



HAL
open science

Agents in a privacy-preserving world

Edgar Galván, Joaquin Garcia-alfaro, Vicenç Torra, Guillermo Navarro-Arribas

► **To cite this version:**

Edgar Galván, Joaquin Garcia-alfaro, Vicenç Torra, Guillermo Navarro-Arribas. Agents in a privacy-preserving world. Transactions on Data Privacy, 2021, 14 (1), pp.53-63. hal-04337700

HAL Id: hal-04337700

<https://hal.science/hal-04337700v1>

Submitted on 12 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Agents in a privacy-preserving world

Edgar Galván^{*}, Joaquin Garcia-Alfaro^{**}, Guillermo Navarro-Arribas^{***},
Vicenç Torra^{****}

^{*}Naturally Inspired Computation Research Group, Department of Computer Science, Maynooth University, Maynooth, Ireland.

^{**}Télécom SudParis—Institut Polytechnique de Paris, 91128 Palaiseau.

^{***}Dept. Information and Communications Engineering – CYBERCAT, Universitat Autònoma de Barcelona, Bellaterra, Catalonia, Spain.

^{****}Department of Computing Science, Umeå University, Umeå, Sweden.

E-mail: edgar.galvan@mu.ie, jgalfaro@ieee.org,
guillermo.navarro@uab.cat, vtorra@ieee.org

Received 15 January 2021; received in revised form 23 April 2021; accepted 24 April 2021

Abstract. Privacy is a fluid concept. It is both difficult to define and difficult to achieve. The large amounts of data currently available at hands of companies and administrations increase individual concerns on what is yet to be known about us. For the sake of penalisation and customisation, we often need to give up and supply information that we consider sensitive and private. Other sensitive information is inferred from information that seems harmless.

Even when we explicitly require privacy and anonymity, profiling and device fingerprinting may disclose information about us leading to reidentification. Mobile devices and the internet of things make keeping our live private still more difficult.

Agent technologies can play a fundamental role to provide privacy-aware solutions. Agents are inherently suitable in the heterogeneous environment in which our devices work, and we can delegate to them the task of protecting our privacy. Agents should be able to reason about our privacy requirements, and may collaborate (or not) with other agents to help us to achieve our privacy goals. We are presented in the connected world with multiple interests, profiles, and also through multiple agentified devices. We envision our agentified devices to collaborate among themselves and with other devices so that our privacy preferences are satisfied. We believe that this is an overlooked field. Our work intends to start shedding some light on the topic by outlining the requirements and challenges where agent technologies can provide a decisive role.

Keywords. Data privacy, user privacy, privacy enhancing technologies, multiagent systems.

1 Introduction

Privacy has become a major concern for companies and individuals. The General Data Protection Regulation (GDPR) in Europe, and similar regulations in the rest of the world, have obliged companies to increase data privacy standards to avoid privacy leakages. Companies need to implement data privacy technologies to avoid disclosure to grant permissions

to data access (e.g., access control [17]), when they share data (e.g., masking methods), as well as when they develop data-driven solutions (e.g., privacy-preserving machine learning to avoid membership attacks).

GDPR as well as other regulations impose substantial fines if sensitive information is leaked. The regulation and increasing public interest has caused a boom in data privacy technologies. Concepts as deidentification, k-anonymity [8, 33, 37], differential privacy [15, 16] are now well-known beyond the data privacy specialists.

Data privacy solutions [14, 19, 39, 41] implemented by companies and organisations mainly assume that data is centralised. Alternatively, some solutions focus on distributed settings but in this case it is usually assumed that the number of databases is reduced. This corresponds to the case of a few companies or organisations sharing their data or agreeing to do a common analysis of their data. Relatively recently, federated learning [26, 44] has emerged as an alternative when a centralised entity wants to compute a data-driven model with data from multiple devices. Federated learning, however, does not avoid privacy leakages.

Individuals' data is distributed and at hands of companies and organisations. Then, individuals need to trust these companies and organisations and expect that they will take care of their data, a difficult endeavour to achieve [28]. Even in the case that their practice is according to law and regulations and privacy by design principles are thoroughly observed (specially the data minimisation requirement), companies and organisations have as their main goal the safeguard of their own interests. For example, they are interested in their intellectual property rights (IPR) and in keeping private any confidential information related to their own practices and protocols. Some initiatives are being developed to provide users a centralised access to their data for any possible data transaction [38]. Our position is that effective protection of individuals' sensitive information needs to be implemented on the individuals' end. Multiagent systems can empower citizens and give them back control of their own data.

GDPR by means of the rights of access, rectification, and erasure, allow some control by the user of this data. GDPR enforcement can benefit from a posteriori actions, specially in situations where sanctions must be imposed to complement accountability measures. For example, situations in which permanent connectivity cannot be guaranteed such as delay-tolerant network and spontaneous communication scenarios, and where privacy violations can be discovered later on, during an audit of monitored logs.

In this paper we discuss the problem of user privacy. This is the case in which the users themselves have full control of their sensitive information. Due to the amount of information out there, and to the fact that our daily activity is tracked, most aware users have difficulties with making learned decisions on which are the best strategies to follow. Multiagent technologies can be very useful for user privacy and it is only by means of sophisticated agent technologies [43] that appropriate decisions can be made. That is, by means of implementing and using privacy agents with full-fledged deductive capabilities, incorporating state-of-the-art machine learning approaches, and multiagent systems designed to incorporate agreement technologies (negotiation and argumentation abilities).

The structure of the paper is as follows. Section 2 provides preliminaries. Section 3 addresses the use of agent technologies for user privacy. Section 4 concludes the paper.

2 Preliminaries

In this section we discuss several issues related to data privacy and disclosure risk that will lay the foundations of our approach detailed in this work.

2.1 Privacy dimensions and user privacy

There are different ways to classify privacy enhancing technologies. One of the classifications is taking into account whose privacy is being sought. We consider (i) respondent privacy, (ii) holder privacy, and (iii) user privacy [12, 39]. Respondents are the individuals that provide the data to an organisation. They correspond to data subjects providing personal data according to GDPR. Data holders are the ones that store and process the data. They correspond to data controller and data processors according to the GDPR.

Respondent privacy and holder privacy have different focus. Individual's privacy preferences are different from organisations keeping data. Organisations are obliged to protect sensitive information about data subjects, but their deepest privacy requirements are more focused on intellectual property rights, and sensitive information about the company that can give economical advantage to competitors (e.g., market information about clients and products). Nevertheless, both respondent privacy and holder privacy shall be implemented by the data holder.

User privacy is about the privacy of a subject. The difference between respondent and user privacy is that subjects have a passive role in the former case and have an active role in the latter case. That is, privacy technologies are implemented by the data holder to ensure respondent privacy, while privacy technologies are implemented by the user themselves to ensure user privacy (e.g., through encryption for private mail transmission).

In fact, user privacy can go beyond what these tools provide. For example, users benefit from tools that permit private distributed computing as e.g. secure multiparty computation (as in [7]).

2.2 Communications and fingerprinting

Privacy in communications has been studied in the context of internet, providing different mechanisms. Widely known is the use of onion routing for anonymous communications (e.g. The Tor Project [11]), or mix networks as routing or message shufflers. These tools, among others, can provide end-to-end anonymity in communications over computer networks and are usually combined with profiling and fingerprinting prevention mechanisms.

It is well-known that privacy is related to our identification in online services. Nevertheless, the use of pseudonyms and other kind of techniques based on hiding our own identity is usually not enough. There are quite a few ways so that business and services providers profile and identify users. Device fingerprints is one of them. That is, the information on the software and hardware of the clients when connected to e.g. a web service. Device fingerprints [2] make most devices unique. Therefore, they are a simple way to reidentify and link data from the same person when using different services. Tools exist to avoid device fingerprinting by means of handling different sets of fingerprints [21].

2.3 Machine learning and disclosure

There is a substantial amount of research [1, 25] on the application of machine and statistical learning to infer user private characteristics from the information available in social

networks. Examples include inferences on user's political preferences, sexual orientation, religious views and even parental separation.

Therefore, what users post, the friends they have, and the interests they express in social networks can make others learn about their private life.

Similar studies have been done on information from different types of devices, including smartphones and IoT devices [34, 40].

2.4 User's privacy in information retrieval

Private information retrieval (PIR) and user privacy in information retrieval provide tools for avoiding to share the identity of the users as well as their queries. Information-theoretic PIR and computational PIR provide the strongest guarantees to avoid the disclosure of the query. This is at the expense of high computational costs and strong requirements on the server side, which makes real implementation difficult. Other approaches exist as well, with smaller privacy requirements and making collaboration with the server unnecessary. These approaches are based on perturbation of queries and on dissociating multiple facets/queries of the same individual. The latter is based on the observation that a person has multiple interests and is the combination of these multiple interests (facets) that makes identification possible. Privacy agents that act on behalf of users to implement different policies for avoiding query disclosure already exist. Additional collaboration among agents would permit to design multiple communication spaces [35] and avoid disclosure of the user identity. This implies defining anonymity sets of users, and agents forwarding queries of other agents within an anonymity set.

2.5 Federated learning

Data-driven models are data intensive. Most machine and statistical learning methods are developed to be applied to centralised databases. In the last years there has been a trend to consider decentralised settings, which are more privacy friendly.

Federated learning [45] considers highly decentralised environments. In a nutshell, data from devices is kept locally. A central authority computes a machine learning model by means of interactions with the devices, but local data is never sent to this central authority. Instead, the central model is shared with the devices, and the devices send differences between their local models and the central model. Transmitting differences are to avoid the sharing of sensitive information.

Leakage of sensitive information can take place because the final model can lead to membership inference [27]. Injection attacks show the difficulties of building models when individual agents intend to fool the protocols and influence the outcome beyond what is necessary. At the same time, privacy standards of federated learning are far from clear. A major concern is that differences between models and local data can lead to disclosure of sensitive information. Research on differential private machine learning models for federated learning tries to solve some of these difficulties [42, 45]. Nevertheless, all information from a single device usually corresponds to a single individual with non-independent records, which makes standard differential privacy protocols unsuitable [13]. In addition, federated learning requires several interactions between the central authority and the devices, which threatens privacy and means that agents either run out of privacy budgets very soon or need to have large privacy budgets. Multiple devices from the same users, make things still more complex. Dependencies between devices will appear and this is not taken into account by current research.

2.6 A Posteriori Access Control

While traditional access control assumes that users must be prevented from gaining access to a resource outside their sphere of access (i.e., to stop unauthorised activity from occurring by blocking it immediately if it is not explicitly allowed in a given policy), the idea of *A Posteriori Access Control* (APAC [3, 10]) can be seen as the combination of monitoring and audit at the same time, i.e., it allows access without imposing prevention constraints; but monitors and audits all the executed actions, in order to apply after-the-fact sanctions (e.g., enforcement of fines due to GDPR infringement), when policy violations are discovered.

APAC can be seen as a specialised version of passive intrusion detection, in which access permissions may evolve over time, fulfilling the objective of not preventing attacks, but rather fixing responsibilities and applying sanctions after a careful analysis of logs. The execution steps of an APAC process also differ from those of passive intrusion detection, since in addition to log processing, normalisation, and analysis, it will also require some data semantic augmentation in order to enforce accountability measures.

3 Agent technologies for user privacy

Respondent and holder privacy is implemented by data holders (data controllers and data processors using GDPR terminology), but in this case individuals need to trust organisations, expect that they apply appropriate technologies, and hope that data management does not lead to unintended leakages.

The only way that individuals may have full control of their data is by implementing themselves user privacy methods. User privacy can empower individuals allowing them to select their own privacy preferences and make their own decisions for attaining appropriate privacy levels. We consider that agent technologies can have a leading role in helping individuals to protect their privacy. That is, in implementing user privacy.

3.1 Challenges to achieve user privacy

We underline some of the major difficulties that multi-agent technologies can help to solve in relation to user privacy.

- Users have access to the connected world through a high variety of devices. For instance, Internet of Things is increasing the number of personal devices connected to the internet. These devices are highly heterogeneous. Any privacy solution needs to take into account this heterogeneity.
 - Each user has access and interacts with the world through multiple devices, instead of a single one. Data protection technologies usually assume that records in a database are independent. It is uncertain that this independence holds under the assumption that each person has a single entry point to internet. E.g., multiple queries to search engines are clearly not independent (recall the AOL case that permitted identification of the user behind a set of queries by the New York Times), multiple contributions to a federated learning algorithm are not independent. Naturally, multiple entry points to internet make this independence assumption more improbable.
 - In a given place such as a household, there are shared devices, which may transmit indistinctly information from several household members.
-

-
- It is difficult for an individual, even a highly motivated one, to be consistent in the use of devices to avoid the disclosure of sensitive information.
 - Private browsing and access to services are expected to provide some privacy guarantees. Nevertheless, devices have fingerprints that can help on reidentification. See Section 2.2.
 - It is also extremely difficult for any individual to know what is inferred from a single online activity, taking into account all previous history. See Section 2.3 on machine learning and disclosure.

Agent technologies can play a fundamental role by means of providing a privacy agent that helps users to ensure a desired privacy level. We can envision,

- a single agent that represents a user in the internet, governing all user devices, and to which the user delegates the task of protecting user's privacy.

3.2 Requirements for privacy agents

To achieve the aforementioned challenges, privacy agents must know the privacy preferences of the user they represent; and should also be aware of all available information about the user. This includes, what information has been transmitted to each service, and what information has been posted in social networks. Agents also will know what can be inferred from these already available information [32, 18], in line with data-driven models extracted from social networks.

Furthermore, agents must be able to act on behalf of users taking into account their preferences. This implies that agents should be able to reason and decide on user's behalf about user's privacy requirements, taking into account user's information, and the services required [9]. They will also take into account e.g. the computation approach (i.e., centralised vs. federated learning) related to a service.

Decision making can take different forms, in general, it may require agents to reach agreements and negotiate [20, 6] so that they can access services or trade. Decisions can be made on different items of interests. They include all kind of decisions with respect to services and trade, but also about deciding to transmit any information, to refrain from transmitting, as well as decisions related to collaborations and agreements with other agents to achieve good privacy levels together. In particular, collaboration with other agents means e.g. sharing information with other agents, building anonymity sets with other agents to hide their own data, mask data in collaboration with other agents, and forward queries to/from other agents to search engines (as discussed in Section 2.4).

During the execution of the agents actions, negotiation activities on behalf of users may benefit from policy refinement approaches that provide semantic enrichment [30, 29]. For instance, use of policy languages enriched with semantics (e.g., to provide augmented privacy concepts). This can be achieved by using, for instance, OWL (Web Ontology Language) and RDF (Resource Description Framework). This will be useful for price negotiation tasks in trading activities (e.g., finding the best price in a negotiation scenario, achieving the *always best deal*) or goods acquisition (e.g., finding the best offer, not necessarily the cheapest one, in a customer-provider scenario). The related literature shows extensive examples of agents negotiating access on behalf of the user, e.g., under the context of mobile phones [5, 4], social media [22, 23, 31, 36], and Internet of Things [24].

Agents could also perform a posteriori actions based on user history, behaviour and preferences. Following the ideas of APAC (cf. Section 2.6), agents will be able to conduct different actions on behalf of the user after their interactions. This can include opt-out actions in internet services or even placing (or advising in) complaints or claims in case a privacy violation is detected.

In addition,

- agents must advise the users on different privacy related questions. Agents will be able to make learned decisions on the transmissions, both with respect to timing and content. Explainable decisions can help users to improve their knowledge;
- agents must be able to deal with multiple profiles, fingerprints, and provide to Internet services appropriate contexts according to user's preferences, user's knowledge and service requirements. Agents need to take into account the case of devices in the same household used by multiple user's;
- agents could provide a mixed local/cloud execution model to suit the user preferences. The main agent core can be externally executed in a trusted cloud provider, but it can also offer a local execution model distributed over the user devices. In this case, there might be some limitations on the autonomy of execution (execute actions asynchronously at any time), computation power (low computation devices) or data storage. If required, a mixed model where the agent executes in a local cloud instance from the user (like a home desktop computer) centralising the computation and data storage for all user instances can also be accommodated.
- agents must provide high-usability user-friendly mechanisms and interfaces to interact with the users regarding the configuration and notification of privacy preferences and information.

Agents themselves can cause privacy leakages and they can be mistrusted by users. Agents will have access to all kinds of user's information and this information can naturally be personal and confidential. Information includes user's preferences. In addition, due to negotiations and agreements, agents may access to personal information supplied by other agents. Agents' design and implementation needs to take into account this fact and apply appropriate measures (e.g., privacy-by-design technology and open source code).

3.3 Agents types of actions

Agent actions with direct implication to the user privacy can be of three different types:

- Preventive: such as configuring the user devices and applications to reflect the user's privacy preferences, installing specific software (e.g. Tor or specific browser extensions), or suggesting actions to the user about how to interact with the devices.
 - Proactive: during the user operations and interactions, the agent actively participates by monitoring and acting to preserve user's privacy. Examples are the generation of noisy traffic or interactions, establishing anonymous connections transparently to the user, blocking actions from applications with potential privacy implications (e.g. sending usage data to the vendor), coordinate with other agents to mask user profiles, etc.
-

-
- A posteriori: after some actions, the agent can detect potential privacy problems, that might require a posteriori actions. These are enforced by a posteriori policies, the analysis of the monitored user actions, and potential cooperation with other agents.

In order to conduct its main actions, the privacy agent might need to conduct additional actions such as autonomous discovery and coordination with other privacy agents for cooperative actions.

3.4 Discussion

A large set of works on both artificial intelligence (AI) and multi-agent systems (some very recent) show that agent technologies may have a leading role in helping individuals to protect their privacy. Moreover, extensive related literature have already settled some fundamental properties and methods to help users manage their privacy based on agent technologies. However, we think that this is an overlooked field, for several reasons.

The most important, lack of acceptance by risk-averse users, especially those aware of privacy implications, who may question the implications for an agent to get granted access to personal data and preferences; as well as trust and securisation issues with respect to the agents. Some insights to address and handle such issues were provided in the previous sections of this paper, trying to cope with the existing limitations.

We envisioned situations in which users will be supported and represented by software assistants. Such assistants can take those necessary steps to protect privacy across all used devices and services. This may be including communication with other agents, installing privacy-coherent software on devices, and even policy-enforcing monitoring and auditing. The objective is to rethink the field from its foundations by proposing a more practical and comprehensive approach to overcome previous shortcomings.

Several factors can help to rapidly develop the field. First of all, the current increase of available resources (e.g., storage, computational and bandwidth), AI evolution and current regulation demands (e.g., GDPR). Second, an increasing attention to the massive collection of sensitive data by dominant companies in the information technology domain (e.g., technological firms like Amazon, Apple, Google, Facebook and Microsoft), as well as issues related to an unethical monetisation of our data. Because of that, a near-future demand for privacy by people of all ages and backgrounds will increase shortly.

4 Conclusion

The need for privacy is well understood. Nevertheless, high connectivity makes that isolated solutions, implemented independently in different devices, do not provide enough privacy guarantees. Users have several entry points to internet, and this means that information transfer can occur due to the use of different services, but also due to the use of these different entry points.

AI can help to ensure a good privacy level to citizens. In addition, privacy is an interesting property to be addressed in multiagent systems, as most of the tools developed in the field can be used to assist citizens to improve their privacy expectations. User privacy agents need to be reactive but also proactive, they need to negotiate with other agents to protect the user. They naturally need to reason and make decisions autonomously taking into account user's preferences, information and activities. They also need to exploit machine learning technologies to learn and infer what others know, as well as to conduct a posteriori actions. We think that this is an overlooked field.

References

- [1] Y. Abid. *Automated Risk Analysis on Privacy in Social Networks. (Analyse automatisée des risques sur la vie privée dans les réseaux sociaux)*. PhD thesis, University of Lorraine, Nancy, France, 2018.
 - [2] G. Acar, M. Juárez, N. Nikiforakis, C. Díaz, S. F. Gürses, F. Piessens, and B. Preneel. Fpdetective: dusting the web for fingerprints. In A. Sadeghi, V. D. Gligor, and M. Yung, editors, *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4-8, 2013*, pages 1129–1140. ACM, 2013.
 - [3] H. Azkia, N. Cuppens-Boulahia, F. Cuppens, G. Coatrieux, and S. Oulmakhzoune. Deployment of a posteriori access control using IHE ATNA. *International Journal of Information Security*, 14(5):471–483, 2015.
 - [4] T. Baarslag, A. T. Alan, R. C. Gomer, I. Liccardi, H. Marreiros, E. H. Gerding, and M. Schraefel. Negotiation as an interaction mechanism for deciding app permissions. In *Proceedings of the 2016 CHI conference extended abstracts on human factors in computing systems*, pages 2012–2019, 2016.
 - [5] T. Baarslag, A. Alper, R. Gomer, M. Alam, P. Charith, E. Gerding, and M. Schraefel. An automated negotiation agent for permission management. In *AAMAS '17: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. ACM, 2017.
 - [6] T. Baarslag, M. Kaisers, E. H. Gerding, C. M. Jonker, and J. Gratch. When will negotiation agents be able to represent us? the challenges and opportunities for autonomous negotiators. In C. Sierra, editor, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 4684–4690. ijcai.org, 2017.
 - [7] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1175–1191, 2017.
 - [8] V. Ciriani, S. D. C. di Vimercati, S. Foresti, and P. Samarati. k -anonymity. In T. Yu and S. Jajodia, editors, *Secure Data Management in Decentralized Systems*, volume 33 of *Advances in Information Security*, pages 323–353. Springer, 2007.
 - [9] A. Das, M. Degeling, D. Smullen, and N. M. Sadeh. Personalized Privacy Assistants for the Internet of Things: Providing Users with Notice and Choice. *IEEE Pervasive Comput.*, 17(3):35–46, 2018.
 - [10] F. Dernaika, N. Cuppens-Boulahia, F. Cuppens, and O. Raynaud. Accountability in the A Posteriori Access Control: A Requirement and a Mechanism. In *International Conference on the Quality of Information and Communications Technology*, pages 332–342. Springer, 2020.
 - [11] R. Dingledine. Tor and circumvention: Lessons learned. In *Annual Cryptology Conference*, pages 485–486. Springer, 2011.
 - [12] J. Domingo-Ferrer. A three-dimensional conceptual framework for database privacy. In *4th VLDB Workshop on Secure Data Management-SDM'2007*, pages 193–202. Springer Berlin Heidelberg, 2007.
 - [13] J. Domingo-Ferrer, D. Sánchez, and A. Blanco-Justicia. The limits of differential privacy (and its misuse in data release and machine learning). *Communications of the ACM, (to appear)*.
 - [14] G. T. Duncan, M. Elliot, and J.-J. Salazar-González. Why statistical confidentiality? In *Statistical Confidentiality*, pages 1–26. Springer, 2011.
 - [15] C. Dwork. Differential privacy. In *33rd International Colloquium on Automata, Languages and Programming, (ICALP 2006)*, volume 4052 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2006.
 - [16] C. Dwork. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, pages 1–19. Springer, 2008.
-

-
- [17] A. A. El-Kalam, R. E. Baida, P. Balbiani, S. Benferhat, F. Cuppens, Y. Deswarte, A. Mieke, C. Saurel, and G. Trouessin. Organization based access control. In *Proceedings POLICY 2003. IEEE 4th International Workshop on Policies for Distributed Systems and Networks*, pages 120–131. IEEE, 2003.
- [18] V. Estivill-Castro and D. F. Nettleton. Can on-line social network users trust that what they designated as confidential data remains so? In *2015 IEEE TrustCom/BigDataSE/ISPA, Helsinki, Finland, August 20-22, 2015, Volume 1*, pages 966–973. IEEE, 2015.
- [19] A. Hundepool, J. Domingo-Ferrer, L. Franconi, S. Giessing, E. S. Nordholt, K. Spicer, and P.-P. De Wolf. *Statistical disclosure control*. John Wiley & Sons, 2012.
- [20] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, M. J. Wooldridge, and C. Sierra. Automated negotiation: prospects, methods and challenges. *Group Decision and Negotiation*, 10(2):199–215, 2001.
- [21] M. Juarez and V. Torra. Dispa: An intelligent agent for private web search. In *Advanced research in data privacy*, pages 389–405. Springer, 2015.
- [22] D. Kekulluoglu, N. Kokciyan, and P. Yolum. Preserving privacy as social responsibility in online social networks. *ACM Transactions on Internet Technology (TOIT)*, 18(4):1–22, 2018.
- [23] N. Kökciyan, N. Yaglikci, and P. Yolum. An argumentation approach for resolving privacy disputes in online social networks. *ACM Transactions on Internet Technology (TOIT)*, 17(3):1–22, 2017.
- [24] N. Kökciyan and P. Yolum. Context-Based Reasoning on Privacy in Internet of Things. In *IJCAI*, pages 4738–4744, 2017.
- [25] M. Kosinski, D. Stillwell, and T. Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 110(15):5802–5805, 2013.
- [26] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.
- [27] L. Lyu, H. Yu, and Q. Yang. Threats to federated learning: A survey, 2020.
- [28] U. Pagallo. The impact of domestic robots on privacy and data protection, and the troubles with legal regulation by design. In *Data protection on the move*, pages 387–410. Springer, 2016.
- [29] S. Preda, F. Cuppens, N. Cuppens-Boulahia, J. Garcia-Alfaro, and L. Toutain. Dynamic deployment of context-aware access control policies for constrained security devices. *Journal of Systems and Software*, 84(7):1144–1159, 2011.
- [30] S. Preda, F. Cuppens, N. Cuppens-Boulahia, J. Garcia-Alfaro, L. Toutain, and Y. Elrakaiby. Semantic context aware security policy deployment. In *Proceedings of the 4th International Symposium on Information, Computer, and Communications Security*, pages 251–261, 2009.
- [31] S. Rajtmajer, A. Squicciarini, C. Griffin, S. Karumanchi, and A. Tyagi. Constrained social-energy minimization for multi-party sharing in online social networks. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 680–688, 2016.
- [32] K. J. Reza, M. Z. Islam, and V. Estivill-Castro. Privacy preservation of social network users against attribute inference attacks via malicious data mining. In P. Mori, S. Furnell, and O. Camp, editors, *Proceedings of the 5th International Conference on Information Systems Security and Privacy, ICISSP 2019, Prague, Czech Republic, February 23-25, 2019*, pages 412–420. SciTePress, 2019.
- [33] P. Samarati. Protecting respondents’ identities in microdata release. *IEEE Trans. Knowl. Data Eng.*, 13(6):1010–1027, 2001.
- [34] C. Stachl, Q. Au, R. Schoedel, S. D. Gosling, G. M. Harari, D. Buschek, S. T. Völkel, T. Schuwerk, M. Oldemeier, T. Ullmann, H. Hussmann, B. Bischl, and M. Bühner. Predicting personality from patterns of behavior collected with smartphones. *Proceedings of the National Academy of Sciences*,
-

117(30):17680–17687, 2020.

- [35] K. Stokes and M. Bras-Amorós. Optimal configurations for peer-to-peer user-private information retrieval. *Comput. Math. Appl.*, 59(4):1568–1577, 2010.
 - [36] J. M. Such and M. Rovatsos. Privacy policy negotiation in social media. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, 11(1):1–29, 2016.
 - [37] L. Sweeney. k-anonymity: A model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl. Based Syst.*, 10(5):557–570, 2002.
 - [38] The Hub of All Things. The hub of all things. (accessed January 2021) <https://www.hubofallthings.com/>.
 - [39] V. Torra. *Data privacy: Foundations, new developments and the big data challenge*. Springer, 2017.
 - [40] I. Ullah, R. Boreli, S. S. Kanhere, S. Chawla, T. A. Ahanger, and U. Tariq. Protecting private attributes in app based mobile user profiling. *IEEE Access*, 8:143818–143836, 2020.
 - [41] J. Vaidya, C. W. Clifton, and Y. M. Zhu. *Privacy preserving data mining*, volume 19. Springer Science & Business Media, 2006.
 - [42] K. Wei, J. Li, M. Ding, C. Ma, H. H. Yang, F. Farokhi, S. Jin, T. Q. S. Quek, and H. Vincent Poor. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security*, 15:3454–3469, 2020.
 - [43] M. Wooldridge. *An introduction to multiagent systems*. John Wiley & Sons, 2006.
 - [44] Q. Yang, Y. Liu, T. Chen, and Y. Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.
 - [45] Q. Yang, Y. Liu, T. Chen, and Y. Tong. Federated machine learning: Concept and applications. *ACM Trans. Intell. Syst. Technol.*, 10(2):12:1–12:19, 2019.
-