

“I need to learn better searching tactics for privacy policy laws.” Investigating Software Developers’ Behavior When Using Sources on Privacy Issues

Stefan Albert Horstmann
stefan-albert.horstmann@rub.de
Ruhr University Bochum
Bochum, Germany

Sandy Hong
sandy.hong@rub.de
Ruhr University Bochum
Bochum, Germany

Maziar Niazian
maziar.niazian@rub.de
Ruhr University Bochum
Bochum, Germany

Cristiana Santos
c.teixeirasantos@uu.nl
Utrecht University
Utrecht, The Netherlands

Alena Naiakshina
alena.naiakshina@uni-koeln.de
University of Cologne
Cologne, Germany

Abstract

Since the introduction of the European General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), software developers increasingly have to make privacy-related decisions during system design and implementation. However, past research showed that they often lack legal expertise and struggle with privacy-compliant development. To shed light on how effective current information sources are in supporting them with privacy-sensitive implementation, we conducted a qualitative study with 30 developers. Participants were presented with a privacy-sensitive scenario and asked to identify privacy issues and suggest measures using their knowledge, online resources, and an AI assistant. We observed developers’ decision-making in think-aloud sessions and discussed it in follow-up interviews. We found that participants struggled with all three sources: personal knowledge was insufficient, web content was often too complex, and while AI assistants provided clear and user-tailored responses, they lacked contextual relevance and failed to identify scenario-specific issues. Our study highlights major shortcomings in existing support for privacy-related development tasks. Based on our findings, we discuss the need for more accessible, understandable, and actionable privacy resources for developers.

CCS Concepts

• Security and privacy → Usability in security and privacy.

Keywords

Privacy, Software Development, Developer Study, Information Sources

ACM Reference Format:

Stefan Albert Horstmann, Sandy Hong, Maziar Niazian, Cristiana Santos, and Alena Naiakshina. 2026. “I need to learn better searching tactics for privacy policy laws.” Investigating Software Developers’ Behavior When Using Sources on Privacy Issues. In *2026 IEEE/ACM 48th International Conference*

on Software Engineering (ICSE ’26), April 12–18, 2026, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3744916.3773150>

1 Introduction

With the introduction of the European General Data Protection Regulation (GDPR) [84] and the California Consumer Privacy Act (CCPA) [76], potential fines [68], customer expectations [39], and the risk of reputational damage from privacy violations [19] have motivated software companies to ensure privacy-compliant development [39]. For example, since 2018, over 6 billion Euros have been given out in fines for GDPR violations alone [34]. As a result, software developers are often tasked with implementing privacy requirements and making privacy-relevant decisions [39]. However, past research found that developers struggle to meet legal requirements, as they often lack privacy expertise [39, 40]. While some companies are able to hire legal experts or include privacy champions [79] to support development teams, communication is often limited and challenging [39, 40, 79, 81]. Consequently, developers frequently have to seek information online independently [80, 82].

With the development of Large-Language Models (LLMs), a novel source to access online legal information has been introduced into the software development process. This shift is reflected in trends, such as Stack Overflow [75] reporting a decline in user activity, which they attributed to the rise of AI systems [17]. The legal information presented in LLMs has the potential to directly impact present-day developers’ privacy decisions. Incorrect information present in such sources will thus have significant consequences with the usage of these new technologies [26, 44, 49, 53, 83]. Therefore, it is essential that online sources, including LLMs, provide correct and helpful legal information for current and potential future usage. To support developers addressing privacy issues during software development, we investigated how they currently seek and use privacy-related information. Further, we examined the challenges they face with existing sources, namely their own knowledge, online sources, and AI assistants, and explored their perspectives on different sources of information, aiming to understand the reasons behind the evolving use of online resources. Through letting the participants use all three sources for a limited time in our study, participants were able to directly compare the different sources and observe the advantages and disadvantages of each. Lastly, we



This work is licensed under a Creative Commons Attribution 4.0 International License. *ICSE ’26, Rio de Janeiro, Brazil*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2025-3/26/04
<https://doi.org/10.1145/3744916.3773150>

evaluated these sources to identify how the privacy information available to developers can be improved. Thus, we addressed the following research questions:

RQ1: What privacy issues do developers identify, and what measures do they propose using their knowledge, online resources, or AI assistance?

RQ2: How do developers use and perceive these sources when addressing privacy requirements?

In this work, we conducted interviews with 30 software development professionals. During think-aloud sessions, participants were presented with a privacy-sensitive scenario and asked to identify privacy issues and possible measures using their knowledge, online sources, and an AI assistant. Afterward, they were asked to reflect on and compare the different sources.

We found that developers prefer using AI assistants over conventional online sources for making privacy-related decisions, as they were perceived as effective and easy to use, despite a general mistrust in the accuracy of these tools. However, our analysis of the AI-generated responses revealed a key limitation: many privacy issues were not reliably identified. The AI assistants often provided vague or generic privacy information and struggled to detect context-specific issues or suggest concrete, scenario-relevant measures. Based on our findings, we provide recommendations on how online sources can improve their content and make the information more relevant to developers' projects and legal information-seeking.

2 Related Work

In this section, we summarize past research on software developers handling privacy requirements, developers using information sources during software development, and the introduction of LLMs into the software development process.

2.1 Developers and Legal/Privacy Requirements

Responsibility for privacy compliance is often shifted to technical roles, as they are the ones responsible for implementing the requirements [39, 77]. However, past research showed ongoing challenges for software developers in understanding and implementing privacy requirements [4, 5, 39, 40, 72]. For example, research by Senarath et al. [72] and Alhazmi et al. [4, 5] emphasized that developers lack formal knowledge of privacy principles such as Privacy by Design, Data Minimization, and Privacy Impact Assessments. Further, Horstmann et al. [39] found a communication gap between developers and privacy experts through interviews. In a follow-up lab study [40], they explored whether access to a privacy expert would support implementation but found that developers were reluctant to make use of it. Kilhoffer et al. [47] also examined how privacy engineers (PE) act as intermediaries between legal and technical teams in large organizations. Their qualitative study found that while PEs help translate legal requirements into actionable technical guidelines, their role is often constrained by organizational priorities, particularly a focus on achieving baseline compliance rather than promoting privacy as a value. Complementing these findings, Khedkar et al. [46] conducted an empirical analysis of Android applications, revealing widespread inconsistencies between what developers report in privacy documentation and what is implemented in source code. By analyzing 70 Android apps, their study

uncovered both over-reporting and under-reporting of privacy-related data collection in the Google Play Store's [37] data safety section. The study highlights that developers' misunderstandings are compounded by unclear regulatory definitions, insufficient tool support, and a lack of standardized guidance.

2.2 Developers and Information Sources

In the absence of clear internal guidelines or accessible legal expertise, especially in smaller companies with limited resources, developers are often required to seek external information to resolve the privacy issues associated with software development, including forums, blogs, and official documentation [81]. Especially, Stack Overflow [75] serves as a key platform where developers seek assistance on a range of privacy-related topics, including privacy policies, access control, and platform version changes [82]. Questions on the platform range from conceptual inquiries in early development phases to direct requests for implementation instructions and help with unexpected errors. While many accepted answers point to official documentation, over half are provided by community members without external references, highlighting a strong reliance on shared peer experience [82]. Other forums, such as Reddit [67], are also used to seek and share information related to privacy regulations. Parsons et al. [63] analyzed threads from developer-focused subreddits, identifying common themes and tracking sentiment around GDPR and CCPA. Their findings reveal a strong focus on permission-related questions and regulatory compliance, particularly GDPR. However, the support these traditional sources offer is often inadequate. Previous research identified a lack of high-quality resources, such as educational videos or example codes, that specifically focus on integrating privacy features [80, 81]. Even when resources are available, developers find it hard to navigate and apply, for example, when development platforms mandate the inclusion of a privacy policy but fail to provide guidance on how to create it. Further, in favor of functionality or revenue, software development tools often marginalize privacy considerations [81]. This can manifest itself in "dark patterns" in developer interfaces that discourage privacy-friendly options.

2.3 LLMs in Software Development

LLMs have emerged as an innovative and effective source and have quickly incorporated themselves into the daily workflows of software professionals [47]. While prior research mainly focused on code generation and completion [27, 93, 94], some engineers are already using tools such as ChatGPT [59] and GitHub Copilot [35], with promising applications for fine-tuned LLMs that can process privacy documents and automatically extract technical and legal requirements [47], suggesting that the capabilities of these technologies can go beyond simple code generation and completion. Studies showed that developers generally trust code completion tools more than generation tools, as they perceive the suggestions as more precise and think that they are derived from reliable and verified API documentation [58]. Still, there is a tendency of developers to "blindly trust" AI-generated output without adequate review or testing, which could negatively impact overall software security [49]. In fact, previous research revealed challenges regarding

Table 1: Study Parameters.

Study Type	Online
Study Task	Think Aloud & Interview
Study Language	English
Length	mean: 71.79 min (md: 70.12, σ : 9.82)
Recruitment Channel	Upwork.com
Recruitment Duration	1 month
Participants	Software Professionals (n = 30)
Compensation	\$60 per participant

the quality and security of the generated code [41, 64, 65]. Providing more context or adjusting model parameters can lead to more secure results, while uncritically accepting suggestions can increase security vulnerabilities, underscoring that the tool's effectiveness depends on the users' skill in prompting and critically reviewing the results [65]. Additionally, using third-party LLMs introduces other privacy risks. Developers are particularly concerned about possible leakage of sensitive data and intellectual property since the information could be used for training future models, or be exposed in a data breach, leading to developers policing themselves by refraining from inputting actual code [49]. Further, developers' perceptions are often tied to security and privacy concerns regarding risks associated with closed-source models and external models [51]. While existing research provides valuable insights into the security implications of AI-assisted code [9, 69], a significant gap exists regarding its impact on privacy compliance, as *"the legal status of code produced by LLMs is an open question"* [69].

In comparison to prior research, we extended the current field by examining the state of privacy information through (i) evaluating information sources based on both the results they deliver and their usability, (ii) understanding developer behavior and identifying the challenges associated with each type of source, and (iii) shedding light on the reasons behind the ongoing shift from online sources to AI assistants as a primary source of information.

3 Research Method

We conducted and recorded semi-structured interviews with 30 developers. Participants were first sent a brief pre-survey to collect demographic information. After completing the survey, they could schedule an appointment using the calendar application Calendly [20]. The interviews were conducted via Zoom [23]. As part of the interview, we used a think-aloud approach and asked participants to find privacy issues and measures for a fitness app scenario. Finally, we conducted a post-survey on their experience with privacy and the different sources. Following best practices [73], the parameters of our study are summarized in Table 1.

3.1 Scenario

3.1.1 Description. Participants were provided a scenario describing the lifecycle of personal data within a hypothetical fitness application. The fitness app scenario was chosen based on related work [40, 48], as it contains a realistic scenario of an application handling sensitive information like health data. Similar to [40], we modeled our scenario based on a real fitness application (e.g., Strava [78] or Nike Training Club [57]). The app *collects* user data,

Table 2: Legal Baseline of Scenario.

Privacy Issue	Suggested Measures
Data Collection	DPIA, Explicit Consent, Encryption
Retention Period	Anonymization after Period
Sharing: Location Data	Anonymizing Data
Sharing: Third Party Hosting Service	Data Processing Agreement, Encryption
Processing: Ads	Explicit Consent, Minimization
Processing: Profiling	Anonymization
Repurposing: Recommendations	Informed Consent
Repurposing: Improvement	Informed Consent
Repurposing: Ads	Informed Consent
Transfer to non-EU countries	Local Storage

including user workout information and, optionally, additional sensitive health data such as blood pressure. This data is then *stored* within the system. The application *processes* the data to provide personalized fitness recommendations, suggest sports products from the company's online store, and display personalized ads to users. Participants were also informed about the *sharing* and *transfer* of data, including the use of AWS for hosting the application, and the fact that it would be available worldwide for all common mobile operating systems. A detailed scenario description can be found in the supplementary material (see Section 7). We introduced privacy issues into the application, which were based on related work [39] or real-life GDPR violations [24].

3.1.2 Legal Background. We base the background of the fitness app under the GDPR, since it applies extraterritorially to all organizations that offer goods or services to, or monitor the online behavior of, individuals within the EU regardless of where it is established (Art. 3(2)), and considering that the GDPR has influenced the adoption of European-style data protection laws globally (known as the Brussels effect) [11, 29]. Considering these data cycle descriptions, jointly with a legal scholar who had 5 years of professional experience as a data protection lawyer, we identified segments of the app description that contained legally relevant information which, if left unaddressed, could result in potential GDPR violations. These passages were then mapped to potential measures to ensure compliance. Privacy issues and examples of possible measures are depicted in Table 2. The privacy issues observed in the scenario were: (i) excessive data collection not necessary for the application, (ii) missing a data retention period, (iii) sharing of location data, (iv) sharing data with the hosting provider, (v) processing personal data for ads and (vi) profiling, using data for purposes different from their initial purpose, namely (vii) recommendations, (viii) improvements of the application, (ix) advertising, and lastly, (x) transferring data outside of the EU. Some possible measures to take were to conduct a Data Processing Impact Assessment (DPIA), obtain consent from the data subjects, encrypt data, anonymize personal data, enter into Data Processing Agreements (DPA) with third parties, employ data minimization, and store data inside the EU.

3.2 Study Procedure

3.2.1 Previous Privacy Knowledge. At the beginning of the interviews, we probed our participants' knowledge of privacy by asking about its role in their professional lives, the regulations and concepts they were familiar with or had applied, and the sources they had previously consulted. To establish a baseline across participants, we provided a definition of privacy based on data protection and data privacy, following [61, 62], before the think-aloud sessions (see supplementary material in Section 7).

3.2.2 Think-Aloud Sessions. We asked all participants to address privacy issues in a provided scenario (see Section 3.1). To avoid priming effects [21], all participants first relied solely on their own knowledge. To address a potential learning effect [38], participants were randomly assigned to use either online sources or an AI assistant as their second source. Afterwards, they switched to the remaining source, resulting in a total of three sessions, using every source once. Based on our piloting results (see Section 3.3), we opted for 10-minute sessions per source, resulting in a total of 30 minutes:

- (1) **Own Knowledge:** Participants were first asked to explore privacy issues in the given scenario without any help. This allowed us to examine if and how they would address privacy issues based on their knowledge and experience before consulting external sources or seeking additional help (see Section 2.1);
- (2) **Online Sources:** Past research showed that developers often rely on online search engines, Stack Overflow, and official documentation as sources of information [2] (see Section 2.2). Therefore, our participants were asked to use the Web to address the given scenario. They could use a web browser installed on their machine and access any online website they preferred. If they attempted to access an AI assistant, we asked them not to use one during this think-aloud session;
- (3) **AI-Assistant:** We included AI assistants as an additional information source, as past research suggests they are increasingly replacing traditional online sources for developers [17, 49] (see Section 2.3). We chose ChatGPT, as it is one of the most widely used AI assistants [42]. Participants were asked to open ChatGPT in a browser to use as they liked. We logged them into a paid ChatGPT-4 account using our own credentials to ensure participants' anonymity and configured the account not to retain or learn from past conversations. All data was deleted after each participant session.

We refer to participants who used online sources first as *O*, and those who used an AI assistant first as *A*. Participants were asked to document all identified issues and proposed privacy measures in an output document. Additionally, after each session, we asked participants open-ended questions adapted from the System Usability Scale (SUS) [16]. Since our study did not aim for a quantitative analysis, we chose an open-ended format to gather more detailed and contextual feedback. For example, instead of asking participants to rate their agreement with the statement “*I found the system unnecessarily complex.*” we asked, “*What are your thoughts on the complexity of the online sources on privacy?*” This approach helped us understand how participants perceived the usability of each

source, allowing them to explain their reasoning and describe specific challenges. We also asked participants whether they trusted the respective sources and how confident they were in the information they found. After the three think-aloud sessions, we asked participants to share their opinions on the three sources and compare them.

3.2.3 Post-Survey. After the interviews, participants were sent a short survey concerning their experience with privacy and their perception of the different sources for privacy in the Zoom [23] chat. The full interview guide and survey can be found in the supplementary material (see Section 7).

3.3 Piloting

We initially designed three different scenarios and planned for participants to solve each one using a different source. The order of scenarios and the assigned information sources were randomized, ensuring that each scenario would be addressed with each source across participants. We piloted this initial study design with three Computer Science student assistants from our research group. However, we noticed that the time participants spent on the tasks varied strongly based on the source used. Further, participants were influenced by the past scenario, making a comparison between the groups difficult. Thus, we created a single scenario and piloted the adapted setup with another researcher from the field of IT security, who had comparable programming experience with the student assistant, to test the setup, framing, questions, and study time.

3.4 Positionally Statement

In line with best practices [60, 70], we include a positionality statement to describe the backgrounds of the researchers involved in this work. Researchers R1, R2, R4, and R5 have background in usable privacy and security research, whereas R3 has over five years of professional experience as a data protection lawyer.

3.5 Analysis

We collected the following data: (i) audio transcripts of the interviews, (ii) search terms used and websites visited by participants when using online sources, (iii) participants' chat with the AI assistant containing prompts and replies, (iv) privacy issues and the measures that were documented by the participants in an output document, and (v) survey results.

3.5.1 Qualitative Analysis. The analysis of i) interviews, ii) used search terms and visited websites, and iii) AI chat was conducted by R1 and R2 using thematic analysis according to Braun and Clarke [12], using MAXQDA [56]. First, both researchers coded two interviews together, using the results of the legal background as a first basis for coding the privacy issues (see Section 3.1.2), but deriving the rest of the codes inductively through the data. Both researchers then independently coded two different interviews, extending the codebook where needed. Afterwards, they met and compared the codebooks, discussed changes, and merged similar codes. Each researcher then coded all interviews individually, meeting regularly to compare and discuss the codebook, ordering codes into themes. For example, codes related to online sources being

Table 3: Demographics of the 30 participants.

Gender			
Man	24 (80%)	Woman	6 (20%)
Age [years]			
Min.	19	Max.	67
Mean (σ)	32.93 (± 9.09)	Median	31
Experience in current field [years]			
Min.	2	Max.	24
Mean (σ)	8.52 (± 5.91)	Median	6
Education			
Bachelor's degree	12 (40%)	Master's Degree	11 (36.67%)
High School Equivalent	2 (6.67%)	Some College	2 (6.67%)
Vocational Degree	1 (3.33%)	Some Grad School	1 (3.33%)
Less than high school	1 (3.33%)		
Current Employment Status			
Full time	14 (46.67%)	Part time	8 (26.67%)
Self-Employed	7 (23.33%)	Prefer not to say	1 (3.33%)
Largest Company Size (Employees)			
Min.	10	Max.	300 000
Mean (σ)	18760.25 (± 64600)	Median	135
Country of Residence			
US: 6	CA: 5	FR: 3	IN: 2
UK: 2	Other Europe: 8	Other Asia: 2	Africa: 2

outdated, untrustworthy, and lacking information were combined into the theme that online sources lack in quality.

The same process was repeated to code AI responses, as well as the websites and search terms. Lastly, similar to related work [39, 49, 50], all conflicts were resolved through discussion among researchers, for example, when a coder overlooked a statement in the text or when clarifying distinctions between data minimization and purpose limitation in cases where participants' descriptions were ambiguous. Thus, we achieved full agreement. We reached saturation after 25 participants, but continued to recruit five additional participants to ensure completeness. We did not calculate the inter-coder reliability, as it is neither recommended by Braun and Clark nor by other research [13, 14, 18].

3.5.2 Output Documents. R1, R2, and R3 conducted the analysis of the output documents where participants noted privacy issues and measures they found regarding the app scenario. For each output document, we coded the following:

- privacy issues noted by participants according to Table 2,
- measures proposed by participants to address the observed privacy issues, and
- other suggestions made by the participants not directly connected to any privacy issues of the scenario.

The coding was performed on the documents after each step (Own Knowledge, Online Sources, AI), i.e., in the state they were after participants had finished using one of the information sources. By examining the new codes added after each step, we could assess the effectiveness of each source, each new code representing an issue or measure identified through that specific source.

3.6 Recruitment and Demographics

We recruited participants through Upwork [89], as past research [45] recommended it as a suitable recruitment platform for developer studies with practical tasks. To address recruitment bias [25], we framed the invitation as a research study on information sources used in software development, but did not specifically mention privacy and data protection. Participants were required to (1) be

Table 4: Overview of the 30 participants.

	Current Job Title	Ctry.	Company Employees	Privacy Tasks	Data Privacy Knowledge
O1	Data Scientist	IN	1-99	All	Very
O2	Senior Web Dev	PK	1-99	Half	Moderately
O3	Lead Software	US	1000+	Some	Slightly
O4	Full-Stack Dev	LV	1-99	Half	Slightly
O5	Software Dev	TR	100-999	Some	Not at all
O6	Software Architect	EE	1000+	Some	Moderately
O7	Full-Stack Dev	US	100-999	Some	Slightly
O8	Software Dev	BA	1-99	Some	Slightly
O9	Software Dev	CA	1-99	Some	Slightly
O10	Senior Dev	VN	1-99	Some	Not at all
O11	Not IT related	CA	100-999	Some	Moderately
O12	Software Dev	US	1-99	Most	Very
O13	Software Dev	CA	1000+	Some	Moderately
O14	Software Dev	GB	1000+	All	Moderately
O15	Software Dev	FR	1-99	Some	Moderately
A1	Automation Eng	LT	1-99	Some	Moderately
A2	SecOps Manager	EG	1000+	Most	Moderately
A3	Full-Stack Dev	IN	1-99	Most	Moderately
A4	Software Dev	UA	100-999	Some	Slightly
A5	Software Dev	ET	100-999	Most	Very
A6	IT Cons.	US	1000+	Half	Moderately
A7	Software Eng	CA	100-999	Some	Not at all
A8	Not IT related	US	100-999	Some	Slightly
A9	Senior Dev	CA	100-999	None	Slightly
A10	CTO	GB	1-99	Most	Very
A11	Full-Stack Dev	FR	1-99	Some	Slightly
A12	Web Dev	PT	100-999	Some	Slightly
A13	Full-Stack Dev	FR	100-999	Some	Slightly
A14	IT Cons.	US	100-999	Some	Not at all
A15	Software Devg	ES	1000+	Some	Slightly

18 years or older, (2) have professionally worked in software development after 2018 to ensure they worked while the GDPR and the CCPA were in force, and (3) describe their last 6 years of software development experience to ensure they were qualified for the study. Participants were paid \$60 for the 60-minute study, as aligned with recommendations from [73]. Overall, Upwork invited 1373 freelancers to apply to the job post. Of those, 378 submitted an application, and Upwork [89] suggested 292 of them. We individually checked all recommended participants' profiles to verify their experience. We contacted 50 freelancers to confirm their continued interest in participating; 30 accepted the invitation. Of the remaining freelancers, two said they were busy with other projects, one was traveling for an extended period, and the remaining 17 did not respond. Participants' demographics can be found in aggregated form in Table 3, with additional information on each participant in Table 4. Our participants represented 17 different countries: 13 were from Europe, 11 from North America, four from Asia, and two from Africa. Of the 30 participants, 24 identified as men and six as women. Participants had, on average, 8.52 (sd: 5.91) years of experience in their current working domain.

3.7 Ethics

Our university did not have an Institutional Review Board (IRB) at the time of this study. However, we did comply with the GDPR obligations and followed the principles outlined in the Menlo report [10]. We provided participants with a consent form, which they were asked to download and review before giving their consent to participate. Prior to the interview, we reiterated information about

Table 5: Privacy issues found by participants using different sources

Privacy Issues	Own Knowledge	Online Sources	AI Assistant
Data Collection	●	○	●
Retention Period	○	○	●
Sharing: Location Data	○	○	○
Sharing: Third Party Hosting	●	○	●
Processing: Ads	●	○	●
Processing: Profiling	○	○	○
Repurposing: Recommendations	○	○	○
Repurposing: Improvement	○	○	○
Transfer to non-EU countries	○	○	○

○ represents the percentage of participants who added the issue to the output document while using the respective source

their rights, data collection, and data processing, and gave them the opportunity to ask questions or raise any concerns. Participants could withdraw from the study at any time without consequences. All participants also gave verbal consent at the beginning of the interview recording. The collected data was anonymized after the study, and all recordings were deleted following transcription and verification of data completeness (e.g., ensuring no missing browser history or ChatGPT responses).

3.8 Threats to Validity

Our study has several limitations that should be considered when interpreting the results. First, our participants were freelancers recruited through Upwork [89], and therefore may not fully represent the global developer population. However, we ensured that each participant had prior experience working in a company setting. Second, we focus on the GDPR as the primary framework for evaluating participants' solutions. We chose the GDPR because of its broad applicability to any company handling the data of EU citizens and its global influence through the so-called Brussels Effect, which has shaped the adoption of European-style data protection laws worldwide (see Section 3.1.2). While focusing on other legislation might have revealed additional or different issues, we expect that the GDPR allowed us to identify the core set of privacy issues likely to appear across regulations and countries. While we randomized the order of using the AI assistant and online sources, a learning bias may still have influenced the results.

4 Results

In this section, we present the analysis of the interview and think-

Table 6: Measures found by participants using different sources

Measures	Own Knowledge	Online Sources	AI Assistant
Transparency	○	○	○
Consent	○	○	○
Encryption	○	○	○
Legal Compliance	○	○	○
User Deletion	○	○	○
Data Minimization	○	○	○
Anonymization	○	○	○
Purpose Limitation	○	○	○
Two Factor Authentication	○	○	○
Age verification	○	○	○
Auditing	○	○	○

○ represents the percentage of participants who added the measure to the output document while using the respective source

aloud sessions, participants' output documents, their browsing history (*Browser History*) when using online sources, and their interactions with the AI assistant (*AI Chat*). Privacy issues and measures proposed are presented in Section 4.1 and the use of privacy sources is reported in Section 4.2. Throughout the section, we differentiate between *reported data*, i.e., thoughts and opinions of the participants, and *observed data*, i.e., sources we saw participants access, and *results* from the analysis of the AI interaction.

4.1 RQ1: Privacy Issues and Measures

We report the effectiveness of the three different sources based on the issues identified and measures proposed by participants in their output documents in Section 4.1.1. We then describe privacy issues in Section 4.1.2 and measures in Section 4.1.3 documented in participants' output documents, as well as those mentioned during interviews and think-aloud sessions. While participants have identified issues and measures, they did not always include them in their written responses. This may suggest difficulties in relating the information from the source to their specific use case, or challenges in understanding it. Therefore, we distinguish between what was *verbally expressed* and what was *documented*. While certain issues arose directly from the scenario description (see Table 2), participants also discussed additional common concerns. Further, participants proposed a wide range of measures—some aligning with those outlined in the legal background, while others were independently suggested.

4.1.1 Effectiveness of Sources. The percentages of participants who identified a specific issue using a particular source are presented in

Table 5. Similarly, Table 6 reports the percentages of participants who suggested measures in their output documents while using each source. To evaluate the effectiveness of each privacy source, we compared the number of new codes that were generated in our analysis after participants used the respective source. Each code represents either a privacy issue or some form of privacy measure found and added by the participant (see Section 3.5.2). Participants performed similarly when applying their own knowledge as a first source, with a mean of 5.2 (min: 2, med: 5, max: 8, sd: 1.66) for the AI-First and 5.27 (min: 1, med: 5, max: 10, sd: 2.63) for the Online-First participants. Looking at the first source of information used, using the AI performed better compared to the online searchers: The AI-First group had 9.73 (min: 0, med: 6, max: 33, sd: 10.75) new codes on average, whereas the group using online first only gained 2.4 (min: 0, med: 3, max: 8, sd: 2.2) new codes on average. When looking at the third iteration of the output document, using the AI assistant again proves more effective: AI-First achieved an average of 1.2 (min: 0, med: 0, max: 6, sd: 1.93) new codes using online sources, whereas the Online-First group achieved 10 (min: 0, med: 4, max: 74, sd: 18.5) on average with the use of ChatGPT. Considering the results independent of their order of application, the use of online sources resulted in the fewest new codes, with 1.8 new codes on average (min: 0, med: 1.5, max: 8, sd: 2.12), followed by use of their own knowledge with a mean of 5.23 (min: 1, med: 5, max: 10, sd: 2.16). Using the AI assistants created the most, with 9.87 new codes on average (min: 0, med: 4.5, max: 74, sd: 14.84). Notably, participants who copied ChatGPT responses into the documents drastically increased the number of coded sections.

This data suggests that online sources require a significantly longer amount of time to use to find privacy issues and measures compared to the participants' own knowledge and AI ChatGPT. However, the number of issues and measures found does not necessarily mean all requirements or the required measures were found. This is especially a concern with the results obtained using AI assistants, as some participants simply copied ChatGPT's response into the output documents, meaning the issues and measures are highly dependent on the correctness of the AI's output.

4.1.2 Identified Privacy Issues. Participants were able to detect some of the *privacy issues* presented in the description of the application. All participants verbally mentioned the issue of data collection for the application, and 23 added it to the output document. *"Yeah, don't collect more than you need."* (O9) Twenty-six mentioned personal information, with 23 referencing health data, twenty citing activity data, and 16 discussing location. Twenty-seven participants noticed the sharing of data with third-party services as an issue: *"So there is a lot of user data being stored on AWS and a lot of it is (...) personal information, age, height, blood pressure, stuff like that."* (A5) Thirteen made note of it. Further, 15 participants documented targeted ads, which were discussed by 13: *"But for including advertising services, showing users personalized ads, that would require its own specific consent agreement."* (A14) Only nine participants noted data transfers as critical, even though nineteen participants discussed the risks associated with cross-border data flows: *"Okay, well, first of all, we are speaking about an old good problem of a US company operating across the globe and hence, EU data has to be stored in*

EU." (A15) Fifteen participants expressed concern about data sharing in general, even when no specific third-party was mentioned. *"We don't want to share the data publicly."* (O12) This concern was registered by five participants. Eight were also concerned about the indefinite storage of data described in the task description: *"How long will the data is retained by app owner, and what happened to that owner when you stop using app to delete their account?"* (O2) One participant raised the issue of repurposing of user data for ads and for recommendations in the app in their output document, which twenty-one participants discussed. Seven of the participants discussed the repurposing of data for improvement and optimization of the app as something to be considered, but only four further wrote it down in their output document. *"Can provide that user data instead of improving, motivating, and optimizing for the future project. Okay, we also should ask for this. Like, can we use the data to improve our services?"* (O7) Four participants discussed profiling and noted it down in the output document: *"So we are processing users' data to recommend a personalized user experience. Let's say it's a problem."* (O5) Eight participants were concerned about the sharing of location data:

"Yeah, so this is also another privacy issue is that your location because your location is different people have access to it and you can have some people, for example, like following you or knowing exactly where you where you run or you live." — O9

However, only two noted it down. Only one participant explicitly mentioned the sharing of health data, which was discussed by six, as an issue in the output document.

4.1.3 Suggested Measures. All participants verbally suggested measures aimed at strengthening user rights, i.e., *"essentially ask for consent"* (O9) or *"be transparent, clear."* (O12) Twenty-one suggested transparency, for example, through a privacy policy in the output document: *"App needs to have privacy policy."* (O5) Further, 21 of the participants added to the document the need to obtain user consent to process their data. Nineteen also suggested user data access, and to give users the option *"to correct or delete their data."* (O5) In this line, 13 participants wrote down about exercising the users' rights, e.g., the deletion of or access to their data. With 26 participants discussing them, 22 noted down additional security measures, often encryption of the data: *"Securely receive & store encrypted PII data when receiving information from the frontend application."* (O3) Purpose limitation was noted by 16 of the participants in the document as well, with 15 introducing some form of data minimization, which was discussed by 21. Eighteen suggested including a retention period, either deleting or anonymizing the data after the period was up, which 12 noted down.

4.1.4 Comparison of Participants. As the GDPR is also relevant for developers outside of the EU, as long as they are handling data of EU citizens, we also included participants from outside of the EU. Overall, we had 10 participants currently residing in the EU and the UK, where the legislation is based on the GDPR, and 20 outside. EU and UK participants performed slightly more effectively when exploring privacy issues and measures, suggesting 18.7 (min: 5, med: 14.5, max: 38, sd: 12.5) on average, compared to an average of 15.9 (min: 3, med: 12, max: 78, sd: 15.95) suggestions by the

non-EU and UK participants. During the interviews, the results were similar as well in almost all cases. When discussing laws, the GDPR was mentioned proportionally more often by EU and UK participants, with 3.5 mentions on average (min: 0, med: 2, max: 12, sd: 3.54), compared to an average of 2.75 (min: 0, med: 2, max: 10, sd: 2.95) mentions by non-EU participants. Instead, the non-EU participants mentioned a wider variety of different laws, such as HIPAA, CCPA, COPPA, and other regional laws. This issue was also observable in the search terms used when using online sources, as well as the chats with the AI assistant. While the data suggests that participants are more likely to be aware of their local legislation, the GDPR remained the most commonly mentioned legislation in every case. However, the qualitative nature of our work to find significant differences in other cases, such as which issues are found, or measures that were suggested by each group. Lastly, as participants reported their data protection knowledge, we compared the number of issues found and measures suggested by the participants. Participants reporting below-moderate data protection knowledge found on average 15.19 (min: 3, med: 13, max: 35, sd: 9.02) issues and measures, which was slightly lower than participants reporting average knowledge or higher, who achieved 18.71 (min: 3, med: 11.5, max: 78, sd: 19.57) on average.

4.2 RQ2: Using Privacy Sources

To evaluate the different sources used by participants, we analyzed their comments during think-aloud sessions as well as their responses to interview questions regarding source usability. Throughout the section, we differentiate between the participants' behavior we were able to observe when interacting with sources, and the thoughts and opinions the participants expressed. We present the results of using their own knowledge in Section 4.2.1, online sources in Section 4.2.2, and an AI assistant in Section 4.2.3.

4.2.1 Own Knowledge. Behavior: Participants read the scenario and noted down issues they saw in the task description. If they knew a privacy measure, they noted it down as well. Since they did not interact with any other source, the information mainly stems from what the participant mentioned during the think-aloud and interview sessions.

Explanations from the Interviews: Ten participants reported frequently relying on their own knowledge of privacy, while nine remarked they would not apply their personal understanding to privacy-related tasks: *"Stuff like this should not be covered by developers. It should be covered by actual people who understand this stuff."* (A7) Five participants explained that they might draw on their knowledge to identify privacy issues in the software they developed: *"I try to go through the steps that I would think or usually they occur as a privacy issue."* (O14) Four participants mentioned they had prior knowledge, e.g., through previous projects, online research, or documentation.

4.2.2 Online Sources. In this section, we examine the participants' usage of online sources, drawing from both their browser history and their self-reported accounts.

Behavior: We examined the online searches and sources that participants accessed. While 28 participants used Google [36] as their search engine, two participants opted to use DuckDuckGo [28].

Concerning search terms, 23 of the participants looked for specific issues related to the context of the fitness tracking app, e.g., *"Potential privacy issues for fitness tracking app in development"* (O13 Browser History), or more specific questions, for example for data storage *"how data storing is organized in fitness app"* (A4 Browser History) or general privacy measures they could take: *"privacy measurements when building a workout app with stores"* (O8 Browser History). Fifteen of the participants looked for answers to specific questions they had, e.g., on health data: *"can you store patient data in the cloud"* (A1 Browser History), or about advertisements: *"how to make sure personalized ads are compliant"* (O4 Browser History). Again, 15 participants specifically looked for information related to laws, for example, the GDPR: *"gdpr regulation requirements for healthcare"* (A5 Browser History), or similar guidelines: *"ftc data privacy"* (A7 Browser History). Four specifically looked for other fitness applications they knew, e.g., *"strava privacy issues"* (O9 Browser History). Three looked for information from specific online forum sources, for example *"privacy and data security health reddit"* (A15 Browser History) and *"processing location data of users stackoverflow"* (O5 Browser History). Six participants looked for privacy requirements from a specific country, mainly the United States: *"requirements in us for storing privacy data"* (A4 Browser History), and two participants also looked for information regarding requirements in Germany.

Additionally, we collected the websites that participants accessed during the study for privacy information. Twenty-six participants accessed blogposts from privacy-focused organizations, e.g., [66, 86], companies [7, 71, 90], or news organizations [33, 87]. Seventeen visited company sites and checked their privacy policy for information, e.g., Uber [88] or Whoop [91]. Only seven visited government websites, such as the Federal Trade Commission [85], the European Council [30], or the Information Commissioner's Office [43]. Seven also got information from community sources, such as Stack Overflow [75] or Wikipedia [92]. Six accessed academic papers during the online research, with three participants accessing a paper on privacy in fitness applications, which suggested blockchain technology as a remedy [3]. Other papers accessed by a participant were related to cloud storage in fitness applications [95], privacy management for fitness applications [1], and on developers' compliance processes for children's applications [6]. Lastly, one developer accessed information on the Spring framework for security information [15].

Explanations from the Interviews: Fifteen of the participants claimed they would use online sources for privacy decisions, and eight explicitly stated they would not use them. Six stated they would use such sources to find measures for identified issues, and three mentioned they would use them to find information on privacy issues. Eight said they would simply use Google to find information, with 11 explicitly stating they would look for projects similar to the scenario they are working on, and checking their privacy policy or statements and commitments the companies have made regarding privacy: *"I think I need to find another fitness app and to check their privacy, to find out maybe I missed something and they found it and already described in their privacy. It should be some famous app."* (A12) Eight would use online sources to research applicable laws: *"I guess I'm just going to Google about GDPR summary and maybe I can find something about that."* (O4) Two mentioned

they would use privacy policy generators, i.e., online tools that automatically generate privacy policies for software or websites (e.g., from Shopify [74]).

4.2.3 AI Assistants. Behavior: We examined how participants interacted with ChatGPT. Twenty-one of them provided the assistant with information about the scenario, either by copying sections from or summarizing the task description:

"I am building an application which tracks users fitness and will be saving this information with a cloud service provider (ie, AWS). The data gathered will be fitness related (ie, activity) and will also monitor and use the users information (ie, age, height, etc) to determine any recommendations. I need to ensure privacy policies are adhered to. Can you give me a high-level understanding of how I can achieve this?" — A9 AI Chat

Eleven asked for general privacy information, e.g., *"What are typical worries or concerns around privacy in fitness tracking applications?"* (O9 AI Chat) Fourteen asked for information on specific questions, e.g., on potential measures: *"will oauth2 technology help with data privacy?"* (O6 AI Chat); on issues they had noticed themselves: *"If the app includes advertisement services, showing users personalized ads, what are potential locations this data could become exploited? And what are the guidelines to prevent that?"* (A14 AI Chat); or on legal questions: *"how to comply with modern data privacy laws? Describe 5 main points"* (O6 AI Chat). Three participants employed prompt engineering to improve the results from the AI: *"You are expired software developer with knowledge about data privacy on US market. I'll ask questions next message"* (O11 AI Chat) One participant copied a paper from Alhajri et al. [3], found online, discussing blockchains as a measure for privacy risks in fitness apps, and asked the assistant to summarize it. Consequently, this was the only time when the assistant started discussing measures using blockchain: *"**Blockchain Application**": Using **smart contracts on a blockchain** could automate these consent grants and revocations, ensuring that changes are transparent, auditable, and enforceable."* (O13 AI Chat), showing that the users' prompts can highly influence the AI assistants' output. One participant asked if they could trust the AI assistant with privacy questions, to which it replied: *"Yes, you can trust these suggestions regarding data privacy. They are based on widely accepted industry standards, best practices, and regulatory frameworks for data security and privacy."* (O16 AI Chat) Our analysis of the AI assistant's responses revealed that GDPR was the most commonly referenced law by the AI assistant, cited 128 times. This was followed by 91 mentions of CCPA, 53 of HIPAA, and 28 of COPPA.

Regarding *privacy issues* identified by the AI assistant, all responses addressed data collection, with twenty-seven specifically mentioning health data. Twenty-four responses referred to sensitive and personal data in general, and 19 also noted location data. Children's data appeared in 15 of the responses, while 11 mentioned personally identifiable information (PII) and 12 mentioned activity data. Biometric data was cited in five cases, and four responses included references to financial data. Fifteen of the replies specifically mentioned targeted ads: *"Provide options to opt-in or opt-out of data sharing and personalized ads."* (O5 AI Chat) Further, ten replies discussed profiling. With regard to data sharing, twenty-five responses

raised concerns about the sharing of data with third parties, such as the hosting service. Similarly, data transfer across country borders was a concern in twenty-two replies. Six responses mentioned the sharing of location data, and three mentioned health data. Lastly, the majority of responses mentioned the general sharing of data as an issue. Concerning data usage, replies addressed the repurposing of data as an issue. Seventeen addressed advertising and 12 recommendations. Further, the missing retention period was criticized in 24 replies. Six replies also mentioned repurposing the data for improvements and optimization as a potential issue. The AI assistant also mentioned issues not directly covered by the app scenario. Twenty-eight discussed data leaks and unauthorized access.

With respect to the *measures* described by the assistant, all suggested measures related to user rights, and stressed the importance of consent and transparency. Next, all replies suggested giving their users control over their data. As the AI assistants had mentioned the issue of underage users before, 15 suggested additional safeguards, i.e., obtaining consent from an adult or age verification before the app could be used. Ten AI assistants suggested educating their users on the privacy settings. In 28 replies, the AI assistant strongly recommended complying with relevant legislation. Further, 28 stressed the need to employ data minimization. Introducing a retention period was suggested by 25. The assistant advised 16 times that companies should evaluate and vet third parties before engaging with them. Nine suggested Standard Contractual Clauses for data transfer outside of the EU. Eleven specifically recommended using privacy by design when building the software. Sixteen recommended ensuring purpose limitation, and 24 suggested access control mechanisms. Further, 25 replies recommended doing regular security or privacy audits. Seven responses emphasized the importance of the company staying up to date with legal developments and educating its team on privacy matters. Eleven also recommended contracting a lawyer or hiring a Data Protection Officer to ensure compliance.

Explanations from the Interviews: Participants reported asking the AI specific questions or giving it a summary of the task. Only three participants talked about employing some form of prompt engineering, for example, O9: *"Also, give a more technical response that would help a developer make appropriate changes"* (O9 AI Chat). Participants went through the list of issues or measures the AI provided: *"I'll just see like for each blue point, provide me a measures as it is in the prompt and if I decide to like rework some stuff, I'll just like rework it, feed it to him again."* (A13) Fourteen of the participants said they had not used AI for privacy before, mainly because those systems are new: *"Because it was more than two years ago and at that time ChatGPT wasn't available for the users probably."* (O7) However, nineteen said they have or would use these tools in real-life: *"Not yet, but I would use them if I have this task as of today."* (O7)

4.2.4 Comparison of Sources. A detailed overview of strengths and challenges reported by participants for each source can be found in the supplementary material (see Section 7). In the *post-survey*, participants rated the usefulness of the different sources. The full distribution of ratings can be found in Figure 1. AI assistants were considered the most useful information source, with 10 participants rating them as 'Extremely Useful' and twelve as 'Very Useful.' Online sources followed, receiving 4 'Extremely Useful' and 16 'Very

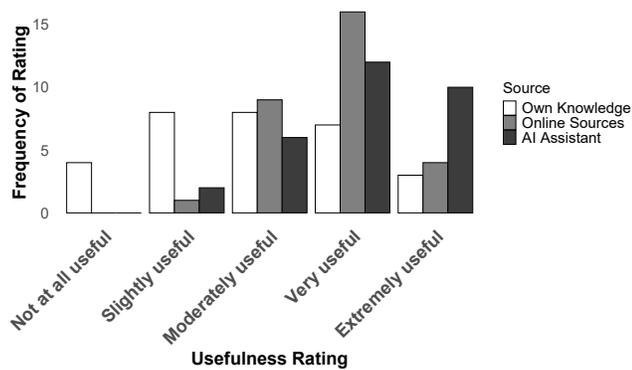


Figure 1: Participants' Rating of the Three Sources

Useful' ratings. Participants' own knowledge was rated the least useful, with only three marking it as 'Extremely Useful' and seven as 'Very Useful.'

In the *interviews*, 14 of the participants explicitly said they preferred using AI assistants for privacy research: "Well, I mean, it's not 100 percent easy, but compared in contrast to Google and searching myself. It's much easier than Google." (O10) Eight found AI assistants to perform significantly better than online sources:

"I think I got more done with ChatGPT, and I was generally more confident with the solutions from ChatGPT in this case. Online, I was really lost, and I did find something, and I don't know if it's even relevant." — A1

Six felt AI assistants and online sources performed similarly:

"This is kind of going down the same road that ChatGPT went. Where's recommending superiority and privacy and whatever. This one, see like this one came exactly out of ChatGPT. Well, ChatGPT stole this probably from this page." — A9

One participant said that AI assistants have started to replace expert advice: "Yeah, I think before, we relied on experts more." (A1) Seven participants preferred online sources, mainly as they saw it as the most trusted source:

"It would take you more work, I suppose, to search it through general Googling and that sort of thing or whatever, but you could find probably a lot more information, a lot more trusted information and just doing a little extra work." — O12

Four participants believed using their own knowledge for such tasks would yield good results, but noted requiring a much experience: "If the task working with this privacy policies concepts is something you do a lot frequently, you can rely on your experience." (O7)

Concerning the use of the sources, 17 participants imagined using them together:

"I suggest using the three of them because it's ideal to know how the AI will respond to you and what knowledge you had before. And like what you searched online or like Googled or something, then you can collect all this information and bring the best output of what you desired to do." — A5

5 Discussion

While our study setup had our participants use each source for a relatively short time, our research does offer fewer insights into single sources; however, it does allow insights to directly compare the participants' perception and use of each source. During the interview, participants often mentioned using the three sources in combination with each other. Thus, our setup is similar to how developers would employ these sources at work.

Overall, developers preferred using AI assistants over relying on online sources or their own knowledge. Although participants acknowledged that some understanding of privacy was required to use both AI assistants and online sources effectively, they generally did not expect to complete privacy-related tasks based solely on their own knowledge. This indicates a strong reliance on external support. When using online sources, developers encountered several challenges in locating and applying relevant information. They frequently found blog posts that addressed isolated issues and were uncertain whether the content was applicable to their specific context. In addition, they considered many sources difficult to interpret and lacking the depth needed to comprehensively cover privacy requirements. By contrast, the adoption of AI assistants appeared to be primarily driven by efficiency. Developers valued their speed, structured responses, and ease of use. Despite persistent concerns regarding the trustworthiness and accuracy of AI-generated content, these perceived advantages made AI assistants the preferred choice for many participants.

Privacy issues. While participants reliably identified certain privacy issues, particularly those related to data collection and targeted advertising, many concerns remained unaddressed, even with the aid of online resources and an AI assistant. Issues involving data repurposing and sharing of location data were especially difficult for participants to detect. In particular, repurposing entails knowledge of the GDPR purpose limitation principle (Art. 5(1)(b)) and the need for another legal basis in case further processing is incompatible with the initial purpose for which data was collected [8]. When developers are unable to identify such complex issues, this can lead to unlawful processing and potential breaches. These issues could not easily be resolved by the sources tested in this study, as online sources were often limited to blog posts and online forums, discussing a few concerns, but were not able to accurately cover privacy issues related to a concrete case, as confirmed in related work [52]. Similarly, AI assistants tended to offer generic best practices and rarely identified scenario-specific issues unless explicitly prompted by participants, lacking contextual and legal relevance.

Measures. While participants were able to recognize some privacy issues and propose standard mitigation measures based on their own knowledge, developers cannot be expected to act as legal experts. Participants' proposed measures were often limited to general best practices and rarely mapped directly to specific legal issues, as outlined in Section 3.1.2. Moreover, the lack of sufficient documentation, often resulting from reliance on informal or unverifiable sources, can hinder the effective implementation of privacy by design principles, as required under Art. 25 of the GDPR. It also weakens an organization's ability to demonstrate compliance, a core component of the accountability principle under Art 5(2).

Our work shows similarities with related work on information sources for software development and security sources; however, we observe some factors unique to privacy. Different from related work on Stack Overflow [75], where the main issues concerned the quality of the provided information [2, 31, 32], our participants had difficulties understanding the information they found and applying it to their own projects. Regarding privacy advice on Stack Overflow, past research showed that developers often rely on online sources for a variety of privacy questions, with many of the answers on Stack Overflow lacking links to official sources, making it challenging to evaluate the quality of replies without any additional knowledge [82]. Our results confirm findings from research on privacy during the software development process, showing privacy as a challenging topic for software developers [39, 40]. Similar to related work [49, 55], our research showed developers favoring the use of AI assistants for their convenience and effectiveness, even though participants reported some degree of mistrust regarding the correctness of the assistants.

Our results suggest that the current state of all three privacy sources does not sufficiently equip developers to independently identify or address significant privacy concerns, making professional legal support necessary, although such support remains infrequently available [39]. In the absence of accessible and tailored legal guidance, and given that direct engagement with legal professionals is both rare and sometimes perceived as adversarial [39, 40], developers can be expected to continue relying on alternative sources of support, particularly online resources and AI-based assistants, to meet privacy requirements. Therefore, ongoing research and policy development are needed to help address the limitations of currently available sources and better support developers in handling privacy-related challenges as follows:

- **Improve Accessibility and Clarity of Official Documentation.** To better support developers during software development, privacy-related information must be more accessible and easier to understand. Legislators could contribute by publishing clear and developer-friendly guidance alongside issued legislation. In its current form, official legal documentations are often avoided, as participants consider it too difficult to interpret and apply.
 - **Provide Practical, Jurisdiction-Specific Privacy Guidance.** Developers need reliable and actionable guidance to identify potential privacy issues in their software when navigating privacy requirements independently. Ideally, such guidance should include concrete, implementable scenarios tailored to the jurisdictions relevant to their user base. However, resources of this kind remain relatively limited [22]. To ensure practical usability and legal accuracy, these should be developed through close collaboration between industry experts and regulatory authorities.
 - **Clarify the Role and Limitations of AI Assistants.** AI assistants should disclaim their limitations clearly, particularly when used for sensitive and/or legally relevant privacy tasks (including special categories of data, profiling). Developers must be critically aware of these limitations, verify the accuracy of AI-generated content, and consult additional
- and reliable sources to ensure comprehensive coverage of privacy requirements.
 - **Curated Knowledge Bases.** AI assistants used in legal and privacy contexts must be curated and fine-tuned on legally vetted sources, as in [54] to ensure trustworthiness and compliance of the information they provide and avoid surface recommendations merely based on popularity, documentation visibility, or commercial marketing materials.
 - **Implication for Industry.** Our results show that, with currently available resources, privacy compliance is challenging for developers, if not supported. Companies should encourage support for open communication with privacy experts, and encourage exchange with more experienced employees to support their development teams.
 - **Future Research.** Future research, both from a legal and computer science perspective, should explore how legal information can be presented in an understandable and actionable way. Further, the integration of AI assistants into the software development process as a primary source of information needs to be carefully studied, as our results already show some of their limitations.

6 Conclusion

As privacy grows in importance for legal compliance, developers are increasingly tasked to integrate privacy into software design. However, due to the high cost and limited availability of legal advice, developers frequently address privacy issues independently. In our study with 30 software professionals, we observed a shift from traditional online searches toward the use of AI assistants as a primary source of information, even for complex and sensitive privacy-related questions. Despite this trend, participants consistently struggled to detect privacy issues and propose appropriate mitigation measures. In many cases, they failed to identify critical issues, and the measures they proposed were often limited to general best practices rather than being tailored to the specific context. To address these shortcomings, we recommend that privacy guidance be presented in clear, accessible language and closely aligned with the practical realities of software development. While AI assistants can offer helpful insights and identify some issues, they should not be relied upon as the sole source of support, as many relevant concerns remain unaddressed in their responses.

7 Availability

To support the reproducibility and transparency of our paper, we provide the following supplementary material: (i) pre- and (ii) post-surveys, (iii) the interview guide, (iv) the task description given to the participants, and (v) the codebooks resulting from our analysis. To preserve our participants' privacy, we do not share the transcriptions of the recordings. The supplementary material can be found here: <https://doi.org/10.6084/m9.figshare.30529859>

Acknowledgments

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2092 CASA - 390781972 and the RENFORCE research program of the School of Law of Utrecht University.

References

- [1] Mohamed Abdelhamid. 2021. Fitness Tracker Information and Privacy Management: Empirical Study. *Journal of medical Internet research* 23, 11 (2021), e23059.
- [2] 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 *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, US, 289–305. doi:10.1109/SP.2016.25
- [3] May Alhajri, Ahmad Salehi Shahraki, and Carsten Rudolph. 2022. Privacy of Fitness Applications and Consent Management in Blockchain. *Proceedings of the 2022 Australasian Computer Science Week (2022)*, 65–73.
- [4] Abdulrahman Alhazmi and Nalin Asanka Gamagedara Arachchilage. 2020. Why are Developers Struggling to Put GDPR into Practice when Developing Privacy-Preserving Software Systems? *arXiv preprint arXiv:2008.02987* (2020).
- [5] Abdulrahman Alhazmi and Nalin Asanka Gamagedara Arachchilage. 2021. I'm all ears! Listening to software developers on putting GDPR principles into software development practice. *Personal and Ubiquitous Computing* 25, 5 (2021), 879–892.
- [6] Noura Alomar and Serge Egelman. 2022. Developers Say the Darndest Things: Privacy Compliance Processes Followed by Developers of Child-Directed Apps. *Proceedings on Privacy Enhancing Technologies 2022 (2022)*, 250–273.
- [7] AO Kaspersky Lab. 2025. <https://www.kaspersky.com/> Last accessed: July 17, 2025.
- [8] Article 29 Data Protection Working Party. 2013. *Opinion 03/2013 on Purpose Limitation*. Technical Report WP 203. European Commission. https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2013/wp203_en.pdf Endorsed by the European Data Protection Board (EDPB).
- [9] Owura Asare, Meiyappan Nagappan, and N Asokan. 2024. A User-centered Security Evaluation of Copilot. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. IEEE/ACM, New York, NY, US, 1–11.
- [10] Michael Bailey, David Dittrich, Erin Kenneally, and Doug Maughan. 2012. The Menlo Report. *IEEE Security & Privacy* 10, 2 (2012), 71–75. doi:10.1109/MSP.2012.52
- [11] Anu Bradford. 2020. *The Brussels Effect: How the European Union Rules the World*. Oxford University Press.
- [12] Virginia Braun and Victoria Clarke. 2006. Using Thematic Analysis in Psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [13] Virginia Braun and Victoria Clarke. 2024. *Got questions about TA? We have prepared some answers to some of the common ones we receive*. <https://www.thematicanalysis.net/faqs/>
- [14] Virginia Braun, Victoria Clarke, and Nikki Hayfield. 2022. 'A starting point for your journey, not a map': Nikki Hayfield in Conversation with Virginia Braun and Victoria Clarke About thematic Analysis. *Qualitative Research in Psychology* 19, 2 (2022), 424–445. doi:10.1080/14780887.2019.1670765
- [15] Broadcom. 2025. <https://spring.io/> Last accessed: July 17, 2025.
- [16] John Brooke et al. 1996. SUS-A quick and dirty usability scale. *Usability evaluation in industry* 189, 194 (1996), 4–7.
- [17] Business Insider. 2023. Death by LLM: Stack Overflow's decline, and its plan to survive, shows the future of free online data in an AI world. <https://www.businessinsider.com/stack-overflow-crisis-future-of-online-data-ai-world-2023-7> Last accessed: Jul 17, 2025.
- [18] David Byrne. 2021. A Worked Example of Braun and Clarke's Approach to Reflexive Thematic Analysis. *Quality & Quantity* 56, 3 (June 2021), 1391–1412. doi:10.1007/s11135-021-01182-y
- [19] Carole Cadwalladr and Emma Graham-Harrison. 2018. Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach. <https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election> Last accessed: July 2025.
- [20] Calendly. 2025. <https://calendly.com> Last accessed: Apr 17, 2025.
- [21] Alain Cohn and Michel André Maréchal. 2016. Priming in economics. *Current Opinion in Psychology* 12 (2016), 17–21.
- [22] Commission nationale de l'informatique et des libertés (CNIL). 2020. GDPR Developer's Guide. <https://www.cnil.fr/en/gdpr-developers-guide> Open-source guide for web/app developers, regularly updated via GitHub; latest major update in 2021.
- [23] Zoom Video Communications. 2025. Zoom. <https://www.zoom.com> Last accessed: Jul 17, 2025.
- [24] Data Protection Report. 2025. First multi-million GDPR fine in Germany: €14.5 million for not having a proper data retention schedule in place. <https://www.dataprotectionreport.com/2019/11/first-multi-million-gdpr-fine-in-germany-e14-5-million-for-not-having-a-proper-data-retention-schedule-in-place/> Last accessed: Oct 24, 2025.
- [25] Miguel Delgado-Rodriguez and Javier Llorca. 2004. Bias. *Journal of Epidemiology & Community Health* 58, 8 (2004), 635–641.
- [26] Zack Delie, Sean Radel, Joe Godinez, Garrett Engstrom, Theo Brucker, Kenzie Young, and Sepideh Ghanavati. 2023. Evaluating Privacy Questions from Stack Overflow: Can ChatGPT Compete?. In *2023 IEEE 31st International Requirements Engineering Conference Workshops (REW)*. IEEE, Hannover, Germany, 239–244. doi:10.1109/REW57809.2023.00048
- [27] Xueying Du, Mingwei Liu, Kaixin Wang, Hanlin Wang, Junwei Liu, Yixuan Chen, Jiayi Feng, Chaofeng Sha, Xin Peng, and Yiling Lou. 2024. Evaluating Large Language Models in Class-Level Code Generation. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. IEEE/ACM, Lisbon, Portugal, 1–13.
- [28] DuckDuckGo. 2025. <https://duckduckgo.com/> Last accessed: July 17, 2025.
- [29] R. Eshkita and E. Stamhuis. 2024. The Influence of the Brussels Effect on the Interpretation of Data Protection Laws in the Gulf. *Global Journal of Comparative Law* 13, 2 (2024), 261–278.
- [30] European Council. 2025. <https://www.consilium.europa.eu> Last accessed: July 17, 2025.
- [31] Felix Fischer, Konstantin Böttinger, Huang Xiao, Christian Stransky, Yasemin Acar, Michael Backes, and Sascha Fahl. 2017. Stack Overflow Considered Harmful? The Impact of Copy&Paste on Android Application Security. In *2017 IEEE symposium on security and privacy (SP)*. IEEE, San Francisco, CA, US, 121–136.
- [32] Felix Fischer, Huang Xiao, Ching-Yu Kao, Yannick Stachelscheid, Benjamin Johnson, Danial Razar, Paul Fawkesley, Nat Buckley, Konstantin Böttinger, Paul Muntean, et al. 2019. Stack Overflow Considered Helpful! Deep Learning Security Nudges Towards Stronger Cryptography. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, Santa Clara, CA, USA, 339–356.
- [33] Forbes Media LLC. 2025. <https://www.forbes.com/> Last accessed: July 17, 2025.
- [34] GDPR Enforcement Tracker. 2024. Statistics: Fines imposed over time. <https://www.enforcementtracker.com/?insights> Last accessed: Jul 17, 2025.
- [35] GitHub. 2025. AI that builds with you. <https://github.com/features/copilot> Last accessed: 2025-07-15.
- [36] Google. 2025. <https://www.google.com/> Last accessed: July 17, 2025.
- [37] Google Play Store. 2025. <https://play.google.com> Last accessed: November 5, 2025.
- [38] F.J. Gravetter and L.A.B. Forzano. 2011. *Research Methods for the Behavioral Sciences*. Cengage Learning. <https://books.google.de/books?id=pl04dzBpHy0C>
- [39] Stefan Albert Horstmann, Samuel Domiks, Marco Gutfleisch, Mindy Tran, Yasemin Acar, Veelasha Moonsamy, and Alena Naiakshina. 2024. "Those things are written by lawyers, and programmers are reading that." Mapping the Communication Gap Between Software Developers and Privacy Experts. *Proceedings on Privacy Enhancing Technologies 2024 (2024)*, 151–170.
- [40] Stefan Albert Horstmann, Sandy Hong, David Klein, Raphael Serafini, Martin Engelbert, Martin Johns, Veelasha Moonsamy, and Alena Naiakshina. 2025. "Sorry for Bugging you so much." Exploring Developers' Behavior Towards Privacy-Compliant Implementation. In *2025 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, US, 1215–1233. doi:10.1109/SP61157.2025.00146
- [41] Saki Imai. 2022. Is GitHub Copilot a Substitute for Human Pair-programming? An Empirical Study. In *2022 IEEE/ACM 44th International Conference on Software Engineering: Companion Proceedings (ICSE-Companion)*.
- [42] Stack Exchange Inc. 2024. 2024 Stack Overflow Developer Survey. <https://survey.stackoverflow.co/2024/> Last accessed: Jul 17, 2025.
- [43] Information Commissioners Office. 2025. <https://ico.org.uk> Last accessed: July 17, 2025.
- [44] Samia Kabir, David N Udo-Imeh, Bonan Kou, and Tianyi Zhang. 2024. Is stack overflow obsolete? an empirical study of the characteristics of chatgpt answers to stack overflow questions. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu, HI, USA, 1–17.
- [45] Harjot Kaur, Sabrina Klivan, Daniel Votipka, Yasemin Acar, and Sascha Fahl. 2022. Where to Recruit for Security Development Studies: Comparing Six Software Developer Samples. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 4041–4058. <https://www.usenix.org/conference/usenixsecurity22/presentation/kaur>
- [46] Mugdha Khedkar, Ambuj Kumar Mondal, and Eric Bodden. 2024. Do Android App Developers Accurately Report Collection of Privacy-Related Data?. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering Workshops*. IEEE/ACM, New York, NY, US, 176–186.
- [47] Zachary Kilhoffer, Devyn Wilder, and Masooda Bashir. 2024. Compliance as Baseline, or Striving for More? How Privacy Engineers Work and Use Privacy Standards. In *2024 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, Vienna, 9–18.
- [48] Jan H Klemmer, Marco Gutfleisch, Christian Stransky, Yasemin Acar, M Angela Sasse, and Sascha Fahl. 2023. "Make Them Change it Every Week!": A Qualitative Exploration of Online Developer Advice on Usable and Secure Authentication. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. ACM, Copenhagen, Denmark, 2740–2754.
- [49] Jan H. Klemmer, Stefan Albert Horstmann, Nikhil Patnaik, Cordelia Ludden, Jr. Burton, Cordell, Carson Powers, Fabio Massacci, Akond Rahman, Daniel Votipka, Heather Richter Lipford, Awais Rashid, Alena Naiakshina, and Sascha Fahl. 2024. Using AI Assistants in Software Development: A Qualitative Study on Security Practices and Concerns. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security (Salt Lake City, UT, USA) (CCS '24)*. Association for Computing Machinery, New York, NY, USA, 2726–2740.

- doi:10.1145/3658644.3690283
- [50] Alexander Krause, Harjot Kaur, Jan H Klemmer, Oliver Wiese, and Sascha Fahl. 2025. "That's my perspective from 30 years of doing this": An Interview Study on Practices, Experiences, and Challenges of Updating Cryptographic Code. In *In 34th USENIX Security Symposium*. USENIX Association, Seattle, WA, USA.
- [51] Hans-Alexander Kruse, Tim Puhlfürb, and Walid Maalej. 2024. Can Developers Prompt? A Controlled Experiment for Code Documentation Generation. In *2024 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, Flagstaff, AZ, US, 574–586.
- [52] Lin Kyi, Cristiana Santos, Sushil Ammanaghatta Shivakumar, Franziska Roesner, and Asia Biega. 2025. Turning to Online Forums for Legal Information: A Case Study of GDPR's Legitimate Interests. In *Annual Privacy Forum*. Springer Nature Switzerland Cham, Frankfurt am Main, Germany, 153–177.
- [53] Florian Leiser, Sven Eckhardt, Valentin Leuthe, Merlin Knaeble, Alexander Mädche, Gerhard Schwabe, and Ali Sunyaev. 2024. HILL: A Hallucination Identifier for Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 482, 13 pages. doi:10.1145/3613904.3642428
- [54] Anna Leschanowsky, Farnaz Salamatjoo, Zahra Kolagar, and Birgit Popp. 2025. Expert-Generated Privacy Q&A Dataset for Conversational AI and User Study Insights. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 233, 12 pages. doi:10.1145/3706599.3720014
- [55] Jenny T. Liang, Chenyang Yang, and Brad A. Myers. 2024. A Large-Scale Survey on the Usability of AI Programming Assistants: Successes and Challenges. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering (Lisbon, Portugal) (ICSE '24)*. Association for Computing Machinery, New York, NY, USA, Article 52, 13 pages. doi:10.1145/3597503.3608128
- [56] MAXQDA. 2025. MAXQDA. <https://www.maxqda.com/> Last accessed: Jul 17, 2025.
- [57] Nike, Inc. 2025. <https://www.nike.com/ntc-app> Last accessed: July 17, 2025.
- [58] Sanghak Oh, Kiho Lee, Seonhye Park, Doowon Kim, and Hyoungshick Kim. 2024. Poisoned ChatGPT Finds Work for Idle Hands: Exploring Developers' Coding Practices with Insecure Suggestions from Poisoned AI Models. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, USA, 1141–1159.
- [59] OpenAI. 2025. Introducing ChatGPT. <https://openai.com/index/chatgpt/> Last accessed: 2025-07-15.
- [60] Anna-Marie Orloff, Matthias Fassl, Alexander Ponticello, Florin Martius, Anne Mertens, Katharina Krombolz, and Matthew Smith. 2023. Different Researchers, Different Results? Analyzing the Influence of Researcher Experience and Data Type During Qualitative Analysis of an Interview and Survey Study on Security Advice. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, Hamburg, Germany, Article 864, 21 pages.
- [61] The European Parliament and the Council of the European Union. 2025. GDPR Chapter 2: Principles. <https://gdpr-info.eu/chapter-2/> Last accessed: Jul 17, 2025.
- [62] The European Parliament and the Council of the European Union. 2025. What is GDPR, the EU's new Data Protection Law? <https://gdpr.eu/what-is-gdpr/> Last accessed: Jul 17, 2025.
- [63] Jonathan Parsons, Michael Schrider, Oyebanjo Ogunlela, and Sepideh Ghanavati. 2023. Understanding Developers Privacy Concerns Through Reddit Thread Analysis. *arXiv preprint arXiv:2304.07650* (2023).
- [64] Hammond Pearce, Baleegh Ahmad, Benjamin Tan, Brendan Dolan-Gavitt, and Ramesh Karri. 2025. Asleep at the Keyboard? Assessing the Security of GitHub Copilot's Code Contributions. *Commun. ACM* 68, 2 (2025), 96–105.
- [65] Neil Perry, Megha Srivastava, Deepak Kumar, and Dan Boneh. 2023. Do Users Write More Insecure Code with AI Assistants?. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (Copenhagen, Denmark) (CCS '23)*. Association for Computing Machinery, New York, NY, USA, 2785–2799. doi:10.1145/3576915.3623157
- [66] Privacy Rights Clearinghouse. 2025. <https://privacyrights.org> Last accessed: July 17, 2025.
- [67] reddit inc. 2025. <https://www.reddit.com/> Last accessed: July 17, 2025.
- [68] Marlene Saemann, Daniel Theis, Tobias Urban, and Martin Degeling. 2022. Investigating GDPR Fines in the Light of Data Flows. In *Privacy Enhancing Technologies*. PETS, Sydney, Australia, 314–331.
- [69] Gustavo Sandoval, Hammond Pearce, Teo Nys, Ramesh Karri, Siddharth Garg, and Brendan Dolan-Gavitt. 2023. Lost at C: A User Study on the Security Implications of Large Language Model Code Assistants. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, USA, 2205–2222.
- [70] Shruti Sannon and Andrea Forte. 2022. Privacy Research with Marginalized Groups: What We Know, What's Needed, and What's Next. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 455:1–455:33.
- [71] Scrut Automation. 2025. <https://www.scrut.io/> Last accessed: July 17, 2025.
- [72] Awanthika Senarath and Nalin AG 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*. ACM, New York, NY, USA, 211–216.
- [73] Raphael Serafini, Stefan Albert Horstmann, and Alena Naiakshina. 2024. Engaging Company Developers in Security Research Studies: A Comprehensive Literature Review and Quantitative Survey. In *33rd USENIX Security Symposium (USENIX Security 24)*. USENIX Association, Philadelphia, PA, USA, 3277–3294.
- [74] Shopify. 2025. Free policy generator. <https://www.shopify.com/tools/policy-generator> Last accessed: Jul 17, 2025.
- [75] Stack Exchange Inc. 2025. <https://stackoverflow.com> Last accessed: July 17, 2025.
- [76] State of California Department of Justice. 2018. California Consumer Privacy Act (CCPA). <https://oag.ca.gov/privacy/ccpa> Last accessed: July 2025.
- [77] Alina Stöver, Nina Gerber, Henning Pridöhl, Max Maass, Sebastian Bretthauer, Matthias Hollick, Dominik Herrmann, et al. 2023. How website owners face privacy issues: Thematic analysis of responses from a covert notification study reveals diverse circumstances and challenges. *Proceedings on Privacy Enhancing Technologies 2023* (2023), 251–264. Issue 2.
- [78] Strava, Inc. 2025. <https://www.strava.com> Last accessed: July 17, 2025.
- [79] 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*. ACM, Yokohama, Japan, 1–15.
- [80] Mohammad Tahaei, Marvin Ramokapane, Tianshi Li, Jason I Hong, and Awais Rashid. 2022. Charting App Developers' Journey Through Privacy Regulation Features in Ad Networks. In *The 22nd Privacy Enhancing Technologies Symposium*. De Gruyter Open Ltd., PETS, Sydney, Australia, 33–56.
- [81] Mohammad Tahaei, Kami Vaniea, and Awais Rashid. 2022. Embedding Privacy Into Design Through Software Developers: Challenges & Solutions. *IEEE Security & Privacy* 21, 1 (2022), 49–57.
- [82] Mohammad Tahaei, Kami Vaniea, and Naomi Saphra. 2020. Understanding Privacy-Related Questions on Stack Overflow. 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, 1–14. doi:10.1145/3313831.3376768
- [83] Minoar Hossain Tanzil, Junaed Younus Khan, and Gias Uddin. 2024. ChatGPT Incorrectness Detection in Software Reviews. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering (Lisbon, Portugal) (ICSE '24)*. Association for Computing Machinery, New York, NY, USA, Article 180, 12 pages. doi:10.1145/3597503.3639194
- [84] The European Parliament and the Council of the European Union. 2023. General Data Protection Regulation (GDPR). <https://gdpr.eu/tag/gdpr/> Last accessed: Jul 17, 2025.
- [85] The Federal Trade Commission. 2025. <https://www.ftc.gov> Last accessed: July 17, 2025.
- [86] The International Comparative Legal Guides. 2025. <https://iclg.com> Last accessed: July 17, 2025.
- [87] The New York Times Company. 2025. <https://www.nytimes.com/> Last accessed: July 17, 2025.
- [88] Uber Technologies Inc. 2025. <https://www.uber.com> Last accessed: July 17, 2025.
- [89] Upwork Global Inc. 2025. Upwork. <https://www.upwork.com/> Last accessed: Mar 04, 2025.
- [90] Usercentrics GmbH. 2025. <https://usercentrics.com/> Last accessed: July 17, 2025.
- [91] WHOOP. 2025. <https://www.whoop.com/> Last accessed: July 17, 2025.
- [92] Wikipedia The Free Encyclopedia. 2025. <https://www.wikipedia.org/> Last accessed: July 17, 2025.
- [93] Frank F Xu, Uri Alon, Graham Neubig, and Vincent Josua Hellendoorn. 2022. A Systematic Evaluation of Large Language Models of Code. In *Proceedings of the 6th ACM SIGPLAN international symposium on machine programming*. ACM, San Diego, CA, USA, 1–10.
- [94] Weixiang Yan and Yuanchun Li. 2022. WhyGen: Explaining ML-powered Code Generation by Referring to Training Examples. In *Proceedings of the ACM/IEEE 44th International Conference on Software Engineering: Companion Proceedings*. ACM/IEEE, Pittsburgh, PA, USA, 237–241.
- [95] Zhou Zheng and Yang Liu. 2022. Design of Cloud Storage-Oriented Sports Physical Fitness Monitoring System. *Computational intelligence and neuroscience* 2022, 1 (2022), 1889381.