not be defined or be negative. In order to resolve this issue, we give a higher ranking in the ordering of SNPs to ratios with null weights with respect to those with positive weights, and we give an even higher ranking to those with negative weights. We furthermore give higher ranking to negative weights with absolute values higher than the others. To enforce this ranking in practice in Section 4, we set $r_j^i = u_j/0.1$ for null p_j^i 's, and $r_j^i = u_j|p_j^i|/0.01$ for negative p_j^i 's. Note that, due to the requirement of integer values for weights, all other (positive) weights p_j^i belong, after the aforementioned multiplication by 10^k and rounding, to \mathbb{N}^+ .

The output of the above optimization algorithm is an optimal solution \mathbf{y}^{*D} that represents the set of SNPs the donor could disclose and an optimal value u^* representing the maximum research utility. We represent the set of the optimal candidate SNPs to be shared as $\tilde{\mathbf{X}}^D \subseteq \mathbf{X}^D$. This is the output we see in state 2 of Fig. 2. We give $\tilde{\mathbf{X}}^D$ as input to the non-linear algorithm described in Subsection 3.3 to eventually reach state 3.

3.3 Non-Linear Extension

3.3.1 Non-Linear Optimization Problem

The LD correlations between the SNPs are not considered in the above optimization problem in order for the con-

⁵For example, assume a child to be homozygous-major at a given SNP and his father to be heterozygous. Then, the estimation error for the child's SNP, thus the child's privacy at this SNP, increases when the father's SNP is observed by the attacker (compared to the case when it is unknown, when only the MAF is used, and this MAF is close to 0).

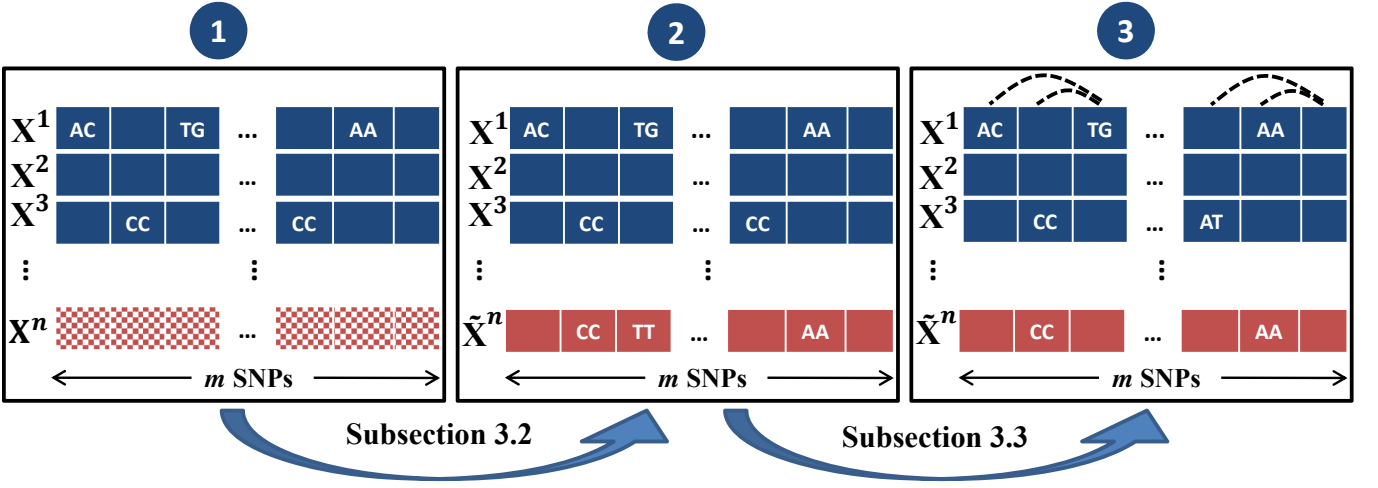


Figure 2: Main steps of the optimization algorithm. Without loss of generality, the donor D is assumed to be the n -th member of the family, thus $\mathbf{X}^D = \mathbf{X}^n$. First, the donor selects a subset $\tilde{\mathbf{X}}^D$ of candidate SNPs to be shared using the optimization algorithm of Subsection 3.2.2, and then reveals less or more SNPs depending on the updated privacy weights computed with LD by relying upon the fine-tuning step of Subsection 3.3.

straints to remain linear. In this subsection, we propose an extension of the branch-and-bound algorithm in order to deal with non-linear constraints.

Whereas in the case without LD, the privacy loss at a given SNP j of individual i depended only on the donor's decision y_j^D regarding SNP j , we have now to consider all the SNPs in LD with j to evaluate the privacy loss at j . Defining \tilde{E}_j^i to be the privacy level of individual i at SNP j quantified by including LD correlations, the privacy loss at SNP j of individual i induced by the disclosure of $\tilde{\mathbf{X}}^D$ is equal to $(\tilde{E}_j^i(\mathbf{y}^D = \mathbf{0}) - \tilde{E}_j^i(\mathbf{y}^{*D}))$. This leads to the following updated privacy weights

$$\tilde{p}_j^i = s_j^i \times (\tilde{E}_j^i(\mathbf{y}^D = \mathbf{0}) - \tilde{E}_j^i(\mathbf{y}^{*D})). \quad (5)$$

Note that now the argument of \tilde{E}_j^i is the entire vector \mathbf{y}^D and not only y_j^D , because of LD. The optimization problem in (4) is reformulated as a non-linear optimization problem:

$$\begin{aligned} & \underset{\mathbf{y}^D}{\text{maximize}} && \sum_{j \in \mathbf{S}} u_j y_j^D \\ & \text{subject to} && \sum_{j \in \mathbf{P}_s^i} \tilde{p}_j^i(\mathbf{y}^D) \leq \text{Pri}(i, \mathbf{P}_s^i), \forall i \in \mathbf{F} \\ & && \sum_{k \in \mathbf{S}_d} \tilde{p}_k^i(\mathbf{y}^D) \leq \text{Pri}(i, d), \forall d \in \mathbf{P}_d^i, \forall i \in \mathbf{F} \\ & && y_j^D \in \{0, 1\}, \forall j \in \mathbf{S}. \end{aligned} \quad (6)$$

Instead of solving this very complex optimization problem, we rely on the optimal solution \mathbf{y}^{*D} computed in Subsection 3.2, embed it into (6), and check whether the privacy constraints are still met with the updated privacy weights \tilde{p}_j^i 's. Let us first study the case when no SNP has been disclosed by any relative before the donor's decision.⁶ If

⁶Without loss of generality, we focus here on the genomic-privacy constraints.

$\mathbf{X}_O = \emptyset$, then

$$\sum_{j \in \mathbf{P}_s^i} \tilde{E}_j^i(\mathbf{y}^D = \mathbf{0}) = \sum_{j \in \mathbf{P}_s^i} E_j^i(\mathbf{y}^D = \mathbf{0}) \quad (7)$$

and, because of LD correlations,

$$\sum_{j \in \mathbf{P}_s^i} \tilde{E}_j^i(\mathbf{y}^{*D}) \leq \sum_{j \in \mathbf{P}_s^i} E_j^i(\mathbf{y}^{*D}). \quad (8)$$

Embedding (7) and (8) in (3) and (5), we get

$$\sum_{j \in \mathbf{P}_s^i} \tilde{p}_j^i(\mathbf{y}^{*D}) \geq \sum_{j \in \mathbf{P}_s^i} p_j^i y_j^D, \quad (9)$$

meaning that, for the same value of $\text{Pri}(i, \mathbf{P}_s^i)$ in (4) and (6), the privacy constraint of family member i in (6) will be violated with high likelihood once LD is taken into account. If $\mathbf{X}_O \neq \emptyset$, then two scenarios can happen. If

$$\underbrace{\sum_{j \in \mathbf{P}_s^i} E_j^i(\mathbf{y}^{*D}) - \tilde{E}_j^i(\mathbf{y}^{*D})}_{\text{Privacy difference using LD or not when } \tilde{\mathbf{X}}^D \subset \mathbf{X}_O} \geq \underbrace{\sum_{j \in \mathbf{P}_s^i} E_j^i(\mathbf{y}^D = \mathbf{0}) - \tilde{E}_j^i(\mathbf{y}^D = \mathbf{0})}_{\text{Privacy difference using LD or not when } \tilde{\mathbf{X}}^D \subset \mathbf{X}_H},$$

then we get the same inequality (9), leading to the same consequences of constraint violation. If, on the contrary,

$$\sum_{j \in \mathbf{P}_s^i} E_j^i(\mathbf{y}^{*D}) - \tilde{E}_j^i(\mathbf{y}^{*D}) < \sum_{j \in \mathbf{P}_s^i} E_j^i(\mathbf{y}^D = \mathbf{0}) - \tilde{E}_j^i(\mathbf{y}^D = \mathbf{0}), \quad (10)$$

then we get

$$\sum_{j \in \mathbf{P}_s^i} \tilde{p}_j^i(\mathbf{y}^{*D}) < \sum_{j \in \mathbf{P}_s^i} p_j^i y_j^D, \quad (11)$$

which might allow the donor to reveal more of his SNPs without violating any of his relatives' privacy constraints. At a first glance, Inequality (11) looks counterintuitive. However, in order to understand it, let us look at Inequality (10), which states that the difference in privacy levels if LD is used or not is smaller after the observation of the

donor's SNPs $\tilde{\mathbf{X}}^D$. This means that, by revealing his own SNPs, the donor reduces the importance of using LD correlations to correctly infer some of the SNPs of his relatives. For instance, let us assume the donor to be the father of a child i whose mother has already revealed SNP j , in LD with another SNP k revealed by the child. Furthermore, assume that the father, mother, and child are homozygous major at SNPs j and k . Now, before the father reveals his SNP j , there is some uncertainty about the child's SNP j ; but by observing SNP k of the child, the attacker improves his estimation if he uses LD correlations and thus reduces his estimation error, meaning $\tilde{E}_j^i(\mathbf{y}^D = \mathbf{0}) < E_j^i(\mathbf{y}^D = \mathbf{0})$. However, once the father decides to reveal his homozygous major SNP j ($y_j^{*D} = 1$), the attacker is certain that the child's SNP j is homozygous major, regardless if LD is used or not, i.e. $E_j^i(\mathbf{y}^{*D}) = \tilde{E}_j^i(\mathbf{y}^{*D}) = 0$. Thus, we have $E_j^i(\mathbf{y}^{*D}) - \tilde{E}_j^i(\mathbf{y}^{*D}) < E_j^i(\mathbf{y}^D = \mathbf{0}) - \tilde{E}_j^i(\mathbf{y}^D = \mathbf{0})$, leading by extension to Inequality (10).

3.3.2 Fine-Tuning Algorithm

Let us first describe how we proceed if one or multiple constraints are violated once LD correlations are considered in the privacy quantification. In this case, we first select the privacy constraint that is not met anymore with the highest difference between $\text{Pri}(i, \mathbf{P}_s^i)$ (or $\text{Pri}(i, d)$) and the newly computed privacy losses. Focusing on the set of genomic-privacy constraints, we thus select the constraint of the family member k , where

$$k = \arg \max_{i \in \mathbf{F}} \left\{ \sum_{j \in \mathbf{P}_s^i} \tilde{p}_j^i(\mathbf{y}^{*D}) - \text{Pri}(i, \mathbf{P}_s^i) \right\}. \quad (12)$$

We want then to hide some SNPs j in $\tilde{\mathbf{X}}^D$ (i.e. where $y_j^{*D} = 1$) in order that the constraint of relative k is satisfied again. For all the SNPs in $\tilde{\mathbf{X}}^D$, we compute a global privacy weight δ_j^k for SNP j of k that includes the privacy loss induced by SNP j on the SNPs $l \in \mathbf{L}$ in LD with j . We compute this global privacy weight at SNP j for individual k as

$$\begin{aligned} \delta_j^k &= \tilde{p}_j^k + \sum_{l \in \mathbf{L}} \tilde{p}_l^k \\ &= s_j^k (\tilde{E}_j^k(\mathbf{y}^D = \mathbf{0}) - \tilde{E}_j^k(\mathbf{y}^{*D})) + \sum_{l \in \mathbf{L}} s_l^k (\tilde{E}_l^k(\mathbf{y}^D = \mathbf{0}) - \tilde{E}_l^k(\mathbf{y}^{*D})). \end{aligned} \quad (13)$$

Then, we compute the ratios of each SNP j (in $\tilde{\mathbf{X}}^D$) for individual k as $\bar{r}_j^k = \delta_j^k / u_j$. The SNPs with the highest ratios represent those where LD correlations cause the highest decrease in the genomic privacy of family member k and/or provide low utility to the optimal solution \mathbf{y}^{*D} computed in Subsection 3.2. Thus, these should be removed first from the set $\tilde{\mathbf{X}}^D$ in order to meet the privacy constraint of individual k again, and to cause the smallest decrease in utility.

To see whether the privacy constraint is met for the family member k , we iteratively remove such SNPs (starting from the one with the highest ratio) from the set $\tilde{\mathbf{X}}^D$ and, after each removal, we input the new solution to the quantification box. We repeat this until all the privacy constraints are satisfied for all family members in \mathbf{F} . Finally, the SNPs left in set $\tilde{\mathbf{X}}^D$ after the final iteration are publicly shared. This case is illustrated in state 3 of Fig. 2.

In the case where including LD correlations in the privacy quantification actually decreases privacy losses, the privacy constraints are still met and can even enable for potential

new SNPs to be included in $\tilde{\mathbf{X}}^D$. In this case, we select the genomic-privacy constraint where the remaining room between the genomic-privacy constraint and the newly computed privacy loss is the smallest, i.e. we select the constraint of the family member k , where

$$k = \arg \min_{i \in \mathbf{F}} \left\{ \text{Pri}(i, \mathbf{P}_s^i) - \sum_{j \in \mathbf{P}_s^i} \tilde{p}_j^i(\mathbf{y}^{*D}) \right\}. \quad (14)$$

For all SNPs *not* in $\tilde{\mathbf{X}}^D$ (i.e., where $y_j^{*D} = 0$), we compute the privacy decrease led by LD for k compared to the privacy level computed without LD. We compute this privacy difference at a SNP j for individual k as

$$\Delta_j^k = E_j^k(y_j^{*D} = 0) - \tilde{E}_j^k(\mathbf{y}^{*D}), \quad (15)$$

where $E_j^k(y_j^{*D} = 0)$ is the privacy value at SNP j for individual k after the linear optimization (without considering LD), and $\tilde{E}_j^k(\mathbf{y}^{*D})$ is the privacy quantified using LD. Then, we compute the ratios of each SNP j (in $\tilde{\mathbf{X}}^D$) for individual k as $\bar{r}_j^k = (u_j \Delta_j^k) / s_j^k$. The SNPs with highest ratios represent those where LD correlations cause the most significant decrease in the genomic privacy of family member k , and/or provide high utility. Thus, these SNPs are the first ones that should be included in $\tilde{\mathbf{X}}^D$, in order to have the smallest difference in privacy loss, thus still meeting k 's privacy constraint and providing maximal utility increase.

We iteratively add SNPs in $\tilde{\mathbf{X}}^D$ and input the new solution to the quantification box to check whether all the privacy constraints are still met for all family members. We repeat this step until one privacy constraint is violated again, and we publicly share the last set $\tilde{\mathbf{X}}^D$ to have satisfied all constraints. In the next section, we briefly show experimentally how close this fine-tuning algorithm is to the maximum found with exhaustive search. The thorough analytical evaluation of the discrepancy between the optimal solution and our approximation is left for future work.

4. EVALUATION

In this section, we evaluate the effectiveness of our optimization algorithm for protecting individual and kin privacy. We study the balance between maximum achievable utility and the privacy of each individual in a family. The results show the total utility we can obtain for different genomic-privacy guarantees.

We make use of the CEPH/Utah Pedigree 1463 [15]. It includes the partial DNA sequences of 17 family members: 4 grandparents, 2 parents, and 11 children. In order to remain at a representative scale, we keep only 5 randomly chosen children out of 11. Fig. 3(a) presents the pedigree structure that we use in our study. We focus on 50 SNPs of chromosome 1 and assume one genomic-privacy constraint, including all the 50 SNPs for each family member. Thus, we have a total of 11 privacy constraints, which represents more constraints than other generic experiments in the optimization literature that included up to 5 or 7 constraints [19]. Considering LD strengths between $r^2 = 0.5$ (medium LD) and $r^2 = 1$ (strongest LD), each SNP is in LD with around 4.5 other SNPs, on average. We set a precision of 0.01 in our privacy weights and tolerance values, thus multiplying these real-valued elements by 10^2 , and rounding them, as explained in Subsection 3.2. Parent P5 is assumed to be the donor in all scenarios presented in this section. In our evaluations, for the sake of simplicity, we assume each SNP

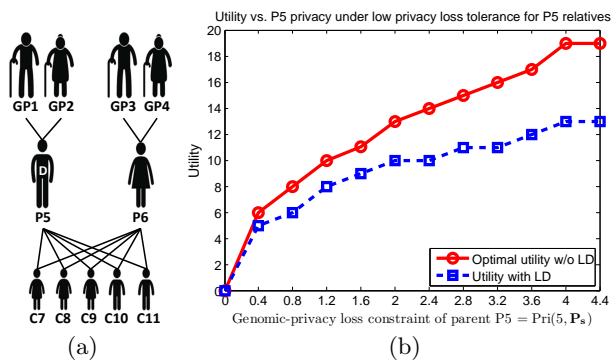


Figure 3: Evaluation of the proposed solution on a real Utah pedigree. (a) Genealogical tree, (b) Utility versus privacy under low tolerance to privacy loss for all relatives except parent P5, and varying values of privacy constraints $\text{Pri}(5, \mathbf{P}_s^5)$ for parent P5 (x-axis). Here, $\mathbf{X}_O = \emptyset$, meaning that no relative has revealed any SNP before P5. Low tolerance is defined as 1/4 of the total privacy loss that a relative would incur if all 50 SNPs of P5 were revealed. Results are shown up to $\text{Pri}(5, \mathbf{P}_s^5) = 4.4$ even if P5’s privacy constraint can go beyond because, from $\text{Pri}(5, \mathbf{P}_s^5) = 4$, the utility stops increasing (capped by other relatives’ low tolerance).

is equally useful for the genomic research, i.e., $u_j = 1$ for all SNPs. We also assume the privacy sensitivities are equal, for all SNPs and individuals, i.e., $s_j^i = s$. Equal values of sensitivities for all SNPs would typically be the default setting if, for example, family members do not want to bother setting their privacy sensitivities themselves. Other distributions over the utility or sensitivity values should not alter the algorithm’s performances significantly.

4.1 No Previous Disclosure by the Family

As of today most people have not publicly revealed their genome, we first analyze the case where no family member has shared any of his SNPs before the donor makes his decision. In other words, we assume that, initially, $\mathbf{X}_O = \emptyset$. We analyze the tension between utility and privacy for different values of parent P5’s privacy constraint. Fig. 3(b) shows the increase in the utility caused by the higher privacy loss tolerance of P5. Because a low tolerance to privacy loss is assumed for all the other relatives in the family in this case, the utility (computed without LD) cannot go beyond 19, even if P5’s constraint increases beyond 4. We also notice that, once the LD is included in the privacy quantification, the utility decreases, reaching a maximum value of 13 instead of 19. This is because LD increases the privacy loss incurred when P5 reveals his SNPs, thus reducing the total number of SNPs parent P5 can reveal without violating the family’s privacy constraints.

4.2 Previous Disclosure by Part of the Family

We want to mimic the situation where some of the family members have already revealed some of their SNPs. We simulate this by randomly selecting (with probability 0.5) some of the family members (except P5, who is the donor) who reveal a subset of their SNPs. Then, for the members

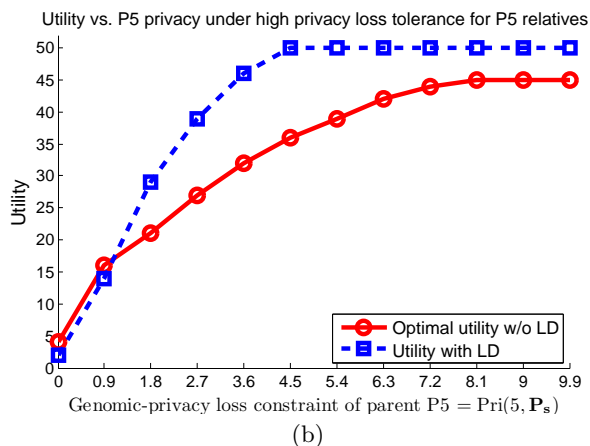
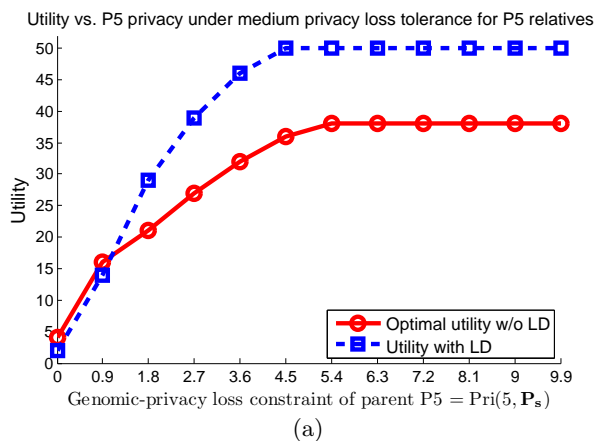


Figure 4: Utility versus privacy under (a) medium, (b) high tolerance to privacy loss for all relatives except parent P5, and varying values of privacy constraints for parent P5 (x-axis). Medium, respectively high, tolerance is defined as around half, respectively 3/4, of the total privacy loss that a relative would incur if all 50 SNPs of P5 were revealed. The x-axis represents the privacy loss constraint of P5, from no privacy loss (strongest constraint) to 9.9 privacy loss (i.e., around 0.2 privacy loss per SNP, which is a weak constraint).

who are selected to reveal their SNPs, we select, uniformly at random, some of their 50 SNPs to reveal. In the scenario we focus on, this leads to the following SNPs being revealed before the donor’s decision: 8 (different) SNPs revealed by GP1 and GP2; 35 SNPs revealed by GP3; 42 revealed by GP4; 0 by P6; 0 SNP by C7, C8, C9, C10; and 30 by C11.

We analyze the relation between utility and privacy for different genomic-privacy constraint values, for each of the eleven individuals, $\text{Pri}(i, \mathbf{P}_s^i)$. Fig. 4(a) and 4(b) illustrate the utility gain with respect to different privacy loss tolerance levels for the donor (P5). The two figures differ essentially in terms of the genomic-privacy constraints of the rest of the family members. In Fig. 4(a), the tolerance is medium; more precisely, the privacy constraint for each individual in the pedigree (except P5) is set to half of the maximum privacy loss that would be incurred by that indi-

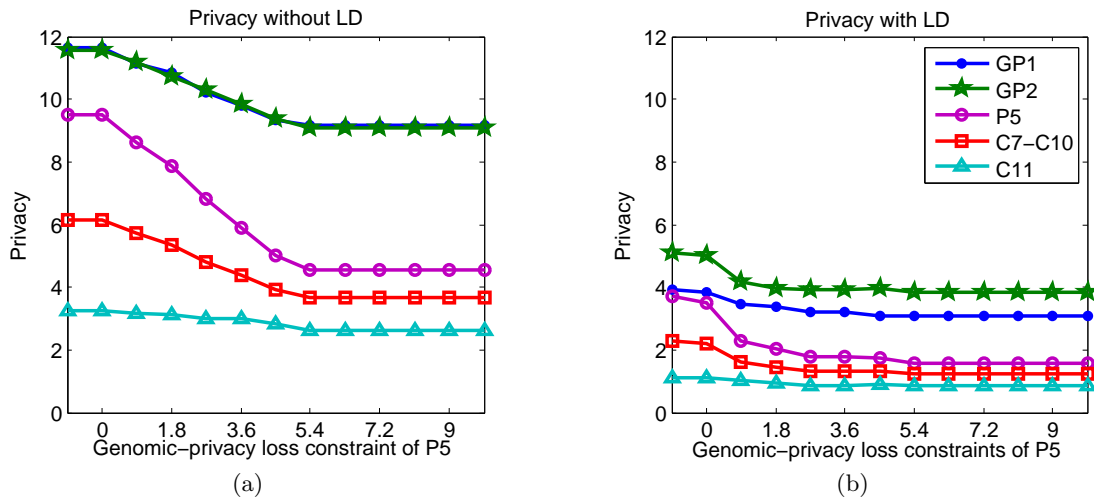


Figure 5: Genomic privacy of all family members given the genomic-privacy constraint of P5, under the same setting as in Fig. 4(a) i.e. under medium privacy loss constraints for P5 relatives: (a) privacy computed without LD, and (b) privacy computed with LD, before the fine-tuning phase. We do not show the privacy levels of GP3, GP4 and P6 as these remain constant. Note the large discrepancy in absolute privacy values and privacy losses between Fig. 5(a) and 5(b). Also notice that GP1 privacy curve is hidden by GP2 privacy curve in Fig. 5(a) (they have same privacy levels w/o LD).

vidual if the donor revealed all his SNPs. In Fig. 4(a), the tolerance is higher, set to 3/4 of the maximum privacy loss.

We first focus on the utility computed using our branch-and-bound algorithm (case w/o LD). In Fig. 4(a), we observe that the utility does not increase beyond 38 when we increase the genomic-privacy loss constraint of the donor more than 5.4. From this point, the increased privacy tolerance of the donor does not enable him to reveal more SNPs, because he is constrained by the rest of the family’s privacy requirements. In Fig. 4(b), we note that the utility keeps increasing with the privacy loss constraint of P5 because his relatives are more tolerant regarding their own privacy losses.

Looking at the utility induced once we include the LD correlations in the privacy quantification, we notice some increase in the utility. In other words, including LD enables the donor to reveal more SNPs than without LD. Utility in both curves reaches 50 SNPs after a 4.5 privacy loss constraint for the donor. This can be explained by the fact that, when LD is considered, we use Equation (5) (privacy loss with LD) instead of Equation (3) (privacy loss without LD) to compute the privacy weights for each SNP in each constraint. And the privacy loss in Equation (5) is actually smaller than in Equation (3) in this scenario, essentially because LD already decreases significantly the relatives’ privacy before the donor reveals any of his own SNPs. This is very visible in Fig. 5(a) and 5(b). In Fig. 5(a), we show the privacy levels for any family member when LD is not included in the privacy quantification. Fig. 5(b) shows the privacy levels when LD correlations are also used in the privacy quantification.

First, we notice that in both figures, it is P5’s privacy level that decreases the most, as he is the one who actually reveals new SNPs in the process. Other relatives’ privacy is only damaged due to familial correlations. At the origin of the x-axis (i.e., on the y-axis), we see the privacy levels before the donor makes a decision, i.e., before the optimiza-

tion algorithm. We notice that, here again, privacy without LD is much higher than privacy once LD is used to infer the SNPs. This is because some relatives have already revealed part of their genomic data. This is the reason, once P5 reveals his own SNPs, the privacy loss is much smaller in Fig. 5(b) than in Fig. 5(a). As a consequence, the donor (P5) can reveal more SNPs while still meeting his family’s privacy constraints, thus leading to the utility increase displayed in Fig. 4(a) and 4(b). We conclude that the values of the privacy-loss constraints have to be carefully determined by the family members or the genetic counsellors, based on family members’ privacy expectations and on whether LD is included or not in the initial inference and privacy quantification. In our case, in order to make use of the linear optimization framework, we defined the privacy loss constraints based on the privacy levels computed without LD.

Finally, we compared the optimal solutions computed with exhaustive search over a subset of 10 SNPs whose privacy weights were computed with LD, with the solutions derived from our optimization algorithm presented in Fig. 2. In the various scenarios we tested, the exhaustive search method could never find higher utility values than our fine-tuning algorithm. In all scenarios, our fine-tuning algorithm reached the maximum utility. Thus, even though we do not have any formal demonstration that the fine-tuning step is optimal, we are confident that in general it provides a very good approximation of the optimum.

4.3 Computational Complexity

As expected, the highest computation time is on average induced by the branch-and-bound algorithm (Subsection 3.2), due to the high complexity of the multidimensional knapsack problem. The non-linear extension (Subsection 3.3) is by design very efficient, as it relies on previous optimal computations and it updates a minimal set of decision variables, trading-off exact optimality for computational ef-

efficiency. This last part only requires quantifying privacy levels twice at the beginning (in the quantification box), to get the $\bar{E}_j^i(\mathbf{y}^D = \mathbf{0})$'s and $\bar{E}_j^i(\mathbf{y}^{*D})$'s, and then quantifying once per update on a decision variable y_j^{*D} .⁷

The multidimensional knapsack problem is NP-complete and admits no fully polynomial-time approximation scheme. From our experiments, we notice that the complexity of the branch-and-bound algorithm highly differs for different settings, e.g., different privacy-loss tolerance values or privacy weights. With 50 SNPs, the vast majority of the solutions were found in less than one second. However, the algorithm did not scale well for more than 50 decision variables. The positive side is that this whole process has to be undertaken only once by the donor and can be run offline. Furthermore, we considered one privacy constraint for each family member, thus eleven constraints in total. In practice, some relatives would certainly not care much about their genomic privacy, hence some constraints could be relaxed, thus enabling us to consider more SNPs in the optimization problem. Also, an advantage of the branch-and-bound algorithm is that it can be parallelized and distributed using a computer cluster. The algorithm's running time then scales linearly with the number of machines and cores [9]. Another way to reduce the complexity is to cluster subsets of SNPs together (based on the diseases they are associated with, or based on the LD correlations between them), thus trading-off the granularity of the obfuscation mechanism for computational efficiency. Note that our optimization problem can easily be adapted to deal with clusters of SNPs: We can simply define the privacy weight of one cluster as the sum of the privacy losses over the SNPs in this cluster. Finally, instead of using an exact optimization method, heuristic approaches [19] could be used to approximate the optimal solution and improve computational efficiency. We intend to study the efficiency of these approaches in future work.

5. RELATED WORK

Stajano *et al.* [32] were among the first to raise the issue of kin privacy in genomics, and to suggest discussing questions such as; Should you be allowed to disclose your genome if other relatives do not want to? Our work aims to address notably this concern. Cassa *et al.* [10] provide a framework for measuring the interdependent privacy risks between two siblings. They show that the inference error is substantially reduced when the sibling's SNPs are known, compared to when only the population frequencies are used. We generalized this evaluation of kin genomic privacy risks by considering any SNP(s) revealed by family members, LD relationships between SNPs, and well-defined privacy metrics [22]. The proposed GPPM builds upon this work.

Homer *et al.* [21] prove that de-identification is an ineffective way to protect the privacy of genomic data, which is also supported by other works [28, 35, 37]. Most recently, Gymrek *et al.* [20] showed how they identified DNAs of several individuals and families who participated in scientific studies. Building upon [21], Sankararaman *et al.* [30] provide quantitative guidelines for researchers willing to make a certain number of SNPs publicly available in GWAS, without revealing the presence of a single individual within a study group. Fienberg *et al.* [18] propose using differen-

tial privacy to protect the identities of participants in scientific study. In the same vein, Johnson and Shmatikov [25] propose privacy-preserving algorithms for computing various statistics related to the SNPs, while guaranteeing differential privacy. However, differential privacy reduces the accuracy of research results and is aimed to be applied on aggregate results. In our work, we focus on protecting individual genomes' privacy.

Some works also focus on protecting the privacy of genomic data and on preserving utility in medical tests such as (i) searching of a particular pattern in the DNA sequence [7, 34], (ii) comparing the similarity of DNA sequences [6, 8, 13, 14, 24, 27], and (iii) performing statistical analysis on several DNA sequences [26]. Furthermore, Ayday *et al.* propose privacy-preserving schemes for medical tests and personalized medicine methods that use patients' genomic data [5]. For privacy-preserving clinical genomics, a group of researchers proposes to outsource some costly computations to a public cloud or semi-trusted service provider [11, 36]. All aforementioned works make use of cryptographic protocols to protect the privacy of genomic data. In this paper, we propose a non-cryptographic approach for protecting genomic privacy.

Finally, Calmon and Fawaz propose an inference framework for evaluating privacy risks under utility constraints in a generic setting [16]. Their goal is to minimize information leakage subject to certain utility constraints. They show that their optimization problem can be cast as a modified rate-distortion problem. They eventually compare their framework with differential privacy.

6. CONCLUSION

There is little doubt that the momentum in genome sequencing will bring new challenges to data security and privacy. In this work, we convey the importance of building mechanisms for preserving genomic privacy. Such privacy goes beyond the protection of genome information of the individual to the consideration of the interests of family members. They might be unwilling to allow predictions of their SNPs based on leakage of information from one or several individuals of the kin. The approach presented here searches for balance between accuracy (utility) of genomic data and privacy by relying on graphical models and optimization.

Acknowledgements

We would like to thank Reza Shokri, Jean Louis Raisaro, and Zhicong Huang for their valuable comments on the submitted manuscript, as well as the anonymous reviewers for their helpful feedback.

7. REFERENCES

- [1] <http://www.nytimes.com/2013/03/24/opinion/sunday/the-immortal-life-of-henrietta-lacks-the-sequel.html?pagewanted=all>.
- [2] <http://www.nytimes.com/2013/08/08/science/after-decades-of-research-henrietta-lacks-family-is-asked-for-consent.html?pagewanted=all>.
- [3] <http://www.ncbi.nlm.nih.gov/projects/SNP/>.
- [4] E. Ayday, J. L. Raisaro, and J. P. Hubaux. Protecting and evaluating genomic privacy in medical tests and personalized medicine. *WPES*, 2013.

⁷Note that the computational complexity of one quantification step is $\mathcal{O}(nm)$ [22].

- [5] E. Ayday, J. L. Raisaro, P. J. McLaren, J. Fellay, and J. P. Hubaux. Privacy-preserving computation of disease risk by using genomic, clinical, and environmental data. *HealthTech*, 2013.
- [6] P. Baldi, R. Baronio, E. De Cristofaro, P. Gasti, and G. Tsudik. Countering GATTACA: Efficient and secure testing of fully-sequenced human genomes. *CCS*, 2011.
- [7] M. Blanton and M. Aliasgari. Secure outsourcing of DNA searching via finite automata. *DBSec*, 2010.
- [8] F. Bruekers, S. Katzenbeisser, K. Kursawe, and P. Tuyls. Privacy-preserving matching of DNA profiles. Technical report, 2008.
- [9] M. Budi, D. Delling, and R. F. Werneck. Dryadpot: Branch-and-bound on distributed data-parallel execution engines. In *IPDPS*, pages 1278–1289. IEEE, 2011.
- [10] C. A. Cassa, B. Schmidt, I. S. Kohane, and K. D. Mandl. My sister’s keeper?: genomic research and the identifiability of siblings. *BMC Medical Genomics*, 1(1):32, 2008.
- [11] Y. Chen, B. Peng, X. Wang, and H. Tang. Large-scale privacy-preserving mapping of human genomic sequences on hybrid clouds. *NDSS*, 2012.
- [12] G. Danezis and E. De Cristofaro. Simpler protocols for privacy-preserving disease susceptibility testing. In *GenoPri*, 2014.
- [13] E. De Cristofaro, S. Faber, P. Gasti, and G. Tsudik. Genodroid: Are privacy-preserving genomic tests ready for prime time? In *WPES*. ACM, 2012.
- [14] E. De Cristofaro, S. Faber, and G. Tsudik. Secure genomic testing with size-and position-hiding private substring matching. In *WPES*. ACM, 2013.
- [15] R. Drmanac, A. B. Sparks, M. J. Callow, A. L. Halpern, N. L. Burns, B. G. Kermani, P. Carnevali, I. Nazarenko, G. B. Nilsen, G. Yeung, et al. Human genome sequencing using unchained base reads on self-assembling dna nanoarrays. *Science*, 327(5961):78–81, 2010.
- [16] F. du Pin Calmon and N. Fawaz. Privacy against statistical inference. In *Allerton*. IEEE, 2012.
- [17] Y. Erlich and A. Narayanan. Routes for breaching and protecting genetic privacy. abs/1310.3197v1, 2013.
- [18] S. E. Fienberg, A. Slavkovic, and C. Uhler. Privacy preserving GWAS data sharing. *Proceedings of the IEEE 11th International Conference on Data Mining Workshops (ICDMW)*, Dec. 2011.
- [19] A. Fréville. The multidimensional 0–1 knapsack problem: An overview. *European Journal of Operational Research*, 155(1):1–21, 2004.
- [20] M. Gymrek, A. L. McGuire, D. Golan, E. Halperin, and Y. Erlich. Identifying personal genomes by surname inference. *Science*: 339 (6117), Jan. 2013.
- [21] N. Homer, S. Szelinger, M. Redman, D. Duggan, and W. Tembe. Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays. *PLoS Genetics*, 4, Aug. 2008.
- [22] M. Humbert, E. Ayday, J. P. Hubaux, and A. Telenti. Addressing the concerns of the Lacks Family: Quantification of kin genomic privacy. *CCS ’13: Proceedings of 20th ACM Conference on Computer and Communications Security*, 2013.
- [23] M. Humbert, E. Ayday, J.-P. Hubaux, and A. Telenti. Reconciling utility with privacy in genomics. Technical report, <https://infoscience.epfl.ch/record/201547>, 2014.
- [24] S. Jha, L. Kruger, and V. Shmatikov. Towards practical privacy for genomic computation. *Proceedings of the 2008 IEEE Symposium on Security and Privacy*, pages 216–230, 2008.
- [25] A. Johnson and V. Shmatikov. Privacy-preserving data exploration in genome-wide association studies. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1079–1087. ACM, 2013.
- [26] M. Kantarcioglu, W. Jiang, Y. Liu, and B. Malin. A cryptographic approach to securely share and query genomic sequences. *IEEE Transactions on Information Technology in Biomedicine*, 12(5):606–617, 2008.
- [27] F. Kerschbaum, M. Beck, and D. Schönfeld. Inference control for privacy-preserving genome matching. In *GenoPri*, 2014.
- [28] N. Masca, P. R. Burton, and N. A. Sheehan. Participant identification in genetic association studies: improved methods and practical implications. *International Journal of Epidemiology*, 40(6):1629–1642, 2011.
- [29] D. Nyholt, C. Yu, and P. Visscher. On Jim Watson’s APOE status: Genetic information is hard to hide. *European Journal of Human Genetics*, 17:147–149, 2009.
- [30] S. Sankararaman, G. Obozinski, M. I. Jordan, and E. Halperin. Genomic privacy and limits of individual detection in a pool. *Nature genetics*, 41(9):965–967, 2009.
- [31] W. Shih. A branch and bound method for the multiconstraint zero-one knapsack problem. *Journal of the Operational Research Society*, pages 369–378, 1979.
- [32] F. Stajano, L. Bianchi, P. Liò, and D. Korff. Forensic genomics: kin privacy, driftnets and other open questions. In *WPES*. ACM, 2008.
- [33] L. Sweeney, A. Abu, and J. Winn. Identifying participants in the personal genome project by name. Available at SSRN 2257732, 2013.
- [34] J. R. Troncoso-Pastoriza, S. Katzenbeisser, and M. Celik. Privacy preserving error resilient DNA searching through oblivious automata. *CCS ’07: Proceedings of the 14th ACM Conference on Computer and Communications Security*, 2007.
- [35] R. Wang, Y. F. Li, X. Wang, H. Tang, and X. Zhou. Learning your identity and disease from research papers: Information leaks in genome wide association study. *Proceedings of the 16th ACM CCS*, pages 534–544, 2009.
- [36] R. Wang, X. Wang, Z. Li, H. Tang, M. K. Reiter, and Z. Dong. Privacy-preserving genomic computation through program specialization. *Proceedings of the 16th ACM CCS*, pages 338–347, 2009.
- [37] X. Zhou, B. Peng, Y. F. Li, Y. Chen, H. Tang, and X. Wang. To release or not to release: Evaluating information leaks in aggregate human-genome data. *ESORICS*, 2011.