GDPR Compliance in Privacy Policies of Mobile Apps: An Overview of the State-of-Practice

Orlando Amaral Cejas*
Luxembourg Institute of Science and Technology
orlando.amaral-cejas@list.lu

Sallam Abualhaija, Nicolas Sannier,
Marcello Ceci, Domenico Bianculli
University of Luxembourg
firstName.lastName@uni.lu

***Abstract*—Mobile apps are ubiquitous in our lives as they provide numerous services to support our daily activities. Personalizing such services entail collecting (possibly sensitive) personal information. Mobile apps must therefore comply with privacy regulations like the General Data Protection Regulation (GDPR) enforced in the European Union (EU). To achieve compliance, an app should implement the legal requirements pertinent to data collection and processing according to the GDPR. Privacy policies associated with apps can serve as intermediary instruments connecting between source code and regulations. They explain to app users how activities involving personal data are implemented and provide a detailed view on how legal requirements are operationalized in the app. Incomplete policies can indicate non-compliant apps.**

This paper sheds light on the state-of-practice of GDPR compliance in two mainstream app markets: the Apple App Store and the Google Play Store. We conducted a study to assess the completeness of 470 apps privacy policies in these stores according to the GDPR. Our analysis shows that, irrespective of the app store, fundamental GDPR requirements (e.g., information pertinent to individuals' rights and details of data transfer outside EU) are missing in ≈92% of the analyzed policies, revealing potential breaches in the respective apps.

***Index Terms*—Legal Compliance, Privacy Policies, General Data Protection Regulation (GDPR), Mobile Apps**

I. Introduction

Mobile apps are ubiquitous in the contemporary world: considering only the two mainstream app stores in Europe, their number at the end of 2023 surpassed 2.5 M in the Google Play Store [1] and 1.6 M in the Apple App Store [2]. The large majority of these apps collects and processes some personal data of its users. For these, the General Data Protection Regulation (GDPR) [3], which came into force in May 2018, applies. This entails the need for app developers to identify the applicable legal requirements of GDPR and determine how they can be operationalized in practice. What once were mere privacy concerns are now full-fledged legal requirements of the GDPR, that must be accounted for when developing mobile apps. Examples include the need to obtain explicit consent from individuals [4] for data collection across mobile apps [5] and for sharing personal data with third-parties [6], [7], [8].

Some of the GDPR requirements—so-called *transparency obligations*, outlined in Articles (Art.) 12, 13 and 14—involve information that must be integrated in the privacy policy associated with the app. Privacy policies are legally binding proposals, through which the users are explicitly informed on the various details related to personal data handling practices adopted in the app. Users have then the choice to agree with the policy terms before using the app. However, privacy policies are often verbose, full of legal jargon, and consequently hard to read, making it difficult for average users to fully understand them before agreeing on the terms contained therein. While the introduction of *privacy labels* in app stores [9] provided an alternative (or additional) source of information, they have also been criticized by regulators [10], [11] for potentially misrepresenting the compliance practice and not covering the full range of concerns covered by GDPR.

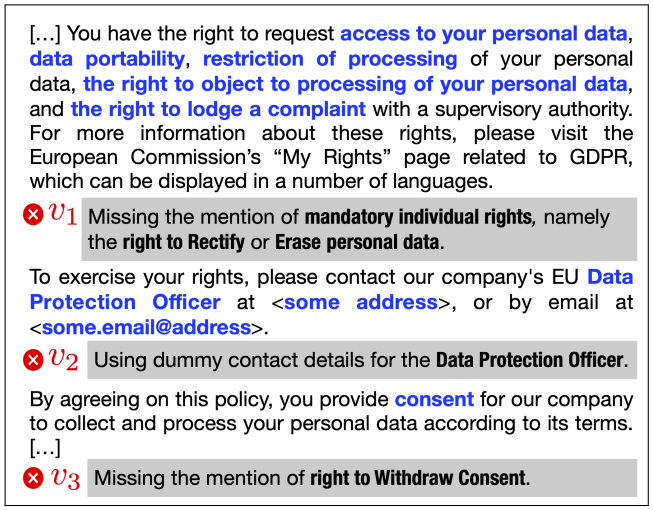

Fig. 1. Excerpt from a Mobile App's Privacy Policy.

Privacy policies should typically reflect an abstract view of the implementation details concerning personal data handling in mobile apps. Requirements engineering (RE) involves ensuring the elicitation of a comprehensive and complete list of legal requirements from the GDPR. Incomplete policies might imply that software requirements related to how personal data is processed in mobile apps could have been missing or remained vague. *Detecting GDPR violations in privacy policies* means determining when information that is required by GDPR is missing from the policy. To illustrate, consider a simplified excerpt of a privacy policy associated with a mobile app, depicted in Figure 1. If the mobile app is collecting data from EU citizens, its privacy policy must comply with the

*This work was done while the author was affiliated with the University of Luxembourg.

content requirements [12] described in Art. 13 GDPR. The primary information types that align with GDPR are set in bold; those included in the policy are highlighted in blue.

Missing an explicit mention of information that should be present in the policy according to GDPR leads to a violation. The example policy has three violations. The first (v_1 in the figure) is caused by not listing the right to *rectification and erasure* among the data subject rights (GDPR, Art. 13(2)(b)). The second violation (v_2 in the figure) is due to inaccurate (dummy) contact details for the *data protection officer* (DPO) (GDPR, Art. 13(1)(b)). Finally, the last violation (v_3 in the figure) is due to not mentioning the right to *withdraw consent* (GDPR, Art. 13(2)(c)). Such a mention is required in the example policy, given that the app is collecting data under the legal basis of "consent" (highlighted in blue in the figure).

Considerable work has been proposed in the literature for checking the completeness of different types of regulated documents against GDPR [13], [14], with an emphasis on detecting violations in privacy policies [15], [16], [17], [18], [19]. Existing work has either a generic scope, such as evaluating the quality of information presented in privacy policies [15], [17], or a more limited scope focusing on specific GDPR-relevant privacy concerns. For instance, among such specific concerns, we mention the following: how the textual content of privacy policies aligns with the data collection practices in mobile apps [19], or whether the information provided in the policies is complete with respect to data collected directly or indirectly from the user [20].

Considering the importance of mobile apps in our daily life and the significant role a privacy policy might have in exposing implementation and operationalization details involving personal data handling, there is a need to observe and understand the state of compliance practice regarding the implementation of GDPR requirements in privacy policies of mobile apps.

In this paper, we report on a comprehensive study that we conducted for assessing the compliance of privacy policies of mobile apps against GDPR. Specifically, we mined a total of 470 privacy policies for 233 mobile apps from the Apple App Store and 237 apps from the Google Play Store. We then leveraged an existing tool (*CompA*ι [21]) to check the compliance of the mined mobile app privacy policies against a set of 19 criteria derived, in close collaboration with legal experts, from GDPR content requirements. The compliance criteria in *CompA*ι cover mandatory content requirements (which, if violated, would result in a clear infringement, such as incompleteness) as well as additional optional requirements, based on best practices according to legal experts, the non-satisfaction of which would raise warnings. The presence or absence of such requirements in a privacy policy can be used as an indicator of the compliance status of the respective app.

Our study reveals a total of 1905 clear violations of GDPR and 363 warnings against the best practices, resulting in a total of 432 (out of 470, ≈92%) privacy policies missing mandatory information. Most common incompleteness issues observed in these policies concern individual rights: ≈64% policies (303/470) miss at least one of the mandatory individuals' rights (e.g., the right to access personal data), while ≈52% (245/470) miss the rights granted to individuals when their data is collected based on informed consent. Other incompleteness issues are related to the absence of details on (i) transferring personal data outside the EU, observed in ≈60% (280/470) policies; (ii) what happens when users fail to provide personal data (280/470); and (iii) categories of (directly or indirectly) collected personal data, in ≈43% (203/470) policies. Further analysis shows that the most common incompleteness issues are related to GDPR requirements that are not well studied in the literature, indicating *a high correlation between the maturity of understanding of privacy requirements in academic versus industrial settings*.

Our analysis did not reveal significant differences between the considered app stores, suggesting lack of general awareness on privacy regulations in the market. While incompleteness issues in a given policy should not necessarily be mapped one-to-one to issues in the source code of the underlying app, they highlight the need for further investigations of the compliance practices in mobile app development. Drafting complete privacy policies, i.e., appropriately implementing the GDPR requirements in the policies, is essential for correctly operationalizing such legal requirements in the apps. Our findings call for more research to improve the community's understanding of GDPR privacy concerns, which would in turn enhance the overall compliance practices in software development.

Data Availability. We have released the scraping tool for extracting privacy policies on Zenodo [22], whereas the dataset of collected policies and the output files of *CompA*ι are available in an online annex [23].

Structure. The remainder of this paper is structured as follows: Section II provides necessary background for our work; Section III illustrates the study design and methodology; Section IV discusses the findings of our study; Section V reviews the relevant literature; Section VI concludes the paper and discusses future work.

II. Background

GDPR. The General Data Protection Regulation (GDPR) came into force in the European Union in 2018. It regulates the data processing activities to ensure that personal data of individuals remain protected throughout the entire data processing chain. GDPR imposes obligations onto organizations inside and outside the EU, as long as personal data of individuals in the EU is being collected or processed. GDPR requires, among other things, that data controllers (i.e., organizations collecting personal data) issue a privacy policy to be agreed upon by the data subjects (i.e., the individuals whose personal data is collected and processed).

Privacy policies are legal agreements that, explain to the user the data collection and processing details, such as what personal data is being collected, for what purpose it will be processed, what rights the users should have on their data, and how personal data will be shared beyond the controller and with whom. Privacy policies should align with the technical

requirements that are implemented in mobile apps [24], [25], [26], [27], thus ensuring the development of GDPR-compliant mobile apps is therefore tightly related to ensuring that their associated privacy policies are complete according to the GDPR content requirements.

CompA*ι *and its usage in our study. The primary goal of our study in this paper is to present an overview of the state-of-practice on GDPR compliance in the mobile apps industry. To achieve this goal, we use *CompA*ι [18], [21] — an existing automated tool for GDPR completeness checking of privacy policies. *CompA*ι leverages natural language processing (NLP) and machine learning (ML) technologies to categorize and classify the textual content of a privacy policy. It implements a set of 23 compliance criteria to assess the presence or absence of the information types according to a comprehensive conceptual model [16], [18] that characterizes the information content (metadata) in privacy policies according to GDPR. The criteria essentially evaluate whether the textual content of a given privacy policy is complete or incomplete according to what is required by GDPR.

Our motivation to select this tool is three-fold. First, the tool has been empirically evaluated over a large set of more than 200 real privacy policies, achieving high precision (92.9%) and recall (89.9%) [18]. These accuracy levels indicate the the tool is suitable for conducting our study. Second, the tool development was informed by the knowledge provided by subject-matter experts since the metadata and compliance criteria were created in close collaboration with legal experts. Third, all necessary artifacts are publicly available, namely the conceptual model, criteria, and the tool.

In its normal operation flow, prior to conducting the analysis, *CompA*ι asks the user to answer a preliminary questionnaire of six questions that capture information which are not derivable from the privacy policy. The answers to these questions determine whether certain information is required and hence whether the respective criteria need to be checked. For instance, the "yes/no" answer to the following question: "Do you plan to transfer the collected personal data outside Europe?" determines whether the information type *Transfer outside Europe* is applicable and must therefore be checked.

Out of the 23 compliance criteria checked by the original tool, only 19 are used in this study. In particular, we excluded the criteria for which the preliminary questionnaire requires open-ended input from the user. For instance, in the preliminary questionnaire the user must input the name of the *Controller* and the *Controller Representative* (where applicable). This information is then used to identify whether relevant information such as the *Identity* and *Contact Details* concerning these entities are present in the privacy policy. Due to the absence of this input, we kept instead other criteria that are triggered by close-ended questions in the questionnaire. Aiming to identify any potential violation, we used a set of default answers that lead to the checking of all criteria (except the four mentioned above), a very likely scenario for mainstream apps. For example, we assume that an app will always transfer personal data outside the EU which triggers the need to provide details on safeguards and measures put in place to ensure that personal data remains protected.

Table I lists the 19 criteria from Amaral et al. [18], which we considered in this work. For each criterion, we provide an ID that will be used in the following sections and a criterion name. We further describe each criterion in terms of preconditions (if applicable) and post-conditions. If a precondition is fulfilled, then the post-condition will be checked. These criteria represent the content requirements to verify whether a privacy policy is complete according to GDPR or not. Some criteria (IDs shaded in red and marked in bold) are mandatory to meet, i.e., not fulfilling them will lead to violations. Other criteria (IDs shaded in orange and italicized) are based on the best practices according to the legal experts' recommendations and are thus not mandatory, i.e., not meeting them raises warnings. We note that the mandatory and optional criteria were determined in the original work [18] in close collaboration with legal experts. In brief, a mandatory criterion originates from the legal interpretation of the GDPR provisions: violating them denotes a clear breach to the GDPR. The optional ones, on the other hand, are related to common practices, indicating that not fulfilling them would require further investigation to determine whether this new practice is still compliant. This distinction was agreed upon among the experts involved in the development of the criteria, including researchers in requirements engineering and legal experts. The criteria pre-conditions and post-conditions are formally represented through activity diagrams [18].

The first five criteria concern data subject rights: Criterion *MR* describes the mandatory rights that individuals have over their personal data in all cases. The remaining criteria (*SA*, *PR*, *OR*, *CR*) are conditionals. For instance, Criterion *CR* states that if personal data is collected based on the legal basis *Consent* and this information is explicitly mentioned in a given privacy policy then, according to GDPR, individuals should have the rights to: erasure of personal data, objection to data processing activities, portability of personal data, and withdrawal of consent. These rights must be explicitly mentioned in the privacy policy.

The remaining criteria are related to transferring personal data outside Europe (*ToE*, *ToED*, *ADD*, *SD*, and *EC*), personal data collection (*IC* and *ICD*), personal data category (*Cat* and *CT*), the recipients who will get personal data besides the data controller (*Rec*), the retention period of storing personal data (*RP*), the purpose for processing personal data (*PP*), what happens when the data subject fails to provide personal data to the controller (*FtD*), and the provision of the DPO's contact details (*DPO*).

III. Study Design

In this section, we present the research questions we investigated, describe the pipeline implemented for conducting our study, and introduce the privacy policies dataset created as part of our work.

TABLE I
GDPR COMPLETENESS CRITERIA OF PRIVACY POLICIES USED IN THE STUDY BASED ON REFERENCE [18]

ID[1]	Criterion Name	Criterion Description
MR	Mandatory Rights	The following rights must always be specified: the right to access and update personal data, the right to restrict the use of personal data, and the right to lodge a complaint with a supervisory authority.
SA	Supervisory Authority	If the right to lodge a complaint with a supervisory authority is specified, the corresponding supervisory authority must also be specified.
PR	Portability Right	If the legal ground for the processing of personal data is based on contract, then data subjects must have the right to obtain and reuse their personal data for their own purposes across different services.
OR	Right to Object	If the legal ground for the processing of personal data is based on either legitimate interests or public function, then the data subjects must have the right to object to the processing of their personal data at any time.
CR	Consent Rights	If the legal ground for the processing of personal data is based on consent, then the data subjects must have the following rights: the right to erase personal data, the right to object to the processing of personal data, the right to obtain and reuse personal data, and the right to withdraw consent at any time.
ToE	Transfer outside Europe	If there is an intent to transfer personal data to third countries outside Europe, then information regarding transferring the personal data outside Europe must be specified.
ToED	Transfer outside Europe Details	If personal data is transferred to third countries outside Europe, then at least one of the following information about the transfer should be specified: applicable adequacy decision, legal safeguards, or specific derogations.
ADD	Adequacy Decision Details	If personal data is transferred to third countries outside Europe and an adequacy decision is in place, then the appropriate information regarding the type of the adequacy decision should be specified. Specifically, the type can refer to the adequacy decision between the EU and a territory (e.g., Andorra, the Bailiwick of Jersey), a specific sector (e.g., the commercial organizations from Canada, Argentina), or a country (i.e., Japan).
SD	Safeguards Details	If personal data is transferred to third countries outside Europe and safeguards are in place, then the appropriate EU contractual clauses or binding corporate rules should be specified.
EC	Explicit Consent	If derogation in specific situations is applicable to the transfer of personal data to third countries outside Europe, then an explicit consent should be specified.
IC	Indirect Data Collection	If personal data is indirectly collected from an individual, then this information must be specified.
ICD	Indirect Data Collection Details	If personal data is indirectly collected from an individual, then the indirect sourced should be specified, namely publicly available and/or third-party sources.
Cat	Personal Data Category	If personal data is indirectly collected from an individual, the categories of the collected personal data must always be specified.
CT	Category Type	If personal data is collected indirectly through publicly available and/or third-party sources, then the categories of the collected personal data should be related to the indirect sources. For example, the category about a person's name is collected from publicly available sources while the category related to a person's medical record is collected from a third-party source.
Rec	Recipients of Personal Data	If there are other recipients (natural or legal person, public authority, agency or another body, to which the personal data is disclosed) of the collected personal data besides the controller, whether a third party or not, those recipients must always be specified.
RP	Retention Period	The data retention period of the collected personal data must always be specified.
PP	Processing Purposes	The purposes for processing personal data must always be specified.
FtD	Failure to Provide Data	If the legal ground for the processing of personal data is based on either contract (more specifically to the steps taken for entering a contract), or on legal obligation, then the data subject is obliged to provide the personal data and must be informed of the possible consequences of failure to provide such data.
DPO	Data Protection Officer	When any of Art. 37.1 (a), (b), or (c) applies, a Data Protection Officer must be specified.

[1] **Mandatory Criteria** / *Optional Criteria (Recommendations)*

A. Research Questions (RQs)

In this study, we investigate the following RQs:

RQ1. What are the most common GDPR-relevant incompleteness issues in privacy policies of mobile apps in the wild? To answer this RQ, we first apply *CompA*ι to check the completeness criteria (listed in Table I) in the collected privacy policies. We then analyze the results and provide a comprehensive overview of the identified violations.

RQ2. How do incompleteness issues vary across categories and app stores? We created our dataset by collecting privacy policies from both the Apple App Store and Google Play Store, two major platforms for hosting mobile apps. In this RQ, we break down the results of RQ1 to better understand incompleteness violations across mobile apps' categories and across app stores.

RQ3. What is the correlation between the awareness manifested in the research landscape on GDPR privacy

in mobile apps and the development practices mirrored in the privacy policies of mobile apps? To answer this RQ, we investigate possible correlations between the incompleteness issues identified in RQ1 versus the research coverage of GDPR privacy concerns in the software engineering literature, outlined in a recent systematic literature review [28]. The goal of this RQ is to identify possible research gaps where certain GDPR privacy concerns are less investigated.

B. Privacy Policy Completeness Checking Pipeline

Figure 2 illustrates our implementation pipeline, composed of four steps. In Step 1, we run a web scraper that mines privacy policies from the considered app stores. In Step 2, we clean the collected privacy policies. The resulting policies are then passed on to Step 3 where we use *CompA*ι to identify the GDPR incompleteness issues (both violations and warnings). Finally, in Step 4, we analyze the results and produce the final output (a CSV file), which outlines the fulfillment status of each criterion in the analyzed policies. Specifically, the file indicates, for each policy, whether a criterion was satisfied, not applicable (N/A), violated, or triggered a warning. Below, we elaborate on the first three (main) steps.

Step 1 - Mining Privacy Policies. We built a scraper to identify and download the privacy policies of mobile apps from the analyzed stores. The scraper takes as input a set of URLs containing the app categories to be mined in the analyzed stores (see Section III-C); it then extracts the privacy policies associated with all the apps available in the input URLs. We designed the scraper to navigate the input URLs only, without crawling subsequent links. The rationale behind this decision is twofold. First, having to navigate through successive possible links and pages requires complex parsing strategies that may not guarantee to retrieve the actual policy without collecting additional noise (i.e., pages not relevant to privacy policies). Second, the fact that a privacy policy is hidden behind multiple consecutive URLs does not align with the GDPR principle on transparency, described in Art. 5(1) [3] and further defined in its recital 58, stating that "any information addressed to the public or to the data subject be concise, easily accessible and easy to understand [...]." Motivated by this transparency principle, despite being not normative, we scope our study to those apps whose policies are directly and easily accessible. Integrating transparency requirements into the assessment of compliance practices is left for future work.

Step 2 - Dataset Cleaning. In this step, we preprocessed the set of scraped privacy policies and prepared them for our analysis. We present in Section III-C the details on how the final dataset was created. The cleaned set of privacy policies was then passed on to Step 3.

Step 3 - Completeness Checking. In this step, we used *CompA*ι, introduced in Section II, to analyze the completeness of the policies in our dataset with respect to GDPR provisions, focusing exclusively on the 19 criteria listed in Table I. *CompA*ι takes as input a privacy policy, the predefined answers to the preliminary questionnaire, and generates as output a report summarizing the identified textual content for each criterion as well as the status of that criterion. We collected such completeness reports for all privacy policies considered in this study and uses them as the basis for our findings on the GDPR compliance in practice, as elaborated in the next section.

Specifically, for each criterion, *CompA*ι provided one of the following possible decisions: (a) *Satisfied*, meaning that the criterion was fulfilled and the required information content was identified in the input policy; (b) *Violation*, meaning that a mandatory criterion was violated and the required information content was not identified in the input policy; (c) *Warning*, meaning that an optional criterion was violated in the input policy; or (d) *Not applicable*, meaning that the criterion was not applicable considering decisions made on previous criteria. For example, if Transfer outside Europe (*ToE*) is not fulfilled, then it is no longer necessary to check the corresponding details of *ToE* such as *ADD*, *SD*, and *EC*. We recall from Section II that *CompA*ι's criteria leading to a violation are derived directly from GDPR, whereas those leading to a warning are recommended as best practices by the legal experts.

C. The Mobile Apps Privacy Policies Dataset

In this section, we present the methodology we followed to create our dataset of mobile app privacy policies, hereafter referred to as *APRICOT* (App PrivaCy pOliciy daTaset). In line with the goal of this study, i.e., to analyze the implementation of GDPR privacy concerns in privacy policies of mobile apps, our priority was to define a representative set of apps. To achieve this, we configured our scraper to target, when possible, the most popular apps, since breaches of GDPR in such apps could have large-scale consequences affecting a significant number of users. This helps avoid mining low-quality and less popular apps such as adwares (advertisement-based apps showing limited functionality and content) where legal breaches would have less practical considerations.

The scraping process took place in November 2023 and targeted privacy policies of apps, organized in different categories, in the Apple App Store and the Google Play Store. Such categories could be accessed by navigating the URLs of the two stores. We considered two categories in the Apple App Store (*Apps* and *Games*) and five in the Google Play Store (*Games*, *Apps*, *Kids*, *Books*, and *Movies & TV*). For illustrative purposes, Figure 3 shows a screenshot of the different categories available in the analyzed stores (note that in Figure 3 the category *Kids* is currently referred to as "Children"). Considering the scope of our study, we excluded the categories *Books* and *Movies & TV*. The first step of the scraping process involved collecting policies of apps listed under the *Apps* and *Games* categories from both stores, and under the category *Kids* only in the Google Play Store. Consequently, we collected a total of 550 candidate privacy policies, of which 287 originated from the Apple App Store and 263 from the Google Play Store.

To ensure the dataset was clean and relevant to our study, in Step 2 we excluded policies according to the following criteria:

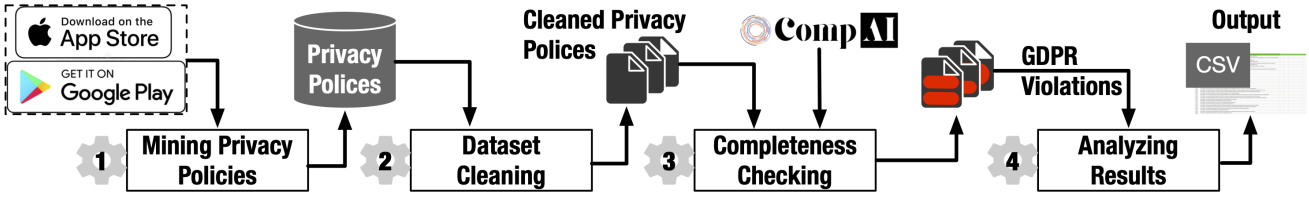

Fig. 2. Pipeline Overview

TABLE II
STATISTICS FOR APRICOT

Category	Apple App Store				Google Play Store				Final
	$\neg EN$	$len < 2$	*Dup*	Rest	$\neg EN$	$len < 2$	*Dup*	Rest	
Apps	7	31	2	154	4	-	-	73	227
Games	5	7	2	79	10	2	8	79	158
Kids	-	-	-	-	-	1	1	85	85
Summary	12	38	4	233	14	3	9	237	470

$\neg EN$, $len < 2$, Dup correspond to the number of privacy policies that are filtered out when their content contained language other than EN, they were too short (less than two pages), or they had duplicate content, respectively.

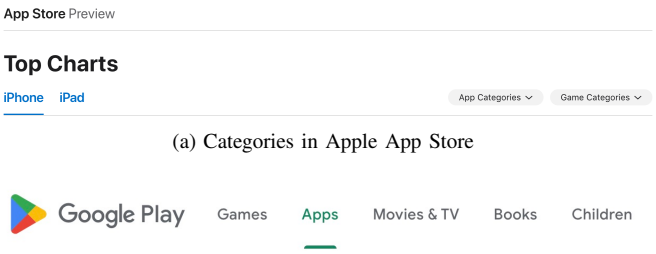

(a) Categories in Apple App Store

(b) Categories in Google Play Store

Fig. 3. A Screenshot of Apps' Categories (date: 11/2024)

(1) Not in English ($\neg EN$): Since *CompAι* supports only the English language, we discarded policies containing (fully or in part) text in a language other than English.

(2) Too short ($len < 2$ pages): We noticed that the content of some policies was too short to allow for a meaningful analysis. For example, a policy can contain a mere link to the actual privacy policy (we recall that our crawling process stops at the first link from which the policy is downloaded). Since the scraping tool (Step 1) produces consistently-formatted Word documents containing the identified privacy policies, we introduced a heuristic to filter out all policies whose length was less than two pages.

(3) Duplicate content (Dup): In some cases, we collected privacy policies that are under different names but have the same content, e.g., when such policies are issued by the same organization to cover multiple apps. We removed these duplicates to ensure that our analysis is representative by covering different policies, and also to avoid amplifying incompleteness issues originated from the same policies.

As a result, we excluded from the *Apps* category in Apple App Store two duplicate privacy policies, seven policies in languages other than English, and 31 "too short" policies. We further excluded from the *Games* category two duplicate privacy policies, five containing languages other than English, and seven "too short".

As for the policies collected from Google Play Store, we filtered out a total of 26 privacy polices across the different categories: nine duplicates, 14 privacy policies in languages other than English, and three "too short" policies. It is worth noting that 20 of the excluded policies belonged to the *Games* category, including ten duplicates, eight not in English, and two "too short".

To validate the automated cleaning process, we manually checked and confirmed that these privacy policies, indeed, did not contain any meaningful information for our analysis. Thus, they were excluded from our final dataset.

Table II shows some statistics of the APRICOT dataset, including the number of policies filtered out for the above reasons (see columns named after the corresponding reasons) and the number of policies that remained after the cleaning step (see the columns named *Rest*). Our final dataset contains 470 privacy policies, of which 233 (49.6%) originate from the Apple App Store and 237 (50.4%) from the Google Play Store.

Implementation. We implemented the pipeline shown in Figure 2 in Java and Python. To develop the scraper (Step 1), we used JSON.simple toolkit (v1.1.1) for preparing the necessary input, jsoup (v1.10.2) and Apache POI (v3.17) for parsing the HTML content and retrieving the policies. We used the Python language detection library (pycld2 v0.41) and the library python-docx (v1.1.2) for creating Word documents containing the policies (the input format required by *CompAι*). In Step 3, we provided as input to *CompAι* the retrieved policies and the default answers to the questionnaire. The same answers were used across all policies.

IV. FINDINGS

In this section, we answer the RQs presented in Section III, present our findings and further discuss them in relation to broader industry practices.

TABLE III
GDPR INCOMPLETENESS ISSUES IN APRICOT DATASET (**RQ1**)

Criterion	Satisfied	N/A	Violated/ Warnings
MR	167 (35.5%)	0 (0.0%)	**303** (**64.5**%)
PR	260 (55.3%)	109 (23.2%)	101 (21.5%)
OR	217 (46.1%)	63 (13.4%)	190 (40.4%)
CR	146 (31.1%)	79 (16.8%)	**245** (**52.1**%)
ToE	190 (40.4%)	0 (0.0%)	**280** (**59.6**%)
IC	455 (96.8%)	0 (0.0%)	15 (3.2%)
Cat	450 (95.7%)	0 (0.0%)	20 (4.3%)
Rec	446 (94.9%)	0 (0.0%)	24 (5.1%)
RP	289 (61.5%)	0 (0.0%)	181 (8.5%)
PP	422 (89.8%)	0 (0.0%)	48 (10.2%)
FtD	139 (29.6%)	51 (10.9%)	**280** (**59.6**%)
DPO	252 (53.6%)	0 (0.0%)	218 (46.4%)
SA	137 (29.2%)	289 (61.5%)	44 (9.4%)
ToED	167 (35.5%)	280 (59.6%)	23 (4.9%)
ADD	6 (1.3%)	419 (89.2%)	45 (9.6%)
SD	142 (30.2%)	292 (62.1%)	36 (7.7%)
EC	20 (4.26%)	450 (95.7%)	0 (0.0%)
ICD	443 (94.3%)	15 (3.2%)	12 (2.6%)
CT	240 (51.1%)	27 (5.7%)	203 (43.2%)

N (P%): The absolute numbers and percentages of privacy policies (out of 470) for which a criterion was satisfied, N/A, or violated.

A. RQ1 - Incompleteness in Privacy Policies

At the privacy policy level, we observe that 38 out of the 470 policies (about 8%) do not contain any violation (i.e., they satisfied all mandatory requirements in GDPR, if applicable), but could have warnings due to not following the best practices. Only eight policies (1.7%) in our dataset have fully satisfied the 19 criteria (if applicable), meaning that they did not contain any violation or warning. These observations indicate that 98.3% of the analyzed privacy policies had at least one warning, and 92% had at least a violation or a warning. The number of violations per policy ranges between zero and nine, with the majority of policies having three, five, or six violations, whereas the number of warnings is at most two per policy. This notable low compliance rate is comparable to what has been reported by other studies [20], [29], [30] and complement these studies with the fine-grained GDPR compliance analysis enabled by *CompAι*.

Table III presents the fulfillment status for the individual 19 criteria (see Section II), providing the absolute numbers of privacy policies and their corresponding percentages under each status, namely *Satisfied*, *Not Applicable* (N/A), *Violated* (in the case of mandatory requirements, listed in the top part of the table) or resulted in a *Warning* (in the case of optional requirements, listed in the bottom part of the table). In each column, the values are aggregated over the privacy policies coming from both stores. We note that the same privacy policy might have multiple violations and warnings at the same time.

We observe that the criteria causing the highest number of violations are *MR*, *CR*, *ToE*, and *FtD* (whose values are set in bold in the table). The first two are concerned with the data subject rights, i.e., the duty to inform the users about all their rights over their personal data. According to GDPR, a data subject (in our case, the user of an app) has the right to access and rectify their personal data, restrict the processing activities, and lodge a complaint. Additionally, when personal data is collected based on the legal basis *Consent* (which is typically the case in mobile apps), the data subject has the rights to erasure, objection, portability, and consent withdrawal. Despite the emphasis on individual rights in GDPR, 303 out of 470 privacy policies in our analysis failed to explicitly mention at least one of the mandatory rights in the first case (thus violating *MR*), and 245 out of 470 failed to mention the necessary rights associated with the legal basis consent, leading to a violation of *CR*.

Regarding the other information types, *ToE* concerns the possibility of transferring personal data outside the EU. This criterion is strictly based on the preliminary questionnaire that *CompAι* asks the user to answer before performing the completeness checks. In our analysis, we assumed the default answer *yes* to the question about the intention to transfer the collected personal data outside the EU (by the app in this case). Given that we mined the policies from the global app stores, we believe that this default answer is plausible. Our results indicate that 280 of the considered privacy policies had no mention of transferring data outside the EU (thus violating *ToE*). While the assumption of additional EU data transfer is unlikely to hold for all apps, we still believe that privacy policies violating this criterion should be further checked.

Finally, criterion *FtD* corresponds to the requirement to explain clearly what happens in case the user fails to provide personal data. For instance, if the user does not share personal data about their location, some services provided by the app based on location tracking will not be available. Such detailed content must be present in the privacy policy. Again, 280 out of the analyzed policies violated this criterion.

With regard to the best practices (optional criteria resulting in warnings), our analysis shows that criterion *CT* caused the highest number of warnings, in 203 of the analyzed privacy policies. *CT* concerns the type of personal data categories that, according to GDPR, must be explicitly listed, in particular when personal data is collected indirectly. For instance, if a mobile app collects personal data on usage analysis produced by third party in order to improve the personalization of the app, then what exactly is collected from those libraries (e.g., usage statistics) must be mentioned in the privacy policy. This criterion is strictly based on the default answer (which we assumed to be *yes*) to the question on whether data are collected not only directly from the user, but also indirectly from other sources. Similar to our analysis provided above for *ToE*, we cannot rule out the possibility that the policies raising this warning are not collecting personal data indirectly. However, the high percentage in relation to this warning highlights, once again, the necessity for further checks.

The answer to **RQ1** is that most of the incompleteness violations that arise in privacy policies mined for mobile apps in the wild are due to missing information concerning: (i) mandatory users' rights (*MR*, identified in 64.5% of

the 470 analyzed privacy policies), (ii) users' rights in relation with the legal basis consent (*CR*, 52.1%), (iii) the mention of transfer outside Europe (*ToE*, 59.6%), and (iv) information about what happens in case of failure to provide personal data by the user (*FtD*, also 59.6%). Missing information concerning the categories of collected personal data (*CT*), on the other hand, has led to warnings in 43.2% of the privacy policies. Our results raise the question about whether such incompleteness issues might as well be propagated to the mobile apps development practices, i.e., less attention given to appropriate integrating features that allow users to exercise their rights.

B. RQ2 - Incompleteness across Categories and Stores

At the privacy policy level, we recall from RQ1 that 38 policies did not contain any GDPR violation in both app stores. This number involves 13 policies under the *Apps* category and six policies under the *Games* category for the Apple App Store, in addition to 11 under *Apps*, three under *Games*, and five under *Kids* for the Google Play Store.

Table IV breaks down the results for each criterion across the different app stores (recall we analyzed a total of 233 privacy policies from the Apple App Store and 237 from the Google Play Store). Like in RQ1, the table shows the absolute numbers and respective ratios of policies containing violations or warnings. In summary, and similarly to the policy level, the analysis reveals that GDPR content requirements are often violated in privacy policies associated with mobile apps nearly equally across the analyzed stores.

However, the table shows a dissimilar distribution of violations and warnings across the two stores. The percentage of privacy policies containing violations is higher in the Google Play Store for six criteria, namely *MR*, *OR*, *CR*, *ToE*, *FtD*, *RP*. The average percentage is 61.1%, compared to 20% in the Apple App Store. In contrast, policies from the Apple App Store, have more violations in the remaining six criteria with an average of 20% compared to 10.3% in Google Play Store.

We also extend our analysis to the app category level. Our results, presented in Figure 4, show the analysis per app category. In the case of the Apple App Store, we observe more violations in the *Apps* category compared to *Games*. For this store, once again, the results suggest that fundamental and mandatory information according to GDPR such as *MR* and *DPO* have been often neglected. In the case of apps from the Google Play Store, we observe that the number of violations in the *Games* and *Kids* categories are much higher than those in the *Apps* category. After a detailed analysis of the criteria, we observe a variation in the levels of emphasis placed on different information content across categories. A case to highlight is the one of individuals' rights which, despite being fundamental in GDPR, have often received very little attention.

The answer to **RQ2** is that there are no significant differences between the number of violation or warnings in the Apple App Store and the Google App Store. Our results with respect to violations observed across categories further indicate the necessity for additional controls in both app stores. With the varying distribution of violations from one app category to another, we conclude that there is a need for more compliance guidelines for both app companies and the two platforms. In particular, information on mandatory rights (*MR*), rights related to consent (*CR*) and Transfer outside the EU (*ToE*) are often missing across all app categories in both app stores.

C. RQ3 - Academic Research versus Industry Practice

In this RQ, we are interested in discovering the possible correlations between the GDPR privacy concerns highlighted in our study and their level of investigation in the software engineering literature. To answer this question, we refer to a recent systematic literature review on GDPR-relevant privacy concerns in mobile apps research [28]. To understand the coverage of GDPR concerns in the literature, the authors reviewed 60 papers published during the period 2018–2023. Then, they mapped the mentioned concepts in the literature to a comprehensive conceptual model [18] characterizing the information content of privacy policies according to GDPR provisions. Our study draws on this SLR and investigates the relation between the GDPR concepts that are reported to be under-explored in reference and the missing information content leading to incompleteness, as discussed earlier.

Table V presents side-by-side the GDPR criteria analyzed in our study and the number of papers from the literature that discuss a relevant concept to each criterion. For the definitions of the GDPR concepts, we refer the reader to the original paper [18] where the conceptual model is introduced.

The table shows that the criteria that are primarily violated by the majority of the privacy policies in our study (namely *MR*, *CR*, *ToE*, *FtD*, and *DPO*) received very little attention in the literature. For instance, *MR*, which concerns the mention of mandatory individuals' rights (i.e., Data Subject Rights) has been investigated only at a very generic level in the literature, without details on such rights [28]. We note that, among these rights, the rights to access and rectification of personal data are highly relevant in the context of mobile app development. This lack of awareness about data subject rights further affected the violations observed for criterion *CR*, despite the research focus on Consent (in 21 papers). Similar correlations can be concluded concerning GDPR concepts associated with *ToE*, *FtD*, and *DPO* that were under-explored in the literature. Other criteria such as *Cat*, *Rec*, *RP*, and *PP* are related to GDPR concepts that have been widely investigated in the literature, thus explaining the relatively lower number of violations associated with them.

A similar pattern can be seen regarding the top-five criteria that lead to N/A and warnings, namely: *SA*, *ToED*, *ADD*, *SD*, and *EC*. These criteria have received very little visibility in the literature. The last two criteria *ICD* and *CT* are mostly satisfied despite the lack of coverage in the literature. *CT* is

TABLE IV
GDPR INCOMPLETENESS ACROSS APP STORES (**RQ2**)

Criterion	Satisfied		N/A		Violated /Warnings	
	AS	GS	AS	GS	AS	GS
MR	99 (42.5%)	68 (28.7%)	0 (0.0%)	0 (0.0%)	134 (57.5%)	**169 (71.3%)**
PR	91 (39.1%)	169 (71.3%)	78 (33.1%)	31 (13.1%)	**64 (27.5%)**	37 (15.6%)
OR	123 (52.8%)	94 (39.7%)	45 (19.3%)	18 (7.6%)	65 (27.9%)	**125 (52.7%)**
CR	78 (33.5%)	68 (28.7%)	57 (24.5%)	22 (9.3%)	98 (42.1%)	**147 (62.0%)**
ToE	104 (44.6%)	86 (36.3%)	0 (0.0%)	0 (0.0%)	129 (55.4%)	**151 (63.7%)**
IC	221 (94.9%)	234 (98.7%)	0 (0.0%)	0 (0.0%)	**12 (5.2%)**	3 (1.3%)
Cat	218 (93.6%)	232 (97.0%)	0 (0.0%)	0 (0.0%)	**15 (6.4%)**	5 (2.1%)
Rec	211 (90.6%)	235 (99.2%)	0 (0.0%)	0 (0.0%)	**22 (9.4%)**	2 (0.8%)
RP	166 (71.2%)	123 (51.9%)	0 (0.0%)	0 (0.0%)	67 (28.8%)	**114 (48.1%)**
PP	195 (83.7%)	227 (95.8%)	0 (0.0%)	0 (0.0%)	38 (**16.3%**)	10 (4.2%)
FtD	79 (33.9%)	60 (25.3%)	37 (15.9%)	14 (5.9%)	117 (50.2%)	**163 (68.8%)**
DPO	105 (45.1%)	147 (62.0%)	0 (0.0%)	0 (0.0%)	**128 (54.9%)**	90 (38.0%)
SA	81 (34.8%)	56 (23.6%)	126 (54.1%)	163 (68.8%)	**26 (11.2%)**	18 (7.6%)
ToED	88 (37.8%)	79 (33.3%)	129 (55.4%)	151 (63.7%)	**16 (6.9%)**	7 (3.0%)
ADD	5 (2.2%)	1 (0.4%)	207 (88.8%)	212 (89.5%)	21 (9.0%)	**24 (10.1%)**
SD	71 (30.5%)	71 (30.0%)	137 (58.8%)	155 (65.4%)	**25 (10.7%)**	11 (4.6%)
EC	12 (5.2%)	8 (3.9%)	221 (94.9%)	229 (96.7%)	0 (0.0%)	0 (0.0%)
ICD	212 (91.0%)	231 (97.5%)	12 (5.2%)	3 (1.3%)	**9 (3.9%)**	3 (1.3%)
CT	147 (63.1%)	93 (39.2%)	21 (9.0%)	6 (2.5%)	65 (27.9%)	**138 (58.2%)**

† *N (P%): The absolute numbers of privacy policies and their percentages in which the criteria were satisfied, N/A, or violated.*
‡ *Percentages are computed out of 233 policies for Apple App Store (AS), and 237 policies for Google Play Store (GS).*

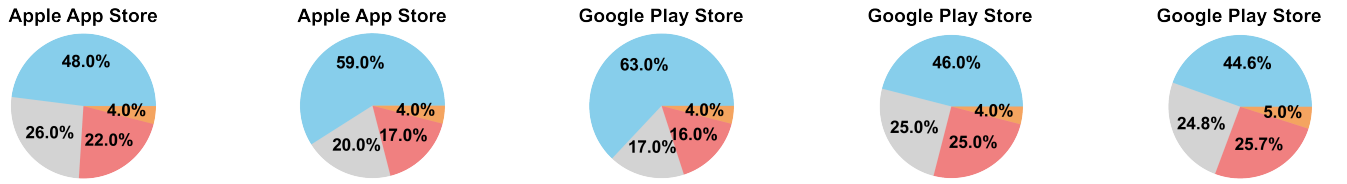

(a) Apps Cat. (154 apps) (b) Games Cat. (79 apps) (c) Apps Cat. (73 apps) (d) Games Cat. (79 apps) (e) Kids Cat. (85 apps)

Fig. 4. GDPR Incompleteness Analysis across App Categories (SATISFIED , VIOLATED , WARNING , NOT APPLICABLE)

TABLE V
AWARENESS OF GDPR PRIVACY IN RESEARCH VERSUS GDPR COMPLETENESS ANALYSIS IN PRACTICE (**RQ3**)

Criterion	%	Relevant Concepts (Number of Papers)
MR	64.5	Data Subject Right (5)
PR	21.5	Legal Basis, Contract (1); Data Subject Right (5)
OR	40.4	Legal Basis, Legitimate Interest, Public Function (1); Data Subject Right (5)
CR	52.1	Legal Basis, Consent (21); Data Subject Right (5); Data Subject Right, Withdraw Consent (5)
ToE	59.6	Transfer outside Europe (1)
IC	3.2	PD Origin, Indirect, Third Party, Cookie (3)
Cat	4.3	PD Category (14)
Rec	5.1	Recipients (26)
RP	38.5	PD Time Stored (9)
PP	10.2	Processing Purposes (8)
FtD	59.6	*PD Provision Obliged* (0)
DPO	46.4	DPO (1)
SA	9.4	Data Subject Right (5)
ToED	4.9	Transfer outside Europe, Safeguards (1)
ADD	9.6	*Transfer outside Europe, Adequacy Decision* (0)
SD	7.7	Transfer outside Europe, Safeguards (1)
EC	0.0	*Transfer outside Europe, Specific Derogation, Unambiguous Consent* (0)
ICD	2.6	PD Origin, Indirect, Third Party (2)
CT	43.2	PD Category, Type (1)

Mandatory Criteria (violation) / *Optional Criteria (warning)*

dependent on *Cat*, which has good coverage in the literature (14 papers). *ICD* is however closely related to *IC*, which has not been investigated very well in the literature. These patterns show that the privacy concept of indirect collection of personal data has matured in practice before getting much attention in the literature. A possible reason is the close relation of indirect collection to the concepts of cookies and third parties, which are important since before the introduction of GDPR (see the amended EU "ePrivacy Directive" or "Cookie Law" [31]). Thus, these concepts are well-known and accounted for in the privacy policies of mobile apps.

The answer to **RQ3** is that there is a clear correlation between academic research and compliance practice. The lack of awareness about certain GDPR concepts such as data subject rights explains the high number of violations in the corresponding criteria. Our findings suggest the need for more academic work investigating GDPR privacy concerns to enhance the understanding of corresponding legal requirements, which would in turn positively affects the overall GDPR compliance practices.

D. Discussion

Our study identified numerous incompleteness issues and highlighted insufficient regulatory mechanisms to prevent launching mobile apps with incomplete privacy policies. These findings show that some fundamental GDPR principles such as data subject rights are still not well understood after years of GDPR enforcement within the regulatory framework. Below, we first discuss the correlation between research and practice and then outline key calls-for-action resulting from our study.

From research on GDPR to compliance practice. Our analysis revealed a clear correlation between academic research and industry practice, where GDPR concepts that are underrepresented in research are not appropriately implemented in practice. Despite the high numbers of publications on GDPR requirements, the majority of existing work is limited in scope and shows shallow or partial understanding of the regulation. For instance, "informed consent" is a concept that has been widely studied, yet many privacy policies in our analysis have violated the closely-related GDPR obligations on the disclosure of "mandatory rights". This suggests that research has often addressed GDPR concepts in isolation, lacking a comprehensive view of the regulation. As a result, there is a gap leading in effectively implementing GDPR legal requirements in practice.

Other examples confirm this correlation. For instance, the right of access and right to portability are two seemingly straightforward concepts but vaguely described in the GDPR. Limited research has explored these concepts in depth, leading to shallow interpretation and non-compliant implementation of these rights. As such, in practice, a direct request to access personal data through the app (e.g., by clicking a button) or direct data portability from one service to another (i.e., without heavy involvement of the user) are hardly implemented in mobile apps. Additionally, even when GDPR highlights PDF as a non-satisfactory format for data portability, no sufficient details are provided related to popular data formats like CSV or JSON. Without adequate legal expertise and in-depth understanding of these rights, identifying a technical and compliant solution remains a major challenge for app development industry.

Call-for-action for requirements engineers. Our study is primarily related to requirements elicitation, which is the activity of gathering the necessary requirements about a system-to-be from stakeholders [32]. As such, requirements engineers involved in mobile app development are informed by our study about the necessity of collecting detailed information on data protection obligations under GDPR in collaboration with legal experts. Since data subject rights (directly impacting users) remain superficially addressed, mining app reviews [33] to analyze user feedback can help better understand the users' awareness of their data protection rights, and how this awareness correlates with app usage patterns and its impact on the compliance status of mobile apps.

Call-for-action for industry practitioners. The numerous violations identified in our study emphasizes the need for practical guidelines with checklists of GDPR requirements and their respective minimal compliance-level for industry practitioners informed by existing guidelines and court decisions. Enforcing additional control mechanisms to ensure the completeness of privacy policies of released apps can increase the awareness of developers about obligations to be addressed in mobile apps. Without such a clear guidance and quality control over privacy policies, insufficient compliance practices are likely to persist.

E. Threats to Validity

Internal Validity. The main threat to internal validity is the possible bias in identifying the violations of GDPR due to subjective interpretation of the text of the privacy policies. We mitigated this threat by using an existing tool which was developed and evaluated independently from this study. Another threat concerns the preliminary questionnaire required to run *CompA*ι, where we provided default answers to most of the questions. In this study we discarded four completeness criteria which were strictly dependent on open-ended answers to the questionnaire, in order to limit potential bias introduced through default answers. We believe that our defined default answers represent highly likely scenarios, and have thus minimal to no impact on the validity of our findings.

External Validity. The generalization of our findings is a main threat related to external validity. The lack of policies from app stores other than Apple's App Store and Google's Play Store could pose an external threat to validity regarding the generalizability of our findings. To reduce this threat, we targeted the most popular apps to ensure that our privacy policies collection is representative to the practical use of the apps, noting that the analyzed stores in our study represent the largest global markets. Another threat is related to the scope of our study, which is limited to privacy policies written in English. While such policies constitute the majority in the global market [20], analyzing policies written in other languages is beneficial to improve this external validity. We believe that our findings provide a significant and representative overview of the common practices in mobile app development.

V. Related Work

There is substantial work on privacy requirements in the RE literature [26], [34]. However, the existing work does not focus specifically on privacy policies of mobile apps. In the following, we review the relevant work on compliance of mobile apps' privacy policies in areas such as mobile health and internet of things (IoT), considering different regulations, including the Health Insurance Portability and Accountability Act (HIPAA) [35] and GDPR [3].

Sunyaev et al. [36] conducted a large scale analysis over 600 apps in the mobile health domain. Similar studies, but at a smaller scale, were performed by Zapata et al. [37] (who analyzed 24 apps) and Benjumea et al. [38] (who analyzed 31 apps). These studies highlighted a global lack of information ranging from absence of the policy itself to missing information and discrepancies against the applicable regulation (HIPAA in this case). That said, these studies are relatively

old and were conducted at the time when privacy and data protection considerations were less standardized. As for IoT services, Paul et al. [15] proposed an assessment framework based on GDPR requirements. Over 94 privacy policies, the authors deemed most of the examined policies insufficient to address certain GDPR requirements. More recently, Xie et al. [19] conducted a systematic mapping study on the compliance of privacy policies of virtual personal assistant apps. After conducting the analysis of all skills (e.g., voice activated apps) listed on the Alexa Store, the authors concluded that a large number of skills suffer from non-compliance issues concerning privacy practices.

In another research strand, some frameworks and tools have been proposed for assuring the consistency between policies and actual app code. For instance, Slavin et al. [39] proposed a framework to detect misalignments between phrases in privacy policies and the code mapped to such phrases. In their evaluation, they detected 341 potential violations on 477 Android apps. Similarly, Hosseini et al. [40] proposed a method to help developers select appropriate unambiguous terms when sharing their data protection practices.

More relevant to our study is the PPChecker tool by Yu et al. [17], which automatically identifies five potential problems leading to non-compliant privacy policies. PPChecker has been evaluated on 2500 privacy policies of real apps. The results showed that about 1850 of the analyzed privacy policies contained at least one of the five problems. Verderame et al. [29] proposed a similar tool (3PDroid) to assess the compliance of privacy policies according to Google Play privacy guidelines. The results suggested that less than 1% of the analyzed apps fully complies with the Google Play privacy guidelines. More recently, Xiang et al. [20] proposed PolicyChecker, a framework similar to *CompA*ι, which evaluates the state of GDPR completeness violations in mobile apps' privacy policies. The tool is based on rules and semantic analysis but, unlike *CompA*ι, it lacks the legal domain knowledge integrated in the compliance checking process. The authors evaluated PolicyChecker on 163 068 policies collected for apps in the UK Google Play Store. Their results showed that 99.3% of the analyzed privacy policies were incomplete according to GDPR, with at least one GDPR requirement not fulfilled. Tan and Song [30] investigated problematic declarations related to user privacy and data collected by third parties. They proposed PTPDroid, an automated tool that helps uncovering undesired disclosure of personal data to third parties. Over 1000 real-world apps, the results revealed that such violations are prevalent in mobile apps with only 61 (3.8%) apps that provided specific details on third parties and 23 (2.3%) that did not share personal data to third parties.

Compared to the aforementioned work, this study presents a deeper and more comprehensive view of GDPR implementation in privacy policies of mobile apps across different categories from two global app stores. Although our dataset is smaller, the legal background of *CompA*ι makes it a strong tool for this kind of analysis. In addition, our analysis focuses on mainstream apps, which are assumed to be more compliant with GDPR than numerous (often low-quality and less-usable) apps in the large Android datasets previously used in the literature.

VI. Conclusion

In this paper, we have investigated the state of practice regarding the implementation of GDPR requirements in privacy policies of mobile apps, focusing on the most common incompleteness issues, their distribution across app categories and stores, and the correlation (if any) with the coverage of the related GDPR concepts in the research literature. More specifically, we have presented a study on the completeness of mobile apps' privacy policies based on 470 mined policies for mobile apps (233 from Apple App Store and 237 from Google Play Store). We used an existing tool (*CompA*ι) to analyze the completeness of the privacy policies against 19 criteria elicited from GDPR and presented in previous work. These criteria concern various GDPR concepts, e.g., the rights of individuals (app users) on their personal data.

Our results show that about 92% of the analyzed policies violate at least one mandatory or optional completeness criterion. More concretely, the most common violated criteria are related to the mention of data subject rights, transfer outside the EU, and the consequences of a failure to provide personal data. Our observations reveal similar compliance practices across the different app stores. Finally, our study concludes that the most frequently violated criteria pertain to GDPR concepts that remain under-explored in the software engineering literature. This calls for more research on certain GDPR concepts which would have positive impact on the compliance practices in mobile development.

In the future, we plan to extend our study with in-depth analysis of the privacy labels provided next to mobile apps. The goal is to identify possible inconsistencies between what a privacy policy states and what the mobile app promises through privacy labels. We also plan to perform a follow-up longitudinal study on the evolution of the state of practice of GDPR implementation in mobile apps privacy policies.

Acknowledgment.

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

References

[1] L. Ceci, "Google play: number of available apps 2017–2024," 2024. [Online]. Available: https://web.archive.org/web/20250621160700/https://www.statista.com/statistics/266210/number-of-available-applications-in-the-google-play-store/

[2] ——, "Apple app store: number of available apps as of q2 2024," 2024. [Online]. Available: https://web.archive.org/web/20250621161732/https://www.statista.com/statistics/779768/number-of-available-apps-in-the-apple-app-store-quarter/

[3] The European Parliament and the Council of the European Union, "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)," 05 2016.

[4] T. T. Nguyen, M. Backes, N. Marnau, and B. Stock, "Share first, ask later (or never?) studying violations of GDPR's explicit consent in android apps," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 3667–3684.

[5] S. Zhang, H. Lei, Y. Wang, D. Li, Y. Guo, and X. Chen, "How android apps break the data minimization principle: An empirical study," in *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2023, pp. 1238–1250.

[6] X. Zhang, X. Wang, R. Slavin, T. Breaux, and J. Niu, "How does misconfiguration of analytic services compromise mobile privacy?" in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, 2020, pp. 1572–1583.

[7] T. T. Nguyen, M. Backes, and B. Stock, "Freely given consent? studying consent notice of third-party tracking and its violations of GDPR in android apps," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 2369–2383.

[8] K. Zhao, X. Zhan, L. Yu, S. Zhou, H. Zhou, X. Luo, H. Wang, and Y. Liu, "Demystifying privacy policy of third-party libraries in mobile apps," in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*, 2023, pp. 1583–1595.

[9] J. Krämer, "The death of privacy policies: How app stores shape GDPR compliance of apps," *Internet Policy Review*, vol. 13, no. 2, 2024.

[10] Competition and Markets Authority (CMA), "Mobile ecosystems market study final report," 2022. [Online]. Available: https://www.gov.uk/government/publications/mobile-ecosystems-market-study-final-report

[11] K. Kollnig, A. Shuba, M. Van Kleek, R. Binns, and N. Shadbolt, "Goodbye tracking? impact of iOS app tracking transparency and privacy labels," in *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, ser. FAccT '22, 2022, p. 508–520.

[12] M. Ceci, D. Bianculli, and L. Briand, "Defining a model for content requirements from the law: an experience report," in *32nd IEEE International Requirements Engineering 2024 conference*, 2024.

[13] O. Amaral, M. I. Azeem, S. Abualhaija, and L. C. Briand, "NLP-based automated compliance checking of data processing agreements against GDPR," *IEEE Transactions on Software Engineering*, vol. 49, no. 9, pp. 4282–4303, 2023.

[14] M. I. Azeem and S. Abualhaija, "A multi-solution study on GDPR ai-enabled completeness checking of DPAs," *Empirical Software Engineering*, vol. 29, no. 4, p. 96, 2024.

[15] N. Paul, W. B. Tesfay, D.-K. Kipker, M. Stelter, and S. Pape, "Assessing privacy policies of internet of things services," in *ICT Systems Security and Privacy Protection*, 2018.

[16] D. Torre, S. Abualhaija, M. Sabetzadeh, L. C. Briand, K. Baetens, P. Goes, and S. Forastier, "An AI-assisted approach for checking the completeness of privacy policies against GDPR," in *28th IEEE International Requirements Engineering Conference, RE 2020*, 2020, pp. 136–146.

[17] L. Yu, X. Luo, J. Chen, H. Zhou, T. Zhang, H. Chang, and H. K. N. Leung, "PPChecker: Towards accessing the trustworthiness of android apps' privacy policies," *IEEE Transactions on Software Engineering*, vol. 47, no. 2, pp. 221–242, 2021.

[18] O. Amaral, S. Abualhaija, D. Torre, M. Sabetzadeh, and L. C. Briand, "Ai-enabled automation for completeness checking of privacy policies," *IEEE Transactions on Software Engineering*, vol. 48, no. 11, pp. 4647–4674, 2022.

[19] F. Xie, Y. Zhang, C. Yan, S. Li, L. Bu, K. Chen, Z. Huang, and G. Bai, "Scrutinizing privacy policy compliance of virtual personal assistant apps," in *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, 2023.

[20] A. Xiang, W. Pei, and C. Yue, "Policychecker: Analyzing the GDPR completeness of mobile apps' privacy policies," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 2023, p. 3373–3387.

[21] O. Amaral, S. Abualhaija, and L. C. Briand, "Compai: A tool for GDPR completeness checking of privacy policies using artificial intelligence," in *2024 39th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2024.

[22] O. Amaral, S. Abualhaija, N. Sannier, M. Ceci, and D. Bianculli, "Artifact associated with "GDPR compliance in privacy policies of mobile apps: An overview of the state-of-practice"," https://doi.org/10.5281/zenodo.15675391, 2025.

[23] ——, "Online annex," https://doi.org/10.6084/m9.figshare.27918729, 2025.

[24] S. Ghanavati, A. Rifaut, E. Dubois, and D. Amyot, "Goal-oriented compliance with multiple regulations," in *Proceedings of 22nd IEEE International Conference on Requirements Engineering*, 2014.

[25] S. Kununka, N. Mehandjiev, and P. Sampaio, "A comparative study of android and iOS mobile applications' data handling practices versus compliance to privacy policy," in *Privacy and Identity Management. The Smart Revolution - 12th IFIP WG 9.2, 9.5, 9.6/11.7, 11.6/SIG 9.2.2 International Summer School, Ispra, Italy, September 4-8, 2017, Revised Selected Papers*, 2017.

[26] J. Bhatia, M. C. Evans, and T. D. Breaux, "Identifying incompleteness in privacy policy goals using semantic frames," *Requirements Engineering*, vol. 24, no. 3, 2019.

[27] M. Fan, L. Yu, S. Chen, H. Zhou, X. Luo, S. Li, Y. Liu, J. Liu, and T. Liu, "An empirical evaluation of GDPR compliance violations in android mhealth apps," in *2020 IEEE 31st International Symposium on Software Reliability Engineering (ISSRE)*, 2020, pp. 253–264.

[28] O. Amaral Cejas, N. Sannier, S. Abualhaija, M. Ceci, and D. Bianculli, "GDPR-relevant privacy concerns in mobile apps research: A systematic literature review," 2024, https://hdl.handle.net/10993/62753.

[29] L. Verderame, D. Caputo, A. Romdhana, and A. Merlo, "On the (un)reliability of privacy policies in android apps," in *2020 International Joint Conference on Neural Networks (IJCNN)*, 2020, pp. 1–9.

[30] Z. Tan and W. Song, "Ptpdroid: Detecting violated user privacy disclosures to third-parties of android apps," in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*, 2023, pp. 473–485.

[31] The European Parliament and the Council of the European Union, "Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications)," 07 2002.

[32] B. Meyer, *Handbook of Requirements and Business Analysis*. Springer, 2022.

[33] W. Maalej, Z. Kurtanović, H. Nabil, and C. Stanik, "On the automatic classification of app reviews," *Requirements Engineering*, vol. 21, pp. 311–331, 2016.

[34] C. Negri-Ribalta, M. Lombard-Platet, and C. Salinesi, "Understanding the GDPR from a requirements engineering perspective—a systematic mapping study on regulatory data protection requirements," *Req. Eng.*, vol. 29, no. 4, pp. 523–549, 2024.

[35] United States Congress, "Health Insurance Portability & Accountability Act. 104th Cong., Public Record 104-191 (1995-1996)," 08 1996.

[36] A. Sunyaev, T. Dehling, P. L. Taylor, and K. D. Mandl, "Availability and quality of mobile health app privacy policies," *Journal of the American Medical Informatics Association*, vol. 22, 2014.

[37] B. C. Zapata, A. Hernández Niñirola, J. L. Fernández-Alemán, and A. Toval, "Assessing the privacy policies in mobile personal health records," in *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2014.

[38] J. Benjumea, J. Ropero, O. Rivera-Romero, E. Dorronzoro-Zubiete, and A. Carrasco, "Assessment of the fairness of privacy policies of mobile health apps: Scale development and evaluation in cancer apps," *JMIR Mhealth Uhealth*, 2020.

[39] R. Slavin, X. Wang, M. B. Hosseini, J. Hester, R. Krishnan, J. Bhatia, T. D. Breaux, and J. Niu, "Toward a framework for detecting privacy policy violations in android application code," in *Proceedings of the 38th International Conference on Software Engineering*, 2016.

[40] M. B. Hosseini, T. D. Breaux, R. Slavin, J. Niu, and X. Wang, "Analyzing privacy policies through syntax-driven semantic analysis of information types," *Information and Software Technology*, vol. 138, 2021.