SoK: Enhancing Privacy-Preserving Software Development from a Developers' Perspective

THARAKA WIJESUNDARA, RMIT, Australia

MATTHEW WARREN, RMIT, Australia

NALIN ARACHCHILAGE, RMIT, Australia

In software development, privacy preservation has become essential with the rise of privacy concerns and regulations such as GDPR and CCPA. While several tools, guidelines, methods, methodologies, and frameworks have been proposed to support developers embedding privacy into software applications, most of them are proofs-of-concept without empirical evaluations, making their practical applicability uncertain. These solutions should be evaluated for different types of scenarios (e.g., industry settings such as rapid software development environments, teams with different privacy knowledge, etc.) to determine what their limitations are in various industry settings and what changes are required to refine current solutions before putting them into industry and developing new developer-supporting approaches. For that, a thorough review of empirically evaluated current solutions will be very effective. However, the existing secondary studies that examine the available developer support provide broad overviews but do not specifically analyze empirically evaluated solutions and their limitations. Therefore, this Systematic Literature Review (SLR) aims to identify and analyze empirically validated solutions that are designed to help developers in privacy-preserving software development. The findings will provide valuable insights for researchers to improve current privacy-preserving solutions and for practitioners looking for effective and validated solutions to embed privacy into software development.

CCS Concepts: • **Security and privacy** → *Software and application security*; *Software security engineering*;

Additional Key Words and Phrases: privacy awareness, embed privacy, software application, support developers, tools, guidelines, methods, methodologies, frameworks

1 INTRODUCTION

In the recent past, various tools, guidelines, methods, methodologies, and frameworks have been proposed, providing different types of support for developers embedding privacy into software applications, including privacy requirement specification, improving privacy awareness, privacy-preserving coding, etc. [73, 78, 83, 95, 97]. Despite the availability of these solutions, developers still face challenges when embedding privacy into software applications. For example, developers struggle to comply with regulations because of a lack of actionable technical guidelines [37, 41, 91, 95, 97]. Additionally, developers often lack privacy awareness, and due to that, they face difficulties in translating privacy principles into software requirements, which leads to deprioritizing privacy in software development [17, 26, 32, 82].

Authors' Contact Information: Tharaka Wijesundara, s4063322@student.rmit.edu.au, RMIT, Melbourne, Victoria, Australia; Matthew Warren, RMIT, Melbourne, Victoria, Australia, matthew.warren2@rmit.edu.au; Nalin Arachchilage, RMIT, Melbourne, Victoria, Australia, nalin.arachchilage@rmit.edu.au.

Even though tools, guidelines, methods, methodologies, and frameworks have been introduced, they come with their own limitations, such as requiring technical expertise to use, limited evaluation with developers, and needing manual intervention of developers [17, 68, 83].

Failing to properly embed privacy into software applications will negatively impact both users and organizations [13, 32, 80]. The Facebook-Cambridge Analytica scandal in 2018 is an example of how poor privacy management can result in major data misuse [56]. It was about utilizing users' data of 50 million Facebook profiles without their consent to better target political messages [56]. It caused a drop in the users' trust in Facebook from 79% to 28% and resulted in a $750 million lawsuit imposed by the European Union (EU) [56]. This incident exemplifies the need for privacy-aware development practices to prevent data breaches, legal penalties, and reputational damage.

Recognizing these risks, governments established strict data protection rules such as General Data Protection Regulation (GDPR) [2], California Consumer Privacy Act (CCPA) [3], the Australian Privacy Act [1], and the New Zealand Privacy Act [4] to ensure consistent privacy-preserving user data handling practices across organizations. However, simply having regulations in place does not guarantee compliance since developers lack the necessary guidance and technical support to implement them effectively [16, 23].

To address these concerns, privacy-by-design (PbD) was introduced, allowing designers and developers to incorporate privacy practices into software systems from the early phase of software development and continuing it throughout the development cycle [33, 35, 95]. PbD is a core concept in GDPR and similar regulations [39], implying that privacy and data protection measures should be integrated into systems from the beginning rather than as an afterthought. However, the lack of practical, developer-friendly guidelines, methods, methodologies, frameworks, and tools to translate its principles into a practical context and assist software developers in embedding privacy into software applications makes implementation difficult [18, 20, 28, 39, 96].

These developer-facing challenges and the limitations of the existing solutions highlight the need for more effective, developer-centric privacy support mechanisms that can be seamlessly incorporated into existing software development practices. Therefore, to help developers embed privacy into software development by developing such support mechanisms, it is crucial to first have a broad understanding of the existing tools, guidelines, methods, methodologies, and frameworks. Additionally, it is important to understand the limitations of those proposed solutions in order to analyze how well those solutions addressed the developer challenges. To do that, we conduct a Systematic Literature Review (SLR) to find answers to these research questions:

RQ1: What tools, guidelines, methods, methodologies, and frameworks are available as solutions to support software developers embedding privacy in software application development?

RQ2: What are the limitations of the solutions identified in RQ1?

The remaining sections of this paper are organized as follows. Section 2 briefly presents an overview of existing research related to the field of this SLR, highlighting the gap that this SLR addresses. Section 3 outlines the approach we used to identify, select, and analyze relevant literature. Then, in Section 4, the key findings from the selected studies are presented. In Section 5, by interpreting the results, we discuss how the solutions contribute to the field, whether they could address the developer challenges, and the suggestions for future research. Finally, Section 6, 7, and 8 respectively address the threats to validity, constraints of the review, and the main insights with key takeaways.

2 RELATED WORK

This section is intended to provide an overview of the current state of secondary studies, which specifically examined the existing developer support in privacy-preserving software development in the current literature.

Trujillo et al. [96] conducted a systematic mapping study to identify the available patterns, models, methods, and tools to support privacy-aware software development. The study highlighted that the majority of primary papers discussed privacy requirements and privacy design patterns, but with limited empirical validation [96]. Chaves and Benitti later extended this study by reviewing papers published after 2018 to examine what has changed in the area [39]. They observed an increase in research that focused on embedding privacy throughout the development lifecycle [39]. However, they emphasized that these solutions have concerns about their practical applicability since they are yet to be validated in the industry settings [39].

Canedo et al. [32] studied a specific area of privacy-preserving software development, which is privacy requirements elicitation. The study identified various methodologies, techniques, and tools developed to help developers in gathering privacy requirements for software systems [32]. However, the study highlighted that these solutions lack empirical evaluation, leaving a concern about their effectiveness in practical scenarios [32].

Overall, the results of the secondary studies highlight that most of the existing solutions that support developers in different stages of privacy-preserving software development are just proofs-of-concept [32, 39, 96]. Also, most of them haven't been implemented in practical scenarios, and there is no empirical evidence [32, 39, 96]. Current SLRs have analyzed those privacy-supporting solutions in a generalized manner, considering tools, guidelines, methods, methodologies, and frameworks collectively, without specifically focusing on empirically evaluated solutions. Furthermore, they do not examine the limitations of those proposed solutions and the limitations identified through empirical evaluation.

To effectively improve existing solutions or propose new developer-supporting tools, guidelines, methods, methodologies, and frameworks, empirically evaluated results may provide valuable insights as they provide insights about practical usability, effectiveness, and the limitations faced by developers. Without having such an SLR, it may not be possible to analyze whether the proposed solutions address practical challenges in software development or are just theoretical contributions.

Therefore, this SLR takes a developer-centric approach, specifically focusing on empirically evaluated tools, guidelines, methods, methodologies, and frameworks that provide practical support for developers in privacy-preserving software development. Through this SLR, we aim to identify the problems in current solutions and provide insights to find gaps for further research.

3 METHODOLOGY

To find out how the literature has answered the research questions *RQ1*, *RQ2*, we conducted an SLR following the guidelines proposed by Kitchenham and Charters, which is a scientific and reproducible approach to conducting SLRs in software engineering [62]. The SLR spans across three stages: *planning, conducting, and reporting*. This section covers the planning and conducting stages through different subsections, and the reporting section is covered in Section 4.

Table 1. PICOC elements, their definitions, and SLR applications [51, 62, 86]. N/A - Not applicable

Element	Definition	SLR Application
Population	What is the population that the SLR is interested in?	entities involved in privacy-preserving software development
Intervention	What are the existing approaches to address the core problem?	tools, guidelines, methods, methodologies, and frameworks
Comparison	What existing approaches can be compared with the intervention?	N/A
Outcome(s)	What are the results or effects of the interventions on the population?	helping to embed privacy in software development, resulting in privacy-aware software development
Context	What is the setting or environment in which the research is conducted?	software application development

Table 2. Derivation of keywords for each PICOC element to formulate research questions and to build search strings. The "Comparison" element is ignored as it is not related to our SLR.

Element	Primary Phrase(s)	Derived Keyword(s)
Population	software developers	developer*, "software engineer", "software engineers", programmer, "software designer", coder, "software practitioner", "software specialist"
Intervention	available developer support to embed privacy	tool*, framework?, guideline?, pattern?, process, method*, technique?, model?, education, educational, knowledge, train*, intervention, awareness, practice?, support*, behaviour*, behavior*
Outcome(s)	privacy-embedded software, privacy awareness	privacy
Context	software development	designing, coding, programming, verification, deployment, integration, "code review", validating, validation, testing

3.1 Planning the SLR

For the planning stage, we established a protocol that included four components: formulating research questions, developing search strings, selecting data sources, and selecting study criteria. The importance of the protocol in this stage was to reduce the researcher bias by establishing components before conducting the SLR [62].

3.1.1 Formulating the research question. The most important part of the SLR, as well as the planning stage, is formulating the research questions [62]. The research questions derive the entire SLR methodology by involving searching, extracting, and analyzing data [62]. In other words,

- *Data Searching*: The search strings should be related to the research questions [62].
- *Data Extracting*: The extracted data should address the research questions [62].
- *Data Analyzing*: The data should be analyzed in a way that can answer the research questions [62].

Therefore, we used a framework called PICOC (Population, Intervention, Comparison, Outcome, Context), which includes five criteria, to formulate research questions [62, 86] as shown in Table 1. The "comparison" element was ignored in this SLR. Since one of the intentions of conducting this SLR was to find existing tools, guidelines, methods, methodologies, and frameworks as discussed in Section 1, there was nothing to compare with the "intervention". Further, the derived keyword(s) under each PICOC element, as shown in Table 2 were used to develop search strings.

3.1.2 Developing the Search Strings. We developed search strings to identify relevant literature using the derived keywords in Table 2 as search terms. First, we derive primary phrases for each PICOC element based on its SLR application as listed under the column "Primary Phrase(s)" in Table 2. Then, we expanded the set of keywords by deriving keywords from each primary phrase (e.g., for software developers, we derived keywords such as software

Table 3. Selected top-tier venues for each discipline

Discipline	Conferences	Journals
Computer Security and Cryptography	- IEEE Symposium on Security and Privacy (S&P) - ACM Symposium on Computer and Communications Security (CCS) - USENIX Security Symposium - Network and Distributed System Security Symposium (NDSS) - Symposium On Usable Privacy and Security (SOUPS) - Proceedings on Privacy Enhancing Technologies Symposium (PoPETS)	- IEEE Transactions on Information Forensics and Security (TIFS) - Computers & Security (Com. & Sec.) - IEEE Transactions on Dependable and Secure Computing (TDSC) - Security and Communication Networks (Sec. & Com. Net.) - ACM Transactions on Privacy and Security (TOPS)
Software Systems	- International Conference on Software Engineering (ICSE) - International Conference on Automated Software Engineering (ASE) - ACM SIGSOFT International Symposium on Foundations of Software Engineering (SIGSOFT)	- IEEE Transactions on Software Engineering (TSE) - ACM Transactions on Software Engineering and Methodology (TOSEM)
Human-Computer Interaction (HCI)	- ACM Conference on Human Factors in Computing Systems (CHI) - Proceedings of the ACM on Human-Computer Interaction (HCI) - Computer Supported Cooperative Work (CSCW)	- IEEE Transactions on Human-Machine Systems (Trans. Hum. Mach. Sys.) - ACM Transactions on Computer-Human Interaction (TOCHI) - IEEE Transactions on Mobile Computing (Trans. Mob. Com.)

engineer, programmer, developer, and so on) and considering synonyms between the derived keywords (e.g., education - knowledge). The final list of derived keywords is listed under the column "Derived Keyword(s)" in Table 2.

Further, we used **asterisk wildcard (*)** for some terms to ensure that the relevant literature covered by different forms of the same word isn't ignored in the search process. For instance, we used "tool*" to cover "tool", "tooling", "tooling", etc. We used **question marks (?)** to extend the possible keywords that are different from one character (e.g., framework? will take into account both framework and frameworks). And, to search with some exact phrases, we used **quotation marks ("")**. During this process, we make sure to remove the duplicate keywords found under each element. Finally, logical operators such as AND and OR were used to refine and advance the search query, which is as follows:

(developer* OR "software engineer" OR "software engineers" OR programmer OR "software designer" OR coder OR "software practitioner" OR "software specialist") AND (tool* OR framework? OR guideline? OR pattern? OR process OR method* OR technique? OR model? OR education OR educational OR knowledge OR train* OR intervention OR awareness OR practice? OR support* OR behaviour* OR behavior*) AND (privacy) AND (software OR designing OR coding OR programming OR develop OR development OR verification OR deployment OR integration OR "code review" OR validating OR validation OR testing) AND (software OR code OR program OR algorithm OR design OR requirement OR test OR validate OR review OR debug OR develop OR development)

3.1.3 Selecting the Data Sources. Once the search strings and the query were finalized, the next step was to identify the data sources from which relevant literature could be extracted. We considered three types of sources: digital libraries, journals, and conference proceedings. As digital libraries, we considered, IEEE Xplore[1], ACM Digital Library[2], Scopus[3], SpringerLink[4], and Google Scholar[5]. To select journals and conferences, first, we categorized the scope of the literature into four disciplines, considering the research questions. The disciplines include computer security, privacy and cryptography, software systems, artificial intelligence, and human-computer interaction (HCI). Then, we selected top-tier venues for each category, as shown in Table 3. In total, we considered 12 conferences and 10 journal publications to conduct the SLR.

3.1.4 Study Selection Criteria. Then, we defined the selection criteria to decide what studies should be included or excluded from the SLR [62] through Inclusion Criteria (IN) and Exclusion Criteria (EX) as shown in Table 4. Criteria were defined in a way that extracts the most relevant publications to address the research questions. To define the criteria, we gain insights from an SLR on privacy-enhancing technologies in software development [30].

Table 4. Inclusion and Exclusion Criteria

Inclusion Criteria
IN1: The publication should address at least one of the research questions.
IN2: The publication should be published in a peer-reviewed journal or conference.
IN3: The publication should include empirical evidence from software developers or students who have a software background. The proposed developer support should be used by the developers (e.g., there is a tool that generates privacy policy statements. The authors generated some statements and got some feedback from developers about the generated statements. However, since the tool has not been used by the developers, it will not be included in our SLR.).
IN4: The entire publication or a portion of it should be privacy-focused. (e.g., publications may address both security and privacy together).
IN5: The publication should address software developers who engage in software development. (e.g., we consider publications that address software developers who do both developing and designing).
IN6: The publication should be written in the English language.
Exclusion Criteria
EX1: The publication focuses on applications outside the software development unless it provides insights about addressing *RQs* of this SLR.
EX2: The publication is a short paper, poster, abstract, tutorial, extended abstract, demonstration, introduction, keynote, opinion paper, concept paper, editorial, work in progress, interview, monograph, secondary study, tertiary study, technical paper, book chapter, ongoing study, white paper.

3.2 Conducting the SLR

Once the protocol was finalized, the next step was conducting the SLR. We conducted the SLR through three stages: data search and selection, data extraction, and data analysis [62].

3.2.1 Study Search and Selection. The identification of relevant publications from data sources discussed in Section 3.1.3 was conducted through three different phases. The study selection was performed in September 2024-January 2025 following three main phases: automatic search on digital libraries [62], manual search on selected journals and conference proceedings [62], and bidirectional snowballing [98], as shown in Figure 1.

For phases I and II, the search process was conducted through the data sources discussed in Section 3.1.3 using a search query constructed in Section 3.1.2. During Phase III, we reviewed the references and citations of the studies identified in Phases I and II to ensure that all the relevant literature was considered for the SLR. Further, we didn't consider any time limitations for our search. Also, to manage publications during all three phases, we used Zotero Reference Management Software [6].

The constructed search query resulted in 7428 publications in phase I. Phase I was conducted in two stages. First, the titles and abstracts of all the resulting publications were reviewed while checking their suitability with inclusion and exclusion criteria. Then, 273 publications were selected for further review. Second, the review was extended towards the introduction, conclusion, and sometimes for the full text of the selected publications in the first step. After the second step, 22 eligible publications were selected for the SLR.

In phase II, the same search query was used at some venues where the search-by-keyword facility was incorporated (e.g., ACM and IEEE venues listed in Table 3 have that facility). Sometimes, we had to do a manual search through all the papers published in different years to select relevant papers (e.g., NDSS, PoPETS, etc.). During the search, we ignored the already selected papers in Phase I. After applying inclusion and exclusion criteria, we identified 02 suitable publications for the SLR.

Once we finished Phases I and II, we conducted the snowballing technique as Phase III of our search based on the finalized papers from Phases I and II. We ignored some previously identified papers during Phase III, just as we did in Phase II. After all, 08 papers were selected.

Once all three phases were conducted, altogether 32 final publications were extracted to conduct the SLR. The overall study search and selection strategy is depicted in Figure 1.
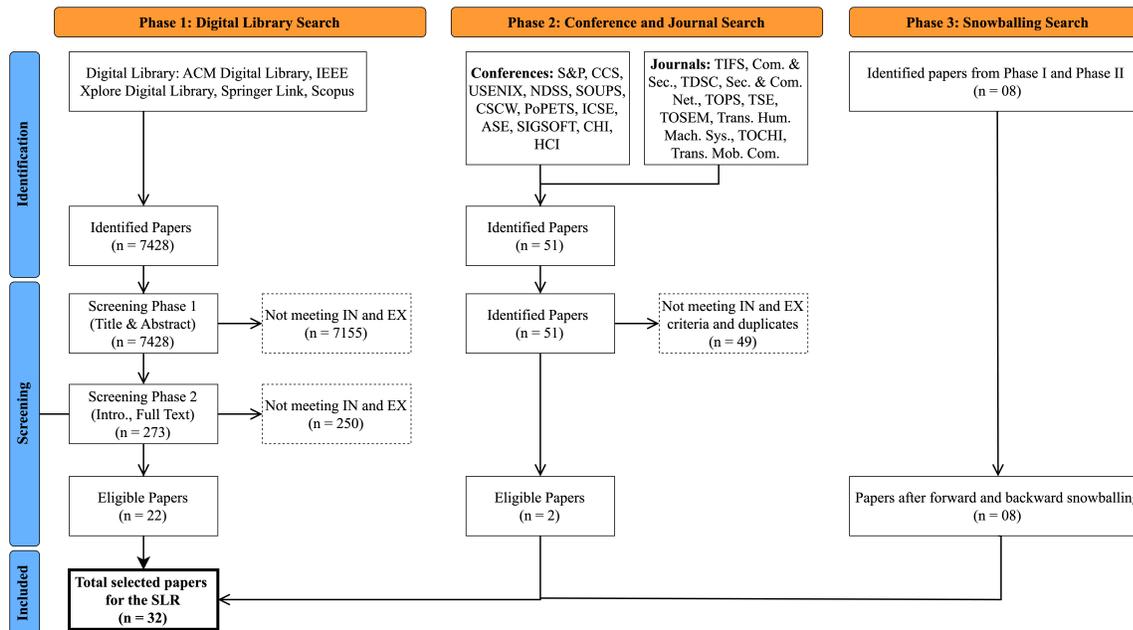
---

[6] https://www.zotero.org/

Fig. 1. The summary of the study search and selection process using the cPRISMA framework. n - number of papers, IN - Inclusion Criteria, EX - Exclusion Criteria.

3.2.2 Data Extraction and Analysis. Once the relevant papers were selected, we extracted some data from each paper and tabularized them in a spreadsheet to familiarize ourselves with the content and the domain of focus of each paper. The table includes the title, authors, publication year, venue, purpose/objective of the paper, contribution, main findings, limitations, and future work.

After the data extraction and familiarization step, we conducted the reflexive thematic analysis [38]. We started coding the selected papers without using a pre-defined coding scheme, but with the research questions in mind [38]. This allowed us to code papers based on a list of codes derived from the research questions and the authors' knowledge of privacy-preserving software development. Then, we expanded the list of codes by adding new codes based on the extracted data [38]. For instance, after going through the extracted data multiple interactions, the phrase "make use of personal data while simultaneously making it easier to analyze how that personal data is processed....." [69] was coded as "personal data", "tool", and "process" by the first author.

Once finished coding, they were grouped under different themes based on the patterns and similarities of the identified codes in order to make use of them to answer the research questions, RQ1 and RQ2. For example, the publication with the above phrase was categorized under *Manage & Secure Personal Data*. Further, we considered the paper's main contribution when assigning them to the themes. For example, Li et al. [69] support managing personal data in privacy-preserving application development while helping developers enhance their privacy education. However, since its main contribution is to provide developer support to manage personal data, we assigned it under the theme *Manage & Secure Personal Data*. Finally, we defined six different themes to answer the research questions, RQ1 and RQ2.

The co-authors validated the generated codes, generated themes' relevance to the research questions RQ1 and RQ2, as well as assigning publication under each theme. This was accomplished through multiple discussions among the

co-authors to determine whether there were any inconsistencies in the coding, theming, or publication assignment process. To accomplish this, the first author created a spreadsheet containing the relevant details of the publications, the coding and theming details, and the assignment of publications to the generated themes. Then, the spreadsheet was discussed during the discussions to check for inconsistencies. During the discussions, the co-authors' diverse perspectives were considered, and thematic development and publication assignments were carried out according to the guidelines proposed by Braun and Clark [31].

To Consider Privacy in the Early-stage of Development (Section 4.1)

Requirement specification	Threat modeling & identification	
Tools	**Tools**	**Methods**
1.PCM-tool 2.RMCM	1.LINDUNN-Go	1.ThreatPoker game
Limitations		
• Limited evaluation in agile and industrial settings [1]. • Need manual intervention of developers [1]. • Need privacy expertise to use [2]. • Difficult to use because of low user experience [1].	• Support for a limited number of threat types & usability in complex systems is less [1]. • Limited evaluation in industrial settings [1].	• Limited evaluation in industrial settings.

To Manage & Secure Personal Data (Section 4.3)

Personal data handling	Location data handling
Tools	**Frameworks**
1.PrivacyStreams	1.Platys
Limitations	
• Does not support managing personal data if complex scenarios are utilized (e.g., machine learning) [1]. • Support for selected platforms (e.g., Android) [1]. • Support for third-party libraries is limited [1].	-

To Address Vulnerabilities in Third-party Dependencies (Section 4.4)

Avoid insecure libraries	Find issues in libraries and SDKs
Tools	**Tools**
1.Up2Dep	1.DataAvalanche.io
Limitations	
• Inherits the limitations of third-party resources used to develop the tool. • Cannot be integrated with IDEs. • Limited evaluation in industrial settings.	• Limited evaluation in industrial settings.

To Consider Privacy During the Design and Development (Section 4.2)

Design & Development				Secure & Privacy-preserving Coding
Tools	**Guidelines**	**Methodologies**	**Frameworks**	**Tools**
1.Canella 2.Parrot	1.POSD 2.PbD guidelines for IoT design	1.PbE 2.CryptSDLC 3.UML-based MDD	1.CIA-level driven SDLC 2.secD4CloudMobile 3.Hails	1.FixDroid 2.PrivacyCAT
Limitations				
• Transferring to unevaluated domains may be challenging [2]. • Support is provided only for selected regulations [1, 2]. • Usability & User-Interface (UI) issues in the developed prototypes [1, 2].	• Validated the usability only for re-engineering existing systems [1]. • The number of supported programming languages is low [1]. • Need manual intervention of developers [1]. • Limited evaluation in agile and industrial settings [1]. • Difficult to follow and time-consuming [2].	• Limited evaluations about the usability in different domains [2]. • Relies on extensive documentation [1]. • Need some expertise in the domain [1]. • Support is provided only for limited privacy principles [3]. • Limited programming language support [3].	• Need developers manual intervention [1]. • Further evaluation is needed [2]. • The number of supported programming languages is low [3]. • Difficult to use because of the lack of support about how to use them [3]. • Limited support for complex systems [3].	• Limited platform support (e.g., Android) [1]. • Limited evaluation in industrial settings [2].

To Make Software Privacy Compliant and Transparent (Section 4.5)

Privacy label generation		Privacy policy generation	Privacy notice generation
Tools	**Methodologies**	**Tools**	**Tools**
1.Matcha 2.Privacy Label Wiz	1.Fine-grained localization of privacy behaviors	1.PrivacyFlashPro 2.Coconut	1.Honeysuckle
Limitations			
• Support for selected platforms (e.g., Android or iOS) [1, 2]. • Code analysis techniques are simple, error-prone, and supported for selected SDKs. Therefore, cannot guarantee the full transparency of application data practices [1, 2]. • Limited evaluation in industrial settings [1, 2].	• Evaluated only for Android platform. • Time-consuming. • Requires developer's manual intervention. • Limited evaluation in industrial settings.	• Support for selected platforms (e.g., iOS) and only for mobile applications [1]. • Code analysis is not entirely accurate, and developer's manual effort is needed [1, 2]. • Evaluated for developers with limited privacy knowledge and further evaluations are needed. [2]. • Users showed a lack of trust in automatically generated fields [2].	• Code analysis technique is simple, and developer's manual effort is needed. • Support only for selected third-party libraries.

To Improve Privacy Education and Awareness (Section 4.6)

Tools	Methods
1.Privacy Ideation Card (PIC)	1.Motivational workshop
Limitations	
• Limited evaluation in industrial software development settings. • The applicability of the gathered privacy awareness was not examined.	• Effective for developers with limited security and privacy knowledge but cannot be generalized. • Effectiveness only when the team leader is supportive.

Fig. 2. The summary of the thematic analysis. In RQ1, we discussed the available tools, guidelines, methods, methodologies, and frameworks as solutions. In RQ2, we discuss their limitations.

4 RESULTS

As shown in Figure 2, this SLR presents an in-depth analysis of existing publications about developer-supporting guidelines, tools, methods, methodologies, and frameworks for incorporating privacy into software applications. Even

Table 5. A summary of empirical evaluation contexts. Here, student refers to those who study at the university or school level, and industry refers to developers with industry experience (i.e., experience in industry-level projects).

Name	Reference	Evaluation context
PCM-tool	[81–83]	Students and industry
RMCM	[72]	Industry
LINDUNN-Go	[99]	Students and industry
ThreatPoker game	[89]	Students
PrivacyStreams	[69]	Students and industry
Platys	[75]	Students
Up2Dep	[77]	Students and industry
DataAvalanche.io	[44]	Industry
Canella	[17]	Students and industry
Parrot	[14, 15]	Students and industry
POSD	[25, 26]	Industry
PbD guidelines for IoT design	[84]	Students
PbE	[28]	Students

Name	Reference	Evaluation context
CryptSDLC	[70]	Industry
UML-based MDD	[64]	Did not clearly mention
CIA-level driven SDLC	[59]	Industry
secD4CloudMobile	[36]	Did not clearly mention
Hails	[54]	Did not clearly mention
FixDriod	[78]	Students and industry
PrivacyCAT	[73]	Industry
Matcha	[66]	Students and industry
Privacy Label Wiz	[52]	Did not clearly mention
Fine-grained localization of privacy behaviors	[58]	Students and industry
PrivacyFlashPro	[100]	Industry
Coconut	[65]	Students and industry
Honeysuckle	[68]	Industry
Privacy Ideation Card (PIC)	[95]	Students
Moticational Workshop	[97]	Industry

though different types of solutions exist, developers struggle with issues such as not having good privacy knowledge to use them, not having proper solutions to be used in the Agile method, and some accuracy issues of the solutions [28, 65, 68, 83, 100].

The results section delves into these issues by discussing the existing developer-supporting mechanisms to address these issues. As shown in Figure 2, to answer the research questions (RQ1 and RQ2), we categorized the identified developer-supporting mechanisms into six different themes (as depicted in yellow color) and sub-themes (as depicted in purple color) based on the type of support they provide for developers. For example, in Section , we discuss the available developer-supporting mechanisms under the theme "consider privacy in the early stage of development" and two sub-themes: "requirement specification" and "threat modeling and identification.". However, please note that Section 4.6 does not have any subthemes because the identified mechanisms were discussed under the main theme.

Further, as shown in Figure 5, we summarized the evaluation context of all the identified developer-supporting solutions, which we discuss further in both the results and discussion sections.

The following sections will mainly discuss the Figure 2 in detail and also Figure 5, answering RQ1, while RQ2 was answered in the discussion section (i.e., Section 5).

4.1 Consider Privacy in the Early-stage of Development

4.1.1 Requirement specification. Privacy is one of the major Non-Functional Requirements (NFRs) in Software Engineering (SE), which is often overlooked or given minimal attention during software development [81]. Addressing privacy requirements early in the development process may help avoid privacy violations that may happen later in the development process [53]. However, Agile software development (ASD) does not have a specific technique to guide developers to privacy requirement specifications in a situation where outdated privacy requirements pose challenges for agile teams [32, 82]. For example, assume an agile team is creating a healthcare app with old privacy rules, like collecting user health data without clear consent. When new laws, such as GDPR, require stricter data protection, the team finds it hard to update their app quickly.

To address this, Peixoto et al. [83] developed the Privacy Criteria Method tool (PCM-tool), which provides structured guidance for privacy requirement specification in Agile. The PCM tool has a user interface with input fields to specify each privacy requirement [83]. It consists of input fields such as a detailed description of privacy requirements, their

priority level, and associated vulnerabilities, legal compliance (e.g., GDPR), ensuring that privacy requirements are updated and aligned with current regulations [83]. Further, some developer guides are available about how to fill these fields [83]. Finally, empirical evaluations with graduate students and industry professionals indicate that the PCM tool aids developers in specifying good-quality privacy requirements [81–83].

Separately, Mai et al. [72] proposed a method called Restricted Misuse Case Modeling (RMCM) that utilizes use cases and misuse cases to model security and privacy requirements. Here, a use case refers to an interaction between a user and the system to achieve a specific goal, while a misuse case refers to a scenario in which an attacker tries to compromise the security and privacy of the system. RMCM provides structured templates and restriction rules (i.e., constraints or conditions that limit how a system, process, or model should behave to ensure security or compliance) to clearly define security threats, threat scenarios, and mitigation strategies [72]. It ultimately helps to model privacy requirements [72]. The approach was successfully applied to an industrial healthcare project, demonstrating its effectiveness [72].

4.1.2 Threat modeling & identification. Threat modeling is another essential technique for early-stage privacy risk assessment [99]. However, current threat modeling methods like LINDUNN are too complex and require a lot of expertise and time for developers to understand [99]. As a solution, Wuyts et al. [99] proposed a simplified card-based toolkit for collaborative threat identification called LINDUNN-Go. LINDUNN-Go is a teamwork that allows people with different skills (developers, privacy experts, etc.) to work together for the privacy threat analysis [99]. LINDUNN-Go consists of threat-type cards that describe potential privacy threats with examples and guidance in a way that can even be understood by non-experts [99]. Each card follows the same template and includes necessary information such as threat source, consequences, the area where the threat occurs, etc [99]. The feedback gathered from students and industry experts showed that LINDUNN-Go is easier to use and understand [99].

Further, to identify threats in the software systems and the effort needed to mitigate vulnerabilities related to those threats during Agile development, Rygge and Jøsang [89] introduced ThreatPoker, a team-based card game [89]. The game consists of two rounds [89]. First is the risk round, where team members discuss and analyze potential security and privacy threats associated with the user stories or features being developed in the current sprint [89]. Also, during the same round, they estimated the severity of security and privacy risks by playing cards [89]. Second is the solution round, where they estimate the effort required to address the threats [89]. Since threat poker is a team game where everyone, including developers, shares their ideas and helps to make security and privacy-related decisions, it helps everyone feel responsible for keeping the project safe [89].

4.2 Consider Privacy During the Design and Development Stages

Ensuring privacy throughout the software design and development process is critical to mitigating the privacy-related problems that might arise later [70, 74]. For example, if a healthcare app fails to encrypt sensitive patient data during development, hackers may later exploit this vulnerability and steal the data. In this section, we discuss the existing tools [14, 17, 73, 78], guidelines [26, 84], methodologies [28, 64, 70], and frameworks [36, 54, 59] to design and develop applications in a privacy-preserving way.

4.2.1 Tool-based design and development. IoT application development is more complicated than desktop, mobile, and web application development because of the collaborative work needed with different development backgrounds, such as software, hardware, and embedded systems [84]. Parrot [14] and Canella [17] are tools specifically designed to help developers in privacy-aware IoT application design and development.

Incorporating privacy features into IoT is challenging for developers because the existing guidelines are difficult to translate into technical implementation due to reasons such as they are written in legal language [14, 90, 94]. For example, terms like "lawful basis for processing" or "data subject rights" from GDPR [2] require legal expertise to fully understand and implement. As a solution, Parrot comes as a developer-friendly tool with an interactive design interface that translates complex legal privacy requirements into actionable technical design choices [14]. For example, PARROT simplifies GDPR concepts like "lawful basis for processing" or "data subject rights" by breaking them down into clear, configurable options and visually highlighting potential privacy risks while providing instant feedback [14]. At the same time, it supports visualizing the system's architecture and data flows [14]. Finally, it showed that PARROT can help novice IoT developers comply with privacy laws and follow privacy guidelines [15].

Separately, Canella supports privacy-aware IoT application development using visual programming tools such as Blockly and Node-RED, which support drag-and-drop-based development [17]. It provides real-time feedback on potential privacy issues in the system while highlighting areas where user data might be at risk and suggesting less risky alternatives for handling personal information [17]. Also, it includes a Privacy Law Validator to ensure that the IoT system complies with the data minimization principle as well as the privacy and data protection rules of the UK [17]. Further, it offers interactive features like warning messages and tooltips to reduce the cognitive load on developers and enhance their understanding of privacy practices. Overall, it helps to protect user data and keep IoT applications up to date with evolving legal requirements [17, 27, 55, 67].

4.2.2 Guideline-based design and development. Privacy guidelines are a set of clear instructions and best practices for developers to follow when embedding privacy into software. It is challenging for developers to translate these guidelines into technical implementations when they do not have the required expertise in privacy [26]. For example, "Minimize the collection of personal data to only what is necessary for the intended purpose" is a guideline based on the data minimization principle of Privacy by Design [34] as well as GDPR [2]. Privacy-Oriented Software Development (POSD) [26] helps developers translate these privacy guidelines into technical implementations by providing developers with a structured, practical, and actionable technical implementation via a Privacy Knowledge Base (PKB) [25]. The PKB acts as a bridge between high-level privacy principles (e.g., Privacy by Design) and practical, actionable, technical implementations [25]. POSD's guidelines proposed a static code analyzer to identify vulnerabilities in the code [26]. PKB helps map these identified vulnerabilities with privacy patterns, which are reusable solutions for common privacy problems (e.g., logging user activities that contain sensitive information for debugging purposes), to mitigate the vulnerabilities [26]. These patterns include detailed descriptions, diagrams, and implementation steps, making it easier for developers to apply them in code [26].

Another set of guidelines was proposed by Perera et al. [84] to assist developers in following the PbD principles when designing IoT applications. This was proposed to support embedding privacy specifically into heterogeneous and distributed IoT systems [84]. This offers systematic guidance through a set of privacy guidelines and a method for applying them during the IoT application design process [84]. It consists of 21 privacy guidelines derived from Hoepman's privacy design strategies, which cover various aspects of privacy, such as data minimization, data anonymization, encryption, etc. [57, 85].

4.2.3 Methodology-based design and development. As a solution to the lack of guidance provided by PbD to implement them in practice, Barbosa et al. [28] proposed a methodology called Privacy-by-Evidence (PbE) [28]. The methodology consists of a clear, step-by-step workflow that can be run in parallel to the regular development cycle [28]. The workflow consists of activities such as identifying the application context and data formats, checking compliance with

regulations, assessing risks, applying privacy techniques (e.g., anonymization, generalization, noise addition), and evaluating potential attacks [28]. Here, evidence refers to documented proofs that explain how privacy risks have been identified, mitigated, and addressed throughout the software development process. PbE emphasizes documenting this evidence using Goal Structuring Notation (GSN) [28].

Even though PbE focuses on systematically incorporating privacy into software systems [28], it does not address the technical complexities of cryptographic security, which requires expertise in the cryptographic engineering domain [70]. For example, k-anonymity is a privacy-preserving technique for anonymization that requires expertise in cryptographic engineering to properly understand and implement its underlying mechanisms. Cryptographic protection is a critical component in privacy-preserving software development, which introduces additional vulnerabilities if implemented incorrectly [70]. For example, improperly applying k-anonymity will lead to re-identification attacks. As a solution, Loruenser et al. [70] proposed CryptSDLC, a methodological approach to ensure secure-by-design software systems by providing software developers with access to advanced cryptographic solutions and minimizing potential errors when integrating cryptographic functionality into software. It was achieved by offering cryptographic solutions in a simplified form by abstracting their complex cryptographic concepts, which eventually reduces the need for deep cryptographic expertise to use them in software systems and minimize integration errors [70]. CryptSDLC facilitates communication and collaboration among different stakeholders, such as cryptographers, security experts, and developers, and bridges the communication gaps that often hinder the effective integration of cryptography into software development [70]. This was achieved by providing communication tools such as Service Level Agreements (SLAs) and cryptographic design patterns to share security requirements, capabilities, and constraints with stakeholders.

Separately, Krstic et al. [64] proposed a Model-Driven Development (MDD) methodology that provides comprehensive support for developers in implementing privacy policies, particularly focusing on purpose limitation and consent as required by regulations like GDPR [64]. Developers work with three models: data model, security model, and privacy model [64]. These models are presented using Unified Modeling Language (UML) diagrams [64]. First, software engineers use the data model to describe how the system's data is organized [64]. Next, security engineers use the security model to specify who can access or change the data [64]. Finally, privacy engineers use the privacy model to ensure that personal data is only used for specific purposes and only with the user's consent [64]. Software engineers, security engineers, and privacy engineers collaboratively work to provide these models [64]. Then, these UML models are translated into actual system implementation codes in C# (ASP.NET) and Python (Flask) [64]. It reduces the manual coding effort and minimizes human errors [64].

4.2.4 Framework-based design and development. Security-by-design is a proactive approach that considers security during the early phase of software development [22, 59]. Secure software development lifecycle (Secure-SDLC) refers to the development process that incorporates security-by-design [22, 59]. Kang et al. [59] proposed a CIA-level driven secure SDLC framework allowing organizations to tailor security practices according to the level of security they require and offer a structured and evidence-based approach to integrating security into the software development lifecycle [59]. It maps 66 security activities across 10 phases of the SDLC, ensuring security is embedded at every stage, from the beginning to the end [59]. For example, it consists of the "Security Requirements Elicitation" activity with the required guidance to implement it [59]. Developers have been provided with detailed security activities, evidence templates, standards that the activity aims to achieve (e.g., ISO/IEC 27001), and tools to document and prove compliance with security requirements [59]. The mapping helps developers to get a clear picture of the activities that must be

completed at each phase to meet security standards. Further, it allows developers to compare their security practices with competitors and identify areas for improvement [59].

Similarly, to specifically address security issues of cloud-based mobile applications through security-by-design, Chimuco et al. proposed a framework called secD4CloudMobile, which includes developer support covering different areas such as security requirement elicitation, attack modeling, and security testing [36]. These sets of solutions facilitate the incorporation of security mechanisms throughout the software development lifecycle [36]. It emphasizes incorporating security-by-design practices that indirectly protect sensitive user data and maintain privacy [36]. For example, helping developers to attack modeling includes helping them to identify potential attacks such as unauthorized access to personal data, which indirectly helps focus on privacy [36]. Overall, the framework provides a simple, text-based interface where developers need to answer a series of questions about their application. Based on the responses provided, the framework generates detailed reports outlining security requirements, best practices (e.g., guidelines for secure coding, authentication, encryption, etc.), attack models (e.g., details about potential attacks), and test specifications (e.g., security testing approaches such as SAST, DAST, etc.).

Separately, Giffin et al. [54] proposed a web framework called Hails, designed to enhance security and privacy in web platforms that rely on third-party applications. In other words, it is designed for developers who are building web platforms like Facebook that allow third-party applications to extend their functionality [54]. One of the main challenges with third-party applications is their unrestricted access to user data: once access to the user data has been granted by the platform, there are no enforced restrictions on how they handle that data, leading to potential misuse or leaks [54]. For example, on a platform like Facebook, a third-party app might gain access to a user's friend list and then share that data with advertisers without the user's knowledge. Therefore, these platforms need to ensure that third-party applications cannot misuse or improperly access user data [54]. Hails mitigates this risk by introducing a new architecture called MPVC (model-policy-view-controller), which is an extended version of MVC by incorporating security policies at the framework level [54]. Hails allows platform developers to define security and privacy policies in a way that is separate from the application logic [54]. For instance, in GitStar, a code-hosting platform built with Hails [54], a third-party Code Viewer app can access a user's private repository. However, Hails makes sure that the app only shows the code to people who are specifically allowed to see it, such as collaborators. In that way, unlike conventional systems where security and privacy policies are deeply connected with application logic, Hails allows platform developers to define and enforce privacy policies separately within the platform, simplifying data access control and policy management [54].

4.2.5   Secure & privacy-preserving coding. In privacy-preserving application development, detecting code defects is crucial since even minor coding defects may lead to unauthorized access, data leaks, and privacy violations. For example, SQL injection occurs when attackers use vulnerable user input to manipulate database queries, allowing unauthorized access to sensitive data. Android Studio comes with some tools, such as Android Lint, to support developers in preventing program applications with privacy vulnerabilities [78]. However, development issues such as requesting more permission than needed [49, 50], incorrect use of cryptographic APIs [43], store sensitive information in non-private locations [47], prove that the available developer support is not sufficient to make secure and privacy-preserving applications [78]. Developers' lack of knowledge and time constraints may be a reason for that, since they also lead to privacy vulnerabilities in applications [12, 43, 47, 49, 50, 61, 78]. For instance, due to the limited time available to develop the applications, developers tend to copy and paste some code to solve security and privacy-related problems [61],

which may raise security and privacy issues. A common example is reusing code snippets from Stack Overflow [93] to implement password hashing that does not have a secure hashing algorithm like Bcrypt.

Considering the above challenges, to support developers in writing secure and privacy-preserving code, Nguyen et al. [78] proposed an Android plugin called FixDroid by improving Lint's limitations. The Android Lint tool has a feature to highlight insecure code, but with a drawback that it uses the same highlighting for all types of warnings [78]. As a solution, to get developers' attention, FixDroid implemented its user interface based on insights from usable security and privacy research as well as the highlighting feature [78]. Additionally, Lint provides only textual tooltips, whereas FixDroid offers code snippets with quick fixes and detailed explanations, enabling developers to easily turn insecure code into secure code [78]. Further, Lint's lightweight code analysis technique cannot detect complex code issues [78]. Therefore, FixDroid has been implemented with more complicated code analysis techniques [78]. Further, if a developer copies a code snippet from somewhere, FixDroid detects its vulnerabilities using an online database of insecure code snippets [78].

Separately, to detect internal code-level privacy vulnerabilities such as unauthorized logging (i.e., recording) of sensitive user data of large-scale software systems (e.g., WhatsApp), an automatic tool called PrivacyCAT has been introduced [73]. PrivacyCAT is used to detect privacy vulnerabilities before they reach the production level [73]. This is because accessing production traffic and fixing privacy issues may raise further privacy issues because it contains actual user information [73]. PrivacyCAT uses both dynamic and static taint analysis to track the flow of data and detect any data leaks [73]. The static analysis is performed on the source code, while dynamic analysis is performed in isolated, sandboxed environments with artificially generated sensitive data. It generates synthesized but realistic data to track how sensitive data flows through the system, ensuring that vulnerabilities are detected during development rather than after deployment [73]. By tracking the flow, it monitors how the system processes data at data sinks (e.g., API end-points) and exchanges through APIs to identify potential leakages [73]. Two years of results of industrial deployment of PrivacyCAT in WhatsApp reports that it effectively finds privacy vulnerabilities and helps prevent them at early stages of development [73].

4.3   Manage & Secure Personal Data

According to privacy laws, personal data can be referred to as Personally Identifiable Information (PII), which can be used to identify a specific person [69]. Developers access and use personal data to create functional, rich, and meaningful applications [69]. Mobile platforms like Android and iOS implement permission-based access control to regulate personal data usage, requiring applications to obtain user consent before accessing data [69]. However, from the users' perspective, the granularity and the purpose of data access are important for users to trust the application behavior [69]. This poses a challenge: "Will the developer be able to provide such granularity?". Asking developers to provide such fine-grained information through a privacy policy requires extra effort, and developers might not have the benefits and incentives (e.g., money) to put that effort into it [69]. Further, the availability of different APIs and data structures makes it difficult for developers to access and use personal data [69], as they may have to follow different protocols (e.g., HTTP, WebSockets), security requirements (e.g., Jason Web Tokens, API keys), data structures (e.g., JSON, XML), etc., to manage the data effectively.

As a solution, an Android library that comes as a functional programming model, PrivacyStreams, has been introduced to access and process different types of personal data available on smartphones, providing such granularity to users [69]. Given that PrivacyStreams is provided as a library, developers can import it and make use of its APIs in the code to request personal data, process it, and take actions on it by centralizing data-handling steps at a single location

[69]. PrivacyStreams includes a static code analyzer that extracts application data processing steps to automatically generate privacy descriptions, as well as to analyze the granularity of data processing [69]. For example, PrivacyStreams helps to expose the granularity of data processing by showing (i.e., generating privacy descriptions) whether an app is using fine-grained data (e.g., exact GPS coordinates) or coarse-grained data (e.g., city-level location). The generated description may be as follows: "This app requests LOCATION permission to get the city-level location.".

Most of the applications that provide location-based services represent the location using latitude and longitude [75]. However, users can be provided an improved experience if users' activities and social context can be incorporated with physical location [75]. For example, consider a coffee shop as the place and consider factors like the type of coffee shop, the atmosphere, the people that the user interacts with, and the activities the user is doing. However, implementing such personalized experiences for users without giving them unnecessary burdens is a challenge, as it is a challenge for developers to determine each user's preferences [75].

This motivates Murukannaiah and Singh to propose a framework called Platys to develop place-aware applications [75]. Here, "place" refers to the combination of factors discussed above, other than the location [75]. Platys extends location-based applications beyond location coordinates by incorporating user-specific contextual information [75]. The core of Platys is its middleware, which consists of a machine learning technique called active learning to learn a user-specific model for places [75]. The middleware exposes the learned model to applications at runtime, allowing developers to use it to develop place-aware applications without making efforts to connect with low-level sensors (e.g., GPS sensors, accelerometer, gyroscope, etc.) [75]. Also, since Platys runs locally on users' devices, it ensures the user's privacy [75]. A developer study proved that the Plays framework reduces the developer's effort to build place-aware applications [75].

4.4 Address Vulnerabilities in Third-party Dependencies

Integration of third-party libraries and SDKs, in other words, dependencies, makes the development process easier by providing reusable functionalities and accelerating the development process [27, 44, 48, 77]. However, integrating these dependencies introduces privacy and security risks, particularly if the libraries are outdated or improperly configured [77]. As a result, previous studies have examined the security and privacy risks associated with third-party libraries and provided them with solutions, which we discuss in this section [44, 77].

When an application relies on dependencies, they should be updated on time regularly to overcome the problems of outdated dependencies [77]. For example, Heartbleed is a vulnerability of an outdated version of the OpenSSL cryptography library, allowing attackers to steal sensitive data from servers [9]. Updating it may ensure protection against known threats. However, developers often fail to update them due to concerns such as extra effort, fear of incompatibilities, and lack of awareness of updates [24, 40, 77]. To address this, Nguyen et al. [77] proposed a tool called Up2Dep that analyzes the libraries used in the code base and helps developers avoid insecure library versions by providing them with feedback about outdated libraries [77]. Further, it warns developers about vulnerabilities and cryptographic API misuse of libraries, allowing them to make informed decisions before integrating a library [77].

Beyond being up-to-date with third-party dependencies, privacy implications in third-party libraries and SDKs should also be considered, specifically when developing children's applications [44]. From the developers' perspective, they face difficulties such as identifying relevant privacy-compliant third-party libraries and SDKs from the extensive range [44, 46, 63], configuring and managing SDK settings correctly [44, 45], and finding privacy-friendly alternative libraries and SDKs [44, 46]. Further, understanding the data-handling process of third-party libraries and SDKs is also difficult due to their lower transparency about the data-handling processes [29, 44, 76].

Considering these challenges, Ekambaranathan et al. [44] proposed a web-based tool called DataAvalanche.io to support developers navigating legal issues and privacy implications associated with third-party libraries and SDKs. Developers can use this tool to get more privacy information about commonly used libraries and SDKs, including data trackers associated with them [44]. Additionally, the proposed tool provides clearer instructions to configure child-specific SDKs and provides a list of privacy-preserving alternative SDKs with documentation of their data-handling practices [44]. A key feature of the tool is its ability to visualize with which countries the SDKs share data, allowing developers to assess whether an SDK complies with regional privacy regulations like GDPR [44].

4.5   Make Software Privacy Compliant and Transparent Using Privacy Statements

This section discusses the developer support available in the existing literature to create privacy statements. We discuss this under three different groups based on the support they provide for developers: privacy label generation [52, 66], privacy policy generation [65, 100], and privacy notice generation [68].

4.5.1   Privacy label generation. Privacy labels are short notices about how applications handle user data, which are often displayed when an application is listed on the app store [60]. However, when implementing them in software development practices, developers face some challenges due to a lack of knowledge about privacy concepts (e.g., data minimization), misunderstanding about technical terms used in privacy labels (e.g., precise, coarse, etc.), and difficulties in privacy label generation when using third-party libraries with black-box data handling practices [52, 58, 66]. Further, differences in privacy labels and the actual application behaviors will lead to losing users' trust in applications [58] and fines for developers for privacy violations [2].

To assist developers in creating accurate privacy labels for Google Play, an Android Studio plugin called Matcha was introduced [66], while Privacy Label Wiz (PLW) was developed for Apple iOS applications [52]. Both use static code analysis techniques to analyze the data-handling practices of the application [52, 66]. Matcha automatically detects third-party SDKs used in the codebase and helps fill parts of privacy labels based on the information provided by the SDKs' privacy documentation [66].

On the other hand, PLW utilizes static code analysis techniques to scan the application codebase based on iOS permissions used in the Swift code [52]. Since some permissions don't align with Apple's data types, it informs developers about potential misalignments regarding data types [52]. For example, let's say an iOS app uses the "Camera" permission to allow users to take photos. In this case, the app might not explicitly collect or store photos. Therefore, the "Camera" permission could be mapped to Apple's "Photos or Videos" data type. Additionally, PLW prompts developers with questions asking developers to confirm or refine the labels, ensuring the final set of privacy labels is an accurate reflection of the application behavior [52].

Matcha has another facility to automatically generate an XML file that consists of all data collection and sharing, and gives the ability for developers to review it and add manual annotations if required [66]. Then, Matcha will generate a CSV file consisting of labels that later can be uploaded into the Google Play developer console [66].

Apart from those plugins (i.e., tools), Jain et al. [58] proposed a novel methodology called "fine-grained localization of privacy behaviors" to accurately identify privacy behaviors (i.e., how an application uses personal information and why) in source code and predict privacy labels. It helps developers create high-quality privacy labels by reducing the time and mental effort required by automating the privacy-label generation process [58]. At the same time, the study demonstrates high accuracy in predicting privacy labels [58]. Since this approach uses individual statements in the

source code for the analysis process by going beyond the class or method level, it avoids the noise and ambiguity that come with analyzing larger code blocks [58]. It eventually improves the accuracy of the generated labels [58].

4.5.2 Privacy policy generation. The privacy policy is a way to convey the application's privacy practices with details such as how the user's and device's data is collected, used, shared, and deleted [5, 6]. However, developers face a challenge in accurately representing how data is collected, used, shared, and deleted [19, 100] because of having an inaccurate understanding of the application behavior [65]. It happens especially when third-party APIs or libraries used in the application have a lack of documentation about their data-handling practices [65]. This issue causes developers to produce inaccurate privacy policy documents. Further, in some situations, multiple data sources are controlled by one permission [69]. For instance, in a situation where developers need the user's name, once the user grants permission, it lets developers access even emails other than the name. These barriers result in producing privacy policies that may not convey the actual data-handling practices to users and may not fully protect users' data or comply with regulations.

In our SLR, we have identified two developer-supporting tools, PrivacyFlash Pro [100] and Coconut [65], that have been proposed to support developers generating privacy policies.

PrivacyFlash Pro [100] has been implemented for iOS applications coded in Swift and integrated Swift and also for libraries coded in Objective-C [100]. PrivacyFlash Pro combines static code analysis techniques and questionnaire-based wizards to generate privacy policies [100]. Fully relying on questionnaire-based policy generators may result in inaccurate policies because the generation is fully dependent on the developers' answers, which may not align with actual data practices of the application [100]. PrivacyFlash Pro is a macOS desktop application that consists of a code analysis logic written in Python and a user interface developed in JavaScript [100].

On the other hand, Coconut uses an annotation-based technique, and developers need to annotate personal data practices used within the application while coding in a predefined format [65]. This process encourages developers to explicitly document how and why personal data is collected, used, and shared [65]. Coconut provides real-time feedback on potential privacy issues, such as violations of privacy principles or unnecessary data collection [65]. Additionally, it provides quick fixes to ease the development process, such as generating annotation skeletons with automatically filled fields (e.g., automatically filling the data type) and recommending more privacy-friendly alternatives (e.g., using coarse-grained location data instead of fine-grained when high precision is not required) [65]. These features help developers adhere to recommended privacy practices while reducing the cognitive workload of manually identifying and addressing privacy concerns [65]. Coconut combines all the annotations into a summary panel, which provides a comprehensive overview of the app's personal data practices [65]. This summary can be directly translated into a privacy policy.

4.5.3 Privacy notice generation. In-app privacy notices provide real-time, brief, and contextual messages about users' data handling when they interact with the applications. However, developers face some challenges when they create privacy notices. One of the main reasons is the lack of awareness among developers about the privacy concept [68] and the design choices and best practices when crafting privacy notices [79], which may lead to poor communication and failure to inform users about how their data is handled in the application. When application development relies on third-party libraries, the design and implementation of in-app privacy notices may further complicate developers because developers are not always aware of third-party applications' data-handling practices when they are not documented [27, 65]. Additionally, developers prioritize security over privacy [55], focusing more on preventing data breaches rather than ensuring users' data-handling transparency.

We found a paper that highlights these challenges and proposes an Android Studio plugin called Honeysuckle [68], which is a programming tool proposed by Tianshi Li et al. to help developers generate in-app privacy notices. Honeysuckle utilizes an annotation-based approach to generate privacy notices where developers need to annotate each data source (where data is collected) and data sink (where data is sent or stored), describing associated data practices [68]. For example, if an app collects the user's location, the developer may annotate the code where the location data is collected and specify the purpose of the data collection (e.g., "for providing nearby restaurant recommendations"). Also, developers can use a configuration file at programming time to fine-tune privacy notices [68]. This configuration file allows developers to customize aspects such as the context information that should be displayed in the notices and the behavior of the notices (e.g., how frequently they are shown or under what conditions) [68]. For instance, a developer may need to configure the system to show a privacy notice only when the app accesses the user's location in the background. Honeysuckle helps developers to generate privacy notices that are more contextualized and user-centered without needing to manually implement each notice from scratch [68].

4.6 Improve Privacy Education and Awareness

Privacy-by-design (PbD) emphasizes integrating privacy into software from the beginning rather than taking it as an afterthought [34]. However, developers face challenges because of its broad interpretation [91, 95]. For example, PbD suggests minimizing data collection broadly, but it does not specify how to achieve it (i.e., there is no proper technical implementation) in software development. Similarly, regulations like GDPR and CCPA lack clear guidance on technical implementation [37, 41, 91, 95, 97]. For instance, GDPR's storage limitation principle requires that personal data should not be kept longer than necessary. However, it does not specify exact retention periods or how developers should enforce deletion. Because of this interpretation and implementation issues of PbD and regulations, enhancing privacy education among developers may be beneficial.

Privacy Ideation Cards (PIC), introduced by Luger et al. [71], help software developers understand how PbD principles and privacy regulations apply to software development. Tang et al. [95] utilized PICs and evaluated practically how PICs can be useful for software engineering students as an educational tool to help them learn privacy. PIC contains a deck of cards with details about PbD and privacy regulation principles (e.g., data minimization, explicit consent, etc., with their detailed explanation). In an ideation session, students draw PIC cards and engage in discussions focusing on how the PbD and regulatory principles on the cards apply to their real-world software projects [95]. The evaluation results revealed that students engaged in deeper discussions about specific aspects of the project from the privacy perspective [95].

Another approach, a workshop-based intervention proposed by Weir et al., addresses the challenge that developers face when effectively communicating security and privacy concerns to product managers who have limited awareness of security and privacy risks [97]. The product manager is responsible for making decisions related to time and resource allocation to address security and privacy concerns [97]. However, the responsibility of implementing security and privacy into software is on the development team [14], implying that it is important to help them communicate when making security and privacy decisions [97]. This intervention includes structured workshops to improve the developers' awareness about identifying and prioritizing security and privacy issues and then framing these issues in terms of business benefits to communicate with product managers [97]. The evaluation results revealed that this workshop motivated developers to engage with product managers in security and privacy-related decision-making [97].

5   DISCUSSION

In the discussion section, we discuss the limitations of the identified solutions (RQ2) that we discussed in the results section (i.e., Section 4) and suggest potential future improvements and future research avenues under the same themes as we reported the results.

5.1   Consider Privacy in the Early-stage of Development

In Section 4.1, we discussed different approaches that support developers integrating privacy in the early stage of software development. The Privacy Criteria Method (PCM-tool) [81–83], Restricted Misuse Case Modeling (RMCM) [72], and collaborative threat modeling and identification techniques such as LINDUNN-Go [99] and ThreatPoker [89] offer structured methods to embed privacy into software requirements, design, and development. However, how far these approaches have addressed the challenges that developers face in early-stage privacy integration should be discussed.

5.1.1   Requirement specification. PCM-tool was developed to be used in an agile environment to provide developers with structured guidance for privacy requirement specification [81–83]. However, evaluation carried out with industry practitioners revealed that the suitability of PCM, especially in agile, is questionable [81]. One potential improvement is integrating the PCM tool with agile project management tools like Jira [21], where the requirements are specified throughout the project cycle. Integrating PCM with Jira-like software may increase the usability of PCM (i.e., increase the user experience (UX)) since managing a separate tool only for privacy requirements may be an extra overhead for the developers. A key limitation of the PCM tool is that it requires the developer's manual intervention to specify the legal basis of privacy requirements (e.g., GDPR) [81–83]. However, since these specifications depend on the knowledge of the developer who does this activity, it may lead to errors (e.g., developers with limited privacy knowledge may provide an incorrect legal basis). Therefore, implementing an intelligent system to cross-check the correctness of the specified legal basis and warn the team if any misalignment is encountered may significantly enhance the accuracy of the privacy requirements' regulatory compliance.

Similarly, we discussed Restricted Misuse Case Modeling (RMCM), which is a structured method for defining privacy and security threats through use case and misuse case diagrams [72]. However, it is required to have privacy expertise to use it [72], hindering its practical applicability among developers who have limited privacy expertise. Since not all organizations have that expertise, integrating RMCM with development environments (e.g., IDE) and then guiding developers through automated tools to identify misuse cases may be a potential improvement.

5.1.2   Threat modeling & identification. Beyond privacy requirement specification, threat modeling and identification is a crucial technique for privacy risk assessment, which is done in the early stage of software development [99]. LINDUNN-Go was developed as a simplified collaborative card-based toolkit for threat identification [99]. Both LINDUNN and LINDUNN-Go cover 7 threat categories. However, the LINDDUN-GO toolkit focuses on the most common or critical threats within each category, while LINDUNN focuses on every possible threat scenario. This implies that the effective applicability of the toolkit in industry practice is still questionable. For example, some complex applications (e.g., healthcare applications) may contain specific privacy threats that are not covered by the most common or critical threats. Therefore, more evaluations are required to validate the effectiveness in industrial applications [99] since they may be complex and need more threat scenarios than LINDUNN-Go covers.

Even though ThreatPoker did not explicitly discuss its limitations, we propose some future experiments based on the findings of its evaluation. The evaluation was conducted with students. However, because the knowledge levels (i.e., privacy and security knowledge) of industry practitioners and students can differ, applicability in the industry cannot be guaranteed based on the current results. Subjectivity may also apply when identifying threats and takes time to make a final decision as a team. Because industry projects may have specific timelines, ThreatPoker should be tested in the industry to determine how subjectivity affects results. As a result, ThreatPoker also requires further empirical evaluations to determine its effectiveness in industry settings [89].

Further, since both LINDUNN-Go and ThreatPoker are card-based games, the usability may be lower when the team consists of a large number of members. Further, playing card games may not be effective in time-constrained software development environments. Therefore, converting these games to digital versions (e.g., software applications) may enhance their usability as well as their scalability.

5.2 Consider Privacy During the Design and Development Stages

We discussed the available developer support in the current literature to embed privacy during the design or development processes in five different categories, as shown in Figure 4. Below, we discuss the implications of the results and potential future directions based on their limitations.

5.2.1 Tool-based design and development. Developers in small teams, small and medium-sized enterprises (SMEs), or individual developers engage in both design and development activities regardless of their position in the team [88]. For example, software developers may be involved in IoT design activities. The identified solutions, Parrot [14] and Canella [17], targeted such teams. Parrot specifically highlighted that it can be used by software developers, senior software developers, or even architects who are involved in the IoT application design process [14].

However, given that Parrot's primary regulatory framework is GDPR [14], its applicability in applications that operate across multiple jurisdictions is questionable. For example, developers may want to build an IoT application that processes users' information who are in Europe and Australia. It necessitates adhering to both GDPR [2] and the Australian Privacy Act [1]. Therefore, future implementations should incorporate other regulatory compliances such as CCPA [3], Australian Privacy Act [1], and the New Zealand Privacy Act [4] to ensure that developers can adhere to regulatory requirements based on their target markets. Additionally, it was developed for healthcare applications where privacy is critical [14]. It may be challenging to transfer it to other domains [14]. Therefore, extended evaluations are needed to evaluate the usability in other domains.

Further, some usability concerns of Parrot have been identified in the evaluation process, highlighting that developers need more responsive feedback mechanisms [14]. For example, if a developer selects a certain data type to be used in the application (e.g., location data), the tool might not immediately show how this selection affects privacy compliance or data security (e.g., the feedback might be delayed or not detailed enough). Improving the user interface (UI) and user experience (UX) by offering quick context-aware suggestions (i.e., tailored to the specific context of the IoT application) and providing pre-configured templates for common IoT use cases may significantly enhance its effectiveness.

Similarly, Canella's limitations impact its usability [17]. Developers raised concerns about its unintuitive interface, particularly due to the small size and colorful blocks [17]. For instance, some developers missed Canella's warning due to their small size and colorful blocks [17]. While visual programming environments may be beneficial for beginner-level developers, experienced developers may prefer traditional code-based approaches. Therefore, incorporating both

block-based and text-based development into Canella will allow developers to choose the best option and may increase Canella's usability.

Further, similar to Parrot, Canella's limited regulatory compliance support may affect when applying it in the broader IoT development domains [17]. Therefore, increasing the support for more privacy regulatory frameworks may make Canella more robust for real-world deployment.

Finally, since both Parrot and Canella educate developers about best privacy-ensured design and development practices [14, 17], expanding them as educational tools by adding interactive tutorials and exercises may educate developers about practical privacy implementations, as we discussed in Section 5.6. It will help developers to stay in one tool without relying on separate educational tools.

5.2.2 Guideline-based design and development. In Section 4.2.2, we discussed sets of guidelines that are privacy-oriented software development (POSD)[26] and PbD-based IoT designing guidelines [84] to help developers integrate privacy into software development. However, their practical effectiveness is questionable because of reasons such as requiring manual intervention, taking too much time to understand, offering limited programming-language support, and being evaluated only for re-engineering a system [26, 84].

Since POSD's support has only been evaluated in re-engineering contexts [26], it suggests that POSD primarily functions as a corrective approach rather than a proactive approach. Developing a system considering privacy as an afterthought may not be effective in both time and cost terms. Therefore, proposing POSD in a way that it is suitable for the forward mode will make this set of guidelines more usable [26].

Fortify SCA is the proposed code analyzer in POSD [26]. Even though it can analyze codes of more than twenty languages, it is in vain since the PKB of POSD supports only exporting Java-based structures for privacy patterns [26], reducing its practical usability among developers who work in languages such as Python, JavaScript, or C#. Therefore, it can be recommended to expand PKB to support more languages to get the benefit of the code analyzer [26] since the tool may remain inaccessible to many development teams without broader language support.

Additionally, in POSD, developers need to manually identify system vulnerabilities using code analysis and then input the detected vulnerabilities to PKB [26]. That process may introduce human errors. Therefore, replacing the manual workload with an automated process may reduce potential human errors. For example, machine learning-based vulnerability detection. However, they might not be 100% accurate. Therefore, implementing a feedback loop to get the developer's feedback for the predictions and then retraining machine learning models based on the feedback may further enhance the prediction accuracy. Further, integrating it with integrated development environments (IDEs) may help developers follow the POSD guidelines while coding without making extra effort. Another issue is POSD's incompatibility with Agile methodologies [26]. Since Agile methodologies prioritize continuous integration, static privacy guidelines like POSD may not be suitable. Without evaluations in real-world Agile environments, the suitability of POSD in Agile environments is unclear. Therefore, the applicability of guidelines in an agile environment should also be evaluated.

On the other hand, the usability of the proposed PbD-based IoT guidelines is less because the way the authors conveyed the guidelines (i.e., printed sheets) for developers was difficult to follow and time-consuming [84]. As a solution, Privacy Ideation Cards (PICs) may be effective for developers to quickly familiarize themselves with the guidelines since PICs have shown their effectiveness in practice [71, 84, 95], as we discussed in section 4.6. Similarly, providing developers with an abstract version of guidelines, which are privacy tactics and patterns, may improve their usability [84], especially among developers who have less privacy knowledge and who face difficulties in following

guidelines. Further, even though utilizing physical cards (e.g., PIC [71, 95]) has shown effectiveness, it may be more effective if a user-friendly interface can be developed using human-computer interaction techniques for developers to follow during their design process [84].

5.2.3 Methodology-based design and development.

The identified methodology-based approaches, Privacy by Evidence (PbE) [28], CryptSDLC [70], and Model-Driven Development (MDD) for GDPR compliance [64], were discussed in detail in Section 4.2.3. However, their effectiveness in practice should be further discussed because of their limitations, such as complexity, reliance on manual processes, and limited real-world validation [28, 64, 70].

As we discussed, PbE relies on extensive documentation, but it is time-consuming [28] and may be impractical for rapid development processes where speed is prioritized. For instance, in agile or DevOps workflows, developers often need to deliver features quickly, and the manual effort required to write detailed privacy-related artifacts (e.g., risk assessments, compliance proofs, and attack simulation reports) may slow down the development process. This makes PbE potentially impractical for teams that are operating under tight deadlines or with limited resources.

Even though PbE was proposed to address the lack of clear implementation guidance for PbD [28], it may not provide a simplified privacy integration for developers with limited privacy expertise. For example, developers may be unfamiliar with privacy models like k-anonymity or differential privacy. As a result, they may struggle to apply these concepts correctly without having proper training or support. Therefore, after a careful analysis of PbE, we suggest developing automated or semi-automated tools to assist developers during the PbE process to reduce the manual intervention of developers and streamline the PbE process. For example, developing a tool that can scan the codebase for sensitive data handling and flag areas where privacy techniques such as anonymization or encryption are needed. It may help to increase awareness about privacy concepts among developers. Further, integrating them with IDEs may streamline the development process.

CryptSDLC [70] was proposed as a solution to the technical complexity of cryptographic engineering. It does not explicitly discuss privacy. However, it lays a strong foundation for privacy-preserving software development. For instance, CryptSDLC emphasizes tools and primitives to ensure the confidentiality of the system [70], which directly supports privacy by protecting personal data from unauthorized access and disclosure. However, the feasibility of CryptSDLC is demonstrated only in the field of cloud computing [70]. Therefore, we suggest the standardization of cryptographic primitives and their properties in CryptSDLC to enhance communication and interoperability between different systems and stakeholders, making it easier to adopt CryptSDLC in diverse software fields. For example, standardized cryptographic protocols may enable seamless integration of privacy-preserving features across different areas such as IoT devices, healthcare, or financial systems, ensuring consistent data protection. Further, standardizing cryptographic primitives may create a common vocabulary (i.e., common standardized cryptographic primitives such as AES-256 encryption or SHA-3 hashing, irrespective of the domain) that all stakeholders who work on different projects can understand. Additionally, integrating automated tools to support the CryptSDLC process [70] may streamline the design process and reduce human errors. It may help to reduce the likelihood of human errors and speed up the development process. For example, automatic tools that generate secure configurations for cryptographic libraries and validate that the libraries meet the required security and privacy requirements.

Further, the organization's privacy culture is improved with proper communication among the team members as per the survey conducted by Tahel et al. [94], which eventually positively affects project development. The architecture used in CryptSLDC facilitates communication and collaboration among stakeholders such as cryptographers, software developers, and application designers [70], which may help to build the organization's privacy culture.

Next, we discussed a Model-Driven Development (MDD) methodology [64], which aimed to address the challenge that developers face when developing systems specifically with purpose limitation and user consent, as mandated by GDPR. However, it is a simplified interpretation of GDPR. Extending its compliance scope beyond purpose limitation and user consent may make the methodology more impactful and scalable. Further, since the model transformation is supported only for C# and Python web applications [64], it covers a small portion of the development stacks. This may limit usability in large-scale systems where developers work with multiple programming languages. Therefore, extending this to other programming languages such as JavaScript for web applications, Java, and Dart for mobile applications may broaden the applicability of the methodology.

5.2.4 Framework-based design and development. Secure SDLC Framework [59], secD4CloudMobile [36], and Hails [54] are frameworks we discussed in Section 4.2.4 that attempted to enhance security and privacy integration in software development.

First, the Secure SDLC Framework [59] primarily focuses on security rather than privacy. Even though its primary focus is security, which is centered around the fundamental security principles of confidentiality, integrity, and availability [59], privacy is a subset of these. However, privacy-by-design (PbD) principles are not explicitly incorporated [59], limiting their effectiveness during the early stage of software development. Additionally, it consists of some manual mappings of secure SDLC standards and guidelines [59], making it more resource-intensive and possibly not possible to apply in large-scale projects with limited time and money constraints. Therefore, automating that process and integrating PbD principles may reduce the developer burden and improve its efficiency. Further, we recommend implementing a feedback loop to collect information from stakeholders such as developers, security, and privacy teams in order to stay up-to-date on emerging threats and understand how competitors have addressed them. Since this may improve the communication among stakeholders, it may improve the organization's security and privacy culture, as discussed in section 5.6.

In a related effort that does not specifically provide privacy support, Chimuco et al. proposed another framework targeting developers who work on cloud and mobile ecosystems [36]. Its validation was conducted on a small scale [36], limiting confidence in its effectiveness for real-world application. Therefore, further evaluations are needed to check the effectiveness in industry settings. Additionally, incorporating AI with all the modules of the framework has been emphasized by the authors [36]. For instance, to improve the accessibility of the text-based console in the framework, a natural language processing-based developer interaction has been proposed [36]. By integrating NLP, developers may be able to interact with the framework using natural language (e.g., typing or speaking in plain English) instead of navigating through a structured questionnaire. Further, the current framework relies on a rule-based system that generates recommendations based on predefined inputs [36]. However, NLP can enable the framework to understand the context of the developer's queries and ensure that the framework provides more tailored recommendations and relevant guidance.

Next, we discussed Hails, where the authors highlighted the risk of data privacy breaches in current web platforms like Facebook that do not have restrictions on how third-party applications handle user data [54]. Hails is built using the Haskell language, which is not as widely used as other languages like JavaScript or Python [54]. It may pose a barrier to adoption for some platform developers. For instance, a platform developer familiar with Python or JavaScript might find it challenging to transition to Haskell. Haskell's ecosystem and libraries for web development are less mature and less widely documented compared to those in more mainstream languages. That makes it harder for newcomers to find resources and community support. Additionally, the authors highlighted that the query system in Hails is simple,

which hinders performance in situations where complex queries are needed [54]. For example, if a platform like GitStar [54] needs to filter, aggregate, or join large datasets for a scenario like finding all repositories owned by users in a specific region, Hails' current query system may not be sufficient. Therefore, we suggest extending the Hails framework to support more queries, such as filtering, aggregations, joins, and techniques for optimizing query execution, such as indexing and caching. Furthermore, Hails does not provide adequate support for platform developers to become familiar with the framework. For example, developers may struggle to write code from scratch for common issues related to security and privacy policies (e.g., access control mechanisms). Therefore, we suggest developing a tool that can generate boilerplate codes for common application components, as well as a debugging tool to debug security and privacy policy issues. Additionally, integrating Hails with popular IDEs (e.g., Visual Studio Code or IntelliJ) will provide better code completion, linting, and debugging support.

5.2.5   Secure & privacy-preserving coding. In our SLR, we found two papers that proposed developer support to write secure code, FixDroid [78], and PrivacyCAT [73]. The primary focus of both solutions is to support writing secure code, but FixDroid provides general code security, while PrivacyCAT specifically detects privacy defects in code [73, 78].

We discussed how FixDroid was developed to enhance Android Lint by providing real-time feedback and quick fixes for developers [78]. However, it was developed to be used in Android environments [78], limiting its usability for iOS, web, and other programming environments. For example, this tool may not be effective for a developer or an organization that has an application written for Android, iOS, and web platforms since the development environments and the programming languages are different. They may need to look for alternative tools (other than for Android) that are compatible with different environments and languages. Therefore, we suggest extending FixDroid for other development environments, such as iOS development, web application development, and other languages, such as C++, JavaScript, and Python, to increase usability. FixDroid does not leverage developer-specific context to provide suggestions [78]. That means that FixDroid provides generic security and privacy feedback to all developers, regardless of their individual coding habits, experience level, or the specific context of the application they are building. Therefore, in the future, we suggest using machine learning to provide personalized developer support utilizing developer-specific contexts such as the developer's coding style, level of experience, historical security and privacy issues, and the context of the application (e.g., healthcare, social media).

On the other hand, PrivacyCAT was designed to detect privacy vulnerabilities in WhatsApp's code using both static and dynamic analysis and is designed to run on both WhatsApp client and server code [73]. PrivacyCAT has been evaluated only in WhatsApp [73], leaving its applicability in other large-scale organizations questionable. For instance, WhatsApp's architecture may not be the same as the structures of other applications. Therefore, its generalizability to different codebases, development teams, and organizational environments should be tested. Technically, applying machine learning along with the already used static and dynamic taint analysis to identify more patterns in privacy violations may improve the system's ability to detect new and emerging threats. Further, since this has been proposed for large organizations that handle large code bases, it may be helpful to use distributed analysis techniques to scale up PrivacyCAT to handle larger code bases distributed among different teams of the organization.

5.3   Manage & Secure Personal Data

We discussed in section 4.3 how PrivacyStreams [69] and Platys [75] offer developer support for managing and securing personal data.

PrivacyStreams comes with a static analyzer that can be used to check how personal data is accessed, processed, and also the granularity of its usage [69]. However, it cannot handle complex data operations such as machine learning (ML) transformations [69], limiting its applicability in scenarios like ML-based personalized recommendations, sentiment analysis, and user behavior predictions. Expanding PrivacyStreams to incorporate ML operations may significantly enhance its usability since it enables developers to perform advanced data processing while maintaining data usage transparency. Additionally, if we can create privacy policies, descriptions, and in-app privacy notices with the help of the static analyzer as used by the tools discussed in section 4.5, it will be useful for developers to provide end-users with detailed information about how their data is handled. Also, if an application relies on third-party libraries that do not use PrivacyStreams to handle personal data, the static analyzer fails to capture complete data processing details [69]. Developing a middleware layer that acts as a bridge between third-party libraries and PrivacyStreams may address this issue. This middleware should be able to intercept data requests and responses from third-party libraries and then automatically translate them into PrivacyStreams-compatible formats. Intercepting data requests can be done using API hooking (e.g., Frida [8]) and Aspect-Oriented Programming (e.g., Spring AOP [11]), while translating can be achieved through simple pre-defined mapping functions. Furthermore, if an application has been coded using code obfuscation (i.e., hiding) techniques, it will be a barrier for PrivacyStreams to analyze the application [69]. We suggest developing a metadata layer within PrivacyStreams to allow developers to annotate the obfuscated data pipelines describing their granularity and purpose. Then, the static analyzer can rely on this metadata to reconstruct the data flow and provide data transparency.

On the other hand, in Platys, we did not discuss the features that end-users have and their contribution to Platys functions in detail because the SLR mainly focuses on developer support. Platys utilizes active learning and semi-supervised learning methods [75]. However, the accuracy of the trained model depends on the accuracy of the place labels provided by the users [75]. If users provide incomplete or inconsistent place labels, the model may misclassify locations, leading to privacy risks or incorrect recommendations. Therefore, the suitability of machine learning should be evaluated again with different scenarios of users (e.g., high-mobility users, low-mobility users, privacy-conscious users, and privacy-indifferent users), and different machine learning techniques (e.g., transfer learning methods, ensemble models - random forests) can also be tested in parallel. Furthermore, the authors mentioned the high battery consumption of Platys [75], which remains unexplored. This is important to check because developers may be hesitant to use it in applications if it consumes too much battery power, as users may reject the application. We suggest comparing battery consumption with different complex and tiny machine-learning models to check whether it has a direct effect on the battery consumption.

5.4 Address Vulnerabilities in Third-party Dependencies

Third-party libraries and SDKs enhance development efficiency but introduce security and privacy risks if they are outdated or misconfigured [44, 77], as we discussed in Section 4.4. Even though the tools discussed, Up2Dep [77] and DataAvalanche.io [44], tried to address these challenges, they have their own limitations.

Since Up2Dep relies on LibScout [87] and Cognicrypt [42] to identify API changes between different versions of libraries and cryptographic misuses of libraries, respectively, Up2Dep inherits the limitations associated with them. For example, LibScout sometimes may not work if the functionality of APIs in libraries changes while the structure stays the same [77]. Integrating machine learning techniques may solve this by helping to recognize semantic changes in APIs. For example, natural language processing (NLP) models may help to analyze API documentation or code comments to detect functional changes that are not reflected in structural changes. Similarly, Cognicrypt sometimes

reports false positives about cryptographic misuse [77]. We suggest expanding the analysis of Up2Dep by incorporating more techniques such as static taint analysis (i.e., to identify sensitive data, track data flows, detect privacy-violating API calls), dynamic taint analysis (i.e., to monitor the application's behavior at runtime), and privacy policy compliance (i.e., to check for compliance with GDPR and CCPA, generate privacy reports). These add-ons may make Up2Dep a more robust tool for improving app security and privacy.

Up2Dep provides developers with information about vulnerabilities detected in third-party libraries using a manually updated database [77], which may become outdated or incomplete. On the other hand, the current version of DataAvalanche.io does not contain a reliable database of alternative libraries and SDKs [44], making it less practical for developers who may seek privacy-centered replacements. To address this, we propose creating a public, developer-maintained repository where developers can contribute vulnerability reports they encounter during development and recommend alternative libraries. For example, if a developer encounters a security or privacy vulnerability in a library like Solid.js [10], they can submit a report to the repository, ensuring that other developers are properly informed, and they can suggest an alternative library such as React.js [7]. This approach may help to keep the database more updated and comprehensive.

Additionally, even though DataAvalanche.io provides alternative libraries and SDKs, it does not provide a standardized framework to select the best-suited one for the application based on privacy implications. Therefore, we suggest creating a standardized framework for evaluating the privacy implications of SDKs and libraries, including criteria such as data-sharing practices, compliance with regulations like GDPR, and user reviews. It may allow developers to systematically compare and contrast different SDKs and libraries and make informed choices.

Furthermore, since neither Up2Dep nor DataAvalanche.io has an integration with IDEs [77], it may require developers to manually check vulnerabilities, their privacy implications, and alternative library versions. Therefore, incorporating them with popular integrated development environments (IDEs) will streamline the development process and encourage developers to use these tools more regularly. For example, providing them with real-time feedback within the IDE, such as highlighting insecure dependencies as developers write code, may significantly enhance usability. For instance, if a developer includes a vulnerable version of a library, the IDE can immediately flag it and suggest a secure alternative.

Lastly, as the authors mentioned, both solutions were evaluated on a small scale [44, 77], leaving concerns about their effectiveness across diverse applications. Therefore, more evaluations are needed in different industry conditions. For example, larger and more diverse groups of developers from different industries (e.g., healthcare, finance, etc.) and regions (e.g., Europe, Australia, etc.) may help identify potential gaps and ensure that the tools are flexible enough to meet the needs of a broader audience.

5.5 Make Software Privacy Compliant and Transparent

As shown in Figure 4, we identified tools from three different categories that help developers make software privacy-compliant and transparent. Below, we discuss the implications of the results and potential future directions based on their limitations.

5.5.1 Privacy label generation. We discussed in Section 4.5.1 how Privacy Label Wiz [52], Matcha [66], and Jain et al.'s methodology [58] helped developers in the privacy label generation process. Some of these solutions utilize static code analysis [52, 66]. However, accurately capturing all data-handling practices needs some further implementation. The authors of both Privacy Label Wiz and Matcha believe that incorporating dynamic analysis techniques to observe runtime behavior may improve their accuracy since static analysis alone cannot capture all data flows or interactions

[52, 66]. For example, static analysis can identify that the application accesses the user's location, but it cannot determine whether the application actually collects location data at runtime, and the frequency of doing it. On the other hand, Dynamic analysis can monitor the application's behavior in real time. It means it can detect when the application accesses location data when the user interacts with a specific feature, such as a map. This may help developers provide more precise information about the privacy labels. Additionally, the authors of Matcha suggested incorporating a machine learning model, Codex (a large programming model), to enhance the analysis process [66]. ML-driven dynamic code analysis and pattern recognition may help both tools identify privacy-sensitive code segments and track data flows more effectively, enhancing the dynamic analysis process.

Further, PLW cannot determine what specific data third-party libraries collect that are integrated into an application [52]. Since privacy violations may occur at the library level, such as third-party libraries collecting or sharing user data without the developer's awareness [44, 77], this limitation may impact the reliability of generated privacy labels. For instance, Firebase might collect user analytics or device identifiers, but the static analyzer of PLW may not be able to identify these practices. Therefore, we suggest using synthetic data generated based on the application context and for different user interactions. For instance, a user logs into the app with synthetic data like email and password. Injecting this data into third-party libraries may help to simulate how libraries behave for different user interactions and identify what specific data the libraries collect. Furthermore, to improve the accuracy of detecting third-party libraries and data collection practices, PLW could be extended to include a centralized repository of privacy-related information (e.g., detailed documentation, privacy policies, and data collection practices provided by third-party library developers) for commonly used third-party libraries [52].

Finally, merging both Matcha and Privacy Label Wiz to make a combined solution for both Apple and Google platforms may help developers generate privacy labels for both platforms within a single solution.

On the other hand, Jain et al.'s fine-grained localization methodology posed some limitations [58], questioning its usability in practical software development scenarios. One of the main concerns is its limited evaluation in industry settings [58]. Usually, developers have limited time to complete the software development project [61]. Since the proposed methodology is time-consuming and requires manual developer intervention [58], it may be effective to evaluate this in different industry settings where developers have short-term and long-term deadlines. In addition, since the evaluations were done for Android applications [58], it may not be effective for developers in organizations where they have the application for both Android and iOS versions, because they require another approach for iOS. Therefore, evaluating it on other platforms and ensuring its effectiveness may improve its usability. This also applies to Matcha and the Privacy Label Wiz.

5.5.2 Privacy policy generation. Under the privacy policy generation, we identified tools, PrivacyFlash Pro [100] and Coconut [65], that especially target individual developers or small organizations. The major difference between these tools is their focus. For instance, Coconut helps developers understand and enhance their app's data-handling practices by providing real-time feedback during development, which helps them write better privacy policy statements. However, it does not automatically generate these statements [65]. In contrast, PrivacyFlash Pro is designed to automate the entire privacy policy generation process [100]. However, both solutions have some limitations that hinder their full effectiveness.

As we discussed, PrivacyFlash Pro combines static code analysis with a questionnaire wizard [100]. It implies that it relies on developer-provided answers. However, reliance on developers' responses may provide inaccuracies that may result in misleading privacy statements because the responses may depend on the developers' privacy knowledge.

Since Coconut can provide real-time feedback and suggestions to improve data handling practices [65], merging that functionality of Coconut with PrivacyFlash Pro may increase the accuracy of developers' responses and ultimately make accurate privacy policies. We suggest improving Coconut's feedback mechanism by implementing more advanced analysis techniques, such as machine learning-based analysis, to give developers feedback based on their privacy knowledge by studying their development behavior. Also, providing feedback for any misalignment with regulatory frameworks like GDPR and CCPA may offer more specific guidance. At the same time, both tools mentioned that the code analysis process requires further improvement [65, 100]. The Python-based code analysis of PrivacyFlash Pro is not entirely accurate and produces some false positives and negatives [100]. Also, developer-based evaluation results of Coconut showed that developers have a lack of trust in automatically filled values of annotations [65]. These imply that the code analysis techniques are not accurate [65, 100]. Therefore, as we discussed in Sections 5.5.1, 5.4, and 5.2.2, we suggest trying ML-based approaches for the code analysis. One potential solution to address PrivacyFlash Pro's limitation of not considering server-side data sharing might be the integration of machine learning (ML) approaches as part of a dynamic analysis. For example, an ML model can be trained to analyze server logs and API calls in real time. Further, extending the PrivacyFlash Pro's platform support may increase the usability as we discussed for some other approaches in Sections 5.2.5 and 5.5.1. Lastly, further evaluations are needed in different industry settings, as we discussed in Sections 5.5.1 and 5.4 to ensure wide usability.

5.5.3 Privacy notice generation. We discussed in Section 4.5.3 a tool, Honeysuckle [68], that helps developers generate privacy notices for Android applications. Honeysuckle relies on developers to get accurate annotations [68], which could lead to misleading notices if developers do not provide accurate information. Additionally, Honeysuckle has limitations in its code analysis capabilities, such as being simple and adapted from the Coconut tool, and also needs developers' manual intervention [65, 68], implying it may fail to detect complex data flows and third-party library interactions. For instance, if the application collects user data in one function, processes it in another, and then sends it to a third-party library, Honeysuckle may miss this data flow, leading to incomplete privacy notices. Developers' manual intervention may lead to errors, as we discussed in Section 5.5.2. At the same time, Honeysuckle provides limited support for third-party libraries (i.e., it might not fully capture the data-handling practice unless the library is explicitly supported) [68], reducing its effectiveness in real-world applications. Therefore, automating some aspects of annotation using code analysis techniques used by Matcha [66], PLW [52], and PrivacyFlash Pro [100] may provide better detection of privacy-related data flows, especially in libraries, and reduce developer workload while improving the accuracy. Additionally, we suggest using ML-driven analysis techniques as we suggested for other approaches in Sections such as 5.5.1, 5.4, 5.2.2, not only to analyze the code but also to generate privacy notice UIs based on the application domain and context. For example, notices (i.e., UI-based notices) for children's applications may differ from those for other applications (e.g., they should be more explainable) to provide them with more contextualized privacy information. Multiple language support may also be useful to help users make more informed decisions using their preferred language. Further, extending the support of Honeycucke for other web and mobile platforms may improve its usability, as we discussed in other Sections (e.g., Sections 5.2.5 and 5.5.1).

5.6   Improve Privacy Education and Awareness

We discussed how Privacy Ideation Cards (PIC) [71, 95] and workshop-based interventions [97] tried to improve the privacy education of software developers in section 4.6. However, their effectiveness across different developer experience levels and industrial settings is questionable because of their limitations.

One of the limitations of PIC is that it has only been tested among software students [95]. Even though the study showed that the students engaged in deeper discussions while using PICs [71, 95], it is still questionable whether industry professionals who face tight deadlines and cognitive tasks [12, 28, 61] would get the benefit of this in industry settings. Similarly, the workshop-based intervention proposed by Weir et al. [97] helped privacy-related communication between developers and product managers who have limited privacy expertise, raising concerns about its applicability among professionals who have some level of privacy expertise. Furthermore, these workshops' effectiveness relies on a supportive leader [97]. It may limit the applicability of this tool among organizations, as not all organizations may have a privacy-oriented leader. Here, the leader in this context refers to the development team's leader (e.g., architect, principal architect, or chief technology officer). Therefore, both solutions should be evaluated further in different industry settings, as discussed in Sections 5.5.1 and 5.4.

Increasing privacy awareness in an organization may be challenging since a team will represent members with different mentalities. For instance, one may have strong privacy attitudes while one may have an "I have nothing to hide" mentality or negative privacy attitudes [92, 94]. If they consider others' privacy in the same direction as they thought, it will affect software as well as the whole team's mentality [94]. Therefore, organizations should focus on educating developers about privacy attitudes, knowledge about human rights, social benefits, and empathy toward users to promote privacy in organizations (i.e., increasing privacy culture) [94].

In addition to that, it is recommended to educate developers about the practical implementation of privacy, targeting their roles in the development team rather than providing them general privacy awareness [94]. Design and code review, as well as mentoring programs, are the recommended approaches to educate practical privacy implementation [94].

With these, we can argue that embedding privacy education into the software development process rather than taking it as an external training activity may help organizations develop a strong privacy culture.

6 THREATS TO VALIDITY

Following the SLR guidelines proposed by Kitchenham and Charters, we identified several potential threats that could affect the validity of this review. Even though we made efforts to conduct the study rigorously, several risks may remain that readers should consider when interpreting the results.

6.0.1 Selection and Inclusion Bias. Although we used clearly defined inclusion and exclusion criteria, there is a possibility that some relevant studies were unintentionally excluded. For example, there may be some studies that addressed developer-supporting privacy solutions but using a different terminology (e.g., keywords) other than what we used in our search strategy. Additionally, the SLR is exposed to a publication bias as we focus solely on peer-reviewed academic sources. This could result in overlooking valuable insights from industry or non-academic sources.

6.0.2 Coder Bias. Although we followed a well-structured protocol (i.e., reflexive thematic analysis) to reduce the subjectiveness in decisions, the study may have introduced a coder bias when classifying the results, since there is a degree of human judgment in the process.

7 LIMITATIONS

This SLR was conducted to offer a comprehensive overview of the topic while ensuring the reproducibility of the reported results in the literature. First, the returned articles for the search query were filtered using their titles and abstracts according to our inclusion and exclusion criteria. In that study selection phase, relevant articles may be overlooked. Therefore, to avoid such scenarios as much as possible, we conducted both forward and backward snowballing searches

[98] for both articles selected from the results of the search query and from top-tier journals and conferences. Once the articles were ready, we thematically analyzed the selected articles to answer the research questions mentioned in Section 1. The first author performed the initial coding process and theme identification. However, the generated codes and themes may be influenced by the coder's experience, knowledge, and perspective, which potentially introduces bias. To minimize this bias, as recommended by Braun and Clarke [31], we incorporated the perspectives of all authors when developing themes, as discussed in Section 3.3. Further, the SLR primarily focused on studies published in academic sources, which may result in the potential overlooking of contributions from industry practice.

8 CONCLUSION AND FUTURE WORK

This SLR explored and analyzed existing solutions, including tools, guidelines, methods, methodologies, and frameworks in the current literature to address the challenges that developers face in integrating privacy into software development while supporting them in the privacy integration process. Through this extensive review, we identified that the existing developer-supporting solutions have been proposed aiming to address a primary set of common developer challenges such as lack of privacy expertise among developers [17, 52, 65, 66, 68, 100], difficulties in translating privacy principles into technical implementations [14, 26, 64, 83, 95, 97], inadequate regulatory guidance [17, 83, 95, 97], and the increasing privacy and security risks associated with third-party libraries and SDKs [44, 54, 77]. In Section 4.1, we discussed tools and methods to support developers in embedding privacy in the early stage of software development. It included tools like PCM-tool [83], RMCM [72], LINDUNN-Go [99], and a method called ThreatPoker [89]. For design and development, tools like Parrot [15] and Canella [17] provide interactive environments for embedding privacy into IoT applications, while POSD (Privacy-Oriented Software Development) [26] offers structured guidelines to facilitate privacy integration. Additionally, methodologies and frameworks have also been proposed to help developers in the software design and development process. It included frameworks such as CIA-level driven SDLC [59], secD4CloudMobile [36], and Hails [54], and methodologies such as PbE [28], CryptSDLC [70], and UML-based MDD [64]. To enforce privacy in coding, tools like FixDroid [78] and PrivacyCAT [73] help detect vulnerabilities and suggest fixes to improve privacy in Android applications, and specifically in the WhatsApp application. Managing third-party dependencies is supported by tools like Up2Dep [77] and DataAvalanche.io [44], which help developers avoid insecure or non-compliant SDKs and libraries. Further, PrivacyStreams [69] and Platys [75] are tools and frameworks, respectively, that have been proposed to help developers manage and secure personal data when developing software applications. Furthermore, to assist developers in understanding privacy principles, regulations, and secure development practices, approaches like Privacy Ideation Cards (PICs) [95] and a workshop-based [97] intervention have been proposed. Finally, generating privacy statements and policies is made easier with tools like Matcha [66], Privacy Label Wiz [52], PrivacyFlash Pro [100], Coconut [65], and Honeysuckle [68] which assist developers in generating accurate privacy labels, policies, and notices to align with compliance requirements.

However, the identified tools, guidelines, methods, methodologies, and frameworks discussed in the results section (i.e., Section 4) offer only partial solutions because of the limitations we discussed in the discussion section (i.e., Section 5). Many solutions require prior expertise in privacy and security [28, 65, 72], and require developers' manual intervention [26, 58, 59, 65, 68, 83, 100], posing a barrier for developers with limited knowledge in these areas to use these solutions, as well as the possibility of human errors. Some solutions do not provide seamless integration with agile and real-world development workflows [26, 83], which causes difficulties in adoption for rapid software development environments. Additionally, many solutions are designed for specific environments (e.g., Android) or programming languages (e.g., Java) [26, 52, 54, 64, 66, 69, 78, 100], limiting their applicability across diverse software development ecosystems. Further, most

of the solutions need further evaluations in different industry settings [15, 26, 36, 44, 52, 58, 66, 70, 73, 77, 83, 84, 89, 95, 99]. Furthermore, some of the solutions provide their support for a selected set of regulations or do not cover the entire regulation, limiting their usability [15, 17]. We discussed in detail how these limitations affect the developers in Section 5. These limitations imply that they fail to address the developer challenges and necessitate that additional research is required.

Therefore, as discussed in the discussion section, we propose future improvements to each developer-supporting approach to make them more effective. First, since we identified that developers' lack of privacy knowledge and inaccuracy of the analysis mechanisms (e.g., code analysis) as critical concerns in the proposed solutions, we suggested to incorporate machine learning approaches (e.g., NLP), to reduce the developers' manual intervention by automating the integrated processes, increase the accuracy (e.g., the accuracy of privacy label generation), and also to reduce the human errors happen because of the lack of privacy knowledge. However, because machine learning solutions are not 100% accurate, developers' feedback may be important at some point. This implies that developers' privacy knowledge may be critical at some point. Therefore, future research should be more focused on enhancing developers' privacy knowledge either as a direct or an indirect goal. Here, indirect refers to the fact that the privacy knowledge is improved while developers engage in some other task, such as privacy policy generation. We discussed how we can provide this for some of the existing solutions in detail in the discussion Section (i.e., Section 5. Next, we suggested expanding some solutions in a way to support more programming languages (e.g., Dart, C++) and platforms (i.e., Android, iOS) in order to increase usability. It helps developers and organizations who work with different tech stacks to rely on the same tool, guideline, method, methodology, or framework regardless of the tech stack. Additionally, we suggested extending some of these solutions to support broader regulatory compliance (e.g., GDPR - Europe, CCPA - USA, etc.) as it is essential for developers and organizations who develop software applications targeting different regions. Further, we suggested conducting extensive empirical evaluations for most of the existing solutions to check their effectiveness and suitability in different industry settings. We explained in detail what kind of evaluations are needed in Section 5.

Finally, regarding the SLR itself, we acknowledge that our study mainly focused on peer-reviewed academic literature. To further improve the comprehensiveness of future SLRs, we suggest incorporating grey literature sources such as technical blogs, industry reports, and open-source repositories. Including these sources will help to explore popular but non-academic solutions (e.g., tools, methods, etc) which offer additional insights about real-world privacy integration challenges and solutions.

In summary, while existing developer-supporting solutions provide developers with assistance in integrating privacy into software development, their limitations highlight that further enhancements are needed. Therefore, addressing these gaps, as we discussed, will be crucial to making these solutions more effective and accessible for developers across diverse software development environments.

REFERENCES

[1] 1988. Privacy Act 1988 (Australia). https://www.legislation.gov.au/Details/C2019C00241
[2] 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX: 32016R0679
[3] 2018. California Consumer Privacy Act (CCPA). https://oag.ca.gov/privacy/ccpa
[4] 2020. Privacy Act 2020 (New Zealand). https://www.legislation.govt.nz/act/public/2020/0031/latest/LMS23223.html
[5] 2022. Google Play Developer Distribution Agreement. https://play.google/developer-distribution-agreement.html Accessed: 2024-09-23.
[6] 2024. Apple Developer Program License Agreement. https://developer.apple.com/support/terms/apple-developer-program-license-agreement/ Accessed: 2024-09-23.
[7] 2024. *React Js.* https://react.dev/ Accessed: 2025-03-11.

[8]   2025. *Frida - A World-Class Dynamic Instrumentation Toolkit.* https://frida.re/ Accessed: 2025-03-11.

[9]   2025. The Heartbleed Bug. https://heartbleed.com/ Accessed: 2025-03-11.

[10] 2025. *SolidJS.* https://www.solidjs.com/ Accessed: 2025-03-11.

[11] 2025. *Spring Framework Documentation - Aspect-Oriented Programming (AOP).* https://docs.spring.io/spring-framework/reference/core/aop.html Accessed: 2025-03-11.

[12] Yasemin Acar, Michael Backes, Sascha Fahl, Doowon Kim, Michelle L. Mazurek, and Christian Stransky. 2016. You Get Where You're Looking for: The Impact of Information Sources on Code Security. In *IEEE Symposium on Security and Privacy, SP 2016, San Jose, CA, USA, May 22-26, 2016.* IEEE Computer Society, 289–305. https://doi.org/10.1109/SP.2016.25

[13] Mamoun Alazab, Seung Hun Hong, and Jenny Ng. 2021. Louder bark with no bite: Privacy protection through the regulation of mandatory data breach notification in Australia. *Future Generation Computer Systems* 116 (March 2021), 22–29. https://doi.org/10.1016/j.future.2020.10.017

[14] Nada Alhirabi, Stephanie Beaumont, Jose Tomas Llanos, Dulani Meedeniya, Omer Rana, and Charith Perera. 2023. PARROT: Interactive Privacy-Aware Internet of Things Application Design Tool. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 7, 1, Article 1 (March 2023), 37 pages. https://doi.org/10.1145/3580880

[15] Nada Alhirabi, Stephanie Beaumont, Omer Rana, and Charith Perera. 2024. Designing Privacy-Aware IoT Applications for Unregulated Domains. 5, 2, Article 11 (April 2024), 32 pages. https://doi.org/10.1145/3648480

[16] Atheer Aljeraisy, Masoud Barati, Omer Rana, and Charith Perera. 2021. Privacy Laws and Privacy by Design Schemes for the Internet of Things: A Developer's Perspective. *ACM Comput. Surv.* 54, 5, Article 102 (may 2021), 38 pages. https://doi.org/10.1145/3450965

[17] Atheer Aljeraisy, Omer Rana, and Charith Perera. 2024. Empowering IoT Developers with Privacy-Preserving End-User Development Tools. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8, 3, Article 90 (Sept. 2024), 47 pages. https://doi.org/10.1145/3678588

[18] Majed Alshammari and Andrew Simpson. 2017. Towards a Principled Approach for Engineering Privacy by Design. In *Privacy Technologies and Policy*, Erich Schweighofer, Herbert Leitold, Andreas Mitrakas, and Kai Rannenberg (Eds.). Springer International Publishing, Cham, 161–177.

[19] Benjamin Andow, Samin Yaseer Mahmud, Justin Whitaker, William Enck, Bradley Reaves, Kapil Singh, and Serge Egelman. 2020. Actions speak louder than words: entity-sensitive privacy policy and data flow analysis with POLICHECK *(SEC'20)*. USENIX Association, USA, Article 56, 18 pages.

[20] Vinícius Camargo Andrade, Rhodrigo Deda Gomes, Sheila Reinehr, Cinthia Obladen De Almendra Freitas, and Andreia Malucelli. 2023. Privacy by Design and Software Engineering: a Systematic Literature Review. In *Proceedings of the XXI Brazilian Symposium on Software Quality* (Curitiba, Brazil) *(SBQS '22)*. Association for Computing Machinery, New York, NY, USA, Article 18, 10 pages. https://doi.org/10.1145/3571473.3571480

[21] Atlassian. 2024. Jira Software. https://www.atlassian.com/software/jira Accessed: 2024-10-29.

[22] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. 2004. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing* 1, 1 (2004), 11–33. https://doi.org/10.1109/TDSC.2004.2

[23] Vanessa Ayala-Rivera and Liliana Pasquale. 2018. The Grace Period Has Ended: An Approach to Operationalize GDPR Requirements. 136–146. https://doi.org/10.1109/RE.2018.00023

[24] Michael Backes, Sven Bugiel, and Erik Derr. 2016. Reliable Third-Party Library Detection in Android and its Security Applications. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (Vienna, Austria) *(CCS '16)*. Association for Computing Machinery, New York, NY, USA, 356–367. https://doi.org/10.1145/2976749.2978333

[25] Maria Teresa Baldassarre, Vita Santa Barletta, Danilo Caivano, Antonio Piccinno, and Michele Scalera. 2022. Privacy Knowledge Base for Supporting Decision-Making in Software Development. In *Sense, Feel, Design*, Carmelo Ardito, Rosa Lanzilotti, Alessio Malizia, Marta Larusdottir, Lucio Davide Spano, José Campos, Morten Hertzum, Tilo Mentler, José Abdelnour Nocera, Lara Piccolo, Stefan Sauer, and Gerrit van der Veer (Eds.). Springer International Publishing, Cham, 147–157.

[26] Maria Teresa Baldassarre, Vita Santa Barletta, Danilo Caivano, and Michele Scalera. 2020. Integrating security and privacy in software development. *Software Quality Journal* 28, 3 (Sept. 2020), 987–1018. https://doi.org/10.1007/s11219-020-09501-6

[27] Rebecca Balebako, Abigail Marsh, Jialiu Lin, Jason Hong, and Lorrie Cranor. 2014. The Privacy and Security Behaviors of Smartphone App Developers. https://doi.org/10.14722/usec.2014.23006

[28] Pedro Barbosa, Andrey Brito, and Hyggo Almeida. 2020. Privacy by Evidence: A Methodology to develop privacy-friendly software applications. *Information Sciences* 527 (2020), 294–310. https://doi.org/10.1016/j.ins.2019.09.040

[29] Reuben Binns, Jun Zhao, Max Van Kleek, and Nigel Shadbolt. 2018. Measuring Third-party Tracker Power across Web and Mobile. *ACM Trans. Internet Technol.* 18, 4, Article 52 (Aug. 2018), 22 pages. https://doi.org/10.1145/3176246

[30] Maisha Boteju, Thilina Ranbaduge, Dinusha Vatsalan, and Nalin Asanka Gamagedara Arachchilage. 2023. SoK: Demystifying Privacy Enhancing Technologies Through the Lens of Software Developers. arXiv:2401.00879 [cs.SE] https://arxiv.org/abs/2401.00879

[31] Virginia Braun and Victoria Clarke. 2012. Thematic Analysis. American Psychological Association.

[32] Edna Dias Canedo, Ian Nery Bandeira, Angelica Toffano Seidel Calazans, Pedro Henrique Teixeira Costa, Emille Catarine Rodrigues Cançado, and Rodrigo Bonifácio. 2022. Privacy requirements elicitation: a systematic literature review and perception analysis of IT practitioners. *Requir. Eng.* 28, 2 (jun 2022), 177–194. https://doi.org/10.1007/s00766-022-00382-8

[33] Ann Cavoukian. 2009. Privacy by design: The 7 foundational principles: Implementation and mapping of fair information practices. *Information and Privacy Commissioner of Ontario, Canada* (2009).

[34] Ann Cavoukian. 2012. Operationalizing Privacy by Design: A Guide to Implementing Strong Privacy Practices. https://gpsbydesigncentre.com/wp-content/uploads/2021/08/Doc-5-Operationalizing-pbd-guide.pdf Accessed: 2024-10-15.

[35] Gauthier Chassang. 2017. The impact of the EU general data protection regulation on scientific research. *ecancer* 11 (2017), 709. https://doi.org/10.3332/ecancer.2017.709

[36] Francisco Chimuco, Bernardo Sequeiros, Tiago Simões, Mário Freire, and Pedro Inácio. 2024. Expediting the design and development of secure cloud-based mobile apps. *International Journal of Information Security* 23 (07 2024), 1–22. https://doi.org/10.1007/s10207-024-00880-6

[37] Luca Compagna, Paul Khoury, Alžběta Solarczyk Krausová, Fabio Massacci, and Nicola Zannone. 2009. How to integrate legal requirements into a requirements engineering methodology for the development of security and privacy patterns. *Artificial Intelligence and Law* 17 (03 2009), 1–30. https://doi.org/10.1007/s10506-008-9067-3

[38] Daniela S. Cruzes and Tore Dyba. 2011. Recommended Steps for Thematic Synthesis in Software Engineering. In *2011 International Symposium on Empirical Software Engineering and Measurement*. 275–284. https://doi.org/10.1109/ESEM.2011.36

[39] Shirlei Aparecida de Chaves and Fabiane Barreto Vavassori Benitti. 2023. Privacy by Design in Software Engineering: An update of a Systematic Mapping Study *(SAC '23)*. Association for Computing Machinery, New York, NY, USA, 1362–1369. https://doi.org/10.1145/3555776.3577626

[40] Erik Derr, Sven Bugiel, Sascha Fahl, Yasemin Acar, and Michael Backes. 2017. Keep me Updated: An Empirical Study of Third-Party Library Updatability on Android *(CCS '17)*. Association for Computing Machinery, New York, NY, USA, 2187–2200. https://doi.org/10.1145/3133956.3134059

[41] Julia Earp. 2003. A Requirements Taxonomy to Reduce Website Privacy Vulnerabilities. *Requirements Engineering - RE* (01 2003).

[42] Eclipse Foundation. 2023. CogniCrypt: Secure Cryptography for Java Developers. https://eclipse.dev/cognicrypt/ Accessed: 2024-09-25.

[43] Manuel Egele, David Brumley, Yanick Fratantonio, and Christopher Kruegel. 2013. An empirical study of cryptographic misuse in android applications. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security* (Berlin, Germany) *(CCS '13)*. Association for Computing Machinery, New York, NY, USA, 73–84. https://doi.org/10.1145/2508859.2516693

[44] Anirudh Ekambaranathan, Jun Zhao, and George Chalhoub. 2023. Navigating the Data Avalanche: Towards Supporting Developers in Developing Privacy-Friendly Children's Apps. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 7, 2, Article 53 (June 2023), 24 pages. https://doi.org/10.1145/3596267

[45] Anirudh Ekambaranathan, Jun Zhao, and Max Van Kleek. 2020. Understanding Value and Design Choices Made by Android Family App Developers. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '20)*. Association for Computing Machinery, New York, NY, USA, 1–10. https://doi.org/10.1145/3334480.3383064

[46] Anirudh Ekambaranathan, Jun Zhao, and Max Van Kleek. 2021. "Money makes the world go around": Identifying Barriers to Better Privacy in Children's Apps From Developers' Perspectives. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 46, 15 pages. https://doi.org/10.1145/3411764.3445599

[47] William Enck, Damien Octeau, Patrick McDaniel, and Swarat Chaudhuri. 2011. A Study of Android Application Security. In *20th USENIX Security Symposium (USENIX Security 11)*. USENIX Association, San Francisco, CA. https://www.usenix.org/conference/usenixsecurity11/study-android-application-security

[48] ENISA. 2018. Privacy and Data Protection in Mobile Applications: A Study on the App Development Ecosystem and the Technical Implementation of GDPR.

[49] Adrienne Porter Felt, Erika Chin, Steve Hanna, Dawn Song, and David Wagner. 2011. Android permissions demystified. In *Proceedings of the 18th ACM Conference on Computer and Communications Security* (Chicago, Illinois, USA) *(CCS '11)*. Association for Computing Machinery, New York, NY, USA, 627–638. https://doi.org/10.1145/2046707.2046779

[50] Adrienne Porter Felt, Helen J. Wang, Alexander Moshchuk, Steve Hanna, and Erika Chin. 2011. Permission Re-Delegation: Attacks and Defenses. In *20th USENIX Security Symposium (USENIX Security 11)*. USENIX Association, San Francisco, CA. https://www.usenix.org/conference/usenixsecurity11/permission-re-delegation-attacks-and-defenses

[51] Iñigo Fernández del Amo, John Ahmet Erkoyuncu, Rajkumar Roy, Riccardo Palmarini, and Demetrius Onoufriou. 2018. A systematic review of Augmented Reality content-related techniques for knowledge transfer in maintenance applications. *Computers in Industry* 103 (2018), 47–71. https://doi.org/10.1016/j.compind.2018.08.007

[52] Jack Gardner, Yuanyuan Feng, Kayla Reiman, Zhi Lin, Akshath Jain, and Norman Sadeh. 2022. Helping Mobile Application Developers Create Accurate Privacy Labels. In *2022 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. 212–230. https://doi.org/10.1109/EuroSPW55150.2022.00028

[53] Mohamad Gharib, John Mylopoulos, and Paolo Giorgini. 2020. COPri - A Core Ontology for Privacy Requirements Engineering. In *Research Challenges in Information Science*, Fabiano Dalpiaz, Jelena Zdravkovic, and Pericles Loucopoulos (Eds.). Springer International Publishing, Cham, 472–489.

[54] Daniel Giffin, Amit Levy, Deian Stefan, David Terei, David Mazières, John Mitchell, and Alejandro Russo. 2012. Hails: protecting data privacy in untrusted web applications. 47–60.

[55] Irit Hadar, Tomer Hasson, Oshrat Ayalon, Eran Toch, Michael Birnhack, Sofia Sherman, and Arod Balissa. 2018. Privacy by designers: software developers' privacy mindset. *Empirical Softw. Engg.* 23, 1 (Feb. 2018), 259–289. https://doi.org/10.1007/s10664-017-9517-1

[56] Katie Harbath and Collier Fernekes. 2023. History of the Cambridge Analytica Controversy. https://bipartisanpolicy.org/blog/cambridge-analytica-controversy/

[57] Jaap-Henk Hoepman. 2014. Privacy Design Strategies. In *ICT Systems Security and Privacy Protection*, Nora Cuppens-Boulahia, Frédéric Cuppens, Sushil Jajodia, Anas Abou El Kalam, and Thierry Sans (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 446–459.

[58] Vijayanta Jain, Sepideh Ghanavati, Sai Teja Peddinti, and Collin McMillan. 2023. Towards Fine-Grained Localization of Privacy Behaviors. In *2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P)*. 258–277. https://doi.org/10.1109/EuroSP57164.2023.00024

[59] Sooyoung Kang and Seungjoo Kim. 2022. CIA-level driven secure SDLC framework for integrating security into SDLC process. *Journal of Ambient Intelligence and Humanized Computing* 13 (03 2022). https://doi.org/10.1007/s12652-021-03450-z

[60] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. 2009. A "nutrition label" for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security* (Mountain View, California, USA) *(SOUPS '09)*. Association for Computing Machinery, New York, NY, USA, Article 4, 12 pages. https://doi.org/10.1145/1572532.1572538

[61] Miryung Kim, Lawrence Bergman, Tessa Lau, and David Notkin. 2004. An Ethnographic Study of Copy and Paste Programming Practices in OOPL. In *Proceedings of the 2004 International Symposium on Empirical Software Engineering (ISESE '04)*. IEEE Computer Society, USA, 83–92.

[62] Barbara Kitchenham and Stuart Charters. 2007. Guidelines for performing Systematic Literature Reviews in Software Engineering. 2 (01 2007).

[63] Konrad Kollnig, Pierre Dewitte, Max Van Kleek, Ge Wang, Daniel Omeiza, Helena Webb, and Nigel Shadbolt. 2021. A Fait Accompli? An Empirical Study into the Absence of Consent to Third-Party Tracking in Android Apps. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*. USENIX Association, 181–196. https://www.usenix.org/conference/soups2021/presentation/kollnig

[64] Srđan Krstić, Hoang Nguyen, and David Basin. 2024. Model-driven Privacy. *Proceedings on Privacy Enhancing Technologies* 2024 (01 2024), 314–329. https://doi.org/10.56553/popets-2024-0018

[65] Tianshi Li, Yuvraj Agarwal, and Jason I. Hong. 2018. Coconut: An IDE Plugin for Developing Privacy-Friendly Apps. 2, 4, Article 178 (Dec. 2018), 35 pages. https://doi.org/10.1145/3287056

[66] Tianshi Li, Lorrie Faith Cranor, Yuvraj Agarwal, and Jason I. Hong. 2024. Matcha: An IDE Plugin for Creating Accurate Privacy Nutrition Labels. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8, 1, Article 33 (mar 2024), 38 pages. https://doi.org/10.1145/3643544

[67] Tianshi Li, Elizabeth Louie, Laura Dabbish, and Jason I. Hong. 2021. How Developers Talk About Personal Data and What It Means for User Privacy: A Case Study of a Developer Forum on Reddit. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW3, Article 220 (Jan. 2021), 28 pages. https://doi.org/10.1145/3432919

[68] Tianshi Li, Elijah B. Neundorfer, Yuvraj Agarwal, and Jason I. Hong. 2021. Honeysuckle: Annotation-Guided Code Generation of In-App Privacy Notices. 5, 3, Article 112 (sep 2021), 27 pages. https://doi.org/10.1145/3478097

[69] Yuanchun Li, Fanglin Chen, Toby Jia-Jun Li, Yao Guo, Gang Huang, Matthew Fredrikson, Yuvraj Agarwal, and Jason I. Hong. 2017. PrivacyStreams: Enabling Transparency in Personal Data Processing for Mobile Apps. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 76 (Sept. 2017), 26 pages. https://doi.org/10.1145/3130941

[70] Thomas Loruenser, Henrich C. Pöhls, Leon Sell, and Thomas Laenger. 2018. CryptSDLC: Embedding Cryptographic Engineering into Secure Software Development Lifecycle. In *Proceedings of the 13th International Conference on Availability, Reliability and Security* (Hamburg, Germany) *(ARES '18)*. Association for Computing Machinery, New York, NY, USA, Article 4, 9 pages. https://doi.org/10.1145/3230833.3233765

[71] Ewa Luger, Lachlan Urquhart, Tom Rodden, and Michael Golembewski. 2015. Playing the Legal Card: Using Ideation Cards to Raise Data Protection Issues within the Design Process. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 457–466. https://doi.org/10.1145/2702123.2702142

[72] Phu X. Mai, Arda Goknil, Lwin Khin Shar, Fabrizio Pastore, Lionel C. Briand, and Shaban Shaame. 2018. Modeling Security and Privacy Requirements: a Use Case-Driven Approach. *Information and Software Technology* 100 (2018), 165–182. https://doi.org/10.1016/j.infsof.2018.04.007

[73] Ke Mao, Cons Åhs, Sopot Cela, Dino Distefano, Nick Gardner, Radu Grigore, Per Gustafsson, Ákos Hajdu, Timotej Kapus, Matteo Marescotti, Gabriela Cunha Sampaio, and Thibault Suzanne. 2024. PrivacyCAT: Privacy-Aware Code Analysis at Scale. In *Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Practice* (Lisbon, Portugal) *(ICSE-SEIP '24)*. Association for Computing Machinery, New York, NY, USA, 106–117. https://doi.org/10.1145/3639477.3639742

[74] Gary McGraw. 2002. Building Secure Software: Better than Protecting Bad Software. *Software, IEEE* 19 (12 2002), 57 – 58. https://doi.org/10.1109/MS.2002.1049391

[75] Pradeep K. Murukannaiah and Munindar P. Singh. 2015. Platys: An Active Learning Framework for Place-Aware Application Development and Its Evaluation. *ACM Trans. Softw. Eng. Methodol.* 24, 3, Article 19 (May 2015), 32 pages. https://doi.org/10.1145/2729976

[76] Finn Myrstad and Ingvar Tjøstheim. 2021. Out of Control: How Consumers Are Exploited by the Online Advertising Industry.

[77] Duc Cuong Nguyen, Erik Derr, Michael Backes, and Sven Bugiel. 2020. Up2Dep: Android Tool Support to Fix Insecure Code Dependencies. In *Proceedings of the 36th Annual Computer Security Applications Conference* (Austin, USA) *(ACSAC '20)*. Association for Computing Machinery, New York, NY, USA, 263–276. https://doi.org/10.1145/3427228.3427658

[78] Duc Cuong Nguyen, Dominik Wermke, Yasemin Acar, Michael Backes, Charles Weir, and Sascha Fahl. 2017. A Stitch in Time: Supporting Android Developers in WritingSecure Code. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (Dallas, Texas, USA) *(CCS '17)*. Association for Computing Machinery, New York, NY, USA, 1065–1077. https://doi.org/10.1145/3133956.3133977

[79] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark Patterns after the GDPR: Scraping Consent Pop-ups and Demonstrating their Influence. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376321

[80] Office of the Australian Information Commissioner. 2019. Data breach preparation and response: A guide to managing data breaches in accordance with the Privacy Act 1988 (Cth). https://www.oaic.gov.au/_data/assets/pdf_file/0017/1691/data-breach-preparation-and-response.pdf

[81] Mariana Peixoto, Tony Gorschek, Daniel Méndez Fernández, Davide Fucci, and Carla Silva. 2024. A natural language-based method to specify privacy requirements: an evaluation with practitioners. *Requirements Engineering* 29 (07 2024), 1–23. https://doi.org/10.1007/s00766-024-00428-z

[82] Mariana Peixoto, Carla Silva, João Araújo, Tony Gorschek, Alexandre Vasconcelos, and Jéssyka Vilela. 2022. Evaluating a privacy requirements specification method by using a mixed-method approach: results and lessons learned. *Requir. Eng.* 28, 2 (Sept. 2022), 229–255. https://doi.org/10.1007/s00766-022-00388-2

[83] Mariana Peixoto, Carla Silva, Ricarth Lima, João Araújo, Tony Gorschek, and Jean Silva. 2019. PCM Tool: Privacy Requirements Specification in Agile Software Development. https://doi.org/10.5753/cbsoft_estendido.2019.7666

[84] Charith Perera, Mahmoud Barhamgi, Arosha K. Bandara, Muhammad Ajmal, Blaine Price, and Bashar Nuseibeh. 2020. Designing privacy-aware internet of things applications. *Information Sciences* 512 (2020), 238–257. https://doi.org/10.1016/j.ins.2019.09.061

[85] Charith Perera, Ciaran McCormick, Arosha K. Bandara, Blaine A. Price, and Bashar Nuseibeh. 2016. Privacy-by-Design Framework for Assessing Internet of Things Applications and Platforms. In *Proceedings of the 6th International Conference on the Internet of Things* (Stuttgart, Germany) *(IoT '16)*. Association for Computing Machinery, New York, NY, USA, 83–92. https://doi.org/10.1145/2991561.2991566

[86] Mark Petticrew and Helen Roberts. 2006. *Systematic Reviews in the Social Sciences: A Practical Guide*. Vol. 11. https://doi.org/10.1002/9780470754887

[87] Reddr. 2023. LibScout: Detect and Analyze Third-Party Libraries in Android Applications. https://github.com/reddr/LibScout Accessed: 2024-09-25.

[88] Lornel Rivas, María Pérez, Luis E. Mendoza, and Anna Grimán. 2008. Towards a Selection Model for Software Engineering Tools in Small and Medium Enterprises (SMEs). In *2008 The Third International Conference on Software Engineering Advances*. 264–269. https://doi.org/10.1109/ICSEA.2008.51

[89] Hanne Rygge and Audun Jøsang. 2018. Threat Poker: Solving Security and Privacy Threats in Agile Software Development. In *Secure IT Systems*, Nils Gruschka (Ed.). Springer International Publishing, Cham, 468–483.

[90] Awanthika Senarath and Nalin A. G. Arachchilage. 2018. Why developers cannot embed privacy into software systems? An empirical investigation. In *Proceedings of the 22nd International Conference on Evaluation and Assessment in Software Engineering 2018* (Christchurch, New Zealand) *(EASE '18)*. Association for Computing Machinery, New York, NY, USA, 211–216. https://doi.org/10.1145/3210459.3210484

[91] Stuart S. Shapiro. 2010. Privacy by design: moving from art to practice. *Commun. ACM* 53, 6 (June 2010), 27–29. https://doi.org/10.1145/1743546.1743559

[92] Daniel J. Solove. 2007. 'I've Got Nothing to Hide' and Other Misunderstandings of Privacy. *San Diego Law Review* 44 (2007), 745. https://ssrn.com/abstract=998565 GWU Law School Public Law Research Paper No. 289.

[93] Stack Overflow Community. 2025. Stack Overflow. https://stackoverflow.com/ Accessed: 2025-03-11.

[94] Mohammad Tahaei, Alisa Frik, and Kami Vaniea. 2021. Privacy Champions in Software Teams: Understanding Their Motivations, Strategies, and Challenges. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 693, 15 pages. https://doi.org/10.1145/3411764.3445768

[95] Ying Tang, Morgan L. Brockman, and Sameer Patil. 2021. Promoting Privacy Considerations in Real-World Projects in Capstone Courses with Ideation Cards. *ACM Trans. Comput. Educ.* 21, 4, Article 34 (Oct. 2021), 28 pages. https://doi.org/10.1145/3458038

[96] Miguel Ehecatl Trujillo, Gabriel García-Mireles, Erick Orlando Matla Cruz, and Mario Piattini. 2019. A Systematic Mapping Study on Privacy by Design in Software Engineering. *CLEI Electronic Journal* 22 (04 2019). https://doi.org/10.19153/cleiej.22.1.4

[97] Charles Weir, Ingolf Becker, and Lynne Blair. 2022. Incorporating software security: using developer workshops to engage product managers. *Empirical Softw. Engg.* 28, 2 (Dec. 2022), 33 pages. https://doi.org/10.1007/s10664-022-10252-0

[98] Claes Wohlin. 2014. Guidelines for snowballing in systematic literature studies and a replication in software engineering *(EASE '14)*. Association for Computing Machinery, New York, NY, USA, Article 38, 10 pages. https://doi.org/10.1145/2601248.2601268

[99] Kim Wuyts, Laurens Sion, and Wouter Joosen. 2020. LINDDUN GO: A Lightweight Approach to Privacy Threat Modeling. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. 302–309. https://doi.org/10.1109/EuroSPW51379.2020.00047

[100] Sebastian Zimmeck, Rafael Goldstein, and David Baraka. 2021. PrivacyFlash Pro: Automating Privacy Policy Generation for Mobile Apps. *Proceedings 2021 Network and Distributed System Security Symposium* (2021). https://api.semanticscholar.org/CorpusID:231878917