

“The Grace Period Has Ended”: An Approach to Operationalize GDPR Requirements

Vanessa Ayala-Rivera, Liliana Pasquale

Lero@UCD, School of Computer Science, University College Dublin, Ireland
e-mail: vanessa.ayalarivera@ucd.ie, liliana.pasquale@ucd.ie

Abstract—The General Data Protection Regulation (GDPR) aims to protect personal data of EU residents and can impose severe sanctions for non-compliance. Organizations are currently implementing various measures to ensure their software systems fulfill GDPR obligations such as identifying a legal basis for data processing or enforcing data anonymization. However, as regulations are formulated vaguely, it is difficult for practitioners to extract and operationalize legal requirements from the GDPR. This paper aims to help organizations understand the data protection obligations imposed by the GDPR and identify measures to ensure compliance. To achieve this goal, we propose GuideMe, a 6-step systematic approach that supports elicitation of solution requirements that link GDPR data protection obligations with the privacy controls that fulfill these obligations and that should be implemented in an organization’s software system. We illustrate and evaluate our approach using an example of a university information system. Our results demonstrate that the solution requirements elicited using our approach are aligned with the recommendations of privacy experts and are expressed correctly.

Index Terms—GDPR; Compliance; Privacy; Requirements;

I. INTRODUCTION

In May 2018, the General Data Protection Regulation 2016/679 (GDPR) came into effect to replace the Data Protection Directive 95/46/EC (DPD95). The GDPR was designed to harmonize data privacy laws across Europe in order to give greater protection and capabilities to individuals for controlling their personal data in the face of new technological developments [1]. The GDPR applies to all the organizations that handle personal data about EU residents, regardless of their physical locations.

Complying with the GDPR has become a top-of-mind for organizations worldwide. Beside other sanctions, infringements of the GDPR can impose fines up to 20 million or 4% of an organization’s global turnover. Moreover, although some of the GDPR obligations were already specified in the DPD95, these have mainly been perceived as “recommendations”. Therefore, most organizations have only started recently to implement measures to comply with the GDPR [2].

However, organizations are facing several obstacles in their journey towards GDPR compliance. Some organizations are not aware or do not understand the changes that the GDPR will bring to their businesses [3]. For example, a survey conducted between July and August 2017 by the Institute of Directors among 869 of its members in the UK revealed that 30% of

company directors have not heard of the GDPR, while 40% were still unsure about whether their company will be affected by the GDPR [4]. Other surveys expose similar problems such as the lack of preparation to meet the GDPR legal obligations and the lack of awareness about the consequences of non-compliance [2], [5], [6].

Most of these problems are rooted in the vague, ambiguous, and verbose nature of regulations, which individuals - who do not possess legal expertise - often find difficult to understand. Likewise, understanding legal requirements is generally time-consuming and cumbersome, thus complicating their operationalization. These problems can jeopardize compliance with the GDPR, especially when this process is not assisted by data protection law experts. This is often the case for small- and medium-sized organizations, or independent researchers and consultants, who usually do not have enough resources to afford legal support [3].

Additionally, extracting requirements from legal texts and interpreting them properly is a complex and error-prone process [7]. Mapping legal obligations into software functionality is also non-trivial [8], [9]. As legal requirements are oftentimes too abstract, they may leave space for multiple interpretations. For example, the GDPR states that companies must provide a reasonable level of protection of personal data, without clarifying what “reasonable” means exactly [10]. Similarly, the GDPR promotes “privacy by design”, without detailing how it should be achieved [11]. Therefore, it is often the case that IT professionals (or those in charge of implementing software changes to comply with the GDPR) lack of guidance to understand what are the requirements that should be operationalized and implemented in an organization’s software system to support compliance [12], [13].

This paper aims to help organizations understand the data protection obligations imposed by the GDPR and identify measures to ensure compliance. To achieve this goal, we propose GuideMe, a 6-step systematic approach that supports practitioners in the elicitation of solution requirements from the GDPR legal obligations. Solution requirements link the GDPR obligations and related business requirements to privacy controls necessary to satisfy them. Privacy controls are also contextualized depending on the stakeholder scenario and the data processing activity to which they should be applied. To increase confidence in the effectiveness of privacy controls,

we created a comprehensive catalog from which such privacy controls can be selected.

We illustrate and evaluate our approach focusing on the data protection obligations stated in Articles 5 and 25 of the GDPR and using a substantive example of a university information system. We validate with privacy experts the appropriateness of solution requirements and verify the correctness of the requirements specification. Our results show that the solution requirements elicited using our approach suggest privacy controls that can satisfy the related GDPR obligations and business requirements and are aligned with the recommendations of privacy experts. Solution requirements are also expressed correctly.

The rest of this paper is structured as follows. Section II presents relevant related work. Section III describes our approach and illustrates it using the university example. Section IV describes the evaluation of our approach. Section V discusses the limitations of our work. Finally, Section VI draws the conclusions and provides pointers to future work.

II. RELATED WORK

Researchers and practitioners have investigated different approaches to support organizations in achieving compliance with the GDPR.

Some work has focused on understanding the implications of the GDPR within organizations. For example, Tikkinen-Piri et al. [3] compared the GDPR with the DPD95 with the aim of analyzing their differences and identifying GDPR's practical implications, specifically for organizations providing services based on personal data. Additionally, a significant number of toolkits and checklists [14]–[19] have been developed (by public agencies and private companies) to support organizations in evaluating their compliance with the GDPR. For example, the Information Commissioner's Office in the UK offers self-assessment tools [14] to examine in a structured manner whether the legal responsibilities of data controllers are met. Similarly, the Data Protection Commissioner in Ireland has also developed GDPR readiness checklists [15], covering a broad range of areas (such as data security, accuracy and retention, or international data transfers) to assist organizations in identifying the personal data they retain and process. Microsoft [16] and Symantec [17] have also provided their own assessment toolkits to assist organizations in evaluating whether they have appropriate measures in place to protect personal and sensitive data. Such firms also offer a broad range of service packages to support customers with the creation and execution of a GDPR implementation programme. These toolkits have been useful to identify gaps in compliance and additional measures organizations should put in place to protect personal data. However, they do not provide concrete suggestions about specific privacy controls that should be implemented in software systems to support compliance. The measures suggested by these toolkits only represent legal obligations, thus still requiring expert knowledge to be contextualized w.r.t. specific data processing activities and/or concrete usage scenarios.

Although previous work has addressed elicitation and modeling of legal requirements, it has not focused on the GDPR. For example, Otto and Antón [7] surveyed 50 years of work about modeling of legal texts for software systems development. Based on this survey, they elicited a set of requirements for a decision support system aimed to help analysts capture legal requirements. Christmann et al. [20] presented a structured method for identifying IT security and legal requirements in cloud services. Ghanavati et al. [21] proposed a systematic method (based on the Goal-oriented Requirement Language) to extract legal requirements from regulations and represented them using a conceptual meta-model, i.e., a legal profile. Boella et al. [22] compared methodologies for extracting legal requirements. They identified that most methodologies do not provide mechanisms to articulate effectively the (potential) unwritten rules that might influence legal reasoning, such as the context of the legal scenario. Finally, Soltana et al. [23] explored the potential usage of legal requirements models in simulated scenarios, such as taxation.

Existing research has also provided concrete data protection techniques to support GDPR compliance. For example, Gjermundrød et al. [8] proposed a technical framework to generate verifiable snapshots of a user's data trails and track disclosure of personal information. Bolognini and Bistolfi [24] investigated the suitability of pseudonymization as a technique to reduce individuals' privacy risks and to help data processors fulfill their data protection obligations. Other approaches have proposed strategies to support privacy-by-design (PbD) [9], [11], [25]. For example, Colesky et al. [9] discussed the use of strategies and privacy patterns to realize PbD as an alternative approach to more traditional requirements engineering methodologies. Koops and Leenes [11] analyzed the challenges of implementing PbD from a coding perspective and conclude that PbD should be complemented with a communication strategy. Cavoukian [25] presented a guide (composed of 7 foundational PbD principles) to implement strong privacy practices. However, although these design guidelines may be useful to enhance the privacy of software systems, they do not provide suggestions that can ensure traceability of privacy controls to the legal text. Thus, these approaches cannot provide information about the (degree of) compliance that suggested privacy controls achieve w.r.t. certain regulations.

III. SOLUTION

In this section, we firstly present an example that we use to illustrate our proposed solution. Then, we discuss some of the assumptions made in this work. Later, we introduce GuideMe, our 6-step approach to support practitioners in the elicitation of solution requirements from the GDPR.

A. University Example

Our example involves an educational institution, University X, which processes personal data about students and staff. As University X handles data from EU residents, it should comply with the GDPR data protection obligations. The information system components of University X are shown

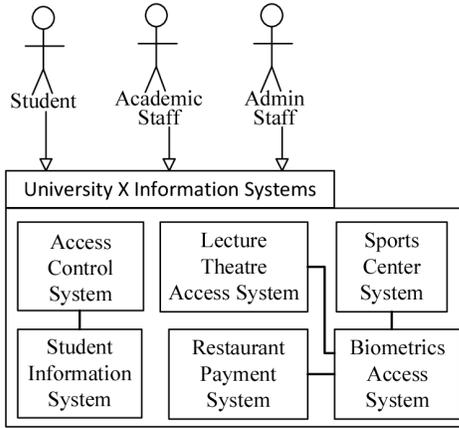


Fig. 1. University X Information Systems.

in Fig. 1. Students and staff can access students’ personal information and academic records managed by the Student Information System. The Access Control System manages university authorization and authentication to control access to the Student Information System. Likewise, students and staff can access physical locations within the university (e.g., lecture theatres, sports center) and use various university services (e.g., restaurants, libraries) by authenticating through fingerprint scanning. The Biometrics Access System authorizes or revokes admission to physical locations and the provision of university services.

B. Assumptions

In the following paragraphs, we introduce the assumptions that we make in this paper.

GDPR Scope: One of the most important provisions of the GDPR in relation to data protection is Privacy by Design and by Default (PbD) stated in Article 25. PbD encourages organizations to consider privacy and data protection rules in any action that involves processing personal data. Our work has focused on this provision given its impact and relevance to existing software systems. However, GuideMe could still be applied, without loss of generality, to operationalize requirements derived from other provisions of the GDPR. The most important PbD obligation applicable to software systems is that the Data Protection Principles (DPRs) are fulfilled. In Article 5 the GDPR states seven DPRs (shown in Table I) with the objective of integrating privacy measures into software effectively. We considered these seven principles as the main legal obligations to be met in order to comply with Articles 5 and 25 in the GDPR. Hence, we illustrate and evaluate our approach focusing on those provisions.

Business Requirements: To remove subjectivity in the interpretation of the DPRs and make them more understandable for a general audience, we expressed DPRs as business requirements. More precisely, we associated each DPR with one or more business requirement using a software requirements specification (SRS) template. Such template can help decompose the DPRs into more granular functionalities that

TABLE I
GDPR DATA PROTECTION PRINCIPLES.

1. Lawfulness, Fairness, and Transparency
2. Purpose limitation
3. Data minimization
4. Accuracy
5. Storage limitation
6. Integrity and Confidentiality
7. Accountability

Requirement ID:	BREQ-5
Requirement Statement:	The university must only collect and process the minimum amount of personal data that is required and relevant to accomplish a specific purpose .
Author:	Alice Brown
Revision Number:	1.0
Release Date:	14-Feb-18
Keywords:	Data minimization, Principle
Legal Compliance:	GDPR Art. 5.1c, Recital 39

Fig. 2. Example of Business Requirement for Data Minimization Principle.

are closer to the system implementation. For example, Fig. 2 presents a business requirement we created to represent the data minimization principle. The template [26] used to specify business requirements includes a unique identifier for the requirement to be used for cross-indexing with other artifacts in the process, the description of the requirement, the person to be contacted for further information, the revision number to track modifications to the requirement, the release date, the keywords associated with the requirement, and the cross-indexes to the related articles and recitals in the GDPR that the requirement aims to meet. The requirement description in the SRS also includes terms highlighted in bold. These terms belong to a glossary, which has the objective to consolidate the definitions of those GDPR terms that may require additional information to be understood by IT professionals. Fig. 3 shows an example of the glossary (contextualized to a university).

Archiving purposes refer to purposes in the public interest, scientific or historical research purposes or statistical purposes
Data Subject means an identified or identifiable natural person such as the staff, students, alumni, visitors, and prospective students of a university.
Legal Basis means the lawful reasons for which the university can process personal data. GDPR Art.6 states six bases: consent, contract performance, legal obligation, vital interest of individuals, public interest, legitimate interests.
Personal Data is any information relating to an identified or identifiable data subject such as names, identification numbers, location data, identifiers directly relating a subject with his/her physical, physiological, genetic, mental, economic, cultural or social identity. These can include student IDs, IP addresses.
Processing means any operation performed on personal data such as collection, recording, storage, retrieval, use, dissemination, combination, restriction, erasure or destruction.

Fig. 3. Example of a Glossary for the Description of Business Requirements Contextualized to the University X Example.

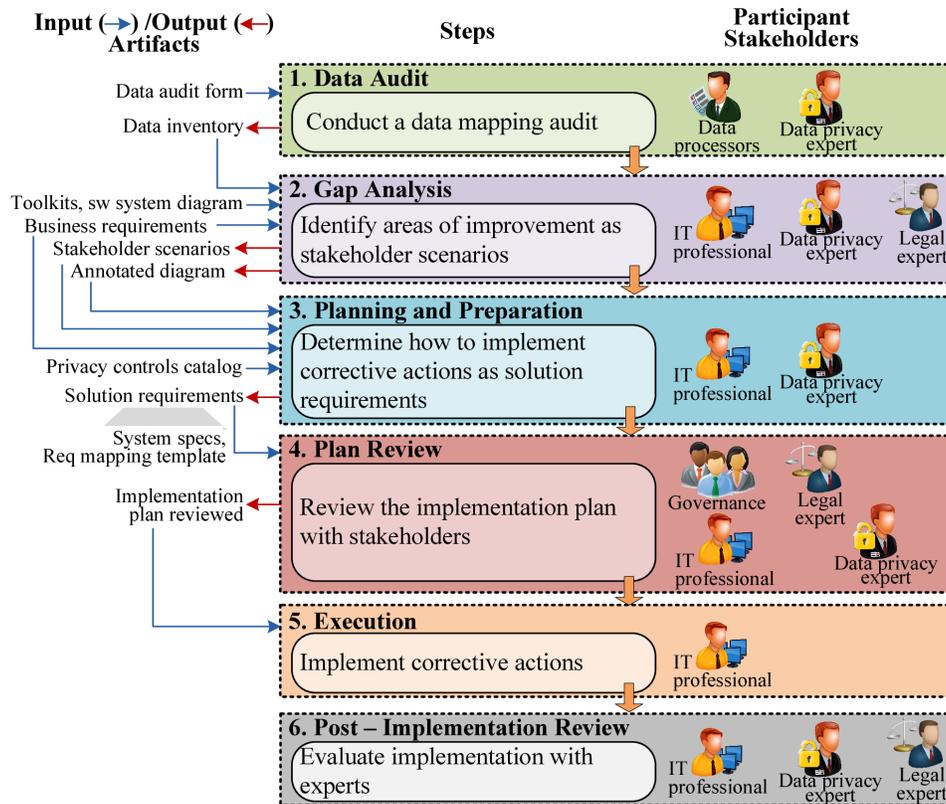


Fig. 4. The GuideMe Approach.

As part of this work, we have created an extensive set of business requirements associated with all DPRs. These serve as input to our approach and are generally applicable to any organization. The specification of these business requirements is available publicly [27].

C. The GuideMe Approach

GuideMe provides a systematic approach to move from abstract legal provisions to solution requirements that can be implemented in software systems by IT professionals. Our approach is inspired by the Business Analysis Body of Knowledge (BABOK) [28], which suggests eliciting requirements progressively, moving from a business perspective down to a solution level.

The GuideMe approach includes 6 steps shown in Fig. 4. In the rest of this section, we explain the activities, artifacts, and the main stakeholders involved at each step of the approach and illustrate them using our university example.

(1) Data Audit: This step requires performing an information audit to assess what data an organization holds, where the data originates, how it is obtained, how it is processed and under what legal grounds, where it ends up being stored, with whom the data is shared, etc. A data audit allows establishing factual context about data processing activities, as well as to identify potential privacy risks in those activities. For example, in University X, different information systems may hold personal information associated with students such

as health information (the university clinic), extracurricular activities (the sports club), family incomes (the office of financial aid), and demographic data (the admissions office). There is no predefined format to adhere to for conducting a data audit. Hence, organizations can utilize different tools such as data audit forms and data mapping templates provided by supervisory bodies (e.g., data protection authorities) or even technological solutions provided by private companies [29]–[31]. Ideally, any staff member handling personal data in the organization should be involved in the data audit (depicted as data processors in Fig. 4) as well as a data privacy expert (e.g., data protection officer), as s/he would have expertise in how to conduct this activity appropriately.

In our example, University X has adopted a data audit form released by the University of Leicester [32]. A sample of this form is shown in Fig. 5. All university staff handling personal data has been involved in the audit exercise such as academics, researchers, marketing, human resources, etc. An example of categories of personal data contained in the data inventory of University X include employment contracts, student records, clinical files, stipend documentation, reference letters, alumni details, details about participants of research studies. More detailed lists of personal data managed within a university can be found in [32], [33].

(2) Gap Analysis: This step requires performing a gap analysis on the completed audit to identify the areas (e.g., flows, processes, systems) that need to be improved through re-

University X Data Protection Audit Form	
Department	_____
Locations of data	_____
Purpose	_____
Legal Ground	_____
Data formats	<input type="checkbox"/> Computer <input type="checkbox"/> Photograph <input type="checkbox"/> Paper <input type="checkbox"/> Video <input type="checkbox"/> Audio <input type="checkbox"/> Other: _____
Data Subjects	<input type="checkbox"/> Students <input type="checkbox"/> Employees <input type="checkbox"/> Alumni <input type="checkbox"/> Clients <input type="checkbox"/> Other: _____
Types of Data	<input type="checkbox"/> Personal details <input type="checkbox"/> Education and Training <input type="checkbox"/> Employment details <input type="checkbox"/> Financial details
Data Recipients	<input type="checkbox"/> Data subjects themselves <input type="checkbox"/> Prospective employers <input type="checkbox"/> Suppliers <input type="checkbox"/> Local government

Fig. 5. Example of Data Audit Form.

medial actions. In other words, this activity allows focusing elicitation of solution requirements only on those scenarios and data processing activities that may violate the GDPR obligations. The stakeholders involved in the gap analysis are IT professionals, such as software engineers, as they have the knowledge about the software systems handling personal data; also data privacy experts and legal and compliance experts can participate to assess the practices in place, give advice, and clarify doubts to ensure GDPR compliance. The gap analysis can be performed using different tools such as the checklists and self-assessment toolkits provided by public agencies and private companies [14]–[19]. Given the set of legal obligations (encapsulated in business requirements documented in Section III-B), the different departments in the organization (or data processors in general) should provide the scenarios describing specific data processing activities violating one or more DPRs. These scenarios are referred to as stakeholder scenarios. The output of this activity is a report containing findings and recommendations, and the scenarios where measures to ensure compliance should be implemented. In our example, IT professionals use model representations (e.g., UML diagrams) of the information systems in University X to identify the parts of the system and data flows that need to be modified. The output artifact is a system diagram annotated with recommendations for changes.

Consider the following scenarios in University X:

Stakeholder Scenario #1. The library has started using the existing biometric access system to provide students physical access to the library building and to borrow books. After the gap analysis, the IT professional and the privacy and legal experts determined that the collection of biometric data violated some DPRs expressed in the elicited business requirements. For example, DPR 3 (i.e., data minimization, expressed by BREQ-5 in Fig. 2) is violated because keeping fingerprints of students accessing library services is not essential to achieve the library’s services purpose. Moreover, DPR 1 (i.e., lawfulness, fairness, and transparency) is also violated as the

Requirement ID:	BREQ-7
Requirement Statement:	The university must not keep <i>personal data</i> for longer than is necessary for the specific <i>purpose</i> that <i>personal data</i> was processed (except in the case there exist a valid legal reason to retain it such as <i>archive purposes</i> subject to the implementation of appropriate measures to protect the privacy of data subjects). The university may keep <i>personal data</i> for an unlimited period only when <i>personal data</i> is irreversibly anonymized.
Author:	Alice Brown
Revision Number:	1.0
Release Date:	14-Feb-18
Keywords:	Storage Limitation, Principle
Legal Compliance:	GDPR Art. 5.1e, Recital 39

Fig. 6. Example of Business Requirement for Storage Limitation Principle.

university did not ask for explicit consent to the processing of that biometric data. Note that biometric data is considered as a special category of personal data in the GDPR [34]. Taking this into account, a stakeholder scenario in which a privacy control should be applied can be “*The university should not collect students’ personal information to provide library services to the students*”.

Stakeholder Scenario #2. The university staff can access the full academic record of any current and former student through the student information system. After the gap analysis, the IT professional and the privacy and legal experts determined that the system was not GDPR compliant as it violated some DPRs expressed in the elicited business requirements. For example, DPR 5 (i.e., storage limitation, expressed by BREQ-7 in Fig. 6) is violated because the university stores the information about students for an indeterminate period of time which indicates that it does not have policies established for data retention. Moreover, DPR 6 (i.e., integrity and confidentiality) is also violated as the system does not prevent access to restricted information or unauthorized activities (e.g., a worker looking at the academic records of their colleagues who are former students of the university). Considering this, a stakeholder scenario in which a privacy control should be applied can be “*The university must not keep students’ personal data for longer than necessary*”.

At the end of the analysis, four DPRs were relevant to each scenario. These are marked with an “X” in Table II.

TABLE II
GDPR DPRs APPLICABLE PER SCENARIO

Scenario	DRP1	DPR2	DPR3	DPR4	DPR5	DPR6	DPR7
#1	X	X	X			X	
#2				X	X	X	X

(3) Planning and Preparation: Based on the gaps identified, it is necessary to elicit solution requirements that determine what privacy controls are necessary to satisfy specific legal obligations. Plans should also be developed to incorporate privacy controls in the organization’s software systems. For example, privacy controls can require changes in the way information is stored, processed, and accessed to support data

TABLE III
SAMPLE OF CATALOG OF SUITABLE PRIVACY CONTROLS PER GDPR DPR

Privacy Control Description	DPR1	DPR2	DPR3	DPR4	DPR5	DPR6	DPR7
1. Access Control: When processing personal data, implement access controls to ensure that personal data is only processed by authorized parties. <i>Problem Addressed:</i> Prevent unauthorized data processing. <i>Benefit:</i> The number of people who have access to personal data (disclosed, processed) is minimized, hence, preventing security breaches and illegal processing.	X	X				X	X
5. Anonymization: When data retention period has expired and wish to keep personal data for further analysis, transform data attributes with the aim of irreversibly preventing the identification of the individual to whom it relates. <i>Problem Addressed:</i> Prevent reidentification and linking attacks. <i>Benefit:</i> The principles of data protection do not apply to anonymous information so the organization can retain personal data for further analysis.			X		X	X	
7. Attribute-Based Credentials: When verifying data, use cryptographic schemes to construct anonymous proofs of ownership of personal attributes so they can be used for verification. <i>Problem Addressed:</i> Prevent leak of information by revealing more information than needed. <i>Benefit:</i> The ownership of attributes can be anonymously verified.			X				
11. Data Track: When users disclose personal data they should be provided with Data Track tools to provide the user with a detailed overview of the personal data released to (send and stored by) communication partners; also, with functions to exercise the rights of access, correct, and delete their data at services sides. <i>Problem Addressed:</i> To avoid data subjects lose track of what personal data they have disclosed, to whom, and under which conditions. <i>Benefit:</i> Data subjects have control over the data released about them by exercising their data protection rights.				X			
19. Logging: Whenever the data controller has to prove that it is in control, implement logging to demonstrate compliance. <i>Problem Addressed:</i> Prevent a non-compliant behavior. <i>Benefit:</i> The organization can demonstrate compliance with information security legislation and prevent fraud and other incidents.							X

subjects' rights. To this end, our approach aims to bridge the gap between the legal obligations and the privacy controls representing the solution.

This is depicted in step 3 of Fig. 4. To support elicitation of solution requirements, IT professionals (e.g., requirements and software engineers) use the set of business requirements associated with the legal obligations of interest, the stakeholder scenarios identified during the gap analysis, the report with a list of recommended changes (in our example, this can be defined in the form of an annotated software system diagram), and a catalog of privacy controls.

This catalog includes a set of privacy controls that serve as potential solutions to satisfy business requirements fully or partially and have been adopted successfully by practitioners in the past. We populated the privacy catalog through an extensive literature review on privacy-enhancing technologies [35], [36], information technology standards for security compliance (e.g., ISOs) [37], [38], and privacy design patterns [9], [39]–[41]. We have currently consolidated 40 privacy controls from the ISO 29100 Privacy Framework [37]; our privacy controls catalog is available publicly [27]. We characterize each privacy control with a short description, an indication of the problem addressed, and its benefits. We also indicate explicitly whether each privacy control contributes to the satisfaction of the GDPR DPRs (marked with an “X” if it contributes). An example of the structure and contents of the catalog is shown in Table III.

Using the catalog classified by DPR, a data privacy expert and an IT professional (e.g., software engineer) can then identify a set of alternative privacy controls that can be applied in the scenarios where DPRs are violated, and identify where these controls can be inserted.

TABLE IV
TEMPLATE TO REPRESENT SOLUTION REQUIREMENTS SATISFYING THE DATA PROTECTION PRINCIPLES IN THE GDPR.

Mapping Template
Under the GDPR, <i><organization></i> is obligated to fulfill <i><GDPRprovision></i> to <i><consequenceOfViolation></i> .
This provision is expressed by requirement <i><BRequirementID></i> , mapped to <i><legalComplianceReference></i> . This requirement specifies that <i><BRequirementDescription></i> .
To help satisfy <i><BRequirementID></i> , in the context of <i><scenarioID></i> , IT professional shall implement <i><privacyControlName></i> (identified by catalog entry ID <i><catalogEntryId></i>) to <i><privacyControlProblem></i> .
This privacy control involves that <i><privacyControlDescription></i> . As a result, <i><privacyControlBenefit></i> .

Then, an IT professional (e.g., requirements engineer) can express the chosen solutions as solution requirements (i.e., system specifications) which are documented using the mapping template shown in Table IV. Our template ensures traceability among business requirements (i.e., legal obligations), the stakeholder scenarios in which such requirements should be satisfied, and the privacy controls that can be applied to satisfy the requirements. The template includes placeholders that can be filled with the information identifying the business requirements (i.e., legal obligations to satisfy), the chosen privacy control, and the scenario in which it will be applied. When all the placeholders are filled, the mapping template becomes a solution requirement (i.e., a system specification). Figures 7 and 8 show examples of solution requirements (i.e., SREQ-5 and SREQ-8) applicable to the stakeholder scenarios 1 and 2,

Requirement ID:	SREQ-5
Requirement Statement:	Under the GDPR, <i>the university is obligated to fulfil the data minimization principle to prevent hoarding redundant data and minimize the risks of a data breach for data subjects.</i>
	This provision is expressed by requirement <i>BREQ-5</i> , mapped to <i>Art. 5.1c, Recital 39</i> . This requirement specifies that the <i>university must only collect and process the minimum amount of personal data that is required and relevant to accomplish a specific purpose.</i>
	To help satisfy <i>BREQ-5</i> , in the context of <i>scenario #1</i> , IT professional shall implement <i>attribute-based credentials</i> (identified by catalog entry <i>ID #7</i>) to <i>prevent leak of information by revealing more information than needed.</i>
	This privacy control involves that <i>when verifying data, use cryptographic schemes to construct anonymous proofs of ownership of personal attributes so they can be used for verification. As a result, the ownership of attributes can be verified anonymously.</i>
Author:	Bob Doe
Revision Number:	1.0
Release Date:	16-Feb-2018
Keywords:	Data Minimization, Principle, Attribute-Based Credentials

Fig. 7. Example of Solution Requirement generated with GuideMe for Data Minimization Principle in Scenario #1.

Requirement ID:	SREQ-8
Requirement Statement:	Under the GDPR, <i>the university is obligated to fulfil the integrity and confidentiality principle to ensure appropriate security of personal data.</i>
	This provision is expressed by requirement <i>BREQ-8</i> , mapped to <i>GDPR Art. 5.1f, 24.1, 25.1-25.2, 28, 39, 32; Recital 29, 71, 156</i> . This requirement specifies that the <i>university must ensure that appropriate technical or organizational measures are in place to safeguard the security and confidentiality of personal data, including prevention of unauthorised access, unlawful processing, and accidental loss, destruction or damage of data.</i>
	To help satisfy <i>BREQ-8</i> , in the context of <i>scenario #2</i> , IT professional shall implement <i>access controls</i> (identified by catalog entry <i>ID #1</i>) to <i>unauthorized data processing.</i>
	This privacy control involves that <i>when processing personal data, implement access controls to ensure that personal data is only processed by authorised parties. As a result, the number of people who have access to personal data (disclosed/processed) is minimized, hence, preventing security breaches and illegal processing.</i>
Author:	Bob Doe
Revision Number:	1.0
Release Date:	16-Feb-2018
Keywords:	Integrity, Confidentiality, Principle, Access Control

Fig. 8. Example of Solution Requirement generated with GuideMe for Integrity and Confidentiality Principle in Scenario #2.

respectively. SREQ-5 enforces attribute-based access control (e.g., using a smart card) to regulate access to sport, restaurant, and lecture facilities because students' personal information is not required to access those facilities. SREQ-8 requires that access to students' information is only given to authorized personnel, i.e., students' lecturers, if a consent to process data is given by the students.

(4) Plan Review: In this step, all the main stakeholders review the plan prepared for GDPR compliance to consider any side effects that the planned changes can bring to the business processes. For instance, even though the privacy controls listed in the catalog provide a set of mechanisms that have proven

to be useful in the past (as per the studied literature), they are not the only way to satisfy a privacy requirement, nor necessarily the best way of action for a particular scenario. The stakeholders thus must conduct an analysis to evaluate the pros and cons of the suggested privacy controls in order to select one of them depending on various factors, such as the specific scope or domain context of the scenario, costs of implementation, strengths of security, performance, effort required to train employees, etc. For example, implementing encryption of data in the university may make archiving more complicated as keys must be stored securely.

(5) Execution: Once the solution requirements are specified and approved for each scenario elicited during the gap analysis, IT professionals (e.g., software engineers) can start implementing the privacy controls indicated in the solution requirements. If the scenarios elicited during gap analysis are sufficiently complete w.r.t. the data and the processing activities collected during the data audit, an organization can have some informal assurances of supporting GDPR compliance. The completeness criteria for the scenarios elicitation may vary between organizations. These can include covering the essential data processing flows, finishing the identification of actors to be served, running out of use cases in the different systems, budget and time boxing, etc. [42]

(6) Evaluation: Finally, organizations need to ensure that all the solution requirements are satisfied. They can do so by evaluating their processes and procedures with IT, legal, and compliance experts. Also, regular audits should be scheduled periodically to identify solution requirements that may need revision, for example, because the data held by an organization or the purpose of the processing activities have changed.

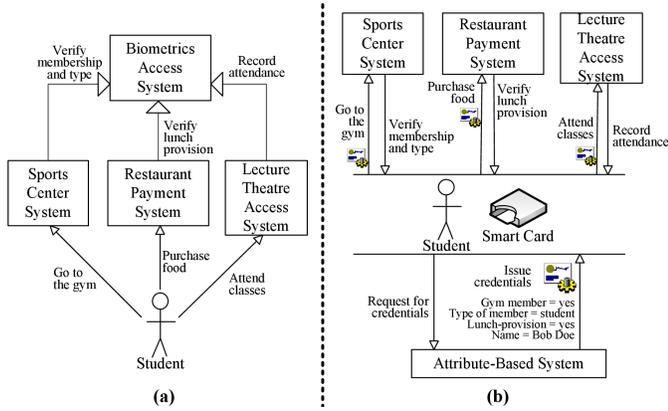
IV. EVALUATION

In this section, we present the evaluation conducted to assess whether GuideMe is able to provide practitioners useful guidance to comply with the GDPR DPRs. Firstly, we validated the privacy controls catalog by assessing the appropriateness of each privacy control in satisfying the DPRs (Section IV-A). Secondly, we verified whether GuideMe allows eliciting correct solution requirements, which have enough detail, are internally consistent, and of high quality (Section IV-B).

A. Privacy Controls Catalog Validation

The privacy catalog was validated using expert judgment gathered from five researchers who have experience in requirements engineering, security and privacy engineering, and human-computer interaction. To this aim, the catalog was shared with each of the researchers who was asked to determine, based on his/her judgment, which privacy controls s/he deemed suitable to address each GDPR DPR.

For this paper, the participants only categorized the suitability of the privacy controls as a binary decision (yes-no) without measuring their level of appropriateness in satisfying the GDPR DPRs. To consolidate the results we considered a privacy control suitable to satisfy a GDPR DPR only if the majority of the participants indicated it. The consolidated



Example 1

The university has implemented a new Biometric Access Control System (BACS) which uses fingerprinting as a means to grant students access to lecture rooms. However, this system has been used as an administrative tool for different purposes than the one stated originally. In the case of the sports center, fingerprinting is used to speed up the access to their facilities, and register to fitness classes based on their membership. In the lecture rooms, academic staff has also leveraged fingerprinting to record student attendance to classes. In the case of the restaurant, fingerprint scans are used as means of payment for the purchased food (i.e., cashless catering). Leveraging the BACS is also a more discreet method for those students that receive financial support and are entitled to free meals as they no longer have to present a special card, which could identify them to other mates. This example is illustrated (in its as-is form) by the figure on the left (a); while figure on the right (b) showcases its enhanced form after taking into consideration some of the solution requirements generated by GuideMe.

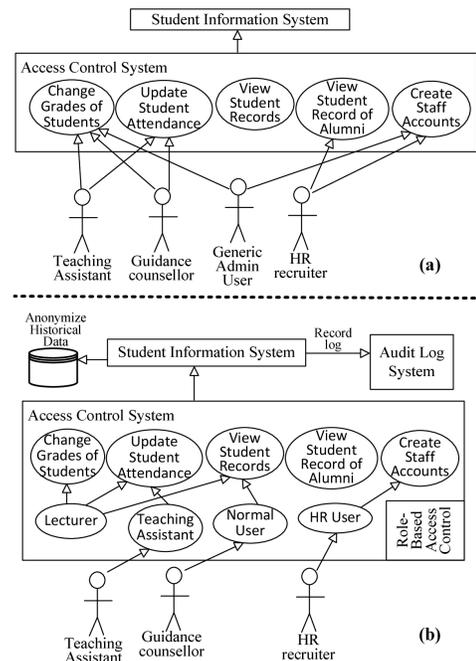
Fig. 9. Example #1

results were used to revise the privacy controls catalog used in GuideMe. It is worth mentioning that all participants involved in the validation were sufficiently familiar with the GDPR terminology/jargon and the DPRs. Moreover, they felt fairly confident about their decisions and spent a reasonable time performing the classification exercise (i.e., between 45 and 105 mins). These aspects were assessed through a questionnaire that we asked all participants to fill (available at [43]). The results we obtained give us confidence that GuideMe suggests a set of privacy controls that are compliant with the measures that a privacy expert would have recommended for the same GDPR DPR.

The privacy controls included in the catalog ensure sufficient coverage of the GDPR DPRs, as there is at least one privacy control identified as appropriate for each DPR. Also, for most requirements, there are many privacy controls that could be suggested to fulfill them. Although this is a positive outcome, the possibility of choosing between multiple options also brings the complexity and additional effort of selecting among alternative privacy controls.

B. Solution Requirements Verification

To verify the solution requirements elicited using GuideMe, we applied our approach to two substantive examples (inspired by real cases) set in a university. The examples are depicted in Figs. 9 and 10. The choice of the university domain allowed us to leverage the domain expertise available in our university to design the study case and conduct the evaluation. However,



Example 2

The Student Information System (SIS) is the application used in the university to assist academic staff in the administration of modules and students throughout their time in the university. Some cases of unnecessary access rights in SIS have been reported. The audit revealed the following: A guidance counselor and a teaching assistant made changes to the grade of a few students. Some employees added new staff user accounts. An HR admin was accessing the archived academic records of a candidate (alumni of this university) applying for a lecturer position to view his disabilities. The university is providing information retrieved from SIS to the department of education and skills about those students having a disability, or those who are part of an ethnic minority to create a program against bullying, however, no explicit consent was obtained. Some employees made changes to update attendance records using two former employee user accounts after the employees left. When officials requested the logs to review users' activities, some departments were unable to provide usable logs. This example is illustrated (in its as-is form) by the figure at the top (a); while figure at the bottom (b) showcases its enhanced form after taking into consideration some of the solution requirements generated by GuideMe.

Fig. 10. Example #2

the business requirements, the privacy catalog and the solution requirements template used in GuideMe can be reused across different types of companies and organizations. The interested reader can find the full list of solution requirements and other supplementary material used in this evaluation online [27].

We asked two IT analysts to review the solution requirements elicited for our university examples using an inspection-like approach similar to the one presented in [44]. The IT analysts (hereafter referred to as reviewers) have more than five years of experience in professional software development; one of them also holds the Certified Business Analysis Professional -CBAP- [45] designation.

The reviewers used the SMART assessment questionnaire, shown in Table V to check whether the solution requirements were correct in terms of how they were formulated. The SMART assessment questionnaire describes the properties that a good requirement should have (i.e., be Specific, Measurable, Attainable, Realizable, and Traceable). For each property, a set of assessment points are suggested in order to determine

TABLE V
SMART ASSESSMENT QUESTIONNAIRE [46]

Property	Assessment Points
Specific	<ul style="list-style-type: none"> a) clear i.e., that there is no ambiguity; b) consistent i.e., that the same terminology has been used throughout the specification to describe the same system element or concept; c) simple i.e., avoid double requirements e.g., X and Y; d) of an appropriate level of detail.
Measurable	<ul style="list-style-type: none"> a) What other requirements need to be verified before this requirement? b) Can this requirement be verified as part of the verification of another requirement? If so, which one? c) How much data or what test cases are required? d) How much processing power is required? e) Can the test be conducted on one site? f) Can this requirement be tested in isolation?
Attainable	<ul style="list-style-type: none"> a) Is there a theoretical solution to the problem? b) Has it been done before? If not, why not? c) Has a feasibility study been done? d) Is there an overriding constraint which prohibits this requirement? e) Are there physical constraints on the size of the memory, processor or peripherals? f) Are there environmental constraints such as temperature, compressed air?
Realizable	<ul style="list-style-type: none"> a) Can we satisfy this requirement given the other system and physical constraints that we have? b) Can we satisfy this requirement given the project resource constraints which we must work to?
Traceable	<ul style="list-style-type: none"> a) Can we know and understand the reason for each requirement's inclusion within the system? b) Can we verify that each requirement has been implemented? c) Can we modify the requirements easily, consistently and completely?

whether a requirement fulfills a particular property. For example, in the case of the *specific* property, a reviewer should check that there is no ambiguity in the text. This involves using a consistent terminology across the set of requirements as well as an appropriate level of detail. Similar assessment criteria are indicated for the other properties characterizing a good requirement. In the remainder of this section, we describe the main observations that the reviewers made about whether and how the solution requirements satisfy the properties indicated in the SMART framework.

Specific: Both reviewers agreed that the elicited solution requirements have an appropriate level of detail as they propose concrete actions to address the privacy needs that motivated them. Fig. 8 shows an example of a solution requirement generated for the integrity and confidentiality DPR. This requirement is also applicable to example #2 (described in Fig. 10). By analyzing the structure of the requirements, it can clearly be seen how each requirement describes the why (i.e., GDPR derived business rule and article), when

GDPR Article:	Art. 5.1f
GDPR Text:	Personal data shall be processed in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures ('integrity and confidentiality').

Fig. 11. Original Text in GDPR for Integrity and Confidentiality Principle.

(i.e., whenever processing personal data), who (i.e., the IT professionals), how (i.e., by implementing access controls), what for (i.e., to prevent unauthorized data processing), as well as the gained business value (i.e., preventing security breaches and illegal processing).

Additionally, the reviewers did not find wrong information in the inspected requirements. More specifically, reviewers searched for ambiguous information, for example, an important term, phrase, or sentence essential to understand the system behavior which was either left undefined or defined in a way that was causing confusion. Reviewers also looked for inconsistent information, i.e., two sentences contained in the requirements contradicting each other, or expressing actions that are mutually exclusive.

Finally, the reviewers considered that the derived requirements were, to a great extent, clearer than their original GDPR counterparts. This can be illustrated by comparing the original text of the GDPR articles (shown in Fig. 11) against the derived solution requirements (shown in Fig. 8).

Measurable: In our evaluation, measurable means that it is possible to verify that a requirement has been fulfilled once the IT system has been developed (or modified). The reviewers considered a requirement to be covered if its corresponding privacy controls are implemented. However, it is worth noticing that solution requirements do not provide details in relation to how privacy controls should be implemented in the software system. Therefore, reviewers indicated that testing is necessary to assess whether solution requirements were satisfied by the underlining software system implementation. For instance, the classic test strategy of trying valid and invalid test inputs for the requirement described in Fig. 8 might involve accessing the personal data with the different combinations of valid -and invalid- system roles. Another observation is that, apparently, there is not an exact (or mandatory) order in which the requirements need to be tested to cover the whole stakeholder scenarios (i.e., they can be tested in isolation). However, the independence of the requirements is a property that can be confirmed in the system testing phase.

Attainable: In our evaluation, attainable means that a requirement can be achieved so that, once it is implemented, the system exhibits the required behavior. The elicited solution requirements satisfy this property because they suggest privacy controls that could be performed, i.e., they refer to practical solutions that are described in standards for security compliance and/or are currently implemented in existing software products. This was also confirmed during the validation of the

privacy controls catalog (Section IV-A) where the participants confirmed suitability of privacy controls in satisfying all the GDPR DPRs.

Realizable: In our evaluation, realizable means that a requirement is possible to be achieved given what it is known about the constraints under which the system (or project) must be developed. All requirements were considered theoretically realizable by the reviewers. More specifically, they can be considered satisfied after implementing the related privacy control and they should be (relatively) straight-forward to achieve. However, as this is only theoretical (due to the lack of additional contextual information, such as the existence of conflicting requirements or budget/schedule project constraints), reviewers marked this point as unachieved (yet). Their advice (which we plan to address as part of our future work) is to extend our evaluation to cover the full software development life cycle to have evidence of the realizability of the generated requirements. Nonetheless, considering that the requirements are derived from legal obligations that must be met (i.e., they are not optional and their violations might involve severe sanctions), it is fair to assume that they should be realizable as they would be a top priority of any project that needs to satisfy them.

Traceable: In our evaluation, traceability is the ability to trace (forward and backward) a requirement through the whole software development life cycle (i.e., from its conception, its specification to its subsequent design, implementation and testing). Our original hypothesis when developing our approach was that it would naturally ensure that there is traceability across the different levels of requirements (i.e., moving from the GDPR original legal text to the solution requirements). This hypothesis was confirmed by the reviewers as they were (individually) able to trace back for each solution requirements, the stakeholder scenarios, business requirements, and the involved GDPR articles, from which the solution requirements were derived. For example, Fig. 8 shows how this solution requirement was derived from Article 5, business requirement BREQ-8, and suitable for the example #2.

In conclusion, these observations allow us to demonstrate—to a certain degree—that the solution requirements elicited using GuideMe are formulated properly from a software requirement perspective.

V. THREATS TO VALIDITY

In this section, we describe the factors that may have affected the validity [47] of our evaluation results.

We relied on human judgment to validate the privacy controls catalog and to verify the correctness of the solution requirements. The researchers recruited for the validation of the privacy catalog also belong to the same research group. These aspects can introduce observation biases and undermine internal validity of our results. To reduce biases during the validation of the privacy catalog, we only took into account the results provided by the researchers that claimed to have sufficient confidence in their judgment. To reduce biases during verification of solution requirements, we recruited more

than one IT analyst and ensured they have sufficient experience in analyzing software requirements.

The researchers involved in the validation of the privacy controls catalog can have also misinterpreted the meaning of the GDPR DPRs and the privacy controls used during the evaluation. This can undermine the construct validity of our results. To address this issue, we ensured that a more intuitive explanation of the GDPR DPRs was provided to the participants. Each principle was expressed avoiding use of complex and abstract legal terminology. Each privacy control was described using a brief explanation, an indication of the problem addressed, and its benefits. To address internal validity threats during the verification of solution requirements, we used the SMART framework, which includes a scientifically-validated assessment questionnaire and has been widely adopted in other research studies.

Threats to external validity may also affect generalizability of our results. In particular, we have only considered 40 privacy controls in the privacy controls catalog and GuideMe was only used in one study case. However, the privacy controls and the examples were sufficient to ensure complete coverage of the DPRs, and application of GuideMe to other domains will be considered in future work.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented GuideMe, a 6-step approach to elicit solution requirements that ensure compliance with the legal obligations imposed by the GDPR. Solution requirements link the GDPR obligations and related business requirements to privacy controls that can be implemented within an organization software system. We have illustrated our approach focusing on the data protection principles of the GDPR and using an example of a university information system. Our results show that the solution requirements elicited using our approach suggest privacy controls that can satisfy the related GDPR obligations and business requirements and are aligned with the recommendations of privacy experts. Solution requirements are also expressed correctly.

In future work, we will further evaluate the applicability of GuideMe to other application domains (e.g., healthcare information systems) to increase confidence in the generalizability of our results. We will possibly involve in the evaluation other stakeholders, such as Data Protection Officers, and individuals in charge of implementing the GDPR inside an organization. We will also use other assessment tools (beside SMART) such as the ISO/IEC/IEEE 29148 to assess the correctness of the solution requirements. To improve the scalability of GuideMe, we will automate some of the steps of our approach, such as the planning step, where privacy controls for specific scenarios and GDPR DPRs are recommended. We will also manage trade-offs of solution requirements elicited using GuideMe with other potentially conflicting requirements, such as usability and performance. Finally, as GDPR violation reference cases emerge, we will consider the interpretations made by courts in relation to what is judged as a violation and will consider this information to revise our approach.

VII. ACKNOWLEDGMENTS

The authors wish to thank Prof. Annie Antón for her fruitful comments and the anonymous reviewers for the suggestions that helped us improve the paper. This work is supported by ERC Advanced Grant no. 291652 (ASAP), and SFI Grants 10/CE/I1855, 13/RC/2094, and 15/SIRG/3501.

REFERENCES

- [1] "EUGDPR.org," <https://www.eugdpr.org/>, Last accessed: 2018-06-28.
- [2] "Privacy and the EU GDPR: 2017 Survey of Privacy Professionals," 2017. [Online]. Available: https://info.truste.com/Web-Resource-PrivacyGDPR-Research-Q22017_LP.html
- [3] C. Tikkinen-Piri, A. Rohunen, and J. Markkula, "Eu general data protection regulation: Changes and implications for personal data collecting companies," *Computer Law & Security Review*, 2017.
- [4] "Survey by the Institute of Directors," <https://www.out-law.com/en/articles/2017/october/survey-reveals-lack-of-awareness-of-gdpr-among-company-directors/>, Last accessed: 2018-06-28.
- [5] "Global Survey from Dell Technologies," <http://www.dell.com/learn/us/en/uscorp1/press-releases/2016-10-11-dell-survey-shows-organizations-lack-awareness>, Last accessed: 2018-06-28.
- [6] "Survey from Trend Micro Research," https://www.trendmicro.com/en_nz/about/newsroom/press-releases/20170906-not-prepared-for-GDPR.html, Last accessed: 2018-06-28.
- [7] P. N. Otto and A. I. Antón, "Addressing legal requirements in requirements engineering," in *International Requirements Engineering Conference (RE)*, 2007, pp. 5–14.
- [8] H. Gjermundrød, I. Dionysiou, and K. Costa, "privacytracker: A privacy-by-design gdpr-compliant framework with verifiable data traceability controls," in *International Conference on Web Engineering*. Springer, 2016, pp. 3–15.
- [9] M. Colesky, J.-H. Hoepman, and C. Hillen, "A critical analysis of privacy design strategies," in *Security and Privacy Workshops (SPW)*. IEEE, 2016, pp. 33–40.
- [10] M. Nadeau, "GDPR requirements, deadlines and facts," 2017. [Online]. Available: <https://goo.gl/yBbMXZ>
- [11] B.-J. Koops and R. Leenes, "Privacy regulation cannot be hardcoded. a critical comment on the privacy by design provision in data-protection law," *International Review of Law, Computers & Technology*, vol. 28, no. 2, pp. 159–171, 2014.
- [12] T. D. Breaux, A. I. Antón, and E. H. Spafford, "A distributed requirements management framework for legal compliance and accountability," *Computers & Security*, vol. 28, no. 1, pp. 8–17, 2009.
- [13] A. Dittel, "Data security requirements under GDPR," 2016. [Online]. Available: <https://www.lexology.com/library/detail.aspx?g=1426e18d-f687-45a0-b779-4aeb362a03ac>
- [14] "ICO - GDPR checklists," <https://ico.org.uk/for-organisations/resources-and-support/data-protection-self-assessment/>, Last accessed: 2018-06-28.
- [15] "DPC Ireland - GDPR readiness checklist tools," <http://gdprandyou.ie/wp-content/uploads/2017/12/A-Guide-to-help-SMEs-Prepare-for-the-GDPR.pdf>, Last accessed: 2018-06-28.
- [16] "Microsoft GDPR Assessment," <https://www.gdprbenchmark.com/>, Last accessed: 2018-06-28.
- [17] "Symantec GDPR Assessment," <https://www.symantec.com/en/uk/campaigns/data-privacy>, Last accessed: 2018-06-28.
- [18] "EuroComply GDPR Assessment," <http://www.eurocomply.com/>, Last accessed: 2018-06-28.
- [19] "ISACA GDPR Readiness Assessment," <https://www.isaca.org/info/gdpr/index.html>, Last accessed: 2018-06-28.
- [20] C. Christmann, J. Falkner, A. Horch, and H. Kett, "Identification of IT security and legal requirements regarding cloud services," *CLOUD COMPUTING 2015*, p. 16, 2015.
- [21] S. Ghanavati, D. Amyot, and A. Rifaut, "Legal goal-oriented requirement language (legal grl) for modeling regulations," in *International workshop on modeling in software engineering*. ACM, 2014, pp. 1–6.
- [22] G. Boella, L. Humphreys, R. Muthuri, P. Rossi, and L. van der Torre, "A critical analysis of legal requirements engineering from the perspective of legal practice," in *International Workshop on Requirements Engineering and Law (RELAW)*. IEEE, 2014, pp. 14–21.
- [23] G. Soltana, M. Sabetzadeh, and L. C. Briand, "Model-based simulation of legal requirements: Experience from tax policy simulation," in *International Requirements Engineering Conference (RE)*. IEEE, 2016, pp. 303–312.
- [24] L. Bolognini and C. Bistolfi, "Pseudonymization and impacts of big (personal/anonymous) data processing in the transition from the Directive 95/46/EC to the new EU General Data Protection Regulation," *Computer Law & Security Review*, vol. 33, no. 2, pp. 171–181, 2017.
- [25] A. Cavoukian, "Operationalizing privacy by design: A guide to implementing strong privacy practices," *Information and Privacy Commissioner, Ontario, Canada*, 2012.
- [26] T. Breaux, Ed., *Introduction to IT privacy: a handbook for technologists*. International Association of Privacy Professionals, 2014.
- [27] "GuideMe Additional Details," <https://drive.google.com/open?id=1hXmr-6OqO9G1ZfKnfnyIX0L5-7tJ30G5>, Last accessed: 2018-06-28.
- [28] International Institute of Business Analysis, *A Guide to the Business Analysis Body of Knowledge (BABOK Guide), Version 3.0*. International Institute of Business Analysis, 2014.
- [29] "Isle of Man Information Commissioner - Data protection audit self-assessment toolkit," https://www.inforights.im/media/1271/gdpr_part-1_toolkit_mapping_may2016.pdf, Last accessed: 2018-06-28.
- [30] "UK ICO - Documentation templates," <https://ico.org.uk/for-organisations/guide-to-the-general-data-protection-regulation-gdpr/accountability-and-governance/documentation/>, Last accessed: 2018-06-28.
- [31] "GDPR data mapping templates," <https://demplates.com/gdpr-data-mapping-template/>, Last accessed: 2018-06-28.
- [32] "University of Leicester - Data protection audit guidance," <https://www2.le.ac.uk/offices/ias/dp/dp-audit-guidance.pdf>, Last accessed: 2018-06-28.
- [33] "Personal Data in School of Computer Science UCD," <https://www.cs.ucd.ie/personal-data/>, Last accessed: 2018-06-28.
- [34] "UK ICO - Special category data," <http://ico.org.uk/for-organisations/guide-to-the-general-data-protection-regulation-gdpr/lawful-basis-for-processing/special-category-data/>, Last accessed: 2018-06-28.
- [35] "CIS PET wiki," <https://cyberlaw.stanford.edu/wiki/index.php/PET>, Last accessed: 2018-06-28.
- [36] A. Cavoukian, "Operationalizing privacy by design: A guide to implementing strong privacy practices," *Information and Privacy Commissioner, Ontario, Canada*, 2012.
- [37] O. Drozd, "Privacy pattern catalogue: A tool for integrating privacy principles of iso/iec 29100 into the software development process," in *IFIP International Summer School on Privacy and Identity Management*. Springer, 2015, pp. 129–140.
- [38] "ISO/IEC 27001," <http://www.iso27001security.com/html/27001.html>, Last accessed: 2018-06-28.
- [39] "Privacy Patterns," <http://privacypatterns.eu/>, Last accessed: 2018-06-28.
- [40] C. Köffel, P. John-Sören, P. Wolkerstorfer, C. Graf, H. Leif Erik, K. Ulrich, H. Hans, K. Benjamin, and P. Stefanie, "HCI pattern collection-version 2," 2010.
- [41] J.-H. Hoepman, "Privacy design strategies," in *IFIP International Information Security Conference*. Springer, 2014, pp. 446–459.
- [42] D. Zowghi and C. Coulin, "Requirements elicitation: A survey of techniques, approaches, and tools," in *Engineering and managing software requirements*. Springer, 2005, pp. 19–46.
- [43] "GDPR Questionnaire," <https://goo.gl/forms/rjaednIvJAm7c5cj1>, Last accessed: 2018-06-28.
- [44] J. Martin and W.-T. Tsai, "N-fold inspection: A requirements analysis technique," *Communications of the ACM*, vol. 33, no. 2, pp. 225–232, 1990.
- [45] "Certified Business Analysis Professional," <http://www.iiba.org/Certification/certificationlevels/level3-cbap.aspx>, Last accessed: 2018-06-28.
- [46] M. Mannion and B. Keepence, "Smart requirements," *ACM SIGSOFT Software Engineering Notes*, vol. 20, no. 2, pp. 42–47, 1995.
- [47] P. Runeson and M. Höst, "Guidelines for conducting and reporting case study research in software engineering," *Empirical software engineering*, vol. 14, no. 2, p. 131, 2009.