



RESEARCH
ARTICLE



OPEN
ACCESS



PEER
REVIEWED

Beyond silos: Bridging the gap between law and software engineering – challenges, successes, and lesson drawing

Martina Siclari *University of Luxembourg*

Salomé Lannier *University of Luxembourg*

Olivier Voordeckers *University of Luxembourg*

Stanisław Tosza *University of Luxembourg*

Sallam Abualhaija *University of Luxembourg*

Marcello Ceci *University of Luxembourg*

Nicolas Sannier *University of Luxembourg*

Domenico Bianculli *University of Luxembourg*

DOI: <https://doi.org/10.14763/2025.4.2042>

Published: 18 November 2025

Received: 27 March 2025 **Accepted:** 13 June 2025

Funding: This research was funded in whole, or in part, by the Luxembourg National Research Fund (FNR), grant reference NCER22/IS/16570468/NCER-FT.

Competing Interests: The author has declared that no competing interests exist that have influenced the text.

Licence: This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 License (Germany) which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. <https://creativecommons.org/licenses/by/3.0/de/deed.en>
Copyright remains with the author(s).

Citation: Siclari, M., Lannier, S., Voordeckers, O., Tosza, S., Abualhaija, S., Ceci, M., Sannier, N., & Bianculli, D. (2025). Beyond silos: Bridging the gap between law and software engineering – challenges, successes, and lesson drawing. *Internet Policy Review*, 14(4). <https://doi.org/10.14763/2025.4.2042>

Keywords: Interdisciplinarity, Methodology, Law, Software engineering, GDPR compliance

Abstract: This paper presents lessons learned and best practices developed for an interdisciplinary research project bridging law and software engineering, in the context of regulatory compliance with the General Data Protection Regulation (GDPR). By exploring challenges and successes encountered in such collaborations, it presents practical tools to support legal scholars in fostering meaningful interdisciplinary cooperation with software engineer researchers. Particularly, this paper brings examples on how to address the gap between the legal discipline and legal provisions, which often lack a nuanced understanding of technical realities, and software engineering, which may overlook critical regulatory contexts. The paper specifically discusses early-stage research challenges in, first, establishing common conceptual and operational ground between disciplines, addressing terminological gaps and divergent methodological assumptions. A second early step in interdisciplinary research involves defining a realistic research scope by balancing normative legal goals with technical feasibility, which should incentivise legal scholars to regard other normative sources than legal provisions. Finally, this paper addresses particularly a hybrid process for translating legal provisions into structured, traceable requirements, incorporating legal design techniques to preserve both legal accuracy and technical usability. By documenting these practices and challenges, the paper fills a gap in existing legal literature on interdisciplinarity methods, empowering legal scholars to engage confidently in collaborative research, advancing both academic inquiry and societal impact.

This paper is part of **The craft of interdisciplinary research and methods in public interest cybersecurity, privacy, and digital rights governance**, a special issue of *Internet Policy Review*, guest-edited by Adam Molnar, Diarmaid Harkin, and Urs Hengartner.

Introduction

Interdisciplinary research has gained significant momentum in recent years as scholars and practitioners increasingly acknowledge the limitations of traditional, single-discipline approaches in addressing complex, multifaceted problems. Indeed, combining different fields, expertise and concepts is more likely to foster innovation to achieve shared objectives for the benefit of society (Vienni-Baptista et al., 2022).

The demand for interdisciplinary collaboration has become more pressing with the advancement of technology, which presents challenges – also labelled as ‘wicked problems’ (Pohl et al., 2017) – that cannot be adequately addressed in disciplinary silos. One-single domain approaches risk producing outcomes that are either legally sound but technically infeasible, or technologically sophisticated yet misaligned with legal and ethical standards. Thus, the involvement of scholars from diverse domains proved essential in the design of new technological solutions (Hoess et al., 2024). Existing literature, for instance, confirms the need for enhanced collaboration between legal and engineering experts in the field of data protection and privacy, for a complete understanding of technical measures (Klymenko et al., 2022). However much of this work emphasises research findings (Ne-

gri-Ribalta et al., 2024), often overlooking the processes and methodologies integral to the research work itself (Germán et al., 2010). Key challenges remain, especially in aligning distinct work approaches and ensuring fruitful collaboration across disciplines with different foundations.

When considered independently, law and software engineering can lead to contrasting outcomes, however, if effectively integrated, they offer a powerful opportunity to develop meaningful solutions that are practically useful and reflect a proper understanding of the law (Abualhaija et al., 2025). Where the interdisciplinary field of Legal Informatics (Erdelez & O'Hare, 1997) deals with the representation of legal knowledge and reasoning in a machine-readable format, it is however mainly concerned with developing abstract models and formalisations of the law, through standards like the Legal Knowledge Interchange Format (LKIF) (Breuker et al., 2006; Hoekstra et al., 2007) or more recently LegalRuleML (Athan et al., 2015) for building and querying legal models and knowledge bases of certain pieces of laws and regulation with the goal of legal search or for reasoning on those formal representations from a logical perspective (Robaldo et al., 2020, Humphreys et al., 2021). However, legal informatics has paid limited attention to the practical implications of legal requirements on any software system to be implemented on the one hand, and also limited reflection on the actual collaboration between legal and engineering experts while building such a system on the other hand. Although researchers have engaged in interdisciplinarity from a computer science perspective (Azeem & Abualhaija, 2024; Hoepman, 2014), there remains limited guidance to support legal scholars to effectively collaborating with software engineers (Bobkowska & Kowalska, 2010; Witt et al., 2024). Generally, existing research has developed general frameworks for interdisciplinarity, particularly within social sciences (Vienni-Baptista et al., 2023; Vienni-Baptista & Thompson Klein, 2022), but the legal literature continues to face challenges in adapting these frameworks to the specificities of legal research.

The paper presents lessons learned from the work process developed in the context of a by-design interdisciplinary project, '*RegCheck: Program analysis for regulatory compliance assessment of FinTech Software*' (RegCheck). Considering this collaborative experience, the paper shares best practices for bridging the gap between law and software engineering and building confidence in engaging in interdisciplinary collaboration. Specifically, the paper presents lessons learned and best practices developed at three early stages of the research project for extracting privacy-related requirements from the GDPR and specifying them into a format that is understandable by software engineers.

This paper is structured as follows. Section 1 gives an overview of the broader regulatory and technical context within which the research project has been developed. Section 2 illustrates the steps followed and associated best practices crafted for interdisciplinary research at the intersection of law and software engineering, particularly in view of guiding legal experts, including scholars in their collaboration towards requirements extraction. Notably, these best practices have been particularly developed when building common ground and shared objectives; defining the research scope; and extracting requirements from the law. Finally, Section 3 reflects on the integration of engineering perspectives in the understanding of the law.

Section 1: RegCheck: interdisciplinarity by-design

RegCheck aims at developing a tool for assessing compliance of data collection and processing in Fintech mobile applications (e.g., mobile banking apps) with the General Data Protection Regulation (GDPR, Regulation (EU) 2016/679). Early stages of the work consisted in extracting GDPR requirements (Abualhaija et al., 2025) and to devise approaches to show evidence of the actual implementation of those requirements in the app code (Alecci et al., 2025). In this article, we are interested in describing the processes and lessons learnt from the first task. The RegCheck project is one of several projects aimed at developing automated tools to streamline compliance checking processes, thus requiring the translation of legal rules into machine-readable formats. This approach frames interdisciplinarity as instrumental, wherein one discipline – law – serves to address the needs of another – software engineering (Thompson Klein, 2017). This exercise, however, is inherently complex, necessitating the integration of expertise in software engineering with a thorough legal knowledge.

This paper shares insights into practical challenges encountered and methodological solutions developed for RegCheck in its initial phase, as an example of successful interdisciplinary collaboration between legal and software engineering. The project prioritises evidence-based research, addressing the gaps in GDPR compliance observed in practice to ensure that the project outcomes address real-world challenges in implementing data protection effectively. For instance, although data subjects' rights are fundamental in the GDPR, they remain underexplored in the software engineering literature, leading to gaps in compliance practices (Amaral Cejas et al., 2024). RegCheck's bottom-up approach focuses on specific use cases to assess the relevance and applicability of legal frameworks. This method aligns with Norström et al.'s principles for co-producing knowledge – adapted to the

GDPR framework – and remains goal-oriented (Norström et al., 2020).

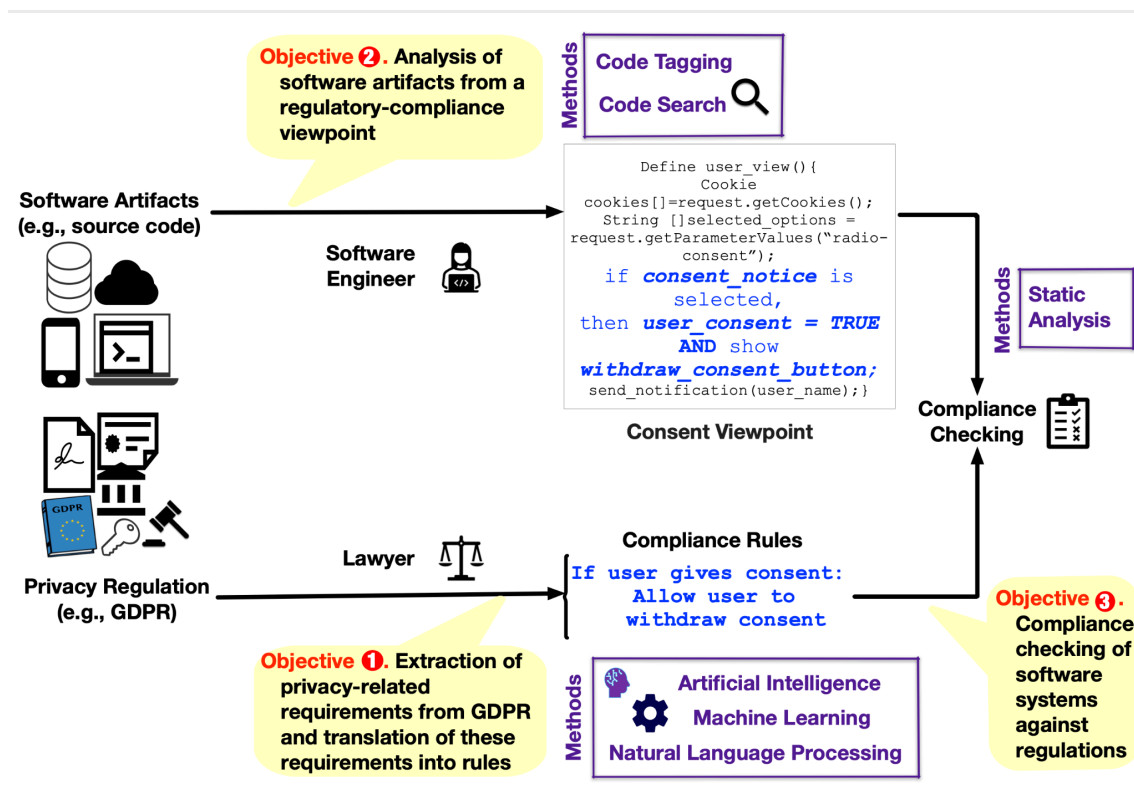


FIGURE 1: Overview of the RegCheck project

To achieve this result, RegCheck's first objective consisted in the extraction of privacy-related requirements from the GDPR for their translation into machine-readable rules. This activity required expertise across privacy regulations, software engineering, regulatory compliance, and artificial intelligence. The necessity to merge software engineering expertise with legal knowledge for the project led to developing interdisciplinary research skills by rethinking legal approaches. Law scholars traditionally focus on the analysis of legal provisions and their application to technologies, yet they may lack a sufficient understanding of technical realities. Conversely, software engineering experts possess in-depth knowledge of the technological environment but might overlook the regulatory context. The presence within the computer science team of a scholar with expertise in legal informatics, regulatory compliance and legislative analysis from a software engineering perspective exemplifies the ability of the group to navigate both legal and technical considerations. Similarly, the law team gathers scholars from various branches of law, some of whom possess industry experience, further enriching the interdisciplinary dialogue. Specifically, RegCheck involved six subject-matter experts: two scholars with substantial experience in requirement engineering and regulatory compliance; a legal informatic scholar with over five years of experience in inter-

disciplinary research combining law and software engineering; a legal scholar with specific expertise on the application of the GDPR for data subjects' protection; an early researcher with professional experience within EU institutions; and a banking law scholar with experience as a compliance consultant and Data Protection Officer in the financial sector.

Section 2: Best practices for instrumental interdisciplinarity

This section illustrates best practices crafted in developing RegCheck, progressively established over time, through an iterative, learning-by-doing process.

Section 2.1: Step 1 – Building common ground

The first, even preliminary, stage of the research project aimed at establishing a common set of goals, from practical to scientific levels. Legal scholars and software engineers may have divergent expectations about the outcomes of their work. Defining a shared research objective was crucial to make sure the whole team was sharing a similar vision of RegCheck: ensuring the compliance of software systems with ever-evolving legal regulations. However, even with this shared purpose, effective collaboration is not automatic.

Interdisciplinarity also requires reconciling differences in academic objectives, publication standards, and dissemination strategies. While legal scholars typically prioritise publishing in law journals, software engineering scholars focus on conference papers and technical journals. Aligning these divergent approaches required explicit discussions on publication strategies, to agree on possible deviations resulting from domain-specific publications, while preserving the overall coherence of the research outputs.

Additionally, transparency regarding workload distribution was fundamental to set realistic expectations, establish a balanced workflow and ensure adaptability in the choice of working tools. Interdisciplinary collaboration revealed distinct preferences shaped by disciplinary practices. Different academic communities rely on distinct platforms and software tailored to their needs. For instance, legal scholars predominantly relied on traditional WYSIWYG (What you see is what you get) word processors such as Microsoft Word, which are well-suited for linear legal drafting, while software engineering scholars preferred LaTeX for its long-established ability to parameterise the formatting of documents as well as handling complex content such as figures, tables, references, and equations. In this work, we relied on Over-

leaf, an online real-time collaborative LaTeX editor. To bridge these differences, a hybrid approach was implemented, allowing each team to work within their preferred environment while adopting interoperable formats and clear protocols for document exchange. Software engineer researchers provided the legal scholars with targeted training on the use of these tools. Nonetheless, maintaining flexibility was crucial to support productive collaboration and accommodate the varying practices and preferences of both communities. Certain tasks may not always be feasible within the technical tools preferred by software engineering scholars. In these cases, alternative methods might still be employed to ensure their smooth execution. This approach proved effective, enabling seamless sharing and integration of contributions despite the heterogeneous toolsets. Drawing on this experience, we recommend that *interdisciplinary teams embrace flexible workflows that respect disciplinary preferences*, complemented by mutual training and the use of interoperable file formats to minimise friction and enhance collaborative productivity.

Despite alignment on these preliminary considerations, smooth collaboration still requires a common understanding of key concepts (Bull & Oughton, 2006). Given that different disciplines operate with different terminologies and conceptual frameworks, creating a shared vocabulary was essential to ensure the consistency of RegCheck's outcomes. For example, the distinction between a 'right' and a 'principle' may not be immediately evident to software engineers. Conversely, terminology in software engineering might be unclear to legal scholars. For instance, while software 'verification' aims to assess whether the behaviour of (individual components of) a system is correct (meaning, it meets the specified requirements), software 'validation' aims to assess whether the solution fulfils its purpose (meaning, it meets the users' needs) (Sommerville, 2011, ch. 8, p. 208).

The challenges of establishing a shared vocabulary are particularly evident when working with the GDPR. Principles, such as data minimisation or purpose limitation (Article 5(1)(c) and (b) GDPR) are inherently abstract and require interpretation before they can be operationalised. The former requires the controller to collect only the data strictly necessary for a specific purpose. Translating the concept of 'necessity' into a machine-readable format involves defining it in measurable and context-specific terms, which can vary significantly depending on the specific setting. In banking software, the necessity of collecting a user's age or marital status might depend on different applicable legislation. Data minimisation requires the exploration of new methods that enable developers to design systems capable of collecting only the necessary data about users. From a software engineering stand-

point, this introduces challenges since collecting more data from users may lead to enhancing system functionality (Senarath & Arachchilage, 2019). Similarly, the principle of purpose limitation, which requires that personal data be used exclusively for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes, introduces further complexities when translated into technical requirements. Implementing this principle might require sophisticated data handling systems. However, these technical considerations are not explicitly addressed in legal texts, leaving room for interpretation and potential inconsistencies. By contrast, rights enshrined in the GDPR, such as the right of access (Article 15) and to data portability (Article 20), tend to be more concrete and thus more suitable for translation into technical functionalities. While these rights are not exempt from technical challenges, they are conceptually more straightforward to interpret and translate into actionable requirements than principles.

Moreover, the abstract nature and complexity of the law further complicate the extraction of precise requirements useful for software engineers. Certain legal concepts remain deliberately vague, to leave a margin of discretion in their interpretation for those who implement it, whereas software engineering typically requires precise definitions to develop technical solutions. For instance, Article 12(3) GDPR requires data controllers to respond to data subject requests ‘without undue delay’. In this case, the legislator intentionally leaves open to interpretation the length of an ‘undue delay’, whose length is to further be decided or influenced by national laws and guidelines from competent supervisory authorities. This deliberate vagueness raises the question of whether legal standards could, or should, be expressed directly in code. Encoding laws might allow lawmakers to shift from broad standards to more fine-grained rules, potentially capturing legislative intent more precisely and operationalising it through algorithms. While this proposition holds promise in domains where legal reasoning can be reduced to quantifiable metrics, such as tax law or certain sentencing frameworks, it remains a considerable challenge in areas like privacy and data protection. Many provisions of the GDPR retain an inherently open-textured character, reflecting deliberate legislative choices to preserve flexibility and contextual sensitivity. This is especially true within multi-level legal systems like the EU, where member states retain autonomy in interpreting and implementing certain standards. Consequently, while some aspects of the law may be formalised into code, a margin of discretion is likely to remain indispensable, particularly in legal disciplines that engage with fundamental rights. From a software engineering perspective, however, even open-textured legal standards must eventually be translated into operational parameters. In the case of

‘undue delay’, this means deriving an explicit duration from the interpretative guidance provided by supervisory authorities and national courts. This translation process underscores the tension between legal indeterminacy and technical determinacy, and highlights the need for interdisciplinary collaboration to bridge the gap between normative frameworks and computational implementation.

Likewise, on the one hand, an ‘incomplete requirement’ for legal experts indicates that a legal provision can only be specified once its application context is established, including the jurisdiction where the software is deployed and its specific legal framework. For example, Article 23 GDPR allows member states to restrict data subject rights under certain conditions. This is particularly relevant in the context of Anti-Money Laundering (AML) compliance legislation, such as the EU Anti-Money Laundering Regulation (AMLR; Regulation (EU) 2024/1624). While the GDPR guarantees a general right to explanation (Articles 15 and 22), the AMLR introduces significant restrictions on the provision of explanations, particularly regarding suspicious transaction reports (STRs), as customers are not entitled to any explanations concerning these reports (Article 76). Moreover, the transposition of the Fifth Anti-Money Laundering Directive (Directive (EU) 2018/843) has led to the exclusion of data subjects’ rights in the context of AML compliance in certain member states. For instance, Article 65 of the Belgian Law on the Prevention of Money Laundering and Terrorist Financing and the Restriction of the Use of Cash (Loi no. 2017013368) explicitly excludes data subjects’ rights when personal data is processed under this law.

Requirements’ incompleteness in software engineering, on the other hand, is a long-standing challenge in the requirements engineering field (Arora et al., 2019). The literature distinguishes between two types of completeness (Zowghi & Gervasi, 2002): internal completeness is concerned with ensuring that no necessary information is missing from any specified requirement, e.g., explicit duration threshold instead of vague values such as an “undue delay”, whereas external completeness is concerned with ensuring that all information relevant for developing a system are captured by the set of specified requirements, e.g., forgetting to include the requirements related to the time to answer that is put on the controller regarding the request and the time information requirement to the subject. In our context, several requirements were created to operationalise such scenarios. In particular, one requirement would state that “*The system shall confirm receipt of the request mentioning that the delay for the delivery of the data starts running*”, and an additional requirement would expose a delay of one month with a follow-up process: “*The system shall, within one month of an access request, either (1) send the*

requested data or (2) send a communication regarding the reasons of the delay or (3) inform the data subject of the reasons for not taking action and on the possibility of lodging a complaint with a supervisory authority and seeking a judicial remedy". In this scenario regulated under Article 12(3) of the GDPR, explicitly acknowledging the request is the beginning of the response process, which provides a concrete interpretation of the "without undue delay" (direct acknowledgement of the request), while using the maximum threshold set by the regulation (30 days) as delay and putting a requirement on the controller to answer within that period.

Differently, some concepts are common to both disciplines, yet their meaning diverges, leading to potential confusion. A prominent example is the much-debated notion of 'artificial intelligence' (AI). The EU definition provided in the AI Act (Article 3(1), Regulation (EU) 2024/1689) is notably broad, encompassing a wide range of AI models and systems, including basic algorithms (Ruscheimer, 2023, pp. 363–364; Veale & Borgesius, 2021, p. 109). Such a broad definition does not fully capture the various categories of AI methodologies from a technical standpoint. While the regulation presents an all-encompassing definition, the AI field maintains a fine-grained distinction between various techniques and approaches, such as rule-based systems, machine learning, deep learning, and symbolic AI. This misalignment may lead to regulatory ambiguities introducing compliance challenges in practice. Therefore, it becomes crucial for effective interdisciplinary collaboration to ensure that all parties involved in the project use concepts consistently or at least understand the standpoint of each discipline.

Furthermore, various approaches are possible to legal compliance, which requires the team to agree on which perspective will be adopted, for example, a formal or a risk-based compliance approach. The former involves systematically identifying discrepancies between a certain reality and the relevant legal framework, and to evaluate such discrepancies based on the authority and hierarchy of the applicable rules. Under such an approach, the extent to which deviations are permissible depends on the normative force of the respective source. For instance, primary legislation normally carries greater legal weight than case law or soft law guidelines. The objective is thus to assess the seriousness of regulatory breaches based on the formal legal nature and hierarchical status of the applicable rule. By contrast, a risk-based approach to legal compliance extends beyond the formal legal nature of the applicable norms, incorporating a range of different elements that may be harder to capture, including for instance the context and severity of the breach, the political priorities and sanctioning practices of the relevant enforcement authorities. Such an approach seeks to evaluate breaches based on the risk of non-com-

pliance, that is, the likelihood and impact of negative events (mainly, fines) that could result from non-compliance. A risk-based approach does not mean that low-risk cases of non-compliance should not be flagged, but rather that an automated assessment of the seriousness of a breach would incorporate factors that go beyond the formal nature of the breach while transparently informing the relevant decision-makers about the integration of those factors - for instance, the likelihood and severity of a sanction - and how they have been taken into account in the assessment. Thus, agreeing on the intended approach to legal compliance is important to define a methodological framework for the interpretation and ranking of the legal norms at stake. In our case, legal compliance is understood as the alignment of the software with a legal framework, including traditional legal rules (primarily law and case law) and additional guidance found in soft law, thus relying on a formal approach to compliance. Legal scholars should not limit themselves to the study of traditional legal sources alone to obtain complete requirements when working with software engineers. This holistic understanding ensures a more comprehensive and precise alignment of technical solutions with the applicable legal norms.

Having established mutual understanding and alignment across disciplines, the subsequent step involved delineating the scope of the research.

Section 2.2: Step 2 – Defining the research scope: balancing ideal goals with pragmatic choices

The second early step of the research project concerned the scope of RegCheck. In fact, the research boundaries might need to be refined to accommodate interdisciplinary thinking and project feasibility. This step involved both the engineering and legal teams. The original concept paper of RegCheck outlined a broad scope focused on the GDPR in its entirety. The GDPR is a comprehensive legal framework that encompasses several fundamental elements, including the core principles of data protection, the rights of data subjects, the obligations for data controllers and processors, as well as provisions for its enforcement and penalties.

In our case, however, the decision was taken to focus exclusively on data subjects' rights. This conclusion was reached based on both legal and software engineering considerations, with a strong emphasis on practical and social implications (Sørum & Presthus, 2020). From a legal perspective, data subjects' rights are fundamental to the empowerment of individuals, as they provide the necessary tools for individuals to assert control over their data. The practical application of these rights directly impacts the extent to which data subjects exercise meaningful informational

self-determination (Pisani, 2024). The right to data portability, for instance, enhances users' autonomy by enabling individuals to transfer their data between service providers. Yet, despite its considerable potential to position individuals at the centre of the data economy, this right remains underused and confined to a limited number of sectors (European Commission, 2020). From a software engineering standpoint, the focus on data subjects' rights offers a manageable entry point, given the well-defined nature of individual rights compared to broader regulatory obligations, which as mentioned above often involve more abstract principles. Moreover, data subjects' rights remain underexplored in the literature (Negri-Ribaltta et al., 2024). This research is particularly evident regarding the right to access personal data. While scholars have analysed the output of a data access request (Bowyer et al., 2022; Bufalieri et al., 2020; Pins et al., 2022; Pöhn et al., 2023; Veys et al., 2021), relatively little attention has been devoted to its technical implementation.

From a practical standpoint, the complexity of the GDPR posed additional constraints on the scope of the project. It became clear early in the process that manually extracting all relevant legal requirements from the regulation within a few months would be unfeasible. This realisation prompted the team to opt for a more targeted approach, focusing on the operationalisation of specific rights rather than attempting to address the full breadth of the GDPR. This resulted in a restrictive delimitation of the research scope. Particularly, the decision was taken to narrow it down to only two specific data subjects' rights: the rights to access and to portability. The choice stemmed from the practical difficulty of addressing the full range of subjects' rights, which also includes the right to rectification, the right to erasure, the right to restrict processing, the right to object and the right not to be subject to a decision based solely on automated processing, in light of the number of sources to analyse (see Figure 2).

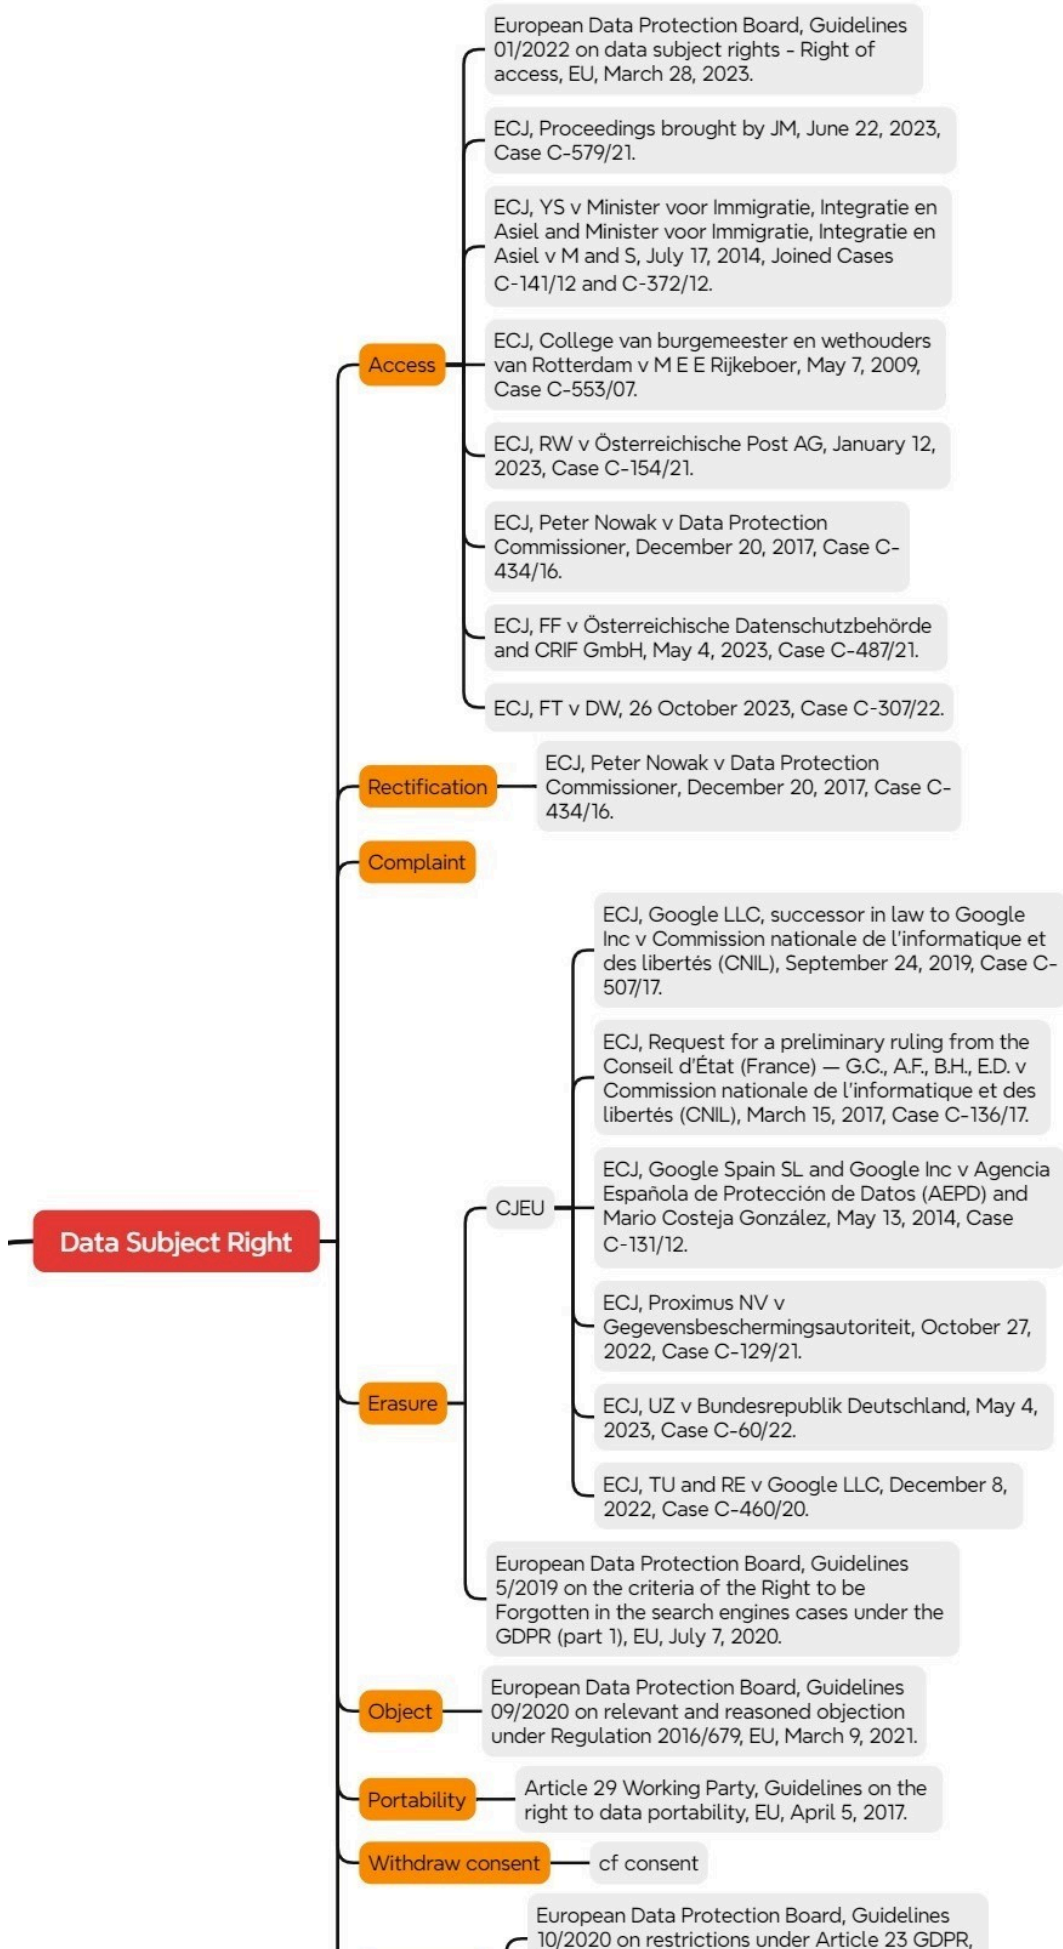


FIGURE 2: Overview of data subjects' rights and relevant sources

The selection of these two specific data subjects' rights was driven by their foundational role in the broader data protection framework. Without access to their data, individuals would be unable to assess its accuracy or challenge its processing. Notably, the right to access is fundamental for the exercise of other rights, as it serves as a prerequisite for the effective exercise of the right to rectification, erasure, and data portability. Despite its centrality, the right to access has not been extensively discussed in the existing literature, particularly regarding its technical implementation (Negri-Ribalta et al., 2024). The complexity of this seemingly straightforward right goes beyond what is stated in the GDPR, making appropriate requirements extraction and implementation a challenge for software engineers. Regarding the right to portability, it is particularly relevant to software applications, as it allows the seamless transfer of personal data when individuals transition between online services or mobile applications offering similar functionalities. Furthermore, its implementation is expected to be reinforced with the EU Data Act (Regulation (EU) 2023/2854), which aims to establish clear and fair rules for accessing and using data within the European data economy. Investigating the right to portability was deemed particularly valuable, as most of the existing literature focused on its theoretical implications (De Hert et al., 2018), while limited attention has been given to its technical implementation. Moreover, it is highly technical in nature, making it particularly well-suited for translation into specific technical requirements.

Simultaneously, the scope of the research project was extensively delineated, particularly regarding the selection of relevant legal sources to be considered. It became evident early on that the text of the GDPR alone did not provide a sufficiently detailed basis for deriving technically actionable compliance requirements. While the GDPR establishes the overarching legal framework, it often articulates principles in abstract or indeterminate terms, leaving substantial room for interpretation. From a technical perspective, software engineers required more granular and operational knowledge than the GDPR itself could offer. At the same time, software engineers faced challenges in independently identifying additional relevant sources without the expertise of legal scholars. This gap highlighted the necessity of drawing upon additional interpretative sources to bridge the distance between legal norms and their implementation in software systems. Accordingly, the research team agreed to integrate a broader set of sources into the analysis. These sources were identified in guidelines issued by the European Data Protec-

tion Board (EDPB), along with relevant academic literature. While these sources do not carry the binding authority of legislation or case law, they nonetheless provide essential interpretative guidance, offering a detailed understanding of legal provisions necessary for their technical implementation. For instance, Article 12(3) of the GDPR allows extending the answer deadline in the case of an access request up to three months in “exceptional cases”. The EDPB clarifies what could amount to an exceptional case, by introducing the term “complex cases”, for instance given the amount of data or additional steps necessary to render the data intelligible (EDPB, 2022, p. 51). Therefore, this requirement had to be connected to two sources. Furthermore, the ECJ plays a crucial role in interpreting the GDPR, clarifying ambiguities and shaping its practical implementation through its binding interpretations. Consideration of the case law of the ECJ, whose decisions serve as a primary source of law, has thus been deemed indispensable in ensuring that technical implementation aligns with the evolving legal framework. As an example, Article 15(3) provides that “The controller shall provide a copy of the personal data undergoing processing”, which has been extracted as one requirement under the right to access. Yet, this requirement had to be complemented by ECJ case law, which clarified that such a copy must be a “faithful reproduction or transcription of [the] original” data and “contain all the personal data undergoing processing” (ECJ, 2023, §§ 21, 32).

In conclusion, a critical lesson for interdisciplinary collaboration is that defining and refining the research scope is essential to reconcile it with practical feasibility. This process requires iterative negotiation among domain experts to identify manageable focus areas that reflect both legal significance and technical constraints. Moreover, exclusive reliance on primary law sources is often insufficient; integrating interpretative guidelines, case-law, and expert knowledge is critical to ensure accurate technical implementation.

Section 2.3: Step 3 – Extraction of requirements

Once the scope of the project was clearly defined, the final step involved the extraction of legal requirements from the selected provisions of the GDPR.

Firstly, this third step involved an in-depth legal analysis of the chosen topics from the selected sources, a task that fell primarily within the expertise of legal scholars, following traditional legal research methodologies.

To improve the transfer of legal knowledge in such a technical context, legal scholars relied on the adoption of visual presentation techniques (see Figure 3).

Legal texts often present dense and highly structured information that can be challenging to navigate. Certain branches of law, such as tax law, which relies on strict rules notably to calculate tax or penalty amounts (Nay et al., 2024; Waidelich et al., 2023), are easier to understand for programming compared to the one analysed in this paper, where more complex interdisciplinary communication is needed. Visual representations, such as flowcharts, diagrams and structured tables, were employed to break down complex legal provisions into more accessible formats. Legal design not only facilitated comprehension for software engineering experts but also allowed legal experts to better conceptualise the technical implications of their interpretations. The incorporation of visual elements enhanced interdisciplinary communication, reduced ambiguities, and fostered a shared understanding of legal provisions, ultimately strengthening the translation of legal norms into technical requirements (Berger-Walliser et al., 2017; Ducato & Strowel, 2021; Perry-Kessarais, 2019).

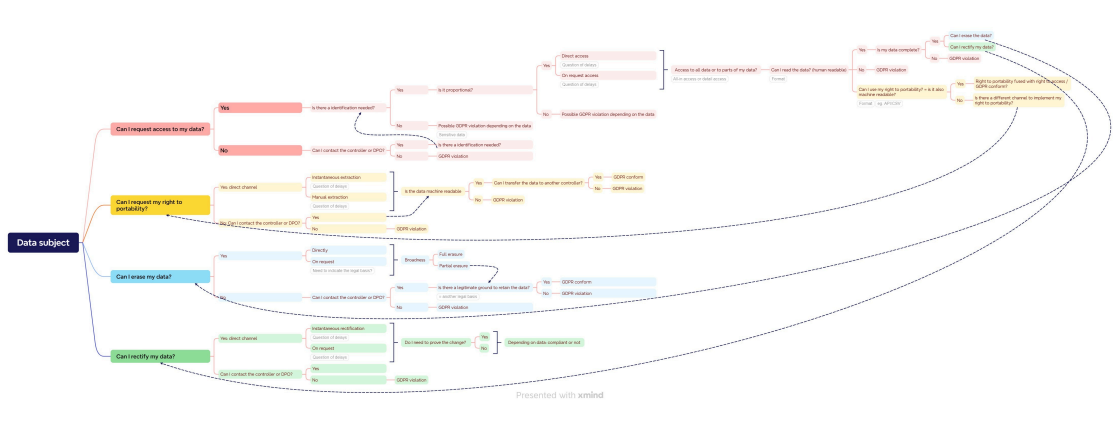


FIGURE 3: Overview of the functioning of data subjects’ rights

Then, the third step involved an intermediary stage in which legal provisions were translated into technical requirements. A crucial aspect of this process was the establishment of clear drafting instructions and a structured format for documenting requirements, without having the burdens of existing requirements template that tend to frame requirements with respect to properties (functional, non functional, conditional, temporal, etc.) that do not always fit well with legal/regulatory requirements. Here, the most important consideration is to have a flexible approach, using plain natural language to enable all parties to analyse, comment, and agree on the drafted requirements with traceability to the sources. Ensuring traceability to the legal source for each requirement was particularly valuable in documenting the legal requirements in a verifiable manner. At this stage of the work, all the drafted requirements come as a list of independent requirements. However, laws and regulations often articulate their norms while stating applicability conditions,

exceptions, prevalence over other norms. Both teams acknowledge the interconnected nature of legal norms and requirements. However, for analysis purposes, it was decided to avoid handling this complexity layer and to leave it to a later stage and future work. This approach was particularly challenging for legal experts accustomed to working with structured legal frameworks. Given the transversal nature of certain provisions, which apply uniformly to both rights, some legal requirements – such as those concerning the information, communication and modalities for the exercise of the rights (Article 12 GDPR) – were duplicated across both the right of access and the right to data portability, to ensure both rights are completely and equally covered with their requirements if taken individually.

From a software engineering perspective, the focus of this phase was on maintaining a uniform and systematic format, ensuring the use of precise and unambiguous language to the greatest extent possible. On the legal side, a major concern was to prevent the oversimplification of legal norms in the pursuit of computability. Legal provisions often contain compound obligations, open-textured terms, and contextual dependencies that resist straightforward operationalisation. For example, the right to data portability encompasses not only the obligation to provide personal data upon request but also conditions related to the type of data, the means of processing, and the potential impact on the rights of others. Extracting requirements from such provisions demands careful legal interpretation to preserve their normative meaning and ensure faithful implementation. The integrity of legal reasoning had to be preserved, avoiding any distortions that could arise from attempting to force legal provisions into rigid technical structures. For instance, a requirement derived from Article 20(1) of the GDPR requires the use of a “structured, commonly used and machine-readable format” to respond to a portability request, but it omits the format details, which may vary across application contexts and regions. To guide the GDPR interpretation, the Article 29 Working Party guidelines merely refers to examples, such as XML, JSON or CSV (Article 29 Working Party, 2017).

This delicate balance between legal precision and technical applicability defined the methodological approach of this phase. The GDPR itself offers limited guidance on how to interpret key terms, such as ‘data provided by the data subject’ in Article 20 GDPR. To address this ambiguity, the Article 29 Working Party – whose guidelines were later endorsed by the EDPB – clarifies that this notion includes not only data actively submitted by the user (e.g. profile information) but also observed data, such as search history, transaction records, or access logs (Article 29

Working Party, 2017). Additionally, it was particularly useful that the EDPB provided, for the right to access, activity diagrams (see Figure 4), which could be translated into requirements (EDPB, 2023). These interpretative resources were essential to fill the normative gaps in the legal text and to ensure that technical implementations remain faithful to the intent and scope of the regulation.

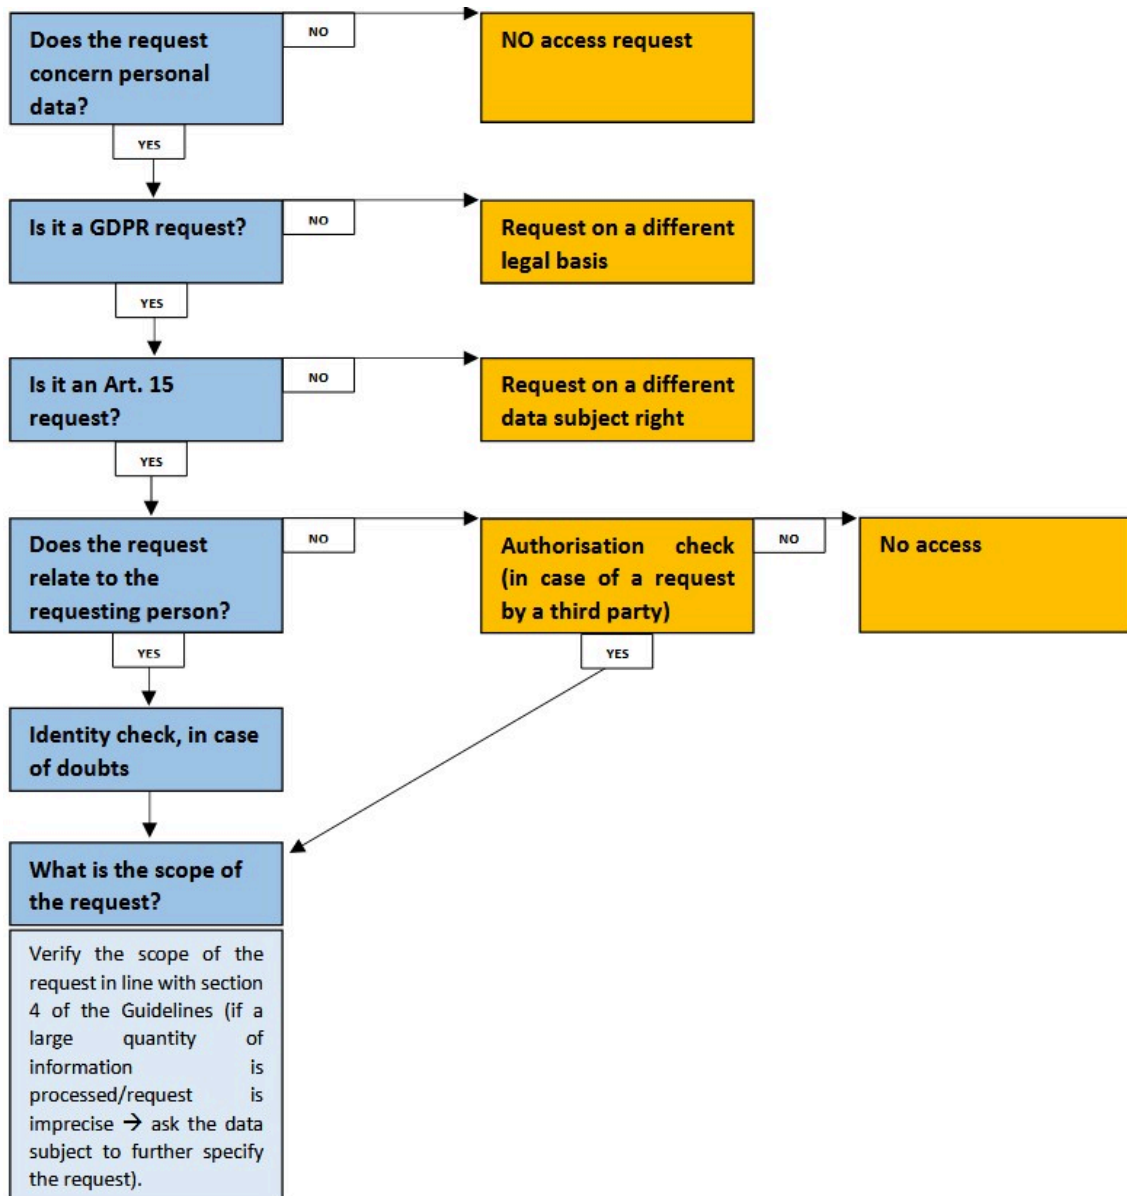


FIGURE 4: Activity diagram “Step 1: How to interpret and assess the request?”, (EDPB, 2023, p. 61)

The final phase of the process involved fine-tuning the drafted requirements, to maintain both legal soundness and technical implementability. This step mostly involved the legal scholars and the researcher in legal informatics. Such a hybrid profile in the team proved crucial for bridging law to software engineering, especially to negotiate among the demands of both disciplines. For example, require-

ments stated from the point of view of data subjects had to be reformulated from the point of view of the system, restrictions had to be reformulated as preconditions, requirements applicable to both the right of access and the right to portability had to be duplicated to appear in both lists (e.g. deriving from Article 12 of the GDPR), and requirements which represented best practices extracted from the EDBP guidelines rather than legal obligations from the GDPR and case law had to be labeled as such. An example of best practice requirement is the following: “The system shall allow the data subject to request access to the data undergoing processing related to them in a clear and visible manner” (Article 29 Working Party, 2013, p. 24).

To summarise, our interdisciplinary collaboration indicated that it is neither necessary nor efficient for all team members to be involved in every stage of the requirement extraction process. Instead, a structured and mixed approach was adopted, clearly delineating the steps where legal experts and software engineering researchers work independently and those requiring close collaboration among them. This structured division of tasks ensures both efficiency and accuracy while acknowledging the distinct expertise required at different stages of the process. It also highlights the importance of clearly defining expectations at each stage of the process and of recognising that not all team members will be able to contribute equally at every phase. In our case, these methodological distinctions were not clearly established from the outset, which led to frequent back-and-forth exchanges between the legal and software engineering teams. Although this iterative process may appear time-consuming in retrospect, it was ultimately beneficial to foster a mutual understanding and establish a shared conceptual foundation. To mitigate avoidable inefficiencies, interdisciplinary teams are advised to implement coordination mechanisms at the outset. These may include joint workshops to establish common concepts and definitions; explicit documentation of roles, responsibilities, and deliverables; and consensus on collaborative tools and communication protocols.

A key best practice emerging from our experience is the value of drafting an initial set of requirements as a team, with input from both legal and technical perspectives, before proceeding to the full-scale extraction process. In this work, we went through small tutorials to requirements drafting based on concrete existing examples, for instance legal requirements from past work of some of the authors (Amaral Cejas, 2023). The goal was to introduce and illustrate good practices regarding legal requirements specification and kickstart the work while avoiding possible misunderstandings that would need to be fixed after. Providing concrete examples

of previously extracted requirements – whether from the same project or from existing literature – can help ensure that all researchers share a common understanding of the expected output. By setting clear methodological expectations from the beginning, future interdisciplinary collaborations can enhance efficiency, reduce miscommunication, and facilitate a smoother transition between legal analysis and technical implementation.

The final step is the implementation of requirements within the compliance software, where they will be tested and operationalised.

Section 3: A discussion for critical interdisciplinarity

Beyond the instrumental interdisciplinary objective of RegCheck, its implementation also fostered a critical reflection on interdisciplinarity from a legal perspective (Thompson Klein, 2017). It resulted in legal scholars questioning what compliance is and interrogating the role of the law and legal practice in achieving compliance. This involves questioning whether compliance refers to a situation that is in conformity with the ‘law in the books’ or whether it embeds consideration of the ‘law in practice’ as expressed through best practices and industry standards.

That questions of whether the achievement of compliance can rely solely on a formal legal interpretation and application of the relevant rules, mainly considering their textual wording and legal status, ranging from hard law such as legislative provisions (in our case, the GDPR) to case law (the ECJ) and soft law such as non-binding recommendations and guidelines (for instance, from the EDPB). By contrast, an alternative, contextual, approach would consider not only the applicable rules but also the risk of non-compliance – that is, the likelihood and impact of a regulatory breach. Such a broader approach requires one to consider factors beyond the legal provision at stake, such as the context and degree of severity of breaches, the political priorities and sanctioning practices of competent enforcement authorities, or the presence of mitigating or aggravating circumstances.

Addressing those issues in the context of an interdisciplinary research project does not serve a purely theoretical purpose. They help to develop a common understanding of the objective of the interdisciplinary research project and the legal methodology that is the most suitable to achieve that objective. In fact, the engineering of a compliance tool should ultimately serve a certain audience, which could be corporate decision-makers, compliance officers, enforcement authorities, judges, etc. Depending on the audience concerned, a different legal methodology may be required. For instance, when boards or management committees decide on

the allocation of resources to compliance measures, they may care more about the (financial) risk than about the purely legal degree of compliance. Thus, the objective of the intended compliance tool should play a role in determining the appropriate legal methodology. For instance, one must decide in advance whether the tool should inform corporate decision-makers on compliance risks, or should it address Data Protection Officers, Data Protection Authorities, or perhaps the judicial authorities in the context of liability claims. While the objective and audience of a certain compliance tool should be kept in mind when determining the appropriate methodology, one should, however, keep in mind this involves a risk of pre-empting human judgments that should arguably be left to the addressees of the tool's output. Indeed, designing an automated compliance assessment to produce value judgments that are not objectively legal involves the risk of influencing decision-making beyond purely legal considerations, thereby potentially affecting the (financial) risk management of the firm in a way that is not aligned with its risk appetite, or that does not adequately inform decision-makers of the rationale behind such value judgments. Those risks should also be taken into consideration when determining the legal methodology that underpins the output of a certain compliance tool, as they may trigger requirements in terms of transparency and explainability.

Conclusion

In conclusion, our experience underscores the critical need to integrate knowledge from diverse disciplines to foster innovative solutions. RegCheck, which aimed to develop a tool for assessing GDPR compliance in FinTech mobile applications, demonstrates this necessity. We presented lessons learned and best practices elaborated during three early stages of an interdisciplinary research project: (1) building common ground and shared objectives, (2) defining the research scope by balancing ideal goals with pragmatic considerations, and (3) extracting requirements from the law using legal design methods. This structured approach delineates the steps where legal experts and software engineers can work independently and those requiring close collaboration.

The example of the RegCheck project offers valuable insights into addressing practical challenges of interdisciplinary research at early stages. The project underscored the importance of a shared vocabulary and consistent approaches to privacy and compliance across disciplines. This echoes the broader need for scholars from different fields to reconcile divergent terminologies and conceptual frameworks to ensure consistency in research outcomes. The necessity of scoping the re-

search to accommodate interdisciplinary thinking and project feasibility reflects the challenges of addressing multifaceted problems that cannot be adequately addressed in disciplinary silos. By focusing on specific data subject rights and relevant legal sources, the project team was able to navigate the complexities of the GDPR and achieve tangible results. Through interdisciplinary collaboration, scholars acquired essential knowledge from each other's disciplines, enabling them to tackle complex compliance issues more effectively. This aligns with the broader goal of interdisciplinarity to foster innovation and achieve shared objectives for the benefit of society. Finally, RegCheck contributed to a richer and more dynamic construction of legal knowledge, by challenging traditional disciplinary boundaries and epistemological assumptions, leading to new insights and a deeper understanding of complex concepts.

In essence, RegCheck demonstrates how a structured and mixed approach to requirement extraction can bridge the gap between legal analysis and technical implementation, offering a practical framework for fostering meaningful interdisciplinary cooperation. By sharing our experiences and lessons learned, we aim to empower legal and software engineering scholars to engage confidently in collaborative research, advancing both academic inquiry and societal impact.

References

- Abualhaija, S., Ceci, M., Sannier, N., Bianculli, D., Lannier, S., Siclari, M., Voordeckers, O., & Tosza, S. (2025). LLM-assisted extraction of regulatory requirements: A case study on the GDPR. *Proceedings of the 33rd IEEE International Requirements Engineering Conference*.
- Alecci, M., Sannier, N., Ceci, M., Abualhaija, S., Samhi, J., Bianculli, D., Bissyande, T. F. D. A., & Klein, J. (In press). Toward LLM-driven GDPR compliance checking for Android apps. *Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering (FSE Companion '25)*. ACM - Association for Computing Machinery.
- Amaral, O., Azeem, M. I., Abualhaija, S., & Briand, L. C. (2022). *NLP-based automated compliance checking of data processing agreements against GDPR* (Version 2). arXiv. <https://doi.org/10.48550/ARXIV.2209.09722>
- Apostel, L., Berger, G., Briggs, A., & Michaud, G. (1972). Interdisciplinarity: Problems of teaching and research in universities. *OECD*.
- Arora, C., Sabetzadeh, M., & Briand, L. C. (2019). An empirical study on the potential usefulness of domain models for completeness checking of requirements. *Empirical Software Engineering*, 24(4), 2509–2539. <https://doi.org/10.1007/s10664-019-09693-x>
- Article 29 Working Party. (2013). *Opinion 02/2013 on apps on smart devices*. EU.
- Article 29 Working Party. (2017). *Guidelines on the right to data portability under Regulation 2016/*

679. EU.

Athan, T., Governatori, G., Palmirani, M., Paschke, A., & Wyner, A. (2015). LegalRuleML: Design principles and foundations. In W. Faber & A. Paschke (Eds), *Reasoning Web. Web Logic Rules* (Vol. 9203, pp. 151–188). Springer International Publishing. https://doi.org/10.1007/978-3-319-21768-0_6

Azeem, M. I., & Abualhaija, S. (2024). A multi-solution study on GDPR AI-enabled completeness checking of DPAs. *Empirical Software Engineering*, 29(4), 96. <https://doi.org/10.1007/s10664-024-10491-3>

Baptista, B. V., & Klein, J. T. (2022). *Institutionalizing interdisciplinarity and transdisciplinarity: Collaboration across cultures and communities* (1st edn). Routledge. <https://doi.org/10.4324/9781003129424>

Board, E. D. P. (2023). Guidelines 01/2022 on data subject rights—Right of access. *EU*.

Bobkowska, A., & Kowalska, M. (2010). On efficient collaboration between lawyers and software engineers when transforming legal regulations to law-related requirements. *2nd International Conference on Information Technology*, 105–109.

Bowyer, A., Holt, J., Go Jefferies, J., Wilson, R., Kirk, D., & David Smeddinck, J. (2022). Human-GDPR interaction: Practical experiences of accessing personal data. *CHI Conference on Human Factors in Computing Systems*, 1–19. <https://doi.org/10.1145/3491102.3501947>

Bracken Bull, L. J., & Oughton, E. A. (2006). ‘What do you mean?’ The importance of language in developing interdisciplinary research. *Transactions of the Institute of British Geographers*, 31(3), 371–382.

Breuker, J., Boer, A., Hoekstra, R., & Berg, K. (2006). Developing content for LKIF: Ontologies and frameworks for legal reasoning. *Proceedings of the 19th Annual Conference on Legal Knowledge and Information Systems, JURIX’06*, 169–174.

Bufalieri, L., Morgia, M. L., Mei, A., & Stefa, J. (2020). GDPR: When the right to access personal data becomes a threat. *2020 IEEE International Conference on Web Services (ICWS)*, 75–83. <https://doi.org/10.1109/ICWS49710.2020.00017>

Cejas, O. A., Sannier, N., Abualhaija, S., Ceci, M., & Bianculli, D. (2024). *GDPR-relevant privacy concerns in mobile apps research: A systematic literature review* (No. arXiv:2411.19142). arXiv. <https://doi.org/10.48550/arXiv.2411.19142>

De Hert, P., Papakonstantinou, V., Malgieri, G., Beslay, L., & Sanchez, I. (2018). The right to data portability in the GDPR: Towards user-centric interoperability of digital services. *Computer Law & Security Review*, 34(2), 193–203. <https://doi.org/10.1016/j.clsr.2017.10.003>

Edwards, L., & Veale, M. (2018). Enslaving the algorithm: From a “right to an explanation” to a “right to better decisions”? *IEEE Security & Privacy*, 16(3), 46–54. <https://doi.org/10.1109/MSP.2018.2701152>

Erdelez, S., & O’Hare, S. (1997). Legal informatics: Application of information technology in law. *Annual Review of Information Science and Technology*, 32, 367–402.

German, D. M., Webber, J. H., & Di Penta, M. (2010). Lawful software engineering. *Proceedings of the FSE/SDP Workshop on Future of Software Engineering Research*, 129–132. <https://doi.org/10.1145/1882362.1882390>

- Hoekstra, R., Breuker, J., Bello, M. D., & Boer, A. (2007). The LKIF core ontology of basic legal concepts. *Proceedings of the 2nd Workshop on Legal Ontologies and Artificial Intelligence Techniques, LOAIT'07*, 43–63.
- Hoepman, J.-H. (2014). Privacy design strategies. In N. Cuppens-Bouahia, F. Cuppens, S. Jajodia, A. Abou El Kalam, & T. Sans (Eds), *ICT Systems Security and Privacy Protection* (Vol. 428, pp. 446–459). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-55415-5_38
- Hoess, A., Hoess, A., Pocher, N., Roth, T., & Fridgen, G. (2024). Towards a design science research process for legal compliance by design. *PACIS 2024 Proceedings*. https://aisel.aisnet.org/pacis2024/rack04_dessci/track04_dessci/3
- Humphreys, L., Boella, G., Van Der Torre, L., Robaldo, L., Di Caro, L., Ghanavati, S., & Muthuri, R. (2021). Populating legal ontologies using semantic role labeling. *Artificial Intelligence and Law*, 29(2), 171–211. <https://doi.org/10.1007/s10506-020-09271-3>
- Klein, J. T. (2017). Typologies of interdisciplinarity: The boundary work of definition. In R. Frodeman (Ed.), *The Oxford Handbook of Interdisciplinarity* (2nd edn, pp. 21–34). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780198733522.013.3>
- Klymenko, O., Kosenkov, O., Meisenbacher, S., Elahidoost, P., Mendez, D., & Matthes, F. (2022). Understanding the implementation of technical measures in the process of data privacy compliance: A qualitative study. *Proceedings of the 16th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement*, 261–271. <https://doi.org/10.1145/3544902.3546234>
- Nay, J. J., Karamardian, D., Lawsky, S. B., Tao, W., Bhat, M., Jain, R., Lee, A. T., Choi, J. H., & Kasai, J. (2024). Large language models as tax attorneys: A case study in legal capabilities emergence. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 382(2270), 20230159. <https://doi.org/10.1098/rsta.2023.0159>
- Negri-Ribalta, C., Lombard-Platet, M., & Salinesi, C. (2024). Understanding the GDPR from a requirements engineering perspective—A systematic mapping study on regulatory data protection requirements. *Requirements Engineering*, 29(4), 523–549. <https://doi.org/10.1007/s00766-024-00423-4>
- Norström, A. V., Cvitanovic, C., Löf, M. F., West, S., Wyborn, C., Balvanera, P., Bednarek, A. T., Bennett, E. M., Biggs, R., De Bremond, A., Campbell, B. M., Canadell, J. G., Carpenter, S. R., Folke, C., Fulton, E. A., Gaffney, O., Gelcich, S., Jouffray, J.-B., Leach, M., ... Österblom, H. (2020). Principles for knowledge co-production in sustainability research. *Nature Sustainability*, 3(3), 182–190. <https://doi.org/10.1038/s41893-019-0448-2>
- Pins, D., Jakobi, T., Stevens, G., Alizadeh, F., & Krüger, J. (2022). Finding, getting and understanding: The user journey for the GDPR'S right to access. *Behaviour & Information Technology*, 41(10), 2174–2200. <https://doi.org/10.1080/0144929X.2022.2074894>
- Pisani, G. (2024). The right to self-determination in the digital platform economy. *Computer Law & Security Review*, 53, 105964. <https://doi.org/10.1016/j.clsr.2024.105964>
- Pohl, C., Truffer, B., & Hirsch-Hadorn, G. (2017). Addressing wicked Problems through transdisciplinary research. In R. Frodeman (Ed.), *The Oxford Handbook of Interdisciplinarity* (2nd edn, pp. 319–331). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780198733522.013.26>
- Pöhn, D., Mörsdorf, N., & Hommel, W. (2023). Needle in the haystack: Analyzing the right of access according to GDPR Article 15 five years after the implementation. *Proceedings of the 18th International Conference on Availability, Reliability and Security*, 1–10. <https://doi.org/10.1145/36001>

60.3605064

Robaldo, L., Bartolini, C., Palmirani, M., Rossi, A., Martoni, M., & Lenzini, G. (2020). Formalizing GDPR provisions in reified I/O logic: The DAPRECO knowledge base. *Journal of Logic, Language and Information*, 29(4), 401–449. <https://doi.org/10.1007/s10849-019-09309-z>

Ruscheimer, H. (2023). AI as a challenge for legal regulation—The scope of application of the artificial intelligence act proposal. *ERA Forum*, 23(3), 361–376. <https://doi.org/10.1007/s12027-022-00725-6>

Senarath, A., & Arachchilage, N. A. G. (2019). A data minimization model for embedding privacy into software systems. *Computers & Security*, 87, 101605. <https://doi.org/10.1016/j.cose.2019.101605>

Sommerville, I. (2011). *Software engineering* (9th edn). Addison-Wesley.

Sørum, H., & Presthus, W. (2021). Dude, where's my data? The GDPR in practice, from a consumer's point of view. *Information Technology & People*, 34(3), 912–929. <https://doi.org/10.1108/ITP-08-2019-0433>

The Court of Justice of the European Union. (n.d.). *Judgment of the court (first chamber) of 4 May 2023: FF v Österreichische Datenschutzbehörde and CRIF GmbH* (No. Case C-487/21). The Court of Justice of The European Union. <https://eur-lex.europa.eu/legal-content/en/TXT/?uri=CELEX:62021CJ0487>

Veale, M., & Zuiderveen Borgesius, F. (2021). Demystifying the draft EU Artificial Intelligence Act—Analysing the good, the bad, and the unclear elements of the proposed approach. *Computer Law Review International*, 22(4), 97–112. <https://doi.org/10.9785/cr-2021-220402>

Veys, S., Serrano, D., Stamos, M., Herman, M., Reitinger, N., Mazurek, M. L., & Ur, B. (2021). *Pursuing usable and useful data downloads under GDPR/CCPA access rights via co-design* (pp. 217–242). <http://www.usenix.org/conference/soups2021/presentation/veys>

Vienni-Baptista, B., Fletcher, I., Lyall, C., & Ohlmeyer, J. H. (Eds). (2023). *Foundations of interdisciplinary and transdisciplinary research: A reader*. Bristol University Press.

Vienni-Baptista, B., Fletcher, I., Lyall, C., & Pohl, C. (2022). Embracing heterogeneity: Why plural understandings strengthen interdisciplinarity and transdisciplinarity. *Science and Public Policy*, 49(6), 865–877. <https://doi.org/10.1093/scipol/scac034>

Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *International Data Privacy Law*, 7(2), 76–99. <https://doi.org/10.1093/idpl/ix005>

Waidelich, L., Lambert, M., Al-Washash, Z., Kroschwald, S., Schuster, T., & Döring, N. (2023). Using large language models for the enforcement of consumer rights in Germany. In J. Maślankowski, B. Marcinkowski, & P. Rupino Da Cunha (Eds), *Digital Transformation* (Vol. 495, pp. 1–15). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-43590-4_1

Weingart, P. (2000). 2. Interdisciplinarity: The Paradoxical Discourse. In N. Stehr & P. Weingart (Eds), *Practising Interdisciplinarity* (pp. 25–42). University of Toronto Press. <https://doi.org/10.3138/9781442678729-004>

Witt, A., Huggins, A., Governatori, G., & Buckley, J. (2024). Encoding legislation: A methodology for enhancing technical validation, legal alignment and interdisciplinarity. *Artificial Intelligence and Law*, 32(2), 293–324. <https://doi.org/10.1007/s10506-023-09350-1>

Zowghi, D., & Gervasi, V. (2002). The three Cs of requirements: Consistency, completeness, and correctness. In *Eighth International Workshop on Requirements Engineering: Foundation for Software Quality*.

Published by



in cooperation with

