How do App Vendors Respond to Subject Access Requests? A Longitudinal Privacy Study on iOS and Android Apps

Jacob Leon Kröger
kroeger@tu-berlin.de
TU Berlin, Weizenbaum Institute
Berlin, Germany

Jens Lindemann
lindemann@informatik.uni-hamburg.de
University of Hamburg
Hamburg, Germany

Dominik Herrmann
dh.psi@uni-bamberg.de
University of Bamberg
Bamberg, Germany

ABSTRACT

EU data protection laws grant consumers the right to access the personal data that companies hold about them. In a first-of-its-kind longitudinal study, we examine how service providers have complied with subject access requests over four years. In three iterations between 2015 and 2019, we sent subject access requests to vendors of 225 mobile apps popular in Germany. Throughout the iterations, 19 to 26 % of the vendors were unreachable or did not reply at all. Our subject access requests were fulfilled in 15 to 53 % of the cases, with an unexpected decline between the GDPR enforcement date and the end of our study. The remaining responses exhibit a long list of shortcomings, including severe violations of information security and data protection principles. Some responses even contained deceptive and misleading statements (7 to 13 %). Further, 9 % of the apps were discontinued and 27 % of the user accounts vanished during our study, mostly without proper notification about the consequences for our personal data. While we observe improvements for selected aspects over time, the results indicate that subject access request handling will be unsatisfactory as long as vendors accept such requests via email and process them manually.

CCS CONCEPTS

• **Security and privacy** → **Social aspects of security and privacy; Privacy protections**; • **Social and professional topics** → **Privacy policies**.

KEYWORDS

GDPR, compliance, subject access request, smartphone, mobile app, privacy

ACM Reference Format:

Jacob Leon Kröger, Jens Lindemann, and Dominik Herrmann. 2020. How do App Vendors Respond to Subject Access Requests? A Longitudinal Privacy Study on iOS and Android Apps. In *The 15th International Conference on Availability, Reliability and Security (ARES 2020), August 25–28, 2020, Virtual Event, Ireland*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3407023.3407057>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
ARES 2020, August 25–28, 2020, Virtual Event, Ireland
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8833-7/20/08.
<https://doi.org/10.1145/3407023.3407057>

1 INTRODUCTION

Many mobile apps collect personal data about their users and share it with third parties, such as analytics services and ad networks [3]. Given the increasing number of apps and the vast amount of data collected – often including data with no apparent relevance to the advertised functionality [30] – it is challenging for smartphone users to keep track of the data that app vendors hold about them.

As the preceding Directive 95/46/EC, the EU General Data Protection Regulation (GDPR [28]), which came into effect on 25 May 2018, affords individuals various rights over their personal data, including the rights to access, rectification, and erasure (Art. 15–17 GDPR). These rights allow smartphone users to demand transparency and regain control over the personal data collected by mobile apps. Numerous empirical investigations have revealed, however, that data subjects face various impediments in exercising their rights, with many data controllers completely refusing to comply with the law [1, 10, 24, 34, 37].

In this paper, we contribute to this line of research by presenting the results of a four-year undercover field study. While most existing work has investigated other types of public and private service providers, we focus on developers and vendors of mobile apps. This industry deserves scrutiny, as smartphones have become the primary device for many users [5].

Contribution. Until now, related studies (cf. Sect. 2) have evaluated data controller behavior based on one-time snapshot data. In contrast, we analyze how the behavior and compliance of a fixed set of app vendors change over time. Our longitudinal study includes observations both before and after the GDPR enforcement date. In three iterations between the years 2015 and 2019, we attempted to exercise the right of access with 225 vendors of mobile apps that were popular in Germany at the beginning of our study. Additionally, our first and our last round of inquiries included a question on third-party data sharing practices. In this paper, we examine in detail the obstacles we encountered as well as the veracity and completeness of the responses received. Secondly, we describe the measures that app vendors apply to verify the requester’s identity and how they choose to protect the confidentiality of transmitted personal information. Thirdly, we shed light on the issue of *dissolution of personal data* that results from vendors going out of business, apps being discontinued, and stale user accounts being deleted without prior notice.

Outline. The remainder of this paper is structured as follows. First, we review related work in Sect. 2. In Sect. 3, we describe our data collection process. Then, we present our methodology for interaction with app vendors in Sect. 4. Section 5 summarizes the results of our study, a discussion of which is provided in Sect. 6. A

reflection on ethical aspects and limitations of our study follows in Sect. 7 before we conclude the paper in Sect. 8.

2 RELATED WORK

Numerous studies have examined how data controllers react to data subject requests in practice. In these studies, test requests for data access, erasure, and/or portability were either sent to specific types of organizations, such as CCTV operators [34], smartphone app vendors and website owners [10], online tracking companies [37], or to a broad range of public and private sector organizations [1, 2, 18, 19, 24, 42, 43].

In [24], investigations were carried out in ten different EU member states and, in addition to the quantitative evaluation, a detailed case-by-case assessment is provided for individual data controllers. With the exception of [1], [37] and [42], the above-referenced studies were conducted prior to the GDPR coming into force. Also, in contrast to our study, none of the existing publications offers a multi-year longitudinal evaluation over a constant set of data controllers. Since 2010, the French Association of Data Protection Officers¹ has published a yearly report on the right of access, including results from probing 150 to 200 service providers [1]. However, their list of examined data controllers is newly compiled each year, which means that behavioral changes and trends within individual organizations are not observed.

In line with our findings, previous studies report various anomalies, poor practices, and severe compliance issues on the side of the data controllers, resulting in low rates of satisfying responses to data subject requests [1, 10, 19, 24, 34, 37]. Besides a widespread unwillingness or inability to provide the requested data in time [2, 10, 34, 37], researchers have observed the use of inappropriate file formats for the transfer of personal data [42], instances of personal information leakage to impostors [7, 10], issues concerning the language and clarity of interactions [24], and unsafe procedures to authenticate data subjects [4, 25]. In some cases, researchers were not even able to locate the contact details of data controllers, rendering any request submission impossible from the outset [17, 24].

While privacy and security aspects of mobile apps have received a lot of research attention in recent years [3, 15, 21, 30, 31], little is yet known about the behavior of app vendors when it comes to fulfilling data subject requests, especially not with the GDPR being in effect.

3 DATA COLLECTION

Our objective was to obtain a comprehensive sample of apps from the iOS and Android app stores, considering a diverse set of app types and vendors. We did not want to focus on the most popular apps, because those might exhibit abnormal behavior as a result of being in the spotlight. On the other hand, we did not want to analyze outdated and niche apps with virtually no users.

To compile a suitable sample of popular mobile apps, we used market research information provided by AppAnnie (<https://www.appannie.com>). AppAnnie monitors the download counts of apps on the app stores of Google and Apple. Specifically, in August 2015,

Vendor's country of residence

?	rest EU	Germany	non-EU countries
8	44	78	95

Operating system

iOS	Android
105	120

App categories (5 apps each)

iOS: Business, Catalogue, Education, Entertainment, Finance, Food, Games, Health, Lifestyle, Medical, Music, Navigation, News, Photo, Productivity, Reference, Social Networking, Sports, Travel, Utility, Weather	Android: Books, Business, Comics, Communication, Education, Entertainment, Family, Finance, Games, Health, Lifestyle, Media, Medical, Music, News, Photography, Productivity, Shopping, Social, Sports, Tools, Transportation, Travel, Widgets
---	---

Figure 1: Popular apps dataset overview; figures refer to number of apps out of 225 apps in total

we obtained AppAnnie's ranking lists containing the most popular apps in Germany. In current mobile operating systems, apps are assigned categories that describe their primary function or subject matter (e. g., Education, Music, or Health). For each of these app store categories, AppAnnie provides a ranked list (of varying length) that contains the most popular apps in the respective category according to the number of downloads within a fixed period. There were 24 categories for Android and 21 categories for iOS apps.

We randomly sampled apps from each of AppAnnie's category lists, subject to the following conditions. We picked a random app, installed it, and checked whether it offered users the possibility to create a personal account or to enter personal information in some other way. If an app did not meet this requirement, we picked another random app from the respective list. We also skipped an app if our sample already contained another app from the same vendor. We sampled apps until we had collected *five apps per category*, amounting to a total sample size of 225 apps (120 Android apps, 105 iOS apps). As a result, apps span a wide number of popularity ranks (avg. rank within a category: 134, std. deviation: 111, min. rank: 1, max. rank 407).

In the following, we characterize the dataset (see also Fig. 1). The largest proportion of the selected apps' vendors is based in Germany (35 %). Vendors located across other EU member states and outside of the EU account for 20 % and 42 %, respectively. For 3 % of the apps in our dataset, we could not find any information on the vendor's residence.

Like the German Federal Data Protection Act [26] and the now superseded EU Directive 95/46/EC [29], the GDPR can also apply to data controllers established in non-EU countries [40]. The GDPR explicitly states that it applies to organizations "not established in the Union, where the processing activities are related to the offering of goods or services, irrespective of whether a payment of the data subject is required, to such data subjects in the Union" (Art. 3 GDPR). While it may be difficult to exercise and enforce data protection rights against data controllers in foreign countries, we include app vendors based outside the EU in our study to compare their behavior with that of domestic vendors (see Sect. 5.5).

¹ Official name in French: Association Française des Correspondants à la Protection des Données à Caractère Personnel (AFCDP)

4 METHODOLOGY FOR INTERACTION WITH THE APP VENDORS

The apps in our sample were downloaded and installed on a smartphone with the corresponding operating system. Where it was possible, a new user account was created by signing up via email (76 %), Facebook (10 %), or Google (10 %); preference was given to signup via email if available. The remaining 4 % of the apps were populated with personal data without account registration.

For account creation and all other app interactions, we acted disguised as an ordinary consumer, always using the same identity of a volunteer (male, of legal age, German nationality). The app vendors were not informed in advance that they would be the subject of a study. We will reflect upon the necessity and ethical aspects of this covert approach in Sect. 7. Using the identity of a real person was necessary to overcome identification barriers (see Sect. 5.4).

After installation, we interacted with each app for about ten minutes, entering as much personal data as possible. Subsequently, over a period of four years, we issued three rounds of data subject requests to the app vendors (from now on referred to as R1, R2, and R3). **R1** requests were issued in November 2015, several weeks after the apps were installed. **R2** requests followed in March 2018, two months before the May 25 GDPR enforcement date. **R3** requests were sent in August 2019, more than one year after GDPR had come into effect.

The subject access requests were written in English or German (depending on the language of the app). We sent them to the vendors' email address specified on the respective page in the app marketplace. If no contact details were provided there, we searched for an email address on the vendor's website or, if it was the only option available on the website, submitted the request via a contact form. Ultimately, we were able to find a point of contact for all apps for R1. We updated our list of vendor email addresses once again before sending out our subject access requests in R2 and R3, respectively.

To account for changing external circumstances and to avoid being recognized as researchers, we deliberately used a different request text for each round of inquiries. While the right of access was the main focus of our study and, therefore, all three rounds contained data access requests, we additionally requested information about data sharing practices in R1 and R3. For R1, we chose a short and informal inquiry text, comprising only seven sentences. The texts for R2 and R3 were more elaborate. They included references to relevant data protection laws (GDPR, EU Directive 95/46/EC, and the German Federal Data Protection Act) as well as a warning that the responsible data protection authorities would be notified in the case of no response (which, in fact, we did not do).

A more formal approach was chosen for R2 and R3 because the introduction of the GDPR and the wide-ranging preparations for the new law aroused increased media attention and public awareness for data subject rights and data protection issues in general [9, 32]. In addition, since 2018, more and more self-help tools and templates for GDPR requests have become available through websites like *datarequests.org* and *gdpr.eu*. We, therefore, assumed that ordinary smartphone users were now better equipped and thus more likely

to make formal requests with legal references than they were in 2015, the year of R1.

To test how a reminder would affect the response rate, we sent a follow-up email to all vendors who had not replied to our request in R3. In the reminder (sent 85 days after the R3 request) we merely stated that we had not received a response to our previous message. Except where specifically indicated, the responses to this reminder will be ignored in the evaluation section, i. e., vendors who only responded *after* receiving a reminder will be counted as non-responders in R3.

5 ANALYSIS AND RESULTS

The received responses were subjected to a qualitative content analysis as proposed by Strauss and Corbin [35]. Following an open coding approach, we built and refined a codebook based on the data collected in R1. First, one coder went through the received responses, creating and applying an initial set of codes to all encountered findings. An independent second coder subsequently applied the resulting list of codes to the same set of responses. The initial coding yielded a Fuzzy Kappa value of 0.819, which indicates a high degree of agreement between the two coders.² Subsequently, the two coders consolidated the codebook, reassessed all responses, and resolved all conflicts. The emerging code categories and overarching themes were used to determine the foci of our paper, including the reachability and responsiveness of app vendors, the provision of the requested data, statements made on data sharing practices, occurring technical and communication problems as well as the security of data transmissions. A single coder applied the consolidated codebook to the responses in R2 and R3. The number of assigned codes per app and iteration varies between 1 and 9.

We have released a sanitized dataset containing our findings for every app [12]. A summary of our results is provided in Table 1, at the end of this section. In the published dataset and in the remainder of this paper, the apps in our sample are referred to using the pseudonyms **A1 to A225**. Where applicable, the corresponding round of requests is indicated through a subscript. **A5_{R1}**, for instance, designates the response from the vendor of App 5 to our first request. Section 5.1 provides an overview of our results, followed by an in-depth analysis of the obtained responses in Sects. 5.2–5.7. Unless specifically stated otherwise, percentage values refer to the whole sample of 225 apps. We will introduce capitalized labels to refer to particular subsets of apps.

5.1 Overview of the Received Responses

While the majority of the investigated app vendors reacted to our inquiries in some way (**RESP**: 78 %_{R1}, 81 %_{R2}, 74 %_{R3}), many did not respond to our requests (22 %_{R1}, 16 %_{R2}, 22 %_{R3}) and some were completely unreachable via email, yielding us only a delivery failure notification (0 %_{R1}, 3 %_{R2}, 4 %_{R3}). The reminder we sent after our last round of requests increased the response rate for R3 slightly from 74 % to 80 %, showing that exercising data protection rights can sometimes require perseverance.

² In contrast to Cohen's Kappa, Fuzzy Kappa by Kirilenko and Stepchenkova [11] can handle codebooks with codes that are not mutually exclusive, which matches our setting.

Out of all examined app vendors, only 14 %_{R1}, 44 %_{R2}, and 37 %_{R3} sent us a response that contained personal data (**DATA**). For reasons outlined in Sect. 5.2, however, many responses were unintelligible and useless. Overall, we received a proper usable export of personal data from merely 12 % of all vendors in R1, and from a significantly greater but still underwhelming proportion of 38 % in R2 and 28 % in R3. Some responses even contained deceptive or misleading statements (7 %_{R1}, 13 %_{R2}, 11 %_{R3}), which will be elaborated on in Sect. 5.6.

In each iteration of our study, the vast majority of responses arrived within five days (82 %_{R1}, 76 %_{R2}, 83 %_{R3}). Apart from a small number of exceptions, the rest arrived between day 6 and day 15 (14 %_{R1}, 22 %_{R2}, 11 %_{R3}). A minor group of outliers (2 %_{R1}, 0 %_{R2}, 2 %_{R3}) replied after more than 31 days, with one month being the time limit for responding to data protection rights requests as per Art. 12 GDPR. Before GDPR enforcement, neither the EU Directive 95/46/EC nor the German Federal Data Protection Act did prescribe such a specific time limit. However, according to legal commentaries by Wolff et al. [41, § 34 Rn. 104-106.1] and Müller-Glöße et al. [22, § 34 Rn. 1], the maximum appropriate response time ranged between two and four weeks. Figure 2 shows the frequency distribution of responses over time.

Out of all received responses (RESP), the proportion that fulfilled our subject access request (**OK**) by either containing a proper export of personal data or a credible statement that the data is no longer stored was 19 %_{R1}, 65 %_{R2}, and 55 %_{R3}. The others were recorded as incomplete responses. When taking all app vendors in our sample into account, including those who were unreachable or did not respond, the proportion of OK responses amounts to 15 %_{R1}, 53 %_{R2}, and 41 %_{R3}. An overview of our evaluation for all three rounds of subject access requests is provided in Fig. 3.

While we cannot precisely determine their individual influence, it can be assumed that both the introduction of the GDPR as well as the more formal and threatening tone of our inquiry in R2 and R3 (cf. Sect. 4) had an impact on the vendors' behavior. In [10], data provision inquiries yielded better response rates when they were written more formally and included explicit references to relevant data protection law.

In cases where data was provided (**DATA**), it was made available either by email (97 %_{R1}, 84 %_{R2}, 79 %_{R3}), by postal mail (0 %_{R1}, 9 %_{R2}, 2 %_{R3}), or through the app itself (3 %_{R1}, 7 %_{R2}, 19 %_{R3}). In most cases, personal data was provided in attachments in various file formats (namely *pdf*, *html*, *json*, *csv*, *jpg*, *png*, *docx*, and *txt*) and as plain text in the email body.

The following sections focus on particular aspects of the responses. First, we will consider deficiencies with format and content in Sect. 5.2. The degree to which the app vendors addressed our additional question on data sharing practices is examined in Sect. 5.3. After that, we consider the security of data in transit (Sect. 5.4) before we compare the responses in terms of their sender's residence (Sect. 5.5). Finally, we report on particularly intriguing aspects in Sects. 5.6 and 5.7, namely deceptive responses as well as the dissolution of personal data when vendors go out of business, apps are discontinued, and accounts are deleted unsolicitedly.

Figure 2: Frequency distribution of responses over time

Figure 3: Evaluation of subject access request responses of all apps (n = 225)

5.2 Insufficient Responses

Failing to fulfil our subject access request, some of the received responses (RESP, $n = 175_{R1}, 183_{R2}, 167_{R3}$) contained only the labels of collected data, e. g., “birth date”, but not our actual data values, e. g., “1977-03-09” (22 %_{R1}, 6 %_{R2}, 5 %_{R3}). The frequency of this phenomenon strongly decreased over the course of our study, with a particularly large drop between R1 and R2.

Even when actual data was transmitted (DATA, $n = 32_{R1}, 100_{R2}, 84_{R3}$), it was often unintelligible due to serious formatting errors or obscure data labels (9%_{R1}, 14%_{R2}, 24%_{R3}). In some of these cases,³ the data consisted of a continuous block of alphanumeric characters without headings, spaces, and line breaks. According to EU law, personal data should be made available “in a structured, commonly used and machine-readable format” (Art. 20 GDPR). Interestingly, despite the relatively short duration of app usage (10 to 15 minutes per app), we sometimes received data in large and confusing folder structures. The responses of A4_{R3} and A15_{R3}, for example, contained 17 folders and 27 individual files, respectively, but only a handful of relevant data points. Also, in several cases, data was only transferred as an image file (e. g., a screenshot of an internal database), which made it impossible to extract information via copy and paste (3%_{R1}, 16%_{R2}, 4%_{R3}).

In other cases, technical or communication problems resulted in relevant information and personal data not being transmitted in the first place. These problems include dead download links and empty email attachments (A11_{R3}, A29_{R3}, A33_{R3}, A146_{R3}), unclear or incomprehensible instructions for how to obtain the data (A120_{R3}, A150_{R3}, A151_{R3}, A225_{R3}), references to inaccessible privacy functions inside the app (A157_{R3}), malfunctioning authentication mechanisms (A87_{R3}, A140_{R2}, A178_{R3}, A181_{R1}, A202_{R3}), and cases where the responsibility for privacy-related inquiries was referred back and forth between departments or affiliated companies (A21_{R1}, A25_{R1}, A118_{R3}). In most of these incidents, pointing out the problem to the app vendor did not help. For example, we exchanged 24 emails with A151_{R3} in an attempt to clarify a technical issue before finally giving up because the responses had become repetitive. These observations are in line with those by Ausloos et al. [2] who found that subject access requests often lead to endless sequences of emails without resulting in the requested transfer of personal data.

Furthermore, some of the investigated app vendors appear to lack essential resources to receive and process subject access requests in a professional manner. For example, from A30_{R3}, we received only a “recipient inbox full” error message. Other vendors explicitly stated that they lack the time, technical means, or personnel to compile and transmit the requested data (A105_{R2}, A172_{R3}), to reply in time (A50_{R3}, A83_{R3}), or to read emails in their inbox (A48_{R3}, A199_{R3}). While the many inadequate responses to our requests indicate widespread unawareness and ignorance about legal duties, some responses revealed a specific deficiency in legal knowledge. The person writing A112_{R2}, for example, did not know that a right of access already existed in EU law before the GDPR. Also, after the commencement of GDPR in May 2018, some vendors stopped offering a login function (A42_{R3}, A58_{R3}, A93_{R3}) or even discontinued the entire app, stating explicitly their inability to comply with the new law (A62_{R3}).

There were also language barriers. Although our requests were written in the corresponding app’s default language (English or German), two vendors replied in Spanish (A108_{R3}, A201_{R3}) and one always replied in Korean (A149). A31_{R3}, A51_{R3} and A174_{R3} changed the language of communication from English to German during

our email correspondence without prior announcement. Another vendor used the online translation service Google Translate to communicate with us, which led to ambiguities and confusion (A150_{R3}). A206_{R2} replied in bad English, and the responses A165_{R1} and A189_{R1} were utterly incomprehensible. Beyond that, some vendors used technical jargon, such as the term “SDK”⁴ (A210_{R3}). While jargon does not render a response useless per se, ordinary users may have difficulty understanding it.

Astonishingly, three vendors stated that data exports are only provided for paid subscriptions, not for users of their free version (A89_{R2}, A137_{R1}, A202_{R2}). And, in another grave violation of the most basic data protection principles, A100_{R2} sent us personal data of *another* user, including the person’s full name, contact details, login credentials, app usage logs, and even location data. In one case, we received the response to our request from a foreign email address with no apparent relation to the app in question (A77_{R3}).

Finally, some responses were impolite or contained careless mistakes. For instance, A161_{R2} declared – without any explanation – to have “no interest in having [us] as a customer.” A193_{R3} and A224_{R3} replied without a formal salutation, A202_{R3} misspelled the name in the salutation, and both A13_{R3} and A151_{R3} sent us the same emails multiple times for no apparent reason. In a strange mistake, the support agent writing A89_{R3} even addressed us with *their* name. Although these incidents were no severe obstacles to communication, they do indicate a lack of care and rigor in answering our requests.

5.3 Data Sharing with Third Parties

Along with the request for data access, in R1 and R3 we also requested information on third-party data sharing practices (cf. Sect. 4). Specifically, we asked the app vendors to name any third-party recipients to whom they had disclosed our personal data. Only a minority of the received responses (RESP, $n = 175_{R1}, 167_{R3}$) addressed this question at all (41%_{R1}, 40%_{R3}).

Out of this subset (DSANS, $n = 72_{R1}, 66_{R3}$), a decreasing proportion stated that personal data is never being shared with third parties (63%_{R1}, 52%_{R3}). The remaining responses vary widely in terms of the length and detail in which data sharing practices are explained. While some state specifically what data is made available to which partners for what purpose (56%_{R1}, 37%_{R3}), others list only a few generic reasons for data sharing (18%_{R1}, 16%_{R3}) or potential data recipients without naming the respective purpose (11%_{R1}, 25%_{R3}). A few vendors disclose only the categories of data being shared (11%_{R1}, 6%_{R3}) or even refrain from providing any specifics beyond mentioning the existence of data sharing (4%_{R1}, 16%_{R3}).

Overall, only 36%_{R1} and 32%_{R3} of the received responses (RESP) adequately addressed our question by either assuring that no data is being shared with third parties or by providing a list of data recipients. And there is still room for doubt, even in these cases. In particular, the proportion of apps that explicitly mention third-party tracking (only 19%_{R1} and 14%_{R3} of DSANS) is far below what one should expect based on the results of large-scale empirical studies on the topic. Binns et al. [3], for example, examined the

³ Out of ethical considerations, we refrain from publishing screenshots and excerpts of the responses (cf. Sect. 7).

⁴ SDK is an abbreviation of “software development kit” (i. e., a collection of software development tools)

prevalence of third-party trackers on 959,000 Android apps and found that 90.4 % included at least one tracker host. Therefore, it is well conceivable that some of the responding app vendors withheld relevant information on their data sharing practices.

5.4 Identity Verification and Security of Transmitted Data

Before responding to a subject access request, data controllers must carefully verify the identity of the requester. In a recent study, Martino et al. [7] demonstrated that unauthorized third parties could fairly easily obtain personal data through fake subject access requests sent from a different email address. In our study, most of the app vendors that sent us a copy of personal data (DATA, $n = 32_{R1}, 100_{R2}, 84_{R3}$) did not ask for prior identity verification (84 %_{R1}, 85 %_{R2}, 76 %_{R3}). Note, however, that we sent our inquiries using the email address that we had used for account registration, which provides a minimal form of authentication by itself.

Only a minority of app vendors in DATA demanded certain pieces of identity-related information. We were asked to provide our birth date, address, and customer number (13 %_{R1}, 2 %_{R2}, 0 %_{R3}) as well as a copy of a utility bill, ID card, or driving license (0 %_{R1}, 5 %_{R2}, 5 %_{R3}). A problem with this type of authentication mechanism is that it often requires those seeking to exercise their rights to provide the data controller with additional personal information. This practice has been recognized as “contrary to the spirit of data protection law” [25]. Implementing a rather unusual method for authentication, A119_{R2} demanded from us to “log into your social media account and post for public viewing a list of random characters.” Surprisingly, in three cases (A5_{R2}, A8_{R2}, A157_{R2}), personal data was disclosed to us before we had even responded to an authentication request.

In accordance with a GDPR best practice recommendation (from Recital 63 [28]), several app vendors did not directly send data to us. Instead, they provided remote access to the data through a self-service system within the password-protected app (3 %_{R1}, 7 %_{R2}, 19 %_{R3} of DATA). For this practice, a positive trend was observed over the three iterations of our study, with a particularly sharp increase between R2 and R3.

Some of the responders (RESP, $n = 175_{R1}, 183_{R2}, 167_{R3}$) failed to use transport-layer encryption via TLS (Transport Layer Security) in at least one of their emails. Transport-layer encryption is a baseline security measure, which has been deemed mandatory for the transmission of personal data via email [36]. For our analysis, we checked whether all relevant *Received* headers within emails contained signs of TLS being used (indicated by “SMTPS” or the statement of the negotiated cipher suite). Here, again, the worst result was recorded in R1, where 13 % of all responses (RESP) were not encrypted. This figure improved and remained stable over R2 and R3 (both 3 %). A21_{R2}, and A75_{R1} even transmitted the requested personal data without encryption. Besides, as indicated by email metadata (again, *Received* headers), many app vendors handle customer support via third-party services, e. g., ticket system software operated by Zendesk, Salesforce, Helpscout, Mandrill and others. As a result, even vendors that are themselves based in the EU may disclose personal data to third parties in other jurisdictions. In these

cases, exercising one’s right to access may disperse one’s personal data to additional data processors.

Where data was transferred (DATA), two other security mechanisms we observed were the expiration of download links after a few days or weeks (0 %_{R1}, 2 %_{R2}, 10 %_{R3}) and the protection of email attachments with passwords (6 %_{R1}, 3 %_{R2}, 12 %_{R3}). Except for one case where the corresponding password was communicated via telephone (A225_{R1}), the access keys were sent via email, sometimes in the same email as the password-protected data (A28_{R3}, A211_{R3}). Finally, for security reasons, several vendors chose to send the data by postal mail (0 %_{R1}, 9 %_{R2}, 2 %_{R3}) – sometimes even by registered mail (A16_{R2}, A81_{R2}, A107_{R2}, A225_{R2}).

Our correspondence with the vendors indicates that the employed authentication mechanisms are not the result of rigorous attacker modeling. Given the prevalence of TLS on mail servers, they provide little to no additional security against passive eavesdroppers. More importantly, they cannot protect against an adversary that has already compromised a user’s email account. Many pieces of information that have been requested from us for identity verification can be extracted from previously exchanged emails. Moreover, adversaries that control the email account of a user receive data and accompanying passwords when these are sent via email. Such adversaries could also compromise a user’s account on the app itself via the commonplace email-based password reset functionality, either to steal data directly from within the app or to change the postal address to which the data is sent.

5.5 Comparison by Vendor’s Residence

In the following, we explore behavioral differences between the app vendors from our sample that are based in Germany (GER, $n = 78$), in other EU member states (EMS, $n = 44$), and in the rest of the world (WLD, $n = 95$). Apps for which the vendor’s residence could not be determined ($n = 8$) are not considered in this analysis.

Across all three iterations of our study, the proportion of vendors responding to our inquiry was always somewhat higher among GER (82 %_{R1}, 88 %_{R2}, 78 %_{R3}) than among EMS (75 %_{R1}, 77 %_{R2}, 77 %_{R3}), and WLD (75 %_{R1}, 77 %_{R2}, 69 %_{R3}). We will refer to these responses as **RESP_{GER}** ($n = 64_{R1}, 69_{R2}, 61_{R3}$), **RESP_{EMS}** ($n = 33_{R1}, 34_{R2}, 34_{R3}$) and **RESP_{WLD}** ($n = 71_{R1}, 73_{R2}, 66_{R3}$). Throughout R1 and R2, the proportion of responses containing an export of personal data was highest among RESP_{GER} (30 %_{R1}, 71 %_{R2}, 61 %_{R3}). RESP_{EMS} caught up towards the end of our study (21 %_{R1}, 62 %_{R2}, 62 %_{R3}), leaving the rate of data transmissions among RESP_{WLD} far behind (7 %_{R1}, 38 %_{R2}, 38 %_{R3}). We name these data containing responses **DATA_{GER}** ($n = 19_{R1}, 49_{R2}, 37_{R3}$), **DATA_{EMS}** ($n = 7_{R1}, 21_{R2}, 21_{R3}$) and **DATA_{WLD}** ($n = 5_{R1}, 28_{R2}, 25_{R3}$).

Where sufficient for comparison, only the percentage figures from the most recent inquiry (R3) will be shown below. If the previous rounds of requests yielded deviating results, we will explicitly mention them.

In terms of OK responses (a usable export of personal data or a credible statement that no personal data is being stored), the three geographical subsets followed a similar development. The results show a sharp improvement between R1 and R2 followed by deterioration in R3 (GER: 51 %_{R3}, EMS: 45 %_{R3}, WLD: 31 %_{R3}). In all three runs, GER had the highest rate of OK responses; the lowest rate

was consistently found among WLD. Contributing to this discrepancy, Germany-based vendors scored better on the structure and intelligibility of transmitted data. Due to a lack of these qualities, 24 %_{R3} and 33 %_{R3} of the data exports provided by DATA_{WLD} and DATA_{EMS} were useless, respectively. Among DATA_{GER}, this figure was relatively low (19 %_{R3}).

And while the geographical subsets had comparable response rates to our question on data sharing practices in R1 (see Sect. 5.3), GER achieved by far the best result in our last round of requests (49 %_{R3}), followed by EMS (25 %_{R3}) and then WLD (17 %_{R3}).

Averaging over all responses received in our three rounds of inquiries, the vendors residing in Germany took longer to respond (5.5 days) than the vendors based in other EU member states (3.2 days) and outside of the EU (2.9 days). However, measured against legally prescribed time constraints (see Sect. 5.1), all of these average durations are acceptable.

The Germany-based vendors exhibited substantially inferior performance in the area of requester authentication. Compared to the already surprisingly small portion of vendors who conducted some form of identity check among DATA_{EMS} (48 %_{R3}) and DATA_{WLD} (28 %_{R3}), authentication requests among DATA_{GER} were much sparser (8 %_{R3}). In our first inquiry, this figure had been 21 %_{R1} for DATA_{GER} and 0 %_{R1} for DATA_{WLD}, but subsequently rose sharply for DATA_{WLD} and, contrary to our expectations, decreased considerably for DATA_{GER}. Before taking the sudden lead in R3, DATA_{EMS} had much weaker authentication results (14 %_{R1}, 10 %_{R2}).

At the end of our study, more apps from GER (10 %) and especially EMS (14 %) had been discontinued than from WLD (7 %). The proportion of disappearing user accounts (see Sect. 5.7) was distributed fairly evenly across the three groups.

5.6 Deceptive and Misleading Responses

Many of the responding app vendors (RESP, $n = 175_{R1}, 183_{R2}, 167_{R3}$) made misleading or demonstrably false statements. Several claimed that the app in question or our user account no longer existed – although we were still able to install, log in to, and use the app on our test devices (6 %_{R1}, 11 %_{R2}, 6 %_{R3}). Some vendors even contradicted themselves. For example, A34_{R3} repeatedly picked up specific points from our email request while pretending that the same message had never arrived. A188_{R3} claimed that no data related to our account had been collected. A213_{R3} stated that, as a matter of principle, copies of personal data would never be provided for security reasons. Both companies, however, had already sent us a copy of stored personal data a few days before. Moreover, some app vendors falsely claimed that they had already replied to our request earlier (A127_{R3}) or falsely promised to get back to us within a specified time (A2_{R3}, A33_{R3}, A59_{R2}, A102_{R2}, A118_{R3}, A126_{R3}).

Besides the capacity limits mentioned in Sect. 5.2, one reason for these observed inconsistencies and contradictions seems to be a lack of standardized processes for dealing with subject access requests. Our inquiries were sometimes processed and answered simultaneously by multiple different departments or employees (A4_{R3}, A21_{R2}, A37_{R3}, A106_{R3}, A153_{R3}, A188_{R3}, A191_{R3}, A197_{R3}, A213_{R3}). Additionally, in some cases where a data export was provided, we noticed that the export was not complete, i. e., some pieces

of personal data stored in the app had been omitted for unknown reasons (A47_{R2}, A158_{R1}, A173_{R1}).

We also suspected that some vendors merely *pretended* to be poorly reachable when they received subject access requests – while others actually *had* insufficient resources to process incoming emails. To confirm this hypothesis, we tested how the vendors that failed to respond to our requests reacted to non-privacy related inquiries. Using another (different) fake identity, we emailed the vendors who had not replied in R1 and R3, expressing interest in promoting their apps on a personal blog (R1) or YouTube channel (R3). Out of the group of initial non-responders, 31 %_{R1} and 22 %_{R3} replied to these dummy requests, many of them within a few hours, proving that their email inbox was in fact being monitored.

5.7 Discontinued Apps and Accounts

Manual checks revealed that prior to R2, 4 % of the initial 225 apps no longer existed. 9 % were gone prior to R3. Some of the corresponding vendors (**GONE**, $n = 10_{R2}, 21_{R3}$) still responded to our requests, informing us that their app had been discontinued (30 %_{R2}, 14 %_{R3}) or providing us with a copy of personal data that was still being stored (10 %_{R2}, 0 %_{R3}). The majority of them, however, did not address our request in their response or remained silent (50 %_{R2}, 57 %_{R3}) and some were not reachable via email anymore (10 %_{R2}, 29 %_{R3}).

Even among the vendors of the apps that could still be installed (**NOTGONE**, $n = 225_{R1}, 215_{R2}, 204_{R3}$), several responded that none of our submitted personal data is stored on their servers (8 %_{R1}, 24 %_{R2}, 17 %_{R3}). While some of these vendors (**NODATA**, $n = 17_{R1}, 51_{R2}, 34_{R3}$) explained that they only process data locally on the user’s device without ever having direct access to it (24 %_{R1}, 25 %_{R2}, 24 %_{R3}), most of them simply stated that a matching record did not exist in their database *anymore* (76 %_{R1}, 75 %_{R2}, 76 %_{R3}).

Based on manual login attempts, 27 % of all user accounts that we had initially created were no longer accessible at the end of our study, including discontinued apps (**GONE**). In a few astounding cases, our user account was deleted or deactivated in direct response to our inquiry (A12_{R3}, A93_{R2}, A133_{R2}, A161_{R2}, A209_{R2}) – either with unconvincing excuses or with no justification at all. A12_{R3}, for example, cited “credit card problems” as the reason, although we had never connected a credit card with any of the investigated apps.

All account discontinuations and erasures of personal data were carried out without any request or consent from our side. It is striking that we found no provisions in the corresponding privacy policies or terms of service that would explain the deletions. Also, apart from two vendors who gave notice that our account would be discontinued due to inactivity between R1 and R2 (A84, A106), we received no account deletion notification outside the responses to our requests. Neither a terminated app nor a deleted account necessarily implies, of course, that all of the user’s data was, in fact, entirely deleted by the app vendor and potential affiliates. This is well illustrated by a few cases where the vendor transferred our personal data despite having discontinued the app or our user account (A93_{R2}, A128_{R2}, A133_{R2}).

Among the vendors who provided an export of personal data (DATA, $n = 32_{R1}, 100_{R2}, 84_{R3}$), many included in their response an

unsolicited offer to delete our user account and the related data – especially in R2 and R3 (9 %_{R1}, 39 %_{R2}, 37 %_{R3}). It seems that the loss of a single user is often perceived as advantageous over bearing the effort and potential legal trouble associated with responding to data subject requests.

Table 1: Summary of results; figures refer to number of apps out of 225 apps in total

	R1 (2015)	R2 (2018)	R3 (2019)
Request Response (Sect. 5.1)			
Vendor unreachable (delivery failure)	1	6	9
Vendor did not respond	49	36	49
Vendor responded (RESP)	175	183	167
Response sufficient (OK)	33	119	92
Response contains data (DATA)	32	100	84
Response deceptive or misleading	13	24	19
Average response time [in days]	3.7	3.8	4.0
Data Sharing Practices (Sect. 5.3)			
Response addresses data sharing (DSANS)	72	n/a	66
Response states that data is shared	27	n/a	32
Response states that data is not shared	45	n/a	34
Sufficient information on data sharing	63	n/a	54
Requester Authentication (Sect. 5.4)			
Authentication mechanism in place	5	15	20
App login required	1	7	16
Identification document required	0	5	4
Identity-related information required	4	2	0
Other authentication mechanism	0	1	0
Security of Transmitted Data (Sect. 5.4)			
Data file(s) protected with password	2	3	10
Download link has expiration date	0	2	8
Data export via postal mail	0	9	2
TLS encryption failure	23	6	5
Existence of App & Account (Sect. 5.7)			
Discontinued apps (aggregated)	0	10	21
Discontinued user accounts (aggregated)	0	49	61

6 DISCUSSION

The results of our four-year undercover study reveal severe obstacles to exercising the right of access in the mobile app space. We found that the documented problems largely persist over time and have not substantially improved since 2014, when Herrmann and Lindemann obtained success rates that are comparable with the ones we obtained in 2019 [10].

In certain areas, we did observe positive trends. Between 2015 and 2019, the investigated vendors seem to have increased their willingness and ability to disclose stored personal information to data subjects upon request. Out of all examined app vendors, the proportion that managed to provide a proper export of personal data rose from 12 %_{R1} to 28 %_{R3}. Overall, the rate of acceptable responses (OK) to our subject access requests went up from 15 %_{R1} to 41 %_{R3} (see Sect. 5.1).

There was also a crucial improvement in the protection of private messages and transmitted personal information. Specifically, among all received replies (RESP, $n = 175_{R1}, 183_{R2}, 167_{R3}$), we observed

a reduction of unencrypted email correspondence from 13 %_{R1} to 3 %_{R3}. And we noticed a growing tendency among the vendors to not send data exports directly through email or postal mail but to make them available via a portal in the password-protected app (see Sect. 5.4). Among all vendors who sent us a copy of personal data (DATA, $n = 32_{R1}, 100_{R2}, 84_{R3}$), the prevalence of this practice increased from 3 %_{R1} to 19 %_{R3}.

Even with these improvements, however, getting access to one’s data is still a frustrating endeavor most of the time. While we observed substantial improvements between 2015 (R1) and 2018 (R2), most of the identified problems (e.g., low overall response rate, incomplete responses, deceptive and misleading statements, unintelligible data exports) remained static or even worsened between R2 and our last round of inquiries in 2019 (R3). Besides the impact of the emerging GDPR, the significant improvements between R1 and R2 can probably be attributed to a considerable extent to the less formally written request in R1 (see Sect. 4), reflecting previous findings from [10]. Although formal and informal requests have the same legal validity, our results suggest that successfully exercising data protection rights in practice may require a certain manner of expression and a degree of legal expertise, potentially going beyond the capacity of average smartphone users.

Why has there been no further improvement between R2 and R3? One explanation may be that app vendors have realized that the mostly underfunded and poorly organized supervisory authorities have insufficient resources to enforce the regulation and penalize misconduct [33, 38, 39]. As a result, there is little incentive for vendors to handle subject access requests professionally. The handling of our requests varied significantly, indicating a lack of well-established best practices in the mobile app industry.

We observed numerous alarming cases of disregard for the most basic information security and data protection principles. Among other failures, 76 to 85 % of the data exports (DATA) were provided without any attempt at verifying the requester’s identity. Where authentication mechanisms were in place, they were sometimes accidentally circumvented by the app vendors themselves. We also received exports of personal data via non-encrypted channels and were even given access to personal data not belonging to us (see Sect. 5.4).

Furthermore, we observed that many of the investigated apps were discontinued (9 %), with some vendors becoming completely unreachable during our study (4 %). The latter case can create unsettling uncertainty for data subjects, leaving them in doubt about the fate of their data. Given this “personal data dissolution,” users cannot get answers to the essential question, “Who knows what about me?” anymore. Answering this question, in turn, is a natural prerequisite for exercising the remaining ARCO rights (access, rectification, cancellation, opposition) in an informed manner, and thus an essential basis for informational self-determination [17]. While the case of unreachable and disappearing data controllers illustrates the problem well and potentially exacerbates it, any non-reaction or incomplete response to a subject access request can, of course, have the same consequence.

Personal data dissolution does not only result from vendors going out of business or apps being discontinued. Another source is user accounts that disappear without prior notice: 27 % of the initial accounts were gone by the end of our study. Of course, blocking or

deleting accounts after long periods of user inactivity is a legitimate concern of app providers. The practice also aligns with the GDPR’s storage limitation principle [8] and is often in the interest of the data subject by helping to prevent the accumulation of unused “zombie accounts” [23]. However, it seems evident that affected users should be informed in advance, be offered a final copy of their data, and be asked for explicit consent if the data is retained for further use. With few exceptions, these basic rules were not followed by the respective vendors in our study. In the vast majority of cases, we did not even receive an account deletion notification.

Overall, our results show very clearly that current processes for receiving and dealing with data subject requests have plenty of room for improvement. Existing research-based suggestions and recommendations for data controllers need to be compiled into actionable guidelines and distributed in a form that makes them digestible for small- and medium-sized organizations, such as app vendors. This includes, in particular, guidance on how to authenticate data subjects safely [4, 7, 25], how to transfer personal data [42], and how to facilitate the submission of requests [2, 10, 17]. It should be a key objective to replace the error-prone manual processing of data subject requests with automated and standardized interfaces for obtaining personal data and other privacy-related information. To incentivize the rapid and broad-scale adoption of such approaches, industry-specific legal requirements along these lines could be helpful.

Examining a fixed group of data controllers over a sustained period of time has proven to be suitable for this type of investigation and should be used more widely. The basic issues and challenges to exercising data protection rights in practice are now well understood. More longitudinal research, however, is needed to monitor whether progress is being made, to assess the impact of legal measures, as well as to track and refine emerging best practices.

Future studies and discussions on the privacy behavior of mobile apps should take into account that the stored personal data is not necessarily limited to recorded actions and inputs of the user. Patterns and correlations in collected data may leak additional information in a way that is not easily understood or anticipated by the user [13, 14, 16]. To achieve an adequate level of transparency, it would be necessary for users to be informed about such forms of data acquisition as well.

7 ETHICAL CONSIDERATIONS AND LIMITATIONS

As explained in Sect. 4, we interacted with the app vendors disguised as an ordinary app user. One crucial ethical aspect of undercover field studies is that resources of the studied subjects are consumed without their consent – namely, in our case, the working time of employees of the examined app vendors. However, as in related studies [1, 2, 10, 24, 37, 42], the covert approach was necessary to test the behavior of the vendors under realistic conditions and to prevent a research participation effect [20]. It can be assumed that employees of the investigated companies would have put more care and diligence into answering our inquiries if they had known about the nature of our study. For this reason, we consider undercover field research the only viable approach to investigate problems related to exercising data protection rights in practice.

It is not the purpose of this paper to name and shame individual companies. Therefore, we only referred to apps from our sample using pseudonyms. Moreover, we have decided not to inform the investigated companies about the study we have conducted to protect the responsible employees from negative consequences. For the same reason, after careful consideration and given the many inadequate responses to our inquiries, we also decided to publish neither the collected postal and email correspondence nor the received data records – not even in pseudonymized form. There would always be a residual risk that a vendor might recognize itself. In the interest of reproducibility and traceability, we did, however, release our *results dataset*, which consists of a sanitized and comprehensive table containing all findings for every app (represented by the numeric ID used in this paper) [12].

Our study has several limitations that need to be recognized. First, we want to stress that we have studied a particular sample of mobile apps. Our results, therefore, cannot be generalized to the whole population of apps, nor to all popular apps. Also, as explained in Sect. 4, we used a different request text for each iteration of our study, varying in length, formality, and in terms of the laws that were cited. Thus, the results obtained from R1, R2, and R3 are not unconditionally comparable.

Furthermore, the sending of three repeated requests, with long pauses between them, after only a few minutes of user activity on each app in the beginning, presumably deviates strongly from normal user behavior. This activity pattern may have aroused curiosity or suspicion among the investigated companies and potentially affected their response to our inquiries. Simulating an active user on 225 apps over a period of four years was not feasible with our available resources. On the other hand, this approach allowed us to measure the extent of *personal data dissolution* (cf. Sect. 6).

Other surrounding factors, such as the increasing public interest in data protection issues [9, 32] and several smartphone-related privacy scandals [6, 27], could also have affected the compliance with subject access requests in the mobile app industry. Therefore, despite the longitudinal nature of our study, we cannot measure the impact of the newly enacted GDPR in isolation.

8 CONCLUSION

In this longitudinal study, we investigated the obstacles faced by data subjects in exercising their right to access with mobile app vendors. Our results indicate positive trends in selected areas. The overall situation, however, is still as unsatisfactory today – with GDPR in force – as it was at the beginning of our study in 2015. Even in the second iteration of our study, in which the reactions to our subject access request were the most promising, we received an acceptable response from only 53 % of the examined vendors. In the first and third iteration, this figure was 15 % and 41 %, respectively. Response rates to questions we asked about third-party data sharing practices were similarly disappointing.

Besides a general lack of responsiveness, the observed problems range from malfunctioning download links and authentication mechanisms over confusing data labels and file structures to impoliteness, incomprehensible language, and even serious cases of carelessness and data leakage. It is evident from our results that there are no well-established and standardized processes for subject

access requests in the mobile app industry. Moreover, we found that many vendors lack the motivation to respond adequately. Many of the responses we received were not only completely insufficient, but also deceptive or misleading. Equally worrisome are cases of unsolicited dissolution of personal data, for instance, due to the apparently widespread practice of deleting stale accounts without prior notice.

With regard to the sensitive personal data that is regularly collected by mobile apps, this deficient and stagnating status quo is hardly tolerable. What could help to improve the situation is a combination of random compliance checks by authorities, coupled with better support for data controllers through industry-specific guidelines and best practices. In particular, there should be mandatory standard interfaces for providing data exports and other privacy-related information to data subjects, obviating the need for the manual processing of GDPR requests.

REFERENCES

- [1] Association Française des Correspondants à la protection des Données à caractère Personnel. 2020. Données personnelles - Index AFCDP 2020 du Droit d'accès. <https://afcdp.net/index-du-droit-d-acces/>
- [2] Jef Ausloos and Pierre Dewitte. 2018. Shattering One-Way Mirrors. Data Subject Access Rights in Practice. *Data Subject Access Rights in Practice (January 20, 2018)*. *International Data Privacy Law* 8, 1 (2018), 4–28.
- [3] Reuben Binns, Ulrik Lyngs, Max Van Kleek, Jun Zhao, Timothy Libert, and Nigel Shadbolt. 2018. Third party tracking in the mobile ecosystem. In *Proceedings of the 10th ACM Conference on Web Science*. 23–31.
- [4] Coline Boniface, Imane Fouad, Natalia Bielova, Cédric Lauradoux, and Cristiana Santos. 2019. Security Analysis of Subject Access Request Procedures. In *Annual Privacy Forum*. Springer, 182–209.
- [5] Christina Bröhl, Peter Rasche, Janina Jablonski, Sabine Theis, Matthias Wille, and Alexander Mertens. 2018. Desktop PC, tablet PC, or smartphone? An analysis of use preferences in daily activities for different technology generations of a worldwide sample. In *International Conference on Human Aspects of IT for the Aged Population*. Springer, 3–20.
- [6] Catalin Cimpanu. 2019. Another Facebook privacy scandal, this time involving its mobile analytics SDK. *ZDNet* (Feb. 2019). <https://www.zdnet.com/article/another-facebook-privacy-scandal-this-time-involving-its-mobile-analytics-sdk/>
- [7] Mariano Di Martino, Pieter Robyns, Winnie Weyts, Peter Quax, Wim Lamotte, and Ken Andries. 2019. Personal Information Leakage by Abusing the GDPR “Right of Access”. In *Proceedings of the Fifteenth USENIX Conference on Usable Privacy and Security (SOUPS’19)*. USENIX Association, USA, 371–386.
- [8] Majid Hatamian. 2020. Engineering Privacy in Smartphone Apps: A Technical Guideline Catalog for App Developers. *IEEE Access* 8 (2020), 35429–35445.
- [9] Alex Hern. 2018. What is GDPR and how will it affect you? *The Guardian* (May 2018). <https://www.theguardian.com/technology/2018/may/21/what-is-gdpr-and-how-will-it-affect-you>
- [10] Dominik Herrmann and Jens Lindemann. 2016. Obtaining personal data and asking for erasure: do app vendors and website owners honour your privacy rights?. In *GI Sicherheit 2016*. Gesellschaft für Informatik e.V., Bonn, 149–160. arXiv:1602.01804
- [11] Andrei P. Kirilenko and Svetlana Stepchenkova. 2016. Inter-Coder Agreement in One-to-Many Classification: Fuzzy Kappa. *PLOS ONE* 11, 3 (03 2016), 1–14.
- [12] Jacob Leon Kröger, Jens Lindemann, and Dominik Herrmann. 2020. Subject Access Request response data - 105 iOS and 120 Android apps. <https://doi.org/10.14279/depositonce-10338>
- [13] Jacob Leon Kröger, Otto Hans-Martin Lutz, and Florian Müller. 2020. What does your gaze reveal about you? On the privacy implications of eye tracking. In *Privacy and Identity Management*. Springer, 226–241.
- [14] Jacob Leon Kröger, Otto Hans-Martin Lutz, and Philip Raschke. 2020. Privacy Implications of Voice and Speech Analysis-Information Disclosure by Inference. In *Privacy and Identity Management*. Springer, 242–258.
- [15] Jacob Leon Kröger and Philip Raschke. 2019. Is My Phone Listening in? On the Feasibility and Detectability of Mobile Eavesdropping. In *Data and Applications Security and Privacy XXXIII*. Springer International Publishing, Cham, 102–120.
- [16] Jacob Leon Kröger, Philip Raschke, and Towhidur Rahman Bhuiyan. 2019. Privacy implications of accelerometer data: a review of possible inferences. In *Proceedings of the 3rd International Conference on Cryptography, Security and Privacy*. 81–87.
- [17] Xavier Duncan L’Hoiry and Clive Norris. 2015. The honest data protection officer’s guide to enable citizens to exercise their subject access rights: lessons from a ten-country European study. *International Data Privacy Law* 5, 3 (2015), 190–204.
- [18] Xavier L’Hoiry and Clive Norris. 2017. Exercising Access Rights in the United Kingdom. In *The Unaccountable State of Surveillance*. Springer, 359–404.
- [19] René Mahieu, Hadi Asghari, and Michel van Eeten. 2018. Collectively exercising the right of access: individual effort, societal effect. *Internet Policy Review* 7, 3 (2018).
- [20] Jim McCambridge, John Witton, and Diana R Elbourne. 2014. Systematic review of the Hawthorne effect: new concepts are needed to study research participation effects. *Journal of Clinical Epidemiology* 67, 3 (2014), 267–277.
- [21] Nurul Momen, Majid Hatamian, and Lothar Fritsch. 2019. Did App Privacy Improve After the GDPR? *IEEE Security & Privacy* 17, 6 (2019), 10–20.
- [22] Rudi Müller-Glögge, Ulrich Preis, and Ingrid Schmidt (Eds.). 2016. *Erfurter Kommentar zum Arbeitsrecht (ErfKoArbR)* (16 ed.). Beck, München.
- [23] David Zeld. 2019. How To Clear Out Your Zombie Apps and Online Accounts. *Wired* (May 2019). <https://www.wired.com/story/delete-old-apps-accounts-online/>
- [24] Clive Norris and Xavier L’Hoiry. 2017. Exercising Citizen Rights Under Surveillance Regimes in Europe—Meta-analysis of a Ten Country Study. In *The Unaccountable State of Surveillance*. Springer, 405–455.
- [25] Chris Norval, Heleen Janssen, Jennifer Cobbe, and Jatinder Singh. 2018. Reclaiming data: Overcoming app identification barriers for exercising data protection rights. In *ACM International Joint Conference and International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*. 921–930.
- [26] German Federal Office of Justice. 2019. Federal Data Protection Act (BDSG). https://www.gesetze-im-internet.de/englisch_bdsge/englisch_bdsge.html#p0014
- [27] Kate O’Flaherty. 2019. Huawei Security Scandal: Everything You Need to Know. *Forbes* (Feb. 2019). <https://www.forbes.com/sites/kateoflahertyuk/2019/02/26/huawei-security-scandal-everything-you-need-to-know/>
- [28] European Parliament and Council of the European Union. 2016. 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). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.L_.2016.119.01.0001.01.ENG
- [29] Article 29 Data Protection Working Party. 2010. Opinion 8/2010 on applicable law (0836-02/10/EN). *WP* 179 (2010).
- [30] Anthony Quattrone, Lars Kulik, Egemen Tanin, Kotagiri Ramamohanarao, and Tao Gu. 2015. PrivacyPalisade: Evaluating app permissions and building privacy into smartphones. In *2015 10th International Conference on Information, Communications and Signal Processing (ICICIS)*. IEEE.
- [31] Abbas Razaghpanah, Rishab Nithyanand, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Mark Allman, Christian Kreibich, and Phillipa Gill. 2018. Apps, Trackers, Privacy, and Regulators: A Global Study of the Mobile Tracking Ecosystem. In *Proceedings 2018 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA.
- [32] Adam Satariano. 2018. G.D.P.R., a New Privacy Law, Makes Europe World’s Leading Tech Watchdog. *The New York Times* (May 2018). <https://www.nytimes.com/2018/05/24/technology/europe-gdpr-privacy.html>
- [33] Adam Satariano. 2020. Europe’s Privacy Law Hasn’t Shown Its Teeth, Frustrating Advocates. *New York Times* (April 2020). <https://www.nytimes.com/2020/04/27/technology/GDPR-privacy-law-europe.html>
- [34] Keith Spiller. 2016. Experiences of accessing CCTV data: The urban topologies of subject access requests. *Urban Studies* 53, 13 (2016), 2885–2900.
- [35] Anselm Strauss and Juliet Corbin. 1990. *Basics of Qualitative Research*. Sage Publications.
- [36] Jörg Thoma. 2014. Datenschutzbeauftragter mahnt mangelnde Verschlüsselung an. *Golem* (Sept. 2014). <https://www.golem.de/news/bayern-datenschutzbeauftragter-mahnt-mangelnde-verschlueselung-an-1409-109260.html>
- [37] Tobias Urban, Dennis Tatang, Martin Degeling, Thorsten Holz, and Norbert Pohlmann. 2019. A Study on Subject Data Access in Online Advertising after the GDPR. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology*. Springer, 61–79.
- [38] Siddharth Venkataramakrishnan. 2020. GDPR accused of being toothless because of lack of resources. *Financial Times* (April 2020). <https://www.ft.com/content/a915ae62-034e-4b13-b787-4b0ac2aaff7e>
- [39] Nicholas Vinocur. 2019. ‘We have a huge problem’: European tech regulator despairs over lack of enforcement. *Politico* (Dec. 2019). <https://www.politico.com/news/2019/12/27/europe-gdpr-technology-regulation-089605>
- [40] Paul Voigt and Axel von dem Bussche. 2017. Scope of Application of the GDPR. In *The EU General Data Protection Regulation (GDPR)*. Springer, 9–30.
- [41] Heinrich Amadeus Wolff and Stefan Brink (Eds.). 2018. *Datenschutzrecht in Bund und Ländern* (23 ed.). Beck, München.
- [42] Janis Wong and Tristan Henderson. 2019. The right to data portability in practice: exploring the implications of the technologically neutral GDPR. *International Data Privacy Law* 9, 3 (2019), 173–191.
- [43] Nils Zurawski. 2017. Exercising Access Rights in Germany. In *The Unaccountable State of Surveillance*. Springer, 109–133.