H. Roßnagel, C. H. Schunck, S. Mödersheim, D. Hühnlein (Hrsg.): Open Identity Summit 2020,
Lecture Notes in Informatics (LNI), Gesellschaft für Informatik, Bonn 2020 121

Automation Potentials in Privacy Engineering

Christian Zimmermann[1]

Abstract: The GDPR enshrines the privacy by design paradigm in law, making sound privacy engineering methods more important than ever. Integrating automation and extensive tool support into the privacy engineering process has the potential to support organizations in streamlining the implementation of privacy and data protection by design and reducing its cost. Based on a privacy engineering reference process, this paper systematically investigates automation potential in privacy engineering. In particular, it discusses potentials and implications of automation in privacy engineering and illustrates directions for future research.

Keywords: Privacy Engineering; Data Protection; Automation

1 Privacy by Design

The GDPR enshrines the "privacy by design"paradigm in law by stipulating "data protection by design and default" in its Article 25. In order to fulfill the data protection by design and default (DPbDD) obligations pursuant Art. 25 GDPR, data controllers need to consider privacy and data protection risks early on in the design of systems for processing personal data. Moreover, privacy and data protection need to be considered in the complete development life-cycle [Eu19]. The implementation of a privacy engineering process can support companies in doing so. Privacy engineering is "the discipline of understanding how to include privacy as non-functional requirement in system engineering" [CSC14] and, hence, a method to integrate the privacy by design paradigm [CSC14] into product development. From a governance perspective, privacy engineering can also be defined as "engineering data governance for personal information into the design and implementation of routines, systems, and products that process personal information" [DFF14].

Developers of systems, devices and software for processing personal data are often no privacy or legal experts and not able to fully consider data protection intricacies and requirements [Ha18]. Consequently, privacy experts need to be involved in systems and privacy engineering to support architects and developers. However, privacy experts are sparse and costly, especially those with a background in both law and computer science. Automating privacy engineering or specific steps of the privacy engineering process seems to be a promising way to mitigate the sparsity of privacy experts and to reduce development cost. Moreover, automation might also help companies establish a consistent minimum

[1] Bosch Research, 71272 Renningen, christian.zimmermann3@de.bosch.com

CC BY NC SA

level of quality with respect to analyses and measures for compliance with data protection legislation.

This paper investigates automation potentials in privacy engineering. In order to discuss privacy engineering in a systematic manner, I first present and discuss a privacy engineering reference process in Section 2. Subsequently, in Section 3, I identify potential for automation, semi-automation or tool support in the individual steps of the reference process and illustrate research streams to be addressed to foster automation of privacy engineering. Section 4 discusses advantages, disadvantages and limits of (semi-)automation in privacy engineering. Section 5 concludes the paper.

2 Privacy Engineering Reference Process

The goal of privacy engineering is to ensure the implementation of appropriate measures and safeguards for specific processing means and purposes. Figure 1 depicts the privacy engineering reference process upon which the discussion in this paper is based. The presented reference process is grounded in and extends the work by Hoepman [Ho14] and his mapping of privacy design strategies and patterns to the software development cycle. I also draw from Gürses et al. [GTD15] and Spiekermann & Cranor [SC09] and take into account the GDPR, the EDPB's Guidelines on Article 25 [Eu19] and the "Standard-Datenschutzmodell" (SDM) [Ko16], the latter of which has been drafted by German DPAs. As can be seen, the privacy engineering reference process can roughly be mapped to the "classic" software development process (see also [Ho14]). The following will briefly introduce the individual process steps and discuss associated challenges.

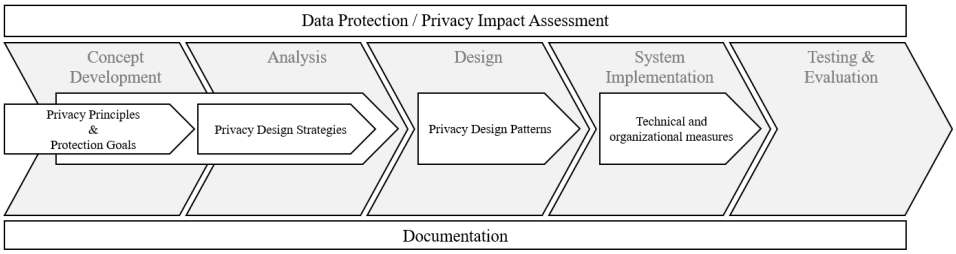

Fig. 1: Generic privacy engineering process

Note that the figure is not intended to imply a necessarily linear, one-time process but a process that might (at least partly) need to be applied iteratively, e.g., within agile development methods. Obviously, privacy engineering is also not an end in itself but a sub-activity of system engineering and product development and a means to design systems and products compliant with data protection legislation and catering to users' needs and demands. Hence, the depicted process is also not to be understood as a stand-alone process but as embedded into a broader system engineering and (product) development process and needs to include interfaces to security engineering (cf. Art. 32 GDPR).

Companies face a variety of challenges when trying to implement privacy by design and privacy engineering processes. Overarching challenges refer to the sparsity of privacy and data protection experts and to the communication-related and cultural challenges that arise in the collaboration between technologists and legal staff. Besides these more general challenges, specific challenges arise in the different process steps. In the following the individual steps of the reference process and associated challenges are briefly illustrated in order to inform the discussion on automation potentials.

2.1 Privacy Principles and Protection Goals

For controllers or processors to comply with the data protection by design and default requirement, systems for processing personal data obviously need to be implemented under consideration of the relevant obligations laid down in the applicable data protection regulations. Consequently, the first step in the depicted privacy engineering process refers to the elicitation of relevant protection goals and privacy-related requirements that the system needs to achieve or fulfill, respectively. Obviously, these encompass primarily the data protection principles laid down in the legislation, e.g., the GDPR. Guidelines by various DPAs and other institutions (e.g. [OE13] or [IS11]) aim to support the translation of these principles and abstract legal requirements into actionable technical and organizational requirements. Notably, the SDM [Ko16] provides a mapping of GDPR articles to the data protection goals proposed by Hansen et al. [HJR15]. It is also advisable to take into account user expectations and demands regarding privacy and to elicit those using, e.g., user studies. Not only might considering user expectations increase user satisfaction and acceptance. Those expectations are also highly relevant in most jurisdictions, e.g. under the GDPR where reasonable expectations of data subjects play a prominent role in assessing lawfulness of data processing based on legitimate interest (cf. Recital 47 GDPR).

The protection goals for privacy engineering [HJR15] are general enough in order to provide guidance for designing systems regardless of their intended domain of deployment. However, translating legal texts and the obligations specified therein into technical requirements is often a daunting task. On the one hand, non-legal staff such as software developers often lack the expertise to interpret legal texts and the knowledge of current legal interpretations of the law. On the other hand, legal staff often lacks the technological expertise to translate legal obligations into technical requirements.

Further problems can arise from the novelty of certain systems, e.g., autonomous systems or IoT systems. In the absence of broad adoption of such systems and the resulting lack of well-defined social norms and expectations regarding their usage, it is hard to formulate reasonable expectations of privacy. Consequently, deploying such systems entails the risk of violating newly forming social norms and expectations of privacy. While this does not necessarily have to amount to a compliance problem, it has the potential to deter potential users from using the systems.

2.2 DPIA and Documentation

The potential risk to privacy and data subjects' rights and freedoms posed by the system to be developed needs to be assessed early on in the development process [Eu19]. In many cases, performing a data protection impact analysis (DPIA) will also be legally required, e.g., in case a planned processing of personal information is likely to pose a high risk to the rights and freedoms of the affected data subjects (Art. 35 GDPR). However, as depicted in Figure 1, it is not sufficient to conduct DPIAs only at the beginning of the engineering process, especially in case agile development practices are used [ZZ20]. Rather, impact assessments need to be conducted repeatedly in order to be able to assess whether changes in the system (either in functionality or in applied measures for data protection) or changes in the state of the art change the identified risk [Eu19]. The (updated) DPIA results need to be reflected in all other process steps.

Several aspects make DPIAs challenging. On the one hand, the need to repeatedly update DPIAs imposes high efforts. This is particularly challenging in case agile development methods are used and changes to the system occur very often [ZZ20] or service-oriented architectures are utilized [GG18]. DPIAs require expert knowledge and assessment, which further increases cost and can delay development when experts are sparse.

DPIA results and information on implemented measures and the actual processing need to be documented (Art. 35 & 5(2) GDPR). Ideally, documentation of DPIA results, design decisions, planned processing steps and implemented measures is conducted in parallel to development. While this will decrease the efforts to be spent after development and during the operation of the system, it imposes a high effort in the development process.

2.3 Privacy Design Strategies & Patterns

The first process steps focus on initial risk and impact assessment and the elicitation of requirements. Subsequently, approaches to satisfying those requirements and mitigating the risks need to be chosen. Privacy Design Strategies "refer to distinct approaches that can be used to achieve privacy protection" [GTD15], e.g., aggregation of information or hiding of information. They describe fundamental approaches that can be implemented using privacy design patterns. A privacy design pattern is "a commonly recurring structure of communicating components that solves a general design problem within a particular context" [GTD15]. Privacy design patterns can also be defined as "design solutions to common privacy problems - a way to translate 'privacy-by-design' into practical advice for software engineering"[2]. For example, encryption can be considered one design pattern for the "information hiding" strategy [Ho14]. Privacy design patterns are similar to software design patterns and more detailed or closer to implementation level than privacy design strategies.

[2] https://privacypatterns.org/

Privacy design strategies can be derived from the data protection principles and protection goals defined in the relevant regulations and best practice guidelines to be adhered to in the development process (cf. [Ho14]). Based on the identified requirements, user and business needs, privacy design strategies should be chosen or developed in the early phases of product concept development. Some (mandatory) strategies can be directly found in regulation, e.g., data minimization as laid down in Art. 5 GDPR. Further, the eight privacy design strategies derived by Hoepman [Ho14] from the OECD privacy guidelines [OE13], Directive 95/46/EC and the ISO 29100 privacy framework [IS11] can be taken into account. Finally, user expectations and desires should be considered in the selection or definition of privacy design strategies.

Challenges related to privacy design strategies refer to the selection of strategies fitting the planned context and scope of the processing, i.e., strategies that provide an optimal balance between effectiveness in reducing the impact on data subjects' rights and freedoms, cost and utility of the system.

Privacy design patterns can be used to implement a chosen privacy design strategy. A broad variety of privacy design patterns have been proposed in the literature. Many of those are collected on the privacypatterns.org website curated by, among others, Jaap-Henk Hoepman, co-author of [Ho14]. The website not only lists privacy design patterns proposed in the literature but also assigns them to privacy design strategies as defined in [Ho14]. However, while privacy patterns are available, it is hard for developers to select and implement fitting patterns as "privacy patterns are scattered, unrelated, inconsistent, and immature" [Co18]. Further, it still needs to be evaluated whether and under which conditions and assumptions "classic" patterns are still viable in new domains such as autonomous systems or the IoT.

2.4 Technical and Organizational Measures

In the final step of the presented process, actual measures for implementing the selected strategies and patterns need to be selected and implemented. The privacy engineering process provided in Figure 1 culminates in the process steps "Technical and organizational measures", whereas the approach presented in [Ho14] puts "privacy-enhancing technologies". In the reference process presented here, a broader perspective is chosen in order to emphasize that technology in general and PETs in particular can not be implemented detachedly from accompanying organizational measures. Moreover, the broader term is used to clearly indicate the inclusion of not only PETs but also transparency-enhancing technologies (TETs) [JWV13; Zi15] as measures for data protection and privacy preservation.

Challenges associated with this step are very similar to those faced in security engineering. Choosing appropriate technology, methods and artifacts is one side of the challenge. The other is the correct implementation of the selected solutions, e.g., selecting appropriate parameters for encryption. Further challenges arise from the application of machine learning and artificial intelligence to personal data [Pa18].

3 Potential for Automation

In the following, automation potentials in privacy engineering are illustrated. The investigation is structured along the steps of the reference process. In particular, I will analyze which aspects of the individual process steps lend themselves to (semi-)automation and discuss avenues for future research. The feasibility and desirability of automation are discussed further in Section 4.

3.1 Privacy Principles and Protection Goals

In this process step, three coarse sub-steps can be delineated. (a) First, relevant legal requirements and protection goals need to be identified. This entails identification of relevant legislation, DPA guidelines, the state of the art and user expectations. (b) Subsequently, relevant parts of these sources need to be identified based on the scope and context of the planned processing. For example, which of the obligations stipulated in the GDPR will apply depends on, i.a., whether data will be transferred to third countries, which types and extent of personal data will be processed or whether the controller will act alone or as a joint controller with others. (c) Finally, the identified (legal) requirements need to be translated into technical or organizational requirements specific to the planned processing. Albeit possibly hard to harness, there is potential for automation or semi-automation in all of these process sub-step.

Sub-step (a) requires knowledge of the context of the planned processing, e.g., applicable jurisdiction and applicable laws. Further, knowledge of relevant case law, legal decisions and DPA opinions and guidelines might be necessary. In the context of international service contracts, Waldburger et al. [Wa10] address the former and propose and implement a modeling method and information model for automated determination of jurisdiction and applicable law. While their approach is not directly transferable to the data protection domain, it illustrates avenues for future research into automated identification of relevant laws. At least semi-automation is conceivable in this sub-step, e.g., based on automated selection of relevant documents based on some input such as a questionnaire.

Once the relevant sources have been identified, the relevant parts, i.e., those including applicable obligations or requirements, need to be identified in sub-step (b). First examples of semi-automation or decision-support tools for this sub-step have already been presented. For example, Colesky et al. [Co19] present a tool to provide information on which recitals and articles of the GDPR are to be considered based on a questionnaire. Work like this can constitute the basis for further automation.

Automated elicitation of technical requirements from identified legal text might be based on work towards formal models of relevant requirement sources, e.g., [Ma08; Me05]. Models of the system as well as the scope and context of the planned processing would also support automation of sub-step (c). Obviously, a chicken-and-egg problem arises here.

Notwithstanding, initial models of the planned system or at least the scope and context of the processing could be used. Several approaches for modeling security-relevant or privacy-relevant system behavior and requirements have been presented [Ah17a; Ah17b; MG07]. Kalloniatis et al. "provide a set of concepts for modeling privacy requirements in the organisation domain and a systematic way-of-working for translating these requirements into system models" [KKG08]. They formally define process patterns for "(1) analysing the impact of privacy requirement(s) on organisational goals, subgoals and processes and (2) suggesting of appropriate system implementation technique(s) for realising these requirements" [KKG08]. Approaches like these might build the basis for automated requirements analysis or translation and merit further research.

3.2 DPIA and Documentation

Performing a DPIA requires an understanding of the system and the scope and context of the planned processing. It further requires analysis and assessment of the impact of the processing on data subjects' rights and freedoms. Consequently, DPIAs are time-consuming and require expert input and, hence, seem desirable candidates for automation. Given their conceptual relation to threat and risk analysis from the field of cyber security and the progress in automation in that area (cf. Threat Dragon[3], MS Threat Modeling Tool[4]), there is at least reason to hope, that DPIAs can at least be (semi-)automated, potentially based on the methods described above. For example, the STRIDE-based LINDDUN [De11] framework might be extendable into a foundation for semi-automated analysis. In fact, some work in the direction of DPIA automation have already been conducted. For example, as already described in [Zi19], the French DPA CNIL provides a tool for performing data protection impact analysis and generating standardized documentation of the analysis results, which comes in the form of an interactive questionnaire with knowledge base[5]. Hence, the tool supports and formalizes DPIAs, but does not provide full automation.

As already described above, potentials for further automation can be found in the formalization of privacy goals (e.g. [MV19]), system behavior and processing context and model-based engineering (see e.g. [Ah17a; Ah17b; KKG08]). In case a DPIA has to be updated during the development process and software code is available, code analysis as applied in the area of cyber security [CM04] but focusing on the flow of personal data might be another research avenue worth following (e.g. [ML00]).

A more detailed analysis of automation potential in privacy impact assessment processes is presented in [Zi19], to which I refer the interested reader. A similar investigation in the area of automation of security engineering is presented in [MF11].

[3] https://owasp.org/www-project-threat-dragon/migrated_content
[4] https://docs.microsoft.com/de-de/azure/security/develop/threat-modeling-tool
[5] https://www.cnil.fr/en/open-source-pia-software-helps-carry-out-data-protection-impact-assesment

3.3 Privacy Design Strategies and Patterns

As already mentioned in Section 2.3, architects and developers often face difficulties selecting suitable privacy strategies and patterns. Clearly, this is an area where automation or decision support tools might be beneficial. (Semi-)Formally described pattern systems can support automated selection of patterns given a set of requirements. Work towards pattern systems has been presented, e.g., [Co18]. In addition to supporting pattern selection with automation, in some instances, it might also be feasible to (semi-)automate the actual implementation of selected patterns [Bu03].

3.4 Technical and Organizational Measures

Technical and organizational measures need to be implemented in order to protect the rights and freedoms of data subjects. While this does not seem as a typical candidate for automation, there are several aspects that exhibit high potential for automation.

Changes to the system will often require the implementation of new measures. Nowadays, many development and release processes take place in a CI/CD fashion. (Frequent) updates to a system might lead to changes that impact users' privacy, be it by design or as a side effect. Consequently, measures might also have to be updated on system updates. Further, some updates might need the implementation of new measures.

Automation potential lies in the automated detection of system changes that require adaption of existing of implementation of new measures. Obviously, this is related to automated continuous DPIA (see above). Further potential lies in the automated selection of measures to be updated or newly applied, based on updated DPIA results. Examples exist in the cyber security area where, e.g., automated code analysis is well established and a variety of tools exist to automatically recommend measures to be taken to make code more secure.

Still, more research into data flow analysis (see, e.g., [ML00]) and automated pattern selection (see 3.4) is necessary in order to investigate methods for supporting automated selection and implementation of technical measures. However, the implementation of organizational measures will usually not be open to automation.

4 Discussion & Limitations

The previous section briefly outlined automation potentials in privacy engineering and suggested avenues for future research. However, only an overview has been presented and many questions deserving further investigation have only been touched upon. For example, the analysis presented in this paper focused primarily on the design phase and not on the operation of systems processing personal information. In the operation phase, an interesting

area for automation is related to data subjects' rights, e.g., to access data or to erasure of data. As there is only a rather short window of time for reacting to data subject requests (DSRs), automation seems highly beneficial. Further, CI/CD and DevOps were discussed only briefly. For privacy engineering to be truly integrated with modern software development approaches, it needs to be integrated into CI/CD and DevOps methods and tools, especially when controllers implement the systems for processing personal data themselves. Future research into Privacy DevOps might be able to draw from the work in the area of SecDevOps [MO16].

Controllers that develop own systems for processing personal data will most likely benefit most directly from automation of privacy engineering, at least from an economic perspective. Automating time consuming tasks requiring expert input can reduce cost and might be able to supports consistency in the engineering process. Still, the actual economic impact of automation in privacy engineering deserves a closer look, as well as the potential of automation to actually support more consistent, compliant or, generally, privacy-preserving results in privacy engineering.

Some of the process steps illustrated above are of less technical nature than others. In particular, privacy and data protection impact assessment often require interpretation and case-specific balancing of technical, economic and legal aspects. Clearly, such a task is a less suited candidate for full automation than more mechanical tasks not requiring balancing decisions. However, besides feasibility, the desirability of automation in privacy engineering also needs to be discussed. In particular, automation in privacy engineering needs to be considered in the light of it's impact on the human rights aspects underlying data protection regulation and the regulator's intentions in stipulating DPIAs, balancing tests and the implementation of measures for ensuring the rights and freedoms of data subjects. As already hinted at in [Zi19], automated DPIAs (in contrast to manual privacy impact assessments) might lead to negligence of relevant privacy aspects and a too narrow focus on compliance [Wr12]. Automation of DPIAs using AI also entails the risk of bias introduced by biased AI [YW18] and, more generally, the codification into technology of "one-size-fits-all" approaches in an area where case-specific deliberation is required [PD16]. Still, automation also has the potential to provide for more secure software products, e.g., through automated security testing as described above. This in turn can prevent privacy violations based on data leaks due to insecure systems. Further, automated DPIAs might also be able to capture a broader spectrum of risks and threats due to a larger knowledge base compared to individual privacy engineering teams.

5 Conclusion

This paper discussed the potential for automation in privacy engineering. To allow for a more systematic investigation, it presented a privacy engineering reference process and discussed the automation potential of the process's steps individually. Based on the discussion, avenues for future research were illustrated .

References

[Ah17a] Ahmadian, A. S.; Peldszus, S.; Ramadan, Q.; Jürjens, J.: Model-based privacy and security analysis with CARiSMA. In: Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering. ACM Press, pp. 989–993, 2017.

[Ah17b] Ahmadian, A. S.; Strüber, D.; Riediger, V.; Jürjens, J.: Model-Based Privacy Analysis in Industrial Ecosystems. In (Anjorin, A.; Espinoza, H., eds.): Modelling Foundations and Applications. Vol. 10376. LNCS, Springer International Publishing, Cham, pp. 215–231, 2017.

[Bu03] Bulka, A.: Design Pattern Automation. In: Proceedings of the 2002 Conference on Pattern languages of Programs - Volume 13. CRPIT '02, Australian Computer Society, Inc., Melbourne, Australia, pp. 1–10, 2003.

[CM04] Chess, B.; McGraw, G.: Static analysis for security. IEEE Security & Privacy 2/6, pp. 76–79, 2004.

[Co18] Colesky, M.; Caiza, J. C.; Del Álamo, J. M.; Hoepman, J.-H.; Martin, Y.-S.: A System of Privacy Patterns for User Control. In: Proceedings of the 33rd Annual ACM Symposium on Applied Computing. SAC '18, Association for Computing Machinery, New York, NY, USA, pp. 1150–1156, 2018.

[Co19] Colesky, M.; Demetzou, K.; Fritsch, L.; Herold, S.: Helping Software Architects Familiarize with the General Data Protection Regulation. In: 2019 IEEE International Conference on Software Architecture Companion (ICSA-C). IEEE, pp. 226–229, 2019.

[CSC14] Cavoukian, A.; Shapiro, S.; Cronk, R. J.: Privacy Engineering: Proactively Embedding Privacy, by Design, 2014, URL: https://iapp.org/resources/article/privacy-engineering-proactively-embedding-privacy-by-design/, visited on: 01/27/2020.

[De11] Deng, M.; Wuyts, K.; Scandariato, R.; Preneel, B.; Joosen, W.: A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. Requirements Engineering 16/1, pp. 3–32, 2011.

[DFF14] Dennedy, M. F.; Fox, J.; Finneran, T.: The Privacy Engineer's Manifesto: Getting from Policy to Code to QA to Value. Apress, 2014.

[Eu19] European Data Protection Board: Guidelines 4/2019 on Article 25 - Data Protection by Design and by Default (Adopted - version for public consultation), Guidelines 04/2019, European Data Protection Board, 2019.

[GG18] Galvez, R.; Gurses, S.: The Odyssey: Modeling Privacy Threats in a Brave New World. In: 2018 IEEE European Symposium on Security and Privacy Workshops. Pp. 87–94, 2018.

[GTD15] Gürses, S.; Troncoso, C.; Diaz, C.: Engineering privacy by design reloaded. In: Amsterdam Privacy Conference. 2015.

[Ha18] Hadar, I.; Hasson, T.; Ayalon, O.; Toch, E.; Birnhack, M.; Sherman, S.; Balissa, A.: Privacy by designers: software developers' privacy mindset. Empirical Software Engineering 23/1, pp. 259–289, 2018.

[HJR15] Hansen, M.; Jensen, M.; Rost, M.: Protection goals for privacy engineering. In: 2015 IEEE Security and Privacy Workshops, SPW 2015. Pp. 159–166, 2015.

[Ho14] Hoepman, J.-H.: Privacy Design Strategies. In: ICT Systems Security and Privacy Protection. Springer, Berlin, Heidelberg, pp. 446–459, 2014.

[IS11] ISO/IEC Joint Technical Committee 1 SC 27: ISO/IEC 29100 Information technology - Security techniques - Privacy framework, 2011.

[JWV13] Janic, M.; Wijbenga, J.; Veugen, T.: Transparency Enhancing Tools (TETs): An Overview. In: 2013 Third Workshop on Socio-Technical Aspects in Security and Trust (STAST). Pp. 18–25, 2013.

[KKG08] Kalloniatis, C.; Kavakli, E.; Gritzalis, S.: Addressing privacy requirements in system design: the PriS method. Requirements Engineering 13/3, pp. 241–255, 2008.

[Ko16] Konferenz der unabhängigen Datenschutzbehörden des Bundes und der Länder: Das Standard-Datenschutzmodell - Eine Methode zur Datenschutzberatung und -prüfung auf der Basis einheitlicher Gewährleistungsziele (V.1.0), 2016.

[Ma08] May, M. J.: Privacy APIs: formal models for analyzing legal and privacy requirements, Dissertation, University of Pennsylvania, 2008.

[Me05] Mercatali, P.; Romano, F.; Boschi, L.; Spinicci, E.: Automatic Translation from Textual Representations of Laws to Formal Models through UML. In: 18. International Conference on Legal Knowledge and Information Systems (JURIX). Vol. 134, pp. 71–80, 2005.

[MF11] Montesino, R.; Fenz, S.: Information Security Automation: How Far Can We Go? In: 2011 Sixth International Conference on Availability, Reliability and Security. Pp. 280–285, 2011.

[MG07] Mouratidis, H.; Giorgini, P.: Secure tropos: a security-oriented extension of the tropos methodology. International Journal of Software Engineering and Knowledge Engineering 17/02, pp. 285–309, 2007.

[ML00] Myers, A. C.; Liskov, B.: Protecting privacy using the decentralized label model. ACM Transactions on Software Engineering and Methodology (TOSEM) 9/4, pp. 410–442, 2000.

[MO16] Mohan, V.; Othmane, L. B.: SecDevOps: Is It a Marketing Buzzword? - Mapping Research on Security in DevOps. In: 2016 11th International Conference on Availability, Reliability and Security (ARES). IEEE, pp. 542–547, 2016.

[MV19] Mödersheim, S.; Viganò, L.: Alpha-Beta Privacy. ACM Transactions on Privacy and Security 22/1, pp. 1–35, 2019.

[OE13] OECD: The OECD Privacy Framework, 2013, URL: https://www.oecd.org/internet/ieconomy/oecd_privacy_framework.pdf, visited on: 01/24/2020.

[Pa18] Papernot, N.: A Marauder's Map of Security and Privacy in Machine Learning: An overview of current and future research directions for making machine learning secure and private. In: AISec '18: Proceedings of the 11th ACM Workshop on Artificial Intelligence and Security. ACM Press, 2018.

[PD16] Pagallo, U.; Durante, M.: The Pros and Cons of Legal Automation and its Governance. European Journal of Risk Regulation 7/2, pp. 323–334, 2016.

[SC09] Spiekermann, S.; Cranor, L. F.: Engineering Privacy. IEEE Transactions on Software Engineering 35/1, pp. 67–82, 2009.

[Wa10] Waldburger, M.; Charalambides, M.; Schaaf, T.; Stiller, B.: Automated determination of jurisdiction and applicable law for international service contracts: Modeling method, information model, and implementation. In: 18th Biennial and Silver Anniversary International Telecommunications Society Conference (ITS 2010). Pp. 1–31, 2010.

[Wr12] Wright, D.: The state of the art in privacy impact assessment. Computer Law & Security Review 28/1, pp. 54–61, 2012.

[YW18] Yapo, A.; Weiss, J.: Ethical Implications of Bias in Machine Learning. In: Proceedings of the 51st Hawaii International Conference on System Sciences. Pp. 5365–5372, 2018.

[Zi15] Zimmermann, C.: A Categorization of Transparency-Enhancing Technologies. In: Amsterdam Privacy Conference. Amsterdam, NL, 2015.

[Zi19] Zibuschka, J.: Analysis of Automation Potentials in Privacy Impact Assessment Processes. In: 1st Workshop on Security, Privacy, Organizations, and Systems Engineering. Luxembourg, 2019.

[ZZ20] Zibuschka, J.; Zimmermann, C.: Lean Privacy by Design. In (Reinhardt, D.; Langweg, H.; Witt, B. C.; Fischer, M., eds.): SICHERHEIT 2020. Gesellschaft für Informatik e.V., Bonn, pp. 125–128, 2020.