Journal of Medical Systems (2022) 46:84
https://doi.org/10.1007/s10916-022-01851-x

IMPLEMENTATION SCIENCE & OPERATIONS MANAGEMENT

New Approach to Privacy-Preserving Clinical Decision Support Systems for HIV Treatment

Gabriele Spini[1] · Emiliano Mancini[2,9,10] · Thomas Attema[1,3,4] · Mark Abspoel[3,5] · Jan de Gier[1] · Serge Fehr[3,4] · Thijs Veugen[1,3] · Maran van Heesch[1] · Daniël Worm[1] · Andrea De Luca[8] · Ronald Cramer[3,4] · Peter M.A. Sloot[2,6,7]

Received: 19 August 2020 / Accepted: 16 August 2022

Abstract

Background HIV treatment prescription is a complex process. Clinical decision support systems (CDSS) are a category of health information technologies that can assist clinicians to choose optimal treatments based on clinical trials and expert knowledge. The usability of some CDSSs for HIV treatment would be significantly improved by using the knowledge obtained by treating other patients. This knowledge, however, is mainly contained in patient records, whose usage is restricted due to privacy and confidentiality constraints.

Methods A treatment effectiveness measure, containing valuable information for HIV treatment prescription, was defined and a method to extract this measure from patient records was developed. This method uses an advanced cryptographic technology, known as secure Multiparty Computation (henceforth referred to as MPC), to preserve the privacy of the patient records and the confidentiality of the clinicians' decisions.

Findings Our solution enables to compute an effectiveness measure of an HIV treatment, the average time-to-treatment-failure, while preserving privacy. Experimental results show that our solution, although at proof-of-concept stage, has good efficiency and provides a result to a query within 24 min for a dataset of realistic size.

Interpretation This paper presents a novel and efficient approach HIV clinical decision support systems, that harnesses the potential and insights acquired from treatment data, while preserving the privacy of patient records and the confidentiality of clinician decisions.

Keywords Clinical decision support systems · Anti-HIV agents · Secure multiparty computation · Privacy · Confidentiality

This article is part of the Topical Collection on *Implementation Science & Operations Management*

We mourn the loss of Prof. Andrea De Luca, co-author of this paper, who tragically passed away in a car accident. We remember his as a kind and generous colleague, who contributed significantly to HIV treatment and drug resistances science.

✉ Gabriele Spini
gabriele.spini@tno.nl

[1] Applied Cryptography and Quantum Algorithms, TNO, 96800, 2509 JE Postbus, The Hague, The Netherlands

[2] Institute for Advanced Study, University of Amsterdam, Oude Turfmarkt 147, 1012 GC Amsterdam, The Netherlands

[3] Cryptology Group, CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

[4] Mathematical Institute, Leiden University, P.O. Box 9512, 2300 RA Leiden, The Netherlands

[5] Philips Research, High Tech Campus 34, 5656 AE Eindhoven, The Netherlands

[6] Complexity Institute, Nanyang Technological University, Academic Building North, Level 1 Section B Unit No. 7 (ABN-01B-07), 61 Nanyang Drive, 637335 Singapore, Singapore

[7] Advanced Computing, ITMO University, Lomonosova street 9, 191002 Saint Petersburg, Russia

[8] Department of Medical Biotechnologies, University of Siena and Siena University Hospital, Viale Mario Bracci 16, 53100 Siena, Italy

[9] Department of Global Health, Amsterdam UMC, Location AMC, 1105 AZ Amsterdam, The Netherlands

[10] Data Science Institute, Hasselt University, Diepenbeek, Belgium

Background and Significance

The constantly rising cost of national healthcare [1] associated to an aging population has highlighted the need for a critical change in traditional healthcare [2, 3]. Most stakeholders (clinicians, healthcare providers, policy makers and patients) agree that the solution lies in new approaches in which technology and health information technology (HIT) play a critical role [4, 5]. HIT services aim to automate and optimize healthcare processes with the overall goal of providing a more effective treatment process for patients. One of the main barriers to the adoption of HIT lies in the challenges associated with the need to preserve the privacy of the patients' data; legislation on the privacy of sensitive data, such as the General Data Protection Regulation (EU) 2016/679 (GDPR), is becoming increasingly more stringent, affecting all parties who handle sensitive data.

In this paper, we focus on one specific category of HIT systems: *Clinical Decision Support Systems* (CDSSs). A CDSS is a system that provides clinicians, patients, and other individuals with intelligently processed disease-specific and patient-specific data. Several different categories of CDSSs can be found in literature, such as diagnostic tools, expert systems, and workflow support. Systematic reviews reported that CDSSs significantly improved clinical practice: a review [6] on one hundred studies reported improvements for more than 62% of the trials on practitioner performance, reminder systems, drug-dosing systems and disease management systems. A review on seventy studies [7] reported a significant improvement of clinical practice in 68% of trials. Recent systematic reviews [8, 9] report an improvement in health care processes in 148 randomized, controlled trials and in 85% of twenty-two studies respectively.

As a use case to present our proposed solution to the problem of preserving the privacy of patients' data, we focus on an expert system for HIV treatment. The prescription of antiretroviral drugs to HIV1 infected patients is a complex process in which clinicians have to take into account several factors in a short amount of time. In particular, clinicians need to choose the most appropriate treatment based on the genotype of each patient's most prevalent strain of the virus in order to minimize drug resistance. A suboptimal treatment will likely result in a more rapid emergence of drug-resistant strains, and, eventually, in increased morbidity and mortality.

CDSSs are used in order to minimize or, ideally, prevent the prescription of suboptimal HIV1 treatments. Some examples of relevant CDSSs range from simple quality improvement consultation programs like HIVQUAL-US [10] that monitors clinical performance, to more sophisticated data-driven systems like Euresist [11] and knowledge-based systems like the HIVdb Program [12]. The main advantage of the use of these CDSSs is that they save a considerable amount of the clinician's time, since it would be impossible for the clinicians to analyze in detail the differences between the HIV genotype that is prevalent in a specific patient in search of critical mutations. However, in this paper we focus on the "comparative Drug Ranking System" (cDRS), a CDSS that helps to minimize the choice of sub-optimal HIV treatments by performing a meta-ranking analysis of three expert systems for HIV-1 genotypic drug resistance interpretation (ANRS, HIVdb, Rega) to resolve possible discordances between them [13–16]. The discordances in drug resistance between the three expert systems are not negligible [17, 18], and are the result of the limited amount of clinical data available for each specific set of mutations and of different methodologies used by the systems. A CDSS able to help clinicians in resolving such discordances is essential to avoid the administration of suboptimal HIV treatments.

Research on the spread of the HIV epidemics has led to the development of tools (such as phylogenetic trees) able to correlate specific viral sequences in different patients and reconstruct with good accuracy the network of infections within a community [19]. In addition, transmission events between patients can be identified by analyzing the viral genotypes, given the uniqueness of specific sets of mutations [20, 21]. Hence, strict privacy regulations prevent the sharing of patient data (e.g., viral genotype) that feed and improve these clinical decision support systems. Moreover, clinicians might not be able, or willing, to openly share their treatment decisions and the resulting outcomes, even though such information might be beneficial for the decision-making process of their colleagues. In conclusion, there is a tremendous amount of valuable information that is unavailable to clinicians because of privacy and confidentiality constraints.

An ideal system should allow clinicians to compare their chosen treatment against the outcome of the treatments chosen by other clinicians for similar genotypes, while solving the issue of utilizing patient and clinicians' data in a secure, privacy-preserving way.

In this exploratory work, we present a solution that uses cryptographic techniques, namely a so-called *secure Multiparty Computation (MPC)* protocol, to achieve this functionality without violating any of the privacy constraints. Informally stated, MPC is a collection of cryptographic techniques that allow several parties, each of which holds some private input, to evaluate a function on those inputs without disclosing any extra information on the input themselves, and without resorting to a trusted external party. Our MPC-based solution would allow clinicians to compare past treatments of 'similar' patients to find the optimal treatment

for new patients preventing any unauthorized party, including the ones performing the computations, to access the input data.

Related Work

Privacy-preserving CDSSs have been presented in recent years [22–25]. However, this line of work focusses on CDSSs for disease-prediction and enables clinicians to securely query remote machine-learning based systems for a given patient's health condition, in a privacy-preserving way. As such, it is not directly comparable with our solution, which has a different scope within the paradigm of privacy-preserving CDSSs.

In more general terms, proposed applications of MPC to the healthcare sector have flourished in recent years. To the best of our knowledge, there exists no article summarizing the scientific literature on MPC applied to the healthcare sector[1]; we provide here a list of recent and relevant work on the topic, but we stress the fact that this list cannot be exhaustive, due to the high number of publications on the topic.

A large sub-domain of the application of MPC (and other related cryptographic techniques) to the medical domain aims to deploy machine-learning techniques on medical datasets held by distinct organizations; examples in this sense include privacy-preserving reinforcement learning [26], Kaplan-Meier survival analysis and genome-wide association studies [27], grid logistic regression for biomedical data [28], training of linear [29] and Lasso [30] regression models on medical data, and computing patient risk-stratification metrics [31].

Other relevant work include medical record searching [32, 33], the study of general methods such as privacy-preserving data mining for joint data analysis between hospitals [34] and branching programs for privacy-preserving classification of medical ElectroCardioGram signals [35], the presentation of specific use case scenarios such as secure disclosure of patient data for disease surveillance [36], R-based healthcare statistics [37], and privacy-preserving genome-wide association study [38], privacy-preserving genome analysis [39] and search of similar patients in genomic data [40].

Finally, iDASH [41] is an important public initiative to stimulate the development of techniques for privacy-preserving sharing of medical data.

[1] One literature overview on privacy-preserving medical data sharing has been produced, but with a focus on blockchain-based applications.

Outline

The rest of the article is organized as follows. In the "Materials and Methods" section we first discuss how to measure the effectiveness of a treatment from patient records and present the method that we propose (setting aside the privacy-preserving aspect); we then provide a brief overview of MPC and of the framework of our choice, SPDZ. In the "Results" section we then explain how the effectiveness measure is securely implemented within SPDZ and present an evaluation of the efficiency of our solution. Finally, the "Discussion and Conclusions" section summarizes the results of the article and provides an appraisal of the achieved results and on possible future work.

Materials and Methods

Measuring Treatment Effectiveness from Patient Records

The viral genotype of a patient refers to the genetic sequence(s) of the HIV-1 virus strain that is most prevalent at the time of the blood test. The HIV-1 virus RNA genome contains 3 key regions that encode for enzymes critical to the life cycle of the virus: *protease (P), integrase (I)* and *reverse transcriptase (RT)*. Each region encodes for enzymes with 99, 288 and 560 amino acids, respectively, all of which could in principle mutate. These mutations play an important role in the drug resistance of the virus strains.

Given an HIV-1 patient, our goal is to obtain treatment results of 'similar' patients, and therefore we need to define a metric or distance function that quantifies the similarity between two patients, or two viral genotypes. Since all expert systems indicate resistance to drugs based on substitutions in the amino acid sequence of the wild-type HIV-1, we need a way to compute the distance in the amino acid sequences of the viral proteins. Metrics of distances between amino acid sequences are fairly complex and often assessed via neural networks [42]. The assignment of a suitable similarity metric is outside the scope of this paper, and for this reason, we have chosen to use a simplified viral genotype representation with a generic metric as a proof of concept. However, our solution is flexible, since it can support other representations and metrics.

From now on we shall represent viral genotypes as bit strings v of a fixed length N, i.e., $v \in \{0,1\}^N$. We can think of each bit in this bit-string as an indicator for the presence or the absence of a specific mutation at a specific position.

Since there are only 97 relevant positions with commonly 1 or 2 resistance-associated substitutions [64] we can expect N to be somewhere between 100 and 200. The

distance between two viral genotypes v_1 and v_2 is defined by the Hamming distance between the bit strings:

$$H(v_1, v_2) = |\{i : v_1(i) \neq v_2(i)\}|,$$

Given this metric we can define two viral genotypes v_1 and v_2 to be similar if their Hamming distance is smaller than a certain threshold B, i.e., $H(v_1, v_2) < B$. Even though this metric is a simplification of the metrics used in practice, it is quite similar to the rule-based metrics used in the CDSSs of [13, 16]. These CDSSs match viral genotypes based on the presence of resistance-associated substitutions in amino acid positions, which can be seen as a Boolean expression. In a clinical setting, these CDSSs compare the two complete viral strings to identify specific insertion, deletions, and substitutions but do not rely on a single threshold value defined as a Hamming distance. In fact, it is well known from specific studies which additions, deletions or substitutions trigger a clinically relevant mutation. In a practical implementation, we would have to look at the difference in specific positions of the sequences of two amino acid strings. The threshold that would be used in that case would be defined by clinicians who set of rules used by the specific CDSS instead of the Hamming distance described in the example. However, the rulesets that would trigger an alert are still Boolean in nature and would fit the proposed secure MPC solution.

Suboptimal treatments of HIV-1 patients result in faster emergence of resistant strains and this emergence renders the treatment ineffective. Hence, a way to measure the effectiveness of a treatment *tr* for genotype *v* is by indicating the *time-to-treatment-failure* $TTF_{tr}(v)$. The $TTF_{tr}(v)$ is defined as the time (in days) between the start of a therapy *tr* and either a therapy switch, a discontinuation of therapy or death [65, 66], for a patient with genotype *v*. Hence, given an HIV-1 patient with genotype *v* we would, for example, like to compute the average $\overline{TTF}_{tr}(v)$ over all patients with similar genotype v_i, as an indication for the unknown true effectiveness measure $TTF_{tr}(v)$:

$$\overline{TTF}_{tr}(v) = \frac{1}{|\{i : H(v, v_i) < B\}|} \sum_{i:H(v,v_i)<B} TTF_{tr}(v_i),$$

Where H denotes, as discussed above, the Hamming distance and B denotes a fixed threshold value.

Secure Multiparty Computation

MPC has been introduced by Yao in the 1980s [43]. Given *n* mutually distrusting parties $P_1, \ldots, P_n$, each holding private inputs x_1,…, x_n, the goal of MPC is to allow the parties to compute the value $f(x_1, \ldots, x_n)$ of a function f on their inputs, without revealing any other information than $f(x_1, \ldots, x_n)$, and without resorting to an external trusted party.

Early research in the 1980s [43–46] established the theoretical feasibility bounds for MPC; informally stated, this line of research proved that any function *f* with finite domain and finite image can be evaluated securely in an MPC fashion. The precise security properties that can be achieved depend on the behavior of players and on the underlying communication model.

Since the first market-ready deployment of MPC in 2008 [47], MPC solutions have been used in various practical contexts, e.g., stock market order matching [48], job market inquiries [49], and frequency bands auctions [50]. Moreover, various software suites and implementation frameworks for MPC have been made available [51–55].

Several considerations have to be made when applying MPC to a given problem. For instance, one may assume that parties $P_1, \ldots, P_n$ will behave semi-honestly (meaning that they may try to learn information on the other parties' inputs, but do follow the protocol), or that it is instead necessary to provide security against fully malicious players that deviate from the protocol instructions. Another important parameter that varies among protocols is the number *t* of corrupted parties that can be tolerated out of the total number *n* of parties.

A remark of notable importance is that many desirable properties of MPC may negatively impact performance, or even be mutually exclusive, which means that the choice of an MPC protocol may be subject to important trade-offs. The reader can refer to [56] for a comprehensive discussion of MPC.

The MPC framework of our choice: SPDZ

We base our MPC solution on the SPDZ protocol [57, 58]. The protocol is distinguished for its fast performance, and is implemented in a freely accessible software suite called SPDZ-2 [52, 54] for UNIX-based systems[2]; SPDZ-2 allow developers to write programs in Python-like syntax, and it then compiles the code to executable format.

SPDZ follows the so-called *share-compute-reveal* paradigm: each input x_i of the function *f* to be computed is 'dispersed' (or, formally speaking, *secret-shared*[3]) into *n* pieces of data, called shares, each of which is assigned to a party;

[2] Notice that support for the SPDZ-2 software suite (implementing the eponymous MPC protocol) has been discontinued. Development moved to forks SCALE-MAMBA and MP-SPDZ, both implementations of the SPDZ protocol.

[3] It is important to notice that 'sharing', here, is by no means a synonym of 'revealing'; on the contrary, it can be seen as a strong form of distributed *encryption.*

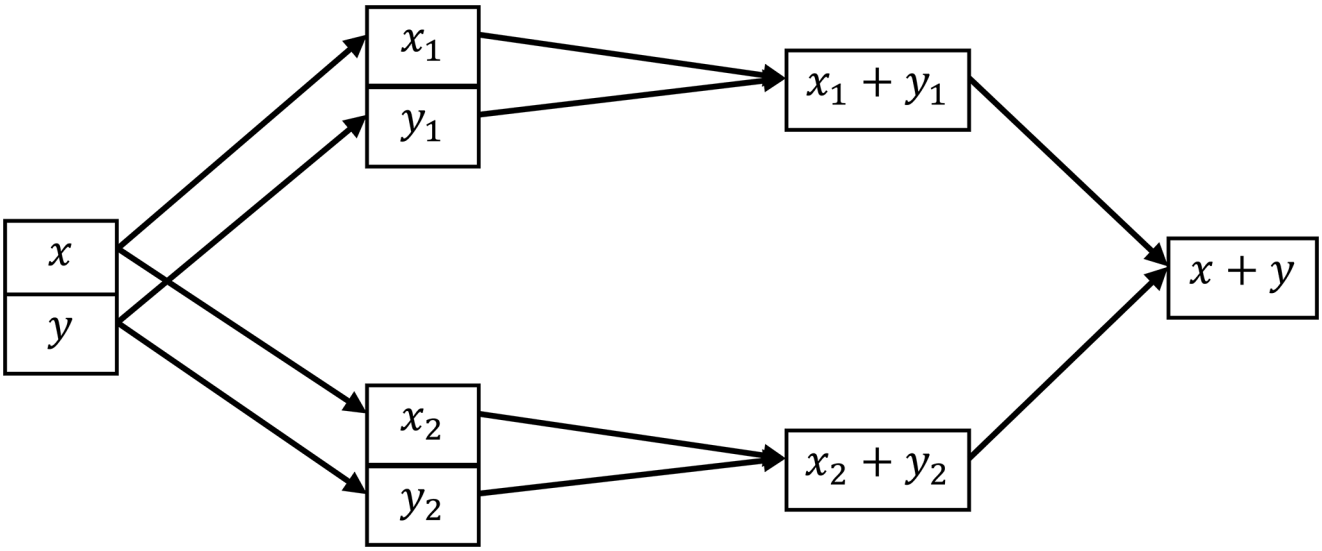

Fig. 1 Example of share-compute-reveal paradigm for addition of two values

this process has the property that no information on x_i can be extracted from a set of shares, unless such a set contains *all* shares (in which case x_i can be completely recovered). Subsequently, parties execute a 'computation' protocol; as a result of this step, each party will have a share of the output $f(x_1, \ldots, x_n)$ of the function. Once all shares have been gathered, the output can then be reconstructed.

A schematic representation of this paradigm for the addition of two values x and y among two parties is provided in Fig. 1. The top row represents the shares held by the first party, while the bottom row represents the shares of the second party. The assumption here is that x_1 and x_2 are two random values subject to the condition that $x_1 + x_2 = x$, and similarly for y; bearing this fact in mind, it is then seen how the process respects the privacy and reconstruction requirements discussed above.

In more general terms, the core idea behind MPC protocols based on the share-compute-reveal paradigm is that the function f to be evaluated on the input values is "decomposed" into basic operations (such as sum and products); these basic operations are then translated into similar operations on the shares and executed in the same order. An important remark is that, in general, these operations on shares require some form of interaction among the parties (for instance, multiplication of two values cannot simply be performed by multiplying the corresponding shares, and requires a more involved and interactive process). The reader can refer to the literature on MPC and on SPDZ that we have provided for a more formal and complete discussion of this topic.

Other cryptographic techniques such as homomorphic encryption [59, 60] could potentially be of relevance for private data analysis, but we ruled out these alternatives, because they would induce a huge computational overhead in our setting.

The share-compute-reveal approach is particularly well-suited for the client-server model we are interested in: the 'input' parties, clinicians (clients) simply need to supply their secret-shared inputs to two or more 'computing' parties (servers), who will execute the computation protocol on these inputs, and then communicate the shares of the output to the input parties, which can thus reconstruct the output.

It is important to remark that the SPDZ protocol does not, *per se*, distinguish between input and computing parties. A framework for MPC in a client-server model was presented in [61]; moreover, in [63] the SPDZ protocol was adjusted to the client-server setting.

The SPDZ protocol is divided into an 'offline' phase and an 'online' phase. The offline phase can be executed before the function inputs $x_1,\ldots,x_n$ are known, and its goal is to produce some secret-shared auxiliary data that will be used in the evaluation of f; producing this data can be a computationally-intensive process, but since secret inputs are not required, this step can be executed during idle time and well before the actual secure computation will take place. Once the auxiliary data has been produced, the evaluation of f can be performed very efficiently: this is of particular relevance for our use case, where input parties (clinicians) need to obtain the output of the function f within a matter of minutes, while preprocessing material can be produced in the background by the computing parties.

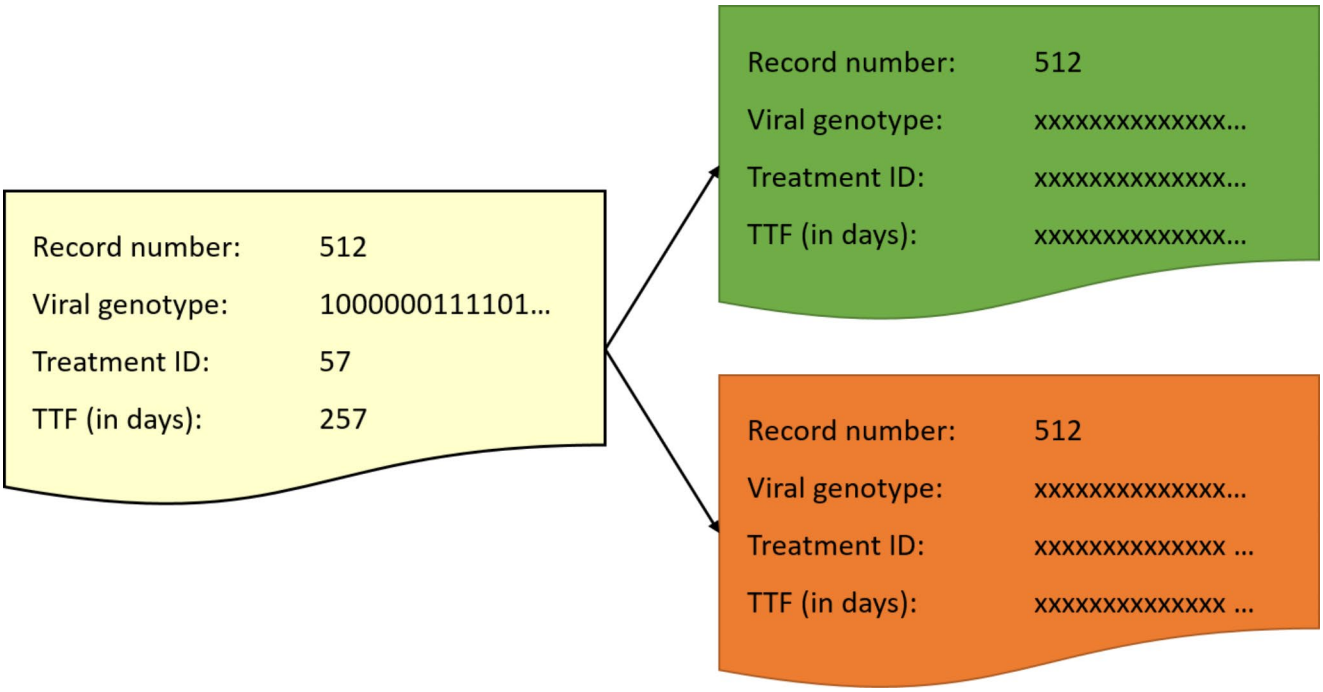

Fig. 2 Secret sharing database records

Results

The functionality we have achieved utilizes HIV patient records to gain new insights in the effectiveness of HIV treatments. The MPC protocol ensures privacy of the patients and the confidentiality of the clinicians' treatment decisions.

The proposed solution distinguishes between 'input' parties, the clinicians supplying the database records, and 'computing' parties running the SPDZ protocol, which can be different medical institutions or IT service providers. The input parties additively secret-share their data records and distribute the shares amongst the computing parties (see Fig. 2).

As a result, the two computing parties each hold a share of all the database records. SPDZ allows the evaluation of queries to this secret-shared database in such a way that only the output of the query (the average time-to-treatment-failure ($\overline{TTF}$) per treatment) is revealed to the clinician, and no additional information is leaked to either the querying clinician, or the computing parties (cf. previous section). In order to protect the private information in the query (the viral genotype), we secret-share the query amongst the computing parties in a similar manner. The computing parties thus take as private inputs their shares of the database records and their share of the query. They do not reconstruct the result of the computation (the average *TTF*) themselves; instead, each of them sends their share of the result to the querying clinician who, in turn, recombines the shares to reconstruct the output. This way the result is only revealed to the clinician, and not to the computing parties (cf. Figure 3).

Our solution allows clinicians to compare their treatment of choice against the outcome of treatments previously chosen by other clinicians for patients with similar genotype, without revealing any private information to the clinicians or the computing parties, who only learn the size and format of the database and the number of queries to the database. This system is secure as long as the two computing parties do not collude.

Performance – Online Phase

In comparison to implementing the functionality without privacy protection, using MPC inherently introduces computational and communication overhead. The main reason for this unavoidable overhead is that, in an MPC protocol, the computation path has to be oblivious, i.e., independent, of the input values, since it would otherwise leak information. Moreover, as explained in the previous section, some basic operations on the input data are translated by MPC into more complex, interactive processes, which lead to unavoidable overhead.

We have evaluated the performance of the online phase of our protocol by deploying the computing parties on two different machines, each using one core of a i7-7567U CPU running at 3.50 GHz and 32 GB of RAM, in a local network with 1 Gbit/s throughput. The system ran on a Fedora operating system and has been developed within the

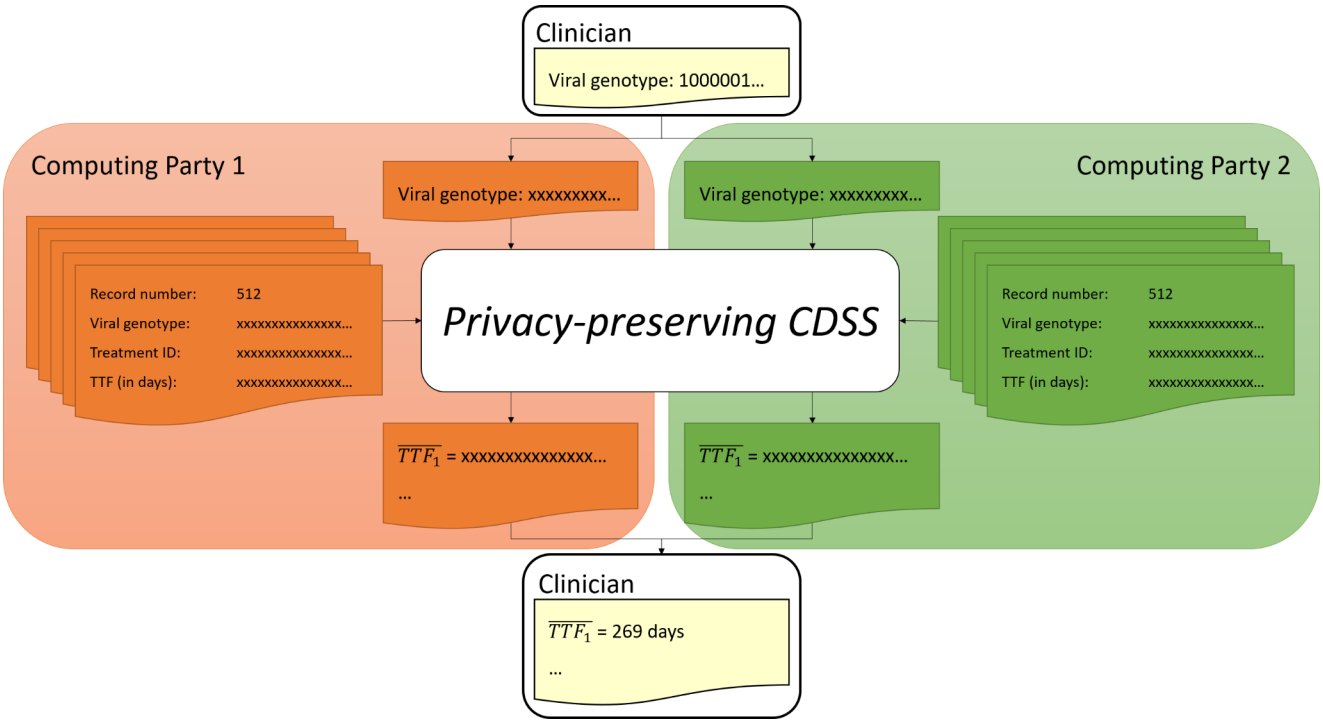

Fig. 3 Query architecture of the privacy-preserving CDSS

Fig. 4 CDSS computation time

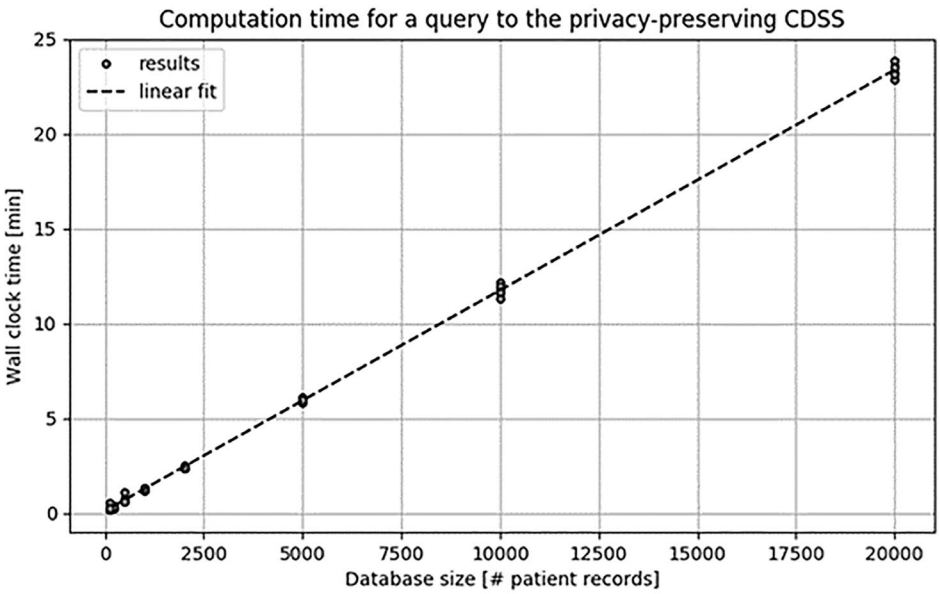

SPDZ-2 software suite discussed in the previous section; the overall orchestration of the scalability experiments has been performed via scripts for the Bash shell. Finally, we have instantiated the SPDZ protocol with 40-bit statistical security, 128-bit computational security and a 128-bit prime field.

The experiments that we have run measure the time it takes for the solution to return the average time-to-treatment-failure $\overline{TTF}_{tr}(v)$ of a given input treatment *tr*, where the additional input value v is the genotype of a given patient. The formal definition of $\overline{TTF}_{tr}(v)$ is presented in the "Materials and Methods" section; as a reminder, it is given by the average over the times-to-treatment-failure of patients with a similar genotype, for the same treatment *tr*.

The results in Fig. 4 show the computation times that are needed for answering one query, for artificially-generated databases with sizes ranging from 100 to 20 000 records. The maximum 20 000 approximates the number of HIV-positive

registered individuals in the Netherlands [62]. The experiment is repeated multiple times, resulting in several data points per database size. Recall that per query we compute the average *TTF* conditioned on 'similar' patients for 100 different treatments. The computational complexity scales linearly in the number of database records. Notice that the threshold value *B* and the number of patient genotypes that have Hamming distance at most *B* from the given input do not affect the running time of the computation: this is inherent to MPC solutions, which have a computation time which does not depend on the input values.

Also notice that these figures refer to the time needed to answer a single query; with the current state of our implementation, the running time would scale linearly in the number of queries. This is an aspect to be kept in mind should, for instance, a practitioner want to query the system for different values of the threshold *B*.

Performance – Offline Phase

In the SPDZ protocol certain computational tasks are executed in the offline phase, that is independent of the MPC use case and that can be implemented with existing protocols. For this reason, we have merely estimated the computational costs of it. The offline phase can be run at any time to generate a large database of preprocessed data which, in turn, is consumed during the online phase.

The performance of the offline phase can be quantified in the number of the so-called multiplication triples that are generated per second. In [63] various approaches for generating multiplication triples in a setting similar to ours were evaluated, generating 30 000 triples/s. To evaluate a single query on a database with 20 000 records approximately 40 million multiplication triples are required. In this setting these triples can thus be generated in approximately 22 min.

Discussion and Conclusions

We presented a novel approach for HIV1 clinical decision support systems, making use of advanced cryptographic techniques to process private information without revealing it. By making use of MPC, we can ensure both the privacy of the clinicians' treatment choices and the privacy of patients.

Towards a fully operational deployment some points are yet to be addressed. Notably, the SPDZ software framework is designed for research purposes only, which means that our implementation should be audited and checked for vulnerabilities. For what concerns efficiency and scalability, we stress the fact that any CDSS for HIV treatment should produce a suggestion within minutes, since practitioners would typically query the system right after visiting a patient and would expect an answer before the patient leaves their office. As shown in Fig. 3, our solution answers a query within 24 min, for a database size roughly matching the number of HIV-positive registered individuals in the Netherlands [62]; while we consider this result to be sufficient for the proof-of-concept presented in this paper, some further work would be needed for a full-scale deployment. The running time of the implementation could be improved by several means, e.g., by using a low-level but very fast programming language such as C, by further parallelizing the computation, or by making use of high-performance computing machines instead of consumer-level hardware.

Acknowledgements The authors would like to thank Pia Kempker for her valuable contributions to the early stages of this research.

Authors' contribution TA contributed to the cryptographical conception of the solution, to the implementation of the proof-of-concept, and to the writing of the manuscript. EM contributed to the conception of the clinical-decision aspects of the solution, and to the writing of the manuscript. GS contributed to the cryptographical conception of the solution, to the implementation of the proof-of-concept, and to the writing of the manuscript. MA contributed to the cryptographical conception of the solution and to the implementation of the proof-of-concept. JdG contributed to the implementation of the proof-of-concept and ran the experiments. SF contributed to the cryptographical conception of the solution. TV contributed to the cryptographical conception of the solution. MvH contributed to the cryptographical conception of the solution and to the implementation of the proof-of-concept. DW contributed to the design of the solution. ADL contributed to the virology parts. RC and PS contributed to the conception of the work. All authors read and approved the final manuscripts.

Funding This work was supported by PPS-surcharge for Research and Innovation of the Dutch ministry of Economic Affairs and Climate Policy and ERC Advanced Investigator Grant 740972 (ALGSTRONGCRYPTO).

Data Availability No real data or material has been used for this article.

Code Availability The code produced in the work described in this article is currently not publicly available.

Declarations

Conflicts of Interest/Competing Interests No conflict of interest or competing interests to report.

References

1. Organisation for Economic Co-operation and Development (OECD), "Fiscal sustainability of health systems: bridging health and finance perspectives." OECD Publishing, 2015.
2. Beard, J. R., Officer, A., de Carvalho, I. A. et al., "The world report on ageing and health: a policy framework for healthy ageing," The Lancet, vol. 387, no. 10033, pp. 2145–2154, 2016.
3. Beard, J. R. and Bloom, D. E., "Towards a comprehensive public health response to population ageing," Lancet (London, England), vol. 385, no.9968, p. 658, 2015.
4. Payne, P. R. O., Lussier, Y., Foraker, R. E. et al., "Rethinking the role and impact of health information technology: informatics as an interventional discipline," BMC Medical Informatics and Decision Making, vol. 16,no. 1, p. 40, 2016.
5. Wachter, R., "Making it work: harnessing the power of health information technology to improve care in England," Report to the National Advisory Group on Health Information Technology in England. London: The Stationery Office, 2016.
6. Garg, A. X., Adhikari, N. K. J., McDonald, H. et al., "Effects of computerized clinical decision support systems on practitioner performance and patient outcomes: a systematic review," Jama, vol. 293, no. 10, pp. 1223–1238, 2005.
7. Kawamoto, K., Houlihan, C. A., Balas, E. A. et al., "Improving clinical practice using clinical decision support systems: a systematic review of trials to identify features critical to success," BMJ, vol. 330, no. 7494, p. 765, 2005.
8. Prgomet, M., Li, L., Niazkhani, Z. et al., "Impact of commercial computerized provider order entry (CPOE) and clinical decision support systems (CDSSs) on medication errors, length of stay, and mortality in intensive care units: a systematic review and meta-analysis," Journal of the American Medical Informatics Association, vol. 24, no. 2, pp.413–422, 2016.
9. Bright, T. J., Wong, A., Dhurjati, R. et al., "Effect of clinical decision support systems: a systematic review," Annals of Internal Medicine, vol. 157, no. 1, pp. 29–43, 2012.
10. Horberg, M. A., Aberg, J. A., Cheever, L. W. et al., "Development of national and multiagency HIV care quality measures," Clinical Infectious Diseases, vol. 51, no. 6, pp. 732–738, 2010.
11. Zazzi, M., Kaiser, R., Sönnerborg, A. et al., "Prediction of response to antiretroviral therapy by human experts and by the euresist data-driven expert system (the EVE study)," HIV medicine, vol. 12, no. 4, pp. 211– 218, 2011.
12. Tang, M. W., Liu, T. F., and Shafer, R. W., "The HIVdb system for HIV-1 genotypic resistance interpretation," Intervirology, vol. 55, no. 2, pp. 98–101, 2012.
13. Sloot, P. M. A., Coveney, P. V., Ertaylan, G. et al., "HIV decision support: from molecule to man," Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, vol. 367, no. 1898, pp. 2691–2703, 2009.
14. Sloot, P. M. A., Coveney, P. V., Bubak, M. T. et al., "Multi-science decision support for HIV drug resistance treatment," in Global Healthgrid: e-Science Meets Biomedical Informatics - Proceedings of HealthGrid 2008, Chicago, IL, USA, 2–4 June, 2008, ser. Studies in Health Technology and Informatics, Solomonides, T., Silverstein, J. C., Saltz, J. H. et al., Eds., vol. 138. IOS Press, 2008, pp. 188–198.
15. Sloot, P. M. A., Coveney, P., Bubak, M. T. et al., "Virolab: a collaborative decision support system in viral disease treatment," Reviews in Antiviral Therapy, vol. 3, pp. 4–7, 2008.
16. Sloot, P. M. A., Boukhanovsky, A. V., Keulen, W. et al., "A grid-based HIV expert system," Journal of Clinical Monitoring and Computing, vol. 19, no. 4–5, pp. 263–278, 2005.
17. Frentz, D., Boucher, C. A. B., Assel, M. et al., "Comparison of HIV-1 genotypic resistance test interpretation systems in predicting virological outcomes over time," PloS one, vol. 5, no. 7, p. e11505, 2010.
18. Ravela, J., Betts, B. J., Brun-Vézinet, F. et al., "HIV-1 protease and reverse transcriptase mutation patterns responsible for discordances between genotypic drug resistance interpretation algorithms," Journal of acquired immune deficiency syndromes, vol. 33, no. 1, pp. 8–14, 2003.
19. Zarrabi, N., Prosperi, M., Belleman, R. G. et al., "Combining epidemiological and genetic networks signifies the importance of early treatment in HIV-1 transmission," PloS one, vol. 7, no. 9, p. e46156, 2012.
20. Leitner, T. and Albert, J., "Reconstruction of HIV-1 transmission chains for forensic purposes," AIDS Rev, vol. 2, pp. 241–251, 2000.
21. Hué, S., Pillay, D., Clewley, J. P. et al., "Genetic analysis reveals the complex structure of hiv-1 transmission within defined risk groups," Proceedings of the National Academy of Sciences, vol. 102, no. 12, pp. 4425–4429, 2005.
22. Liu, Ximeng, et al. "Hybrid privacy-preserving clinical decision support system in fog–cloud computing." Future Generation Computer Systems 78 (2018): 825–837.
23. Liu, Ximeng, et al. "Privacy-preserving outsourced clinical decision support system in the cloud." IEEE Transactions on Services Computing (2017).
24. Rahulamathavan, Yogachandran, et al. "Privacy-preserving clinical decision support system using gaussian kernel-based classification." IEEE journal of biomedical and health informatics 18.1 (2013): 56–66.
25. Liu, Ximeng, et al. "Privacy-preserving patient-centric clinical decision support system on naive Bayesian classification." IEEE journal of biomedical and health informatics 20.2 (2015): 655–668.
26. X. Liu, R. H. Deng, K. -K. Raymond Choo and Y. Yang, "Privacy-Preserving Reinforcement Learning Design for Patient-Centric Dynamic Treatment Regimes," in IEEE Transactions on Emerging Topics in Computing, vol. 9, no. 1, pp. 456–470, 1 Jan.-March 2021, doi: https://doi.org/10.1109/TETC.2019.2896325.
27. Froelicher, D., Troncoso-Pastoriza, J.R., Raisaro, J.L. et al. Truly privacy-preserving federated analytics for precision medicine with multiparty homomorphic encryption. Nat Commun 12, 5910 (2021). https://doi.org/10.1038/s41467-021-25972-y.
28. Shi, H., Jiang, C., Dai, W. et al. Secure Multi-pArty Computation Grid LOgistic REgression (SMAC-GLORE). BMC Med Inform Decis Mak 16, 89 (2016). https://doi.org/10.1186/s12911-016-0316-1.
29. Dankar FK, Madathil N, Dankar SK, Boughorbel S. Privacy-Preserving Analysis of Distributed Biomedical Data: Designing Efficient and Secure Multiparty Computations Using Distributed Statistical Learning Theory. JMIR Med Inform 2019;7(2):e12702. doi: https://doi.org/10.2196/12702.
30. van Egmond, M.B., Spini, G., van der Galien, O. et al. Privacy-preserving dataset combination and Lasso regression for healthcare predictions. BMC Med Inform Decis Mak 21, 266 (2021). https://doi.org/10.1186/s12911-021-01582-y
31. Dong X, Randolph DA, Weng C, Kho AN, Rogers JM, Wang X. Developing High Performance Secure Multi-Party Computation Protocols in Healthcare: A Case Study of Patient Risk Stratification. AMIA Jt Summits Transl Sci Proc. 2021 May 17;2021:200–209. PMID: 34457134; PMCID: PMC8378657.
32. Rogers, J., Adetoro, E., Bater, J., Canter, T., Fu, D., Hamilton, A., … Kho, A. (2022). VaultDB: A Real-World Pilot of Secure Multi-Party Computation within a Clinical Research Network. arXiv preprint arXiv:2203.00146.
33. Y. Sun, J. Liu, K. Yu, M. Alazab and K. Lin, "PMRSS: Privacy-Preserving Medical Record Searching Scheme for Intelligent Diagnosis in IoT Healthcare," in IEEE Transactions on Industrial

Informatics, vol. 18, no. 3, pp. 1981–1990, March 2022, doi: https://doi.org/10.1109/TII.2021.3070544.
34. Lindell, Y. and Pinkas, B., "Secure multiparty computation for privacy-preserving data mining," IACR Cryptology ePrint Archive, vol. 2008, p. 197, 2008.
35. Barni, M., Failla, P., Kolesnikov, V. et al., "Secure evaluation of private linear branching programs with medical applications," in Computer Security - ESORICS 2009, 14th European Symposium on Research in Computer Security, Saint-Malo, France, September 21–23, 2009. Proceedings, ser. Lecture Notes in Computer Science, Backes, M. and Ning, P., Eds., vol. 5789. Springer, 2009, pp. 424–439.
36. Emam, K. E., Hu, J., Mercer, J. et al., "A secure protocol for protecting the identity of providers when disclosing data for disease surveillance," JAMIA, vol. 18, no. 3, pp. 212–217, 2011.
37. Chida, K., Morohashi, G., Fuji, H. et al., "Implementation and evaluation of an efficient secure computation system using R for healthcare statistics," JAMIA, vol. 21, no. e2, pp. e326–e331, 2014.
38. Bonte, C., Makri, E., Ardeshirdavani, A. et al., "Privacy-preserving genome-wide association study is practical," IACR Cryptology ePrint Archive, vol. 2017, p. 955, 2017.
39. Kim, M. and Lauter, K., "Private genome analysis through homomorphic encryption," BMC Medical Informatics and Decision Making, vol. 15, no. 5, p. S3, Dec 2015.
40. Asharov, G., Halevi, S., Lindell, Y. et al., "Privacy-preserving search of similar patients in genomic data," Proceedings on Privacy Enhancing Technologies, vol. 2018, no. 4, pp. 104–124, 2018.
41. Ohno-Machado, L., Bafna, V., Boxwala, A. A. et al., "iDASH: integrating data for analysis, anonymization, and sharing," Journal of the American Medical Informatics Association, vol. 19, no. 2, pp. 196–201, 2012.
42. Waterman, M. S., Smith, T. F., and Beyer, W. A., "Some biological sequence metrics," Advances in Mathematics, vol. 20, no. 3, pp. 367–387, 1976.
43. Yao, A. C., "Protocols for secure computations (extended abstract)," in 23rd Annual Symposium on Foundations of Computer Science, Chicago, Illinois, USA, 3–5 November 1982. IEEE Computer Society, 1982, pp. 160–164.
44. Goldreich, O., Micali, S., and Wigderson, A., "How to play any mental game or A completeness theorem for protocols with honest majority," in Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA, Aho, A. V., Ed. ACM, 1987, pp. 218–229.
45. Ben-Or, M., Goldwasser, S., and Wigderson, A., "Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract)," in Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2–4, 1988, Chicago, Illinois, USA, Simon, J., Ed. ACM, 1988, pp. 1–10.
46. Chaum, D., Crépeau, C., and Damgård, I., "Multiparty unconditionally secure protocols (extended abstract)," in Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2–4, 1988, Chicago, Illinois, USA, Simon, J., Ed. ACM, 1988, pp. 11–19.
47. Bogetoft, P., Christensen, D. L., Damgård, I. et al., "Secure multiparty computation goes live," in Financial Cryptography and Data Security, 13th International Conference, FC 2009, Accra Beach, Barbados, February 23–26, 2009. Revised Selected Papers, ser. Lecture Notes in Computer Science, Dingledine, R. and Golle, P., Eds., vol. 5628. Springer, 2009, pp. 325–343.
48. Partisia, "Secure order matching," https://partisia.com/order-matching/, 2018, accessed: 2018-06-27.
49. Sharemind, "Big data analytics protection," https://sharemind.cyber.ee/ big-data-analytics-protection/, 2018, accessed: 2018-06-27.
50. Partisia, "Spectrum auctions," https://partisia.com/spectrum-auctions/, 2018, accessed: 2018-06-27.
51. VIFF Development Team, "VIFF: Virtual Ideal Functionality Framework," http://viff.dk/, 2018, accessed: 2018-06-27.
52. Bristol Crypto, "SPDZ-2: Multiparty computation with SPDZ, MASCOT, and Overdrive offline phases." https://github.com/bristolcrypto/ SPDZ-2, 2018, accessed: 2018-06-27.
53. FRESCO Development Team, "FRESCO: FRamework for Efficient and Secure COmputation." https://github.com/aicis/fresco, 2018, accessed: 2018-06-27.
54. COSIC KU Leuven, "Secure Computation Algorithms from LEuven (SCALE) and Multiparty AlgorithMs Basic Argot (MAMBA)," https://github.com/KULeuven-COSIC/SCALE-MAMBA, 2018, accessed: 2018-06-27.
55. Schoenmakers, B. "MPyC-secure multiparty computation in Python, v0. 4.7. GitHub (2018)". Accessed: 2019-11-25
56. Cramer, R., Damgård, I., and Nielsen, J. B., Secure Multiparty Computation and Secret Sharing. Cambridge University Press, 2015.
57. Damgård, I., Pastro, V., Smart, N. P. et al., "Multiparty computation from somewhat homomorphic encryption," in Advances in Cryptology - CRYPTO 2012–32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2012. Proceedings, ser. Lecture Notes in Computer Science, Safavi-Naini, R. and Canetti, R., Eds., vol. 7417. Springer, 2012, pp. 643–662.
58. Damgård, I., Keller, M., Larraia, E. et al., "Practical covertly secure MPC for dishonest majority - or: Breaking the SPDZ limits," in Computer Security - ESORICS 2013–18th European Symposium on Research in Computer Security, Egham, UK, September 9–13, 2013. Proceedings, ser. Lecture Notes in Computer Science, Crampton, J., Jajodia, S., and Mayes, K., Eds., vol. 8134. Springer, 2013, pp. 1–18.
59. Katz, J. and Lindell, Y., "Introduction to Modern Cryptography," Second Edition. CRC Press, 2014.
60. Gentry, C., "Fully homomorphic encryption using ideal lattices," in Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009, Mitzenmacher, M., Ed. ACM, 2009, pp. 169–178.
61. Veugen, T., de Haan, R., Cramer, R. et al., "A framework for secure computations with two non-colluding servers and multiple clients, applied to recommendations," IEEE Trans. Information Forensics and Security, vol. 10, no. 3, pp. 445–457, 2015.
62. Sighem, A., Boender, S., Wit, F. et al., "HIV monitoring report 2017," Stichting HIV Monitoring (SHM), Tech. Rep., 2017.
63. Keller, M., Pastro, V., and Rotaru, D., "Overdrive: Making SPDZ great again," in Advances in Cryptology - EUROCRYPT 2018–37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III, ser. Lecture Notes in Computer Science, Nielsen, J. B. and Rijmen, V., Eds., vol. 10822. Springer, 2018, pp. 158–189.
64. Shafer, R. W., "Rationale and uses of a public HIV drug-resistance database," The Journal of infectious diseases, vol. 194, no. Supplement 1, pp. S51–S58, 2006.
65. Pönisch, W., Mitrou, P. S., Merkle, K. et al., "Treatment of bendamustine and prednisone in patients with newly diagnosed multiple myeloma results in superior complete response rate, prolonged time to treatment failure and improved quality of life compared to treatment with melphalan and prednisone—a randomized phase iii study of the east German study group of hematology and oncology (OSHO)," Journal of cancer research and clinical oncology, vol. 132, no. 4, pp. 205–212, 2006.
66. Hicks, C. B., Cahn, P., Cooper, D. A. et al., "Durable efficacy of tipranavir-ritonavir in combination with an optimised background regimen of antiretroviral drugs for treatment-experienced HIV-1-infectedpatients at 48 weeks in the randomized evaluation

of strategic intervention in multi-drug resistant patients with tipranavir (resist) studies: An analysis of combined data from two randomised open-label trials," The Lancet, vol. 368, no. 9534, pp. 466–475, 2006.