"Revoked just now!" Users' Behaviors toward Fitness-Data Sharing with Third-Party Applications

Noé Zufferey
University of Lausanne
Switzerland
noe.zufferey@unil.ch

Kavous Salehzadeh Niksirat
University of Lausanne
Switzerland
kavous.salehzadehniksirat@unil.ch

Mathias Humbert
University of Lausanne
Switzerland
mathias.humbert@unil.ch

Kévin Huguenin
University of Lausanne
Switzerland
kevin.huguenin@unil.ch

ABSTRACT

The number of users of wearable activity trackers (WATs) has rapidly increased over the last decade. Although these devices enable their users to monitor their activities and health, they also raise new security and privacy concerns, given the sensitive data (e.g., steps, heart rate) they collect and the information that can be inferred from this data (e.g., diseases). In addition to sharing with the service providers (e.g., Fitbit), WAT users can share their fitness data with third-party applications (TPAs) and individuals. Understanding how and with whom users share their fitness data and what kind of approaches they take to preserve their privacy are key to assessing the underlying privacy risks and to further designing appropriate privacy-enhancing techniques. In this work, we perform, through a large-scale survey of $N = 628$ WAT users, the first quantitative and qualitative analysis of users' awareness, understanding, attitudes, and behaviors toward fitness-data sharing with TPAs and individuals. By asking these users to draw their thoughts, we explore, in particular, users' practices and *actual* behaviors toward fitness-data sharing and their *mental models*. Our empirical results show that about half of WAT users underestimate the number of TPAs to which they have granted access to their data, and 63% share data with at least one TPA that they do not actively use (anymore). Furthermore, 29% of the users do not revoke TPA access to their data because they forget they gave access to it in the first place, and 8% were not even aware they could revoke access to their data. Finally, their mental models, as well as some of their answers, demonstrate substantial gaps in their understanding of the data-sharing process. Importantly, 67% of the respondents think that TPAs cannot access the fitness data that was collected before they granted access to it, whereas TPAs actually can do this.

KEYWORDS

privacy, fitness trackers, fitness tracking, wearables, wearable data, third-party applications, user survey

Proceedings on Privacy Enhancing Technologies 2023(1), 1–21

https://doi.org/10.56553/popets-2023-0001

1 INTRODUCTION

The number of wearable activity trackers (WATs), such as wrist-worn fitness trackers and smartwatches, has grown rapidly over the last decade [11]. These devices collect a wide variety of personal data, including physiological (e.g., heart rate) and contextual data (e.g., the time and place where an activity was conducted).[1] WATs are used to monitor (parts of) their users' lives (quantified-self [14])—often related to health—including physical activities, sleep patterns [77], and stress levels [64]. Earlier studies have shown that fitness data can help researchers and practitioners detect, early on, diseases such as Parkinson's [64], sclerosis [8, 32, 80], and SARS-CoV-2 infection [38]. Data collected by WATs can also be used to infer user activities [45, 78, 90], food and alcohol consumption [34, 90], and smoking habits [79].

Although some of the aforementioned usages of fitness data can be perceived as beneficial, fitness data can also be used for malicious and curious purposes, which raises security and privacy issues. For instance, it can be used to infer what a user is typing (e.g., on a computer keyboard, smartphone keypad, or an ATM pin code) [49, 54–56, 72], which can leak user passwords or mount impersonation attacks on biometric authentication systems [18, 19]. Bike-speed data shared on activity-based social networks (e.g., Strava) can drive thefts by enabling thieves to identify and locate high-priced bikes [10]. Finally, the aggregated-heat maps of users' geotagged activities reveals the location of secret military bases all over the world [37]. Therefore, users have to find a balance between utility and security & privacy; unfortunately, oftentimes, they end up compromising the latter [46, 86].

Such attacks can be mounted by any individual and/or entity who has access to users' fitness data. Naturally, this includes the WAT service providers (e.g., Apple, Fitbit—owned by Google—, and Garmin) that collect the data from the trackers by uploading it to their servers, typically through companion mobile apps installed on smartphones paired with trackers and, by extension, their business partners with whom they share data (e.g., advertisers, data brokers), hence even hackers. In these last two examples, the users might not agree with or even know about the access to their data. Beyond these usual suspects, data is often made available voluntarily by users to some individuals (e.g., family, friends, co-workers, healthcare professionals [4, 27]) and entities (e.g., employers [71],

[1] Henceforth, we will refer to the data collected by WATs as fitness data.

insurance companies [12], service providers), typically through third-party applications (TPAs) or social network profiles. Users do so for increased social or financial benefits (e.g., projected image, decreased premiums) and/or for additional features not offered by the original services or applications.

Understanding how users share their fitness data, and more generally *who* has access to their data, is paramount to properly assessing the security and privacy risks associated with fitness data and to developing effective privacy-enhancing technologies (PETs). Although WAT users' *attitudes* toward fitness-data sharing has been widely studied (e.g., [28, 50, 74, 86]), users' *actual behaviors* have so far received, to the best of our knowledge, little attention. In particular, fitness-data sharing through TPAs has been mostly overlooked, although it has received substantial attention in the context of social network data [30, 42, 43, 91] and from the point of view of the security of the associated protocols (i.e., OAuth) [47, 48].

In this work, we fill this gap by addressing the following research questions:

- **RQ1.** To what extent and how do WAT users use and manage the access of fitness-related TPAs? To what extent are they aware of the data shared with these TPAs?
- **RQ2.** To what extent are users aware of the availability of their PII and fitness data on their fitness-tracking profiles (data types and visibility/audience)? Which types of data do they share, and with whom?
- **RQ3.** What are users' attitudes toward existing and potential (e.g., granular sharing) PETs for controlling their fitness data shared with TPAs?
- **RQ4.** What are users' mental models regarding fitness-data collection and sharing processes between WATs and TPAs?

We designed a questionnaire that we deployed through a large-scale survey, in the US ($N = 628$), of WAT users equipped with a device from Apple, Fitbit, or Garmin. We explored users' *behaviors*, especially with TPAs, toward data sharing. We surveyed users' general understandings of how data sharing works, including an analysis of respondents' mental models [59] by asking volunteer respondents to draw the data flow between WATs, TPAs, and other components. We also assessed their understandings about how they can monitor data sharing with their main companion app paired with their WAT (especially access revocation). We evaluated how convenient it is for them to use these functionalities. Last, we measured their attitudes toward different PETs.

Our results show that 70% of WAT users share fitness data with at least one TPA. About half of them underestimate the number of TPAs to which they grant access to their data, and 63% share data with at least one TPA that they do not actively use (anymore). Not surprisingly, 29% of users did not revoke TPA access to their data because they forgot they had given access to it in the first place, and 8% of them were not even aware they could revoke access to their data. Finally, there is a substantial mismatch between the data that users think the TPA can access and that data it can in fact access: 67% think the TPA cannot access fitness data that was collected before they granted access to it, whereas it actually can. Such gaps in users' understandings were also highlighted after we analyzed their mental models.

Outline. The rest of the paper is organized as follows. In Section 2, we present the related work. In Section 3, we describe the system and the threat models. In Section 4, we detail our methodology including participant selections, survey designs, and procedures. We describe the results of our analysis of the survey data in Section 5. We discuss the design implications of our findings in Section 6, and the study limitations in Section 7. Finally, we conclude the study in Section 8.

2 RELATED WORK

Prior research explored WAT users' (self-reported) behaviors, habits, concerns, and attitudes regarding fitness-data sharing. However, to the best of our knowledge, no such study has been conducted about fitness-data sharing with TPAs. More specifically, though some studies previously analyzed smartphone permissions in generic TPAs [65] and users' perception regarding these permissions [16, 20, 61, 92], no study has focused on TPAs for WATs. A different line of research explored the data-sharing intentions and practices of WAT users, but without specific focus on TPAs [4], and the relations between WATs and TPAs to build privacy-enhancing techniques [83] were modeled. However, there is no study specifically about WAT users' *behaviors*, *understandings*, and *attitudes* toward data-sharing with TPAs.

2.1 General Fitness-Data Sharing

Bilogrevic and Ortlieb [9] explores the attitudes and concerns of users of search engines, online shopping, and online social networks toward data sharing, in particular regarding cross-platform data combination privacy risks (user survey with $N = 918$ and interviews with $N = 14$). They show that the type of data and services have an important impact on the users' comfort toward fitness-data sharing. Liao [50] measured the effect of different factors on online data-sharing behaviors (user survey with $N = 553$). They focused, more specifically, on health conditions (self-stigma), on privacy attitudes, and on the different factors (e.g., types of data, online platforms) that affect health and fitness-data sharing. Their results show a negative correlation between self-protection factors (e.g., data security concerns) and data-sharing self-reported behavior. They also show that data-sharing is correlated with gratification for their use of social platforms.

More relevant to WATs and fitness data, the work of Schneegass et al. [75] analyzes how WAT users' willingness to share fitness data changes depending on the data type (incl. the sensors used to collect them), derived data, and the data recipients (user survey with $N = 249$). Their results show a negative correlation between their willingness to share and the size of the audience. Moreover, users prefer to share specific derived data rather than sensor raw data. Furini et al. [26] analyze WAT users' willingness to share fitness data *for altruistic reasons*, more specifically, to help fight the COVID-19 pandemic (user survey with $N = 76$). Their results show that, when given a strong altruistic motivation, individuals tend more to agree with data sharing.

Gabriele and Chiasson [27] study WAT users' general attitudes, knowledge, and behaviors toward fitness-tracking devices (survey with $N = 212$). They focus on the user's awareness regarding the effect of fitness-data collection on their privacy, their sharing

intentions and behaviors, and their general feelings toward data sharing. Their results show that users' concerns and behaviors depend mostly on the data type and data recipients. They also suggest that users need more/finer sharing options.

More recently, Velykoivanenko et al. [86] have studied WAT users' perceptions of the utility of the different features offered by WATs, the associated applications and services (including data sharing) and the associated privacy risks (survey with $N = 227$ coupled with a 4-month field experiment and interviews with $N = 19$). Their results show that respondents are generally aware of the possibility of inferring sensitive information from fitness data, but only physiological information. However, the respondants have rather strong concerns about some of these inferences (e.g., personality). The results also show that respondents are open to some data generalization-based PETs, including those that we study in this work.

Lupton [52] study how WAT users share fitness data with other individuals and privacy concerns (interview with $N = 40$). They report that most of their respondents consider only privacy in a "social privacy" point of view and do not view how their data can be used by third parties (e.g., advertisers, health insurers). Pinchot and Cellante [69] study the factors that affect data-sharing perceptions (survey with $N = 325$). They measure privacy concerns, data-sharing habits, the understanding of privacy settings and privacy policies, the perception of data sensitivity, and the perception of data values. Their results show that self-reported data-sharing behaviors are negatively correlated, with statistical significance, to the understanding of privacy settings, to the understanding of privacy policies, and to the perception of personal-data values. Murmann et al. [63] study the possible adoption of privacy notifications for WAT usage (survey with $N = 304$), where most of their respondents found privacy notifications useful for monitoring their data-sharing and for increasing their privacy.

Finally, a few works propose PETs. For instance, Epstein et al. [22] developed a fine-grained step-count sharing system that enables the user to delete, modify, and aggregate step-count data to prevent privacy risks yet to preserve some information (e.g., total daily step count).

2.2 Fitness-Data Sharing with TPAs

More relevant to fitness TPAs, the work of Alqhatani and Lipford [4] qualitatively analyzes fitness-data sharing behavior and the concerns of WAT users (interview with $N = 30$). Five of the participants reported sharing fitness data with TPAs such as health insurance companies (to reduce their premiums). Regarding the TPAs' behavior, Nobakht et al. [65] developed a software to analyze the code of TPAs compatible with `Google Fit` in order to determine whether they are over-privileged and whether the data requested by the TPA is indeed needed. Their analysis of 20 free fitness-related TPAs shows that some of these TPAs indeed request far more data than they really need to provide their services.

Finally, through a case study (i.e., Fitbit-data shared with the `Lose it!` app), Torre et al. [84] modeled the complexity of the relationship between WATs and TPAs and assessed the privacy risks associated with TPAs. They used a Bayesian network to compute the probability of inferring different information from data tracked or collected by WATs and smartphones. They also developed a system, called AID-S, for helping users manage their privacy. They concluded that users can easily lose track of the accesses they granted to TPAs.

In conclusion, despite all the relevant studies about (a) behavior, attitudes, and concerns regarding online social networks [30, 42, 43, 91], (b) privacy-protection mechanisms [3, 5, 13, 17, 76], and (c) risk assessment [2, 23, 87, 88], WAT users' *behaviors*, *understandings*, and *attitudes* toward data-sharing with TPAs have not yet been studied. Furthermore, the *actual behaviors* of WAT users toward data-sharing with other individuals—i.e., by asking respondents to check what they *actually* do in their account settings—has never been studied.

3 SYSTEM AND THREAT MODELS

Most major WAT providers such as Apple, Fitbit, and Garmin enable data sharing with TPAs by using application programming interfaces (APIs). These APIs enable third-party services to access a part of users' data, provided that the users consent. Fitness data are usually stored on the users' WATs (temporarily), and on the smartphone paired with the WATs. [2] and/or on the WAT provider's servers. Sharing data with a TPA (i.e., granting it access to the data) enables it to access this data, or at least a part of it, until users revoke the access. During this period of time, a given TPA is technically able to store the collected data on their own server and to keep them for as long as they want.

When the data is stored on the WAT provider's servers—as it is the case for Fitbit and Garmin [25, 29]—, user consent is collected through a webpage, and the TPA subsequently obtains access tokens to make requests to the WAT provider's web API that is usually secured with the `Oauth2.0` protocol [15]. Such requests can be made by a mobile app or by a server controlled by the TPA. When the data is stored on the smartphone—as it is the case for Apple—, the TPA can make requests to the local API of the operating system of the smartphone (e.g., iOS) that is secured with mobile permissions for which the user is prompted. Such requests can be made only via a mobile app. Yet, the data collected by the mobile app can be subsequently sent to a server. In both cases, the user can grant access selectively to their fitness data, by choosing in the list of data types requested by the TPA. For instance, they can grant access to step data but not to heart-rate data.

TPAs are known to ask users for access to far more data than they really need and to use to provide their services [65]. Such TPAs can use the data for their own profit either by tracking or inferring new information about the users beyond their services, or by sharing them with other companies without notifying the user [21, 58]. Also, it is possible that some TPAs change their privacy policies, without the users noticing. Individuals who use a large number of functionalities through different TPAs might just not notice the changes or accept the privacy change notifications, without properly reviewing them. Previous research argued that, due to the large number and availability of TPAs, users can easily lose track of their granted accesses [84]. Last, to cease the data-sharing, a user must actively revoke the access permissions by using the WAT

[2] The fitness data stored on the smartphone can be synchronized, possibly encrypted, with a server (e.g., iCloud).

provider's platform; this is not necessarily easy to do for every user, as suggested by our results.

Another way to access WAT users data is to use users' (public) profiles. Users' PII (e.g., birthdate, e-mail address), as well as, to some extent, fitness data (e.g., average step count, list of achievements) might be publicly available on the service provider web platform or using the social functionalities of the companion app. Depending on the privacy settings, potential adversaries can access sensitive information without any authorization and/or consent. Furthermore, as the API data access used by TPAs—and using the `Oauth2.0` protocol— needs only the user's validation (by clicking on a link) and does not necessarily require any account creation or notification, an adversary could use social engineering techniques, such as phishing [85], to gain access to user data.

4 METHODOLOGY

In order to answer our research questions, we collected quantitative and qualitative data about WAT users' data-sharing practices, through a questionnaire we designed and deployed in an online user survey ($N = 628$). Given the exploratory nature of the study, we did not run any statistical power analyses a priory to set the number of respondents. However, considering previous survey studies published on fitness-data sharing (e.g., Liao [50], $N = 553$), we recruited around 600 individuals. Furthermore, we ran an a posteriori power analysis which revealed a high level of power (1.0). The study was approved by the institutional review board (IRB) of our university.

4.1 Recruitment

We recruited our survey respondents via `Prolific` that was assessed as a reliable crowdsourcing platform for scientific research [67]. We first ran a screener survey to select the respondents eligible for our main survey. We relied on `Prolific`'s native screening feature to target individuals who (a) *use* a WAT (i.e., either a fitness tracker or a smartwatch) and (b) live in the US and speak English fluently. We asked respondents four screening questions: (1) the brand of their WAT, (2) the operating system of the smartphone paired with their WAT (if any) (3) the frequency at which they wear their WAT (i.e., number of days per week), and (4) whether they ever shared their fitness data with TPAs. We collected the data of $N = 2504$ respondents. This enabled us not only to select eligible respondents but also to compute general statistics on the market shares of WAT brands and on the use of TPAs.

For our main survey, we selected the respondents who reported using a WAT manufactured by Apple, Fitbit, or Garmin, paired with an Android or iOS smartphone with the official companion app (i.e., Apple Health, Fitbit, and Garmin Connect, respectively). We chose these manufacturers as they are the three market leaders in the US.[3] We excluded those who reported not wearing their devices at least one day per week. We further excluded those who reported having never granted access to their fitness data to a TPA. The screener took 53 sec on average. Following Prolific's recommendations, we paid the respondents USD 0.12. We selected 1461 eligible respondents that we contacted for participating in the main survey.

[3]Apple is the leader in the US WAT market with a share of 37.6% in terms of sales volume. Fitbit is second with 19.3%, followed by Garmin with a 8.1% [81]. This was confirmed by the results of our screener survey.

4.2 Design of the Survey Questionnaire

We designed the questionnaire to collect various information about WAT users' behavior, awareness, understandings, and attitudes toward fitness-data sharing. In addition to demographics and general WAT usage data, we collected information related to fitness-data sharing with individuals and TPAs and information about their general understandings of the fitness-data sharing ecosystem and the respondents' willingness to use new features that could help them better preserve their privacy in the data-sharing process with TPAs. The questionnaire was composed of between 40 and 51 items spread over seven sections. For some sections of the survey, the number of items varied depending on the respondent's WAT brand, smartphone operating system, and previous answers. The questionnaire was designed to take around 30 minutes to complete. Next, we explain each survey section in detail. All supplementary materials of the paper are available in the Open Science Framework (OSF) repository.[4] The questionnaire is available in Supplementary Material 1.[5]

Sec. A: Introduction. The respondents had to confirm consent to participate in the study and they had to meet all the requirements. For a quality check, they were asked to answer again the same (small) set of questions as in the screener survey. Next, we asked a question about their WAT's utility. The respondents were asked which functionalities of their device they generally use (i.e., related to the data collected by their WAT), such as step tracking, sleep tracking, or stress monitoring.

Sec. B: Data Sharing Using TPAs. We polled the respondents' behaviors concerning and awareness of data sharing with TPAs (see RQ1). To assess the respondents awareness regarding their own data-sharing behavior, we repeated several questions in the survey (what they think they do vs. what they actually do). The first time, we asked the respondents to answer the question "off the top of their heads", and the second time, we asked them to answer the same question after checking their mobile apps (i.e., Apple Health, Fitbit, or Garmin).

We asked them to answer "off the top of their heads" about the number of TPAs they currently use and about the names of the TPAs. Then we asked them to answer the same question after checking the privacy settings of their apps. To facilitate answering these questions and to reduce their cognitive effort, we provided a step-by-step visual guideline to help them navigate through their apps to find the requested information. We also provided the respondents with a list of well-known TPAs that we selected by using the ranking from `data.ai` (i.e., formerly `App Annie`). For each mobile platform (i.e., Android or iOS), we selected the ten apps in the "Health & Fitness" category with the highest number of active users at the time when we deployed the survey. In order to reduce the respondents' cognitive effort, we limited the number of proposed options to ten. We did not include either the fitness-tracker companion apps (i.e., Fitbit, Garmin, and Apple Health) in the app list, or the apps that do not use data collected with Apple, Fitbit, or Garmin WATs (e.g., `Oura` can be only linked to a specific connected ring).

[4]See https://osf.io/z6fw9 (DOI: 10.17605/OSF.IO/Z6FW9), last accessed September 2022.
[5]Note that, as some questions can directly provide information about the data-sharing process hence about prime the respondents, they are not displayed in the same order as presented in this article and are not necessarily ordered by information type.

Finally, we asked the respondents about their general usage of these TPAs (e.g., whether they still use them actively). We also asked them how they generally choose which data to share, among those requested by the TPAs. Indeed, during the data-sharing process (i.e., granting access to a TPA), the user has to select, for each data type requested by the TPA, which ones they agree to grant access to. Because some TPAs request access to more data types than they actually need to provide their services [65], we asked our respondents if they usually share all requested data types, if they share everything only when it is necessary to use the app, or if they share selectively.

Sec. C: Data Sharing via Public Profile. We also probed the respondents'about their behaviors concerning fitness-data sharing via their public profile[6] and their awareness regarding the types of information that are accessible via their public profile (see RQ2). Similarly to the previous section about data sharing using TPAs, we asked the respondents to select, from a list, the types of data that are publicly available on their fitness companion app profiles. We explicitly asked them to do it "off the top of their heads", then we asked them to check their apps' privacy settings. Thus, we could estimate the difference between what they *think* they are publicly sharing and what they *actually* share. Finally, we asked the respondents if they had ever modified the default privacy settings of their app to change the availability of some of the data on their profile.

Sec. D: Data Sharing with Others. We asked if they share their fitness data with other individuals or entities such as their friends, employer, health insurers (see RQ2). We asked the respondents to check their apps and to select the types of data that they share with other individuals, the number of individuals they share with, and the types of relationships with those individuals. We selected the following types of data recipients based on a previous study [4]: friends, family, strangers, physicians (or health professionals), co-workers, and financial-incentive programs. We replaced the "financial-incentive program" with "employer" as most of these programs are set up in collaboration with employers [24], especially in the US where the employers pay for health insurance. Furthermore, an employer is more likely to represent a *natural* person, compared to an organization that represents a *legal* person. Hence, we also asked the respondents if they share their fitness data in the framework of any health programs (e.g., with employer or health insurers) [31, 60, 82].

Sec. E: Attitudes toward Privacy-Enhancing Technologies. We evaluated the willingness of the respondents to use new PETs for data-sharing practices with TPAs (see RQ3). We present three different functionalities: (1) reduce time granularity, (2) to reduce spatial granularity (i.e., data precision), and (3) remind users to monitor TPA access to their data (i.e., "privacy checkup reminder"). For each of these functionalities, we asked them to evaluate how likely they would use it on a seven-point Likert scale from *extremely unlikely* to *extremely likely*.

The first solution (i.e., time-granularity reduction) enables users to choose the level of time granularity with which their fitness data should be shared. The second solution (i.e., data-precision reduction) enables users to choose the level of precision with which their fitness data are shared. The solutions are illustrated in Figures of Appendix A. The solution let users choose to share data as it is, rounded to the tens, rounded to the hundreds, or rounded to the thousands. The last solution (i.e., access-monitor reminder) is shown in Figure 1. It sends users recurrent privacy notifications reminding them to check and revise their previously granted access to TPAs. Users can customize the notification period to receive it either weekly, monthly, or every three months. To the best of our knowledge, none of these functionalities were currently tested in the earlier studies or implemented in the existing fitness platforms (Although, Facebook and Google do encourage–with reminders—their users to conduct so-called security/privacy checkups). However, since version 11, a similar mechanism is used by Android to revoke permission granted to apps that are no longer used [68].

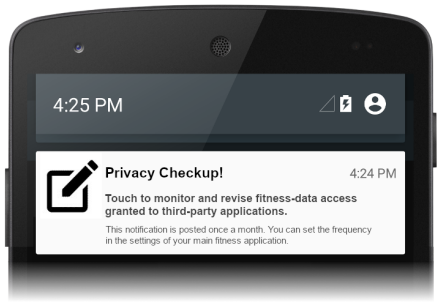

Figure 1: Picture presented to the respondents to illustrate the proposed TPA access monitoring reminder.

For the respondents who answered that they are not using actively all their installed TPAs, we also included an open-ended question: *"Why did you not revoke their access ?"* We asked them to evaluate how easy did they find the whole data-sharing process. Finally, we asked one last open-ended question about the usability of the data-sharing monitoring process in order to collect respondents' suggestions.

Sec. F: Understanding of Data Sharing. We assessed the respondents' understandings of the data-sharing process (see RQ4). We asked them to evaluate (i.e., mark as true/false) different statements about what happens to their shared data (from technical and legal aspects) after they grant access to TPAs and after they revoke it.

Furthermore, we probed respondents' *mental models* by asking them to draw their thoughts. Mental models are explanations of individuals' subjective and implicit assumptions (i.e., tacit knowledge) about how they perceive and conceptualize different phenomena [40] and how they think different technologies work [59]. Given that verbalizing such tacit knowledge might be difficult for individuals (respondents in our case) [53], recent studies on security and privacy [41, 44, 53, 66, 86] asked their participants to draw their thoughts. Following these studies, we asked the respondents to *"Draw a picture representing how you think the access granting to TPAs is processed, and how your fitness data is transferred between different entities."* We recommended they consider including all relevant elements in their drawing, including their WAT, their smartphone, the WAT providers' servers, the TPAs, and any other elements they deemed relevant. We also instructed them to use lines/arrows to connect these entities (i.e., typically for data flows) and to use text to label them. We did not provide any template

[6]Only applicable for Fitbit and Garmin users, as Apple Health does not provide any public profile functionalities.

drawings so as to avoid priming respondents' and limiting their creativity.

The respondents were asked (1) to not spend more than five minutes on the drawing, (2) to take a clean sheet of paper and a pen or pencil, and (3) to take a good-quality photo with their smartphone. Last, they were informed that their drawing skills would not be judged or evaluated by the researchers. Making the drawing was optional, and the respondents were informed that by submitting a drawing they would automatically be enrolled in a lottery for an extra 10$ bonus payment (1 bonus per 5 participants). We collected a total of 142 drawings.

Sec. G: Additional Questions. We included some questions that were not directly related to data sharing. These questions were asked either to collect demographic information that is not provided by Prolific, to verify that the respondents correspond to all the criteria (i.e., screening questions), or to personalize the survey (e.g., questions about the device usage and companion app). Finally, we measured the data-collection concern by using the Collection part (four items) of the Internet Users' Information Privacy Concerns (IUIPC), as well as the Global Information Privacy Concern of the respondents by using three items (i.e., items 2, 3, and 6) of the corresponding scale described in Malhotra et al.' article on IUIPC [57].

4.3 Procedure

Before deploying the survey, we conducted cognitive pretests in order to address potential problems in the survey design. We asked five researchers, who were not involved in this research project, from our university to take the survey. They were all WAT users (2 Apple, 2 Fitbit, 1 Garmin), and all met our selection criteria. One pretest was conducted in person, whereas the others were conducted remotely via Zoom. For each pretest, the first author carefully observed the test subject taking the survey. The test subject was instructed to rephrase the questions, in their own words and out loud, to describe what they think was asked, then to answer it. At the end of the pretest, the author and the test subject were debriefed about the subject's understanding and answers. The pretests showed that the survey instructions and questions were overall clear. Few understanding issues were raised and addressed. For example, we removed the negative forms in some questions, put some important elements in bold, and added instructions to specify when the respondents could validate a multiple choice question without selecting any options.

Out of the 1461 eligible (potential) respondents we contacted (from the screener), 745 started the main survey. To reach our objective, we contacted them in batches. Ultimately, 660 completed the main survey (slightly above our objective). It took, on average, 16 min and 14 sec to complete (SD: 10 min and 28 sec, Min: 3 min and 33 sec, Max: 86 min and 17 sec). The respondents were paid USD 5.

4.4 Data Reliability

Although Prolific is a reliable crowdsourcing platform, it cannot prevent undesirable behavior from some respondents, such as speeders and straightliners. Thus, we applied some strategies to clean the data. First, the individuals who answer "no" to the question on the use of TPAs in the main survey and "yes" in the screener survey were redirected to the end of the survey and their data was discarded (as they gave inconsistent information). Second, we analyzed the answers of the speeders who completed the survey in less than five minutes. We decided to consider such respondents as reliable only if their answers were consistent and if their answers to the open-ended questions made sense [62]. Third, we analyzed inconsistent answers, where the answers to some questions contradict the answers to other questions. For example, some Apple Watch users indicated that they share some type of data with their family but, in the subsequent question, they indicated that they share data with no one from their family. We decided to remove such answers, yet we kept their mental model, if they submitted one. As a result, we removed the answers of 32 respondents.

4.5 Coding Process

We collected 142 drawings that represent the respondents' mental models. We first applied a quality check to ensure that all the drawings have proper quality and include (relevant/meaningful) content. We excluded 6 drawings (4.2%): those whose photos were of low quality, did not include any relevant content, or were copied from the Internet. For the remaining 136 drawings, we focused on two aspects. First, we studied the technical understanding and correctness of respondents, in terms of the information flow within the ecosystem of WATs and TPAs. Second, we studied the contextual information, such as their understanding of data-sharing and privacy concerns, they included in their drawings. The mental model dataset (i.e., all drawings) and two codebooks (i.e., technical codebook and contextual codebook) are available in Supplementary Material 2.

For the respondents' technical understanding, we excluded 4 drawings (2.9%), as the respondents illustrated high-level abstract drawings and did not represent the low-level details. Among the remaining 132 drawings, we checked the types of the elements (WAT, smartphones, connected devices, WAT servers, TPAs, *etc.*) depicted in the drawings and the way these elements were connected to each other. Accordingly, we clustered the mental models and identified the main types of models. Also, following previous studies [1, 44, 53, 86, 89], we labeled respondents' mental models as either correct, inaccurate, or incorrect.

For the contextual information displayed in the drawings, we reviewed (1) the textual information and labels that indicate users' actions, attitudes, and understanding (e.g., access revoking, reporting privacy consequences), (2) the recipient types (e.g., advertisers, hackers, public), and (3) the data types (e.g., step, location, heart rate). Out of 136 drawings, we identified 73 (53.7%) that illustrate contextual information. We developed a codebook by using open coding [73], where we coded 113 elements in the drawings. In total, we identified 20 distinct codes categorized in four themes.

Finally, for the analysis of the answers to the open-ended questions, we used the affinity diagramming method [51] to organize and sort the ideas and thoughts raised in the answers. The second author proceeded to the coding of open-ended questions, then the first author reviewed and provided feedback. The codebooks for three open-ended questions are available in Supplementary Material 3.

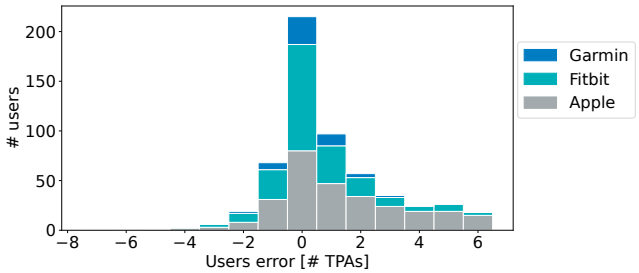

Figure 2: Distribution of the number of users in terms of the difference between the number of TPAs they really have and the number of TPAs they think they have. A positive difference means that they underestimated the number of TPAs, whereas a negative difference means that they overestimated it.

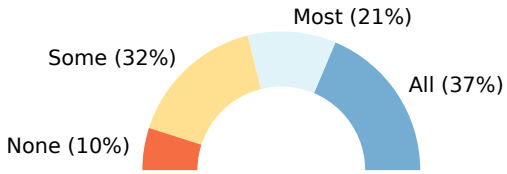

Figure 3: Ratio of respondents who (still) actively use all, most, some, or none of their TPAs.

4.6 General Statistics

Regarding the WAT brand, 53% of our respondents own an Apple Watch, 38% own a Fitbit device, and 9% are Garmin users. Regarding gender, 61% of our respondents are women, 37% are men, and 2% are non-binary. This roughly corresponds to the general population of fitness-tracking users [70]. The average age of the respondents was 35 years old (SD: 11, Min: 18, Max: 73) distributed in the following ranges: 18-29: 37%, 30-39: 34%, 40-49: 16%, 50-59: 9%, 60+: 4%.[7] The respondents reported that they wear their devices 6.4 days a week on average (SD: 1.1, Min: 1, Max: 7), and daily for 1-6 hours (7%), 7-12 hours (24%), 13-18 hours (30%), 19-24 hours (39%). 17% of the respondents reported that they have had their current device for less than a year, 41% for 1 to 3 years, 28% for 3 to 5 years, and 14% for 5 years or more. As for their privacy concerns (assessed using IUIPC items), the collection scores are the closest to a truncated normal distribution (IUIPC Collection score: $\mu = 5.4, \sigma = 2.3, a = -0.05, b = 6.05$; the Global Information Privacy Concern score: $\mu = 3.5, \sigma = 1.6, a = -0.05, b = 6.005$), with μ the mean score, σ the standard deviation, and a and b the bounds. Figures 13 and 14 in Appendix B show the score distributions.

5 RESULTS

In this section, we present the results and findings from our survey, according to the ordering of the questions as presented in Section 4.2.

5.1 Users Rend to Forget About Their TPAs

The data collected in the screener survey shows that the majority of the US-based WAT users (70.2%) share some of their fitness data with TPAs. Using TPAs for fitness data is therefore a common practice and it is paramount that users understand the functioning of this ecosystem (WAT-TPA) and its privacy implications. Among the respondents of the main questionnaire, "MyFitnessPAL", "Strava", and "Achievement" were the three most frequently installed TPAs with fitness-data access. Figure 2 shows the distribution of the respondents' errors when estimating the number of their TPAs that have access to their fitness data. The error is computed for each respondent and is defined as the difference between the actual number of TPAs that have access to their fitness data (obtained by asking the respondents to verify in their companion app settings) and the estimated number of their TPAs that have, according to them, access to their fitness data ("off the top of their heads", before verification). We can see that the number of such TPAs is clearly underestimated by respondents (the difference is significant with $t(627) = 12.85, p < .001$, Cohen's $d = 0.51$, paired sample t-test), which confirms Torre et al. [84]'s statement that, due to the large number and availability of TPAs, users can easily lose track of the TPAs to which they granted access to their fitness data. Although one-third of the respondents (35%) correctly estimated the number of TPAs, almost half of them (49%) underestimated it, and only 16% overestimated it. As shown in Figure 3, two-thirds of the respondents reported that they do not actively use some of their TPAs. Such behavior confirms that a large proportion of WAT users share their data with service/app providers, without benefiting from the service/app (as they do not use it), and sometimes even without being aware of it. Moreover, 64% of the respondents reported that they have never revoked data access, and 8% did not even know it was possible.

In order to better understand how WAT users share their fitness data with TPAs, we looked at the type of data that they agreed (by selecting them, when asked) to share with their TPAs. As explained in Section 3, when giving access to a TPA, the user can choose the type(s) of data they want to share among those that are requested by the TPA (the types are defined by the companion app). TPAs are known to ask for far more data than they really need to provide their services [65]. 32% of the respondents declared that they share everything; 45% of them share only the data necessary for the use of the TPA; and 23% share selectively, despite a potential decrease in the utility of the TPA. Note that the number of users who agree to share all the requested data is substantially higher among owners of Apple devices (39%), compared to owners of Fitbit (25%) and Garmin (21%) devices.

We looked at the reasons the respondents who reported not actively using some of their TPAs did not revoke their access. Table 1 shows the results. First, some respondents reported they usually do not bother with access management. They reported that they have never thought about such actions, and some of them mentioned they do not perceive fitness data as sensitive hence would not care about doing any privacy-preserving actions. [M, 30-39 y.o., Apple]: *"I just never think about it and do not think it is an issue to leave them on."* Second, many respondents simply did not revoke any accesses, as they forgot that they had installed these TPAs. A few of them mentioned they remembered their TPAs, only after answering our survey, and they plan to revoke their access later.

[7] The age information for three respondents was not available in Prolific's statistics.

Category	Freq.
comfortable to share data (not interested in access management)	29.7%
forgot about installed TPAs (might revoke later)	29.4%
contemplate using the TPA (actively) again in the (near) future	26.7%
not familiar with data sharing and access management	18.7%
perceive access management as complex / difficult (hassle)	3.9%
want to get more benefits (health or monetary)	1.1%
trust TPAs	0.8%
others	2.7%

Table 1: Main reasons respondents do not revoke access to their data to the TPAs that they no longer use actively.

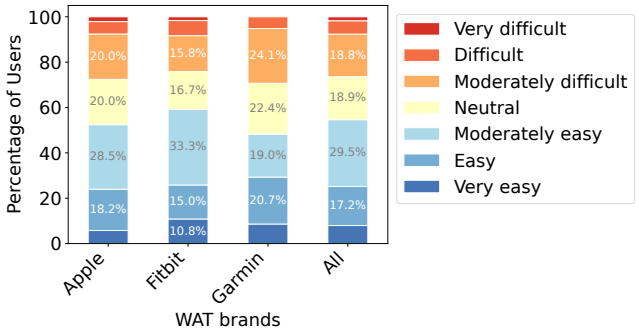

Figure 4: Evaluation of the complexity / difficulty of the TPA fitness-data sharing monitoring process.

[W, 18-29 y.o., Apple]: *"I forgot and didn't realize the apps had access until completing this survey."* This confirms the aforementioned findings that using many TPAs and forgetting them is a common (privacy) issue among WAT users. Third, several respondents did not revoke access as they thought they might use the TPA later in the future. Fourth, around one-fifth of the respondents reported they did not know that TPAs collect their data or did not know how to manage these accesses. Finally, a few respondents perceived access management as a hassle. [M, 18-29 y.o., Apple]: *"I find it troublesome to revoke their access."* This is confirmed by the results in Figure 4 that shows that around one-fifth of the respondents consider the TPA data-sharing monitoring process as moderately difficult to very difficult.

Conversely, we looked at the reasons the respondents who reported revoking access did so. More than four-fifths of the respondents reported revoking access after not using their TPAs. 64.9% did not use the app for a long time hence stopped the data collection, 13.5% were not satisfied with the app or had technical issues, and 2.3% used a new TPA and revoked the access of the older one. A total of 27 respondents (15.8%) reported revoking access due to privacy concerns as they felt uncomfortable with data collection. [W, 30-39 y.o., Garmin]: *"I was nervous about the data they were accessing."*

5.2 Users Generally Overestimate the Amount of Data They Share on Their Public Profiles

Unlike Apple, Fitbit and Garmin include social network functionalities in their applications, where users have a public profile on

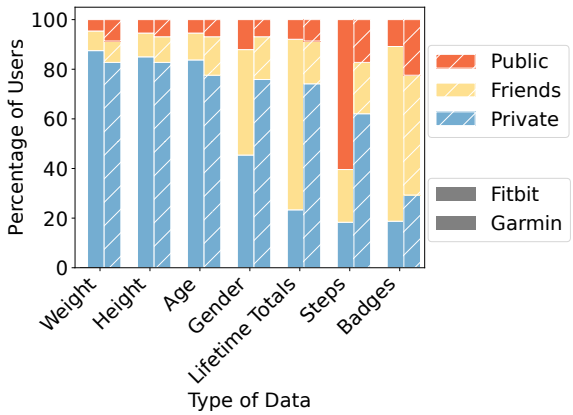

Figure 5: Privacy level of different profile information for Fitbit and Garmin users. We show only the types of data that are available on both Fitbit and Garmin users' profiles.

which they can share certain personal data. In its settings, Fitbit defines nine different types of data for which the users can choose three privacy levels: "private", "my friends", or "public." However, since a recent update, a user's average daily steps can no longer have the "private" level . Garmin defines four different types of data for which the users can choose four privacy levels: "only me," "my connections," "my groups and connections," and "everyone." A fifth level is available for activities (namely "custom"), but none of our respondents used it. Moreover, the users can also select among nine types of data that one can be displayed on their profile.

Figure 5 shows, for each type of data that can be made available on Fitbit and Garmin user profiles, the proportion of users that selected each level of privacy. Here, we refer to concepts of both service providers: We used (1) Garmin's labels (e.g., "Badges" and "Badges and Trophies"), (2) Fitbit's privacy labels, and (3) both Garmin's "my connections" and "my groups and connection" as "friends". We also removed all types of information that are not available in both Fitbit and Garmin profiles. More details are available in Figure 15 of Appendix C. It can be observed that, in general, Fitbit users tend to share more information via their public profile. This might be caused by the difference of the default privacy settings in both apps. Indeed, though all profile information are by default set to private, Fitbit set the privacy level of "Lifetime Steps, Distance, and Floors" (called "Lifetime Totals"), and "Badges & Trophies" (called "Badges") to "Friends" and the privacy level of "Average Daily Step Count" (called "Steps") to "public". Moreover, 43% of the respondents declared never having changed their privacy settings.

We also looked at the information Fitbit and Garmin users thought "off the top of their heads" were publicly available on their profiles before they checked their settings. As shown in Figure 6, Fitbit and Garmin users highly overestimate the public accessibility of their data, except for the friends' list. This means that, for a large number of users, they well overestimated the amount of information that is actually publicly available.

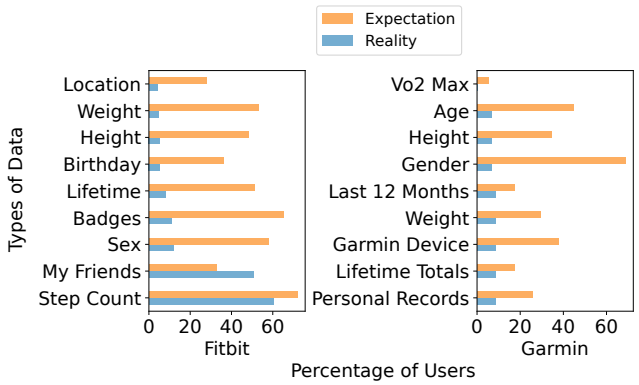

Figure 6: Expected vs. real proportion of public availability of specific data types (Fitbit and Garmin users).

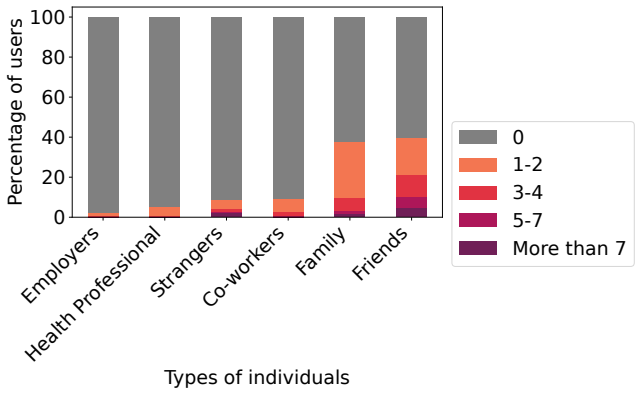

Figure 7: Number of individuals in each relationship group with whom our respondents share their fitness data.

5.3 Friends and Family Are Favorite Data Recipients

As seen before, WAT users have the possibility to share some of their fitness data with individuals. Although Fitbit and Garmin provide privacy levels for each type of data, Apple provides the possibility to define which type of data they want to share with each of their contacts. We asked our respondents, among a list of social relationships, with how many of them they share at least one type of fitness data. Figure 7 shows that WAT users tend to share their fitness data with friends and family more than with other groups of individuals. Indeed, 40% of the respondents declared sharing data with at least one friend and 38% with at least one family member, whereas less than 10% share data with the other groups of individuals (only 2% with employers). Furthermore, 1% of them declared sharing their fitness data in the framework of a health program (e.g., with their employer and/or health insurance company). This corroborates Gabriele and Chiasson [27]'s findings about users' privacy concerns and willingness to share. However, the *actual* sharing behavior that we measured is far lower than their willingness to share, as well as their comfort in sharing, measured

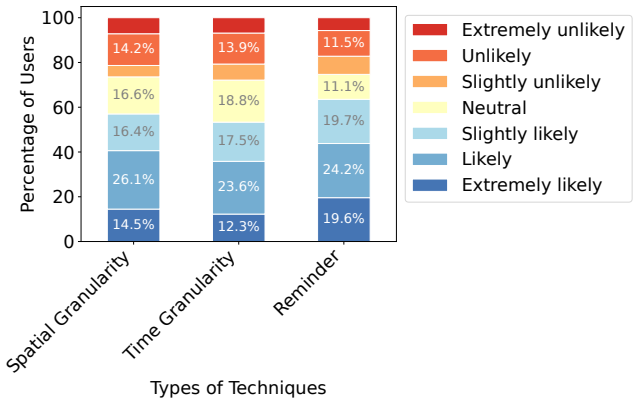

Figure 8: Self-declared likelihood to use the three different proposed PETs.

by Gabriele and Chiasson. This shows that, even if users are ready to share their data, they do not necessarily do so.

5.4 Users Are Enclined to Use PETs

We looked at the (self-reported) likelihood that respondents would use the different PETs we proposed. The results, depicted in Figure 8, suggest that most of the respondents are (slightly to extremely) likely to use each of the three techniques. The likelihood is even higher for the access monitoring reminder, for which 63.5% of the respondents declared that they would likely use. Therefore, we recommend fitness-data service providers to offer the reminder technique that is rather straightforward to implement. As for the other techniques, Velykoivanenko et al. [86] show that they can be implemented with a modest decrease in (perceived) utility.

We also looked at the participants' suggestions on how to facilitate the TPA access management process (i.e., granting, monitoring, or revoking access). We collected 480 meaningful (open-ended) answers and categorized them into three main families of solutions.

First, the majority of our respondents (53.5%) proposed access monitoring solutions. In line with our earlier finding, the most promising solution (39.8%) was the use of **periodic reminders** in the form of pop-up notifications. The respondents were in favor of a system that could review TPAs, flag those that have not been used for a certain amount of time, and remind users to reconsider the accesses they granted. [W, 30-39 y.o., Fitbit]: *"I think the reminders are great! I allowed access to some app and totally forgot about it. I'm not sure if they're still collecting data, but had I remembered, I would have revoked it."* A few respondents (3.5%) even proposed a more proactive solution: **a privacy check-up** feature that could automatically revoke access for unused TPAs and then provide users a list of TPAs whose accesses were revoked [M, 30-39 y.o., Garmin]: *"Garmin should automatically revoke access every few months (such as every six months) and ask me again whether I should grant access to the third-party apps. Then I can decide whether I am still interested in those apps and whether it is worth sharing the data."* Some respondents (4.0%) asked for a specialized app or a feature in the phone operating system to handle the access management procedure. They (6.3%) proposed a consolidated feature that can

present a list of TPAs, including the types of data they collect and where they store the data. [M, 60+ y.o., Apple]: *"Place the permissions in a consolidated location, rather than skipping around to apps that may or may not be reading data."* Note that, for Fitbit and Garmin, we should distinguish between the use of the fitness tracking service's (i.e., Fitbit and Garmin) API to access fitness data and the use of the TPA's mobile application associated with the TPA (e.g., Strava). Indeed, the API calls could be made from Strava's servers, regardless of whether the Strava mobile app is actually used. The fitness tracking service knows only when API is used, whereas the mobile operating system knows only when the TPA's mobile app is used.

Second, several respondents (12.9%) suggested solutions for improving the access granting procedure. They asked for clear, transparent, and easy to understand **privacy policies** (8.5%). [M, 40-49 y.o., Garmin]: *"I would like to see everything laid out in plain English, no lawyer-speak. I would like it to be clear whether they can keep my data forever, sell it data, collect it after I revoke access, etc. I would also like to know who and why is potentially buying my data."* Note that Harkous et al. [35] proposed a similar solution named Polisis. A few respondents suggested a **time-framed sharing** feature where users can decide to share only data collected in a given time frame (1.0%).

Third, many respondents (31.7%) proposed generic solutions. In particular, they (20.8%) asked that the access management procedure be facilitated and that the user interface be made easier to interact with. They mainly found that the information about collected data types, sharing conditions, and sharing consequences were not clearly stated when granting access and/or that they were hidden in the interface. They asked for **better visibility** to help them make informed decisions when granting access, and for usable interfaces to facilitate revoking access. [M, 40-49 y.o., Fitbit]: *"Don't bury the feature under multiple levels of the app's user menu. Place it front and center at the top level under My Account."* A few respondents asked that users be informed about TPAs (4.0%) and that there be better legislation (privacy rules) and law enforcement for some TPAs that infringe user privacy (2.5%). The rest of the suggestions (4.4%) were about ensuring the deletion of previously-stored data after an access is revoked, and the automatic revoking of an access when uninstalling a TPA.

5.5 Users Lack Knowledge About Data Sharing

We evaluated our respondents' awareness and understanding of fitness-data sharing with TPAs, both qualitatively and quantitatively. We asked them questions, for which we knew the ground truth, and we requested that they draw (facultative) on paper how they picture the data flow when sharing fitness data with TPAs.

5.5.1 Mental Models – Technical Understanding. We first present our findings on respondents' technical understanding of the information flow in the WAT and TPA ecosystems. In terms of the elements drawn, most of the drawings illustrated the main elements of the ecosystem (i.e., WATs: 92.4%, connected devices: 84.8%, WAT servers: 81.1%, TPAs: 97.0%), where 65.9% of the drawings included all together these four elements. Among the drawings with TPAs, 56.3% included one TPA or a third-party server, and the rest included two or more TPAs or third-party servers. A few drawings (10.2%) included additional elements such as databases, other smart devices (e.g., scales), satellites, API, PC, GPS, etc.

name	description	Apple	Fitbit	Garmin
mm_1	online, using a phone	✗	✓	✓
mm_2	online, without using a phone	✗	✗	✗
mm_3	offline, using a phone	✓	✗	✗

Table 2: Ground truth for mental models.

category	Apple	Fitbit	Garmin	total
mm_1	$n = 18$ (29.0%)	$n = 24$ (45.3%)	$n = 6$ (35.3%)	36.4%
mm_2	$n = 2$ (3.2%)	$n = 6$ (11.3%)	$n = 3$ (17.7%)	8.3%
mm_3	$n = 23$ (37.1%)	$n = 10$ (18.9%)	$n = 2$ (11.8%)	26.5%
$mm_{1\&3}$	$n = 3$ (4.8%)	$n = 1$ (1.9%)	$n = 1$ (5.9%)	3.8%
mm_4	$n = 16$ (25.8%)	$n = 12$ (22.6%)	$n = 5$ (29.4%)	25.0%
correct	$n = 23$ (37.1%)	$n = 24$ (45.3%)	$n = 6$ (35.3%)	40.2%
inaccurate	$n = 3$ (4.8%)	$n = 1$ (1.9%)	$n = 1$ (5.9%)	3.8%
incorrect	$n = 36$ (58.1%)	$n = 28$ (52.8%)	$n = 10$ (58.8%)	56.1%
total	$n = 62$	$n = 53$	$n = 17$	$n = 132$

Table 3: Mental model results.

We identified four main patterns in the drawings: the different types of mental models.

- mm_1. Online data synchronization where the data is transmitted from a WAT to a TPA via a connected device and a WAT server.
- mm_2. Online data synchronization where the data is transmitted without passing through a connected device: Directly from a WAT to a WAT server and then to a TPA server.
- mm_3. Offline data synchronization where the data is transmitted locally on a connected device between a WAT app (e.g., Apple Health app) and a TPA—without requiring data transmission through their respective servers.
- mm_4. Drawings that we could not attribute to mm_1–mm_3 (other).

Before evaluating these models, we checked the ground truth by carefully reviewing the privacy policies and technical documents of Apple, Fitbit, and Garmin [6, 25, 29]. We also contacted the Garmin support team to confirm our findings related to their devices. Table 2 summarizes the ground-truth findings showing that Apple devices have a different ecosystem, compared with Fitbit and Garmin devices. Whereas Apple devices exchange information with TPAs locally and not via their servers (i.e., mm_3), Fitbit and Garmin devices do it via their servers (i.e., mm_1).[8] We also found that the data (for all devices) is always transmitted through the smartphone, hence mm_2 is an "incorrect" model.

In summary, we evaluated the drawings for each respondent considering exclusively their device brand. To wit, we considered mm_1 as "correct" for Fitbit and Garmin owners and "incorrect" for Apple owners. Similarly, mm_3 was considered "correct" for Apple owners and "incorrect" for others. We also labeled the drawings that included both mm_1 and mm_3 "inaccurate". Finally, we considered other drawings (i.e., mm_4) "incorrect", as they usually missed the main elements, and they did not correctly connect them.

[8] Note that Apple users can back up their data on iCloud. Also, TPAs can store their data on their servers. However, the primary connection between the Apple Health app and TPA is held locally.

Table 3 summarizes the findings. The first type (i.e., mm_1) was the most frequently seen mental model where 36.4% of respondents drew it (e.g., Figures 16a and 16b in Appendix D). We found that 45.3% of the Fitbit owners and 35.3% of the Garmin owners correctly drew mm_1. However, 29.0% of the Apple owners incorrectly thought that their Apple device transmits their health data by using Apple servers.

The second type of mental model (i.e., mm_2) was related to those respondents who incorrectly thought that the online synchronization occurs without passing through a phone. This model was seen for 8.3% of the respondents (see Figures 16c and 16d).

The third type (i.e., mm_3) was for those respondents who connected their WAT mobile app and TPA locally, without using any online path using the WAT server (e.g., see Figures 17a and 16b). 26.5% of the drawings were related to mm_3. Apple owners (correctly) shared this mental model more than other brand owners (e.g., 37.1% for Apple vs. 11.8% for Garmin). All these respondents also connected their phones or TPAs to the servers of WATs and/or TPAs. This indicates that respondents thought that, despite the local synchronization, their data could also be stored on servers.

A few respondents (i.e., 3.8%) had an inaccurate understanding of the information flow, i.e., mixed mm_1 with mm_3 (see Figure 17c). Hence, we considered these models as inaccurate. Finally, 25.0% of the drawings belonged to the "other" category (i.e., mm_4) and were considered as incorrect (see Figures 17d and 18a).

In conclusion, these findings show that more than half of the respondents (56.1%) had incorrect mental models. Among this group, 44.3% mistakenly drew a mental model that belonged to a device different than the device they owned. The others, with incorrect mental models (55.7%), either thought their device could connect to servers without using a connected device or drew irrelevant infrastructures. These respondents did not have the essential understanding of the main elements and their respective connections in the WAT-TPA ecosystem. Such incorrect mental models can cause users to make dangerous decisions when sharing their data hence compromise their privacy. Lastly, in terms of the brands, our findings show that Fitbit owners had a relatively better understanding of the ecosystem compared with the other device owners (i.e., 45.3% for Fitbit vs. 36.2% for others). Also, Apple users confused their ecosystem with that of other brands more than the other device owners (i.e., 33.9% for Apple vs. 19.2% for others).

5.5.2 Mental Models – Contextual Information. We identified four main themes in order to summarize the contextual information included in 73 drawings as follows. Respondents expressed their **lack of trust in TPAs** in 64.4% of the drawings. They voiced their concern that TPAs would share their data to make profits (38.4%). They thought that TPAs could share the data with entities interested in users' data such as companies working in market analysis and advertisement, developers, other TPAs, giant tech companies, scientific institutes, and governments (e.g., see Figures 18b and 18c). The respondents (19.2%) also drew that their data is stored on the third-party servers (see Figure 16b) and might be further analyzed (see Figure 18d). A few respondents particularly mentioned 'information analysis' (6.8%) and wrote about 'user profiling' (5.5%) (see Figure 19a). Finally, a few participants (8.2%) reported being concerned on whether TPAs can keep their data safe and secure (see Figure 19b). In conclusiong, these findings indicate that some users (i.e., 35.3% of the total sample), despite using TPAs, have serious privacy and security concerns about them.

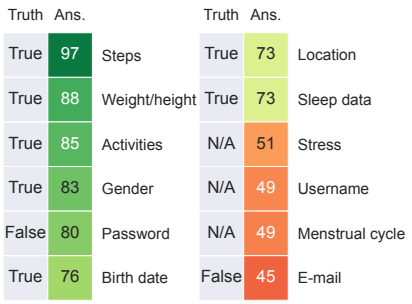

Figure 9: Proportion of correct answers regarding the data shared with TPAs. For each type of data, the ground truth is provided. N/A means that we cannot define a common ground-truth for all respondents as it depends on their device brand.

Some respondents (16.4%) reflected on their **general privacy concerns**, in particular about **the WAT services**. A few respondents expressed concerns that Apple and Fitbit might share their data, without their consent. A respondent reported that Fitbit might share the data with affiliated companies (i.e., Google, see Figure 19c).

Interestingly, about half of the respondents (47.9%) pointed to actions related to **access management** in their drawings. Most of the respondents (42.5%) drew some elements about 'granting or revoking access' in their drawings (e.g., see Figure 19d). A few respondents (8.2%) also sketched 'selective sharing' showing that they could share some data types and avoid sharing others (e.g., see Figure 17b). Although these drawings show that some respondents (i.e., 25.7% of the total sample) are knowledgeable about PETs, such as revoking access or partial sharing, these findings could also be biased as the respondents already received informed about such practices while answering the survey, and this might not reflect their actual practices in their everyday life.

Finally, only a few respondents **reflected trust and comfort** in their drawings (5.5%) where they reported feeling safe about their privacy and being comfortable with the WAT and TPA companies. Two respondents drew that the data collected by WATs could be further analyzed to improve their services and products. One respondent also reported believing that the data would be deleted by a TPA after they revoke their access (see Figure 19d), which is not necessarily the case.

5.5.3 Quantitative Measurement of the Users' Understandings. As for quantitatively measuring our respondents' understandings about fitness-data sharing with TPAs, we asked them to answer two types of questions. For the first, we provided a list of data types and asked them to select, as if they had granted access for all possible types of data, which one could be shared with TPAs. For the second, we provided five statements about what TPAs can technically and legally do after a user grants them access. Then, we provided three statements about what TPAs can technically do after access is revoked.

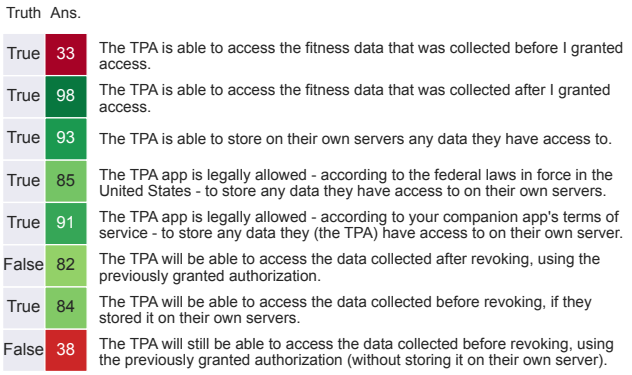

Figure 10: Proportion of correct answers regarding the (legal and technical) feasibility of data access by TPAs.

Figure 9 shows the proportion of correct answers for each type of data. We can see that, in particular, 20% of the respondents believe that the password of their companion app account is shared with TPAs, whereas 55% of them believed that the e-mail address linked to their account is shared. Both are not shared by any of the studied WAT brands. Indeed, sharing such user information can be considered to be a high privacy and security threat. However, we also observe that a non-negligible fraction of the respondents underestimated the information that can be shared with TPAs. For example, more than one fourth of the respondents believed that location or sleep data cannot be shared with the TPAs, whereas in fact, they can.

Figure 10 shows the percentage of correct answers for each provided statement about fitness-data sharing with TPAs. We can observe that, in particular, most of the respondents (i.e., two-thirds) falsely believed that the data collected before they granted access cannot be accessed by the TPAs; this is false. Indeed, granting fitness-data access permits the TPAs to access every data of a specified type stored either on a server when using APIs or on a smartphone, when using local sharing. In addition to this statement, most of respondents (i.e., almost two-thirds) also falsely believe that a TPA, for which the data access has been revoked, can still access the fitness data collected before the access revocation, even if they did not store it.

In summary, a large majority of WAT users do not completely understand the actual process of data sharing with TPAs. Such a limited understanding could lead to an uninformed user making a decision that could have serious privacy implications. For example, a given user could share every type of data, without checking what a TPA actually does, while thinking that no previously collected data would be shared. In this way, the TPA will be able to collect much more fitness data than expected by the user in the first place, and even without their knowledge of it.

6 DISCUSSION

Our findings show that around seven out of ten WAT users in the US shared their fitness data, with at least one TPA (see RQ1). In line with the findings of a previous study [84], about half of the users underestimated the actual number of the TPAs to which they granted access to their fitness data. The two main reasons for not revoking accesses are due to the lack of concern about privacy and basic forgetfulness. Many respondents reported that they forgot about the accesses that they have previously granted, especially as they probably have stopped using the TPA (due to utility-related reasons). Indeed, after realizing that they were sharing more data than they thought, many respondents reported they plan to revoke some of their previously granted access authorizations.

Our results show that WAT users highly overestimate the availability of their personal information on their public profile (see RQ2). However, such lack of knowledge about their own privacy settings should not be too harmful, as their actual privacy levels tend to be higher than their estimations. Furthermore, the default privacy settings of their companion apps seem to substantially influence their current settings. Therefore, we recommend that WAT providers increase the default privacy level, as much as possible, in order to help their users preserve their privacy (i.e., opt-in). As for data sharing with other individuals, as expected given the existing literature on the topic, they tend to share data with friends and family members more than with other types of individuals (e.g., co-workers).

Our respondents positively perceived all three PETs we proposed in the survey (see RQ3), which is consistent with Murmann et al. findings [63]. However, when we asked them for their design suggestions, they only highlighted the importance of reminders and privacy checkups. They thought such reminders could effectively help them to recall and review their TPAs and to revoke the accesses they no longer use. A few respondents asked for more proactive and specialized privacy checkups, such as TPA access managers that periodically revoke access from unused TPAs then ask users to reconsider them to either renew or leave them (i.e., similar to what recent versions of Android do: they revoke permissions from unused apps [68]). Yet, some of the proposed solutions highlighted users' misconceptions about the functioning of the WAT-TPA ecosystem and were in fact not feasible. For example, the solution about privacy-checkup is feasible for Apple more than for Fitbit/Garmin, as Apple Health can interact with iOS to monitor the usage of both mobiles apps and TPAs. Finally, a few respondents mentioned interesting solutions about time-framed sharing for enabing users to define the time frame for the data they share.

Our findings on users' knowledge of data sharing (see RQ4) show that most of the WAT users have a limited understanding of the WAT-TPA ecosystem. Many respondents had incorrect mental models or they confused this eco-system with that of devices from other brands. Such incorrect mental models can induce other risky behaviors for privacy, such as sharing more data than is actually required or not regularly checking the previously granted permissions. The respondents were mainly confused about the temporal dimension of access management, they were uncertain about what could be done with their data before they grant accesses and after revoking them. This is a particularly risky belief, as many WAT users can grant access to their previously collected sensitive data, thinking that the TPAs will access only their new data. Our findings regarding mental models are relatively positive, compared to those from an earlier study [86]. The respondents in Velykoivanenko et al. [86]'s study were *fresh* WAT users (i.e., they began using WATs for the experiment and filled the questionnaire a few months afterwards),

where our respondents were experienced WAT users.[9] Our findings show that WAT users have poor knowledge about the data-sharing process. The implementation of transparency-enhancing technologies (TETs) [39] could be helpful in such case. For example, to help users improve their mental models when using their app, service providers could display visual information as drawings, thus representing where and how the collected data is transferred. Such a visualization method has been used in the past, for example, to help users understand privacy policies (Poli-see) [33]. Another solution would be to use our results to highlight the most problematic areas and to add information to help users better understand specific points about data sharing (e.g., clearly state that *"granting access to a TPA will cause sharing all the collected data without taking into account the sharing date."*)

Finally, more than one third of the respondents who submitted their drawings demonstrated their privacy concerns and their lack of trust in TPAs. Unfortunately, despite these privacy concerns, most WAT manufacturers (i.e., with their companion apps) do not take responsibility for actively supporting users against privacy threats with TPAs. Exceptionally, Apple is relatively restrictive about which TPAs their users can share their data with (e.g., they have to be fitness-oriented and have a clear privacy policy) [7]. However, the companion app's service provider does not provide substantial technical or legal support. For example, about data sharing with TPAs, Garmin privacy policy states only that *"once you direct us to share data with a third party, the third party's handling of your personal data is the responsibility of that third party, and you should carefully review the third party's privacy policy."*

In the case of data sharing, users' privacy is directly related to their behavior, as they voluntarily choose to share their data. However, we demonstrated users' general lack of awareness about how they should manage their TPAs (as they tend to forget what they granted access in the past), as well as their lack of knowledge about the functioning of WATs. Furthermore, our respondents demonstrated privacy concerns and a positive attitude toward PETs, which suggest that they want to improve their privacy. As their lack of awareness and knowledge is, at least partially, the reason for their risky behavior, helping them to improve their understanding of the whole data-sharing process (e.g., by implementing TETs in WAT apps) could be a promising approach for the adoption of less risky behavior.

7 LIMITATIONS

Our work has some limitations. First, all the respondents were TPA users, and as 70.2% of the WAT users are also TPA users, our respondents do not represent all of the WAT users. This should be noted, in particular, for questions related to data sharing on public profiles. Second, we asked the respondents to draw their mental models at the end of the survey; the drawing was optional. Our findings about mental models are relatively correct, compared to an earlier study [86]. This could be due to the self-selection bias problem [36], because our mental-model question was not mandatory, hence the respondents who were less confident or knowledgeable might have skipped this question. It is also possible that answering the survey could have influenced the respondents' knowledge about the system architecture (e.g., some questions refer to "servers"). This could have affected the respondents' mental model, but only in a positive way. Our results revealed an important lack of knowledge. Therefore, mental models would have probably been even worse if we had collected the drawings at the beginning of the survey and from all survey respondents. Third, when asking the respondents about how many of their TPAs they were actively using, they had to choose between "None," "Some," "Most," and "All." The boundary between "Some" and "Most" could lack clarity, as these terms do not represent a specific number or a ratio. However, "All" and "None" are distinct enough to support the presented results. Fourth, we should have calculated the minimum sample size by using power analysis to conduct the statistical analysis. But we relied on earlier similar studies and recruited slightly more. Nevertheless, we believe that our statistical test is valid, as an a posteriori power analysis by using G*Power 3.1 for a paired t-test revealed a high level of power (1.0), which means that it is highly likely we did not commit a type II error. Finally, the way we advertised the study (by referring to "fitness-data sharing") could have slightly biased the respondents and the recruitment process. However, we mentioned only data-sharing with TPAs, and avoided using the terms "privacy" and "security" to not prime the respondents.

8 CONCLUSION

Through a large-scale survey with $N = 628$ Apple Watch, Fitbit, and Garmin users in the US, this work contributes to the research area of wearable privacy by qualitatively and quantitatively analyzing WAT users' perceptions and data-sharing behaviors with third-party applications and individuals. Our analysis provides valuable insights to privacy researchers and practitioners to better understand WAT users and to design novel PETs for fitness-data sharing with TPAs.

As part of our future work, we will design (with a participatory approach) and evaluate such PETs, including—but not limited to—granularity reduction, time-framed sharing, automated access revocation, and access-monitoring reminders. We will also further explore the types of fitness data shared with TPAs. Finally, we will investigate, through longitudinal studies, how users grant and revoke accesses to their TPAs over time and will put it into perspective with their actual use of the TPAs.

ACKNOWLEDGMENTS

This work was partially funded by the Swiss National Science Foundation with Grant #200021_178978 (PrivateLife), by armasuisse S+T with Grant #CYD-C-2020007, and by the HEC Research Fund. We thank Pierre Huber for his help in designing the first draft of the survey. We thank Holly Cogliati and Vincent Vandersluis for proofreading this article. We also thank Gaël Bernard, James Tyler, Lev Velykoivanenko, Pooja Rao, and Yamane El Zein for participating in the cognitive pre-tests for the survey. Finally, we thank the Garmin Developer Program Support's team who kindly answered our questions.

[9] Note also that, in Velykoivanenko et al.'s study [86], they did not consider TPAs.

REFERENCES

[1] Ruba Abu-Salma, M. Angela Sasse, Joseph Bonneau, Anastasia Danilova, Alena Naiakshina, and Matthew Smith. 2017. Obstacles to the Adoption of Secure Communication Tools. In *IEEE Symp. on Security and Privacy (SP)*. IEEE, San Jose, CA, USA, 137–153. https://doi.org/10.1109/SP.2017.65

[2] Seyed Hossein Ahmadinejad and Philip W.L. Fong. 2013. On the Feasibility of Inference Attacks by Third-Party Extensions to Social Network Systems. In *Proc. of the ACM on Asia Conf. on Computer and Communications Security (AsiaCCS)*. ACM Press, Hangzhou, China, 161. https://doi.org/10.1145/2484313.2484333

[3] Seyed Hossein Ahmadinejad, Philip W.L. Fong, and Reihaneh Safavi-Naini. 2016. Privacy and Utility of Inference Control Mechanisms for Social Computing Applications. In *Proc. of the ACM on Asia Conf. on Computer and Communications Security (AsiaCCS)*. ACM, Xi'an China, 829–840. https://doi.org/10.1145/2897845.2897878

[4] Abdulmajeed Alqhatani and Heather Richter Lipford. 2019. "There Is Nothing That I Need to Keep Secret": Sharing Practices and Concerns of Wearable Fitness Data. In *Proc. of the USENIX Symp. on Usable Privacy and Security (SOUPS)*. USENIX Association, Santa Clara, CA, USA, 421–434.

[5] Pauline Anthonysamy, Awais Rashid, James Walkerdine, Phil Greenwood, and Georgios Larkou. 2012. Collaborative Privacy Management for Third-Party Applications in Online Social Networks. In *Proc. of the Workshop on Privacy and Security in Online Social Media (PSOSM)*. ACM Press, Lyon, France, 1–4. https://doi.org/10.1145/2185354.2185359

[6] Apple. 2020. HealthKit | Apple Developer Documentation. https://developer.apple.com/documentation/healthkit.

[7] Apple. 2022. Legal - Data & Privacy - Apple. https://www.apple.com/legal/privacy/data/en/health-app/.

[8] Amid Ayobi, Paul Marshall, Anna L. Cox, and Yunan Chen. 2017. Quantifying the Body and Caring for the Mind: Self-Tracking in Multiple Sclerosis. In *Proc. of the CHI Conf. on Human Factors in Computing Systems (CHI)*. Association for Computing Machinery, New York, NY, USA, 6889–6901. https://doi.org/10.1145/3025453.3025869

[9] Igor Bilogrevic and Martin Ortlieb. 2016. "If You Put All The Pieces Together...": Attitudes Towards Data Combination and Sharing Across Services and Companies. In *Proc. of the CHI Conf. on Human Factors in Computing Systems (CHI)*. ACM, San Jose California USA, 5215–5227. https://doi.org/10.1145/2858036.2858432

[10] Alex Bowden. 2018. Cyclist Who Had Five Bikes Stolen Says Thieves Are Looking for Quick Times on Strava to Try and Find High-End Bikes – Warns Other Users to Check Their Privacy Settings. https://road.cc/content/news/248798-cyclist-who-had-five-bikes-stolen-says-thieves-are-looking-quick-times-strava.

[11] Business Wire. 2020. Shipments of Wearable Devices Leap to 125 Million Units, Up 35.1% in the Third Quarter, According to IDC. https://www.businesswire.com/news/home/20201202005304/en/Shipments-of-Wearable-Devices-Leap-to-125-Million-Units-Up-35.1-in-the-Third-Quarter-According-to-IDC.

[12] Angela Chen. 2018. What Happens When Life Insurance Companies Track Fitness Data? https://www.theverge.com/2018/9/26/17905390/john-hancock-life-insurance-fitness-tracker-wearables-science-health.

[13] Yuan Cheng, Jaehong Park, and Ravi Sandhu. 2013. Preserving User Privacy from Third-Party Applications in Online Social Networks. In *Proc. of the International Conference on World Wide Web - WWW '13 Companion*. ACM Press, Rio de Janeiro, Brazil, 723–728. https://doi.org/10.1145/2487788.2488032

[14] Eun Kyoung Choe, Nicole B Lee, Bongshin Lee, Wanda Pratt, and Julie A Kientz. 2014. Understanding Quantified-Selfers' Practices in Collecting and Exploring Personal Data. In *Proc. of the Conf. on Human Factors in Computing Systems (CHI)*. ACM, Toronto Ontario Canada, 1143–1152. https://doi.org/10.1145/2556288.2557372

[15] Blaine Cook and Chris Messina. 2012. OAuth 2.0 — OAuth. https://oauth.net/2/.

[16] Kenan Degirmenci. 2020. Mobile Users' Information Privacy Concerns and the Role of App Permission Requests. *International Journal of Information Management* 50 (Feb. 2020), 261–272. https://doi.org/10.1016/j.ijinfomgt.2019.05.010

[17] Jaime Delgado, Eva Rodríguez, and Silvia Llorente. 2010. User's Privacy in Applications Provided through Social Networks. In *Proc. of the ACM SIGMM Workshop on Social Media (WSM)*. ACM Press, Firenze, Italy, 39. https://doi.org/10.1145/1878151.1878163

[18] Simon Eberz, Giulio Lovisotto, Andrea Patane, Marta Kwiatkowska, Vincent Lenders, and Ivan Martinovic. 2018. When Your Fitness Tracker Betrays You: Quantifying the Predictability of Biometric Features Across Contexts. In *S&P*. IEEE, San Francisco, CA, 889–905. https://doi.org/10.1109/SP.2018.00053

[19] Simon Eberz, Nicola Paoletti, Marc Roeschlin, Andrea Patani, Marta Kwiatkowska, and Ivan Martinovic. 2017. Broken Hearted: How To Attack ECG Biometrics. In *Proc. of the Network and Distributed System Security Symp. (NDSS)*. Internet Society, San Diego, CA. https://doi.org/10.14722/ndss.2017.23408

[20] Haroon Elahi, Guojun Wang, and Dongqing Xie. 2017. Assessing Privacy Behaviors of Smartphone Users in the Context of Data Over-Collection Problem: An Exploratory Study. In *2017 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computed, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCom/IOP/SCI)*. IEEE, San Francisco, CA, 1–8. https://doi.org/10.1109/UIC-ATC.2017.8397613

[21] William Enck, Peter Gilbert, Byung-Gon Chun, Landon P. Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N. Sheth. 2014. TaintDroid: An Information Flow Tracking System for Real-Time Privacy Monitoring on Smartphones. *Commun. ACM* 57, 3 (March 2014), 99–106. https://doi.org/10.1145/2494522

[22] Daniel A. Epstein, Alan Borning, and James Fogarty. 2013. Fine-Grained Sharing of Sensed Physical Activity: A Value Sensitive Approach. In *Proc. of the ACM Int'l Joint Conf. on Pervasive and Ubiquitous Computing (UbiComp)*. Association for Computing Machinery, New York, NY, USA, 489–498. https://doi.org/10.1145/2493432.2493433

[23] Shehroze Farooqi and Zubair Shafiq. 2019. Measurement and Early Detection of Third-Party Application Abuse on Twitter. In *The World Wide Web Conference on - WWW '19*. ACM Press, San Francisco, CA, USA, 448–458. https://doi.org/10.1145/3308558.3313515

[24] Christina Farr. 2019. Fitbit Has a New Health Tracker, but You Can Only Get It through Your Employer or Insurer. https://www.cnbc.com/2019/02/08/fitbit-releases-insprire-for-employers.html.

[25] Fitbit. 2020. Fitbit SDK. https://dev.fitbit.com/.

[26] Marco Furini, Silvia Mirri, Manuela Montangero, and Catia Prandi. 2020. Can IoT Wearable Devices Feed Frugal Innovation?. In *Proc. of the Workshop on Experiences with the Design and Implementation of Frugal Smart Objects (FRUGALTHINGS)*. Association for Computing Machinery, New York, NY, USA, 1–6. https://doi.org/10.1145/3410670.3410861

[27] Sandra Gabriele and Sonia Chiasson. 2020. Understanding Fitness Tracker Users' Security and Privacy Knowledge, Attitudes and Behaviours. In *Proc. of the Conf. on Human Factors in Computing Systems (CHI)*. ACM, Honolulu HI USA, 1–12. https://doi.org/10.1145/3313831.3376651

[28] Sandra Gabriele and Sonia Chiasson. 2020. Understanding Fitness Tracker Users' Security and Privacy Knowledge, Attitudes and Behaviours. In *Proc. of the CHI Conf. on Human Factors in Computing Systems (CHI)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3313831.3376651

[29] Garmin. 2020. Overview | Garmin Connect Developer Program | Garmin Developers. https://developer.garmin.com/gc-developer-program/overview/.

[30] Jennifer Golbeck and Matthew Mauriello. 2016. User Perception of Facebook App Data Access: A Comparison of Methods and Privacy Concerns. *Future Internet* 8, 4 (March 2016), 9. https://doi.org/10.3390/fi8020009

[31] Nanna Gorm and Irina Shklovski. 2016. Sharing Steps in the Workplace: Changing Privacy Concerns Over Time. In *Proc. of the Conf. on Human Factors in Computing Systems (CHI)*. ACM, San Jose, California, USA, 4315–4319. https://doi.org/10.1145/2858036.2858352

[32] Gabriel Guo, Hanbin Zhang, Liuyi Yao, Huining Li, Chenhan Xu, Zhengxiong Li, and Wenyao Xu. 2022. MSLife: Digital Behavioral Phenotyping of Multiple Sclerosis Symptoms in the Wild Using Wearables and Graph-Based Statistical Analysis. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 4 (Dec. 2022), 158:1–158:35. https://doi.org/10.1145/3494970

[33] Wentao Guo, Jay Rodolitz, and Eleanor Birrell. 2020. Poli-See: An Interactive Tool for Visualizing Privacy Policies. In *Proc. of the Workshop on Privacy in the Electronic Society (WPES)*. Association for Computing Machinery, New York, NY, USA, 57–71. https://doi.org/10.1145/3411497.3420221

[34] Mario A. Gutierrez, Michelle L. Fast, Anne H. Ngu, and Byron J. Gao. 2016. Real-Time Prediction of Blood Alcohol Content Using Smartwatch Sensor Data. In *Smart Health*, Xiaolong Zheng, Daniel Dajun Zeng, Hsinchun Chen, and Scott J. Leischow (Eds.). Springer International Publishing, 175–186.

[35] Hamza Harkous, Kassem Fawaz, Rémi Lebret, Florian Schaub, Kang G. Shin, and Karl Aberer. 2018. Polisis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning. In *27th USENIX Security Symposium (USENIX Security 18)*. 531–548.

[36] James J Heckman. 1990. Selection bias and self-selection. In *Econometrics*. Springer, 201–224.

[37] Alex Hern. 2018. Fitness Tracking App Strava Gives Away Location of Secret US Army Bases. https://www.theguardian.com/world/2018/jan/28/fitness-tracking-app-gives-away-location-of-secret-us-army-bases.

[38] Robert P Hirten, Matteo Danieletto, Lewis Tomalin, Katie Hyewon Choi, Eddye Golden, Sparshdeep Kaur, Drew Helmus, Anthony Biello, Alexander Charney, Riccardo Miotto, Benjamin S Glicksberg, Ismail Nabeel, Judith Aberg, David Reich, Dennis Charney, Laurie Keefer, Mayte Suarez-Farinas, Girish N Nadkarni, and Zahi A Fayad. 2021. Physiological Data from a Wearable Device Identifies SARS-CoV-2 Infection and Symptoms and Predicts COVID-19 Diagnosis: Observational Study. *Journal of Medical Internet Research* (2021), 36.

[39] Milena Janic, Jan Pieter Wijbenga, and Thijs Veugen. 2013. Transparency Enhancing Tools (TETs): An Overview. In *2013 Third Workshop on Socio-Technical Aspects in Security and Trust*. IEEE, New Orleans, LA, 18–25. https://doi.org/10.1109/STAST.2013.11

[40] David Jonassen and Young Hoan Cho. 2008. Externalizing Mental Models with Mindtools. In *Understanding Models for Learning and Instruction*, Dirk Ifenthaler,

Pablo Pirnay-Dummer, and J. Michael Spector (Eds.). Springer US, Boston, MA, 145–159. https://doi.org/10.1007/978-0-387-76898-4_7

[41] Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara Kiesler. 2015. "My Data Just Goes Everywhere:" User Mental Models of the Internet and Implications for Privacy and Security. In *Proc. of the USENIX Symp. on Usable Privacy and Security (SOUPS)*. USENIX Association, Ottawa, Canada, 39–52.

[42] Jennifer King, Airi Lampinen, and Alex Smolen. 2011. Privacy: Is There an App for That?. In *Proc. of the USENIX Symp. on Usable Privacy and Security (SOUPS)*. ACM Press, Pittsburgh, Pennsylvania, 1. https://doi.org/10.1145/2078827.2078843

[43] Hanna Krasnova, Nicole Eling, Oleg Schneider, Helena Wenninger, Thomas Widjaja, and Peter Buxmann. 2013. Does This App Ask For Too Much Data? The Role Of Privacy Perceptions In User Behavior Towards Facebook Applications And Permission Dialogs. *ECIS* (2013), 14.

[44] K. Krombholz, K. Busse, K. Pfeffer, M. Smith, and E. von Zezschwitz. 2019. "If HTTPS Were Secure, I Wouldn't Need 2FA" - End User and Administrator Mental Models of HTTPS. In *Proc. of the IEEE Symp. on Security and Privacy (S&P)*. IEEE, San Francisco, CA, USA, 1138–1155. https://doi.org/10.1109/SP.2019.00060

[45] Preeti Kumari, Lini Mathew, and Poonam Syal. 2017. Increasing Trend of Wearables and Multimodal Interface for Human Activity Monitoring: A Review. *Biosensors and Bioelectronics* 90 (April 2017), 298–307. https://doi.org/10.1016/j.bios.2016.12.001

[46] He Li, Jing Wu, Yiwen Gao, and Yao Shi. 2016. Examining Individuals' Adoption of Healthcare Wearable Devices: An Empirical Study from Privacy Calculus Perspective. *International Journal of Medical Informatics* 88 (April 2016), 8–17. https://doi.org/10.1016/j.ijmedinf.2015.12.010

[47] Wanpeng Li and Chris J. Mitchell. 2014. Security Issues in OAuth 2.0 SSO Implementations. In *Information Security (Lecture Notes in Computer Science)*, Sherman S. M. Chow, Jan Camenisch, Lucas C. K. Hui, and Siu Ming Yiu (Eds.). Springer International Publishing, Cham, 529–541. https://doi.org/10.1007/978-3-319-13257-0_34

[48] X. Li, J. Xu, Z. Zhang, X. Lan, and Y. Wang. 2020. Modular Security Analysis of OAuth 2.0 in the Three-Party Setting. In *Euro S&P*. 276–293. https://doi.org/10.1109/EuroSP48549.2020.00025

[49] Y. Liang, Z. Cai, J. Yu, Q. Han, and Y. Li. 2018. Deep Learning Based Inference of Private Information Using Embedded Sensors in Smart Devices. *IEEE Network* 32, 4 (July 2018), 8–14. https://doi.org/10.1109/MNET.2018.1700349

[50] Yuting Liao. 2019. Sharing Personal Health Information on Social Media: Balancing Self-presentation and Privacy. In *Proc. of the Int'l Conf. on Social Media and Society (SMSociety)*. Association for Computing Machinery, New York, NY, USA, 194–204. https://doi.org/10.1145/3328529.3328560

[51] Andrés Lucero. 2015. Using affinity diagrams to evaluate interactive prototypes. In *IFIP conference on human-computer interaction*. Springer, 231–248.

[52] Deborah Lupton. 2021. "Sharing Is Caring:" Australian Self-Trackers' Concepts and Practices of Personal Data Sharing and Privacy. *Frontiers in Digital Health* 3 (2021). https://doi.org/10.3389/fdgth.2021.649275

[53] Alexandra Mai, Katharina Pfeffer, Matthias Gusenbauer, Edgar Weippl, and Katharina Krombholz. 2020. User Mental Models of Cryptocurrency Systems - A Grounded Theory Approach. In *Proc. of the USENIX Symp. on Usable Privacy and Security (SOUPS)*. USENIX Association, USA, 341–358.

[54] Anindya Maiti, Oscar Armbruster, Murtuza Jadliwala, and Jibo He. 2016. Smartwatch-Based Keystroke Inference Attacks and Context-Aware Protection Mechanisms. In *Proc. of the ACM on Asia Conf. on Computer and Communications Security (AsiaCCS)*. ACM, Xi'an, China, 795–806. https://doi.org/10.1145/2897845.2897905

[55] Anindya Maiti, Murtuza Jadliwala, Jibo He, and Igor Bilogrevic. 2015. (Smart)Watch Your Taps: Side-channel Keystroke Inference Attacks Using Smartwatches. In *Proc. of the ACM Int. Symp. on Wearable Computers (ISWC)*. ACM, Osaka, Japan, 27–30. https://doi.org/10.1145/2802083.2808397

[56] Anindya Maiti, Murtuza Jadliwala, J. He, and I. Bilogrevic. 2018. Side-Channel Inference Attacks on Mobile Keypads Using Smartwatches. *IEEE Transactions on Mobile Computing* 17, 9 (Sept. 2018), 2180–2194. https://doi.org/10.1109/TMC.2018.2794984

[57] Naresh K. Malhotra, Sung S. Kim, and James Agarwal. 2004. Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model. *Information Systems Research* 15, 4 (Dec. 2004), 336–355. https://doi.org/10.1287/isre.1040.0032

[58] Anna Maria Mandalari, Daniel J. Dubois, Roman Kolcun, Muhammad Talha Paracha, Hamed Haddadi, and David Choffnes. 2021. Blocking Without Breaking: Identification and Mitigation of Non-Essential IoT Traffic. *Proceedings on Privacy Enhancing Technologies* 2021, 4 (Oct. 2021), 369–388. https://doi.org/10.2478/popets-2021-0075

[59] K. I. Manktelow and Man Cheung Chung (Eds.). 2004. *Psychology of Reasoning: Theoretical and Historical Perspectives* (first ed.). Psychology Press, Hove ; New York.

[60] Stefania Marassi and Philippa Collins. 2021. Is That Lawful? Data Privacy and Fitness Trackers in the Workplace. *International Journal of Comparative Labour Law and Industrial Relations* 37, 1 (Feb. 2021).

[61] Maximilian Marsch, Jens Grossklags, and Sameer Patil. 2021. Won't You Think of Others?: Interdependent Privacy in Smartphone App Permissions. *Proc. of the ACM on Human-Computer Interaction* 5, CSCW2 (Oct. 2021), 437:1–437:35. https://doi.org/10.1145/3479581

[62] Tenga Matsuura, Ayako A. Hasegawa, Mitsuaki Akiyama, and Tatsuya Mori. 2021. Careless Participants Are Essential for Our Phishing Study: Understanding the Impact of Screening Methods. In *European Symposium on Usable Security 2021*. ACM, Karlsruhe Germany, 36–47. https://doi.org/10.1145/3481357.3481515

[63] Patrick Murmann, Matthias Beckerle, Simone Fischer-Hübner, and Delphine Reinhardt. 2021. Reconciling the What, When and How of Privacy Notifications in Fitness Tracking Scenarios. *Pervasive and Mobile Computing* 77 (Oct. 2021), 101480. https://doi.org/10.1016/j.pmcj.2021.101480

[64] K. Niazmand, K. Tonn, Y. Zhao, U. M. Fietzek, F. Schroeteler, K. Ziegler, A. O. Ceballos-Baumann, and T. C. Lueth. 2011. Freezing of Gait Detection in Parkinson's Disease Using Accelerometer Based Smart Clothes. In *IEEE Biomedical Circuits and Systems Conf. (BioCAS)*. 201–204. https://doi.org/10.1109/BioCAS.2011.6107762

[65] Mehdi Nobakht, Yulei Sui, Aruna Seneviratne, and Wen Hu. 2020. PGFit: Static Permission Analysis of Health and Fitness Apps in IoT Programming Frameworks. *Journal of Network and Computer Applications* 152 (Feb. 2020), 102509. https://doi.org/10.1016/j.jnca.2019.102509

[66] Maggie Oates, Yama Ahmadullah, Abigail Marsh, Chelse Swoopes, Shikun Zhang, Rebecca Balebako, and Lorrie Faith Cranor. 2018. Turtles, Locks, and Bathrooms: Understanding Mental Models of Privacy Through Illustration. *Proc. on Privacy Enhancing Technologies (PoPETs)* 2018, 4 (Oct. 2018), 5–32. https://doi.org/10.1515/popets-2018-0029

[67] Stefan Palan and Christian Schitter. 2018. Prolific.Ac—A Subject Pool for Online Experiments. *Journal of Behavioral and Experimental Finance* 17 (March 2018), 22–27. https://doi.org/10.1016/j.jbef.2017.12.004

[68] Rajesh Pandey. 2020. Android 11 Will Automatically Revoke Permissions from Unused Apps. https://www.neowin.net/news/android-11-will-automatically-revoke-permissions-from-unused-apps/.

[69] Jamie Pinchot and Donna Cellante. 2021. Privacy Concerns and Data Sharing Habits of Personal Fitness Information Collected via Activity Trackers. *JISAR* 14, 2 (June 2021), 4.

[70] Rocket Fuel. 2014. *'Quantified Self' Digital Tools: A CPG Marketing Opportunity*. Technical Report. Rocket Fuel.

[71] Christopher Rowl. 2019. With Fitness Trackers in the Workplace, Bosses Can Monitor Your Every Step - and Possibly More. https://www.washingtonpost.com/business/economy/with-fitness-trackers-in-the-workplace-bosses-can-monitor-your-every-step–and-possibly-more/2019/02/15/75ee0848-2a45-11e9-b011-d8500644dc98_story.html.

[72] Mohd Sabra, Anindya Maiti, and Murtuza Jadliwala. 2018. Keystroke Inference Using Ambient Light Sensor on Wrist-Wearables: A Feasibility Study. In *Proc. of the ACM Workshop on Wearable Systems and Applications (WearSys)*. ACM, Munich, Germany, 21–26. https://doi.org/10.1145/3211960.3211973

[73] Johnny Saldana. 2021. *The Coding Manual for Qualitative Researchers* (4th ed ed.). SAGE Publishing, Thousand Oaks, California.

[74] Stefan Schneegass, Romina Poguntke, and Tonja Machulla. 2019. Understanding the Impact of Information Representation on Willingness to Share Information. In *Proc. of the CHI Conf. on Human Factors in Computing Systems (CHI)*. Association for Computing Machinery, New York, NY, USA, 1–6. https://doi.org/10.1145/3290605.3300753

[75] Stefan Schneegass, Romina Poguntke, and Tonja Machulla. 2019. Understanding the Impact of Information Representation on Willingness to Share Information. In *Proc. of the Conf. on Human Factors in Computing Systems (CHI)*. ACM, Glasgow, Scotland Uk, 1–6. https://doi.org/10.1145/3290605.3300753

[76] Mohamed Shehab, Said Marouf, and Christopher Hudel. 2011. ROAuth: Recommendation Based Open Authorization. In *Proc. of the USENIX Symp. on Usable Privacy and Security (SOUPS)*. ACM Press, Pittsburgh, Pennsylvania, 1. https://doi.org/10.1145/2078827.2078842

[77] Anita Valanju Shelgikar, Patricia F. Anderson, and Marc R. Stephens. 2016. Sleep Tracking, Wearable Technology, and Opportunities for Research and Clinical Care. *Chest* 150, 3 (Sept. 2016), 732–743. https://doi.org/10.1016/j.chest.2016.04.016

[78] Sheng Shen, He Wang, and Romit Roy Choudhury. 2016. I Am a Smartwatch and I Can Track My User's Arm. In *Proc. of the Annual Int. Conf. on Mobile Systems, Applications, and Services (MobiSys)*. ACM, Singapore, Singapore, 85–96. https://doi.org/10.1145/2906388.2906407

[79] Muhammad Shoaib, Ozlem Durmaz Incel, Hans Scholten, and Paul Havinga. 2018. SmokeSense: Online Activity Recognition Framework on Smartwatches. In *Mobile Computing, Applications, and Services*, Kazuya Murao, Ren Ohmura, Sozo Inoue, and Yusuke Gotoh (Eds.). Vol. 240. Springer International Publishing, Cham, 106–124. https://doi.org/10.1007/978-3-319-90740-6_7

[80] Stephanie L. Silveira, Jessica F. Baird, and Robert W. Motl. 2021. Rates, Patterns, and Correlates of Fitness Tracker Use among Older Adults with Multiple Sclerosis. *Disability and Health Journal* 14, 1 (Jan. 2021), 100966. https://doi.org/10.1016/j.dhjo.2020.100966

[81] Statista. 2022. Wearable Band Market Share in North America by Vendor 2018-2020. https://www.statista.com/statistics/1042044/north-america-quarterly-wearable-band-market-share-by-vendor/.
[82] Etye Steinberg. 2021. Run for Your Life: The Ethics of Behavioral Tracking in Insurance. *Journal of Business Ethics* (June 2021). https://doi.org/10.1007/s10551-021-04863-8
[83] Ilaria Torre, Frosina Koceva, Odnan Ref Sanchez, and Giovanni Adorni. 2016. Fitness Trackers and Wearable Devices: How to Prevent Inference Risks?. In *Proc. of the EAI Conf. on Body Area Networks (BODYNETS)*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), ICST, Brussels, Belgium, Belgium, 125–131.
[84] Ilaria Torre, Odnan Ref Sanchez, Frosina Koceva, and Giovanni Adorni. 2018. Supporting Users to Take Informed Decisions on Privacy Settings of Personal Devices. *Personal and Ubiquitous Computing* 22, 2 (April 2018), 345–364. https://doi.org/10.1007/s00779-017-1068-3
[85] J D Tygar and Marti Hearst. 2006. Why Phishing Works. (2006), 10.
[86] Lev Velykoivanenko, Kavous Salehzadeh Niksirat, Noé Zufferey, Mathias Humbert, Kévin Huguenin, and Mauro Cherubini. 2021. Are Those Steps Worth Your Privacy?: Fitness-Tracker Users' Perceptions of Privacy and Utility. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 4 (Dec. 2021), 1–41. https://doi.org/10.1145/3494960
[87] Na Wang. 2012. Third-Party Applications' Data Practices on Facebook. In *Proc. of the CHI Conf. on Human Factors in Computing Systems (CHI)*. ACM, Austin Texas USA, 1399–1404. https://doi.org/10.1145/2212776.2212462
[88] Na Wang, Heng Xu, and Jens Grossklags. 2011. Third-Party Apps on Facebook: Privacy and the Illusion of Control. In *Proc. of the ACM Symposium on Computer Human Interaction for Management of Information Technology (CHIMIT)*. ACM Press, Cambridge, Massachusetts, 1–10. https://doi.org/10.1145/2076444.2076448
[89] Rick Wash and Emilee Rader. 2011. Influencing Mental Models of Security: A Research Agenda. In *Proc. of the Conf. New Security Paradigms Workshop (NSPW)*. Association for Computing Machinery, Marin County, California, USA, 57–66. https://doi.org/10.1145/2073276.2073283
[90] Gary M. Weiss, Jessica L. Timko, Catherine M. Gallagher, Kenichi Yoneda, and Andrew J. Schreiber. 2016. Smartwatch-Based Activity Recognition: A Machine Learning Approach. In *IEEE-EMBS Int. Conf. on Biomedical and Health Informatics (BHI)*. IEEE, Las Vegas, NV, USA, 426–429. https://doi.org/10.1109/BHI.2016.7455925
[91] Pamela Wisniewski, Heng Xu, Heather Lipford, and Emmanuel Bello-Ogunu. 2015. Facebook Apps and Tagging: The Trade-off between Personal Privacy and Engaging with Friends: Facebook Apps and Tagging: The Trade-off Between Personal Privacy and Engaging with Friends. *Journal of the Association for Information Science and Technology* 66, 9 (Sept. 2015), 1883–1896. https://doi.org/10.1002/asi.23299
[92] Verena M. Wottrich, Eva A. van Reijmersdal, and Edith G. Smit. 2018. The Privacy Trade-off for Mobile App Downloads: The Roles of App Value, Intrusiveness, and Privacy Concerns. *Decision Support Systems* 106 (Feb. 2018), 44–52. https://doi.org/10.1016/j.dss.2017.12.003

A PRECISION FUNCTIONALITY

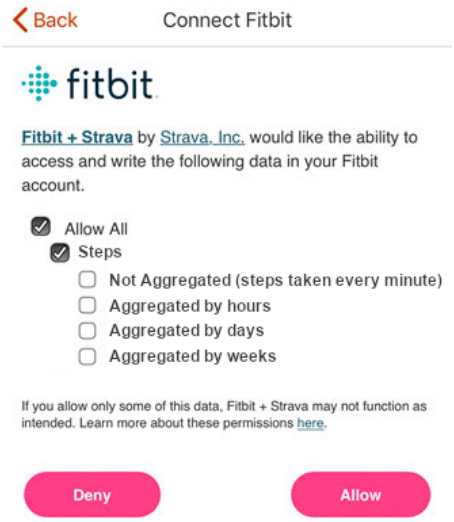

Figure 11: Illustration of the time granularity feature where a user can choose the aggregation level of the data they share with TPAs. "Not Aggregated (every minute)" is the default option on most WAT apps.

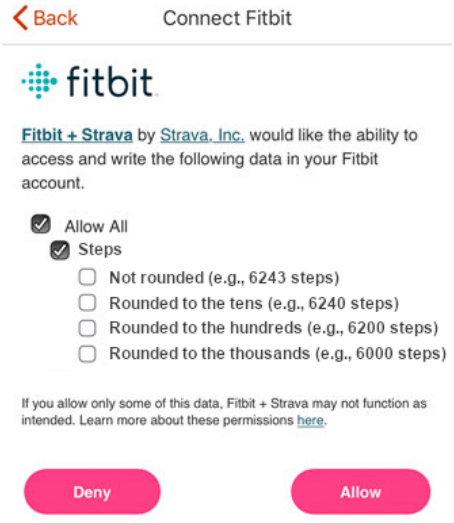

Figure 12: The figure shows the precision functionality where a user can choose the precision level of the data they want to share. "Not rounded" is the default option of most on the WAT apps.

B RESPONDENTS' PRIVACY CONCERN

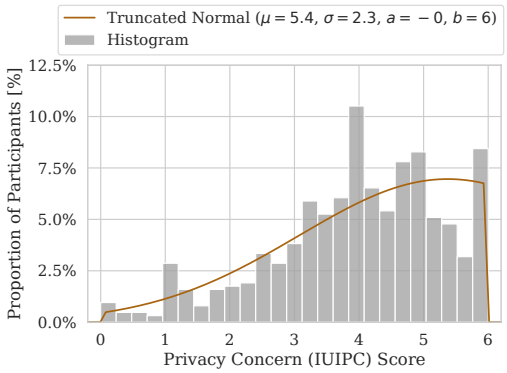

Figure 13: Data collection privacy concern (IUIPC) w/ fit.

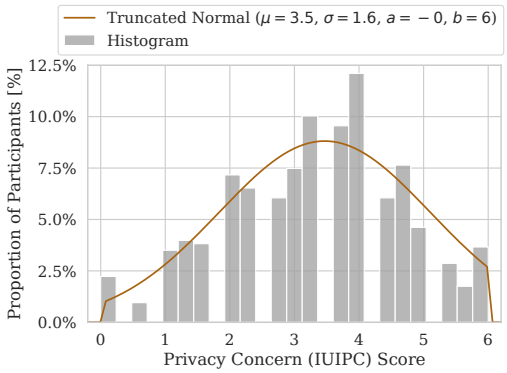

Figure 14: Global information privacy concern (IUIPC) w/ fit.

C FITNESS DATA SHARING

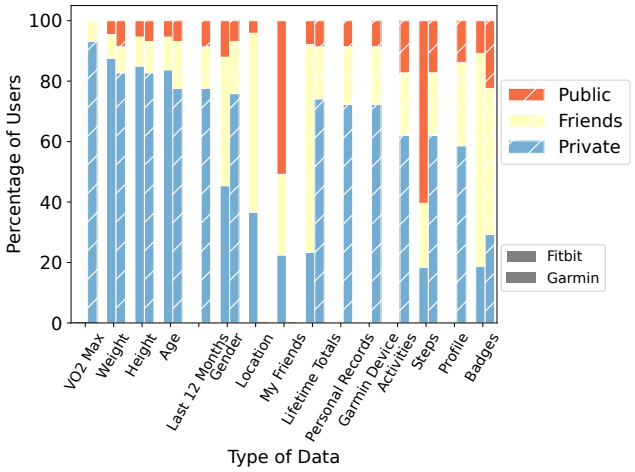

Figure 15: Privacy level of different profile information for Fitbit and Garmin users. Here, we decided to refer to similar concepts of both service providers using the Garmin's labels (e.g., "Badges" and "Badges and Trophies"), and to use Fitbit's privacy labels and to refer to both Garmin's "My Connections" and "My Groups and Connection" as "Friends".

D MENTAL MODELS

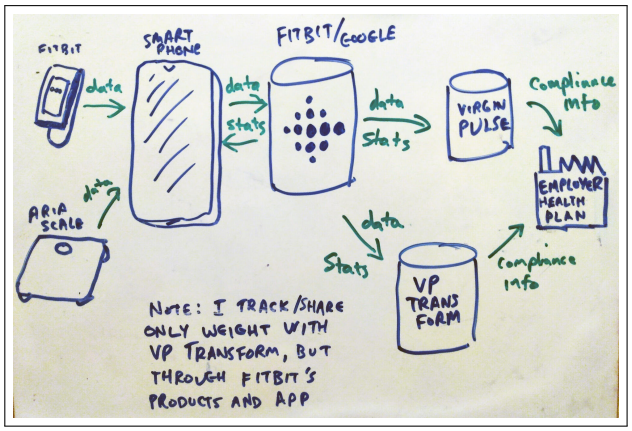

(a) The first type of the mental models (mm_1): The fitness data is transmitted to TPAs via a connected device and the WAT server.

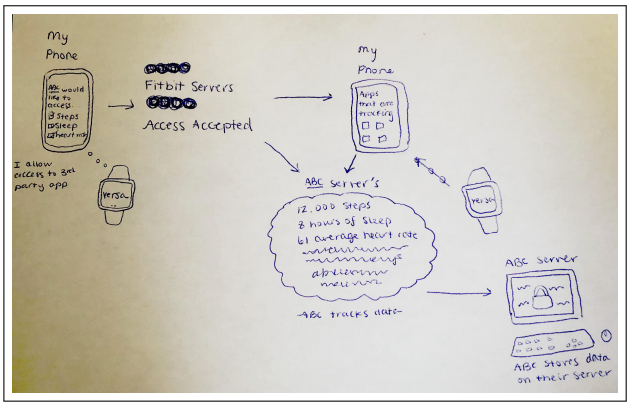

(b) The first type of the mental models (mm_1): The fitness data is transmitted to TPAs via a connected device and the WAT server.

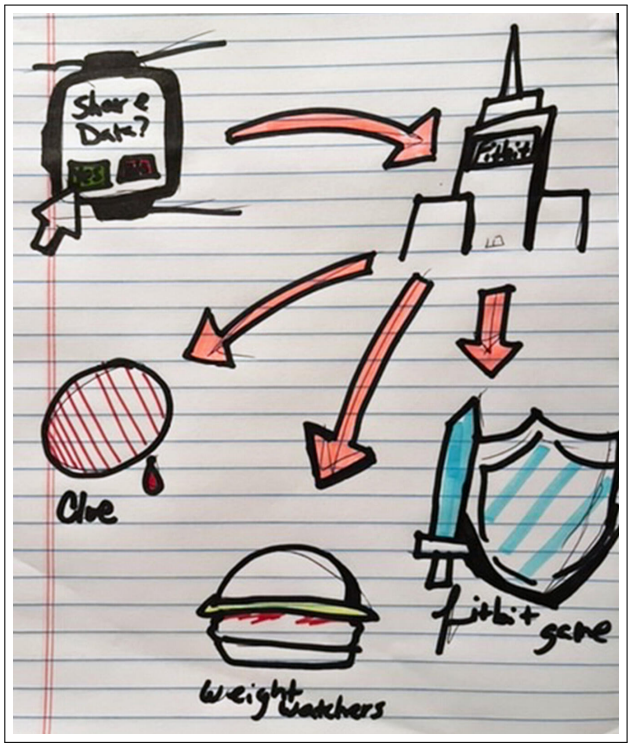

(c) The second type of the mental models (mm_2): The fitness data is transmitted without passing via a connected device.

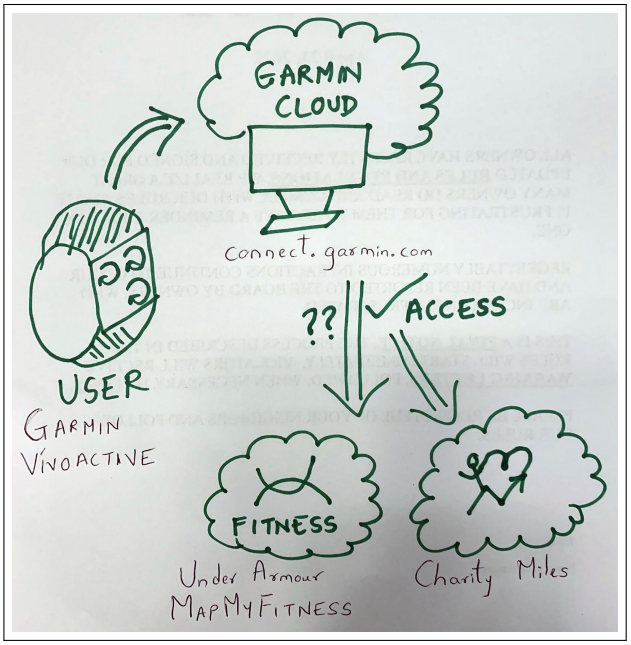

(d) The second type of the mental models (mm_2): The fitness data is transmitted without passing via a connected device.

Figure 16: Examples of users' mental models - 1

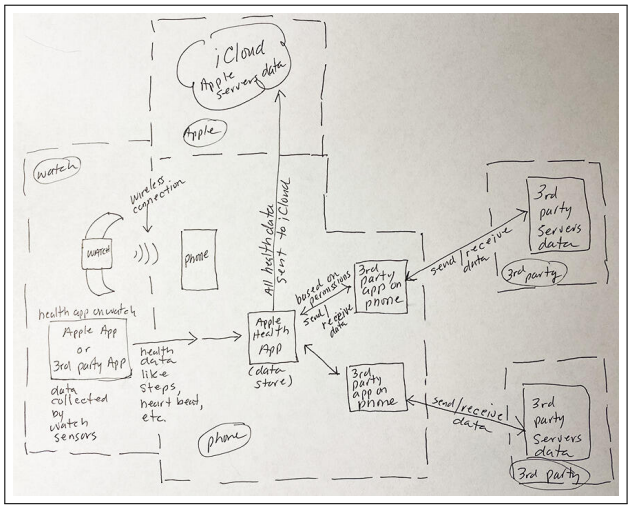

(a) The third type of the mental models (mm_3): a local synchronization between the TPA and the companion app.

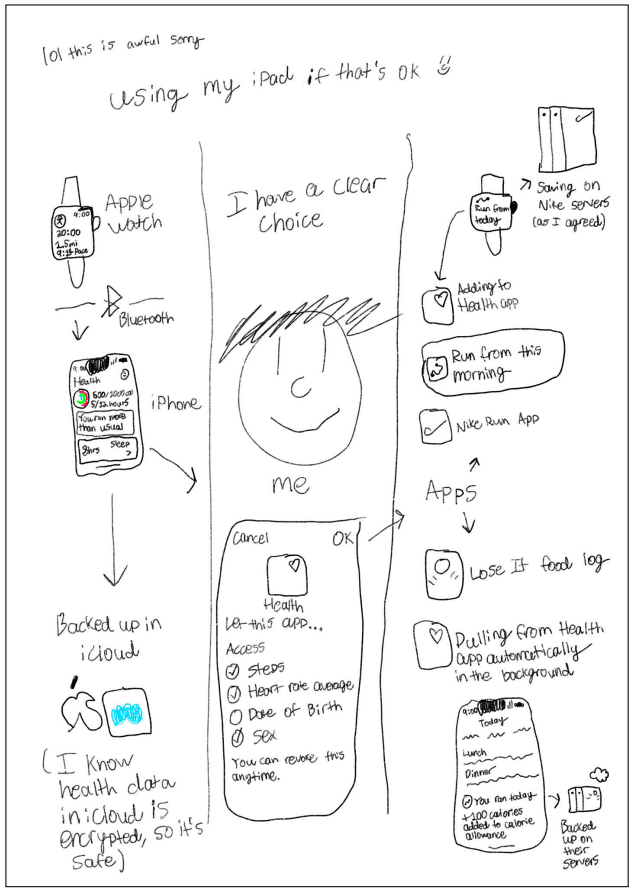

(b) The third type of the mental models (mm_3): a local synchronization between the TPA and the companion app. The respondent also is aware of selective sharing.

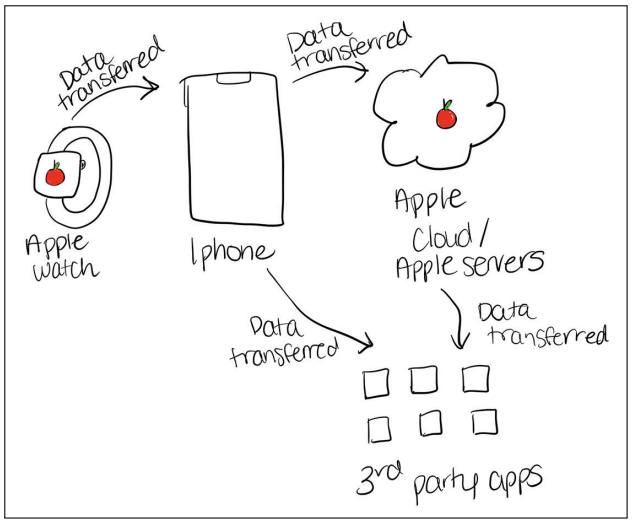

(c) An example of an inaccurate mental model that combines mm_1 with mm_3.

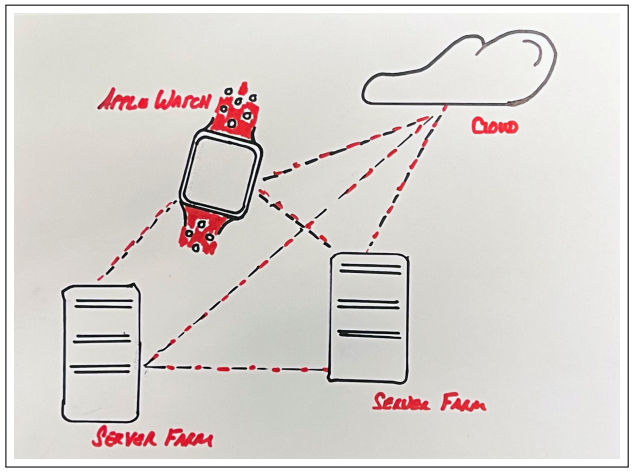

(d) An example of mm_4 (i.e., an incorrect mental model): This drawing cannot be attributed to any of the mm_1, mm_2, mm_3 models.

Figure 17: Examples of users' mental models - 2

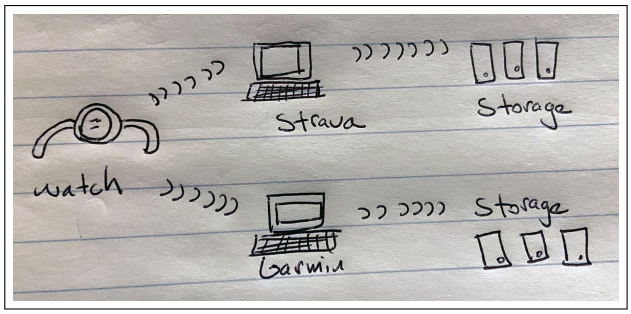

(a) An example of mm_4 (i.e., an incorrect mental model): This drawing cannot be attributed to any of the mm_1, mm_2, mm_3 models.

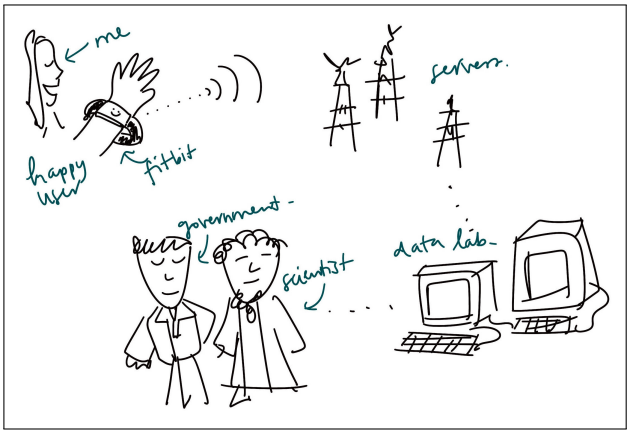

(b) An example mental model that shows a respondent thinks the fitness data is shared with 'scientists', 'data labs', and 'government.'

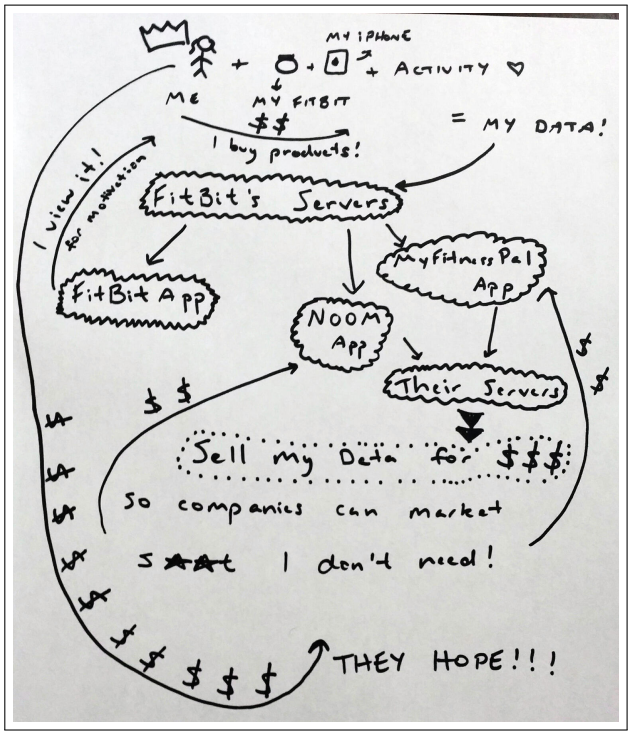

(c) An example mental model that shows a respondent thinks TPAs sell data for monetary benefits.

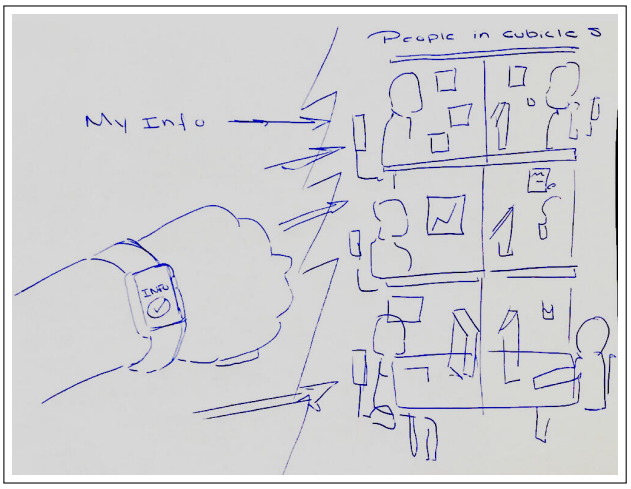

(d) An example mental model that shows a respondent thinks fitness data is further analyzed and scrutinized by a TPA company.

Figure 18: Examples of users' mental models - 3

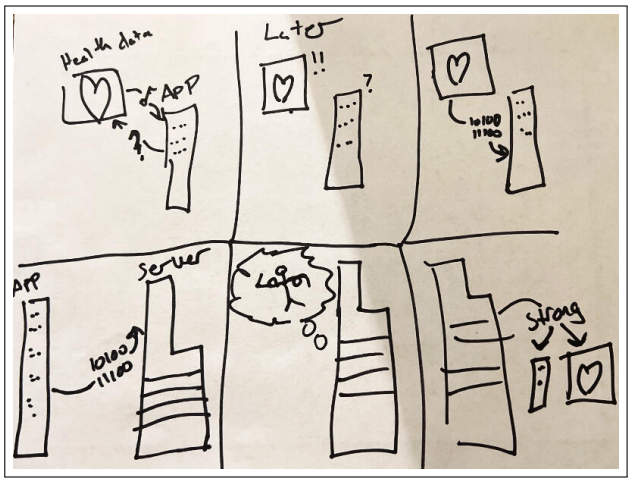

(a) An example mental model that shows a respondent thinks TPA will make their profile based on the fitness data.

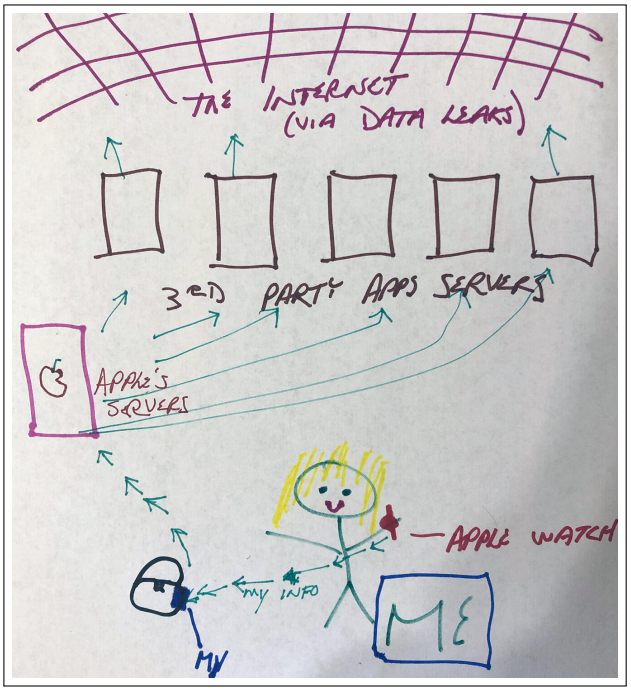

(b) An example mental model that shows a respondent is concerned about the network security of TPAs (i.e., possible privacy breach).

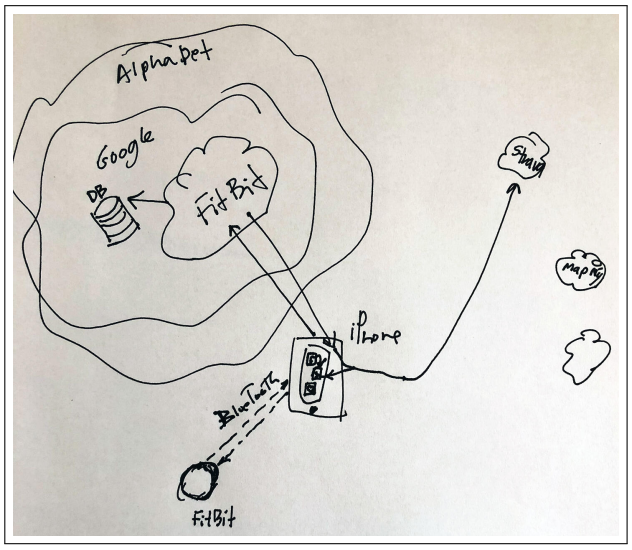

(c) An example mental model that shows a respondent thinks that the WAT company (i.e., Fitbit) can share the data with its affiliated giant company (i.e., Alphabet's Google).

(d) An example mental model that shows a respondent is informed about granting and revoking access. The example also shows that respondent believes the data will be deleted from TPA servers after they revoke the access.

Figure 19: Examples of users' mental models - 4