

# **Secure Discovery of Genetic Relatives across Large-Scale and Distributed Genomic Datasets**

by

Matthew M. Hong

B.Eng., Tsinghua University, 2020

Submitted to the Department of Electrical Engineering and Computer Science  
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2024

© 2024 Matthew M. Hong. This work is licensed under a [CC BY-NC 4.0](#) license.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Authored by: Matthew M. Hong  
Department of Electrical Engineering and Computer Science  
May 17, 2024

Certified by: Bonnie Berger  
Simons Professor of Mathematics  
Thesis Supervisor

Accepted by: Leslie A. Kolodziejski  
Professor of Electrical Engineering and Computer Science  
Chair, Department Committee on Graduate Students



# **Secure Discovery of Genetic Relatives across Large-Scale and Distributed Genomic Datasets**

by

Matthew M. Hong

Submitted to the Department of Electrical Engineering and Computer Science  
on May 17, 2024 in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

## **ABSTRACT**

Finding relatives within a study cohort is a necessary step in many genomic studies. However, when the cohort is distributed across multiple entities subject to data-sharing restrictions, performing this step often becomes infeasible. Developing a privacy-preserving solution for this task is challenging due to the significant burden of estimating kinship between all pairs of individuals across datasets. In this thesis, we introduce SF-Relate, a practical and secure federated algorithm for identifying genetic relatives across data silos. SF-Relate vastly reduces the number of individual pairs to compare while maintaining accurate detection through a novel locality-sensitive hashing approach. We assign individuals who are likely to be related together into buckets and then test relationships only between individuals in matching buckets across parties. To this end, we construct an effective hash function that captures identity-by-descent (IBD) segments in genetic sequences, which, along with a new bucketing strategy, enable accurate and practical private relative detection. To guarantee privacy, we introduce an efficient algorithm based on multiparty homomorphic encryption (MHE) to allow data holders to cooperatively compute the relatedness coefficients between individuals, and to further classify their degrees of relatedness, all without sharing any private data. We demonstrate the accuracy and practical runtimes of SF-Relate on the UK Biobank and All of Us datasets. On a dataset of 200K individuals split between two parties, SF-Relate detects 94.9% of third-degree relatives, and 99.9% of second-degree or closer relatives, within 15 hours of runtime. Our work enables secure identification of relatives across large-scale genomic datasets, and thus a wide range of downstream privacy-preserving collaborative studies.

Thesis supervisor: Bonnie Berger

Title: Simons Professor of Mathematics



# Acknowledgments

This thesis was made possible by many people. Bonnie Berger invited me to the amazing world of computational biology, inspired me with her acute vision, and has been an excellent advisor and a true leader. She provided me with freedom to explore my ideas and offered immeasurable support and kindness for me through many challenges. My mentor Yael Kalai's guidance on cryptography research and inspiring suggestions have been indispensable to the completion of the thesis. Her energy, intellectual curiosity and dedication to research are infectious. Hyunghoon Cho has been an extraordinary mentor, and is instrumental in shaping the development of the thesis. David Froelicher sparks many of the key ideas and is a superb collaborator, and offers enormous help. Ricky Wagner and Victoria Popic have provide invaluable resources and suggestions. I also thank Manaswitha Edupalli, Simon Mendelsohn and Matthew Mosca for their help in processing the biobank datasets, and for integrating SF-Relate into the sikit web server. The RECOMB reviewers' suggestions are extremely helpful on revising the thesis. In addition, members of the Berger lab offers valuable discussions and feedback.

I am immensely grateful to my previous research mentors for their vital and continuing guidance, support and encouragement. Yu Yu initiated my interests in the beauty of cryptography. Yuval Ishai introduced me to multi-party computation and is incredibly knowledgeable, and his deep insights and suggestions really hone my analytical skills. Amit Sahai and Aayush Jain hugely nurtured my interests in foundational cryptography research. Zhaohui Wei kindly supervised my undergraduate thesis despite the difficult COVID pandemic. A delightful collaboration opportunity with Henry Corrigan-Gibbs, Sarah Meiklejohn, Alexandra Henzinger, Vinod Vaikuntanathan crucially equipped me with the knowledge and skills to develop the applied cryptographic aspects of the thesis.

Finally, I could not express how grateful I am to my parents and sisters for their unwavering support, and my friends, for being there for me.



# Contents

<b>Title page</b>	<b>1</b>
<b>Abstract</b>	<b>3</b>
<b>Acknowledgments</b>	<b>5</b>
<b>List of Figures</b>	<b>9</b>
<b>List of Tables</b>	<b>11</b>
<b>1 Introduction</b>	<b>13</b>
<b>2 Background and Preliminaries</b>	<b>17</b>
2.1 Overview of the cross-dataset kinship estimation problem . . . . .	17
2.2 Preliminaries . . . . .	18
2.2.1 Kinship coefficients . . . . .	18
2.2.2 Locality sensitive hashing . . . . .	19
2.2.3 Homomorphic encryption . . . . .	20
<b>3 Results</b>	<b>21</b>
3.1 Overview of SF-Relate . . . . .	21
3.2 Datasets and evaluation settings . . . . .	22
3.3 SF-Relate accurately and efficiently detects close relatives between large datasets .	23
3.4 SF-Relate's accuracy-runtime tradeoffs and parametrization . . . . .	25
3.5 SF-Relate's IBD-based hashing strategy improves detection accuracy over the KING estimator . . . . .	28
3.6 SF-Relate supports alternative output settings . . . . .	29
3.7 A case study: SF-Relate reduces false positives in genome-wide association studies	32
<b>4 Algorithmic Details</b>	<b>39</b>
4.1 Existing cross-dataset approaches and their limitations . . . . .	39
4.2 Our approach to secure cross-dataset kinship estimation . . . . .	40
4.3 Dataset preprocessing . . . . .	50
4.4 Ground truth kinship preparation . . . . .	52
4.4.1 Kinship estimation using alternative non-secure methods . . . . .	52
4.5 Phenotype simulation for the GWAS case study . . . . .	53

<b>5 Discussion</b>	<b>55</b>
<b>Data and Code Access</b>	<b>57</b>
<b>References</b>	<b>59</b>



# List of Figures

1.1	SF-Relate overview.	15
3.1	SF-Relate achieves higher accuracy for samples with closer kinship and enables a trade-off between accuracy and runtime.	26
3.2	SF-Relate's runtime scales linearly in database dimension in practice.	27
3.3	Increasing sketching ratio $s$ allows almost perfect precision.	27
3.4	SF-Relate excludes the spurious 4th-degree relatives detected by KING.	29
3.5	SF-Relate and PC-Relate exclude spurious 4th-degree relatives detected by KING, when compared to RAFFI as the ground-truth.	30
3.6	SF-Relate's alternative output mode computes the revealed kinship accurately on the subsampled set of SNPs.	31
3.7	SF-Relate's alternative output modes compute accurate individuals-level statistics with customizable thresholds.	32
3.8	SF-Relate reduces false positives in multi-site GWAS.	33
4.1	The approach based on Homer et al.'s attack, as used in the iDASH 2023 Competition, does not produce a statistic separating samples with relatives and samples without relatives.	41
4.2	SF-Relate's workflow	42
4.3	SF-Relate's hashing and micro-bucketing strategies effectively assign close relatives to the same bucket.	43
4.4	The LSH for Hamming similarity retains more high-probability pairs (candidate IBD segments) than the LSH for Jaccard similarity (MinHash).	44
4.5	KING can be accurately estimated on subsampled subsets of SNPs.	51



# List of Tables

3.1	Symbols, parameters and default values. . . . .	34
3.2	SF-Relate achieves near-perfect accuracy for identifying close relatives in UK Biobank and All of Us datasets. . . . .	35
3.3	SF-Relate effectively detects relatives across heterogeneous datasets. . . . .	35
3.4	SF-Relate's recall remains high across different subpopulations. . . . .	36
3.5	SF-Relate scales efficiently to large datasets. . . . .	36
3.6	SF-Relate achieves high recall across various parameter settings. . . . .	37
3.7	SF-Relate outperforms KING and achieves comparable results to the more advanced methods, PC-Relate and RAFFI. . . . .	37
3.8	Assignment of UK Biobank assessment centers to geographic areas. . . . .	38



# Chapter 1

## Introduction

Collaborative studies that aim to jointly analyze genomic data from multiple parties are essential for increasing the sample sizes to enhance the discovery of biomedical insights. However, when sharing individual-level genetic data is not feasible due to privacy concerns (e.g., [1]), the range of joint analyses that can be performed is severely limited. As a result, many existing collaborations have relied on simplified analysis pipelines where some key analysis steps, such as cohort identification, quality control procedures, and correction for confounding factors (e.g., population structure) are performed independently by each party on their respective datasets without considering the pooled data. This presents a key barrier to realizing the full potential of collaborative genomics research.

An important analysis task that is commonly omitted in collaborative studies is the identification of genetic relatives across isolated datasets. Identifying and excluding close relatives within a study cohort is a standard step in many genetic analyses (e.g., genome-wide association studies [2]), because the presence of relatives can introduce bias and confounding that undermine the accuracy of study results [3]–[11]. For large-scale biobanks, a substantial portion of study participants may be biologically related; an estimated 32.3% of the individuals in the UK Biobank dataset [12] have a third-degree or closer relative in the same dataset. Thus, controlling for relatedness can have a major impact on the size and composition of the analysis cohort, and thereby affect the final analysis results. Removal of duplicate individuals across datasets is a special case of detecting relatives, which our work also addresses.

There are several key hurdles to identifying related individuals across datasets. Unlike other analysis tasks that derive aggregate-level insights from the pooled data, such as association tests,

finding relatives is an inherently sensitive task, directly operating at the level of individuals. Consequently, most existing approaches for cross-site analysis, e.g., meta-analysis or federated learning, cannot be applied in our setting, as they rely on sharing aggregate-level data between the parties. Furthermore, despite the growing literature on cryptography-based secure computation algorithms for biomedicine [13]–[16], which allow joint computation without sharing private data between parties, to our knowledge no practical solution exists for relative detection. This is mainly because standard tools for evaluating kinship require all pairs of individuals between two datasets to be compared [17]–[19] or involve complex combinatorial operations (e.g., string matching) [20]–[22], which incur an overwhelming cost when implemented using cryptographic operations.

In this thesis, we introduce SF-Relate, a scalable and privacy-preserving solution for identifying relatives across distributed datasets, as illustrated in Figure 1.1. Our novel approach entails each party locally assigning their samples to buckets such that related samples are more likely to be assigned to the same bucket, via *locality sensitive hashing* (LSH) [23], and then securely estimating kinship only between samples that end up in the same bucket across parties. We devise a data encoding scheme for LSH aimed at capturing identity-by-descent (IBD) segments in genetic sequences, thus effectively grouping together samples that are likely to be related. Furthermore, we introduce a new strategy for bucket assignment, in which buckets obtained through multiple LSH trials are merged and filtered to obtain a set of size-one buckets (referred to as *micro-buckets*), enabling efficient comparison between parties. We illustrate how our techniques guarantee a high probability of detecting related samples, minimize computational overhead, and ensure that no private information is revealed between the parties. Moreover, to estimate kinship coefficients ([Algorithmic Details](#)) for pairs of samples in the same bucket between parties without sharing data, we introduce a provably end-to-end secure approach that leverages homomorphic encryption, a cryptographic technique allowing for direct computation on encrypted data, combined with efficient distributed computation strategies. SF-Relate keeps each party’s data confidential throughout the computation, revealing only the final output to each party, which is exactly the list of their own samples that have at least one relative in another dataset.

Our results show that SF-Relate remains practical even for large-scale biobank datasets, with a runtime of 14.5 hours to identify the relatives among a pair of datasets consisting of 200K real genomic samples in total. Our privacy-preserving solution enables us to identify more than 97% of

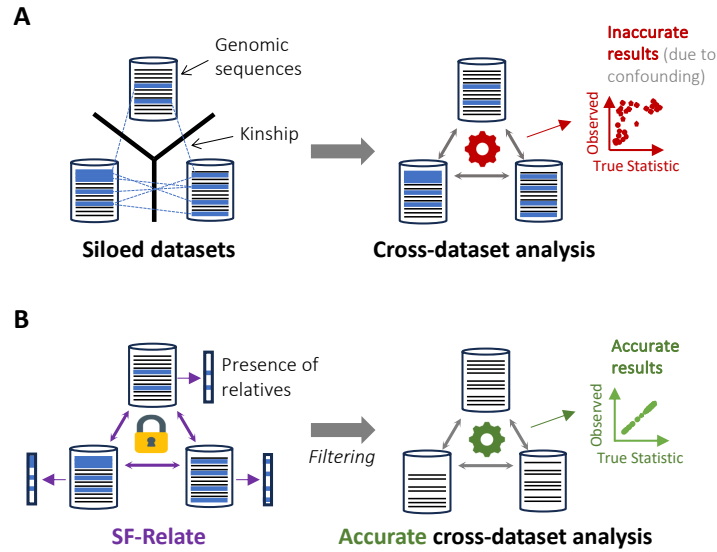


Figure 1.1: **SF-Relate overview.** When genetic relatives across datasets cannot be identified due to restricted data sharing, joint studies can suffer from bias and confounding (**A**). SF-Relate allows parties to securely identify and correct for cross-dataset relatives to enhance downstream analysis (**B**).

the related samples (third-degree or closer) by performing only 0.13% of the computation required by the naive solution in which all pairs of samples are compared between the parties.

We expect SF-Relate to enable a wide range of privacy-preserving collaborative studies. This thesis is based on a co-authored paper to be orally presented at 28th Annual International Conference on Research in Computational Molecular Biology (RECOMB 2024) [24], and its full version [25] has been invited to the corresponding special issue of Genome Research.





# Chapter 2

## Background and Preliminaries

### 2.1 Overview of the cross-dataset kinship estimation problem

We consider the setting in which multiple parties (e.g., researchers from different institutions) wish to identify genetically related individuals between their private datasets while protecting the confidentiality of their data (**Figure 1.1**). The goal of the parties is to use this information to facilitate downstream collaborative analysis by excluding duplicate or related individuals from the analysis to minimize bias in statistical analyses. For simplicity, we focus on the setting with two parties, each holding a dataset including phased haplotypes of  $n$  individuals over  $m$  genetic variants, such as single nucleotide polymorphisms (SNPs). The desired output for each party is a list of individuals in their private dataset who have at least one “close” relative in the other party’s dataset with respect to some relatedness threshold. In our work, we mainly consider the detection of third-degree or closer relatives, which are the closest relations most commonly used in genomic studies and data releases (e.g., UK Biobank [12]). We assume each input dataset to be locally phased by each party (i.e., each individual’s genome is represented as two haplotype sequences), which is crucial for capturing identity-by-descent (IBD) sharing patterns, as we describe later. We further consider a threat model where the parties are *honest-but-curious*; i.e., they faithfully follow the protocol as specified, but might attempt to infer private information about the other parties’ datasets based on the data observed during the process. Given this model, we aim to provide formal privacy guarantees for each party’s input dataset, ensuring that no information is revealed to other parties except for what can be gleaned from each party’s respective output and global parameters of the problem

(e.g., dataset dimensions and security parameters).

## 2.2 Preliminaries

### 2.2.1 Kinship coefficients

For every individual, each of their two *haplotypes* is represented a string  $\sigma \in \Gamma^m$  of  $m$  variants. The alphabet  $\Gamma \subset \mathbb{N}$  is chosen such that different SNPs are encoded as distinct integers. An individual's *genotype vector* is a vector  $\mathbf{x} \in \{0, 1, 2\}^m$ , where each element  $x_i$  represents the number of minor alleles at the  $i$ -th variant.

The *kinship coefficient*  $\phi$  between a pair of individuals is defined as the probability that a pair of randomly sampled alleles is identical by descent (IBD), i.e., when the pair of alleles is identical due to genetic inheritance rather than by chance. For example, since human genomes are diploid, a direct descendant inherits exactly one set of chromosomes from each parent; in this case, the kinship coefficient between an individual and his or her parent is, in principle,  $0.5 \cdot 0.5 = 0.25$ .

Existing approaches for estimating the kinship coefficients typically fall into one of two classes: *distance-based methods*, which use a notion of distance between two genotype vectors that often incorporate information about minor allele frequencies (MAF) [17], [19]; and *IBD-segment-based methods*, which first identify long shared segments between individuals that are likely due to IBD [21], [22], [26] and estimate kinship based on the extent of these shared segments. In a benchmarking study, Ramstetter et al. [27] show that both approaches achieve high accuracy for up to third-degree relationships, while agreement becomes weaker for more distant relationships, which we also observe between PC-Relate and KING (Figure 3.4). Although IBD-segment-based methods generally offer a more accurate estimation of kinship by analyzing IBD sharing patterns, distance-based approaches represent an efficient alternative that does not involve costly string matching, which often leads to substantially higher runtime (e.g., days vs. minutes [27]). As a result, distance-based methods have been more commonly applied to large datasets [12]. Recently proposed methods (e.g., RaPID [21]) introduce hashing techniques to improve the scalability of IBD segment detection, which in some cases exceed the efficiency of distance-based methods due to the quadratic scaling of pairwise distance calculation. However, IBD segment finding methods are

still combinatorial in nature and cannot be efficiently implemented using existing secure computation techniques. In our work, we adopt a distance-based approach to minimize the computational overhead associated with secure computation but simultaneously exploit IBD sharing patterns to significantly improve the scalability of our approach.

Our work addresses the problem of applying the widely adopted distance-based method for kinship estimation, the KING-robust estimator (referred to as KING in what follows for simplicity) [17], to find relationships between two datasets. This method is implemented in several standard genomic analysis toolkits, such as Hail [28] and PLINK [18], and recently the UK Biobank released relatedness data for individuals in the dataset using this estimator [12]. KING estimates the kinship coefficient between two genotype vectors  $\mathbf{x}$  and  $\mathbf{y} \in \{0, 1, 2\}^m$  using the following formula:

$$\phi(\mathbf{x}, \mathbf{y}) = \frac{1}{2} - \frac{1}{4} \cdot \frac{\|\mathbf{x} - \mathbf{y}\|^2}{\min\{h_{\mathbf{x}}, h_{\mathbf{y}}\}}, \quad (2.1)$$

where each element in  $\mathbf{x}$  and  $\mathbf{y}$  represents the minor allele dosage of a genetic variant, and  $h_{\mathbf{x}}$  and  $h_{\mathbf{y}}$  represent the fraction of heterozygous loci in each vector (i.e., the heterozygosity of the individual).

### 2.2.2 Locality sensitive hashing

Locality sensitivity hash (LSH) functions [23] map similar items to the same value more frequently based on a similarity notion. The Hamming LSH specifically relies on the Hamming similarity (defined as the number of equal coordinates between vectors), projects a vector onto one random coordinate and uses its value as the output. The probability that two vectors produce equal values is proportional to their Hamming similarity in this case. Another frequently employed LSH function, MinHash [29], captures The Jaccard similarity, i.e. the ratio of the size of the intersection to the size of the union, between sets. It achieves this by applying the same random permutation to each set, and using the identity of the first element as the output. We evaluate both LSH functions in Chapter 4.

### 2.2.3 Homomorphic encryption

Homomorphic Encryption, first conceived in the 1970s [30], allows computation over encrypted data without decryption. Modern schemes, such as the BGV [31] and CKKS [32] schemes, encode a vector of real number values in a single ciphertext and is well-suited for calculations where a small amount of noise can be tolerated. They provide operations for addition, multiplication, and rotation (i.e., permutation of elements in a vector) of encrypted values in a ciphertext while providing the single instruction, multiple data (SIMD) property. In such schemes, operations involving plaintext (unencrypted) data are substantially more computationally efficient; e.g., ciphertext-plaintext multiplication is seven times faster than ciphertext-ciphertext multiplication based on our parameter setting. Multiparty homomorphic encryption (MHE) [33], [34] extends the CKKS scheme [32] to the setting with multiple parties, by secret-sharing the decryption key and constructing a shared collective encryption key. Any party can encrypt data and perform homomorphic computations locally, but decryption can be performed only if all parties cooperate. In SF-Relate, all exchanged data are encrypted under the collective encryption key and only the final result can be decrypted with the cooperation of all parties.

Although MHE is efficient on operations consisting of simple algebra, e.g. vector distance computations, SF-Relate needs to employ additional algorithmic techniques, including locality sensitive hashing, for efficient relative detection at scale. This is because MHE per se does not guarantee the protection of the access patterns to individual data in the database.

# Chapter 3

## Results

### 3.1 Overview of SF-Relate

SF-Relate enables multiple parties to detect cross-site relatives in their joint dataset without having to share any sensitive information (Figure 1.1). The input dataset for each party includes phased haplotype sequences from individuals within that party’s cohort. We consider the parties to be honest-but-curious, meaning that they follow our analysis protocol faithfully but might try to infer information about other parties’ datasets based on what they observe individually during the protocol execution. Based on this model, SF-Relate guarantees end-to-end confidentiality for each party’s input dataset, protecting it from other parties in the protocol. During the protocol, any data exchanged between the parties are encrypted such that decryption requires the participation of *all* parties, thus ensuring a high level of protection. This approach allows parties to disclose only the information they agree to reveal, such as the final output.

To efficiently scale to large datasets, SF-Relate follows a two-step approach (Chapter 4). In *Step 1: Hashing and Bucketing*, each party locally evaluates a series of hash functions on each individual’s haplotype sequences to assign the individual to buckets across a collection of hash tables, such that related individuals are more likely to be assigned to the same bucket index. For this purpose, we devise a novel encoding scheme that splits and subsamples genotypes into *k*-SNPs (similar to *k*-mers, but non-contiguous; SNP: single nucleotide polymorphism), such that the similarity between *k*-SNPs reflects extended runs of identical genotypes, typically indicative of relatedness. We then leverage locality sensitive hashing (LSH; Section 2.2.2) to derive bucket

indices from the k-SNPs. To capitalize on the fact that related samples will likely be assigned to the same buckets multiple times, SF-Relate merges buckets with the same indices across multiple hash tables (produced by different subchromosomes) and then filters every bucket down to a *single* element, thus minimizing the number of costly kinship evaluations. We refer to this as a micro-bucketing strategy. *A somewhat surprising finding of our work is that, despite the extreme level of filtering applied to each bucket during this process, our strategy enables accurate detection of relatives with remarkable efficiency.* At the end of this process, each party obtains a single hash table with size-one buckets, effectively an ordered list of samples.

In *Step 2: Secure Kinship Evaluation*, the parties securely perform element-wise comparisons between their ordered lists of samples from Step 1. Each comparison involves evaluating a standard estimator of the kinship coefficient KING-robust (Section 2.2.1). To calculate the estimator without revealing private information between the parties, we employ multiparty homomorphic encryption (MHE; Section 2.2.3). Multiparty homomorphic encryption [33], encrypts data such that ciphertext can be directly used in computation without needing to be decrypted first, and decryption requires the cooperation of all parties. To minimize the computational overhead of MHE, SF-Relate uses sketching techniques on input haplotypes to reduce data dimensionality before performing kinship computations. Furthermore, our protocol is optimized to maximize the use of operations on local, non-encrypted data, which are significantly more efficient than operations on encrypted data. Finally, the encrypted results are compared to relatedness thresholds and aggregated. For each individual, each party obtains an indicator that reflects the presence or absence of a close relative in the other dataset.

We detail our algorithms and novel techniques in Chapter 4.

## 3.2 Datasets and evaluation settings

To evaluate SF-Relate, we obtain three genomic datasets of varying sizes, including a dataset of 20K samples (individuals) with 1M SNPs from the All of Us Research Program (AoU) [35] and two datasets from the UK Biobank (UKB) [12] including 100K and 200K samples, respectively, both with 650K SNPs. The two UKB datasets are uniformly sampled from the full UK Biobank release v3 ( $n = 488,377$ ), and the AoU dataset comprises the first 20K individuals in the All of Us release v5

( $n = 98,590$ ). We then evenly split each dataset into two parts to emulate a cross-dataset analysis involving two parties. We compute the ground-truth by evaluating all pairwise kinship coefficients using the KING approach ([17]; see Chapter 4) in plaintexts on a set of ancestry-agnostic SNPs, as in UK Biobank’s pipeline [12]. Further details on dataset preparation is in Chapter 4.

We evaluate the accuracy of our method in detecting close relatives between two datasets using the standard metrics of *recall* and *precision*. Recall represents the fraction of samples with a close relative in the other dataset (as determined by the baseline KING method given a threshold) that SF-Relate successfully identifies. Precision represents the fraction of samples identified by SF-Relate as having a close relative in the other dataset that actually have such a relationship according to the baseline method. For evaluation of computational costs, we measure the elapsed wall-clock time and the total number of bytes sent from one party to another (given the symmetry of SF-Relate’s computation) for runtime and communication costs, respectively.

We perform all of our experiments on virtual machines (VMs) on the Google Cloud Platform (GCP). This represents a realistic setting where parties use a cloud service provider to access high-performance computing resources that may not be readily available in the local environment. Furthermore, many biobank datasets, including AoU and UKB, are now hosted on cloud-native environments for data analysis. For UKB, we use two VMs (one for each party) with 128 virtual CPUs (vCPUs) and 856 GB of memory (n2-highmem-128) co-located in the same zone in GCP. For AoU, we emulate the two parties in a single VM with 96 vCPUs and 624 GB memory due to the constraints of the provided data analysis platform. We implement our protocols using the CKKS implementation in Lattigo [36]. Table 3.1 summarizes all symbols and default parameters.

### 3.3 SF-Relate accurately and efficiently detects close relatives between large datasets

We summarize our results on the AoU and UKB datasets in Tables 3.2 and 3.5. Across all three datasets (AoU-20K, UKB-100K, and UKB-200K), SF-Relate obtains near-perfect recall and precision (both exceeding 97% in all cases) for detecting the presence of 3rd-degree or closer relationships between two parties. Calculating the recall separately for each relatedness degree from 0th (monozygotic twins) to 3rd, we observe that most missing relationships are for the 3rd degree;

SF-Relate finds *all* existing relationships up to the 2nd degree in all three datasets, except for the 2nd degree in UKB-200K, for which it missed 2 out of 1711 individuals with a relative. The recall metric for third-degree relationships remains high—above 94% for all three datasets. Note that the more distant the relationship, the more difficult it is to detect, because the IBD segments become more scattered and reduced in quantity, which in turn results in a lower rate of surviving the filtering step in micro-bucketing. In UKB-200K, a small fraction (5%) of 3rd-degree relatives, missed by SF-Relate, correspond to those with kinship coefficients near the 4th-degree threshold (Figure 3.1), suggesting that some of them may not be real 3rd-degree relationships considering the stochastic nature of the kinship estimator.

Furthermore, SF-Relate consistently achieves high detection accuracy across a variety of populations with distinct ancestry backgrounds. SF-Relate maintains a recall rate exceeding 98% for a dataset comprising individuals of African ancestry (Table 3.3). SF-Relate largely remains effective in multi-ancestry datasets, achieving a recall higher than 80% across all subpopulations (Table 3.4). Ancestry groups with the lowest recall (Indian and Other with 84.4% and 82.4%, respectively) are associated with small sample counts, suggesting that the slight reduction in recall may be due to sampling noise. Taken together, these results demonstrate SF-Relate’s accurate relative detection performance across a range of datasets, which is achieved without revealing any private information between the two parties due to SF-Relate’s use of secure computation techniques when jointly analyzing the two datasets.

Despite the overhead of cryptographic protocols for secure computation, the runtime of SF-Relate remains practical for all three datasets, resulting in 5.8, 7.3, and 14.5 hours of runtime for AoU-20K, UKB-100K, and UKB-200K, respectively. We note that the doubling of runtime from UKB-100K to UKB-200K reflects the linear scaling of SF-Relate in the number of individuals in the dataset, since these datasets were analyzed using the same computing environment, unlike AoU. More precisely, the computational cost of the MHE calculation of pairwise kinship coefficients, which is the main computational bottleneck of SF-Relate, grows linearly in both the number of SNPs after sketching and the size of the hash table. Although both parameters can be adjusted by the user, to maintain accurate performance, these parameters need to be linearly scaled with the total number of SNPs and individuals in the original input dataset, respectively. We provide a systematic demonstration of the linear scaling of SF-Relate’s runtime in Figure 3.2. The observed communica-

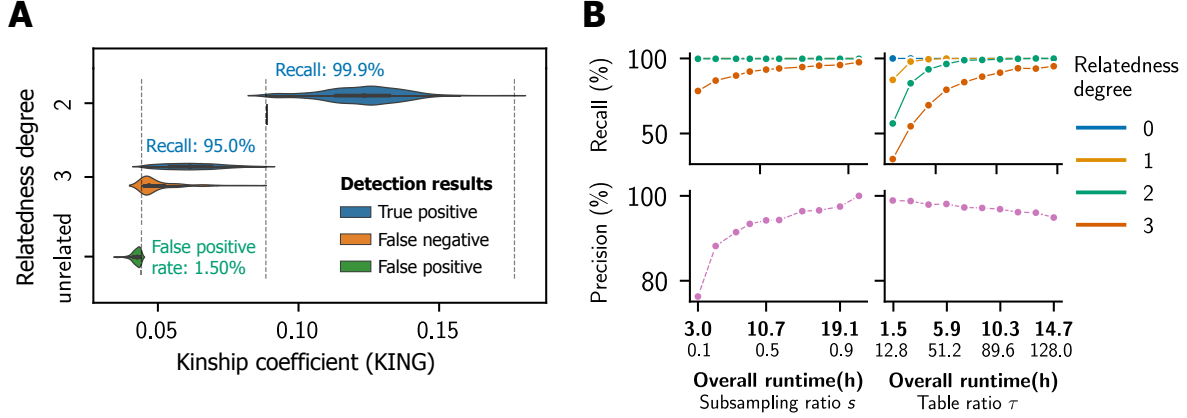


tion costs in the order of tens of terabytes (e.g., 93.2 TB for UKB-200K) are not small, but our results demonstrate that transferring such large amounts of data does not lead to impractical runtimes. In addition, we note that more than 99% of the communication bandwidth is due to the exchange of encrypted hash tables including the sketched haplotypes, which can, in principle, be transferred in a single round of communication. Therefore, we expect the impact of high communication on runtime to be minimal even in a wide-area network (WAN) setting with high communication latency (round-trip delay).

Furthermore, we highlight that without the hashing and bucketing strategy we introduced in SF-Relate, it would not be feasible to securely detect relatives between datasets by all-pairwise computation of the kinship coefficient (**All-pairwise** in Table 3.5). Even with our efficient MHE implementation of the kinship calculation over the sketched haplotypes, performing all-pairwise comparisons for the UKB-200K dataset is estimated to take 1.3 years based on the same computational setting. On the contrary, SF-Relate obtains practical runtimes by significantly reducing the number of candidate individual pairs to test without compromising accuracy through our novel use of LSH hash tables. Remarkably, SF-Relate makes only 1.28%, 0.26%, and 0.13% of pairwise comparisons compared to the total number of individual pairs between the two datasets for AoU-20K, UKB-100K, and UKB-200K, respectively (Table 3.2). This drastically reduces not only the runtime but also the communication costs; e.g., our MHE implementation of the all-pairwise computation would require 65 PB of communication without our hashing techniques (Table 3.5).

### 3.4 SF-Relate’s accuracy-runtime tradeoffs and parametrization

The runtime and accuracy of SF-Relate are primarily influenced by the hash table size  $N = \tau \cdot n$ , determined by the dataset size  $n$  and table ratio ( $\tau$ ; Chapter 4), and the sample size for comparison, determined by the subsampling ratio ( $s$ ) and the number of SNPs (Chapter 4). As shown in Figure 3.1, increasing  $s$  improves the overall recall and precision, while increasing  $\tau$  enables the detection of more distant relationships, also increasing the overall recall. However, the runtime depends linearly on both  $s$  and  $\tau$ , highlighting the trade-off between SF-Relate’s accuracy and



**Figure 3.1: SF-Relate achieves higher accuracy for samples with closer kinship and enables a trade-off between accuracy and runtime.** (A) We plot the distribution of kinship coefficients (KING) stratified by the (closest) relatedness degree of the relative pairs and by whether they were detected by SF-Relate as related. Misclassifications by SF-Relate are concentrated around kinship thresholds for different relatedness degrees, indicated by vertical dashed lines. (B) We vary the subsampling ratio ( $s$ ) and the table ratio ( $\tau$ ) parameters in SF-Relate and report the resulting precision and recall for different relatedness degrees. For precision, only the overall metric is shown for detecting 3rd-degree or closer relatives. By default,  $s = 0.7$  and  $\tau = 128$ . These parameters determine the trade-off between the runtime and accuracy of SF-Relate.

runtime. We expect the optimal trade-off to depend on the application setting.

For example, if users want to focus on identifying relatives up to the 2nd degree within the UKB-200K dataset, they could set the table ratio  $\tau$  to 64 and the subsampling rate  $s$  to 0.7, instead of  $\tau = 128$  and  $s = 0.7$  in our experiments (Chapter 4). This results in a two-fold improvement with respect to our experiments due to halving of the hash table size. Even in this scenario, users would maintain an effective detection rate of over 95% for individuals with relationships closer than the 2nd degree (Figure 3.1B). Alternatively, if users want to achieve perfect accuracy, they can increase  $s$  and  $\tau$ . Increasing  $s$  from 0.7 to 1 (i.e., no sketching), improved SF-Relate’s overall recall on UKB-200K from 97.0% to 98.7% and precision from 98.5% to 99.9% (Table 3.2 and Figure 3.3). The runtime increased from 14 to 21 hours. Furthermore, by doubling the table ratio  $\tau$ , SF-Relate achieves perfect accuracy for relations up to the 3rd degree, while doubling its runtime. Overall, SF-Relate’s recall remains consistently high, above 95%, across a wide range of parameters, and only starts to decrease when the parameters significantly deviate from the default setting (Table 3.6).

To choose suitable values for  $s$  and  $\tau$  in practice, we recommend that users first determine the farthest relationships they wish to detect and an acceptable level of recall. Using Figure 3.1, they can then determine the required  $s$  and hash table size  $N = \tau n$ . The expected runtime can be es-




Figure 3.2: **SF-Relate's runtime scales linearly in database dimension in practice.** In **A**, we perform a fixed number of kinship computations  $N$  while varying the number of SNPs by varying the subsampling rate  $s$ . In **B** and **C**, we increase the number of kinship computations  $N$  while keeping the subsampling rate  $s$  at  $.7$ . We report the runtime of the MHE-Phase 1 and MHE-Phase 2 in (B) and (C) separately.




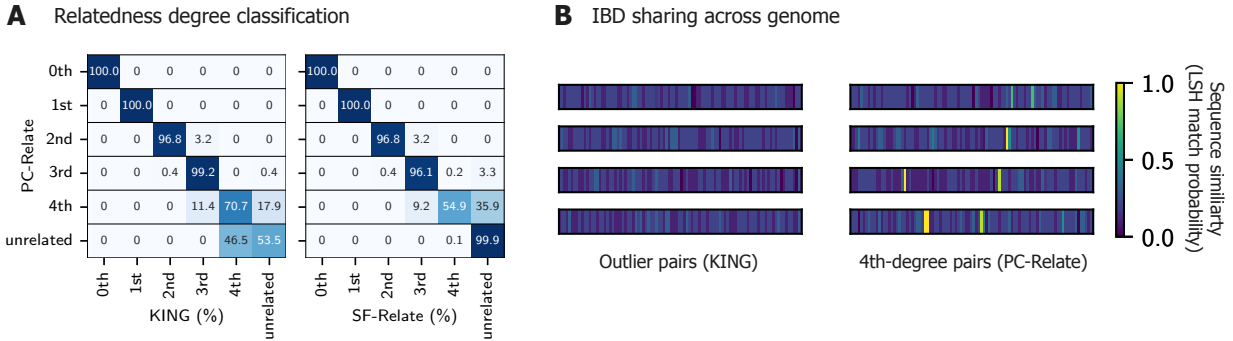
Figure 3.3: **Increasing sketching ratio  $s$  allows almost perfect precision.** We perform SF-Relate on UKB-200K under the alternative sketching (subsampling) ratio  $s = 1$ , and shows counts of individuals based on their closest relations and SF-Relate's detection results. Using the full set of SNPs in SF-Relate allows significant reductions in false positives and false negatives due to the sketching noise, with a moderate increase of runtime from 14.5 to 21 hours.

timated by considering the linear relationship between these parameters and SF-Relate’s runtime. To select the *Hashing and Bucketing* parameters, the users may initially opt for our recommended parameters (Table 3.1), which consistently achieve accurate and efficient performance across all our datasets. Optionally, users can conduct a local assessment using their own dataset to verify and fine-tune the chosen parameters, which involves examining the number of shared genomic segments between local relatives and their matching probabilities (as depicted in Figure 4.3 in Chapter 4) and estimating the detection rate among local relatives.

### 3.5 SF-Relate’s IBD-based hashing strategy improves detection accuracy over the KING estimator

SF-Relate leverages a secure implementation of the KING formula [17] to estimate relatedness (Chapter 4). Nevertheless, SF-Relate can sometimes lead to even more accurate identification of relatives than KING. This is due to SF-Relate’s IBD-based approach to bucketing the samples, which helps filter out spurious pairs of samples that are identified by KING as being related, but in fact do not share any long IBD segments. We confirm this in our comparisons with PC-Relate [19] and RAFFI [37], recent methods for kinship detection designed to improve upon KING’s accuracy by correcting for population structure and incorporating IBD segment detection, respectively. We evaluate all methods on a subset of 20K samples from the UKB-200K, distributed across two parties. As expected, they produce highly similar results for up to the 3rd degree relatives (Figures 3.4 and 3.5 and Table 3.7). However, for detecting 4th-degree relatives, standard KING erroneously identifies numerous individuals as being related due to the presence of an outlier sample that is detected as related to *thousands* of samples in the other dataset; SF-Relate, akin to the more advanced tools PC-Relate and RAFFI, successfully avoids these errors. Similar outlier-related issues regarding KING have been noted in UK Biobank’s official report on relatedness inference (Supplementary in [12]). In Figure 3.4, we visualize sequence similarity between four pairs of haplotypes involving the outlier sample, compared with typical 4th-degree relative pairs identified by PC-Relate. We observe light yellow bands of high sequence similarity regions exclusively in PC-Relate’s pairs, which signifies real IBD segments. This suggests that SF-Relate’s bucketing approach based on

IBD segments can effectively distinguish outlier pairs from real ones, thus leading to more accurate detection of relatives.



**Figure 3.4: SF-Relate excludes the spurious 4th-degree relatives detected by KING.** (A) We show confusion matrices assessing the relatedness classification accuracy of KING (left) and SF-Relate (right), comparing with the output of PC-Relate as the ground-truth. SF-Relate is performed in plaintext (i.e., without MHE), focusing on the evaluation of the bucket assignment. Unlike SF-Relate, KING classifies many unrelated samples as 4-th degree relatives. Most of these pairs involve the same outlier sample, which has many spurious relationships. (B) We verify that pairs involving the outlier do not exhibit IBD sharing patterns (left), evident in 4th-degree pairs from PC-Relate (right). Four example pairs are shown for both cases. For each pair, we compute the Hamming similarity between the two samples of genomic segments across the genome. Bright yellow bands represent likely IBD segments. The locations of the bands are randomly permuted to obscure their positions.

### 3.6 SF-Relate supports alternative output settings

In SF-Relate’s default setting, each party learns only whether each local individual has a close relative within the joint dataset. SF-Relate offers an option to output kinship computation results with more detailed granularity, summarizing them at various levels to address a range of analysis needs. Alternative outputs include the closest relatedness degree for each individual, the maximum kinship coefficient for each individual (discretized), and the full list of computed kinship coefficients. The SF-Relate output remains accurate in all settings: The individual kinship coefficients computed by SF-Relate exhibit a small average absolute error of  $5.8 \times 10^{-4}$  when compared to KING (Figure 3.6). The closest degree reported for each individual matched with the ground truth for 99.9% of the individuals (Figure 3.7). Calculating the maximum kinship coefficient, discretized using the smallest bin width of 0.016, SF-Relate accurately assigns more than 85% of the samples to the correct bins, and more than 99.9% were within one bin of the true output (Figure 3.7).

Comparison of relatedness degree classification (against RAFFI as the ground truth)

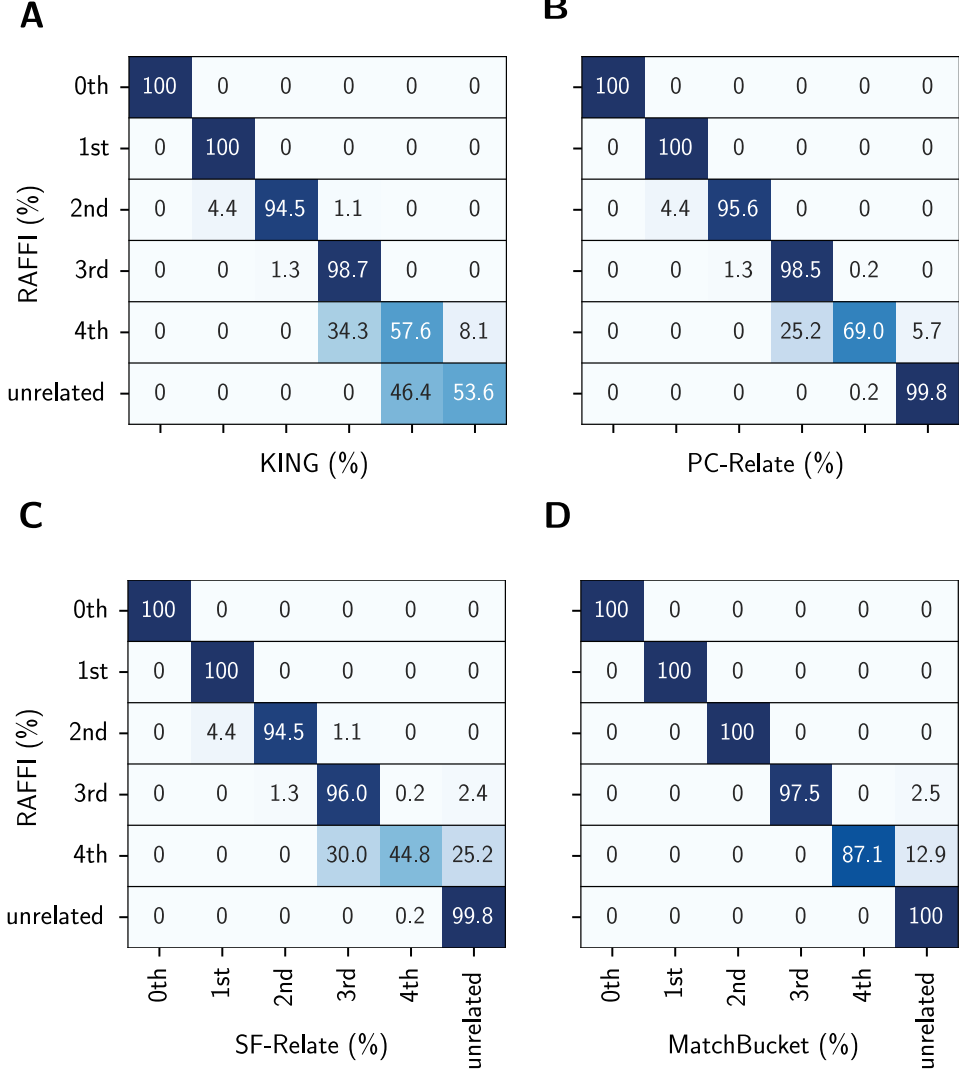


Figure 3.5: **SF-Relate and PC-Relate exclude spurious 4th-degree relatives detected by KING, when compared to RAFFI as the ground-truth.** On a subset with 20K samples from UKB-200K, we present the confusion matrices assessing the relatedness classification accuracy of KING (**A**), PC-Relate (**B**) and SF-Relate (**C**), comparing them with the output of RAFFI as the ground-truth. *MatchBucket* (**D**) denotes the (non-private) hybrid approach where RAFFI is performed in plaintext on pairs in SF-Relate’s corresponding buckets and serves as a reference. Both SF-Relate and PC-Relate label RAFFI-unrelated individuals as unrelated, unlike KING, which labels them as 4th-degree relatives. This suggests that both methods avoid the spurious relationships identified by KING.


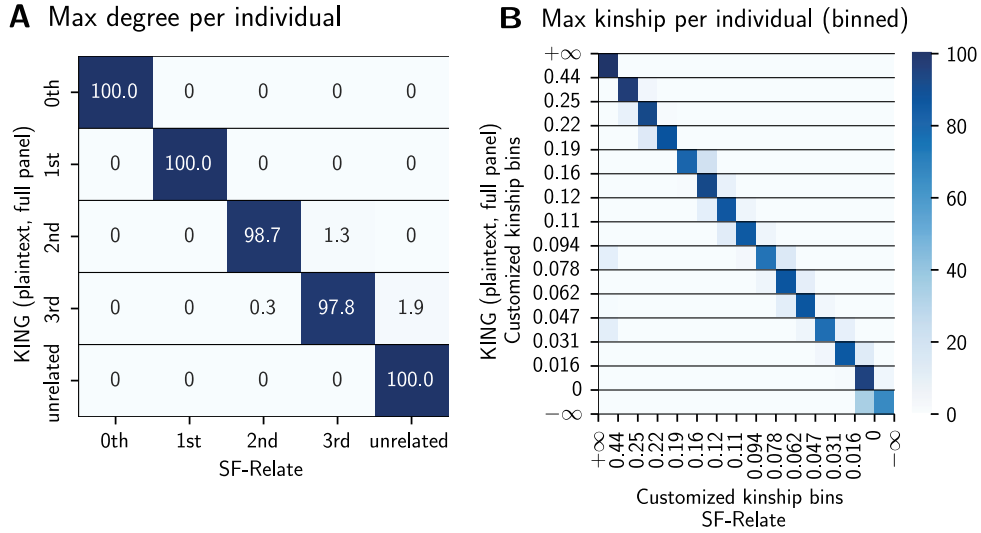


Figure 3.6: **SF-Relate’s alternative output mode accurately reports all computed kinship coefficients.** On a 10% random subset of kinship coefficients that SF-Relate computed on UKB-200K, we evaluate the accuracy of the alternative setting of SF-Relate where the full list of kinship coefficients is revealed. As shown in the plot, SF-Relate’s output accurately matches with KING, indicating that the MHE noise is negligible with respect to the kinship coefficients being computed on the subsampled set of SNPs.

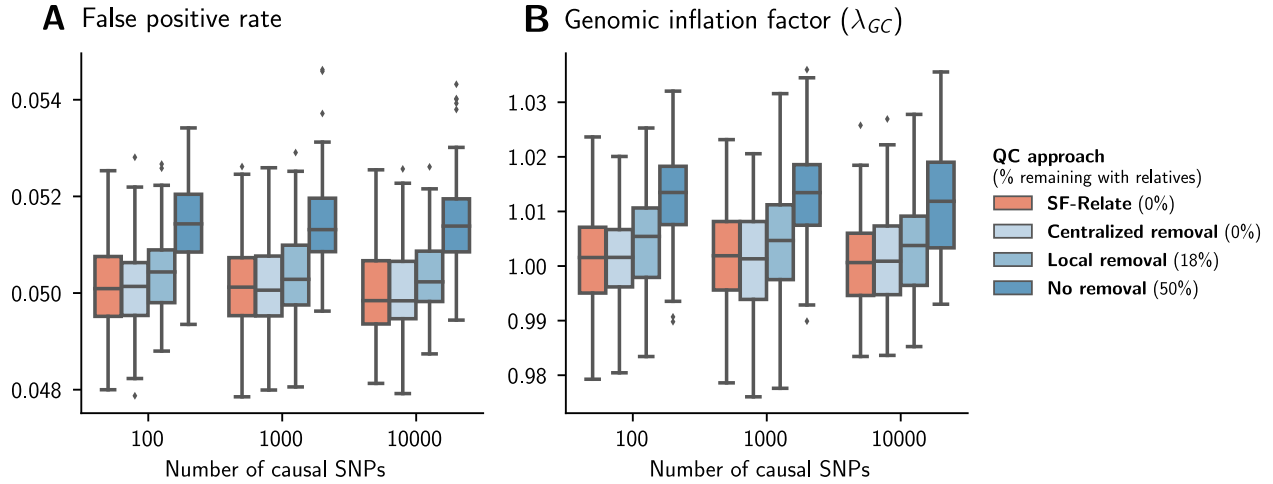


**Figure 3.7: SF-Relate’s alternative output modes compute accurate individuals-level statistics with customizable thresholds.** On a subset of 10% kinship coefficients computed in UKB-200K’s hash tables, we evaluate SF-Relate’s alternative output modes that support the comparison of the computed kinship coefficients with a sequence of thresholds. Numbers in the cells or their colors indicate the corresponding recall rates. In **(A)**, we choose the thresholds to be the recommended kinship cutoff by KING in Manichaikul et al. [17]. In **(B)**, the thresholds defining the refined bins are marked on the axes. The predictions perfectly align with the ground-truth KING predictions computed on the full SNP panel on more than 99.9% and 85% individuals in **(A)** and **(B)**, respectively. More than 99.9% of the predictions in **(B)** are accurate, with deviations at most shifted by a single bin. This result indicates that both modes produce reliable assessment of the maximum kinship level, and highlights the utility of SF-Relate in various workflows.

### 3.7 A case study: SF-Relate reduces false positives in genome-wide association studies

SF-Relate can improve the accuracy of downstream studies without requiring the sharing of sensitive information. Using the GWAS workflow, we demonstrate SF-Relate’s effectiveness in mitigating confounding from cryptic relatedness by enabling parties to detect and remove relatives from their joint dataset prior to conducting the GWAS. We simulate a multi-site GWAS using a subset of 100K samples from white British participants in UKB, distributed geographically among six parties (Chapter 4 and Table 3.8). 50% of these samples have at least one relative of 3rd-degree or less in the dataset. We then simulate 100 phenotypes and perform a linear regression-based GWAS with the top principal component as a covariate (Chapter 4). On average, using a nominal significance cut-off  $p$ -value of 0.05, SF-Relate removes 2.60% of the falsely identified loci (false positives) with





**Figure 3.8: SF-Relate reduces false positives in multi-site GWAS.** We vary the number of causal SNPs used in simulating the phenotypes, and compare four quality control (QC) approaches for excluding related individuals: (i) *Centralized removal* (non-private): all relatives are removed from the pooled dataset, (ii) *SF-Relate*: relatives are removed using our secure approach, (iii) *Local removal*: each party filters relatives from its local dataset independently, and (iv) *No removal*: no relatives removed. Initially, 50% of samples have relatives, and *Local removal* results in 18% of remaining samples still having a relative in the joint dataset. We plot the fraction of significant loci ( $p$ -value  $< 0.05$ ) on even numbered chromosomes in that are designed to be non-causal in the simulation (**A**), and the genome inflation factor  $\lambda_{GC}$  in (**B**). The filled boxes represent interquartile ranges of statistics across 100 simulated phenotypes. While local removal of relatives help reduce the confounding to some extent, SF-Relate significantly mitigates confounding, comparable to centrally-coordinated sample removal.

a drop in false positive rate (FPR) from 5.14% to 5.01%, compared to when relatives are not removed from the dataset (Figure 3.8). When parties independently remove local relatives, 18% of the remaining samples still have relatives in the joint dataset, and 2.03% of false positives are removed with an FPR drop from 5.14% to 5.04%. The one-sided Mann-Whitney U test  $p$ -value that SF-Relate produces lower FPRs across all phenotypes compared to the local removal of relatives is  $1.25 \times 10^{-5}$ . Thus, SF-Relate significantly mitigates confounding, producing an FPR near the nominal cutoff  $p$ -value of 0.05, comparable to centrally-coordinated sample removal (Figure 3.8).

Symbol	Definition	Default
$\phi$	KING kinship coefficient	—
$d$	relatedness degree	—
$\theta$	KING coefficient threshold	$2^{-d-1}$
cMLEN	subchromosome length	8cM
cMSTEP	overlap between subchromosomes	4cM
TARGETLEN	# SNPs after random projection	80
$w'$	size of projection window	—
$k$	# SNPs in k-SNPs	8
$n$	# local individuals	—
$m$	# SNPs	—
$\tau$	table ratio	1 ~ 128
$s$	subsampling rate (for SNPs)	0 ~ 1
$N$	# buckets (table size)	$\tau \cdot n$
$S$	bucket capacity	1
$\ell$	# LSH values in each index	4
$L$	# repetitions of hashing ( <i>Step 1: Hashing and bucketing</i> )	3
$M$	# (subsampled) SNPs	$s \cdot m$
$C$	bucket capacity used in micro-bucketing	1
$B$	Number of values in each ciphertext	8192

Table 3.1: **Symbols, parameters and default values.** The *Default* column indicates the optimal values for the parameters across all datasets in our experiments. For the encoding pipeline (see Chapter 4), we divide each chromosome into subchromosomes of length in equal genetic distances cMLEN, with adjacent segments overlapping in a fixed distance cMSTEP. Empirically, 3cM to 20cM are reasonable values, with 8cM being the best for both UK Biobank and All of Us datasets. After that, we randomly project each subchromosome vector down to a vector of a fixed length TARGETLEN. The projection randomly selects one variant out of every window of size  $w'$  SNPs, with probability proportional to their minor allele frequencies. Empirically TARGETLEN is set to 80 to ensure enough SNPs are chosen. To make each k-SNP correspond to roughly the same genetic distance,  $w'$  is chosen differently for each subchromosome, and it is computed as the ratio between the actual number of SNPs in the subchromosome and TARGETLEN. The parameter  $k$  in k-SNP can be between 5 and 30, and 8 is the best value for all datasets. The repetition parameter  $L$  specifies the number of times parties should repeat *Step 1: Hashing and bucketing* (Chapter 4). To ensure resulting merged table is highly utilized (> 99.9% non-dummies),  $L$  can be chosen on the fly independently by each party.

Dataset	Recall (%, counts)					Precision (%, counts)	% of comparisons w.r.t. all-pairwise
	Relatedness degree				Overall		
	0th	1st	2nd	3rd			
UKB-200K	100.0% 16/16	100.0% 4702/4702	99.8% 1709/1711	94.9% 8475/8925	97.0% 14902/15354	98.5% 14902/15129	0.13%
UKB-100K	100.0% 6/6	100.0% 1243/1243	100.0% 404/404	95.1% 2169/2279	97.2% 3822/3932	98.7% 3822/3872	0.26%
AoU-20K	100.0% 14/14	100.0% 209/209	100.0% 93/93	94.1% 145/154	98.0% 461/470	100.0% 461/461	1.28%

Table 3.2: **SF-Relate achieves near-perfect accuracy for identifying close relatives in UK Biobank and All of Us datasets.** Ground-truth relatedness degrees for recall and precision metrics are obtained using the KING method and assigning each sample to the lowest degree of relatedness observed. SF-Relate obtains accurate results while performing only a small fraction of comparisons compared to **all-pairwise** comparison between datasets.

Dataset	Recall (% , counts)				
	Relatedness degree				Overall
	0th	1st	2nd	3rd	
Geographically Distributed	100.0% 22/22	100.0% 3370/3370	100.0% 1483/1483	99.5% 10591/10640	99.7% 15515/15466
Single African Ancestry	100.0% 104/104	100.0% 1278/1278	100.0% 850/850	98.8% 530/542	99.6% 3262/3274
Locally-Phased UKB	100.0% 16/16	99.5% 3999/4017	98.2% 1598/1626	93.2% 3604/3863	96.7% 9217/9522
Locally-Phased AoU	100.0% 14/14	100.0% 209/209	97.8% 91/93	87.0% 134/154	95.3% 434/468

Table 3.3: **SF-Relate effectively detects relatives across diverse datasets.** We showcase the robustness of SF-Relate’s in two scenarios. (1) *Geographically Distributed*: a collaborative setting involving six centers collecting patient data within their respective geographic region (see Table 3.8). The dataset comprises 100K individuals of white British ethnicity sourced from the UK Biobank. (2) *Single African Ancestry*: a collaborative setting with two parties each holding 10K samples of African ancestry from All of Us. Additionally, we confirm that using locally-phased data does not affect SF-Relate’s ability to identify relatives. We consider two examples, each involving 20K patients split between two parties and sourced either from the UK Biobank (*Locally-Phased UKB*) or All of Us (*Locally-Phased AoU*). We test two prominent phasing tools that we use independently on each site, by applying SHAPEIT 5 (phase\_common) on UKB and Eagle v2.4.1 with the option `-maxMissingPerSnp 0.5` on AoU. Across all settings, SF-Relate’s recall remains consistently high, all above 87%. We observe a possible minor reduction in the third-degree recall on locally phased AoU, likely due to phasing errors associated with the small dataset size. We expect local phasing to be more accurate in a larger cohort or when large public reference panels are used.

Dataset	Race/ethnicity label	Relatedness degree	Recall (counts, %)	
AoU-20K	Black or African American	3	39/44	88.6%
	None Indicated	3	16/17	94.1%
	White	3	77/78	98.7%
UKB-200K	Any other mixed background	3	9/10	90.0%
	Any other white background	3	77/86	89.5%
	British	3	8033/8499	95.1%
	Caribbean	3	22/22	100.0%
	Indian	3	27/32	84.4%
	Irish	3	209/217	96.3%
	Other ethnic group	3	14/17	82.4%
	Pakistani	3	18/20	90.0%
	Prefer not to answer	3	28/29	96.6%

Table 3.4: **SF-Relate’s recall remains high across different subpopulations.** We apply SF-Relate to the multi-ethnicity datasets with 20K individuals from AoU and 200K individuals from UKB, each split between two parties. The recall for the 0th, 1st and 2nd degrees is nearly 100% for all subpopulations in both datasets, and thus excluded from the table. We separately calculate and report the third-degree recall for each subpopulation with at least 20 third-degree samples.

Dataset	SF-Relate								All-pairwise	
	Runtime				Communication				Runtime (est. total)	Comm. (est. total)
	Step 1	Step 2 (MHE)		Total	Step 1	Step 2 (MHE)		Total		
		Phase 1	Phase 2			Phase 1	Phase 2			
UKB-200K	1.8 m	14.0 h	0.5 h	14.5 h	—	46.6 TB	0.5 GB	46.6 TB	1.3 y	32.5 PB
UKB-100K	49.5 s	7.05 h	0.23 h	7.29 h	—	23.85 TB	241.7 MB	23.85 TB	112 d	9.8 PB
AoU-20K	18.6 s	5.65 h	0.11 h	5.79 h	—	6.2 TB	77.6 MB	6.2 TB	18.8 d	2.31 PB

Table 3.5: **SF-Relate scales efficiently to large datasets.** We report the runtime and communication costs for individual steps of SF-Relate described in Chapter 4. The runtime and communication costs for setting up the cryptographic keys are 40.4 s and 1.7 GB, respectively, constant across all experiments. We also show the estimated total costs of running **all-pairwise** comparisons and determining the closest relationship for each individual both using MHE.

cMLEN	5.0								
$k$	6.0			8.0			10.0		
$\ell$	3.0	4.0	5.0	3.0	4.0	5.0	3.0	4.0	5.0
Recall (%)	<b>84.2</b>	<b>96.6</b>	<b>97.5</b>	<b>96.1</b>	<b>97.5</b>	<b>98.3</b>	<b>97.6</b>	<b>98.0</b>	<b>98.3</b>
cMLEN	8.0								
$k$	6.0			8.0			10.0		
$\ell$	3.0	4.0	5.0	3.0	4.0	5.0	3.0	4.0	5.0
Recall (%)	<b>82.7</b>	<b>96.0</b>	<b>97.3</b>	<b>96.0</b>	<b>97.9</b>	<b>98.1</b>	<b>97.7</b>	<b>98.0</b>	<b>98.5</b>
cMLEN	11.0								
$k$	6.0			8.0			10.0		
$\ell$	3.0	4.0	5.0	3.0	4.0	5.0	3.0	4.0	5.0
Recall (%)	<b>77.3</b>	<b>95.7</b>	<b>97.2</b>	<b>96.0</b>	<b>97.3</b>	<b>97.5</b>	<b>97.4</b>	<b>97.5</b>	<b>97.9</b>

Table 3.6: **SF-Relate achieves high recall across various parameter settings.** On UKB-200K, we select a combination of reasonable parameters and compute the third-degree recall achieved by SF-Relate. The overlap between chromosomes cMSTEP is set to equal half of the subchromosome length cMLEN (see Table 3.1). We vary the values of  $k$ , the number of SNPs in  $k$ -SNPs and  $\ell$  the number of LSH values in each index, and maintain the remaining parameters at their default values provided in Table 3.1. SF-Relate consistently achieves near-perfect recall ( $> 95\%$ ), except on the lower end of the parameters, when there is not enough entropy among the subchromosomes for the LSH to separate samples into buckets.

Degree combinations →	<u>A</u>	<u>B</u>	<u>C</u>	D	<u>E</u>	F	G	H	I	J	<u>K</u>
SF-Relate	<b>1st</b>	<b>2nd</b>	<b>3rd</b>	3rd	<b>4th</b>	4th	U	U	<b>U</b>	U	<b>U</b>
KING	<b>1st</b>	<b>2nd</b>	<b>3rd</b>	3rd	<b>4th</b>	4th	3rd	4th	<b>4th</b>	U	<b>U</b>
RAFFI	<b>1st</b>	<b>2nd</b>	<b>3rd</b>	4th	<b>4th</b>	U	3rd	4th	<b>U</b>	U	<b>U</b>
PC-Relate	<b>1st</b>	<b>2nd</b>	<b>3rd</b>	3rd	<b>4th</b>	U	3rd	4th	<b>U</b>	4th	<b>U</b>
Number of individuals	<b>508</b>	<b>173</b>	<b>863</b>	48	<b>93</b>	23	22	23	<b>8410</b>	23	<b>9716</b>

Table 3.7: **SF-Relate outperforms KING and achieves comparable results to the more advanced methods, PC-Relate and RAFFI.** On a subset of 20K samples from UKB-200K, we compute the maximum relatedness degree for each individual using the different methods, and report the number of individuals falling into each classification type (when this count exceeds 20). The abbreviation U represents *Unrelated*. **Underlined columns** correspond to the set of individuals with identical classifications across all methods. **Bold-faced columns** denote substantial sets containing at least 100 individuals. Most methods have consistent results in the largest categories, except in column I. Whereas SF-Relate and KING agree on most columns, SF-Relate outperforms KING by aligning with the more accurate methods in column I, effectively excluding the thousands of likely-spurious 4th-degree pairs identified by KING.

UK Biobank Assessment Center	Geographic Area	Number of Samples
Edinburgh Glasgow	Scotland	7051
Middlesbrough Newcastle	Northeast England	14185
Liverpool Bury Stockport Manchester Leeds	Northwest England	25096
Birmingham Stoke Sheffield Nottingham Wrexham	Southeast England	23362
Barts Hounslow Croydon Reading Oxford	Central England	16836
Swansea Cardiff Bristol	Wales	13470
<b>Total Number of Samples</b>		100,000

Table 3.8: **Assignment of UK Biobank assessment centers to geographic areas.** To simulate a federated study, we use a subset of 100K white British individuals from UK Biobank dataset. We organize the 22 data collection centers (Data-Field 54) into six study groups according to their geographic locations within the UK.

# Chapter 4

## Algorithmic Details

### 4.1 Existing cross-dataset approaches and their limitations

Wang et al. [38] propose a homomorphic encryption method for identifying genetic relationships across parties, but their approach requires kinship computation for *all pairs* of samples, which does not scale to large datasets. Other previous approaches [39]–[42] rely on sharing a limited amount of processed data between parties to find related samples, which sacrifices both privacy and accuracy to some extent. For instance, Dervishi et al. [39] introduce a solution in which the parties reveal a subset of SNPs in a shuffled order for their respective samples to estimate the kinship coefficients. Robinson and Glusman [42] and Glusman et al. [40] propose to compare “fingerprints” obtained by applying a random projection to genomic samples to infer relatedness. He et al. [43] and Hormozdiari et al. [41] use error-correcting codes and fuzzy encryption to compare genotype vectors such that the comparison result can be decoded only if the two vectors are similar enough. These solutions require comparison between all pairs of samples between datasets, and the processing of genotype vectors into limited representations that can be shared leads to loss of precision.

In a recent competition organized by the iDASH Workshop 2023 [44], identifying the presence of relatives in encrypted datasets was posed as one of the challenge tasks. The challenge considered a setting that is different from our work. It involved a client-server scenario with small datasets (e.g., 2K individuals and 16K variants) and restricted any data preprocessing for evaluation purposes, thus limiting the applicability of the proposed solutions in a real use-case scenario. The best-

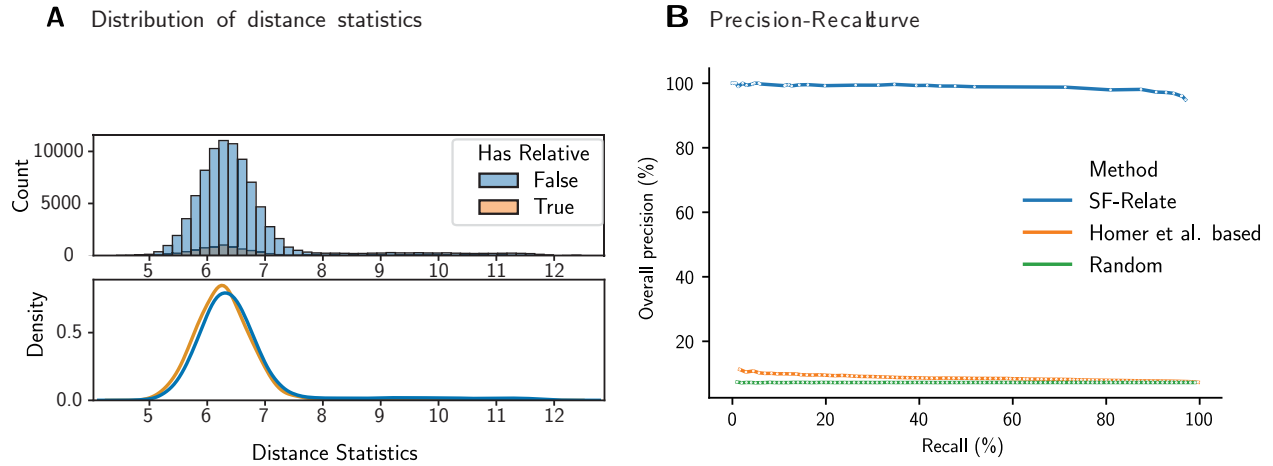
performing solutions to the challenge generally followed a similar approach adapting the work of Homer et al. [45]. In this approach, the presence of a relative in a dataset was inferred based on a statistic evaluating whether the individual’s genotype vector was closer to the allele frequencies in the dataset than to background frequencies computed in a reference population. We show in Figure 4.1 that computing this statistic securely does not lead to accurate results in realistic settings, involving complete genomes and tens of thousands or more individuals.

## 4.2 Our approach to secure cross-dataset kinship estimation

We develop SF-Relate to allow secure and efficient detection of relatives between large-scale datasets by drastically reducing the number of kinship computations, from quadratic to linear, while preserving accuracy. We present a graphical overview of SF-Relate’s workflow in Figure 4.2.

To achieve this solution, SF-Relate draws ideas from both *distance-based* and *IBD-segment-based* kinship estimation methods (Section 4.1). It first identifies pairs of samples to be compared between parties using a locality sensitive hash function (LSH; Section 2.2.2), which we adapt to ensure that both (or all) parties assign individuals with shared IBD segments to the same bucket with a higher probability than for unrelated individuals. SF-Relate sets the capacity of each bucket to one, discarding duplicate hits, which we refer to as a *micro-bucketing* strategy. As we demonstrate, this novel approach is key to minimizing the number of comparisons while maintaining accuracy. Next, SF-Relate securely estimates and thresholds the kinship coefficients between pairs of samples in buckets with the same index across parties then aggregates the results per sample using our secure implementation of a distance-based kinship estimator (KING-robust; Section 2.2.1). To perform these operations while keeping each party’s data confidential from other parties, we develop efficient two-party computation protocols based on multiparty homomorphic encryption (MHE; Section 2.2.3) techniques [33]. We additionally incorporate sketching techniques to further reduce the computational cost of MHE computations. In the following, we provide details of each step of our algorithm.





**Figure 4.1: Relative detection based on Homer et al.’s attack results in near-random performance.** Homer et al.’s attack [45] predicts whether an individual contributes to a genetic dataset by statistically testing whether the individual’s genotype count vector is closer to the allele frequencies in the dataset than some alternative reference frequencies. The same attack can potentially reveal the presence of a relative in a dataset due to shared genetic sequences. We evaluate the efficacy of this approach on the UKB-200K dataset split among two parties. In **(A)**, we compute (1) the  $L^1$  distance of every sample’s genotype count vectors to the mean genotype count vector from the local dataset (excluding relatives), representing the background statistic, and (2) the distance between the sample genotype count vector and the mean count vector in the other party’s dataset, representing the target statistic. We then subtract the two distances to see which dataset the sample is closer to. The figure shows that this estimator does not effectively separate samples that have a relative in the other dataset from samples that do not. In **(B)**, we show the precision-recall curves for various methods for detecting 3rd-degree or closer relatives in UKB-200K. For Homer et al., we vary the distance threshold used as a cut-off to determine whether a sample is closer to the target dataset. For SF-Relate, we plot the curve by varying the table size parameter ( $\tau$  in Chapter 4). Homer et al.’s approach obtains precision comparable to the level of random guessing, whereas SF-Relate achieves near-perfect precision.

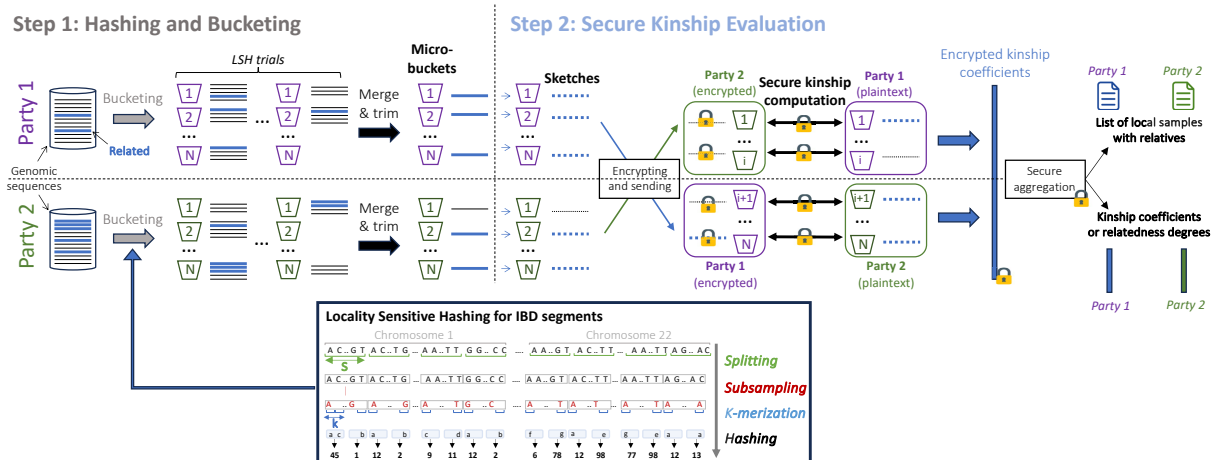


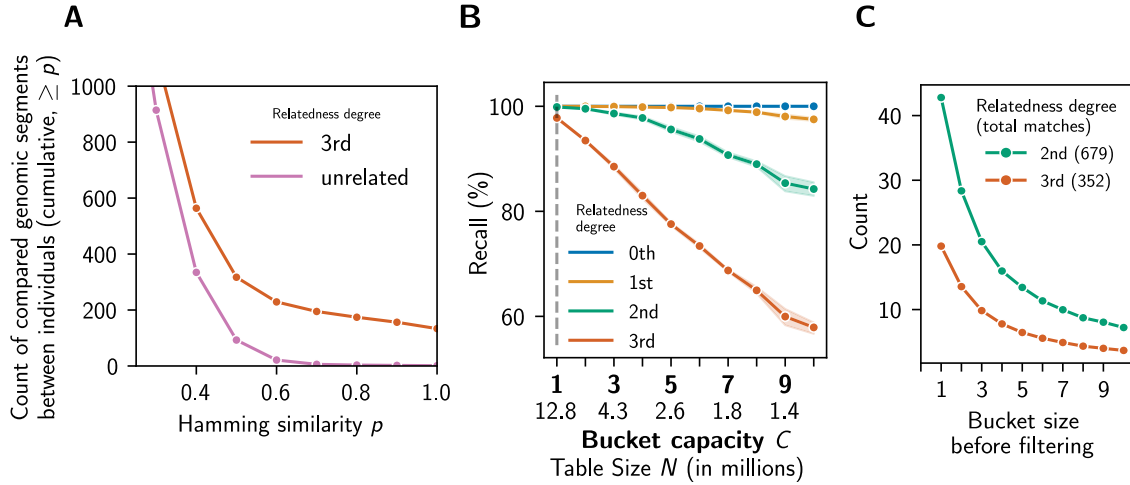
Figure 4.2: **SF-Relate's workflow**. In Step 1, the parties perform multiple trials of LSH to bucket samples before merging and filtering to obtain buckets of size 1 (micro-buckets). In Step 2, each sample is sketched and securely compared against the other party's sample in the same bucket to evaluate kinship (MHE-Phase 1). Finally, parties securely aggregate the results to obtain per-sample output (MHE-Phase 2).

## Step 1: Hashing and bucketing

In this step, each party locally evaluates a series of hash functions on each individual's haplotype sequences to assign the individual to buckets across a collection of hash tables, such that related individuals are more likely to be assigned to the same bucket index. Only individuals in the same bucket across parties are compared in a later step.

**Locality sensitive hashing to capture IBD segments.** SF-Relate assigns individuals to buckets using a locality sensitivity hash (LSH). In our setting, applying the Hamming LSH (or other LSH methods, such as MinHash [29]) directly to each sample, encoded as a genotype vector, for bucket assignment would not work in practice. This is because the difference between related and unrelated individuals, with respect to the distance between samples' encodings, is too small for LSH to distinguish; the average relative Hamming distance between third-degree relatives and unrelated individuals in UK Biobank are 22% and 23%, respectively. Hence, SF-Relate applies an encoding scheme that results in highly similar Hamming vectors for related samples. This encoding captures the biological signal of IBD distributions, and can be seen as a variant of the encoding in [22]. It applies LSH on split chromosomes, exploiting the key insight that IBD segments are unevenly distributed on the genome. The subchromosomes are further divided into short string of genotypes (similar to k-mers), to extract long runs of identical genotypes that are unlikely by chance.

The first three steps of our hashing approach encode IBD segments: (1) *Splitting*: Haplotypes are divided into genomic segments of fixed-length in genetic distance (centi-Morgans). (2) *Sub-sampling*: Each segment is randomly projected down to a fixed number of SNPs, where the SNPs are sampled with probability proportional to their minor-allele frequency (MAF). This reduces the impact of genotyping errors and rare variants on hashing and unifies the SNP density across different segments and datasets, inspired by Naseri et al. [21]. (3) *K-merization*: Subsampled genotypes for several contiguous segments are concatenated to form a k-SNP (akin to  $k$ -shingles in [22]), i.e., a sequence of genotypes for  $k$  SNPs, which helps to identify matches of long genomic segments. Through these steps, we encode each sample as a list of subchromosome k-SNP vectors. Indeed, we confirm that related pairs share significantly more Hamming-similar subchromosomes under this encoding scheme (Figure 4.3).



**Figure 4.3: SF-Relate’s hashing and micro-bucketing strategies effectively assign close relatives to the same bucket.** (A) Hamming LSH with SF-Relate’s k-SNP encoding scheme enables separation of pairs of genomic segments with high Hamming similarity (likely IBD) between close relatives from those pairs between unrelated individuals. (B) Setting the bucket capacity  $C = 1$  achieves the highest recall in relative detection compared to larger values of  $C$ . For comparison, we adjust the table size  $N$  accordingly to keep the number of comparisons  $NC^2 = 1.28$  million constant. The recall for each relatedness degree is the fraction of relative pairs of that degree that are assigned to the same bucket. (C) Close relatives (of 2nd and 3rd degrees) who are assigned to the same bucket are most often found in size-1 buckets before filtering. All results are based on UKB-200K.

The final step of hashing applies LSH on the subchromosome vectors to obtain the actual bucket index. To utilize the raw probability gap between non-IBD segments and IBD-segments produced by LSH (such as the gap between 0.5 and 1.0 in Figure 4.3) we need to amplify it. For this, we choose a concatenation parameter  $\ell$  and define the bucket index as the FNV-1 hash [46] on the

concatenation of the outputs of  $\ell$  independent Hamming LSH applied to the subchromosome vector. This in turn boosts the gap by raising it to the power of  $\ell$ . An alternative (as in [22]) would be to apply the LSH function MinHash [29], which captures the set-based Jaccard similarity. Nevertheless, Hamming similarity is more natural for IBD detection, as IBD segments are by definition matching k-SNPs at the same genetic positions, while the Jaccard similarity discards positional information of the two sets of k-SNPs. Indeed, Hamming similarity detects more highly similar subchromosomes (Figure 4.4). The final output of this procedure is a list of hash tables, each consisting of buckets storing sample IDs.




Figure 4.4: **The LSH for Hamming similarity retains more high-probability pairs (candidate IBD segments) than the LSH for Jaccard similarity (MinHash).** On the UKB-200K dataset, we count the number of subchromosome pairs between individuals based on the probability that the pairs would produce the same LSH index with an LSH amplification with  $\ell = 4$ . A higher matching probability indicates higher similarity between segments and thus more likely to be in IBD. The raw similarity score  $0 \leq p \leq 1$  between vectors is transformed to the amplified LSH match probability  $p^\ell$  by the amplification described in *Step 1: Hashing and bucketing*. Each curve shows the count of segments with probability higher than  $p^\ell$  averaged over pairs of related samples in the respective relatedness degree classes. The plot shows that Hamming LSH identifies more segments with high matching probability after the LSH amplification, suggesting its ability to detect more IBD segments.

**Micro-bucketing strategy.** To prevent leakage of private information, the parties need to compare the samples in corresponding buckets without exposing any additional information about their datasets, such as the distribution of non-empty buckets and their sizes. This requires that the buckets created in the previous step be padded to a fixed size by adding dummy samples. However, to keep the sample identities hidden, all pairs of samples in a bucket between the parties need to be compared, which leads to both quadratic scaling of comparisons with the bucket size and a large

amount of wasted computation involving dummy samples. To address this issue, SF-Relate introduces a *micro-bucketing* strategy, which merges buckets across multiple hash tables (produced by different subchromosomes) and then filters every bucket down to a *single* element. Dummy samples are added only at the end to pad empty bucket to size one. This effectively transforms the parties' local bucket assignments into an ordered list of samples to be securely compared against the corresponding list obtained by the other party in an element-wise fashion. This approach avoids the quadratic scaling while minimizing the addition of dummy samples due to the merging of buckets (i.e., a bucket is filled if at least one sample is assigned to it in one of the hash tables). Despite the extreme level of filtering applied to each bucket during this process, our strategy enables accurate detection of relatives with remarkable efficiency.

SF-Relate chooses a **table size parameter**  $N$ , and a **bucket size parameter**  $C$ . It aligns the different hash tables and merge buckets with the same remainder modulo  $N$  into one. This ensures the number of buckets is at most  $N$ . In practice, on a dataset with  $n$  local samples,  $N$  is determined by first choosing a **table ratio**  $\tau$ , and letting  $N = \tau n$ . (We provide a table of the main symbols and their default values in Table 3.1.)

After this merge, buckets with more than  $C$  samples are filtered until  $C$  samples remain. In this filtering step, samples from buckets built by smaller subchromosome index are given higher preference, but otherwise a uniformly random filtering is performed. The preference towards smaller subchromosome indices is to ensure samples in corresponding buckets in hash tables more likely originates from the same hash table. The parties then repeat the entire hashing step locally  $L$  times (each with new randomness) until 99% of the buckets are full. Only at the end, dummy samples are inserted to ensure a constant bucket size.

On the realistic dataset UKB-200K, we evaluate the best bucket capacity  $C$ , by keeping the number of comparisons  $NC^2$  fixed and checking the fraction of related pairs that appear in at least one corresponding bucket after micro-bucketing. Surprisingly, the minimal capacity  $C = 1$  gives the highest recall. This is because related samples share many rare IBD segments (many of which is unique) that cause them to end up in small-sized bucket (Figure 4.3). When this happens, both samples remain in the corresponding bucket with high probability (as there is no competition within that bucket). The average fraction of shared unique IBD segments is not high for 3rd degree (20/352), but given that our goal is to use the list of buckets to trigger subsequent kinship com-

putations (which operates on the entire genome), and not identifying all IBD segments between pairs (unlike [21], [22]), as long as one of the many small-sized buckets due to rare IBD segments survive micro-bucketing, the pair of relatives would be discovered. This explains the effectiveness of micro-bucketing when restricted micro capacities  $C$ . In fact, for the rest of the thesis, we always set  $C = 1$ .

In sum, each party in SF-Relate obtains a single hash table with  $N$  buckets (each with  $C = 1$  samples), and micro-bucketing ensures the hash tables are highly utilized. The parties hence perform  $C^2N = N$  secure kinship computation in *Step 2: Secure kinship evaluation*.

## Step 2: Secure kinship evaluation

In this step, the parties perform element-wise comparisons between their ordered list of samples (representing elements in a merged hash table with size-one buckets), obtained in Step 1. At a high level, they first jointly evaluate the kinship coefficient for each pair (**MHE-Phase 1**), before aggregating the results to obtain an indicator for each individual reflecting the presence or absence of a close relative in the other dataset (**MHE-Phase 2**).

Our efficient two-party protocols for secure kinship evaluation (MHE-Phase 1 and 2) apply multi-party homomorphic encryption (MHE) to computations over the bucketed samples from Step 1. We exploit key properties of MHE to minimize the cryptographic overhead of our protocols by maximizing the use of locally available plaintext data and balancing the workload between the parties. Although we focus on the two-party setting, our protocols naturally extend to settings with more than two parties, since we can execute the protocols between all pairs of parties and then aggregate the results.

**MHE-Phase 1: Secure computation of kinship coefficients.** Given a list of  $N$  samples on both sides, where each sample is associated with a vector of  $M$  SNPs, the parties collaborate to calculate kinship coefficients for the  $N$  pairs of samples between them and compare each to a threshold. The desired comparison test given a threshold  $\theta$  is  $\phi = \frac{1}{2} - \frac{1}{4} \cdot \frac{\|\mathbf{x}-\mathbf{y}\|^2}{\min(h_x, h_y)} \geq \theta$ . For efficient evaluation under encryption, we can rewrite it as  $(2-4\theta) \min(h_x, h_y) - \|\mathbf{x}-\mathbf{y}\|^2 \geq 0$ , thus avoiding the division operation. The comparison test passes when both  $h_x$  and  $h_y$  satisfies it, so we compute

$$\text{SIGN} \left( (2-4\theta)h_x - (\|\mathbf{x}\|^2 - 2\langle \mathbf{x}, \mathbf{y} \rangle + \|\mathbf{y}\|^2) \right) \cdot \text{SIGN} \left( (2-4\theta)h_y - (\|\mathbf{x}\|^2 - 2\langle \mathbf{x}, \mathbf{y} \rangle + \|\mathbf{y}\|^2) \right),$$

which evaluates to 1 if the coefficient is above the threshold and 0 otherwise. Boxed terms represent encrypted data, and  $\text{SIGN}(\boxed{v})$  is the indicator function for  $v \geq 0$ . Note that we assign the evaluation of this expression to the party that holds  $x$  (Party 1) and have the other party (Party 2) transfer the encrypted  $y$  for computation. This allows most operations to be performed efficiently using the plaintext  $x$ .  $h_y$  (the number of heterozygous in the genotype) and  $\|y\|^2$  can be computed locally by Party 2 before encryption. Hence, the most expensive operation is the inner product  $\langle x, \boxed{y} \rangle$ , which requires a plaintext-ciphertext multiplication between vectors of size  $M$  and a summation of elements of the resulting vector. We use a polynomial approximation of the sign function to homomorphically evaluate  $\text{SIGN}(\cdot)$ . The cost of this function, despite requiring homomorphic evaluation of a high-degree polynomial, is dwarfed by the cost of computing the inner products. Finally, in addition to the SIMD property of MHE operations, we process batches of coefficients in parallel and evenly distribute the workload between parties by alternating their roles across batches with respect to who holds the plaintext vector  $x$ .

**MHE-Phase 2: Secure aggregation of results for individual samples.** Next, the parties aggregate the comparison results for each individual to compute a single binary indicator for the presence or absence of a relative. For this, they first perform a linear scan over the comparison results, which selects the results corresponding to the same individual and masks the rest. The selected results are then accumulated, after which another sign test is performed to obtain a binary value as desired, hiding the number of identified relationships as a result. At the end of the protocol, the parties decrypt the vectors and each obtain a list of indicators, determining whether each sample has at least one close relative in the other dataset; this is the only information that is revealed to each party. Note that the complexity of this step depends only on the number of comparisons and individuals, while MHE-Phase 1 also scales with the number of SNPs, which is typically large (e.g., >500K).

**Accelerating kinship computation using sketching.** To further reduce the cost of secure kinship evaluation, SF-Relate first reduces the size of each sample through *sketching*. In particular, given the **subsampling ratio** parameter  $0 < s \leq 1$ , SF-Relate randomly chooses an  $s$  fraction of SNPs to use for kinship evaluation. This provides a natural approximation for the KING kinship estimator, which includes squared Euclidean distance between two genotype vectors that can be estimated using a random subset of coordinates in an unbiased manner (see Figure 4.5). Our results show

that this approach enables a meaningful trade-off between accuracy and efficiency; a minor loss in precision introduced by sketching allows us to obtain a substantial reduction in computational cost while maintaining near-perfect detection accuracy.

**Alternative output modes.** By default, SF-Relate computes a list of indicators representing whether each sample has at least one close relative in the other datasets. SF-Relate also supports securely computing other types of output, including the closest relatedness degree for each individual, the maximum kinship for each individual (discretized), and the full list of computed kinship coefficients.

For outputting all kinship coefficients, we replace the final sign test computations with the computation of the kinship coefficients in **MHE-Phase 1**. That is, the parties homomorphically compute  $\phi = \frac{1}{2} - \frac{1}{4} \cdot \frac{\|x-y\|^2}{\min(h_x, h_y)}$ . To speed up the computation, they precompute the values  $h_x^{-1}$  and  $h_y^{-1}$  locally in plaintext, which allows replacing the division with a multiplication by  $\max(h_x^{-1}, h_y^{-1})$ , which can be more efficiently computed.

For both the closest relatedness degree and the maximum kinship coefficient, SF-Relate computes multiple comparison tests with respect to a series of kinship thresholds, then executes **MHE-Phase 2** in parallel to accumulate the resulting comparison results. The decryption of these results reveals the largest threshold at which the comparison tests succeed. Based on this information, the parties can determine the closest degree or the maximum kinship as desired.

Note that, in any of these settings, the final results can also be kept in encrypted form and utilized in subsequent analysis steps without revealing the results of relative detection. Additionally, the complexity of **MHE-Phase 1** is constant across all output modes with  $N$  plaintext-ciphertext multiplications between vectors of length  $M$ , whereas the complexity of **MHE-Phase 2** increases linearly with the number of thresholds  $t$ . The default setting of SF-Relate corresponds to  $t = 1$ , while the version that reveals all coefficients corresponds to  $t = 0$ .

## Detailed MHE protocols for secure kinship evaluation

We detail here SF-Relate protocols for computing and detecting kinship coefficients that are above a predefined threshold (MHE-Phase 1) and for aggregating these results per individual (MHE-Phase 2).



We denote a matrix  $A$  with  $N$  rows and  $M$  columns as  $A^{N \times M}$  and a vector  $x$  with  $B$  elements as  $r^{B \times 1}$ . Under the 0-based indexing, i.e. columns and rows are numbered starting from 0.  $x[a : b]$  denotes the subvector from row  $a$  to  $b$  (including  $a$  but excluding  $b$ ). We omit  $a$  when it is 0, and omit  $b$  if it equals  $N - 1$ . Similarly, for matrices, we use  $A[a : b, c : d]$  to denote the submatrix specified by the ranges. For a vector  $x^{N \times 1} = (x_0, \dots, x_{N-1})$ , we use  $x^2$  to denote the vector  $(x_0^2, \dots, x_{N-1}^2)$ , i.e. the one where elements are squared. The notation  $\text{Sign}(E)$  denotes the indicator of the event  $E$ . When applied to a vector of events as in  $\text{Sign}(x^{n \times 1} == 1)$ , it corresponds to the vector of indicators. In other words,  $\text{Sign}(x == 1)$  is equivalent to  $(\text{Sign}(x[i] == 1), \dots, \text{Sign}(x[n - 1] == 1))$ . Finally,  $\mathbf{0}^{B \times 1}$  and  $\mathbf{1}^{B \times 1}$  denotes the vector  $(0, 0, \dots, 0)^{B \times 1}$  and  $(1, 1, \dots, 1)^{B \times 1}$ , respectively.

Every ciphertext under the CKKS encryption [32] encrypts a vector with length  $B$ , the CKKS block length, which is typically a power of 2 like 8192. We denote the encryption of a vector  $x$  as  $\boxed{x}$ . We simplify the function interfaces of the CKKS implementation in Lattigo [36] as follows. Cryptographic keys are omitted in the function calls.

- $\text{ENC}(x)$  takes in a plaintext vector  $x$  and returns an encrypted ciphertext  $\boxed{x}$ .
- $\text{COLLABORATIVEDECRYPT}(\boxed{x})$  takes in an encrypted vector  $\boxed{x}$  and returns the plaintext result  $x$ . Note that all parties would need to collaborate for this operation.
- Homomorphic Single-Instruction Multiple-Data (SIMD) operations, including coordinate-wise addition, subtraction and multiplication between any combination of ciphertexts encrypting vectors and plaintext vectors are denoted by  $+$ ,  $-$ ,  $\cdot$ , respectively.
- The  $\text{ROTATE}(\boxed{m}, i)$  operation that receives as input  $\mathbf{m}^{B \times 1} = (m_0, \dots, m_{B-1})$  outputs a ciphertext encrypting  $(m_i, m_{i+1}, \dots, m_B, m_0, m_1, \dots, m_{i-1})$ .

We also implement the following helper functions:

- $\text{SIGN}(\boxed{v})$  makes use of the Chebyshev polynomial interpolation and Newton's method to approximately compute the sign function, using combinations of the homomorphic SIMD operations  $+$ ,  $-$  and  $\cdot$ . It returns a ciphertext  $\boxed{r}$  such that if one decrypts it, the result equals the Boolean-valued indicator for  $\text{Sign}(v)$ .
- $\text{EXTRACT}(\boxed{r^{B \times 1}}, i)$  returns a ciphertext encrypting the vector  $(0, \dots, 0, r[i], 0, \dots, 0)^{B \times 1}$ . We implement this by homomorphically multiplying the  $i$ -th basis vector with  $\boxed{r^{B \times 1}}$ .

- $\text{INNERSUM}(\boxed{\mathbf{m}^{B \times 1}})$  transforms the vector  $\mathbf{m} = (\mathbf{m}[0], \dots, \mathbf{m}[B-1])$  into the vector  $\mathbf{s} = (s, s, \dots, s)^{B \times 1}$  where  $s := \sum_{i=0}^{B-1} \mathbf{m}[i]$ . Adding a length- $B$  vector with its powers-of-2 rotated copy (see  $\text{ROTATE}(\boxed{\mathbf{m}}, i)$  above), for all the powers of 2 at most  $B$ , namely  $2^1, 2^2, \dots, B$  achieves this. Hence, we implement it efficiently by iterating over  $0, \dots, \log B$  rotating  $\boxed{\mathbf{m}}$  accordingly, and then homomorphically add the results together.

We display here the two protocols used in Chapter 4, namely MHE-Phase 1 and MHE-Phase 2, for secure kinship evaluation between two parties.

---

### MHE-Phase 1 Distributed relative detection

---

**Input:** Party 1 and Party 2 have the genotype counts matrices  $\mathbf{A}^{N \times M}, \mathbf{D}^{N \times M} \in \{0, 1, 2\}^{N \times M}$ , respectively.  $\theta$  denotes the kinship detection threshold, and  $B$  denotes the CKKS block length. The rows in the matrices are organized in the order specified by SF-Relate's hash tables.

**Output:** A block-encrypted vector of the  $N$  detection results,  $\boxed{\mathbf{r}_1^{B \times 1}}, \dots, \boxed{\mathbf{r}_{\lfloor N/B \rfloor}^{B \times 1}}$ .

```

1: for  $b = 0, \dots, \lceil (\lfloor N/B \rfloor - 1)/2 \rceil$  do ▷ Remaining blocks computed in parallel by switching roles
2:   Party 1:  $\mathbf{X}^{B \times M} \leftarrow \mathbf{A}[bB : (b+1)B, :]$ 
3:   Party 2:  $\mathbf{Y}^{B \times M} \leftarrow \mathbf{D}[bB : (b+1)B, :]$ 
4:   Party 1:  $\mathbf{x}_{sq}^{B \times 1} \leftarrow \sum_{i=1}^M (\mathbf{X}[:, i])^2$ , and  $\mathbf{h}_x^{B \times 1} = \sum_{i=1}^M \text{Sign}(\mathbf{X}[:, i] == 1)$  ▷  $\text{Sign}(v)$  is the indicator of  $v \geq 0$ 
5:   Party 2:  $\mathbf{y}_{sq}^{B \times 1} \leftarrow \sum_{i=1}^M (\mathbf{Y}[:, i])^2$  and  $\mathbf{h}_y^{B \times 1} = \sum_{i=1}^M \text{Sign}(\mathbf{Y}[:, i] == 1)$ 
6:   Party 2: send  $\boxed{\mathbf{y}_{sq}} \leftarrow \text{Enc}(\mathbf{y}_{sq})$  and  $\boxed{\mathbf{h}_y} \leftarrow \text{Enc}(\mathbf{h}_y)$  to party 1.
7:   Party 1:  $\boxed{\mathbf{p}^{B \times 1}} \leftarrow \text{Enc}(\mathbf{0}^{B \times 1})$ . ▷ Initiate an encrypted vector of  $B$  zeros
8:   for  $j = 0, \dots, M-1$  do
9:     Party 2: send  $\boxed{\mathbf{y}_j^{B \times 1}} \leftarrow \text{Enc}(\mathbf{Y}[:, j])$  to party 1
10:    Party 1:  $\mathbf{x}_j^{B \times 1} \leftarrow (\mathbf{X}[:, j])$ , and  $\boxed{\mathbf{g}} \leftarrow \boxed{\mathbf{p}} + 2\mathbf{x}_j \cdot \boxed{\mathbf{y}_j}$  ▷ SIMD coordinate-wise multiplication
11:    Party 1:  $\boxed{\mathbf{t}_1^{B \times 1}} \leftarrow \text{SIGN}((2 - 4\theta)\mathbf{h}_x - (\mathbf{x}_{sq} + \boxed{\mathbf{y}_{sq}} - \boxed{\mathbf{g}}))$  ▷  $\text{SIGN}(\boxed{v})$  is the (encrypted) indicator  $\boxed{v \geq 0}$ 
12:    Party 1:  $\boxed{\mathbf{t}_2^{B \times 1}} \leftarrow \text{SIGN}((2 - 4\theta)\boxed{\mathbf{h}_y} - (\mathbf{x}_{sq} + \boxed{\mathbf{y}_{sq}} - \boxed{\mathbf{g}}))$ 
13:    Party 1: Save the batch  $\boxed{\mathbf{r}_b} \leftarrow \boxed{\mathbf{t}_1} \cdot \boxed{\mathbf{t}_2}$  and share with party 2. ▷ Results are shared to execute MHE-Phase 2

```

---

## 4.3 Dataset preprocessing

We utilize three datasets sampled from prominent genomic data consortia: UK Biobank (UKB) and All of Us (AoU). Specifically, we extract two datasets from UKB, one comprising 100K samples (UKB-100K) and the other comprising 200K samples (UKB-200K). In both cases, we randomly split the datasets among two sites. From AoU, we extract a dataset of 20K samples. Due to the smaller size of this dataset and to avoid a highly imbalanced distribution, we first split the individuals with




Figure 4.5: **KING can be accurately estimated on subsampled subsets of SNPs.** A subsampling ratio at  $s = .7$  achieves reliable relatedness degree classification, where accurate kinship is not necessary. We apply SF-Relate on UKB-200K and compute the kinship coefficients on the list of micro-bucket pairs, under subsets of SNPs with different subsampling ratios. The  $x$  axis shows the kinship on the full set of SNPs. Points near the degree boundaries have a higher chance of classification error. The scatter plots show that sketching introduces noise to the kinship, but reliable relatedness degree classification is possible when points are not too close to the thresholds.

## MHE-Phase 2 Accumulation to per-sample output

**Input:** Party 1 has a list of block-encrypted boolean values  $\mathbf{R}^{N \times 1} = (\mathbf{r}_0^{B \times 1}, \dots, \mathbf{r}_{\lfloor N/B \rfloor}^{B \times 1})$ , each storing  $B$  indicators of positive detection for close relations (i.e. the corresponding KING coefficient passes the threshold  $\theta$ ), and a list  $D = (D_1, \dots, D_N)$  storing the local IDs party 1 used in the comparisons.

**Output:** A list of (decrypted) boolean values  $b_1, \dots, b_n$ , one per local ID, signifying whether each ID appears in at least one high-kinship comparison.

```

1: for  $i = 0, \dots, \lfloor n/B \rfloor - 1$  do
2:    $\mathbf{o}^{B \times 1} \leftarrow \text{ENC}(\mathbf{0}^{B \times 1})$  ▷  $\mathbf{0}^{B \times 1}$  means the vector of  $B$  zeros
3:   for  $\text{id} = iB + 1, \dots, \min((i + 1)B - 1, n)$  do ▷ Iterate by blocks of size  $B$ 
4:      $\mathbf{e}^{B \times 1} \leftarrow \text{ENC}(\mathbf{0}^{B \times 1})$ 
5:      $\text{locs}_{\text{id}} \leftarrow \text{list of locations } j \text{ in } 1, \dots, N \text{ where } D_j = \text{id}$ 
6:     for  $j \in \text{locs}_{\text{id}}$  do
7:        $\mathbf{u}^{B \times 1} \leftarrow \text{EXTRACT}(\mathbf{r}_{\lfloor j/B \rfloor}, j \bmod B)$  ▷ Extract the boolean result through multiplication with a one-hot encoded plaintext vector
8:        $\mathbf{e} \leftarrow \mathbf{e} + \mathbf{u}$ 
9:        $\mathbf{e} \leftarrow \text{INNERSUM}(\mathbf{e})$  ▷  $\text{INNERSUM}(\mathbf{m})$  transforms  $(\mathbf{m}_1, \dots, \mathbf{m}_B)$  into  $(\sum_{i=1}^B \mathbf{m}_i) \cdot \mathbf{1}^{B \times 1}$ 
10:     $\mathbf{o} \leftarrow \mathbf{o} + \text{EXTRACT}(\mathbf{e}, \text{id} \bmod B)$ 
11:     $\mathbf{f} \leftarrow \text{SIGN}(\mathbf{o} - 0.5 \cdot \mathbf{1}^{B \times 1})$  ▷ Compute whether the result is (much) larger than 0
12:     $b_{i*B}, \dots, b_{\min((i+1)*B, n)} \leftarrow \text{COLLABORATIVEDECRYPT}(\mathbf{f})$  ▷ Decrypt ciphertext collaboratively

```

close relations on the two sites, and then randomly split the set of unrelated individuals across the two sites. On the UKB dataset, we use phased autosomal haplotypes officially released in UK Biobank v3 as input to the hashing (Chapter 4), while for All of Us, we phase a batch of 20K samples from AoU using Eagle 2 [47]. Independently phasing the data at each site does not affect the accuracy of SF-Relate (Table 3.3).

## 4.4 Ground truth kinship preparation

To compute the ground truth of related individuals in our datasets, we follow the approach proposed in UKB documentation [12]. We first filter the SNPs based on their implications in population structure, before computing the kinship using the KING approach [17]. To determine the set of SNPs to retain, we conduct a PCA on a publicly-available dataset (i.e., 1000 Genomes) using the intersection of loci with our dataset. Utilizing a reference dataset ensures that our method is not tailored to the processed dataset and effectively generalizes to other datasets. We then exclude SNPs that exhibit high PC loadings in the top three PCs, using a threshold set at the 75th percentile of these loadings. This strategy enables us to filter out SNPs with heavy loadings while retaining sufficient ancestry-agnostic autosomal SNPs for kinship inference.

Applying this approach to the UKB datasets results in selecting 90K SNPs, upon which the KING estimator predominantly identifies the same related pairs as those in UK Biobank’s relatedness release. The ground-truth relatedness degrees in our experiments are based on these KING coefficients, utilizing the recommended thresholds  $2^{-d-1.5}$  for degree  $d$  [17].

### 4.4.1 Kinship estimation using alternative non-secure methods

Even though SF-Relate builds upon the KING estimator [17], it outperforms it due to its novel approach in pre-selecting individuals likely to be related, utilizing an encoding and hashing scheme specialized to capture IBD signals. To showcase this, we compare SF-Relate and standard KING estimator alongside two advanced relative detection tools: PC-Relate [19] and RAFFI [37]. For PC-Relate, we rely on the hail implementation [28] and consider only bi-allelic variants from the UK Biobank SNP panel to compute all pairwise coefficients. We set the minor allele frequency and the number of PCs to 0.1 and 10, respectively, and remove variants with missing rate higher than 5%.

Additionally, we perform ld-pruning on the SNP set, reducing from 600K to 300K variants using parameters  $r^2 = 0.05$  and `bp_window_size = 500000`. For kinship estimation with RAFFI, we first run the IBD-finding tool RAPID with the parameters `-r3 -s1 -d5 -w3`, followed by executing RAFFI (v0.1) for kinship estimation. However, we observe that the initial segment of length 10 cM on chromosome 15 is only covered by 16 base pairs in the UKB dataset, resulting in excessive candidate IBD segment pairs processed by RAPID. As the 10 cM sharing negligibly affects the overall kinship coefficient, we remove these base pairs when running RAPID on UKB.

## 4.5 Phenotype simulation for the GWAS case study

To evaluate the effectiveness of SF-Relate in mitigating the confounding effects of cryptic relatedness for GWAS, we simulate a GWAS study on a subset of 100K samples from UK Biobank using simulated phenotypes, following the methodology proposed in REGENIE [48] as described next. We select a set of random variants from odd-chromosomes to serve as causal variants, reserving the even-chromosomes to assess the level of false positive associations. We exclude extremely rare SNPs (minor allele count  $< 5$ ), before randomly selecting  $P$  SNPs located in odd-numbered chromosomes, for  $P \in \{100, 1000, 10000\}$ . These selected SNPs are designated as causal. For each causal SNP, we sample its effect size  $\beta_j$  with the constraint that the total variance (i.e., narrow-sense heritability) is  $h^2 = 0.2$ . We then use a linear model with top principal component correction to simulate the phenotype  $Y_i$  for each individual  $i$  as

$$Y_i = \sum_{j=1}^P G_{i,j} \beta_j + A_i + \epsilon_i,$$

where  $G_{i,j}$  is the standardized genotype of individual  $i$  at SNP  $j$ ,  $A_i$  is the first PC score of individual  $i$  (scaled to have variance 0.05), and  $\epsilon_i$  is a Gaussian noise variable with variance 0.75, representing environmental effects.



# Chapter 5

## Discussion

We present SF-Relate, a secure federated algorithm for identifying close relatives between isolated genomic datasets. Using a novel strategy for hashing and bucketing individuals to capture shared IBD segments between relatives, SF-Relate achieves near-perfect detection while maintaining a practical runtime (i.e., less than a day) even on a large dataset including 200K individuals. To the best of our knowledge, our work is the first to demonstrate secure relative detection at scale while also ensuring a strong, formal notion of privacy for each input dataset. We expect SF-Relate to be a useful tool for the growing networks of collaborating institutions, which currently lack the tools to jointly perform a variety of genetic analyses without sharing data. To facilitate the use of SF-Relate, we provide automated deployment workflows on the `sokit` web server [49], which streamlines the collaborative execution of a range of secure federated tools such as SF-Relate.

There are several directions that we would like to pursue in future research. First, while SF-Relate identifies relatives based on the standard KING-robust estimator [17], there are other approaches that may provide more robust estimation, especially for more distant relatives beyond the third degree, in terms of both correcting for population structure (e.g., PC-Relate [19]) and detecting IBD segments to allow a more direct calculation of the proportion of IBD sharing (e.g., RAFFI [37]). Although we have demonstrated that our method can often mirror the behavior of these advanced methods (Figures 3.4 and 3.5 and Table 3.7), directly implementing these approaches may be more effective for identifying distant relatives. Integrating our approach with a recently proposed secure federated algorithm for principal component analysis [50], may help to address the former. For the latter, we posit that an extension of our hashing strategy to quantify the rate of collision, which

represents the sharing of a short IBD segment between individuals, may be possible.

Extending SF-Relate to accommodate a broader range of scenarios represents another key direction of future work. For scenarios involving more than two parties, developing a more efficient strategy than the straightforward all-pairwise execution of SF-Relate would be beneficial. Additionally, enabling the detection of relatives for a single query individual within a large (distributed) database would be useful for services that help individuals find lost biological relatives (e.g., MyHeritage [51]). There is also a need to facilitate similarity computations for other data types, including medical records. In any of these scenarios, it would be meaningful to further explore potential information leakage in the output and devise strategies to mitigate any remaining risk. Overall, our work offers technical insights that are broadly applicable to discovering relations across siloed datasets.



# Data and Source Code Access

Data access application to genotypes and haplotypes from the UK Biobank can be submitted at the [UK Biobank website](#). The All of Us (AoU) Controlled Tier Dataset v5 is available through the Controlled Tier of the AoU Researcher Workbench. The application to access the AoU dataset can be submitted at the [All of Us Website](#).

Our open-source software, SF-Relate, along with a demo using a public dataset, are available at [Github](#). Additionally, SF-Relate can be conveniently executed through sfkit, a web server for secure collaborative genomic studies, accessible at the [Sfkit website](#).

## Data acknowledgments

The All of Us Research Program is supported by the National Institutes of Health, Office of the Director: Regional Medical Centers: 1 OT2 OD026549; 1 OT2 OD026554; 1 OT2 OD026557; 1 OT2 OD026556; 1 OT2 OD026550; 1 OT2 OD 026552; 1 OT2 OD026553; 1 OT2 OD026548; 1 OT2 OD026551; 1 OT2 OD026555; IAA #: AOD 16037; Federally Qualified Health Centers: HHSN 263201600085U; Data and Research Center: 5 U2C OD023196; Biobank: 1 U24 OD023121; The Participant Center: U24 OD023176; Participant Technology Systems Center: 1 U24 OD023163; Communications and Engagement: 3 OT2 OD023205; 3 OT2 OD023206; and Community Partners: 1 OT2 OD025277; 3 OT2 OD025315; 1 OT2 OD025337; 1 OT2 OD025276. In addition, the All of Us Research Program would not be possible without the partnership of its participants.



# References

- [1] Y. Erlich, T. Shor, I. Pe'er, and S. Carmi, "Identity inference of genomic data using long-range familial searches," *Science*, vol. 362, no. 6415, pp. 690–694, 2018.
- [2] C. A. Anderson, F. H. Pettersson, G. M. Clarke, L. R. Cardon, A. P. Morris, and K. T. Zondervan, "Data quality control in genetic case-control association studies," *Nature Protocols*, vol. 5, no. 9, pp. 1564–1573, 2010.
- [3] W. Astle and D. J. Balding, "Population structure and cryptic relatedness in genetic association studies," *Statistical Science*, vol. 24, no. 4, pp. 451–471, 2009.
- [4] G. Bhatia, A. Gusev, P.-R. Loh, H. Finucane, B. J. Vilhjálmsón, S. Ripke, Schizophrenia Working Group of the Psychiatric Genomics Consortium, S. Purcell, E. Stahl, M. Daly, *et al.*, "Subtle stratification confounds estimates of heritability from rare variants," *BioRxiv*, 2016. doi: [10.1101/048181](https://doi.org/10.1101/048181).
- [5] B. Devlin and K. Roeder, "Genomic control for association studies," *Biometrics*, vol. 55, no. 4, pp. 997–1004, 1999.
- [6] J. N. Hellwege, J. M. Keaton, A. Giri, X. Gao, D. R. Velez Edwards, and T. L. Edwards, "Population stratification in genetic association studies," *Current protocols in human genetics*, vol. 95, no. 1, pp. 1–22, 2017.
- [7] H. M. Kang, J. H. Sul, S. K. Service, N. A. Zaitlen, S.-y. Kong, N. B. Freimer, C. Sabatti, and E. Eskin, "Variance component model to account for sample structure in genome-wide association studies," *Nature genetics*, vol. 42, no. 4, pp. 348–354, 2010.
- [8] D. L. Newman, M. Abney, M. S. McPeck, C. Ober, and N. J. Cox, "The importance of genealogy in determining genetic associations with complex traits," *The American Journal of Human Genetics*, vol. 69, no. 5, pp. 1146–1148, 2001.
- [9] K. Shibata, A. Hozawa, G. Tamiya, M. Ueki, T. Nakamura, H. Narimatsu, I. Kubota, Y. Ueno, T. Kato, H. Yamashita, *et al.*, "The confounding effect of cryptic relatedness for environmental risks of systolic blood pressure on cohort studies," *Molecular Genetics & Genomic Medicine*, vol. 1, no. 1, pp. 45–53, 2013.
- [10] B. F. Voight and J. K. Pritchard, "Confounding from cryptic relatedness in case-control association studies," *PLoS genetics*, vol. 1, no. 3, e32, 2005.
- [11] A. I. Young, S. Benonisdottir, M. Przeworski, and A. Kong, "Deconstructing the sources of genotype-phenotype associations in humans," *Science*, vol. 365, no. 6460, pp. 1396–1400, 2019.

- [12] C. Bycroft, C. Freeman, D. Petkova, G. Band, L. T. Elliott, K. Sharp, A. Motyer, D. Vukcevic, O. Delaneau, J. O'Connell, *et al.*, "The UK Biobank resource with deep phenotyping and genomic data," *Nature*, vol. 562, no. 7726, pp. 203–209, 2018.
- [13] M. Blatt, A. Gusev, Y. Polyakov, and S. Goldwasser, "Secure large-scale genome-wide association studies using homomorphic encryption," *Proceedings of the National Academy of Sciences*, vol. 117, no. 21, pp. 11 608–11 613, 2020.
- [14] H. Cho, D. J. Wu, and B. Berger, "Secure genome-wide association analysis using multiparty computation," *Nature Biotechnology*, vol. 36, no. 6, pp. 547–551, 2018.
- [15] D. Froelicher, J. R. Troncoso-Pastoriza, J. L. Raisaro, M. A. Cuendet, J. S. Sousa, H. Cho, B. Berger, J. Fellay, and J.-P. Hubaux, "Truly privacy-preserving federated analytics for precision medicine with multiparty homomorphic encryption," *Nature communications*, vol. 12, no. 1, pp. 1–10, 2021.
- [16] H. Cho, D. Froelicher, J. Chen, M. Edupalli, A. Pyrgelis, J. R. Troncoso-Pastoriza, J.-P. Hubaux, and B. Berger, "Secure and Federated Genome-Wide Association Studies for Biobank-Scale Datasets," *bioRxiv*, 2022. doi: [10.1101/2022.11.30.518537](https://doi.org/10.1101/2022.11.30.518537).
- [17] A. Manichaikul, J. C. Mychaleckyj, S. S. Rich, K. Daly, M. Sale, and W.-M. Chen, "Robust relationship inference in genome-wide association studies," *Bioinformatics*, vol. 26, no. 22, pp. 2867–2873, 2010.
- [18] S. Purcell and C. Chang, *PLINK 2.00*, <https://www.cog-genomics.org/plink/2.0/>, accessed on 01.2023.
- [19] M. P. Conomos, A. P. Reiner, B. S. Weir, and T. A. Thornton, "Model-free estimation of recent genetic relatedness," *Am. J. Hum. Genet.*, vol. 98, no. 1, pp. 127–148, 2016.
- [20] A. Gusev, J. K. Lowe, M. Stoffel, M. J. Daly, D. Altshuler, J. L. Breslow, J. M. Friedman, and I. Pe'er, "Whole population, genome-wide mapping of hidden relatedness," *Genome research*, vol. 19, no. 2, pp. 318–326, 2009.
- [21] A. Naseri, X. Liu, K. Tang, S. Zhang, and D. Zhi, "Rapid: Ultra-fast, powerful, and accurate detection of segments identical by descent (ibd) in biobank-scale cohorts," *Genome biology*, vol. 20, no. 1, pp. 1–15, 2019.
- [22] R. Shemirani, G. M. Belbin, C. L. Avery, E. E. Kenny, C. R. Gignoux, and J. L. Ambite, "Rapid detection of identity-by-descent tracts for mega-scale datasets," *Nature communications*, vol. 12, no. 1, pp. 1–13, 2021.
- [23] P. Indyk and R. Motwani, "Approximate nearest neighbors: Towards removing the curse of dimensionality," in *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, 1998.
- [24] M. M. Hong, D. Froelicher, R. Magner, V. Popic, B. Berger, and H. Cho, "Secure Discovery of Genetic Relatives across Large-Scale and Distributed Genomic Datasets," in *Research in Computational Molecular Biology - 27th Annual International Conference (to appear)*, 2024.
- [25] M. M. Hong, D. Froelicher, R. Magner, V. Popic, B. Berger, and H. Cho, "Secure discovery of genetic relatives across large-scale and distributed genomic datasets," *bioRxiv*, 2024. doi: [10.1101/2024.02.16.580613](https://doi.org/10.1101/2024.02.16.580613).

- [26] J. Nait Saada, G. Kalantzis, D. Shyr, F. Cooper, M. Robinson, A. Gusev, and P. F. Palamara, "Identity-by-descent detection across 487,409 british samples reveals fine scale population structure and ultra-rare variant associations," *Nature communications*, vol. 11, no. 1, pp. 1–15, 2020.
- [27] M. D. Ramstetter, T. D. Dyer, D. M. Lehman, J. E. Curran, R. Duggirala, J. Blangero, J. G. Mezey, and A. L. Williams, "Benchmarking relatedness inference methods with genome-wide data from thousands of relatives," *Genetics*, vol. 207, no. 1, pp. 75–82, 2017.
- [28] *Hail: Relatedness*, <https://hail.is/docs/0.2/methods/relatedness.html>, accessed on 01.2023.
- [29] A. Z. Broder, "On the resemblance and containment of documents," in *Proceedings. Compression and Complexity of SEQUENCES 1997 (Cat. No. 97TB100171)*, 1997, pp. 21–29.
- [30] R. L. Rivest, L. Adleman, M. L. Dertouzos, *et al.*, "On data banks and privacy homomorphisms," *Foundations of secure computation*, vol. 4, no. 11, pp. 169–180, 1978.
- [31] Z. Brakerski, C. Gentry, and V. Vaikuntanathan, "(leveled) fully homomorphic encryption without bootstrapping," in *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, 2012.
- [32] J. H. Cheon, A. Kim, M. Kim, and Y. Song, "Homomorphic encryption for arithmetic of approximate numbers," in *Advances in Cryptology – ASIACRYPT 2017*, 2017.
- [33] C. Mouchet, J. R. Troncoso-pastoriza, J.-P. Bossuat, and J. P. Hubaux, "Multiparty homomorphic encryption from ring-learning-with-errors," in *Proceedings on Privacy Enhancing Technologies Symposium*, 2021.
- [34] D. Froelicher, J. R. Troncoso-Pastoriza, A. Pyrgelis, S. Sav, J. S. Sousa, J.-P. Bossuat, and J.-P. Hubaux, "Scalable privacy-preserving distributed learning," in *Proceedings on Privacy Enhancing Technologies Symposium*, 2021.
- [35] All of Us Research Program Investigators, "The "All of Us" research program," *New England Journal of Medicine*, vol. 381, no. 7, pp. 668–676, 2019.
- [36] *Lattigo v4*, <https://github.com/tuneinsight/lattigo>, EPFL-LDS, Tune Insight SA, Aug. 2022.
- [37] A. Naseri, J. Shi, X. Lin, S. Zhang, and D. Zhi, "RAFFI: accurate and fast familial relationship inference in large scale biobank studies using RaPID," *PLoS Genetics*, vol. 17, no. 1, e1009315, 2021.
- [38] S. Wang, M. Kim, W. Li, X. Jiang, H. Chen, and A. Harmanici, "Privacy-aware estimation of relatedness in admixed populations," *Briefings in Bioinformatics*, vol. 23, no. 6, bbac473, 2022.
- [39] L. Dervishi, X. Wang, W. Li, A. Halimi, J. Vaidya, X. Jiang, and E. Ayday, "Facilitating federated genomic data analysis by identifying record correlations while ensuring privacy," in *AMIA Annual Symposium proceedings 2023*, 2023.
- [40] G. Glusman, D. E. Mauldin, L. E. Hood, and M. Robinson, "Ultrafast comparison of personal genomes via precomputed genome fingerprints," *Frontiers in genetics*, vol. 8, p. 136, SEP 2017.

- [41] F. Hormozdiari, J. W. J. Joo, A. Wadia, F. Guan, R. Ostrosky, A. Sahai, and E. Eskin, "Privacy preserving protocol for detecting genetic relatives using rare variants," *Bioinformatics*, vol. 30, pp. i204–i211, 2014.
- [42] M. Robinson and G. Glusman, "Genotype fingerprints enable fast and private comparison of genetic testing results for research and direct-to-consumer applications," *Genes*, vol. 9, no. 10, p. 481, 2018.
- [43] D. He, N. A. Furlotte, F. Hormozdiari, J. W. J. Joo, A. Wadia, R. Ostrovsky, A. Sahai, and E. Eskin, "Identifying genetic relatives without compromising privacy," *Genome Research*, vol. 24, pp. 664–672, 2014.
- [44] *iDASH Privacy & Security Workshop 2023 - secure genome analysis competition*, <http://www.humangenomeprivacy.org/2023/competition-tasks.html>, accessed on 11.2023.
- [45] N. Homer, S. Szelinger, M. Redman, D. Duggan, W. Tembe, J. Muehling, J. V. Pearson, D. A. Stephan, S. F. Nelson, and D. W. Craig, "Resolving individuals contributing trace amounts of dna to highly complex mixtures using high-density snp genotyping microarrays," *PLoS genetics*, vol. 4, no. 8, e1000167, 2008.
- [46] D. Eastlake, T. Hansen, G. Fowler, K.-P. Vo, and L. Noll, *The fnv non-cryptographic hash algorithm*, <https://web.archive.org/web/20231024061616/https://datatracker.ietf.org/doc/html/draft-eastlake-fnv-17.html>, accessed on 10.2023.
- [47] P.-R. Loh, P. Danecek, P. F. Palamara, *et al.*, "Reference-based phasing using the haplotype reference consortium panel," *Nature Genetics*, vol. 48, no. 11, pp. 1443–1448, 2016.
- [48] J. Mbatchou, L. Barnard, J. Backman, *et al.*, "Computationally efficient whole-genome regression for quantitative and binary traits," *Nature Genetics*, vol. 53, no. 7, pp. 1097–1103, 2021.
- [49] S. Mendelsohn, D. Froelicher, D. Loginov, D. Bernick, B. Berger, and H. Cho, "sfkit: a web-based toolkit for secure and federated genomic analysis," *Nucleic Acids Research*, vol. 51, no. W1, W535–W541, 2023.
- [50] D. Froelicher, H. Cho, M. Edupalli, J. S. Sousa, J.-P. Bossuat, A. Pyrgelis, J. R. Troncoso-Pastoriza, B. Berger, and J.-P. Hubaux, "Scalable and privacy-preserving federated principal component analysis," in *2023 IEEE Symposium on Security and Privacy*, 2023.
- [51] *My Heritage*, <https://www.myheritage.com/>, accessed on 11.2023.