

COPING WITH THE GENERAL DATA PROTECTION REGULATION; ANONYMIZATION THROUGH MULTI-PARTY COMPUTATION TECHNOLOGY

Wouter van Haaften / Alex Sangers / Tom van Engers / Somayeh Djafari

Ph.D. researcher, University of Amsterdam, Leibniz Institute Nieuwe Achtergracht 166, 1018 WV, Amsterdam, The Netherlands vanHaaften@uva.nl; <http://www.leibnizcenter.org>

Senior researcher, TNO, Leibniz Institute New Babylon, Anna van Buurenplein 1, 2595 DA The Hague, The Netherlands Alex.sangers@tno.nl; <http://www.tno.nl>

Professor, University of Amsterdam, TNO, Leibniz Institute Nieuwe Achtergracht 166, 1018 WV, Amsterdam, The Netherlands vanengers@uva.nl; <http://www.leibnizcenter.org>

Ph.D. researcher, TNO, Leiden University New Babylon, Anna van Buurenplein 1, 2595 DA The Hague, The Netherlands somayeh.djafari@tno.nl; <http://www.tno.nl>

Keywords: *General Data Protection Regulation, Anonymisation, Privacy by Design, Multi-Party Computation*

Abstract: *Analysing combined data sets can result in significant added value for many organisations, but the GDPR has put strict constraints on processing personal data. Anonymization by using Multi-Party Computation (MPC) however may offer organizations some relief of the perceived burden of GDPR under specific conditions. In this paper, we will explain the mechanisms behind this technology and illustrate its use by a health care case where medical data have to be combined for creating a prediction model, without revealing any sensitive personal data. We will argue why the use of this type of MPC would allow us to anonymize the highly sensitive personal data within the specific boundaries of the case and conclude our paper with some reflection on MPC in the context of the GDPR.*

1. Introduction

The GDPR aims to protect personal data from abuse of these data thus violating the fundamental right to protection of the private and family life, the home and the correspondence of individual subjects¹. In particular the GDPR forbids the processing of sensitive personal data, except under certain conditions. In order to be able to process sensitive personal data, like health data, for research purposes it would be convenient to anonymise these data first, so that all risks of unlawful processing can be eliminated.

In this paper we will explain MPC, a set of *cryptographic* techniques that enables data analytics to be applied without the need to share the underlying personal data; in fact, MPC is an interpretation of the concept of Privacy by Design² when processing sensitive data. Based on a use case, we will present a legal analysis of the use of this anonymisation technique within the framework of the GDPR. Our claim is that MPC enables data controllers to transfer anonymized data to other processors. Subsequently, these processors can analyse that data, while it is technically and organizationally guaranteed that the data cannot be de-anonymised. The paper ends with some conclusions and considerations regarding GDPR.

¹ European Convention on Human Rights Article 8.1.

² Art. 25 GDPR.

2. Pilot

The BigMedilytics project aims to enhance patient outcomes and increase productivity in the health sector by applying big data technologies to complex datasets while ensuring security and privacy of personal data. In the alleged pilots, The Netherlands Organisation for applied scientific research (TNO), health insurer Zilveren Kruis and hospital Erasmus MC cooperate in a pilot to improve health care for heart failure patients. The independent third party ZorgTTP is also included in this pilot and acts as a computing party. In this pilot, TNO developed proof-of-concept software that allows Erasmus MC and Zilveren Kruis to train a prediction model on their combined data, without actually sharing these datasets. The scope of the pilot is restricted to a proof-of-concept of software that works in the operational environment of the parties involved by using synthetic data. This means that no personal data are processed altogether and the GDPR is not applicable at any moment during the pilot. The challenge of the pilot is to show that the developed software can be qualified as anonymisation technique within the framework of the GDPR. In the BigMedilytics case this may be hard to imagine, coming from a fully legitimate processing of the original personal datasets by the controllers Erasmus MC and Zilveren Kruis. Nevertheless, we will show that no natural persons can be identified from the data at the end of the process.

3. Legal framework

The legal context of MPC has two components, the application of the MPC technique and the BigMedilytics pilot with three parties cooperating in a certain way in order to avoid the abuse of personal data. The cooperating parties Erasmus MC and Zilveren Kruis have personal data of their own, as well as a legitimate purpose and a legal basis for processing those data. This means that from the perspective of the pilot the data are personal and are being pseudonymised within the MPC process.

Although the technical process leads to a dataset which is irreversibly non identifiable, one could state that the organisational measures to be taken cannot entirely exclude identification of natural persons when the parties don't follow the right procedure or even collude. The first obviously is a much higher risks, as parties could potentially fake to have applied MPC in order to avoid enforcement of GDPR violations, while unlawful data processing by data controllers will always be possible.

3.1. From pseudonymisation to anonymisation

When looking for a legal qualification within the GDPR regarding techniques like MPC the big question is whether MPC is truly anonymising or if it is merely pseudonymising. What MPC shows close resemblance to the definition of pseudonymisation³ in the GDPR. So, it is likely the data processing within the BigMedilytics pilot will at least qualify as pseudonymisation.

In that respect it is an instrument of Privacy by Design⁴, [...] implement appropriate technical and organisational measures, such as pseudonymisation, which are designed to implement data-protection principles, such as data minimisation, in an effective manner [...] and security⁵, [...] the controller and the processor shall implement appropriate technical and organisational measures to ensure a level of security appropriate to the risk.[...] Since MPC can achieve all the necessary requirements for pseudonymisation, «as will become apparent from the systems» description, anonymisation remains the challenge. Not in order to be able to get

³ GDPR Article 4.5) «pseudonymisation» means the processing of personal data in such a manner that the personal data can no longer be attributed to a specific data subject without the use of additional information, provided that such additional information is kept separately and is subject to technical and organisational measures to ensure that the personal data are not attributed to an identified or identifiable natural person.

⁴ Article 25 GDPR.

⁵ Article 32 GDPR.

careless with the data. To keep data anonymous strict and extensive security measures will be required. But more from an administrative burden relief point of view, e.g. less information will have to be saved and stored for the purpose of informing the subject.

So, what circumstances would be necessary to reach anonymity status? The crucial part of the definition of pseudonymisation in this respect is in the wording [...] .. technical and organisational measures to ensure that the personal data **are not** attributed to an identified or identifiable natural person;... If it is possible to go one step further in the sense that the personal data **cannot, or can no longer**, be attributed to an identified or identifiable person, then the «anonymisation» will be achieved.

3.2. Anonymisation

The process of anonymising data means that all identifying elements are eliminated from a set of personal data so that the data subject is no longer identifiable.⁶ In its Opinion 05/2014, the Article 29 Working Party analyses the effectiveness and limits of different anonymisation techniques.⁷ It states:

«An effective anonymisation solution prevents all parties from singling out an individual in a dataset, from linking two records within a dataset (or between two separate datasets) and from inferring any information in such dataset. Generally speaking, therefore, removing directly identifying elements in itself is not enough to ensure that identification of the data subject is no longer possible. It will often be necessary to take additional measures to prevent identification, once again depending on the context and purposes of the processing for which the anonymised data are intended.»⁸

In the Opinion 5/2014 on anonymisation techniques some clear boundaries are set regarding the anonymisation phenomenon. The WP29 formulated three basic checkpoints for a robust anonymisation technique based on three criteria, captured in three questions that will have to be answered negatively in order to *pass the anonymity test*:

- (i) is it still possible to single out an individual,
- (ii) Is it still possible to link records relating to an individual, and
- (iii) can information be inferred concerning an individual?⁹

These checkpoints will be leading when scrutinizing the MPC solution in the BigMedilytics pilot for anonymity. Also according to the Opinion the original legal basis for processing the personal data concerning health can cover the data processing towards anonymisation. And thus, it can also cover for performing the anonymization process within the framework of BigMedilytics pilot. This means that the parties involved, the hospital, the health insurance company and the third party all have to comply with the complete set of rules from the GDPR before anonymisation. They will have to take both technical and organizational measures to guarantee the confidentiality of the personal data as long as they qualify as such. When the personal data have been successfully anonymised, they are no longer personal data and data protection regulation will no longer apply. We will come back to this legal basis in the next paragraph.

Now the question arises at which particular moment the medical data are being anonymized. In order to find out we followed the technical steps one by one, each time looking at the process and the outcome in order to establish whether the technical status «anonymized» has been reached or not. In the legal considerations in between the description of the steps of the MPC protocol we have pointed out the legal implications of those steps and the specific character that actually leads to anonymisation. This anonymisation ability puts the MPC protocol in a good position to support research on combined datasets that basically are to be qualified as per-

⁶ Recital 26 GDPR.

⁷ Opinion 5/2014 on Anonymization.

⁸ Opinion 5/2014 on Anonymization, page 9.

⁹ Opinion 5/2014 on Anonymization, page 3.

sonal data. In section 4 we describe the proof-of-concept of software developed by TNO and indicate when the data protection regulation no longer applies.

The anonymity capabilities of MPC will have to be found in the intelligent stacking of anonymisation techniques to one that will eventually provide the anonymisation that is desired. But getting there starts with a legal basis, since the input data for MPC will still be personal data.

3.3. Legal basis for anonymisation

In order to process the personal data to the level of anonymised data a legal basis will be required. The GDPR distinguishes, apart from the standard personal data (Art. 4(1)), also special categories of personal data which, by their nature, may pose an extended risk to the data subjects when processed and therefore need enhanced protection. Such data are subject to a prohibition of processing (Article 9.1. GDPR) and there is a limited number of conditions under which such processing is allowed. *Lawful processing*, however, also requires a legal basis in itself (Art 6.1).

An example of special categories of personal data is personal data concerning health. In article 9.2 GDPR an exhaustive list of exemptions to the prohibition can be found. Application of an exemption is a condition for the processing of sensitive data. In the Bigmedilitics use case the most relevant exemptions for processing personal data concerning health include situations where processing is necessary: (i) «for the purposes of preventive or occupational medicine, for the assessment of the working capacity of the employee, medical diagnosis, the provision of health or social care or treatment or the management of health or social care systems and services on the basis of Union or Member State law or pursuant to contract with a health professional»¹⁰ or (ii) «for public interest reasons in the area of public health»¹¹. In this case (i) is the most appropriate exemption rule since the data will eventually be used for occupational medicine.

Pre-condition for the application of Article 9.2.h is that there is a legal basis for processing of the data, i.e. the data are '...processed by or under the responsibility of a professional subject to the obligation of professional secrecy¹² under Union or Member State law'(article 9, paragraph 3 GDPR). According to Dutch law¹³ both the hospital and the health insurance company are to be considered as «[...] another person also subject to an obligation of secrecy under Union or Member State law or rules established by competent national bodies». Article 9.2 sector h¹⁴, in conjunction with article 9.3 will provide sufficient legal basis for the original treatment.

4. Application of GDPR on the MPC -protocol in the BigMedilitics pilot

As explained above TNO has developed a proof-of-concept software that allows Erasmus MC and Zilveren Kruis to train a prediction model, without sharing datasets. Under GDPR, both parties can be considered as data controllers. Besides these parties, ZorgTTP was introduced as an independent third party to act as a computing party (MPC node) whose role is to facilitate fast computations in the encrypted domain.

The objective of the three parties is to comply with the GDPR, therefore a precondition for the operation is that no other data exchanges will take place but the data processing within the MPC protocol. It is also es-

¹⁰ Article 9, paragraph 2, under h GDPR.

¹¹ Article 9, paragraph 2, under j GDPR.

¹² Art. 88 Wet op de beroepen in de individuele gezondheidszorg. Art 7–457 BW, Wet inzake de geneeskundige behandelingsovereenkomst.

¹³ Parliamentary Document II 1997/98, 25 892, nr. 3, p. 114.

¹⁴ In case the initial processing is based on Article 9.2. of the GDPR, the subsequent anonymization for research purposes must have been communicated when granting permission. (Official investigation by the CBP into the processing of geolocation data by TomTom N.V. PUBLIC VERSION 20 December 2011).

sential in the context of the GDPR that the data will not be fed back to the original patients' files other than as generic research results, i.e. abstracted from individual patients' data.

4.1. Assumptions

The cryptographic security model is the passive setting (semi-honest). This means the participating parties in the protocol follow the exact prespecified protocol. This security model is chosen since the parties have mutual interest, can easily use organizational measures to decrease the risks and the passive security model allows for faster computations.¹⁵ The assumption is that Erasmus MC, Zilveren Kruis and ZorgTTP have an agreement to follow the exact prespecified protocol.

4.2. IT set up in the BigMedilytics pilot

The solution sets up secure communication channels between the three parties involved using the HTTPS¹⁶ protocol. Each party hosts a computer that is externally reachable via HTTPS. On each computer, three Docker containers will be installed:

- Docker container with HTTPS server functionality
- Docker container with the MPC module
- Docker container with HTTPS client functionality

The HTTPS server receives incoming HTTPS communication and forwards this to the MPC module. The MPC module listens to the HTTPS server, does some computations and as soon as communication to other external hosts is required, the MPC module activates the HTTPS client to set up a connection to the HTTPS server of the corresponding external host.

4.3. The MPC process steps in the BigMedilytics pilot

In this paragraph, we describe the technical steps in the MPC-protocol developed in BigMedilytics use case. At every step, the status of the data in terms of anonymization is briefly explained.

Step 1. Preparing the data

Firstly, both data controllers prepare the data for further processing. This involves: (i) collecting the relevant data, (ii) filtering persons from the data that have missing attribute values and (iii) arbitrarily (randomly) sorting the persons in the data. It should be noted that the data in this form are not shared with any other party. In the tables the results of the processes are illustrated.

Identifier	Attribute 1	Attribute 2
234567	1.2	3.2
123456	4.1	7.2
345678	2.1	4.1

Table 1: An example of the data of a data controller. The identifier is directly traceable to an individual and is agreed upon and present in the data of both data controllers.

¹⁵ Y. AUMANN & Y. LINDELL. «Security against covert adversaries». Tcc 2007.

¹⁶ Hypertext T LAUSLAHTI/MATTILA/HUKKINEN/SEPPÄLÄ Transfer Protocol Secure (HTTPS) is an extension of the Hypertext Transfer Protocol (HTTP). It is used for secure communication over a computer network, and is widely used on the Internet. (Source: Wikipedia)..

Step 2. Keyed-hashing of the identifiers

Both data file controllers hash the identifiers in the data using *keyed-hashing* with the same key. The data controllers construct this key together by both generating a large random number and adding these numbers, without revealing it to the independent third party. The keyed-hashing results in two sets of hashed identifiers in which similar identifiers will be similarly hashed. But it is impossible to go back to the original identifier. It is only possible to compare the data records that belong to similar hashed identifiers. When the keyed-hashing of the identifiers is finished, the key and any derivatives of the key are removed, so that the hashing cannot be reperformed. Again, it should be noted that the data in this form are not shared with any other party.

Hashed identifier	Attribute 1	Attribute 2
2a4[...]e2d	1.2	3.2
49a[...]5fe	4.1	7.2
957[...]2ae	2.1	4.1

Table 2: An example of the data of a data controller after hashing the identifiers.

The legal implication of these first two steps is that identification has been seriously complicated, but yet is not altogether impossible. One could say that, although very unlikely in a setting of parties that have a crucial interest in anonymity, the data may still be personal because of the possible reversibility using the attribute values.

Step 3. Additive homomorphic encryption¹⁷ of the attribute values.

All attribute values (not being identifier values) are encrypted using so-called additive homomorphic encryption. *Additive homomorphic encryption* is an encryption scheme that allows additions and subtractions in the encrypted domain. That is, given encryption of m_1 and m_2 , one can compute the encryption of $m_1 + m_2$. Firstly, both data controllers generate a key pair with a public and a private key and they both use their own public key to encrypt the attribute values of their data (excluding the hashed identifiers). The two resulting datasets with hashed identifiers and encrypted attribute values will be called the scrambled datasets.

Hashed identifier	Enc(Attribute 1)	Enc(Attribute 2)
2a4[...]e2d	7518[...]3424321	7385[...]34587
49a[...]5fe	7654[...]5738934	0598[...]54843
957[...]2ae	3541[...]2830472	1474[...]74933

Table 3: An example of the scrambled data of a data controller.

In this case the data are being encrypted, and the question is whether the scrambled datasets that are a result of step 2 can be reversed. Since the data controllers can decrypt their own encrypted data, the data is still personal to the controllers which hold the private key of their own encrypted data due to the possible reversibility of the encryption. It is not personal data to any other party that does not hold the private key to that data, since they cannot decrypt the data without having access to the private key. The party having the private key is not «reasonably likely»¹⁸ to collaborate with the other party against their agreement not to share data, regarding their eventual research interests.

¹⁷ PAILLIER, PASCAL (1999). «Public-Key Cryptosystems Based on Composite Degree Residuosity Classes». EUROCRYPT. Springer. pp. 223–238. doi:10.1007/3-540-48910-X16.

¹⁸ Recital 26 of the GDPR: «[...] To determine whether a natural person is identifiable, account should be taken of all the means reasonably likely to be used, such as singling out, either by the controller or by another person to identify the natural person directly or indirectly. To ascertain whether means are reasonably likely to be used to identify the natural person, account should be taken of

Step 4. Share the scrambled dataset with an independent third party

The scrambled datasets can now be shared with an independent third party. The third party only learns how large the datasets are, but it cannot learn any identifier or attribute value in the scrambled datasets. However, it can still find hashed identifiers that are identical in both datasets. The third party finds the encrypted attribute values of the matching hashed identifiers, stores them and returns merely the number of matched hashed identifiers to the data controllers.

Hashed identifier	Enc(Attribute 1)	Enc(Attribute 2)
2a4[...]e2d	7518[...]3424321	7385[...]34587
49a[...]5fe	7654[...]5738934	0598[...]54843

Table 4: An example of the stored encrypted dataset of the matching hashed identifiers. Note that the Hashed identifier column will never be stored, but is only shown here for explanatory purposes.

This step makes the process irreversible. By introducing an independent third party not only the technical anonymity is fulfilled, but also the organizational aspects of the data processing exclude the processing of personal data by any of the three organizations involved in this stage.

Step 5. Generating random shares

Both data controllers generate a number of sufficiently large random numbers, as many as the number of matched hashed identifiers, times the number of attributes of the other data controller. Each data controller stores the generated random numbers, which we will refer to as random share. Next, both data controllers encrypt these random shares with the public key of the other data controller (generated in step 3). These encrypted random shares are shared with the independent third party.

Random (unencrypted) shares for attribute 3
16354728.2
964543782.3

Table 5: An example of the generated random shares of a data controller. The size of this example table is 2 times 1: the number of matched hashed identifiers times the number of attributes of the other data controller.

This step doesn't change the anonymous status of the data.

Step 6. Computing with the encrypted data

The third party utilizes the homomorphic encryption property. It takes the stored encrypted attribute values of Erasmus MC and subtracts the encrypted random shares of Zilveren Kruis. The resulting encrypted shares are forwarded to Erasmus MC. This party is able to decrypt these encrypted shares with its private key. This process is repeated for Zilveren Kruis. The result is that the data controllers each have an additive secret share of data that corresponds to the intersection of both data sets. This means that in case both data controllers would add up their secret shared tables, the attribute values of the set intersection could be revealed (which is not done of course, since the agreement prescribes not to reveal any personal data). Next, since the private keys do not have a function anymore, the private keys of the data controllers are removed.

all objective factors, such as the costs of and the amount of time required for identification, taking into consideration the available technology at the time of the processing and technological developments [...].»

Attribute 1	Attribute 2	Attribute 3
87785483.9	5789471.2	-16354726.5
84053906.6	6769073.0	-964543780.3

Table 6: An example of the shares of the data controller.

Attribute 1	Attribute 2	Attribute 3
-87785482.7	-5789468.0	16354728.2
-84053902.5	-6769065.8	964543782.3

Table 7: An example of the shares of the other data controller

Note that both data controllers do not know anything about which identifiers or attribute values are in this data (since it is secret shared), and neither does the third party, since this party is unable to decrypt any data.

In this stage the third party takes the stored encrypted attribute values of Erasmus MC and subtracts the encrypted random shares of Zilveren Kruis, and forwards the result to Erasmus MC. And vice versa. Both data controllers do not know anything about which identifiers or attribute values are in this data (since it is secret shared), and neither does the third party, since this party is unable to decrypt any data. So, the data can still be considered anonymous.

Step 7. Transforming the two additive secret shares into three Shamir secret shares

Each data controller divides the additive shares into three parts based on Shamir secret sharing. Both data controllers divide the three parts among the three parties (keep one to themselves) and all parties add up the two Shamir shares belonging to the additive shares. Each party now has a Shamir secret share with threshold 2, meaning that any two parties together are able to unlock the secret using polynomial interpolation.

The still anonymized (secret shared) data set can now be analysed for medical research purposes.

Step 8. Performing secure lasso regression

What rests is an anonymized (secret shared) dataset that can be analysed for medical research purposes, using any library that supports Shamir secret sharing with three parties¹⁹. For example, the parties cooperatively perform a lasso regression analysis, which finds a linear relation between the number of hospitalization days, and explanatory attributes. Only the number of iterations of the optimization algorithm and the results of the analysis are revealed during this MPC process step. The results of this analysis are the regression coefficients for each explanatory attribute and a measure of fit, which eventually can be used for an individual medical risk assessment. These results are an aggregate of all patients’ data that were used.

The results of the secure lasso regression cannot be restored into identifiable data. Patients whose data were used or any other patients can only profit from the trained prediction model via the application of the generic research results.

From the previous steps it can be deduced that the ultimate running of the data through the MPC process does imply anonymization of the underlying data.

¹⁹ Note that this secret sharing protocol requires three or more parties.

5. Conclusions and considerations

In this paper we presented the eight steps of the MPC protocol as developed in the BigMedilytics heart failure pilot resulting in a fully anonymised data set ready for further analysis. Because of the generic use of MPC, it is possible that it may result in anonymised data in other applications as well. Besides applications within the health domain one could think of applications in e.g. law enforcement, combining data of different governmental agencies that want to better understand what circumstances will increase the likelihood of non-compliant behaviour, or within the banking sector looking for indicators for financial fraud, just to mention a few examples. Obviously, using MPC means that once certain patterns are discovered one could apply these results for case assessment, i.e. making decisions on individual cases, within the framework of the legal basis that applies in the specific situation. Applying the results of a prediction model to individual cases concerns the processing of personal data (obviously not anonymized) and will, off course, need a legal basis. Even in case of data that is not strictly personal data, MPC may have useful applications, for example for organisations that not have enough data of their own to have a good basis for data analytical approaches, but would if they combined their data with the data of similar, even competing, organisations. MPC could allow them to collaboratively analyse their data without sharing datasets.

The application of AI-technologies has great potential in various domains including the medical domain. Particularly data analytics using machine learning has already contributed to a much better understanding of complex constellations of factors that impact the development of diseases and effective treatments. The use and combination of various data sources is essential for being able to develop models that enable us to find relations between these factors or to predict the most effective treatment to a particular disease. Since that data typically concerns highly sensitive personal data the processing thereof should be handled with greatest precautions. Many argue that the GDPR comes with huge hurdles that prevent researchers in the medical domain to combine and analyse data, because of the prohibition to process sensitive personal data. However, the GDPR allows for anonymizing data from different sources, even in case data is highly sensitive personal data, as long as the processing of that data leads to irreversible anonymity.

With those anonymous data processes can be performed leading to e.g. finding better remedies against diseases. MPC is a combination of technologies enables organizations to collaboratively analyse data without sharing it. Parties that need to cooperate with others, sharing data while optimizing data protection and minimizing data infringement risks, can use MPC to produce anonymised datasets. Those sets preserve the mathematical characteristics necessary for being able to do the analysis, while completely hiding identification of the data subjects as well as original data attribute values that could be used for identifying data subjects by either de-anonymisation by reversed coding or singling out mechanisms.

MPC is a powerful data anonymisation technique but MPC is not the only technology that can be used to allow for connecting data without sharing. The federated learning technology, as developed in the Personal Health Train project (see <https://www.dtls.nl/fair-data/personal-health-train/>), and the Digital Market Place technologies (see Deljoo et.al. 2018) are alternative technologies that are currently being further developed and tested in a wide variety of application domains, including logistics in the DL4LD project, the health domain in the VWData and EPI projects and the banking sector in the SPPDDP project. The protection of the rights of data subjects is essential in our society in which data and data processing have become valuable assets driving the (data) economy.

The protection via legislation such as the GDPR and appropriate privacy enhancing technologies should go hand in hand. Technologies such as MPC help us to benefit from the tremendously increased analytical power of data-driven technologies without data subjects having to fear infringements of their fundamental rights.

Acknowledgements

Thanks to Achmea Zilveren Kruis, ZorgTTP and Erasmus MC for discussions on the use case and the technical requirements.

The BigMedilytics project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 780495.

Thanks to NWO for sponsoring and the colleagues of the UvA SNE team and all other partners for collaborating in the DL4LD, VWData and EPI projects.

References

- DELJOO, A.; VAN ENGERS, T.; VAN DOESBURG, R.; GOMMANS, L. and DE LAAT, C. (2018). A Normative Agent-based Model for Sharing Data in Secure Trustworthy Digital Market Places. In Proceedings of the 10th International Conference on Agents and Artificial Intelligence – Volume 2: ICAART, ISBN 978-989-758-275-2, pages 290–296. DOI: 10.5220/0006661602900296
- A DELJOO, T VAN ENGERS, R KONING, L GOMMANS, C DE LAAT, 2018, Towards trustworthy information sharing by creating cyber security alliances, 17th IEEE International Conference On Trust, Security And Privacy In, 2018
- A. SHAMIR, How to share a secret, *Communications of the ACM* **22** (1979), 612–613.
- Y. AUMANN & Y. LINDELL. «Security against covert adversaries». Tcc 2007.
- PAILLIER, PASCAL (1999). «Public-Key Cryptosystems Based on Composite Degree Residuosity classes». EU-ROCRYPT. Springer. pp. 223–238. doi:10.1007/3-540-48910-X_16