



UNIL | Université de Lausanne

Unicentre

CH-1015 Lausanne

<http://serval.unil.ch>

---

Year : 2023

## WEARABLE ACTIVITY TRACKERS AND PRIVACY: ASSESSMENT OF THE RISKS, THREATS, AND COUNTERMEASURES

Zufferey Noé

Zufferey Noé, 2023, WEARABLE ACTIVITY TRACKERS AND PRIVACY: ASSESSMENT OF THE RISKS, THREATS, AND COUNTERMEASURES

Originally published at : Thesis, University of Lausanne

Posted at the University of Lausanne Open Archive <http://serval.unil.ch>

Document URN : urn:nbn:ch:serval-BIB\_374AE52D5DAF1

### **Droits d'auteur**

L'Université de Lausanne attire expressément l'attention des utilisateurs sur le fait que tous les documents publiés dans l'Archive SERVAL sont protégés par le droit d'auteur, conformément à la loi fédérale sur le droit d'auteur et les droits voisins (LDA). A ce titre, il est indispensable d'obtenir le consentement préalable de l'auteur et/ou de l'éditeur avant toute utilisation d'une oeuvre ou d'une partie d'une oeuvre ne relevant pas d'une utilisation à des fins personnelles au sens de la LDA (art. 19, al. 1 lettre a). A défaut, tout contrevenant s'expose aux sanctions prévues par cette loi. Nous déclinons toute responsabilité en la matière.

### **Copyright**

The University of Lausanne expressly draws the attention of users to the fact that all documents published in the SERVAL Archive are protected by copyright in accordance with federal law on copyright and similar rights (LDA). Accordingly it is indispensable to obtain prior consent from the author and/or publisher before any use of a work or part of a work for purposes other than personal use within the meaning of LDA (art. 19, para. 1 letter a). Failure to do so will expose offenders to the sanctions laid down by this law. We accept no liability in this respect.



UNIL | Université de Lausanne

---

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES  
DÉPARTEMENT DES SYSTÈMES D'INFORMATION

**WEARABLE ACTIVITY TRACKERS AND PRIVACY:  
ASSESSMENT OF THE RISKS, THREATS, AND  
COUNTERMEASURES**

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales  
de l'Université de Lausanne

pour l'obtention du grade de  
Doctorat en systèmes d'information

par

Noé ZUFFEREY

Directeur de thèse  
Prof. Kévin Huguenin

Co-directeur de thèse  
Prof. Mathias Humbert

Jury

Prof. Christian Zehnder, président  
Prof. Stéphanie Missonier, experte interne  
Prof. Michelle L. Mazurek, experte externe  
Prof. Thorsten Strufe, expert externe

LAUSANNE  
2023







UNIL | Université de Lausanne

---

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES  
DÉPARTEMENT DES SYSTÈMES D'INFORMATION

**WEARABLE ACTIVITY TRACKERS AND PRIVACY:  
ASSESSMENT OF THE RISKS, THREATS, AND  
COUNTERMEASURES**

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales  
de l'Université de Lausanne

pour l'obtention du grade de  
Doctorat en systèmes d'information

par

Noé ZUFFEREY

Directeur de thèse  
Prof. Kévin Huguenin

Co-directeur de thèse  
Prof. Mathias Humbert

Jury

Prof. Christian Zehnder, président  
Prof. Stéphanie Missonier, experte interne  
Prof. Michelle L. Mazurek, experte externe  
Prof. Thorsten Strufe, expert externe

LAUSANNE  
2023

# IMPRIMATUR

La Faculté des hautes études commerciales de l'Université de Lausanne autorise l'impression de la thèse de doctorat rédigée par

**Noé Zufferey**

intitulée

***Wearable Activity Trackers and Privacy: Assessment  
of the Risks, Threats, and Countermeasures***

sans se prononcer sur les opinions exprimées dans cette thèse.

Lausanne, le 06.10.2023



Professeure Marianne Schmid Mast, Doyenne



Members of the thesis committee:

Prof. Kévin HUGUENIN  
University of Lausanne  
Thesis supervisor

Prof. Mathias HUMBERT  
University of Lausanne  
Thesis co-supervisor

Prof. Stéphanie MISSONIER  
University of Lausanne  
Internal expert

Prof. Michelle L. MAZUREK  
University of Maryland, College Park  
External expert

Prof. Thorsten STRUFE  
Karlsruhe Institute of Technology  
External expert

Prof. Christian ZEHNDER  
University of Lausanne  
President



University of Lausanne  
Faculty of Business and Economics

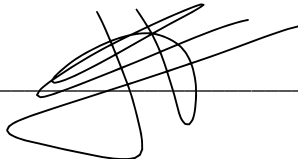
PhD in Information Systems

I hereby certify that I have examined the doctoral thesis of

**Noé Zufferey**

and have found it to meet the requirements for a doctoral thesis.  
All revisions that I or committee members made during the doctoral  
colloquium have been addressed to my entire satisfaction.

Signature: \_\_\_\_\_

A handwritten signature in black ink, consisting of several overlapping loops and lines, positioned over a horizontal line.

Date: 07.09.2023

Prof. Kevin Huguenin  
Thesis supervisor



University of Lausanne  
Faculty of Business and Economics

PhD in Information Systems

I hereby certify that I have examined the doctoral thesis of

**Noé Zufferey**

and have found it to meet the requirements for a doctoral thesis.  
All revisions that I or committee members made during the doctoral  
colloquium have been addressed to my entire satisfaction.

Signature: \_\_\_\_\_

A handwritten signature in black ink, appearing to read 'M. Humbert', written over a horizontal line.

Date: 08.09.2023

Prof. Mathias Humbert  
Thesis co-supervisor






University of Lausanne  
Faculty of Business and Economics

PhD in Information Systems

I hereby certify that I have examined the doctoral thesis of

**Noé Zufferey**

and have found it to meet the requirements for a doctoral thesis.  
All revisions that I or committee members made during the doctoral  
colloquium have been addressed to my entire satisfaction.

Signature:  \_\_\_\_\_

Date: 07.09.2023

Prof. Stéphanie Missonier  
Internal expert



University of Lausanne  
Faculty of Business and Economics

PhD in Information Systems

I hereby certify that I have examined the doctoral thesis of

**Noé Zufferey**

and have found it to meet the requirements for a doctoral thesis.  
All revisions that I or committee members made during the doctoral  
colloquium have been addressed to my entire satisfaction.

Signature: 

Date: 09/07/2023

Prof. Michelle Mazurek  
External expert



University of Lausanne  
Faculty of Business and Economics

PhD in Information Systems

I hereby certify that I have examined the doctoral thesis of

**Noé Zufferey**

and have found it to meet the requirements for a doctoral thesis.  
All revisions that I or committee members made during the doctoral  
colloquium have been addressed to my entire satisfaction.

Signature:  \_\_\_\_\_

Date: **7.9.2023** \_\_\_\_\_

Prof. Thorsten Strufe  
External expert



## Acknowledgement

First and foremost, I would like to thank my thesis supervisor, Kévin Huguenin who offered me the opportunity to conduct interesting research in the ISP Lab. Secondly, I would like to thank my co-supervisor Mathias Humbert, who joined the University of Lausanne during my Ph.D. The expertise of my two supervisors has greatly helped me in my research projects and the numerous relevant feedbacks have highly contributed to improving the quality of my thesis.

I also thank Stéphanie Missonier, Michelle Mazurek, and Thorsten Strufe, for having accepted to sit on my committee and providing valuable comments that helped me to improve the relevance of my manuscript, as well as Christian Zehnder who accepted to chair the public defense.

Besides my supervisors and thesis committee members, I am particularly grateful to Kavous Salesadeh Niksirat who also gave me valuable advice and support regarding multiple of my research projects. I also thank all my other co-authors with whom I published articles: Romain Tavenard, Lev Velykoivanenko, and Mauro Cherubini, and the other researchers with whom I have been able, even briefly, to collaborate on projects that will hopefully lead to publications in the coming month: Arianna Boldi, Amon Rapp, Marc-Olivier Boldi, Maurizio Caon as well as Arnaud Bonvin, whom I had the pleasure of supervising during his master thesis.

Many thanks to the Labex team, and especially to Anina Eggenberger, who helped me with the experiments involving participants.

I would also like to thank all the past and present members of the ISP Lab with whom I have shared an office: Yamane El Zein, Lev Velykoivanenko, Benjamin Trubert, Didier Dupertuis, Alexandre Meylan, Alain Mermoud, Bertil Chapuis, Dimitri Percia David, Rémi Coudert, and Gaël Bernard, as well as other people who helped me at some points with a research or teaching project, either to advise students, welcome participants, test a methodology, improve a data processing or simply proofreading: Alpha Diallo, Pierre Huber, James Tyler, Pooja Rao, Rita Abi Akl, Marie Reignier, Patrick Rousseau, Robin Zufferey, Lahari Goswami, Holly Cogliati, and Vincent Vandersluis.

Special thanks also to Désirée Krejci and Caroline Violot, who have successively been at my side as Ph.D. representatives at the department council, as well to all the members of this council and members of the Doctoral Program in Computer Science of CUSO. I also want to thank all the candidates of “Ma Thèse en 180 secondes 2023” and the event organization staff (it was a really



fun experience).

Thanks also to Sarah Duplan, executive secretary in charge of the doctoral school. Also, I thank all the members of the Information System department, as well as the University of Lausanne administrative, technical, security, cafeteria, and cleaning staff, without whom none of this could function properly.

Besides the University of Lausanne, I also want to thank Laetitia Salamin, the mother of my child, Lycia (whom I also thank, even though at three years old she was not able to help me in full knowledge of the facts) for her support, as well as my whole family. I also want to thank all my friends and in particular, all those who were Ph.D. students at the same time as me and to whom I attended the defense, or who will be defending soon: Jacky Casas, Guillaume Buro, Mauro Salomon, and Laurianne Pillet. I would also like to thank all the other people I have worked with in the past and with whom I have learned about computer science, and in particular RTFM Corp.

Finally, thanks to the Swiss National Science Foundation, armasuisse S+T, and the HEC Research Fund who partially funded many of my research projects.

Wearable Activity Trackers and Privacy:  
Assessment of the Risks, Threats, and  
Countermeasures

Noé Zufferey  
University of Lausanne

October 9, 2023



## Abstract

Wearable devices, such as wearable activity trackers (WATs), are increasing in popularity. Although they can help improve a person's quality of life, they also raise serious privacy issues. Although security aspects of WATs have been widely studied (e.g., Bluetooth security, inference of password or biometrics), as well as privacy-related aspects such as users' attitudes and concerns, we lack knowledge about the privacy of WAT users. Indeed, the security aspects that were studied in prior work are not enough to build a realistic adversary model, as these studies focus mostly on communication protocols and not on large-scale data collection. Furthermore, previous work related to data inference by using WATs focuses on only functionalities rather than on privacy (e.g., better monitoring of activity or health to improve user experience). Moreover, these studies focus only on the inference of behavioral patterns (e.g., activities, consumption) or conditions (e.g., diseases), but none of them investigate the inference of users' personal attributes (e.g., personality, religion, political views).

In this thesis, composed of three research papers and a literature review, we contribute to the WAT security & privacy research field by analyzing how the data of WAT users can be accessed at a large scale by many potential adversaries, by evaluating how such data can be used to infer users' personal attributes and, finally, by proposing privacy enhancing technologies (PETs) to protect their privacy. Concretely, after analyzing the current literature about WAT security & privacy, we conduct a user-survey study to better understand the WAT user's behaviors towards data sharing, especially with respect to third-party applications (TPAs) that can easily be used by adversaries to collect data. We then use a rigorous machine-learning approach to evaluate to what extent users' psychological profiles (Big 5) can be inferred from WAT data, and we discuss the related consequences on the users' privacy and society as a whole. Finally, to propose effective and likely-to-be-adopted protection mechanisms, we conduct a user-centered design study by using a participatory design methodology before analyzing and evaluating the proposed designs in order.



# Contents

|          |   |           |
|----------|---|-----------|
| <b>1</b> | <b>Introduction</b>   | <b>1</b>  |
| 1.1      | Information Security & Privacy . . . . .  | 2         |
| 1.2      | Research Ethics . . . . .   | 3         |
| 1.3      | Wearable Activity Trackers (WATs) . . . . .   | 4         |
| 1.4      | The WAT Ecosystem . . . . .   | 5         |
| 1.5      | Adversarial Model . . . . .   | 7         |
| 1.6      | Research Scope and Methodology . . . . .  | 9         |
| 1.7      | Contributions . . . . .   | 10        |
| 1.7.1    | List of Publications . . . . .  | 11        |
| <b>2</b> | <b>Literature Review</b>  | <b>13</b> |
| 2.1      | Privacy . . . . .   | 14        |
| 2.1.1    | Are WATs Risky for Users' Privacy? . . . . .  | 14        |
| 2.1.2    | Research on WAT Users . . . . .   | 16        |
| 2.1.3    | Privacy Policies, Regulations, and Ethics . . . . .   | 22        |
| 2.1.4    | Health at Work or Workplace Surveillance? . . . . .   | 27        |
| 2.1.5    | Privacy-Enhancing Technologies (PETs) . . . . .   | 28        |
| 2.2      | WAT Security . . . . .  | 30        |
| 2.2.1    | WAT–Phone Communication . . . . .   | 30        |
| 2.2.2    | Phone–Server Communication and Data Storage . . . . .   | 33        |
| 2.2.3    | Side-Channel Attacks . . . . .  | 35        |
| 2.2.4    | Authentication . . . . .  | 36        |
| 2.2.5    | Threat Assessment and Mitigation, and Security Protocols . . . . .  | 37        |
| 2.3      | Overview and Research Gaps . . . . .  | 38        |
| <b>3</b> | <b>“Revoked just now!” Users’ Behaviors toward Fitness-Data<br/>Sharing with Third-Party Applications</b> | <b>41</b> |
| 3.1      | Introduction . . . . .  | 42        |

---

|          |   |           |
|----------|---|-----------|
| 3.2      | Related Work . . . . .  | 44        |
| 3.3      | Methodology . . . . .   | 45        |
| 3.3.1    | Recruitment . . . . .   | 46        |
| 3.3.2    | Design of the Survey Questionnaire . . . . .  | 46        |
| 3.3.3    | Procedure . . . . .   | 52        |
| 3.3.4    | Data Reliability . . . . .  | 53        |
| 3.3.5    | Coding Process . . . . .  | 53        |
| 3.3.6    | General Statistics . . . . .  | 55        |
| 3.4      | Results . . . . .   | 56        |
| 3.4.1    | Users tend to forget about their TPAs. . . . .  | 56        |
| 3.4.2    | Users generally overestimate the amount of data they<br>share on their public profiles. . . . .       | 59        |
| 3.4.3    | Friends and family are favorite data recipients. . . . .  | 60        |
| 3.4.4    | Users are inclined to use PETs. . . . .   | 62        |
| 3.4.5    | Users lack knowledge about data sharing. . . . .  | 64        |
| 3.5      | Discussion . . . . .  | 70        |
| 3.6      | Limitations . . . . .   | 72        |
| 3.7      | Conclusion . . . . .  | 73        |
| <b>4</b> | <b>Watch your Watch: Inferring Personality Traits from Wear-<br/>able Activity Trackers</b> . . . . . | <b>75</b> |
| 4.1      | Introduction . . . . .  | 75        |
| 4.2      | Background . . . . .  | 78        |
| 4.3      | Adversarial Model . . . . .   | 79        |
| 4.4      | Data Collection and Statistics . . . . .  | 80        |
| 4.4.1    | Data-Collection Campaign . . . . .  | 80        |
| 4.4.2    | Descriptive Statistics . . . . .  | 83        |
| 4.4.3    | Participants' Privacy Concerns . . . . .  | 86        |
| 4.5      | Inference . . . . .   | 87        |
| 4.5.1    | Methodology . . . . .   | 87        |
| 4.5.2    | Feature Extraction . . . . .  | 89        |
| 4.6      | Results . . . . .   | 92        |
| 4.7      | Related Work . . . . .  | 102       |
| 4.8      | Discussion . . . . .  | 103       |
| 4.9      | Limitations and Generalization of the Results . . . . .   | 107       |
| 4.10     | Conclusion and Future Work . . . . .  | 108       |

|          |  |            |
|----------|--|------------|
| <b>5</b> | <b>Our Data, Our Solutions: A Participatory Approach for Enhancing Privacy in Wearable Activity Tracker Third-Party Apps</b> | <b>111</b> |
| 5.1      | Introduction . . . . .   | 112        |
| 5.2      | Methodology . . . . .  | 114        |
| 5.2.1    | Recruitment . . . . .  | 115        |
| 5.2.2    | Session Procedure . . . . .  | 116        |
| 5.2.3    | Room Layout . . . . .  | 122        |
| 5.2.4    | Participants & Groups Composition . . . . .  | 123        |
| 5.2.5    | Coding Process . . . . .   | 124        |
| 5.2.6    | Expert Review Meeting . . . . .  | 125        |
| 5.3      | Results . . . . .  | 126        |
| 5.3.1    | Feature 1 - Partial Sharing . . . . .  | 126        |
| 5.3.2    | Feature 2 - Visualization . . . . .  | 129        |
| 5.3.3    | Feature 3 - Centralization . . . . .   | 130        |
| 5.3.4    | Feature 4 - Reminders . . . . .  | 131        |
| 5.3.5    | Feature 5 - Revocation Assistance . . . . .  | 132        |
| 5.3.6    | Feature 6 - Education & Sensitization . . . . .  | 133        |
| 5.3.7    | Feature 7 - TPAs Limit . . . . .   | 134        |
| 5.4      | Discussion . . . . .   | 135        |
| 5.5      | Limitations . . . . .  | 139        |
| 5.6      | Conclusion . . . . .   | 140        |
| <b>6</b> | <b>Conclusion</b>  | <b>143</b> |
| 6.1      | Contributions . . . . .  | 143        |
| 6.2      | Future Work and Perspectives . . . . .   | 144        |
| <b>A</b> | <b>Appendix</b>  | <b>181</b> |
| A.1      | Questionnaire of Chapter 3 . . . . .   | 181        |
| A.2      | All Mental Models . . . . .  | 232        |
| A.3      | Technical Codebook . . . . .   | 268        |
| A.4      | Contextual Codebook . . . . .  | 268        |
| A.5      | Why Revoking Access Codebook . . . . .   | 269        |
| A.6      | Why Not Revoking Access Codebook . . . . .   | 270        |
| A.7      | Suggestions Codebook . . . . .   | 271        |
| A.8      | WAT Data Sharing . . . . .   | 272        |
| A.9      | Mental Models . . . . .  | 273        |
| A.10     | Design Feature Coding . . . . .  | 277        |





# Chapter 1

## Introduction

The number of users of wearable devices and, in particular, (wrist-worn) wearable activity trackers (WATs) increases daily. It reached 218 million in 2022 and is projected to reach over 320 million in the next five years [1], and there are more than one billion wearable devices worldwide [2]. These devices collect large amounts of physiological and contextual data, such as step counts, heart rate (for those equipped with the appropriate sensors), activities, and sleep. Such data can help WAT users better monitor their physical activities and health, following a *quantified-self* [3] approach. However, wearable devices raise new privacy and security issues. For instance, Eberz et al. [4] show that data collected from wearable devices can be used to bypass biometric authentication systems by using accelerometer data to impersonate users. Furthermore, accelerometer data can be used to infer keystrokes (e.g., on pinpads) [5, 6, 7]. Moreover, WAT data can be used to infer daily activities and habits [8, 9, 10, 11] (e.g., eating) and drug usage [12] (e.g., cocaine), and even to identify SARS-CoV-2 infections [13]; such inferences are highly sensitive from a privacy perspective. Finally, WAT data, such as running routes, can be used to infer sensitive locations (e.g., user’s home), even when they use protection mechanisms [14, 15, 16]. Aggregated location data have even been used to locate military bases and to infer their internal structures [17], specifically in remote areas where unusual activity patterns were observed.

In the context of the *quantified-self*, questioning the effect of such data collection (and sharing) on people’s privacy is becoming increasingly relevant, especially as many users express concerns about the misuse of their data [18, 19]. Personal information, such as personality, socioeconomic status, sexual orientation, and religion can probably be inferred from data collected by wearable

devices, similarly to the possibilities to do the same with location and social network data (e.g., [20, 21, 22]). Moreover, third-party entities such as advertisers, marketers, health insurers, employers, and governments might have an interest in learning sensitive information derived from the data collected by WATs [23]. For example, an employer could offer free WATs to their employees if they agree to share the collected data with their employer, hence the employer could monitor their employee’s health (in US, individuals’ insurance health plans are generally covered by their employer) and activities. Indeed, some organizations, encouraged in particular by Fitbit (one of the market leaders for WATs [24, 25]), now offer their employees tracking devices through health programs [26]. An insurance company (e.g., health insurance ) could also directly provide tracking devices to their policyholders to better analyze risks. For instance, Google acquired Fitbit [27] and Alphabet, Google’s parent company, and their influence is growing rapidly in the health insurance market [28]. Furthermore, they plan to force Fitbit users to migrate to their Google accounts [29]. A government, for national security reasons, could also gain access to the data of a WAT service provider. For example, in 2019, former US President Trump suggested using data from wearable devices for national security purposes, essentially to preemptively detect mass shooters [30].

## 1.1 Information Security & Privacy

There are multiple different concepts using the notion of privacy. The Office of the United Nations High Commissioner for Human Rights (OHCHR) defines privacy as “*the ability of an individual or group to seclude themselves or information about themselves, and thereby express themselves selectively*” [31]. Whereas, Westin defines in 1967 privacy as “*the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others*”. Such definitions correspond to a *privacy by control* approach, hence we can claim that an entity’s privacy is about who is authorized, by this entity, to access which personal data. In addition to this usual approach, in this work, we invoke multiple, not exclusive, concepts or methods related to privacy. *Privacy by design* is about taking users’ privacy into consideration upstream, during the development of an information and communications technology (ICT), by taking into account the current state of the art and principles of data-protection [33]. *Privacy by default* is about setting all the parameters of an ICT so that they guarantee by default the maximum possible level of privacy for the users [33]. Finally, the

concept of *privacy impact assessment* (PIA) was introduced in the European General Data Protection Regulation (GDPR) and refers to “*the obligation of the controller to conduct an impact assessment and to document it before starting the intended data processing*” [34]. Basically, it consists of a concept stating that any entity that has access to personal data and intends to process this data should evaluate how their processes can affect the privacy of the involved individuals and use the appropriate mitigation techniques accordingly.

As for security, and in particular information security, the US National Institute of Standards and Technology (NIST), as well as the European Union Agency for Cybersecurity (ENISA), defines it as “*the protection of information and information systems from unauthorized access, use, disclosure, disruption, modification, or destruction in order to provide confidentiality, integrity, and availability*”. This concept is therefore highly related to privacy, and we will often refer to both by the expression *(information) security & privacy*.

## 1.2 Research Ethics

A large part of the research in security and privacy-related topics consists of embodying a given adversary and conducting specific attacks in order to uncover vulnerabilities. Typically, in this specific thesis, we use an inferential privacy methodology to show that WAT data can be used to infer – with some accuracy – the personality traits of WAT users, potentially without them being aware of it and/or against their consent. The question then becomes: What are the ethical consequences of such an approach, particularly as discovering and disclosing such vulnerabilities could help adversaries to conduct related types of attacks? To address this question, experts in computer security have developed the concept of coordinated vulnerability disclosure (or responsible disclosure) model [35], which consists of the uncovered vulnerability being publicly disclosed only when the parties involved have had enough time to proceed to a remediation [36]. However, in a more privacy-oriented research field, we adopt a full-disclosure model consisting of publicly disclosing an uncovered vulnerability as early as possible. In such a way, we ensure that the knowledge of this vulnerability is rapidly disseminated, so that as many potential targets as possible are aware of the risks. This model has multiple advantages as, for example, it can provide users with leverage to demand that the vulnerability is patched when the parties involved (e.g., the service provider) have no other incentive to do so (privacy issues are indeed often only harmful to the users). Furthermore, vulnerabilities are often not only known by the parties who un-

covered them and plan to disclose them, they may have been uncovered and exploited for years by third parties who never intended to disclose them, either publicly or to the service provider. Furthermore, many privacy risks are not necessarily direct vulnerabilities but are just the consequences of using a given service. Taking into consideration the concept of privacy calculus and the fact that individuals consider the trade-off between privacy risks and receiving relevant services [37], it is important that they are fully aware of the risks before adopting a given technology.

### 1.3 Wearable Activity Trackers (WATs)

In the current literature, there is no consistent definition of a wearable activity tracker (WAT). Furthermore, as there are multiple types of devices that could be considered a WAT, we first need to define precisely the type of devices studied in this work. Despite the multiple definitions existing in the literature, there are many commonalities between them. Therefore, in order to create a standardized definition, we identified the essential and optional properties of WATs, as described in the literature. To be considered a WAT, a given device must have all of the essential properties and can also have the optional ones [38]. We identified essential and accidental properties of WATs, based on the studies of Becker et al. [39], Hoy [40], and Pingo and Narayan [41]. These properties were initially optional in the context of a survey on WAT utility, privacy, and security [42]. The essential properties are as follows:

- be worn on the body
- has sensors that record physiological/environmental data
- be an electronic/digital device
- provides data analysis that is available to users, without the need for a health professional

For example, a smartphone, although corresponding to most of these criteria, is not considered as a WAT as it is not designed to be worn on the body. The optional properties are as follows:

- uploads data to a server or connected device (e.g., using Bluetooth)
- uses a docking station to sync with a PC, or WiFi to upload directly
- enables users to visualize data in graphical format on a companion app or website
- enables users to visualize some of the data on the WAT itself
- provides immediate feedback, and

- provides general/numerical feedback (after an activity).

The most common sensors used in WATs are accelerometers, gyroscopes, photoplethysmograms (used for measuring heart rate and respiration), pulse oximeters (blood oxygenation), altimeters, and GPSs. More advanced and recent models tend to include a compass, a thermometer, a microphone, a magnetometer, an ambient light sensor, and an electrodermal activity sensor. Therefore, we consider smartwatches as WATs because, even if they offer more functionalities than some fitness trackers, they still fit our definition. All the devices that we studied in this work correspond to the aforementioned definition. Medical-connected devices (e.g., insulin pumps) and wearable devices with very specific purposes (e.g., connected shoes or e-glasses) are not considered WATs. Moreover, in this specific work, we focus only on wrist-worn devices, as they are the most common types.

## 1.4 The WAT Ecosystem

Figure 1.1 depicts the typical WAT ecosystem. Generally, a WAT ecosystem is composed of a WAT paired with a connected device (e.g., smartphone, tablet). The WAT can store only data that was collected in the past few days. The personal data is regularly transmitted from the WAT to the connected device via a Bluetooth communication protocol, Bluetooth low energy (BLE). A companion app provided by the WAT's vendor (i.e., the service provider) is installed on the connected device to monitor the pairing and to visualize the collected data. The connected device can generally store only recent data, as older data needs to be stored on the cloud and downloaded if needed. Hence, to store the collected data on the cloud, the companion app regularly synchronizes, through the Internet, with the service provider's servers. The servers that store the users' WAT data can process raw WAT data and perform various analytics [45], for example, in some cases, data stored on the servers are processed to compute further information that is automatically sent back to the companion app. In some cases (Apple), the data stored on the service provider's servers is encrypted and can be accessed only by the user. The connected device can also send data to the WAT, such as firmware updates or notifications [46]. In a few cases (depending on the model), the WAT can use direct communication with the service provider's servers for firmware update purposes. However, in most cases, the WAT does not directly transmit data (i.e., fitness data) to the service provider servers.

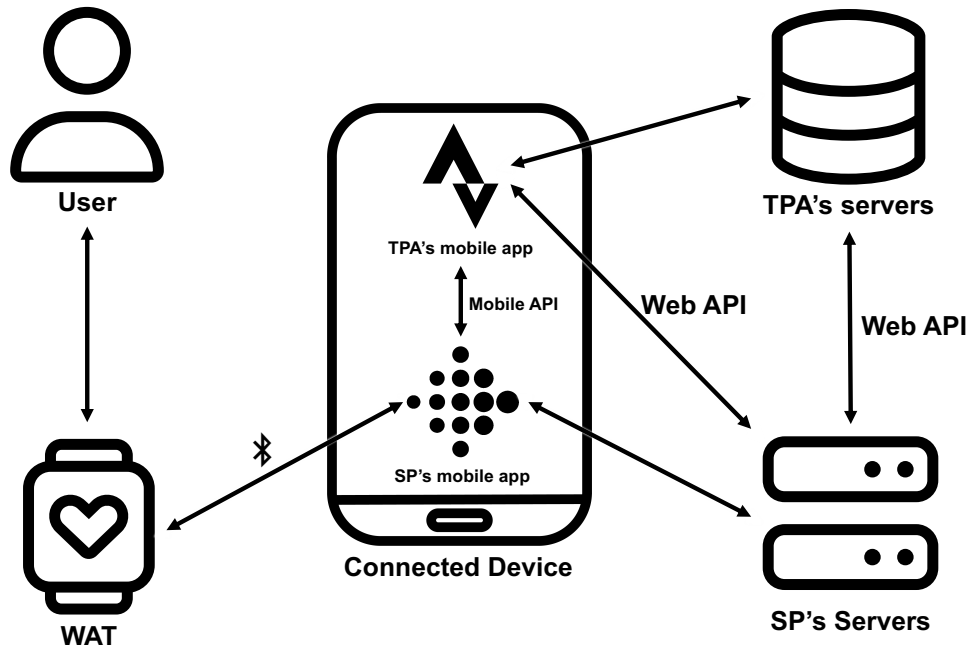


Figure 1.1: WAT ecosystem. The user wears the WAT that transmits their data to the service provider's servers (SP's servers) via their phone (connected device). The user can grant authorization (e.g., using a protocol as OAuth2) for a given third-party application (TPA) to access their data, generally by using a web API. The user can select the different types of data that they agree to share with the TPA, then the TPA receives a pair of tokens that they can use to request the user's data through the API, either from their own servers (TPA's servers) or from an app (e.g., Strava) installed on the user's phone. In some cases (e.g., Apple), the data is directly transmitted, through a local mobile API, from the companion app to a TPA's mobile app installed on the user's phone. All these APIs are permitted [43, 44]

Moreover, the companion app generally offers more functionalities such as social network features and data sharing with third-party applications (TPAs). A user can grant authorization to a TPA to access their data by using a specific authorization protocol (e.g., OAuth2). By doing so, the TPA will receive a token (or a pair of tokens). With this token, the TPA can request data from the service provider's server by using a dedicated API. This request can come from either the TPA's server or a TPA app (e.g., Strava) installed on the user's smartphone (connected device). In some cases (e.g., Apple) the

TPA can directly access WAT data stored on the connected device by using the app installed on it. Yet, the data collected by the mobile app can be subsequently sent to a server. The user will then generally be able to access their data by using the TPAs functionalities (e.g., the TPA's app installed on their smartphone). Such a data-sharing method can also work in the opposite way, meaning the companion app and/or the service provider can access data that were collected by a TPA. Most of the market leader WATs available, such as Apple Watch, Fitbit, and Garmin, as well as TPAs such as Strava or MyFitnessPAL, match this model. At any moment, the user can revoke previously granted access. Such revocation would remove the access privilege; after which, the token can no longer be used by the TPA to access the user's data. However, during the period of time the token is valid (after having granted access authorization and before revoking it), 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.

Beyond data sharing with TPAs, a WAT user can also share their data on a dedicated social network (e.g., Fitbit Community). Such data sharing can be done according to different audiences (public, friends, groups). In other cases, a user can directly share their data with another user with the same companion app (e.g., Apple Health).

## 1.5 Adversarial Model

In this thesis, we focus on an adversary that can access some or all of a user's data processed by the service provider. Therefore, we focus on devices where data are collected by the service provider and are not stored only locally. As we will see in Chapter 2, many studies were conducted on the security of WATs and their communication protocols, as BLE for WAT-smartphone communication or any protocol used to transmit users data to the cloud. However, an adversary that seeks access to users' data in such a way (e.g., intercepting BLE or HTTP communication) will have a very limited scope of action as they would need, for example, to capture signals of a specific user (e.g., with a specific antenna), or to eavesdrop on HTTP(S) traffic. Another type of adversary, complementary to the previous one in the sense that they access WAT users by different means, is the one that has direct access to already processed data (e.g., step count, heart rates). This type of adversary constitutes, in our opinion, a particularly realistic model as it does not necessarily require physical proximity to the device, and can therefore particularly easily collect data



from a larger number of users. Furthermore, in the context of WATs, considering TPAs as an adversary actually brings to light a new type of adversary model that users are not fully aware of. Indeed, whereas, it is perfectly clear when they start using a WAT that they will, at some point, share personal data with the service provider, users may overlook the fact that, depending on their usage of the WAT, their data may be accessible to other companies (i.e., companies providing them with TPAs), who have their own terms of services. Moreover, WAT users tend to underestimate the impact their processed data can have on their privacy, especially related to inference threats [19, 46]. This is the reason why, in this work, we focus on such adversaries. There exist multiple adversaries who correspond to this description. One such adversary is typically the service provider itself, such as Fitbit. But most of all, it can be any TPA (or their business partners) to whom many users have granted, knowingly or not, access to their data (e.g., have given a token pair through OAuth 2.0 [43]). Users may want to share their data with TPAs for multiple reasons, generally, they do it for additional functionalities not offered by the original services or applications, but it can also be to share their data with companies that base their business on WAT-data collection such as WeWard [47], that offers their users to be paid according to the number of steps they take, or Actifit [48] that offers their users to have free access to fitness-related services as nutrition consultation according to their activity count, or other similar companies based on cryptocurrency such as Fitcoin [49]. It could also be any of their business partners [50]. Such an adversary (i.e., a company providing TPAs) would have the possibility to obtain years of data collected from millions of users. For example, there were 31 million Fitbit users in 2020 [51], and 20 million for WeWard in 2023 [47] (this includes users that use only the step-count feature of their smartphone).

TPAs are known to ask users to access far more data than they actually need to provide their services [52]. 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 [53, 54]. 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 simply 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 [55, 56]. Finally, to cease the data sharing, a user must actively revoke the access permissions by

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

Another way to access the data of WAT users is to use users’ (public) profiles. Users’ PII (e.g., birth date, e-mail address), as well as, to some extent, fitness data (e.g., average step count, list of achievements) could be publicly available on the service provider web platform or could be accessed by 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. Even if leaving a given type of information publicly available may be considered as an authorization, and is considered as an exception to the prohibition of processing personal data under GDPR<sup>1</sup>, it is important to highlight that privacy settings are often set to the highest visibility by default, and many users, due to a lack of knowledge and awareness of data sharing and privacy, do not modify them [57]. Furthermore, sometimes, it is even not possible for a user to set the lowest visibility for a given type of data. For example, although Fitbit offers three different modes of visibility for their users’ public profiles (i.e., “Private”, “Friends”, “Public”), the lowest visibility option (i.e., “Private”) is not available for daily step counts and the user can only choose between the two other options (i.e., “Friends” or “Public”) for that specific type of data. Moreover, as the API data access used by TPAs—and that uses 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 [58], to gain access to user data.

## 1.6 Research Scope and Methodology

As explained above, WATs can raise multiple security & privacy issues for the users. Although a large amount of research studies focus on the security aspects of WATs (e.g., Bluetooth security, inference of password and/or biometrics), the overwhelming majority of studies about personal information inference from WAT data are not privacy-oriented. Indeed, the studies related to, for example, inference of activities, consumption, or diseases using WATs are all about creating new functionalities to help the users or a related

---

<sup>1</sup>“Processing of personal data [...] shall be prohibited [except if it] relates to personal data which are manifestly made public by the data subject.” <https://gdpr.eu/article-9-processing-special-categories-of-personal-data-prohibited/>

third party to better monitor their activity or health. Most studies about inferences made by using WAT data focus on inferring behavioral data (e.g., activities, consumption) or conditions (e.g., diseases), but none of them are about inferring users' personal attributes (e.g., personality, religion, political views).

Another research gap, more related to the HCI approach, is that, whereas WAT users' *attitudes* toward fitness-data sharing has been widely studied (e.g., [19, 46, 59, 60]), no study has focused on the *actual behaviors* of data sharing by users of WATs.

Moreover, as most of the provided solutions related to the privacy of WAT users are device- or data-oriented, none of them are based on a user-centric approach, in particular, we did not find any studies with participatory design or co-design approaches.

In this work, we intend to contribute to the WAT security & privacy research field relative to these three identified research gaps. In the next chapters, we will indeed (1) conduct a systematic literature review about WAT security & privacy and highlight these three aforementioned gaps, (2) conduct a user-survey study to better understand the *actual* WAT user's behaviors towards data sharing (assessment of the risks), (3) use a machine-learning approach to evaluate how users' personal attributes (i.e., personality) can be inferred from WAT data and discuss the related consequences on the users' privacy and society as a whole (assessment of the threats), and (4) we will conduct a user-centric design study by using a participatory design methodology in order to propose effective and likely-to-be-adopted protection mechanisms (development of draft countermeasures).

## 1.7 Contributions

For this thesis, three main research projects were completed. The first focused on the data-sharing behavior of WAT users (Chapter 3) and was published at PETS 2023 [61]; we deployed a large online survey (n=628) for polling users about (1) the third-party applications, (2) the contacts to which the users granted access to their WAT data, and (3) their understanding (incl. mental models) regarding the way third-party applications can access their fitness data. The second project focused on the extent to which personality (viz. Big-5 traits, aka OCEAN [62]) can be inferred from the data collected by fitness trackers (Chapter 4) and was published at USENIX Security 2023 [63]; for this, we organized an in-situ experiment ( $n = 204$ ) where participants were

provided with Fitbit Inspire HR bracelets that they actively used for four months. Finally, the third project was about designing new privacy-enhancing technologies for WATs with a user-centric approach; for this, we organized participatory design sessions ( $n = 26$ ) with WAT users and supervised discussions about WAT-data sharing and the related privacy risks. Then, the participants proposed and evaluated multiple designs that, in order to increase the privacy of WAT users, aim to help them better manage their data sharing. This work has been accepted as a poster at SOUPS.

Three additional research projects related to WATs and based on the in-situ experiment described in Chapter 4 were conducted by the author of this thesis, but he is not the first author. The first additional project is about users' perceptions of the privacy risks related to fitness tracking and was published at IMWUT (UbiComp 2022) [46]. The second one is a project on the effect of fitness trackers on self-perception and body image (in submission at the time of writing), and the third is about WAT-data series (i.e., heart-rate series) identification using step counts (in progress).

A part of this thesis also includes a literature survey on the costs (w.r.t. privacy and security) and the benefits (w.r.t. utility) of wearable activity trackers (submitted to ACM Computing Surveys). The privacy part and the security part of this survey constitute the literature review section of this thesis (Chapter 2).

This work, therefore, contributes to the security & privacy research field related to WATs in the following way. We analyze the privacy & security literature related to WATs in Chapter 2 and identify multiple research gaps. In Chapter 3 we motivate the work by identifying specific types of privacy risks related to the way WAT users share their data. Then, in Chapter 4 we explore a specific threat model and show how WAT data can be used to infer personal sensitive information and discuss the related consequences on user privacy and society at large. Finally, in Chapter 5 we propose different privacy-enhancing technologies that can be used to minimize the risks, and, therefore, help the users protect their privacy, before discussing future work and concluding in Chapter 5.

### 1.7.1 List of Publications

1. L. Velykoivanenko, K. S. Niksirat, **N. Zufferey**, M. Humbert, K. Huguenin, and M. Cherubini. 2022. Are Those Steps Worth Your Privacy?: Fitness-Tracker Users' Perceptions of Privacy and Utility, *Proc. of the ACM on*

*Interactive, Mobile, Wearable and Ubiquitous Technologies*

2. **N. Zufferey**, K. Salehzadeh Niksirat, M. Humbert, and K. Huguenin. 2023. ‘Revoked just now!’ Users’ Behaviors Toward Fitness-Data Sharing with Third-Party Applications, *PoPETs*
3. **N. Zufferey**, M. Humbert, and K. Huguenin. 2023. Watch your Watch: Inferring Personality Traits from Wearable Activity Trackers, *Proc. of the USENIX Security Symposium*
4. **N. Zufferey**, K. Salehzadeh Niksirat, M. Humbert, K. Huguenin. 2023. Personal Data, Personal Solutions: A Participatory Approach to Design Privacy-Enhancing Technologies for Fitness-Data Sharing, *SOUPS poster, and in preparation for submission at the time of writing*
5. K. S. Niksirat, L. Velykoivanenko, **N. Zufferey**, K. Huguenin, M. Humbert. Wearable Activity Trackers: A Survey on Utility, Privacy, and Security, *under review at the time of writing*
6. A. Boldi, A. Silacci, M. Boldi, Ma. Cherubini, M. Caon, **N. Zufferey**, K. Huguenin, A. Rapp. On the Effect of First-Time Use of Fitness Tracker on Users’ Body Representation: A Mixed-Method Controlled Experimental Study, *under review (major revision) at the time of writing*
7. A. Bonvin, **N. Zufferey**, K. Huguenin, M. Humbert. Reidentification of wearable activity tracker data series in anonymized dataset. *in preparation for submission at the time of writing*

# Chapter 2

## Literature Review

There has been extensive research on WATs. Therefore, conducting a systematic literature review is important to identify the different approaches, and methodologies used in this research field. Moreover, it helped us to identify research gaps and to evaluate the potential of our research projects. In the framework on this thesis, we conducted a literature review about the utility, privacy, and security aspects of WATs. In this chapter, we include parts related to privacy & security. The part about WAT utility is not included in this thesis.

To conduct this literature review about WAT privacy & security, we followed the methodology of Kitchenham et al. [64]. After defining our keywords<sup>1</sup>, we searched in ACM DL, IEEE Xplore, AIS library, USENIX, PoPETs, Science Direct, and Springer Link. We also used Google Scholar to include papers from other databases and publishers (e.g., Taylor & Francis). During the review process, we also kept track of the most recent and relevant published proceedings (e.g., CHI 2023) to update our paper database. We excluded the papers that (1) are not written in English, (2) were published in 2012 or earlier, and (3) are not peer-reviewed (e.g., position papers, letters, editorials, prefaces, article summaries, theses, patents, or books). We included only the papers (1) that are about WATs and/or have implications for WATs (according to the definition in the introduction of this thesis) and (2) that have direct

---

<sup>1</sup>We used the following search strings: “physical activity data” OR “physical activity tracker” OR “fitness data” OR “fitness tracker” OR “wearable activity tracker” OR “fitness tracking” OR “wearable activity tracking” AND “utility” OR “privacy” OR “security” OR “perception” OR “understanding” OR “experience” OR “expectation” OR “sharing” AND “system” OR “device” OR “application” OR “app” OR “service” OR “bracelet” OR “wrist-worn”.

relevance to the privacy, and/or security of WATs.<sup>2</sup> For that last criterion, we decided to include borderline cases (e.g., articles about inference of personal information but not explicitly privacy/security oriented). To synthesize the findings, we followed the JBI Manual for Evidence Synthesis [65], where we reviewed the paper summaries, identified the common patterns and homogeneity between the papers and, in terms of their findings, highlighted the heterogeneity and diversity.

## 2.1 Privacy

In this section, we first review the proven privacy risks of using WATs. Then, we delve into users' perceptions of privacy and behaviors. We later discuss privacy policies, regulations, the forensic use of WATs, and ethics. Next, we investigate the privacy consequences of using WATs in workplaces. We conclude the section with privacy-enhancing technologies (PETs).

### 2.1.1 Are WATs Risky for Users' Privacy?

Machine-learning (ML) models using WAT data can be used to monitor ECG waveforms [66], post-surgery complications [67], multiple-sclerosis symptoms [68], SARS-COV-2 infection [69], mental health states such as stress resilience [70] or depression [71], and to predict the readmission of cancer patients [72]. Combining WAT data with other resources can help build a health persona [73]. In an edge case, WAT data could help specialists better understand social engagements between autistic children who have difficulties in making non-verbal communication [74].

However, WAT data can also be used in **adversarial settings** with potentially negative consequences for users. Especially as WAT service providers, such as Fitbit, are known to share users' data with their business partners [50] and as developers do not always know how to protect users' privacy [75]. Here, we review different types of information that can be inferred and can violate user privacy. Studies about human activity recognition (HAR) show how various types of activities can be inferred using data collected by WAT sensors. Several novel algorithms and frameworks were developed for HAR (e.g., [76, 77]). Dietrich and van Laerhoven [78] propose a typology for classifying the different contexts of WAT usage. Activities can be successfully

---

<sup>2</sup>The list of the reviewed papers is available on [OSF](#)

recognized by using data collected with off-the-shelf WATs and even for short-duration data (i.e., short and quick movements) [79]. Such inferences are generally more efficient than using data from other common devices (i.e., smartphones) [80]. This is because WATs, unlike phones that are usually in users' purses or pockets, are worn close to the body (i.e., worn on the user's wrist). Although these HAR studies propose new functionalities, their findings can be utilized by attackers.

The most frequently inferred activity types are **eating** and **drinking**. Thomaz et al. [81] and Weiss et al. [82] explore eating/drinking detection by using WAT data, whereas Biel et al. [83] also used contextual data (e.g., time, geolocation) to infer the type of meal. Also, in relation to consumption, a model for detecting users' **drunkenness** in real-time was developed Gutierrez et al. [84]. Shoaib et al. [85] use WAT data to detect **smoking** events. WAT data can also be used for other purposes such as tracing the **geometric motion of a user's arm** [86], for recognizing objects moved by users and the identity of the users who moved them [87], and for preventing **pedestrian distractions** [88]. WAT data can also be used for more than predicting activities: to predict users' **moods** and to **recommend music** [89]. The movement of a user's WAT, when using NFC payment terminals, can help to infer their **height** [90]. Finally, information shared by WAT users on social media can be used to infer personal information, such as **weight** [91].

A few studies focus on **location inference** have been conducted. Hassan et al. [14] and Dhondt et al. [16] study bypassing endpoint privacy zones (EPZs) to infer users' locations; they could infer more than four-fifths of the locations. EPZ is a mitigation technique that consists in defining a private zone within which some data are not revealed (e.g., to protect users' exact location). Meteriz et al. [92] also showed that location inference is possible, with certain previous knowledge and by using the elevation profile. **Handwriting** recognition is a particular case of HAR, where inference is made not only to detect the event but also to infer the written letters and words. For example, WAT data can be used to recognize air-writing gestures (and words) [93], finger-writing gestures [94], and gestures of writing on a whiteboard [95]. Xia et al. [96] showed that they can infer one-third of hand-written words. Wijewickrama et al. [97] replicate the four previously described studies and obtain lower accuracy than in the original studies, thus reminding us of the necessity of the replication studies in HAR research. Therefore, considering that HAR literature reports handwriting-event detection accuracy is between 65% and 90% (depending on the context), and knowing that most users do not wear the



device on their dominant hand, they concluded that handwriting recognition is unlikely to pose an important threat to users.

In conclusion, much research has been done on WAT inference, where a large majority of these inferences raise privacy issues. For example, all information about consumption (e.g., eating, drinking, smoking), activities (e.g., sport), location (e.g., city name), or disease (e.g., cancer), can be directly used by adversaries (e.g., health insurers, employers, and advertisers) to target their customers and/or even to discriminate against them.

### 2.1.2 Research on WAT Users

We review WAT privacy studies conducted with users. This review provides a comprehensive understanding of users' awareness & knowledge about privacy, their concerns, attitudes & behaviors, the roles of individual differences on users' perceptions, and the utility-privacy trade-offs.

#### WAT users' Privacy Awareness and Knowledge

Many studies have assessed users' awareness of and knowledge about privacy. Overall, WAT users have **limited knowledge** about the privacy policies of service providers [98, 99]. Most users are not aware of who has access to their data, and that their data are transmitted, stored, and used [99, 100, 101]. Vitak et al. [98] find that, after they were asked to read the relevant part of the terms of service, most users are not aware of what they have given consent to and were surprised about the extent of access they provided to service providers. Several misconceptions have been identified, such as not being able to distinguish privacy from security and being overconfident about privacy knowledge [102]. Some users think WATs are secure because they do not have an "input" device (e.g., a keyboard), hence users cannot input sensitive information such as passwords [102]. Most users also cannot judge the difference between storing data on a cloud and a device [103]. This lack of awareness could be due to a lack of interest in learning about how their data is used [101]. Finally, most users are not aware that data from motion sensors could be used to infer passwords entered on WATs and that, as people tend to use the same code for diverse applications and devices (e.g., ATM PIN codes), the risk of such attacks increases [104]. Privacy awareness is negatively associated with data-sharing habits, whereas non-aware users tend to share more [105].

Gabriele and Chiasson [19] show that WAT users tend to believe that most privacy risks are **unlikely to materialize**. Unfortunately, WAT users first consider the likelihood of being subject to a privacy risk, and only then do they contemplate its severity [106]. Therefore, not knowing about the likelihood of such threats prevents them from thinking about their severity. Gerber et al. [107] show that users perceive privacy risk scenarios as *likely* if they are written in an abstract form. Many users consider privacy only from a “social privacy” point of view and do not think how their data could be used by third parties (e.g., advertisers, health insurance) [108]. Users’ knowledge about privacy also depends on the **type of information** collected by WATs. Rader and Slaker [109] argue that WAT users recognize sensors that they can see and verify (i.e., those that are physically visible). Velykoivanenko et al. [46] show that users think that sensitive information not directly related to a specific sensor cannot be inferred from their data.

### WAT Users Privacy Concerns

Privacy concerns affect the usage of WATs[110]. Users aware of privacy risks tend to be more concerned about their privacy. These users use coping mechanisms [111] and/or contemplate abandoning their devices [112]). Therefore, it is important to better understand their concerns. Due to low level of awareness of most users, most of them are not concerned about privacy. Alqhatani and Lipford [18] show that their participants had mainly utility-related concerns (e.g., to have a better self-image by data sharing) and not privacy-related ones. Several studies [102, 106, 18, 113, 114, 115] show that most users express only minor privacy concerns. The majority perceive their WAT data as harmless, innocuous, and not sensitive [102, 113, 115], and they report that they would share their data, without requiring that the privacy boundaries be managed [102, 113]. Lidynia et al. [116] show that their study participants did not consider storing data on the server (compared to their device) as a critical issue. However, such attitudes and (lack of) concerns could be due to a misunderstanding about the WAT ecosystem and the lack of a correct mental model [117].

Aktypi et al. [118] highlight that multiple factors reassure users about their privacy, especially the fact that they tend to trust WAT companies. However, there is no consensus about this trust. Although some studies show that users trust companies to handle their data [113] and that they believe in companies’ technical capabilities to preventing privacy breaches [115], others [46, 102, 119]

do not show this same confidence. Given the huge amount of data collected from millions of individuals, some users cannot see how their data can be used against them: “... *just a drop in the ocean*” [118, p. 8]. For some WAT users, privacy concerns evolve over time. Some start being concerned if their data is misused or after their privacy is violated (e.g., after a privacy breach) [114]. In the workplace context (for more details, see Sec. 2.1.4), at first, some users perceive their data as harmless; but over time, they report different concerns as their data creates many inter-colleague discussions that reveal their private-life activities and cause social pressure [120]. Interestingly, with the participants of research experiments, those who usually are unconcerned about privacy expressed their concerns about WATs after being confronted with questions about their private life [101, 102]. This could be due to the well-known privacy paradox [121], where users report having privacy concerns, but then they behave as if they do not have these concerns. Finally, Vitak et al. [98] shows that the more users perceive their WAT data as valuable for third-parties, the more privacy concerns they have.

Earlier studies identified **concerned users** who prioritize their privacy and use thicker privacy boundaries to protect their information [102, 103, 113]. Three types of concerns are recognized among such users. **(1) Data Collection and Storage:** concerns about the anonymization of data [119] and the location where the data is stored [116, 119]. **(2) Control over Data:** concerns about the data being used for purposes other than for the main purpose or being shared with third parties [102, 119]. Some users think they have limited control over disclosing their own data [102, 122, 123]. They also mention the *forced-choice dilemma* where they have to decide between using the device (and facing the consequences) and not using it. Lastly, they mention the *post-purchase lock-in effect* where privacy policies might change after agreeing to them. **(3) Storage Security:** Some users are concerned about their devices or the service providers’ platforms being hacked. They think that security breaches could lead to negative consequences [123, 124].

### WAT users’ Sharing Attitudes

WAT users’ willingness to share WAT data is strongly related to the **type of data** and the **audience** they intend to share their data with [19, 60]. If more data than step counts are shared, users worry. They perceive location data to be the most sensitive data type [113, 116, 125, 117, 126]. They are concerned about the negative consequences of sharing location-data, such as

home burglary and bike theft [126]. But if they sell their data, they would ask for significantly more money for their location data than for health-related data [125]. Also, WAT users are more reluctant to share movement data, other than step data [114]. Weight and sleep data [116] and any data related to personally identifiable information (PII) and financial information [18] are perceived as particularly sensitive.

However, even for the most sensitive data, WAT users change their sharing decisions, based on the intended recipients. They generally seem willing to share their location with their friends, whereas they do not want to share with online advertisers [19]. Schneegass et al. [60] found that users' willingness to share is inversely proportional to the size of the recipient group they share the data with. This finding is in line with other studies [19, 18, 115, 116], wherein users would be willing to share their data with small groups of people, such as their family, friends and colleagues, and/or with health practitioners if they ask for it; but they would not share with the general public, employers, insurance companies, banks, and advertisers. Finally, Alqhatani and Lipford [18] study WAT users' motivations for sharing data with different types of recipients. For example, they share data with friends to compete or to show a positive image of themselves, whereas they share data with family members to encourage and motivate one another to be healthier.

### Individual Differences

Individual differences play an important role in WAT users' privacy awareness, concerns, and attitudes. For example, older users tend to be more relaxed about data sharing [60] and perceive their data as less valuable to third-parties [98], although they give their data more personal value. Women tend to share more data than men do [60]. Future studies should replicate these differences due to age and gender, and they should investigate the underlying reasons behind such differences. The findings of studies about the differences between users from different regions are rather inconsistent. Ilhan and Fietkiewicz [127] find significant differences, regarding their level of concern and awareness, between WAT users from the US and Germany. Whereas, the same group of researchers did not observe any differences between users from the US and Europe [128].

Earlier studies categorized WAT users into **different classes** such as (1) non-sensitive and (2) sensitive users [116], (1) unconcerned, (2) somewhat concerned, and (3) highly concerned [102], and as (1) data protectors (i.e., those

concerned with privacy), (2) benefit maximizers (i.e., those concerned with utility), and (3) fact enthusiasts (i.e., those most concerned with motivational design) [103]. Individuals can also be differentiated as users, former users, and non-users. This can help us to understand their reasons for using technology or abandoning it, and to understand if they would contemplate using such technology in the future. Previous studies [129, 128] show that *non*-users of WATs are more concerned than users about the collection of WAT data. Surprisingly, former users are less concerned about privacy than actual users [128]. In contrast, Bélanger et al. [130] do not find any significant difference between privacy concerns of users and non-users. This difference may be due to the fact that the population studied was not the same as the first study, which is mostly about European Union citizens ( $\sim 80\%$  of their respondents) [128] while the former is about US citizens only [130]. Lastly, in a more quantitative approach, Fietkiewicz and Ilhan [128] show that it is possible to categorize WAT users using clustering techniques (e.g., k-means).

### WAT users' Behaviors, before and after Privacy Violation

Overall, WAT users take limited actions to protect their privacy [113]. Coping strategies vary depending on their concerns and threat perceptions [131]. Some report adjusting the privacy settings of their WATs only immediately after setting up their device (i.e., after unboxing), whereas others could not remember when they changed them, and still others thought they were using the default settings [113]. In the context of the workplace, users might consider partial sharing if they could exclude specific parts of their data related to private situations [114]. Velykoivanenko et al. [46] unveiled that a minority of WAT users ( $\sim 5\%$ ) consider removing their device for privacy-related reasons (e.g., before engaging in sexual activities). Some users reported that privacy settings are complex and that they have difficulties adjusting them [113, 126].

Besides the privacy management behavior, several researchers studied users' coping behavior **after they faced privacy breaches**. Lehto and Miikael [115] asked WAT users what they would do if their service provider had a security breach. Surprisingly, none mentioned that they might stop using their WAT; however, they said this might affect their future WAT purchases. Other studies showed that though users' privacy perceptions do not have an effect on their avoidance motivation (i.e., privacy management behavior) [106], their perceptions can affect their coping behavior [132]: Higher privacy concerns increase users' threat perception, which has an effect on an individual's coping

behavior. Theoretical studies found that users use two main coping mechanisms [111, 132, 133]: (1) emotion-focused coping when the perceived level of threat is high and the level of efficacy is low, and (2) problem-focused coping when the perceived level of threat is low and the level of efficacy is high. Therefore, in the event of a privacy breach, users will likely not be able to show rational behavior and would instead seek emotional support.

### Trade-Offs between Utility and Privacy

According to privacy calculus theory [37, 134], technology users always weigh the perceived benefits and risks. Perceived utility and privacy concerns affect users' intentions to use their devices [135, 37]. Several studies [113, 124, 130, 136] found that **WAT users prefer to take a utilitarian approach and that the perceived benefits can outweigh their privacy concerns**. They usually perceive a fairly positive effect from data sharing [18, 137]. However some users, especially older adults [138], do not make rational trade-offs by ignoring/underestimating the risks [101]. Furthermore, daily WAT users often willingly share data, despite compromising their confidentiality, as they find the health and social benefits worth the risk [102]. They sacrifice privacy to receive immediate financial benefits, such as a reduction in insurance fees [126] or a higher wage [117]. Although users tend to express concerns when they carefully read previously-agreed-to data-collection policies, they would not change their usage behavior [100].

Following Nissenbaum [139]'s definition of privacy (a.k.a. contextual integrity), earlier studies [122, 140, 130] show that **WAT users' utility–privacy trade-off depend on context**. Ebert et al. [129] show that WAT users are marginally concerned about privacy more than loyalty-card users are. Lehto and Miikael [115] discuss that individuals consider their health data (collected by their doctors) as private/sensitive, unlike data collected from WATs; and they consider financial information as the most sensitive. Furini et al. [125] show that, when given a strong altruistic motivation (e.g., sharing data for contact tracing for COVID-19), users tend to agree to share their data. Similarly, research participants might be willing to share their data, as they consider it a donation and contribution to science [117]. Finally, Velykoivanenko et al. [46] argue that users' concerns about the inference of certain types of information (e.g., religion and sexual orientation) are heavily dependent on the social norms and conditions in their country of residence.

Although the utility-privacy trade-off is often imbalanced toward the side

of utility, it can be further explored by researchers and designers in order to create privacy-enhancing solutions. For instance, where users do not use a particular feature, turning off that feature (i.e., data minimization) could help in privacy protection (see Section 2.1.5).

### Monitoring Family Members

Several studies analyze users' privacy in the context of using WATs in families, between different generations, and between couples. Kuzminykh and Lank [141] show that parents are interested in monitoring their children's health and activity levels, but not to the extent that it would compromise their relationships or prevent children from developing self-sufficiency. However, Jørgensen et al. [142] show that usage of WATs by parents for monitoring their children can deteriorate trust in both directions. Similarly, Li et al. [143] find that younger users worry about their family members' opinions about them, based on their WAT data. Potapov and Marshall [144] reveal children's concerns about their data being misused by their teachers in a school context. In a different context, Leitão [145] shows that WATs can be used by abusive partners for stalking, threatening, and harassing (a.k.a. intimate-partner abuse).

## 2.1.3 Privacy Policies, Regulations, and Ethics

### Privacy Policy

As a means of communication between service providers and users, privacy policies are used to inform WAT users about the data collection and usage practices and to obtain their permission. However, their usability and compliance with users' privacy needs and data-protection regulations (e.g., GDPR<sup>3</sup>) is still under debate. Many studies have reviewed WAT-related privacy policies. Braghin et al. [146] argue that privacy policies are of “*dubious validity*.” Users report **a lack of (legal) accountability** in cases of privacy breaches [118]. Paul and Irvine [147] reveal many statements that have the potential to violate user privacy in the privacy-policy content of four market leaders in 2012.<sup>4</sup> Several studies present heuristic frameworks for evaluating privacy policies. Katurura and Cilliers [148] show that both Fitbit and Apple

---

<sup>3</sup>General Data Protection Regulation, see <https://gdpr-info.eu/>, last accessed: Dec. 2022.

<sup>4</sup>Note these findings are from almost a decade ago. Some of these products are no longer sold, and some policies might have been amended.

did not provide minimal protection for choice or consent: Before they collect data, these companies ask for consent; but after the collection, the users were not permitted to enforce how their data is used. Hutton et al. [149] compare the privacy policies of self-tracking apps in different domains and show that apps related to WATs generally met fewer heuristics compared with apps related to other types of tracking (e.g., time management, cost management). Becker et al. [137] show that the type of statements used in privacy policies can influence WAT users' decisions about disclosing their health information (e.g., policies framed positively).

Another issue with privacy policies is the **usability problem**. They are lengthy, complex, and annoyingly profuse, thus users often do not read them. Users report not reading the privacy policies of their WATs to avoid cognitive load; furthermore, they perceive their acceptance as a binary choice (i.e., forced choice dilemma) hence as a necessary condition to use the device [126]. Gluck et al. [150] show that shortening the privacy policies to some extent can be an effective way to increase user awareness. Guo et al. [151] propose a visualization tool, named Poli-see, for helping users understand WAT privacy policies. Drozd and Kirrane [152] present CURE, a GDPR-compliant consent-collection system that obtains users' partial consent in a more usable fashion and that provides the users a better explanation of the consent they have given. Murmann et al. [153] study the adoption of **privacy notifications** (e.g., notifying users when their data is stored on a cloud or when it is transferred to another country) and show that most of their respondents perceived notifications as useful. Masuch et al. [154] show that confidence-building mechanisms (i.e., statements by service providers about how data will be treated securely) resulted in an increase of the users' expectations about the security of the service. However, users observed a large discrepancy between expectation and reality; this negatively influenced their satisfaction and intentions to continue using their WATs. Thuraisingham et al. [155] propose a (hypothetical) privacy-aware data-management framework to enable users to manage the collection, storage, sharing, and analysis of their own data.

### Protective Law for WAT users

There have been several studies about existing regulations, laws, and policies that could protect WAT users' privacy. Most of these works study the regulations in the US and in Europe. In **the US**, there are several relevant regulations, however, none are effective [156, 157, 158, 159]. More specifically,



WAT users are not affected by federal legislation, such as the Health Insurance Portability and Accountability Act (HIPAA) or the Health Information Technology for Economic and Clinical Health Act (HITECH Act), as they are not expansive enough to address WAT data. WAT data is not counted as protected health information (PHI) because service providers are not covered entities, unlike hospitals or clinics [160]. In the case of WAT data being stored by a covered entity, HIPAA is applicable only for data processing and disclosure and not for data collection [161]. Similarly, the Food and Drug Administration (FDA) classifies WATs as low-risk wellness products [158, 162]. As a result, WAT data is not protected by the US Federal Food, Drug, and Cosmetic (FD&C) Act [156, 157]. The Privacy Act of 1974 is another relevant law that regulates the collection, usage, and disclosure of PII. But the definition of PII in this act is rather limited [157], as it includes only information such as names, e-mail addresses, and social security numbers. Similarly, WAT data is not protected by the Electronic Communication Privacy Act (ECPA) [156, 157], as the ECPA does not include devices that use radio frequency identification (RFID).

Researchers advocate for new WAT regulations, recommended including WAT data in existing frameworks, such as the Privacy Act of 1974 [157], and expanding terminologies such as “covered entities” and “third parties” to include service providers [158]. Brinson and Rutherford [157] also developed a portal to help users and data brokers interact and determine the use of their data.

Several studies on legislation in **other countries** have been conducted. Daly [163] discuss that the most important source of WAT regulation in Australia is the Therapeutic Goods Administration (TGA). However, the TGA’s regulations can be easily avoided if WAT manufacturers do not intend for their WATs to be classified as medical devices (as defined by the TGA). Similarly, Katurura and Cilliers [148] show that the Protection of Personal Information Act (POPIA) in South Africa cannot force foreign manufacturers to comply. Compared to other countries, the GDPR provides better protection for users in the **European Union (EU)** [159, 161, 164]. The GDPR has several advantages. First, it forbids processing personal data, except in far-reaching conditions (i.e., if they are anonymized) [159]. Second, it forbids the processing of data concerning health, unless the patient has explicitly consented.<sup>5</sup> This

---

<sup>5</sup>Other exceptions include when the processing is necessary to protect the vital interests of the patient or of another person, to perform another contract for the patient, to carry out a task of public interest or of other legitimate interests, except when such interests are

affects the collection of health-related data such as heart-rate data. Third, it is an enforceable law and is applicable to foreign manufacturers who export their products to the EU [161, 164].<sup>6</sup> This is further supported by the Privacy Shield 2.0.<sup>7</sup> Fourth, it permits the use of *anonymized* data for science and research purposes and for the sake of technological development and demonstration [159].

### Use of WAT Data in Investigations

WAT data, such as profile information and activity data, can be used as evidence in **forensic** investigations regarding, for example, suspicious deaths, airplane crashes, malpractice [165], or even detecting police brutality [165], especially in cases of racial injustice [166]. WAT data integrity can also be assessed, for example, insurance companies can check whether a reported activity was created artificially [167]. Several studies present software tools for forensic science [167, 168] and guidelines for investigators [168, 169]. Other studies [168, 169, 170] show the forensic soundness of their tools or guidelines by using existing WATs, such as Fitbit, Xiaomi, and Huawei. Only, one study fails to recover information, after a forensic analysis [171]. It used a real-life scenario instructing a participant (with a Fitbit) to walk to a specific location and hit the ground several times then to return to their point of departure. Future studies should use similar real-life scenarios to validate the reliability of forensic methods.

Courts and forensic investigators can face **several challenges** that reduce the objectivity of judicial decisions [158, 165] in order: (1) to ensure the accuracy of measured metrics, (2) to ensure data integrity by confirming that the data were not changed after an incident and that WAT was not worn by other individuals, and (3) to handle massive amounts of data and still create precise statistical/inference models, even if part of the data is missing. Finally, it is necessary to maintain WAT users' privacy during forensic investigations; in particular, in interdependent privacy situations [172]. Hassenfeldt et al. [173] show that using web scraping and leaderboard information from Strava, they can access other users' information, regardless of whether their data was pri-

---

overridden by the interests and personal data protection rights of the patient. For details, see Art. 6 GDPR: <https://gdpr-info.eu/art-6-gdpr/>, last accessed: Dec 2022.

<sup>6</sup>Here, the product includes all equipment and covers mobile devices, applications, services, and wearable devices.

<sup>7</sup>EU-US Privacy Shield framework, see <https://www.privacyshield.gov>, last accessed: Dec 2022.

vate or public. Kumari and Hook [158] argue that courts should try to obtain data from the users themselves, or from their acquaintances. Accordingly, asking service providers to share data should not be the first option.

## Ethics

In addition to analyzing legislation, **ethical implications** of using WATs for their users are analyzed in several studies [162, 174, 175]. Lupton [174] uses the term “*dataveillance*” (i.e., digital surveillance of individuals) to explain how WAT use can lead to “*function creep*” (i.e., using data for purposes other than living a healthy and active lifestyle). Tuovinen and Smeaton [162] define the term “*wearable intelligence*” as the convenience and simplicity of using WATs. They discuss that, unlike in the context of a black box, users need to know that the information presented to them is an approximation generated by computational models and not absolutely accurate. Also, they warn about a potential *power imbalance* between non-expert users and expert data-analyst entities; as this imbalance can cause further privacy and trust issues. Steinberg [175] discuss the fairness of insurance companies that use WATs as incentive programs, where WAT users can receive a discount on their premiums if they choose to share their data with their insurer.

In addition to taking ethics into consideration for WAT users, researchers should be also mindful of **research ethics**. The collection of WAT data can serve in the development of ML models to infer users’ states and to propose proper interventions for them. It has become common practice to collect such datasets and to share them with the public to support open science. Publicly sharing such a large volume of datasets has privacy risks for the data owners and ethical risks for designers (i.e., designing interventions based on biased datasets). Lee et al. [117] conduct a risk-benefit assessment with WAT data owners. The results show that financial compensation was the main incentive for data owners. Some data owners accept to provide even more data in order to receive even more money. Among those who refused the offer, some mentioned they could accept, but only if the compensation amount was higher. Less than half of the data owners thought they were subject to surveillance. Some also mentioned a lack of trust about how data would be handled by researchers. Given these vulnerabilities, it is important to protect WAT-data owners after data collection. We recommend, beyond routine practices, such as using informed consent and anonymization, researchers should consider data

sharing with *restricted access*. Among the FAIR<sup>8</sup> open science repositories, Zenodo provides an option for restricted access,<sup>9</sup> where data can be stored privately on the platform, and researchers can share access to it only after certain agreements.<sup>10</sup>

#### 2.1.4 Health at Work or Workplace Surveillance?

In the context of workplaces, existing studies show that employers have a vested interest in promoting the use of WATs for their employees [114, 161]. This creates a profitable business for WAT manufacturers, as they can sell more of their products (and additional services) to companies.<sup>11</sup> Companies intending to adopt WAT-based wellness programs follow either **wellness model** or **performance management model** [164]. Whereas the former is used to promote healthy lifestyle habits and to enhance the well-being of the employees, the latter aims to increase efficiency, productivity, and safety.<sup>12</sup> The concern with the first model is employees' privacy. Whereas the second is even more serious, as data can be used to monitor and detect misconduct, hence it could have a long-term impact on employees' careers.

Most studies focus on the first model [120, 178, 179]. Many employees report perceiving wellness programs positively. They usually participate in such programs to improve their awareness of their activity levels, to become more physically active [178], or to socialize [120]. During the campaigns, employees can become concerned about the erosion of the boundary between their work and personal life. However, they also tend to discuss their WAT data with colleagues (as an ice breaker for conversations during breaks). In the workplace, discussions about step counts or activities can increase social pressure, breach privacy boundaries, hence raise tensions. Studies show that not all employees are happy to join such campaigns and some decide to not join [120, 178]. Furthermore, after the end of the campaign, employees usually return to their previous activity routines.

Given the lack of evidence of the long-term benefits of wellness campaigns and the social distance created between participants and non-participants,

---

<sup>8</sup><https://www.go-fair.org/fair-principles/>, last accessed: Dec 2022.

<sup>9</sup><https://about.zenodo.org/policies/>, last accessed: Dec 2022.

<sup>10</sup>For example, a data recipient must agree to not make data public and not infer data owners' identities.

<sup>11</sup><https://healthsolutions.fitbit.com/corporatewellness/>, last accessed: Dec 2022.

<sup>12</sup>To read comprehensive surveys on the use of wearables for safety at work, see [176, 177].

Gorm and Shklovski [179] suggest reconsidering the notion of “*success*” in such campaigns. Marassi and Collins [164] discuss the privacy and autonomy concerns of wearing WATs in the workplace and express many reservations, especially about the employees’ **right to bodily integrity**, **life-work boundaries**, and the **power imbalance** between employers and employees. In the US, there is no legislation that protects employees’ privacy [156].<sup>13</sup> In the EU, GDPR does not permit employers to monitor their employees. To address these issues, previous studies recommend (1) clarifying the terms and implications of information disclosure to employees [156], (2) proposing new laws that limit data collection by employers [156], and (3) using a coaching-based approach, where employers can use third-party services as mediators that provide health advice to their employees [164].<sup>14</sup>

### 2.1.5 Privacy-Enhancing Technologies (PETs)

As an addition to Alqhatani and Lipford [181]’s work that reviews existing PETs provided by known WAT brands, our work reviews the PETs proposed by the literature.<sup>15</sup>

#### Anonymization Techniques

Given the high dimension and sequential time-series nature of WAT data, anonymizing such data-sets is challenging. Na et al. [182] show that accelerometer data can be de-anonymized with high accuracy. Multiple studies focus on methods for effectively anonymizing WAT data. Parameshwarappa et al. [183] use a multi-level clustering anonymization technique to prevent the re-identification of WAT users. Gong et al. [184] propose a theoretical framework for federated learning that preserves individuals’ privacy and trains an ML model by using multiple WATs’ data. Garbett et al. [185] designed ThinkActive: an activity-sharing platform for classrooms that enables students to use pseudonymized avatars to share WAT data, without exposing their identity.

---

<sup>13</sup>For example, there was a case in California where an employee’s claim that their employer violated their privacy by linking their Apple account to a work-related device was rejected by a court [180].

<sup>14</sup>In such a case, the employers should continue to be the data controller (rather than processor), and the coaching service should be the data processor.

<sup>15</sup>PETs for usable privacy policies [151, 152, 155] are already discussed in Section 2.1.3.

### Limited Sharing and Data Minimization

Wang et al. [186] study user preferences and sharing behavior related to partial-data release. Epstein et al. [187] investigate if fine-grained step-count sharing can help WAT users preserve users' privacy while they share activities. Velykoivanenko et al. [46] assess users' utility perceptions to inform future PET design. They also show that there is a high potential for implementing *data minimization* that can avoid certain privacy risks. Finally, Kalupahana et al. [188] propose a framework to use random noise from WAT sensors in order to generate noise for differential privacy protection.

### Pedagogical Solutions

Torre et al. [56] model the complexity of WATs and TPAs in order to compute the probabilities of inferring different information from WAT data. Their model is designed to show WAT users that they can protect their privacy by not sharing certain data. Aktypi et al. [118] design a pedagogical tool that informs WAT users of the risks they are exposed to when sharing certain WAT data (e.g., running route), together with other information (e.g., the information available on their social media). Alvarez et al. [189] show that watching a video about privacy and security risks of collecting and sharing WAT data can significantly improve attitudes toward cybersecurity, whereas a text version of the information has no significant effect. Sanchez et al. [190] model the privacy preferences of WAT users and developed a system for recommending personalized privacy settings for users in different scenarios.

### Others

Data integrity is critical for healthcare providers and insurance companies that are interested in users' WAT data. du Toit [191] designed PAUDIT, a decentralized data architecture that enables users to store their WAT data in a personal online data store and permits healthcare providers to read data and audit the logs (i.e., changes made to the access control list). Ghazinour et al. [192] propose an access-management tool that enhances users' decision-making by enabling them to share their WAT data after considering four aspects: purpose (why), visibility (who), granularity (how), and retention (when). Liu et al. [193] propose a machine-learning framework to provide WAT users with personalized fitness recommendations without collecting personal information. Finally, Kazlouski et al. [50] analyzed unnecessary communication from the

Fitbit companion app (as well as six of the most used TPAs) to their business partners and propose an easy-to-use methodology to block them.

## 2.2 WAT Security

We review attacks on WATs' Bluetooth communication and discuss various vulnerabilities related to companion apps. Then, we investigate how WAT data can be used to bypass security systems. We also review different authentication methods related to WATs. Lastly, we analyze security protocols and threat assessments.<sup>16</sup>

### 2.2.1 WAT–Phone Communication

A large amount of research has been conducted on WATs and Bluetooth security. Multiple attacks, privacy issues, and mitigation techniques were identified. Table 2.1 shows all the studies related to Bluetooth and Bluetooth Low Energy (BLE) security and WATs. By analyzing these studies, we identified six main types of attacks: tracking, eavesdropping, injection, denial of service (DoS), traffic analysis, and firmware modification.

**Tracking** is being able to locate or identify the presence of a specific device. Several studies [194, 196, 202, 146] analyze how WATs, from multiple vendors, communicate with the companion app (generally installed on a smartphone). They show that **all of the tested WATs use permanent BLE addresses, which makes them vulnerable to tracking attacks**. Although these previous studies state that using address randomization should mitigate the tracking attack, recent studies [203, 204] show how generic attribute (GATT) profiles<sup>17</sup> can be used to build unique fingerprints. Becker et al. [205] developed a method to track BLE devices by using features extracted from the payload of advertising messages.

An **eavesdropping attack** consists in intercepting data communication between two devices, whereas an **injection attack** consists in sending additional (i.e., fake) data to a specific device. Except for one of them, all the analyzed studies describing eavesdropping attacks are also about data-injection attacks. Both types of attacks can be performed using similar techniques, such

---

<sup>16</sup>Note that multiple attacks that we review are not necessarily specific to WAT devices (e.g., attacks on Bluetooth or HTTP communication).

<sup>17</sup>GATT profiles are available without any authentication and contain basic information about features and services.

Table 2.1: All articles about Bluetooth security and WATs. For each paper, the table shows what type of attacks is described and tested.

| Article                        | Tracking | Eavesdrop. | Data Inject. | DoS | Traffic Anal. | Firm. Mod. | Passive | Active |
|--------------------------------|----------|------------|--------------|-----|---------------|------------|---------|--------|
| Das et al. [194](2016)         | •        |            |              |     | •             |            | •       |        |
| Lotfy and Hale [195](2016)     |          | •          | •            |     |               |            | •       | •      |
| Goyal et al. [196](2016)       | •        |            |              | •   |               |            | •       | •      |
| Rahman et al. [197](2016)      | •        |            | •            | •   |               |            | •       | •      |
| Zhang and Liang [198](2017)    |          | •          | •            | •   |               |            | •       | •      |
| Fafoutis et al. [199](2017)    |          |            |              |     | •             |            | •       |        |
| Shim et al. [200](2017)        |          |            |              |     |               | •          |         | •      |
| Classen et al. [201](2018)     |          | •          | •            | •   |               | •          | •       | •      |
| Mendoza et al. [202](2018)     | •        | •          | •            |     |               |            | •       | •      |
| Braghin et al. [146](2018)     | •        | •          | •            |     |               |            | •       | •      |
| Celosia and Cunche [203](2019) | •        |            |              |     |               |            | •       |        |
| Zuo et al. [204](2019)         | •        | •          | •            |     |               |            | •       | •      |
| Becker et al. [205](2019)      | •        |            |              |     |               |            | •       |        |
| Hale et al. [206](2019)        |          | •          | •            | •   |               |            | •       | •      |
| Wang et al. [207](2020)        |          | •          | •            |     |               |            | •       | •      |
| Gouda et al. [208](2020)       |          | •          | •            |     |               |            | •       | •      |
| Barman et al. [209](2021)      |          |            |              |     | •             |            | •       |        |
| Fúster et al. [210](2023)      | •        | •          | •            | •   |               |            | •       | •      |

as *a man in the middle* (MitM) attack. Several studies [195, 198, 146, 204] show that multiple WATs use un-encrypted communication, either while already paired or during the pairing process with a smartphone. They even permit pairing without authentication. Therefore, an attacker can retrieve information about the devices, and then they can proceed toward more sophisticated attacks, such as a MitM, which can lead to eavesdropping and data injection (even after pairing). Rahman et al. [197] reverse-engineered two WATs (Fitbit and Garmin) and built a framework that can perform various attacks, such as injecting data into the devices. Other studies [201, 207, 208] performed attacks that force a device to be paired with a fake companion app that grants access to all transmitted data before redirecting it; it was also able to inject data and commands. Mendoza et al. [202] analyze one of the most



popular WATs and show that its communication with a paired smartphone does not follow the BLE security specifications and that the device accepts connections from unknown smartphones.

Hale et al. [206] developed an open-source platform that aims to be used by security & privacy researchers to facilitate wearable security investigations. The platform could be used to collect data, conduct attacks, and identify security vulnerabilities. They used their platform to analyze BLE communications of multiple WATs and observed that all of them use encryption protocols to communicate with their companion apps. There has been a large amount of research on eavesdropping on and injecting data through WAT Bluetooth communication. **Overall, WATs tend to not use any protection mechanisms. Most WATs do not implement protection mechanisms, such as basic cryptographic schemes** [206], and they send un-encrypted traffic, mainly for optimization reasons (e.g., save battery). Both eavesdropping and data injection attacks are successful.

Four of the aforementioned studies also describe **denial of service (DoS) attacks**.<sup>18</sup> Goyal et al. [196] performed a DoS attack on a Fitbit Charge by spamming it with requests that prevent it from communicating (to synchronize or send other data) with the companion app on a paired smartphone. Rahman et al. [197] developed two different DoS attacks against Fitbit and Garmin devices. They show that it is possible to quickly drain the WATs' batteries by spamming them with BLE requests. Furthermore, they show that an attacker can overload the WAT's storage by injecting fake data; this overloading can cause various problems (e.g., being unable to store newly collected data, display issues, etc.). Zhang and Liang [198] also show that attackers can conduct DoS by continuously sending fake commands (as it is possible to inject fake commands to WAT devices as Syntrino's TW64 and LifeSense's Mambo HR). Classen et al. [201] demonstrate that DoS attacks can be performed on Fitbit WATs by injecting commands to enable the alarm clock or disabling the WAT's functionalities. They also show that it is possible to disable pairing and data synchronization.

**Traffic analysis** consists in trying to bypass encryption by using meta-data and signal treatment to infer some characteristics of the un-encrypted message. Das et al. [194] show that, by analyzing BLE traffic patterns, it is possible to identify individual users, with high accuracy. Fafoutis et al. [199] analyze, by using BLE, the correlation between activity levels (based on

---

<sup>18</sup>DoS attacks occur when access to a service is temporarily blocked by overloading the host machine or network with requests.

device acceleration) and the received signal strength (RSS) in the context of a WAT communicating with a smart home system. The results show that the RSS and the un-encrypted data are strongly correlated. Finally, Barman et al. [209] report that **a large amount of information can be inferred from encrypted Bluetooth traffic between a WAT and its paired smartphone**, such as the type of device, actions, and the type of data.

There are several studies about **firmware modification**. Shim et al. [200] analyze a WAT and its companion app's APK. Using reverse-engineering, they analyzed the BLE communication (to understand the communication protocol) when the companion app attempts a firmware update of the WAT. This enabled them to create a fake gateway for injecting malicious firmware updates. Similarly, Classen et al. [201] reverse-engineered Fitbit's firmware to study how to modify it in order to build custom firmware. They show that attackers can use un-encrypted BLE communication to flash modified firmware onto Fitbit devices. As explained above, most WATs use un-encrypted communication.

In conclusion, our review shows that WAT Bluetooth security has been studied extensively. Studies find that the public attributes of the transmitted packets can be used to track the devices and that communication is often not encrypted, which can lead to eavesdropping, data injection, or firmware modification. MitM attacks can be performed to bypass (basic) encryption mechanisms. An attacker can inject fake commands to conduct DoS attacks. Traffic analysis can disclose sensitive information, even if the communication between a WAT and its paired smartphone is encrypted. Some studies (e.g., [206]) also provide tools that researchers and developers can use for their work, while others conduct a comparative analysis of multiple models and brands available on the market by testing a large number of different types of attacks [210]. Finally, whereas some studies proposed mitigation techniques, only a few of them actually evaluated those techniques.

### 2.2.2 Phone–Server Communication and Data Storage

Several studies analyze the security of WAT companion-apps that are usually installed on the users' smartphones and are paired with a WAT to process, store, and to transfer online the data generated by the WAT. Companion apps transfer WAT data to the server or store it on the smartphone.

Goyal et al. [196] analyze the communication protocol and the data storage used by two models of WATs (a Fitbit Charge and a Jawbone UP Move). They analyze the code of the companion apps, how the data is stored on

the paired smartphone, the privacy policies, and the communication between the app and the service providers' servers (by sniffing the HTTP/HTTPS communication). They show that for both devices the data stored on the smartphone is not encrypted and some of it is even shared with third parties. Rahman et al. [197] analyze the HTTP communication between Fitbit and Garmin devices and their servers. They show that the data was not encrypted, including user credentials for Fitbit. Fereidooni et al. [211] consider WAT users as potential adversaries. Users might want to send fake data to their service provider's cloud for financial gain.<sup>19</sup> They analyze multiple WATs and successfully injected, using MitM attacks, fake data into their corresponding servers. By reverse-engineering the companion apps, they show that multiple companion apps that only store data on the smartphone do not encrypt the data, which makes it easily readable and writable.

To inject fake data, Fereidooni et al. [211] also conducted MitM attacks between the companion app and the service provider. They performed a new attack directly on the WAT by reverse-engineering the hardware system and directly accessing the device's memory to inject fake data [212]. After synchronization, the fake data was correctly encrypted and registered by the companion app.

Mendoza et al. [202] analyze how the Fitbit companion-app communicates with Fitbit's servers by sniffing HTTP/HTTPS communication and how TPAs can access data using Fitbit's API. They show that authentication credentials are sent un-encrypted and that the OAuth 2.0 protocol<sup>20</sup> is not correctly implemented. This creates vulnerabilities that an attacker can use to gain access to or modify the data. Classen et al. [201] reverse-engineered the Fitbit companion app to study how to modify it. Modifying the app could enable attackers to associate it with another account in order to download a user's data. Finally, Kazlouski et al. [213] analyze the communication between two well-known (yet anonymized) WAT companion-apps and servers. They collected ground truth by using a MitM setup and sniffed the encrypted packets by using Wireshark. Then they computed correlations of the size and frequency of the packets with the activities, heart rate, and step count. Their results show that activities and meta-data of encrypted packets are strongly correlated and that it is possible to use meta-data to identify the occurrence and duration of several activities and even to estimate other information (e.g., estimating the heart rate). **As**

---

<sup>19</sup>This applies to cases where users participate in financially incentivized data-sharing schemes, such as corporate wellness programs.

<sup>20</sup>OAuth 2.0 is used to enable TPAs to access some of the data.

we can see, multiple devices do not implement adequately secure phone storage and communication with the server, which can lead to a threat as serious as eavesdropping or data injection.

### 2.2.3 Side-Channel Attacks

Multiple prior studies analyzed how WAT data can be used to perform side-channel attacks (i.e., conducting an attack based on extra available information, instead of using vulnerabilities of the security protocols). As the main purpose of WATs is to track users' movements, it is possible to use the sensor data to infer sensitive information, such as the words a user writes, their typing on a keyboard, or even, their biometrics.

To this end Maiti et al. [6] studied how WAT (Samsung Gear Live) sensor data can be used to recognize typing patterns hence to infer which words are typed on a computer keyboard. Such attacks can be used by adversaries to collect passwords for bypassing authentication systems. Similarly, Maiti et al. [214] used smartwatch sensor data to infer which keys were typed on a 10-digit keypad and a QWERTY keypad on a smartphone. They reached an accuracy of 74% for the 10-digit keys and had a mean accuracy of 30% for the QWERTY keypad. Sabra et al. [215] and Wang et al. [216] show how similar attacks can be conducted to infer ATM PIN codes. The former obtains an accuracy of 80% for 6-digit PIN codes; this increases to 93% with 5 tries. Lu et al. [104] aimed to infer PIN codes and android pattern lock (APL) patterns. They find that it is possible to infer the APL pattern two-thirds of the time, within the first 20 guesses. Maiti et al. [7] studied the inference of rotary combination lock passcodes. Their results show that WAT sensor-data (especially gyroscope data) can be used to greatly increase the likelihood of inferring the lock combination. In addition to password retrieval, impersonation attacks are focused on by Eberz et al. [4], and they show that WAT sensor-data can be used to mimic an individual's biometrics (e.g., gait) hence to bypass such authentication systems.

In general, **we can affirm that using WAT sensor-data to bypass a security system is a potential threat that should be considered by vendors.** Several mitigation techniques are proposed. For example, WATs could deactivate sensors when they detect real-time activities such as typing [6]. Alternatively, WATs could add fine-grained noise to sensor data, in such a way that activities such as walking or swimming are still recognized, but where fine hand movements such as typing are not [216]. Or, users can simply remove their devices when they are typing.

### 2.2.4 Authentication

WAT data can be used to enhance security systems by using the collected data to authenticate users, and by either substituting or complementing other credentials.

Cola et al. [217] and Johnston and Weiss [218] study how gait (i.e., walking style) can be used as an authentication factor for WAT users, and show low error-rates (between 2% and 3% in both cases). Vhaduri and Poellabauer [219] propose a method for authenticate the users by uniquely identifying them based on both the physiological and activity data collected by WATs. They show that is it possible to recognize users, with high accuracy, by using their WAT data. Tehranipoor et al. [220] study how electrocardiogram (ECG) data can be used to authenticate a WAT users. The results show that the average entropy of an ECG-based key is 0.98, thus close to the maximum (i.e., 1.0). Therefore, ECG-based keys can be effectively used to identify users. Chen et al. [221] propose a novel authentication system that mixes credentials and biometrics by simulating a twelve-key keypad on a user's fingers. A user has to type a four-digit code on this virtual keypad and is authenticated if the code is correct and the biometrics correspond to the WAT's owner. Their results show this method is particularly effective and resilient against attacks. Moreover, they conducted a user study to evaluate the usability of the system and show that their system was most often the favorite one compared to all other proposed authentication methods.

Li et al. [222] implemented a robust authentication system using WATs to authenticate who is using a given IoT device. This software-oriented system is composed of an external server that securely communicates with both the tracker and the IoT device; and, to authenticate users, it compares movement data collected from the tracker and some basic IoT usage input (e.g, touching a button, turning a knob). To open a secure communication channel without transmitting encryption keys between two WATs, Shen et al. [223] developed a protocol based on handshaking patterns. This protocol is therefore secured, needs minimal extra effort from the users, and is designed to run without additional devices. Shrestha and Saxena [224] propose an authentication system for web service accounts based on users' activities. It uses WATs to compare wrist movements with the web service usage-data (e.g., keyboard and mouse movements) and rejects non-matching users. Sturgess et al. [225] developed an authentication system for NFC payments with smartwatches. This system detects the intent to pay and then authenticates the user when they want to proceed with a payment using their smartwatch and an NFC terminal. This

system prevents attackers from paying with stolen devices or from executing unwanted payments with unlocked devices worn by the user. However, in another study, the same authors showed that an attacker of approximately the same height as the user has a 20.6% higher likelihood of impersonating the user [90]. **Although WATs are, by design, equipped with multiple sensors, they are privileged devices for biometric authentication, either for the WAT firmware itself or for third-party services.**

### 2.2.5 Threat Assessment and Mitigation, and Security Protocols

Although a large number of studies related to security are about weaknesses, attacks, and privacy leaks, some of them are about **new protocols** and tools that can help preserve the security of systems. To protect against different attacks, Rahman et al. [197] propose an encryption protocol based on symmetric keys. They show that their solution has little effect on the device's performance. Using a system of tagged packets, Skalka et al. [226] developed a framework to manage and filter private data at the edge-router level.

Yan et al. [227] propose an ML-based method that uses received signal strength indicators (RSSI) to detect, with high accuracy, spoofing attacks from peripheral devices (e.g., additional sensors worn on the foot). Finally, Xin et al. [228] show that their new framework (SIAMESE\_MIL) is effective at detecting when data is injected in WAT sensor-data streams through specific data variations.

A few studies aimed to **identify and assess** the different types of **existing attacks**. To classify attacks, Mnjama et al. [229] developed a conceptual WAT threat assessment framework. They base their work on the CIA triad (i.e., confidentiality, integrity, availability) and on Microsoft STRIDE (i.e., spoofing, tempering, repudiation, information disclosure, denial of service, the elevation of privilege). They analyze different phases of WAT-data transmission and storage and the current health-wearable literature. They propose a framework for assessing the different existing types of threats, based on six core elements: authentication, authorization, availability, confidentiality, non-repudiation, and integrity. Moganedi and Pottas [230] identify all known vulnerabilities affecting WATs, from a privacy and security perspective. They discuss these vulnerabilities with regard to their corresponding parts of the WAT ecosystem and to how they are classified according to various existing standards. To classify the different currently known vulnerabilities, they

identify five main components in the WAT ecosystem (the WAT, Bluetooth, smartphone companion app, WiFi, and cloud storage) and six control families (access control, audit and accountability, identification and authentication, system and communication protection, system and information integrity, and PII processing and transparency).

## 2.3 Overview and Research Gaps

In this literature review, we have provided comprehensive information about the security and privacy of WATs, and revealed several related open issues. Interestingly, this review revealed several research streams that have been extensively studied. In particular, users privacy concerns have been substantially covered, as well as studies related to security breaches, and, in particular, related to Bluetooth communication. The literature shows that privacy risks are huge, diverse, and widespread, in terms of the information that can be inferred and of the consequences. Many communications and storage protocols are still vulnerable to different types of attacks such as eavesdropping or side-channel attacks. Service providers tend to implement as little protection as possible. This might be because service providers give high priority to device features and the final price of the product rather than to implementing security and privacy techniques that could be costly and less visible.

Whereas there are multiple studies about user concerns and attitudes toward such risks (e.g., data-sharing attitudes), there is no study, to our knowledge, about the *actual* behavior of WAT users, in particular toward data-sharing. Furthermore, while few studies evaluate WAT users understanding of the WAT ecosystem, they are not focused on how third parties (and in particular TPAs, which represent one of the main threats in our adversarial model) access WAT user data. We have also reviewed many studies on human activity recognition (HAR) and inference. Most of these HAR studies are however rather functionality-oriented in that they mainly highlight HAR benefits and focus on achieving high performance, and therefore do not focus on the risks related to such inferences and their consequences. Furthermore, they mainly focus on user activities and health and none of them are about user personal attributes (e.g., personality, religion, political views, or consumption habits). Moreover, privacy-oriented inference papers do not even consider users' activities or health, as most of them study the inference of data such as passwords or other types of information that could be used for authentication, but are not directly harmful to the users' privacy. Finally, regarding PETs, we noticed

that most of the provided solutions are device- or data-oriented. Moreover, some of the proposed mitigation solutions lack proper evaluations. Consequently, their effectiveness is questionable. Therefore, there is a need for more effective PETs. Following the principles of usable security, researchers should consider and study more user-centric solutions, as from a methodological point of view, most of the proposed PETs are not designed in a user-centric manner. In particular, we could not find any studies with participatory design or co-design approaches [231].

To address these research gaps, in this thesis we (1) conduct a study about the understanding and *actual* behavior of WAT users toward data sharing, (2) conduct a privacy-oriented study about inferring WAT user's personal attributes (i.e., personality traits) and discuss its consequences for their privacy, and (3) conduct a user-centered study to discuss and propose multiple solutions to minimize the risks for privacy when using a WAT.





## Chapter 3

# “Revoked just now!” Users’ Behaviors toward Fitness-Data Sharing with Third-Party Applications

**Abstract.** Although WATs enable their users to monitor their activities and health, they also raise new security & 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 chapter, 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.

## 3.1 Introduction

Attacks, and in particular inference attacks using WAT data 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 [18, 19]) and entities (e.g., employers [26], insurance companies [232], 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 & 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., [59, 60, 19, 46]), 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 [233, 234, 235, 236] and from the point of view of the security of the associated protocols (i.e., OAuth) [237, 238].

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 [239] by asking volunteer respondents to draw the data flow between WATs, TPAs, and other components. We also assessed their understanding of 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. This last point, related to RQ3, is particularly important as in this thesis, we eventually intend to propose multiple solutions to help users in the data-sharing process to better protect their privacy. Therefore, we aim to assess (1) how users perceive the solutions that we propose in privacy research, and (2) the potential for conducting future work about designing new PETs related to data-sharing.

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 the 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’ understanding were also highlighted after we analyzed their mental models.

**Outline.**

The rest of the chapter is organized as follows. In Section 3.3, we detail our methodology including participant selections, survey designs, and procedures. We describe the results of our analysis of the survey data in Section 3.4. We discuss the design implications of our findings in Section 3.5, and the study limitations in Section 3.6. Finally, we conclude the study in Section 3.7.

## 3.2 Related Work

Data sharing and access permissions have been widely studied in other fields, in particular with regard to online social networks and mobile permissions (mostly on Android). In their article, King et al. [233] explored what Facebook users understand about their data-sharing with TPAs and how they interact with them. Wang et al. [240] analyzed a large number of Facebook TPAs and their user-data collection behavior, then, they reviewed the permission process to show it can override users' general privacy settings and how the permission box dialog reflects the true app behavior. Krasnova et al. [234] studied the users' privacy concerns and attitudes toward data-sharing with TPAs on Facebook, whereas Wisniewski et al. [235] focus on how their concerns and attitudes are related to Facebook users engagement with their "Facebook friends". Arias-Cabarcos et al. [241] studied the effect of transparency on users' attitudes toward data sharing by confronting them to Facebook TPAs' behavior toward data sharing. Multiple studies have proposed different protection mechanisms to improve the online social network data-sharing ecosystem. Delgado et al. [242] developed a policy file-oriented solution to better manage data-sharing with TPAs on Facebook. After analyzing Facebook TPA data-collection behavior the permission user-interface, Wang et al. [240] proposed various solutions (i.e., alternative design for the panel) and helped the user to better manage their permissions. Shehab et al. [243], Anthonysamy et al. [244], and Cheng et al. [245] developed solutions to enable more flexibility in online social network data-sharing with TPAs (e.g., fine-grained data-sharing). Kontaxis et al. [55] developed a framework to only share a minimal amount of information with TPAs when opening a Facebook session on other websites. Ahmadinejad et al. [246] developed a framework to formally quantify the privacy and utility implications of sharing data with TPAs for online social network users.

As for mobile permission, Felt et al. [247] examined how Android permis-

sions are efficient in helping users to be attentive and to understand data access authorization during the installation of an app. Kraus et al. [248] analyzed how statistical information (e.g., the number of requested permissions compared to other apps) can help users in the privacy-utility trade-off. Shklovski et al. [249] conducted a qualitative study about users’ reactions when confronted with the data-collection behavior of their smartphone apps. Andriotis et al. [250] studied how users reacted and adapted to Android permission-system change. Liu et al. [251], Olejnik et al. [252], and Wijesekera et al. [253] developed frameworks to automatically assess new permissions according to those the user has granted in the past to help users in the permission-granting process, whereas Tsai et al. [254] used a user-centered approach to propose a feedback system to better involve users in the process of automatic permission and correct possible errors. Olejnik et al. [252] also propose a solution for fine-grained sharing related to mobile permission. Finally, Cao et al. [255] measured the actual behavior of Android users toward permissions.

Despite the previously cited work, there are still missing related works specifically about WAT users. The highly numerical and physiological aspect of WAT data opens the door to new threats that may require different types of PETS that can not be used for online social networks on mobile phone data. Furthermore, as a WAT is worn on the body, it collects data “passively” (without the user actively using the device) as long as the user keeps wearing it and thus collects more (and different types of) data than mobile phones or online social networks which (in most of the case) only collects data related to active usage. This is why, it is crucial to extend scientific knowledge about user behavior to WAT users.

### 3.3 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 priori to set the number of respondents. However, considering previous survey studies published on fitness-data sharing (e.g., Liao [59],  $N = 553$ ), 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.

### 3.3.1 Recruitment

We recruited our survey respondents via **Prolific** that was assessed as a reliable crowdsourcing platform for scientific research [256]. 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.<sup>1</sup> 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.

### 3.3.2 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

---

<sup>1</sup>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% [257]. This was confirmed by the results of our screener survey.

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. The questionnaire is available in Appendix A.1.<sup>2</sup>

**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).

Finally, we asked the respondents about their general usage of these TPAs

---

<sup>2</sup>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 herex and are not necessarily ordered by information type.



(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 [52], 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<sup>3</sup> 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, and 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 [18]: 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 [258], especially in the US where 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) [120, 164, 175].

---

<sup>3</sup>Only applicable for Fitbit and Garmin users, as Apple Health does not provide any public profile functionalities.

**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 3.1 and 3.2. The solution lets 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 3.3. 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 the permission granted to apps that are no longer used [259].

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

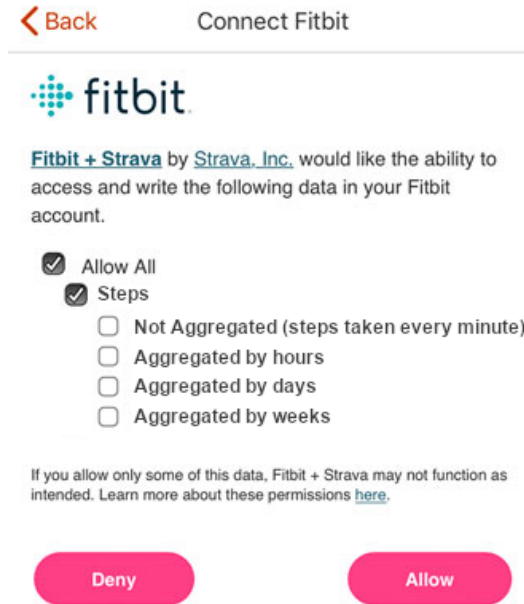


Figure 3.1: Illustration of the time granularity feature where a user can choose the aggregation level of the data they share with TPAs. “Not Aggregated (every minute)” is the default option on most WAT apps.

conceptualize different phenomena [260] and how they think different technologies work [239]. Given that verbalizing such tacit knowledge might be difficult for individuals (respondents in our case) [261], recent studies on security & privacy [262, 263, 46, 261, 264] 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 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.

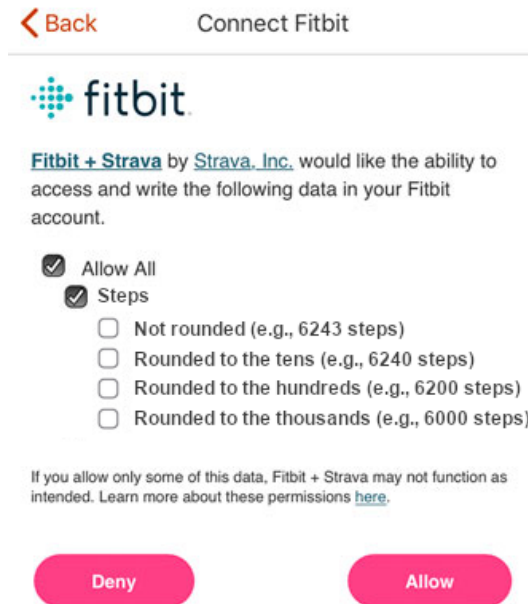


Figure 3.2: Illustration of the time granularity feature where a user can choose the aggregation level of the data they share with TPAs. “Not Aggregated (every minute)” is the default option on most WAT apps.

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 [265]. Figure 3.4 show the score distributions.

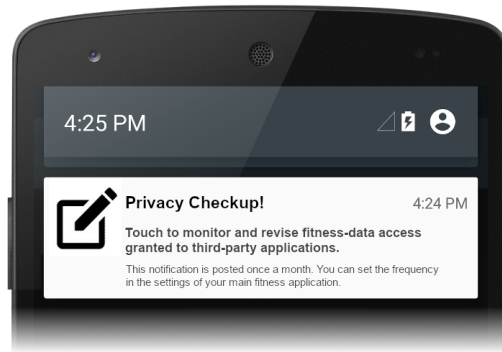


Figure 3.3: Picture presented to the respondents to illustrate the proposed TPA access monitoring reminder.




Figure 3.4: Global information privacy concern (IUIPC) w/ fit.

### 3.3.3 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 main investigator 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 main investigator 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. A 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.

### 3.3.4 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 [266]. 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. As this inconsistency suggests that they may have answered randomly, we decided to remove such answers. Yet we kept their mental model if they submitted one (and potentially removed it during the coding process as explained in the next sub-section). As a result, we removed the answers of 32 respondents.

### 3.3.5 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 is available in Appendix A.2 (i.e., all drawings<sup>4</sup>) and two codebooks (i.e., technical codebook and contextual codebook) are available in Appendix A.3 and A.4 respectively.

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 [267, 263, 261, 268, 46], 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 [269], 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 [270] to organize and sort the ideas and thoughts raised in the answers. One of the researchers working on that project proceeded to the coding of open-ended questions, then the main investigator reviewed and provided feedback. The codebooks for three open-ended questions are available in Appendix A.5, A.6, and A.7 respectively.

---

<sup>4</sup>We removed all sensitive content (e.g., location information from the phone).




Figure 3.5: Data collection privacy concern (IUIPC) w/ fit.

### 3.3.6 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 [271]. 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%.<sup>5</sup> 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  $\mu$  the mean score,  $\sigma$  the standard deviation, and  $a$  and  $b$  the bounds. Figure 3.5 shows the score distributions.

<sup>5</sup>The age information for three respondents was not available in Prolific’s statistics.






Figure 3.6: 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.

## 3.4 Results

In this section, we present the results and findings from our survey, according to the ordering of the questions as presented in Section 3.3.2.

### 3.4.1 Users tend 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 3.6 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. [56]’s statement that, due to the large number and availability of TPAs, users can




Figure 3.7: Ratio of respondents who (still) actively use all, most, some, or none of their TPAs.

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 3.7, 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 Chapter 1, 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 [52]. 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 3.1 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. [Man, 30-39 y.o., Apple]: *“I just never think about it and do not think it is an issue to leave*

| 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 3.1: Main reasons respondents do not revoke access to their data to the TPAs that they no longer use actively.

*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. [Woman, 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. [Man, 18-29 y.o., Apple]: *“I find it troublesome to revoke their access.”* This is confirmed by the results in Figure 3.8 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. [Woman, 30-39 y.o., Garmin]: *“I was nervous about the data they were accessing.”*




Figure 3.8: Evaluation of the complexity / difficulty of the TPA fitness-data sharing monitoring process.

### 3.4.2 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 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 3.9 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 A.1 of Appendix A.8. 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,




Figure 3.9: 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.

whereas all profile information are by default set to “Only me” (i.e., private) for Garmin, Fitbit set the privacy level of most types of data 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 3.10, 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.

### 3.4.3 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 3.11 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




Figure 3.10: Expected vs. real proportion of public availability of specific data types (Fitbit and Garmin users).




Figure 3.11: Number of individuals in each relationship group with whom our respondents share their fitness data.

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




Figure 3.12: Self-declared likelihood to use the three different proposed PETs.

### 3.4.4 Users are inclined to use PETs.

We looked at the (self-reported) likelihood that respondents would use the different PETs we proposed. The results, depicted in Figure 3.12, 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. [46] 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. [Woman, 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 [Man, 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. [Man, 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%). [Man, 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. [272] 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. [Man, 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.

### 3.4.5 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.

#### 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.

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

| 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 3.2: 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 3.3: Mental model results.

Garmin [44, 273, 274]. We also contacted the Garmin support team to confirm our findings related to their devices. Table 3.2 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.,  $mm_1$ ).<sup>6</sup> 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.

Table 3.3 summarizes the findings. The first type (i.e.,  $mm_1$ ) was the

<sup>6</sup>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.

most frequently seen mental model where 36.4% of respondents drew it (e.g., Figures A.2a and A.2b in Appendix A.9). 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 A.2c and A.2d).

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 A.3a and A.2b). 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 A.3c). 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 A.3d and A.4a).

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 an 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 compromising their privacy. For example, while  $mm_3$  (offline sharing, i.e., on the phone) suggests that if the user deletes the TPA’s mobile app from their phone sharing is no longer possible, it is not necessarily the case for  $mm_1$  (online sharing, e.g., from the service provider’s servers to the TPA’s servers). 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).

### 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 A.4b and A.4c). The respondents (19.2%) also drew that their data is stored on the third-party servers (see Figure A.2b) and might be further analyzed (see Figure A.4d). A few respondents particularly mentioned ‘information analysis’ (6.8%) and wrote about ‘user profiling’ (5.5%) (see Figure A.5a). Finally, a few participants (8.2%) reported being concerned on whether TPAs can keep their data safe and secure (see Figure A.5b). In conclusion, 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.

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 A.5c).

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 A.5d). 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 A.3b). 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

| Truth | Ans. |               | Truth | Ans. |                 |
|-------|------|---------------|-------|------|-----------------|
| True  | 97   | Steps         | True  | 73   | Location        |
| True  | 88   | Weight/height | True  | 73   | Sleep data      |
| True  | 85   | Activities    | N/A   | 51   | Stress          |
| True  | 83   | Gender        | N/A   | 49   | Username        |
| False | 80   | Password      | N/A   | 49   | Menstrual cycle |
| True  | 76   | Birth date    | False | 45   | E-mail          |

Figure 3.13: 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.

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 A.5d), which is not necessarily the case.

### 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 3.13 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

| Truth | Ans. |   |
|-------|------|---|
| True  | 33   | The TPA is able to access the fitness data that was collected before I granted access.  |
| True  | 98   | The TPA is able to access the fitness data that was collected after I granted access.   |
| True  | 93   | The TPA is able to store on their own servers any data they have access to.   |
| True  | 85   | The TPA app is legally allowed - according to the federal laws in force in the United States - to store any data they have access to on their own servers.    |
| True  | 91   | The TPA app is legally allowed - according to your companion app's terms of service - to store any data they (the TPA) have access to on their own server.    |
| False | 82   | The TPA will be able to access the data collected after revoking, using the previously granted authorization.   |
| True  | 84   | The TPA will be able to access the data collected before revoking, if they stored it on their own servers.  |
| False | 38   | The TPA will still be able to access the data collected before revoking, using the previously granted authorization (without storing it on their own server). |

Figure 3.14: Proportion of correct answers regarding the (legal and technical) feasibility of data access by TPAs.

cannot be shared with the TPAs, whereas in fact, they can.

Figure 3.14 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.

## 3.5 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 [56], 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 [153]. 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 and 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 [259]). 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 mobile

apps and TPAs. Finally, a few respondents mentioned interesting solutions about time-framed sharing for enabling 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 [46]. The respondents in Velykoivanenko et al. [46]’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.<sup>7</sup> Our findings show that WAT users have poor knowledge about the data-sharing process. The implementation of transparency-enhancing technologies (TETs) [275] 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) [151]. 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) [276]. However,

---

<sup>7</sup>Note also that, in Velykoivanenko et al.’s study [46], they did not consider TPAs.



the companion app’s service provider does not provide substantial technical or legal support. For example, about data sharing with TPAs, Garmin’s 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 suggests 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.

## 3.6 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 [46]. This could be due to the self-selection bias problem [277] 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. How-

ever, “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.

### 3.7 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.

In Chapter 5, we will design, with a participatory approach, and evaluate such PETs, including—but not limited to—time-framed sharing, automated access revocation, and access-monitoring reminders.



# Chapter 4

## Watch your Watch: Inferring Personality Traits from Wearable Activity Trackers

**Abstract.** One particularly sensitive type of information has recently attracted substantial attention, namely personality, as it provides a means to influence individuals (e.g., voters in the Cambridge Analytica scandal). This chapter presents the first empirical study to show a significant correlation between WAT data and personality traits (Big Five). We conduct an experiment with 200+ participants. The ground truth was established by using the NEO-PI-3 questionnaire. The participants' step count, heart rate, battery level, activities, sleep time, *etc.* were collected for four months. By following a *principled* machine-learning approach, the participants' personality privacy was quantified. Our results demonstrate that WATs data brings valuable information to infer the openness, extraversion, and neuroticism personality traits. We further study the importance of the different features (i.e., data types) and found that step counts play a key role in the inference of extraversion and neuroticism, while openness is more related to heart rate.

### 4.1 Introduction

One particularly valuable type of personal information, as illustrated by the Cambridge Analytica scandal [278], is personality. Personality is often charac-

terized by the Big Five OCEAN traits (openness, conscientiousness, extraversion, agreeableness, neuroticism) [279], and it is known to influence behavior. Information about an individual’s personality enables others to manipulate this individual more efficiently by sending them appropriate signals (e.g., targeted advertisements), thus raising serious ethical concerns. For instance, Cambridge Analytica used data from social networks to infer the personality traits of US voters and to influence them during the 2016 Presidential Election [280, 281]. Similarly, credit card companies have exploited clients’ purchase history to profile debtors and craft the appropriate strategies to recover their debts [282] (e.g., by determining whether a specific client would respond better to a comforting or threatening message). As a result, individuals are increasingly worried about the potential misuses of automatic personality assessment [114]. Besides social networks, prior work has demonstrated that personality could be inferred from the data collected by individuals’ (smart)phones [283, 284, 285, 286, 287].

In this work, we focus on the problem of personality inference in the context of WATs. As such devices collect a large amount of behavioral and physiological data, they bring valuable information to infer personality. Indeed, behavioral indicators are one of the three types of indicators that are used to assess an individual’s personality [288]. Furthermore, previous research extensively studied the relationship between personality and physical activity [289] and identified multiple correlations between the two. Moreover, recent works show that WAT data can be used to infer characteristics related to personality, such as stress resilience [70] and mood [89]. It has also been shown that some personality traits are correlated to sleep [290]. Finally, WAT data can also be combined with other types of data that are already known to be helpful for personality inference (e.g., social network behavior, smartphone usage). Data brokers can indeed easily link different types of users’ data from different databases [291] and build accurate inference models using such a larger and more diverse data set. To the best of our knowledge, this is the first work to address the concrete (privacy) risks of personality inference from data collected by WATs.

## Contributions and Results

We present the first study on the inference of personality traits from data collected by WATs. We equipped 200+ volunteers with Fitbit wearable devices (namely, Fitbit’s Inspire HR WAT) and captured their step counts, heart rate,

battery level, activity, and sleep time over the course of a four-month period, as well as data available on their user profile, such as gender. To determine the personality profile of our participants, we used the Big Five personality scores captured through the standardized NEO-PI-3 questionnaire [62]. Our longitudinal data collection enabled us to precisely evaluate to what extent data collected by wearable devices are correlated to personality traits, and thus may be used together with other types of data, to conduct personality inference attacks.

In particular, we rely on a machine learning model trained on the data collected by the wearable devices to predict the given personality trait tercile. Although our model does not reach high levels of accuracy for any Big Five personality trait, it is evaluated using a rigorous leave-one-out (LOO) cross-validation, and we show that it can classify WAT users according to openness, extraversion, and neuroticism with statistical significance compared to the random-guess baseline. We also report on the most relevant features by analyzing those that are the most used by the inference model.

Furthermore, we collected our participants' concerns and perceptions regarding personality inference from their WAT data in an exit questionnaire. Nearly half of our participants thought that such inference would not be possible at all, while nearly two-third of them reported that they would be worried if it was. This is in line with a recent qualitative study, using interviews, that shows that a substantial fraction of users are worried about personality assessment and about its potential misuse [114]. Finally, we analyzed related prior work based on phone and smartphone data, discuss their methodologies, and compare our results and methodology to theirs. We show that the accuracy level achieved by our model outperforms that of the current state-of-the-art found in literature about (smart)phone-based inference using similar methodologies (ternary classification) [286] for all five personality traits. Furthermore, we are the first to show, with a rigorous evaluation process, correlations between wearable data and neuroticism and openness. Based on our analysis, we also discuss the design of potential privacy-preserving solutions.

In summary, in this work, we provide answers to the following questions: **(RQ 1)** To what extent do the data collected from WATs help infer their users' personality traits? **(RQ 2)** For each personality trait, which types of WAT data bring more information for the inference? **(RQ 3)** How does aggregating data or removing access to multiple types of collected data affect the inference accuracy?

## Outline

The rest of the chapter is organized as follows. In Section 4.2, we introduce the relevant background on personality measurement and on fitness tracking. Section 4.3 we further discuss our adversarial model in the specific angle of personality inference. We present the previous work related to personality inference in Section 4.7. In Section 4.4, we describe our data-collection campaign, provide general statistics about the collected data, and report the privacy concerns and opinions of the study participants. We introduce our personality inference attack framework in Section 4.5 and present the results in Section 4.6. We further discuss our results in Section 4.8 before concluding this chapter in Section 4.10.

## 4.2 Background

In this section, provide the necessary background regarding one key aspects of this chapter: personality assessment. The assessment of an individual’s personality is generally based on the Big Five personality traits, also known as the five-factor model. The Big Five personality traits constitute a psychological model that defines an individual’s personality through five main traits (specifically openness, conscientiousness, extraversion, agreeableness, and neuroticism; conveniently abbreviated OCEAN) that are subdivided into six subtraits each [62]. This model, which has been proven to be robust and stable over time [292], is structured as follows [279]:

- **Openness** to experience — Individuals who score high on this dimension tend to be intellectual, imaginative, sensitive, and open-minded. Those who score low tend to be down-to-earth, insensitive, and conventional.
- **Conscientiousness** — Individuals who are high in conscientiousness tend to be careful, thorough, responsible, organized, and scrupulous. Those low on this dimension tend to be irresponsible, disorganized, and unscrupulous.
- **Extraversion** — Individuals who score high on extraversion tend to be sociable, talkative, assertive, and active. Whereas, those who score low tend to be retiring, reserved, and cautious.
- **Agreeableness** — Individuals who score high on agreeableness tend to be good-natured, compliant, modest, gentle, and cooperative. Individuals who score low on this dimension tend to be irritable, ruthless, suspicious, and inflexible.
- **Neuroticism** — Individuals high on neuroticism tend to be anxious, de-

pressed, angry, and insecure. Those low on neuroticism tend to be calm, poised, and emotionally stable. Neuroticism is sometimes referred as emotional stability; it represents the exact same facet of personality, excepting that the score is reversed.

The NEO-PI-3 (third version of the NEO-PI) is a standardized questionnaire for assessing an individual's personality, along the five aforementioned traits. It is considered to be a reference in the personality assessment research field [62]. The NEO-PI-3 is a 240-item questionnaire describing and analyzing the five main aforementioned personality traits. This questionnaire delivers, for each of the five personality traits, a score between 0 and 192. The Big Five personality traits and the NEO-PI questionnaires are deeply related and have been developed mainly by Costa and McCrae [293]. Official translations of this questionnaire exist in many languages. In this work, we used the official translation of the full NEO-PI-3 questionnaire, in French, the local language at our institution.

### 4.3 Adversarial Model

As described in the introduction of this manuscript (Chapter 1), we mainly focus on an adversary that can access data processed by the service provider (i.e., Fitbit). Accessing the raw data would require either setting up a very specific environment with very specific devices (or even custom-made ones) and/or being in physical proximity to the devices. As we wanted to study a largely scalable attack with of-the-market devices, we decided to focus on already processed data that is easily available using a web API. One such adversary is typically the service provider itself. In this case, the risks we can measure represent a lower bound of the actual risks as the service provider has access to the raw WAT data and the smartphone data collected by the companion mobile app. Such an adversary also typically corresponds to any entity that manages third-party apps (TPAs). In Chapter 3, we show that 70% of WAT users share their data with at least one TPA and that users who share their data with TPAs tend to forget that they do. Furthermore, it also shows the users' lack of knowledge about the data-sharing process and demonstrates that they are not aware of the actual amount of data they share. Moreover, 9% of the participants in this study claimed to grant Fitbit access to at least one of their social media accounts, so that Fitbit can automatically make posts on their social media profiles related to their activity (e.g., step



counts). An adversary could use such information, alone or combined with other information available on the social profiles [294, 295, 296], to infer users' personality. Also, an employer could offer free WATs to their employees if they agree to share the collected data with their employer. Over the past few years, US companies have engaged in such corporate-wellness programs using Fitbit devices [26]. A government could gain access to a WAT service provider's data, for national security reasons, as recently suggested by a former US president [30]. An insurance company (e.g., health) could provide tracking devices to their policyholders to better analyze risks. For instance, Google acquired Fitbit [27] and Alphabet, Google's parent company, is growing rapidly in the health insurance market [28], furthermore, they plan to force Fitbit users to migrate to their Google accounts [29]. Finally, other adversaries could obtain such information by accessing tracking-device companies' leaked databases or by using eavesdropping techniques, as WATs and their related mobile applications are known to use poorly protected wireless communication protocols and data storage [194, 195, 198, 201, 208, 209].

In this work, we consider one such adversary who subsequently uses the collected data to infer the psychological profiles of the associated users. As explained in the introduction of this chapter, such personal information is highly sensitive, from a security & privacy point of view. This information is highly valuable for adversaries, thus pushing them to conduct such attacks. In particular, psychological profiling enables discrimination and manipulation. Indeed, assessing an individual's personality can help influence their behavior. For instance, it can be used to influence consumers' choices through targeted advertisements [282, 297] and even voters' choices [298] as in the Cambridge Analytica scandal related to the 2016 US presidential election [278], and thus have an impact beyond manipulating individuals, by influencing global politics.

## 4.4 Data Collection and Statistics

We describe our data collection campaign and we report on the general statistics regarding our participant pool.

### 4.4.1 Data-Collection Campaign

Evaluating the privacy of WAT users, with respect to their personality, requires having access to both WAT and personality data for a number of indi-

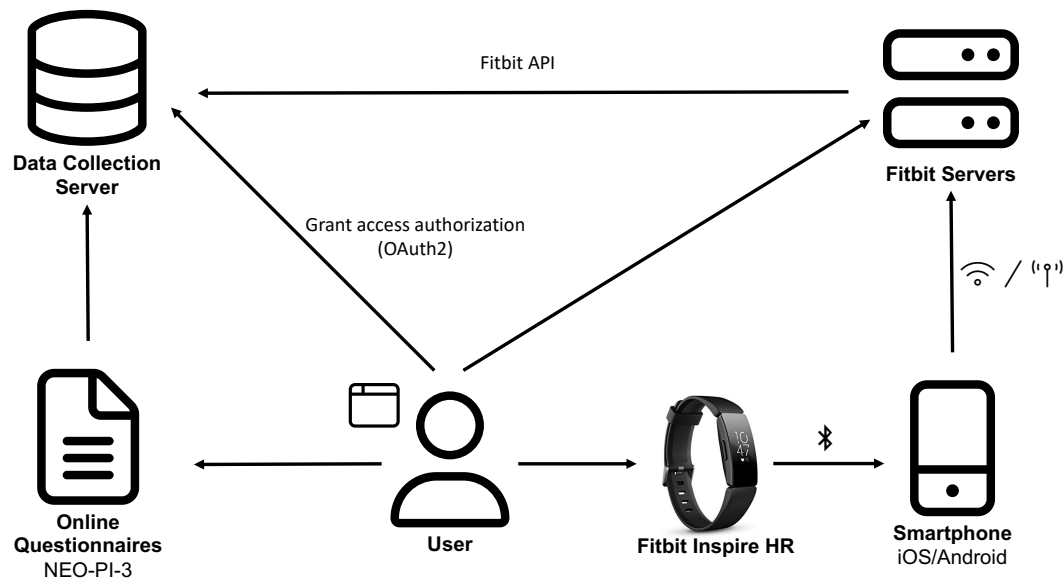


Figure 4.1: Data-collection architecture. Users wear the device that transmits their data to the Fitbit servers via their smartphone. At the beginning of the experiment, they grant us authorization to access their data, from the Fitbit server, using the Fitbit API. For this purpose, the protocol OAuth2 is used. That protocol allows users to select the different types of data that they agree to share with a given entity, then this entity receives a pair of tokens that they can use to request the user’s data through a provided API. Users also answer the online NEO-PI-3 personality questionnaires through which the ground truth is established.

viduals. In order to collect such data, we organized a large-scale experiment. We recruited the participants through LABEX, a dedicated structure of the University of Lausanne (UNIL); it manages a pool of around 8’000 students from two universities (a technical one, i.e., EPFL, and a general one, i.e., UNIL itself, that covers a broad range of disciplines). Those who were interested in our experiment responded to a screener questionnaire that we used to verify their eligibility for participating. 981 individuals answered the screener questionnaire and 429 were compatible with the experiment criteria: to own a smartphone compatible with the Fitbit application, to speak French (i.e., the local language at the universities; the questionnaires were in the local language), and to not already own a WAT. Given the number of devices at our disposal and the withdrawals during the participants selection phase, we

finally recruited 230 individuals.<sup>1</sup>

In order to ensure a better diversity of personality profiles, we selected the participants from different academic institutions and various study disciplines (see Figure 4.3). Every selected participant received a Fitbit Inspire HR bracelet. We chose to use a Fitbit device because Fitbit is one of the leaders in the WAT market [24] and because it provides a well-documented and effective API [299] to collect users' data. Moreover, the Fitbit Inspire HR is a high-end general-purpose WAT; as such, it gave us access to a wide range of data types (including step count, activities, sleep time, and heart rate) while still being used by a large user base. Using Fitbit trackers introduced some minor limitations such as the limited accuracy of some of their sensors (compared to higher-end devices) [300] as well as limited access to the data that they collect (only processed data, unlike specialized devices).

We only recruited new users because we wanted to provide them all the same WAT model, for data homogeneity and data collection infrastructure (Fitbit API). Furthermore, recruiting individuals who already owned a WAT could have increased the dropout rate as they would have been tempted to switch back to their own devices during the data collection.

Once they received their devices, the participants had to install the Fitbit application on their smartphones. They were instructed to wear the bracelet daily and all day long (they were free to remove it for comfort reasons, for example, at night) and to regularly synchronize with the Fitbit app running on their smartphones. They also had to answer a questionnaire that consisted of demographic questions and the NEO-PI-3 standardized personality assessment items [62] (see Section 4.2), which were used to compute their Big Five scores.<sup>2</sup> We chose that specific questionnaire because it is a reference questionnaire and because it provides results with high confidence and fine granularity. The purpose of this questionnaire was to collect the necessary ground truth. As most of the individuals in the LABEX participants pool are native French speakers, and as the French version of the test was available at the psychology department of our university, we decided to use the test in French for this study. The participants also had to answer an exit questionnaire, at the end of the experiment, which consisted of questions about how they used the application and device, what their privacy concerns were, and what they understood about

---

<sup>1</sup>Part of the participants agreed to share their WAT data. The dataset is available at <https://dx.doi.org/10.5281/zenodo.7621224>

<sup>2</sup>The questionnaire is available on <https://www.parinc.com/Products/Pkey/275>, unfortunately, we cannot directly share it due to copyright issues.

their data processing and storage. Except for their answer about their privacy concerns related to personality inference presented in Section 4.4.3, the data collected through this exit questionnaire are not directly used in this work.

The WAT data were collected for four months (between May and September 2020)<sup>3</sup> using the Fitbit API (the participants had to grant us access authorization by using the OAuth2 protocol).<sup>4</sup> We collected the step count for every one-minute interval; the average heart rate for every one-minute interval; the sleep related data such as the bedtime, wake-up time, sleep quality or the number of times that a user was restless during their sleep (for those who wore the device at night); as well as the sports activities (e.g., running, biking) that were automatically detected by Fitbit. Finally, in order to ensure high data-utility of our dataset, we decided to only keep the 204 individuals who wore their devices at least 50% of the time. Figure 4.1 depicts the global hardware and software architecture of our data-collection campaign.

### Ethical Considerations

During the distribution of the devices, the participants had to sign a consent form that described the conditions of participation, the data being collected (and the associated data management plan), the procedure to withdraw from the study, and information about the financial incentive. The institutional review board at our university validated the consent form and approved the entire experiment. As a reward, participants were paid 60 CHF ( $\sim$  60 USD) at the end of the data collection campaign, and they were allowed keep their device for personal use, which they all did.

### 4.4.2 Descriptive Statistics

Among the 204 selected participants, 64.7% were women, 34.8% were men, and 0.5% (1 participant) preferred not to indicate their gender. The women/men ratio is representative of the Fitbit user base. Indeed, we can observe that, in the general population, two-thirds of Fitbit users are women [301]. 72% of our participants are students from the general university (where a majority of students are women), and 28% are from the technical university. They are on average 22.6 years old with a standard deviation of 2.7 years. The

---

<sup>3</sup>The data collection campaign was conducted during the COVID-19 pandemic. However, there was no lockdown or restriction from May to September in Switzerland; only large events were canceled.

<sup>4</sup>Our access was revoked shortly after the end of the experiment.




Figure 4.2: Cumulative distribution function of the personality score for each of the five main traits (each score is between 2 and 192). The solid lines are the terciles, we display them because we proceed to a ternary classification (see Section 4.5). The gray line (in the middle) is the median.

Table 4.1: Distribution of the number of samples for each tercile and each personality trait.

|        | <b>O</b> | <b>C</b> | <b>E</b>   | <b>A</b> | <b>N</b> |
|--------|----------|----------|------------|----------|----------|
| Low    | 71 (35%) | 68 (33%) | 72 (35%)   | 69 (34%) | 69 (34%) |
| Medium | 70 (34%) | 71 (35%) | 64 (31.5%) | 69 (34%) | 67 (33%) |
| High   | 63 (31%) | 65 (32%) | 68 (33.5%) | 66 (32%) | 68 (33%) |

youngest is 18 years old and the oldest is 33 years old. Note that even if the age range is not representative of the general population, as the Big Five model is known to be stable over time [292], this should not substantially influence our results. Regarding the national statistics in our country, the age distribution corresponds to the student population. However, the proportion of women is slightly higher in our dataset than in the global student population. The scores for all personality traits correspond to a normal distribution. The medians of the scores for the five different personality traits have values between 96.5 and 125 points, depending on the trait. As shown in Figure 4.2, the scores for all personality traits are bell-shape distributed. This figure also shows the different terciles for each personality trait. Terciles are relevant as we focus on ternary classification in our experiments, as explained in detail in Section 4.5.1. With terciles of 84 and 109 points, neuroticism has the highest score variability, which helps us to better infer that personality trait (this is confirmed by our results; see Section 4.6), as the difference between individuals appears to be substantial. Table 4.1 shows the distribution of participants across each tercile of each personality trait. We can observe that the distribution is globally well balanced with no majority class containing more than 35% of the samples. Because participants can have the exact same scores, the terciles are not always of size exactly 33%. The participants wore their devices during 88% of the data collection period on average. The individual with the




Figure 4.3: Distribution of the study fields among the participants.

lowest wearing percentage wore it 50% of the time, and the individual with the highest percentage wore it 99% of the time. They have an average heart rate of 75 bpm (beats per minute) with a standard deviation of 7 bpm. The highest average heart rate is 94 bpm, and the lowest one is 59 bpm. During the data collection period, the participants took 1,066,263 steps on average, with a standard deviation of 337,049. The highest number of steps taken is 2,637,922 and the lowest one is 360,133. During the data collection period, the participants took 8,669 steps per day on average with a standard deviation of 2,740; a minimum of 2,928 steps, and a maximum of 21,447 steps. They slept for 8 hours and 17 minutes per day on average with a standard deviation of 2 hours and 4 minutes. Physical activities are automatically detected and recorded by the device, however, it only takes into account activities lasting 15 minutes or more. As shown by Figure 4.4, walking was, by far, the most practiced activity (63% of the activities). This is explained by the fact that, if other physical activities are generally considered for sports, walking can also be included by simply moving from one place to another. As the participants were free to sometimes remove their bracelets, they probably took steps, slept, or did activities that were not taken into account by the device, therefore, the previously discussed statistics about Fitbit collected data could be slightly underestimated. Details about the distribution of study fields and activities are shown in Figure 4.3 and 4.4.

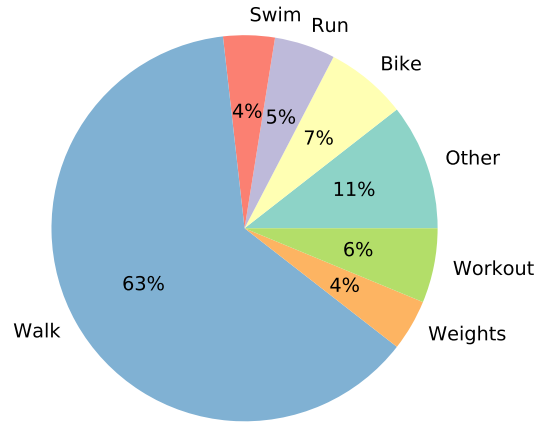


Figure 4.4: Breakdown (by types) of the activity practiced by the participants during the data collection campaign.

### 4.4.3 Participants' Privacy Concerns

In the exit questionnaire, we asked the participants to evaluate on a 5-point Likert scale: (1) *To what extent (that is, with what precision) can personality be inferred based on the data collected from your Fitbit tracker?* (from “Not at all precise” to “Extremely precise”) and (2) *To what extent would you be worried if the user’s personality could be inferred accurately based on the data collected by your Fitbit tracker?* (from “Not at all worried” to “Extremely worried”). For the first question, 47% of the participants answered “Not at all precise” or “Slightly precise”, 34% answered “Moderately precise” and 19% answered “Very precise” or “Extremely precise”. For the second question, 38% of the participants answered “Not at all worried” or “Slightly worried”, 26% answered “Moderately worried” and 36% answered “Very worried” or “Extremely worried”. Our participants also ranked personality as one of the most concerning types of information in a proposed list<sup>5</sup> (the top-3 was: political views, personality, and socio-economic status), and they were more concerned with personality being inferred than religion or sexual orientation (both have less than 30% of participants who either are “Very worried” or “Extremely

<sup>5</sup>age, alcohol, and tobacco consumption, illegal drugs consumption, menstrual cycles, political views, religion, sexual activity, sexual orientation, socio-economic status

worried”).

## 4.5 Inference

Privacy is commonly characterized as the (in)accuracy of an inference process [302], conducted by an adversary, that takes user data as input (data collected from WATs in our case) and outputs (a probability distribution across possible) values for some private attributes of the users (scores for the OCEAN personality traits in our case). In order for the privacy quantification to be fair and unbiased, it is paramount to properly design the inference framework and methodology, as shown by Mønsted et al. [286].

In this section, we describe the machine-learning-based inference methodology, the data extracted from the WATs for the inference (i.e., the features), and we report on our empirical results regarding the quantification of the privacy of WAT users, with respect to their personality.

### 4.5.1 Methodology

We define an inference framework which consists in training and testing a machine-learning (ML) model for predicting the scores for each of the OCEAN personality traits, for a given user and the WAT data associated to them. Based on the participants’ “actual” scores, computed from their responses to the NEO-PI-3 questionnaire [62] by following a standardized methodology, we establish the ground truth for the personality traits. We use this ground truth to train the ML model, in a supervised manner, and to evaluate its performance in terms of accuracy.

#### Inference Method

We chose to rely on classification methods because (1) the category within a general population to which an individual belongs to is the most important aspect from a psychological point of view [62] (as explained in Section 4.2) and (2) it is the most common method used in prior work [284, 285, 286]. Classes can be defined based on quantiles in order to get evenly sized groups (in terms of their number of individuals). For example, in the case of two classes (i.e., binary classification), the first class is defined as the individuals whose score is lower than the median and the second class as those whose score is higher than the median. In the case of three classes (i.e., ternary



classification), the class boundaries correspond to terciles. A common problem of using the aforementioned technique with an even number of classes is that, for bell-shaped distributions of scores, it splits participants in classes in the middle of the bell, where most of the participants lie. To minimize this issue, we defined the inference attack as a ternary classification process, similarly to previous works [285, 286]. Therefore, the classification problem consists in inferring, for each individual and each personality trait, if they belong to the bottom, middle, or top personality score class (regarding the score terciles), with respect to the whole dataset. Similarly to the related work [284, 285, 286], we directly computed the terciles on the participant dataset. We considered computing the terciles according to official statistics (national, for example), but on the one hand, such information is not necessarily easily available (or even exists), and on the other hand, this would not have solved the problem of choosing the population for which these terciles should be computed (e.g., Swiss students, Swiss citizens, European citizens). By calculating the terciles on the dataset, we at least guarantee that they correspond to the population directly studied in this work. Furthermore, personality is not usually measured in absolute terms by psychologists [303], but relative to a given population in space and time (terciles change with time and culture).

## Evaluation

We evaluated privacy for each of the five main personality traits (OCEAN) independently. For each trait, we defined three classes from the whole dataset as explained above, and we conducted the inference and the evaluation. In order to train and evaluate the model, we proceeded to a nested Leave-One-Out (LOO) cross-validation. More specifically, for a dataset  $S = \{x_i | i \in [1..N]\}$ , where  $x_i$  denotes the data of participant  $i$ , the model was trained and evaluated  $N$  times using  $S \setminus \{x_i\}$  as training set and  $\{x_i\}$  as testing set for each  $i \in [1..N]$ . Moreover, for each of the  $N$  iterations, the feature selection strategy and its hyper-parameters (i.e., number of selected features) as well as the hyper-parameters of the model were chosen using a grid search with LOO cross-validation on the  $N - 1$  elements of the training set.

By proceeding this way, we make sure that the results presented are fair in the sense that information leakage (e.g., when the feature selection is done on the entire dataset) is prevented. As pointed out by Mønsted et al. [286], sharing data between model selection and model evaluation steps leads to overestimating performance of the models at stake. In particular, they show

that some of the works related to ours [284, 285] are subject to such methodological biases. We use the accuracy (i.e., the proportion of correctly classified instances) as our evaluation metric. This metric is the most suitable for comparing different models, and it provides a clear understanding of their performance. Moreover, it is the only metric that is used in all prior work performing classification [284, 285, 286]. However, we are aware that accuracy is limited since, as it aggregates the confusion matrix into a single value, it does not distinguish between different types of errors and their associated magnitudes (e.g., misclassifying a participant as “bottom” instead of “top” is worse than misclassifying them as “middle”). Finally, we compare our results to the baseline defined by a uniformly-random naive classifier (the probability of inferring the correct class for each trait and each test individual is therefore 33%). Due to slight differences between the class sizes, we decided not to use the majority baseline. When the difference between two class sizes is zero or one, holding-out a single sample from the training set would result in the corresponding class being under-represented in the training set and the majority-class classifier would then underperform the random baseline.

The inner loop of this nested cross-validation performs both feature and model selection. The feature selection strategy is cross-validated among (1) univariate feature selection, (2) a greedy feature elimination strategy, and (3) a model-based feature importance approach. The models at stake in this inner cross-validation loop are Support Vector Machines (SVM) and Random Forests (RF). Cross-validated hyper-parameters for SVMs are the kernel (Gaussian and linear kernels are considered),  $C$  and  $\gamma$  (for Gaussian kernels), while for RFs, we have cross-validated the number of trees in the forest and the split criterion. For all traits, in all iterations of the inner loop, the selected model is an SVM. Note that, as it can be observed in Table 4.5, SVM is the most common ML method used in prior work for solving similar problems. For the implementation, we have relied on the `scikit-learn` [304] machine learning library for Python.

## 4.5.2 Feature Extraction

We collected different types of data through the Fitbit API: time series (steps, heart rate, battery level), events (sleep, activities) and standalone features (gender, resting heart rate). The extraction of most of our features consisted of aggregating time-series data over time intervals, with some periodicity using the following method: for each day of the week, we aggregated data according

Table 4.2: List of all features used in the evaluation. “Std.” stands for standard deviation. The “+” operation for data aggregation means that both aggregating methods were used to obtain the given feature. The dots in the last 5 columns indicate that the corresponding features of this data type were selected by the model for inferring the corresponding trait in our evaluation.

| Data Type           | Statistics                  | Aggregation Method               | O | C | E | A | N |
|---------------------|-----------------------------|----------------------------------|---|---|---|---|---|
| Step count          | Mean, Std.                  | Days of the week + 4-hour period | • | • | • | • | • |
| Step goals          | Nb. of occurrences          | The whole data collection period | • |   |   |   | • |
| Heart rate          | Mean, Std.                  | Days of the week + 4-hour period | • | • | • | • | • |
| Sleep time          | Mean, Std.                  | Days of the week + 4-hour period | • | • |   | • | • |
| Other sleep details | Mean, Std.                  | No aggregation                   |   |   |   | • | • |
| Activity time       | Mean, Std.                  | Days of the week + 4-hour period | • | • |   | • |   |
| Activity types      | Entropy, Nb., Proportion    | Activity type                    | • | • | • | • | • |
| Battery charging    | Entropy, Nb. of occurrences | Days of the week, 4-hour period  |   | • |   |   |   |
| Gender              | Category                    | N/A                              | • |   |   | • | • |

to predefined periods of the day. To this end, we partitioned the day into six periods of four hours with boundaries at: 2AM, 6AM, 10AM, 2PM, 6PM and 10PM. Previous studies highlighted that personality is correlated with individuals’ circadian rhythm (natural process that regulates a 24-hour biological cycle) [305, 306]. We thus defined  $6 \times 7 = 42$  different periods (e.g., “Monday between 10AM and 2 PM”) for aggregating the data into features. We then computed features corresponding to their two first statistical moments (i.e., the mean and the variance for the heart rate and step count taken across each of these periods).

Note that, although the extracted features refer to physiological and behavioral information, they are not as rich as those that can be collected from a (smart)phone [283, 307, 285, 286, 308]. They could also contain errors as, for example, the sensor signal analysis might sometimes not detect the right activity or confuse a step with certain arm gestures.

Furthermore, they are particularly centered on the user’s activities and, unlike phone data, contain no direct social information, even though multiple personality traits have a strong social component.

## Steps and Heart Rate

Steps and heart rate have the same data structure: they are sequences of pairs  $(t, x)$ , where  $t$  a timestamp, and  $x$  a scalar value. The sampling period is of one minute. We extracted features from the data of both types by using the periodic aggregation method explained above. As Fitbit “rewards”, on a daily basis, its users whose step counts exceed a certain so-called “daily

step goal” (set to 10,000 by default), we added the following three related features: the number of times this goal is achieved, the number of times it is just achieved (up to 5% *more* than the step goal), and the number of times it is almost achieved (up to 5% *less* than the step goal). Furthermore, the Fitbit API directly provides the resting heart rate for each user, which we used as such as a feature. As mentioned in Section 4.2, a relatively high score in extraversion is, for example, linked to sociable and active individuals whose traits could influence the step count. One of the extraversion sub-traits is excitement seeking, which can lead to an augmentation of an individual’s heart rate. Neuroticism is linked to impulsivity and stress, which can also cause variations in heart-rate. Moreover, it has been shown that heart-rate variability and an individual’s personality are correlated [309].

### Sleep and Activities

Sleep data are composed of a start time, a duration, and other information such as the sleep quality, the number of times the user wakes up during their sleep, and the number of times they are agitated. We built features of the same structure as steps and heart rate. We generated, the mean and standard deviation of sleep time, for each four-hour and day-of-the-week periods. We also computed the mean and standard deviation of the awake duration during sleep, the awaking count, the sleep duration, the time (in minutes) it takes to fall asleep, the restless-moment count and duration, and the sleep efficiency (all these details are directly provided by Fitbit). The data structure of the activities is similar to that of sleep data. We therefore built similar features. We computed the number and proportion of each practiced activity, as well as the entropy of the distribution of practiced activities. As mentioned previously, active individuals tend to obtain higher scores in extraversion. As for sleep, previous studies established that an individual’s sleep quality and habits are correlated with their personality [310, 290, 311].

### Battery

The “current” battery level of the device is available at any point in time through the profile endpoint of the Fitbit API. To eventually obtain a battery data time series for each participant, we collected this twice a day, at a fixed time. Note that the API returns the battery level at the time of the last synchronization (together with the time of the last synchronization). Then, we extracted the average battery level right before and after a charge, as well as

its standard deviation. We also computed how many times each participant charged their device and the entropy of the time elapsed between these events, for each day of the week. We also created similar features by using only the six previously defined periods of the day (without again aggregating with the days of the week). However, the Fitbit API provides only the battery level at the time of the last synchronization between the bracelet and the smartphone. Therefore, we might have lost information if users had not synchronized their data regularly (e.g., if Bluetooth was not continuously activated on their phone).

### Gender

As gender is known to be correlated with the score of some personality traits [312], and as such information is often available through the profile endpoint of the Fitbit API, we included gender data as a feature. All the participants specified a gender in their profiles. We observed a mismatch between the gender they specified in their Fitbit profiles and that specified in their responses to our questionnaire for only 0.98% ( $n = 2$ ) of the participants. Self-reported gender data can therefore be considered as a readily-available (to an adversary) and trustworthy data in the inference process.

## 4.6 Results

### Inference Accuracy

As shown in Figure 4.5, we obtained results that are statistically significantly better than the baseline<sup>6</sup> for openness ( $p < 0.01$ ), extraversion, and neuroticism ( $p < 0.001$ ). The trained model correctly classified 45% of the participants' scores in openness (+36% with respect to the baseline), 52% of the participants' scores in extraversion (+58% with respect to the baseline), and 50% of the participants' scores in neuroticism (+52% with respect to the baseline). We further observe that Fitbit data brings some valuable information for the inference of other traits, such as agreeableness and conscientiousness, but these results are not statistically significant. Regarding the definition of each personality trait, it is relatively intuitive that WAT data are less informative for a trait such as agreeableness than for neuroticism or extraversion.

---

<sup>6</sup>All statistical tests for model comparison were conducted using McNemar's test, with Bonferroni correction.

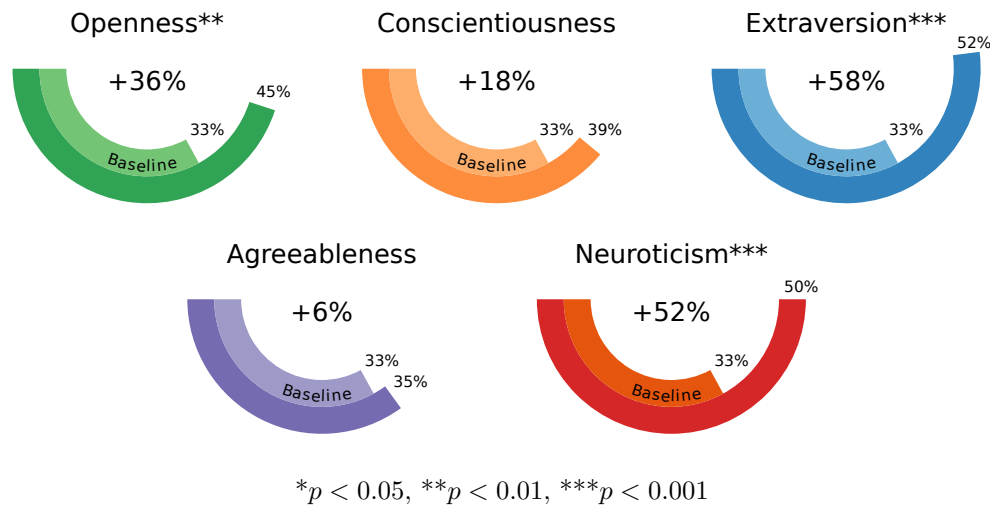


Figure 4.5: Accuracy of the ternary classification with respect to the baselines for each of the five main traits. For each trait, we display the increase of accuracy (in percentage) compared to the random baseline, the accuracy of the baselines and the accuracy of the prediction. Percentages are rounded to the unit. The accuracy of the prediction outperforms both baselines with statistical significance with Bonferroni correction (i.e., using an  $\alpha$  value of  $0.05/m$  with  $m$  the number of inferences, 5 in our case) for openness ( $p < 0.01$ ), extraversion, and neuroticism ( $p < 0.001$ ).

Table 4.3 provides more performance metrics, namely precision, recall and f1-scores for each tercile. For openness, extraversion, and neuroticism, the weighted mean of the f1-score (respectively 0.45, 0.51, and 0.50) is clearly higher than the baseline (0.33), which confirms the results presented above.

### Influential Features

In Table 4.2, we can see which general-data types were used to extract the relevant features for inferring each personality trait. For each inference, we looked at the three most informative features. We considered the features selected more times during the inner loop of our cross validation as more informative. For features used to infer openness, extraversion, and neuroticism, we conducted statistical tests (Kruskal-Wallis with Bonferroni correction) to reject the natural null hypothesis that the differences between terciles are incidental to the collected data. We show that we can reject the null hypothesis for all of these features with  $p < 0.05$ (\*),  $p < 0.01$ (\*\*),  $p < 0.001$ (\*\*\*) or  $p < 0.0001$ (\*\*\*\*). Figures 4.6, 4.7, and 4.8 show the distribution of the most informative features over the terciles for openness, extraversion, and neuroti-

Table 4.3: Precision, recall and f1-score for each class.

| <b>Openness</b>     | <b>Prec.</b> | <b>Rec.</b> | <b>f1-score</b> | <b>B. f1-score</b> |
|---------------------|--------------|-------------|-----------------|--------------------|
| Low                 | 0.47         | 0.39        | 0.43            | 0.34               |
| Medium              | 0.48         | 0.60        | 0.53            | 0.34               |
| High                | 0.39         | 0.35        | 0.37            | 0.32               |
| Weighted Mean       | 0.45         | 0.45        | 0.45            | 0.33               |
| <b>Conscien.</b>    | <b>Prec.</b> | <b>Rec.</b> | <b>f1-score</b> | <b>B. f1-score</b> |
| Low                 | 0.39         | 0.43        | 0.41            | 0.33               |
| Medium              | 0.33         | 0.31        | 0.32            | 0.34               |
| High                | 0.44         | 0.26        | 0.43            | 0.33               |
| Weighted Mean       | 0.39         | 0.39        | 0.39            | 0.33               |
| <b>Extraversion</b> | <b>Prec.</b> | <b>Rec.</b> | <b>f1-score</b> | <b>B. f1-score</b> |
| Low                 | 0.54         | 0.61        | 0.57            | 0.34               |
| Medium              | 0.44         | 0.31        | 0.37            | 0.32               |
| High                | 0.56         | 0.63        | 0.59            | 0.33               |
| Weighted Mean       | 0.51         | 0.52        | 0.51            | 0.33               |
| <b>Agreeab.</b>     | <b>Prec.</b> | <b>Rec.</b> | <b>f1-score</b> | <b>B. f1-score</b> |
| Low                 | 0.35         | 0.36        | 0.36            | 0.34               |
| Medium              | 0.39         | 0.41        | 0.40            | 0.34               |
| High                | 0.31         | 0.29        | 0.30            | 0.33               |
| Weighted Mean       | 0.35         | 0.35        | 0.35            | 0.33               |
| <b>Neuroticism</b>  | <b>Prec.</b> | <b>Rec.</b> | <b>f1-score</b> | <b>B. f1-score</b> |
| Low                 | 0.55         | 0.59        | 0.57            | 0.34               |
| Medium              | 0.41         | 0.42        | 0.41            | 0.33               |
| High                | 0.53         | 0.49        | 0.51            | 0.33               |
| Weighted Mean       | 0.50         | 0.50        | 0.50            | 0.33               |

cism. Even if our approach would likely lead to similar results with other types of population (i.e., than the one we collected the data from) and the performance would be comparable, we would expect the influential features to be quite different. For instance, going out late at night has a different meaning for students and for middle-aged adults with children. The three most informative features for each inference process are (when there are more than three features, all the presented features are considered as equally important by the model):

- **Openness\*\***
  - Step-goals ( $\geq 10k$  steps) just achieved.\*\*
  - Number yoga activities.\*
  - HR std from 2AM to 6AM on Thu.\*\*
  - HR std from 10AM to 2PM on Fri.\*\*
  - HR std from 2PM to 6PM on Thu.\*
- **Conscientiousness**
  - Std of HR btw Wed. and Thu. (10PM-2AM)
  - Sleep-time mean from 10AM to 2PM on Sun.
  - Sleep-time mean from 2AM to 6AM on Sat.
- **Extraversion\*\*\***

- Step mean btw Fri. and Sat. (10PM-2AM).\*\*\*\*
- Step mean on Mon. btw 6PM and 10PM.\*\*\*\*
- Step mean btw Thu. and Fri. (10PM-2AM).\*\*\*\*
- Number of distinct activities.\*\*\*
- HR mean btw Sun. and Mon. (10PM-2AM).\*\*\*\*
- **Agreeableness**
  - Steps std on Sun. btw 6PM and 10PM.
  - Sleep-time mean (global).
  - Std of HR on Thu. btw 10AM and 2PM.
- **Neuroticism\*\*\***
  - Gender.\*\*\*\*
  - Steps mean on Mon. btw 6PM and 10PM.\*\*
  - Sleep-time mean from 10AM to 2PM on Sun.\*


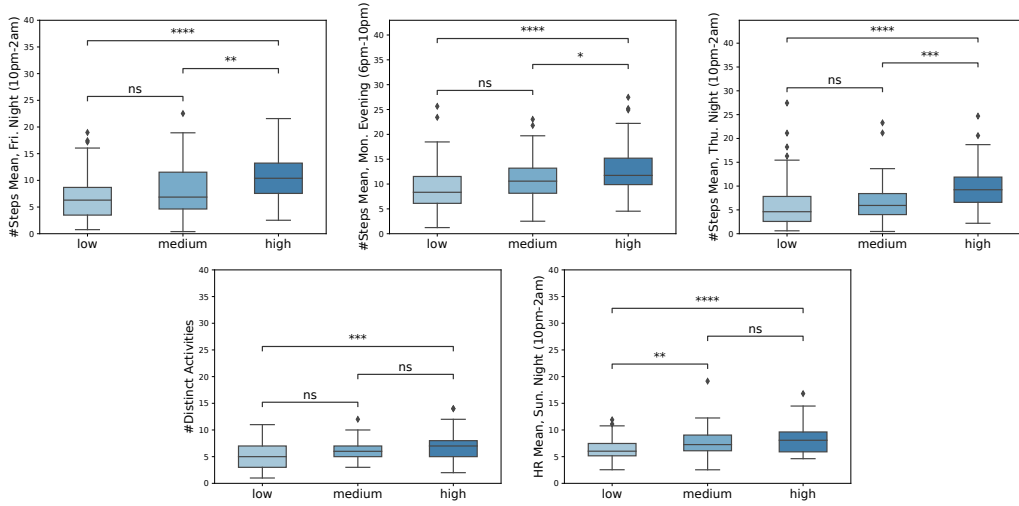


Figure 4.6: Distribution of four of the main features used for openness inference for each tertile. HR mean is weighted regarding the individual’s resting HR.

Interestingly, we can see that the practice of yoga is highly informative for the inference of openness. This is coherent as users with high openness tend to seek new experiences and to engage in self-examination and individuals who practice yoga are known to obtain higher score in openness [313]. However, we cannot make a general conclusion here with that information as only eight participants recorded yoga activities during the data collection. Among those participants, only one was not classified in the high openness tertile.





ns:  $p > 0.05$ , \* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$ , \*\*\*\* $p < 0.0001$

Figure 4.7: Distribution of the five main features used for extraversion inference for each tertile. Step count means are weighted regarding the bracelet wearing time, HR mean is weighted regarding the individual’s resting HR.

Table 4.4: The obtained inference accuracy using different combinations of data sources. The increase in accuracy is computed using the random baseline. The last line corresponds to aggregations by day (i.e., not 4-hours time slots) for heart rate and steps.

| Data source                       | O            | C          | E            | A         | N            |
|-----------------------------------|--------------|------------|--------------|-----------|--------------|
| All data sources                  | 45% (+36%)** | 39% (+18%) | 52% (+58%)** | 35% (+6%) | 50% (+52%)** |
| All data but gender               | 44% (+33%)*  | 39% (+18%) | 52% (+58%)** | 35% (+6%) | 47% (+42%)** |
| All data but heart rate           | 35% (+6%)    | 32% (-3%)  | 50% (+52%)** | 34% (+3%) | 50% (+52%)** |
| All data but heart rate and sleep | 34% (+3%)    | 33% (+0%)  | 50% (+52%)** | 33% (+0%) | 48% (+45%)** |
| Only step count                   | 34% (+3%)    | 32% (-3%)  | 47% (+42%)** | 34% (+3%) | 44% (+33%)*  |
| All data (aggregated) but gender  | 38% (+15%)   | 35% (+6%)  | 33% (+0%)    | 34% (+3%) | 34% (+3%)    |

\* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$

As shown in Figure 4.6, HR-related features are important for the inference of openness. Psychology studies have shown that features related to cardiac activity (including heart rate), are correlated with openness [314, 315].

This is confirmed by Table 4.4 which shows that without HR-related features, our model is not able to correctly classify individuals according to their openness level significantly higher than the baseline. Note that most of these HR-related features are relative to Thursday and Friday afternoons. One possible reason is that openness is related to art sensitivity and creativity and that these time slots are the most favorable for such activities (museums or art galleries, for example, are often closed at the beginning of the week). Thursday




Figure 4.8: Distribution of the three main features used for neuroticism inference for each tertile. Step count means are weighted regarding the bracelet wearing time. The sleep time is in hours. The area of each circle in the gender plot is proportional to the number of participants who corresponds to the given gender.

and Friday evenings/nights or Saturday, however, are time periods related to extravert-oriented activities (e.g, clubbing). We can also observe that steps goals are used to infer the score of openness, however, there is no previous research that can help us understand the reason of this correlation.

Looking at Table 4.2, we can first observe that, information related to steps, heart rate, and activities are used to infer extraversion. This can be explained by the fact that people with higher scores in extraversion tend to be more active, assertive, and sociable (see Section 4.2). Three of the most informative features relate to the average step count *at night*, thus showing that the level of (social) activity plays a key role in the inference of extraversion. This is confirmed in Figure 4.7 Indeed, the more extraverted a participant is, the more steps they take at night (especially at night between Thursday and Friday, on Monday evenings, and at night between Friday and Saturday). This could be explained by the fact that the more extraverted the individual, the more they go out at night (e.g., to meet friends, to go clubbing, etc.). That may also be supported by the mean heart-rate on Sunday night being higher for the most extraverted individuals. Furthermore, we observe that the most extraverted individuals tend to do more distinct activities, which corresponds to the activity and excitement seeking component of extraversion as described in Section 4.2. Moreover, to assess personality traits, standard tests combine behavioral, cognitive, and affective indicators [288], and behavioral indicators are the most informative to assess extraversion [316]. This explains why WAT data, which are almost exclusively related to behavior, are the most informative for this trait.

Steps, heart rate, and activities are also used to infer neuroticism. However,

we observe that HR-related features do not appear to be the most informative features for neuroticism. Instead, these features relate to gender, sleep, and steps. Previous works show that information such as step count, heart rate, or duration of sleep are indicators of stress resilience, which by definition is highly correlated with neuroticism [70]. Step count, heart rate, and duration of sleep have also been used in previous studies to predict depression [71]. We also observe that sleep and gender are used to infer neuroticism. Both are indeed known to be correlated with this personality trait [290, 312], and this is confirmed in Figure 4.8. We can observe that the mean for sleep hours on Sunday midday for both participants with a low and high neuroticism score is significantly higher than that for the participants with a medium score of neuroticism. In fact, these two groups differ mainly in gender, as men tend to have a low neuroticism score and women a high score. More specifically, there is a significant difference among the terciles regarding the sleep time (here on Sunday between 10am and 2pm). It also shows that there is a significant difference between genders regarding their neuroticism score. As gender is correlated with neuroticism, we trained and evaluated a simple decision tree to infer the neuroticism class from gender only with the same methodology as described before. Such a model reaches an accuracy score of 48%. Additionally, we also evaluated our model without using gender and showed that it reaches 47% of accuracy. Therefore, a model using WAT data is similar, in terms of accuracy, to a model based on gender for inferring neuroticism. However, considering that WAT users can easily lie about their gender without decreasing their utility, which is not the case with step count or sleep data, a model based on WAT data (possibly helped by gender), is therefore more reliable than a model based on gender only.

Note that the list of informative features for the conscientiousness and agreeableness traits should be considered with caution, because it corresponds to prediction tasks for which our models do not significantly outperform the baseline.

### Sensitivity Analysis

We evaluated the inference performance by using a subset of data sources. Indeed, when giving access to the API, WAT users can choose to restrict access to some information by selecting only some types of data or, simply, by choosing to not report personal information (i.e., gender). Furthermore, some devices can simply not collect certain data due to the lack of sensors (e.g., un-




Figure 4.9: Evolution of the performance of the inference with training data collected for the first 8, 12, and 16 weeks. As it does not evolve over time, gender is not used as a feature.


like the Fitbit Inspire HR, the Fitbit Inspire does not collect heart-rate data). Table 4.4 summarizes the results obtained by evaluating the inference model which uses different data source combinations. The accuracies of the extraversion and neuroticism inferences are still significantly higher than the baseline when using only step-count-related features. This demonstrates that even devices that do not collect the heart rate, such as the Fitbit Inspire bracelet, can be used to accurately infer the personality of their users. However, the results from Table 4.4 suggest that heart rate data is essential to infer openness as the inference accuracy significantly declines when we remove this data source from the features set.

### Performance Evolution over Time and Training Set Size

Additionally, we analyzed how the inference performance evolves with training data collected over an increasing period of time. As it does not evolve over time, we did not use gender as a feature. Figure 4.9 shows, for each trait, how the inference accuracy evolves using training data collected for 8, 12, and 16 weeks.

We can observe that only 8 weeks are necessary to obtain an accuracy significantly better than random for neuroticism while 16 weeks are required to significantly outperform the baseline for openness and extraversion. We can

also postulate that the inference performance would be better with a few more months of data (which would capture additional seasonal phenomena), especially for extraversion, that shows the highest growth with time. We observe that the inference of extraversion is highly dependent on the data collection duration. This is probably due to seasonal behavior change (e.g., people tend to go out more often during the summer), and due to the fact that the most important features are probably related to social events, and thus that more time is necessary to collect enough data related to these specific, and possibly short, events. However, the results tend to show that an augmentation of data collection duration would not highly impact the inference of conscientiousness and agreeableness. Note that we use the same set of participants for all inferences, which may introduce a bias due to the fact that we selected the ones who wore their devices at least 50% of the time during the whole four-month period. Results with fewer months could so be slightly underestimated considering that some participants may have been selected while they were not wearing the device much during that specific period.



\* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$

Figure 4.10: Evolution of the performance of the inference with training dataset size by evaluating the model with  $k$ -fold cross validation with  $k \in \{2, 3, 4, 5, 10\}$

Finally, we also evaluated our model using  $k$ -fold cross validation with  $k \in \{2, 3, 4, 5, 10\}$ , details are available in Figure 4.10, and show that, especially for openness, neuroticism, and extraversion, the inference accuracy tends to increase with the size of the training set. For all traits, the accuracy does

Table 4.5: Comparative table of the most relevant publications. The ‘year’ is the year of publication, the ‘source’ represents the data source used to build the features for the inference process, ‘N’ is the number of participants, ‘var.’ means that the data collection duration is not fixed among the different participants, ‘CDR’ stands for Call Detail Records, the inference type is either regression or classification,  $k$  is the number of classes in case of classification, ‘SVR’ stands for Support Vector Regression, ‘SVC’ for Support Vector Classification, ‘RF’ for Random Forest, and ‘LOO’ stands for Leave-One-Out evaluation. Finally, the ‘Results’ column shows, in bold, which traits were inferred statistically significantly better than their respective baseline.

| Article                   | Year | Source     | N   | Dur. | Inference          | Model | Eval.   | Results       |
|---------------------------|------|------------|-----|------|--------------------|-------|---------|---------------|
| de Oliveira et al. [283]  | 2011 | CDR        | 39  | var. | Regression         | SVR   | 10-fold | <b>OCEAN*</b> |
| Chittaranjan et al. [284] | 2011 | Smartphone | 83  | 8 m  | Class. ( $k = 2$ ) | SVC   | LOO     | <b>OCEAN*</b> |
| de Montjoye et al. [285]  | 2013 | CDR        | 69  | 16 m | Class. ( $k = 3$ ) | SVC   | 10-fold | <b>OCEAN*</b> |
| Mønsted et al. [286]      | 2018 | CDR        | 636 | 24 m | Class. ( $k = 3$ ) | SVC   | 10-fold | OCEAN         |
| Stachl et al. [287]       | 2020 | Smartphone | 624 | 30 d | Regression         | RF    | 10-fold | <b>OCEAN</b>  |
| → This chapter            | 2023 | WAT        | 204 | 4 m  | Class. ( $k = 3$ ) | SVC   | LOO     | <b>OCEAN</b>  |

\* Mønsted et al. [286] showed that these articles suffer from test-data leakage (i.e., when data from the test data is used for training, for instance, in the feature selection step), which leads to overfitting. Therefore, the performance reported in those articles is largely overestimated. For example, according to Mønsted et al. [286], if de Montjoye et al. had used a rigorous experimental setup, they would have only obtained statistically significant results for extraversion (leading to **OCEAN** instead of **OCEAN** in the table).

not plateau for larger training sets, which indicates that the accuracy would increase if the sample included data of more individuals.

## Obfuscation

Finally, we evaluated the inference performance using heart rate and step count data *aggregated by day* (instead of 4-hour intervals), mimicking the case where the adversary would only have access to the average daily heart rates and total daily step counts (other features such as sleep and activities are used in the same way as described previously). Indeed, previous research suggests that such aggregation may be used as an obfuscation technique to reduce privacy risks and shows high acceptance among WAT users [46]. Table 4.4 shows that aggregating heart rate and step count results in an important drop in accuracy and that none of the inferences are significantly better than the baseline in this case. Note that we also removed gender from the features to properly evaluate the impact of such an obfuscation technique on neuroticism.

## 4.7 Related Work

Recently, Meegahapola et al. highlighted the correlation between data from wearable for pets and dog’s personality using a custom-built device equipped with an accelerometer and a gyroscope and using as ground truth a Big5-like model specifically adapted for dogs.

More related to human beings, prior studies about mobile-phone-related data highlighted the link between collected personal data and personality traits. Table 4.5 compares all the related-work experimental layout and results that we discuss in detail next.

de Oliveira et al. studied to which extent it is possible to infer personality traits from call-detail records using regression. Their model obtained mean square errors (MSE) significantly ( $p < 0.05$ ) lower than the baseline (MSE of 1.184) for openness (MSE of 0.670), extraversion (MSE of 0.650), and agreeableness (MSE of 0.615) [283].

Chittaranjan et al. evaluated the accuracy of personality-trait inference from smartphone data by using binary classification methods [284, 307]. They obtained an average accuracy of 72% (+25% of accuracy compared to the baseline on average) for all traits.

de Montjoye et al. evaluated the accuracy of personality-trait inference from phone-based metrics by using ternary classification methods [285]. They obtained an average accuracy of 53% (+42% of accuracy compared to the baseline on average) for all traits.

However, Mønsted et al. show that the inference results were overestimated in the aforementioned articles [283, 284, 285]. More specifically, the authors of these works optimized some parameters (e.g., feature, model, and hyperparameter selection) based on the *entire* dataset instead of doing so based on only the training set considered in each iteration of the cross-validation loop; this corresponds to the common pitfalls P3 and P5 listed in Arp et al.’s recent work on the dos and don’ts of machine learning in computer security [318]. Mønsted et al. further proceed to a ternary classification of the five traits by using the same models, features, and approach as de Montjoye et al.’s article [285]. They show that, based on their correlation with the trait to infer without using cross-validation (i.e., on the entire dataset), previous research about inferring personality from phone data overestimated model performance by selecting certain features. After following the same approach and obtaining similar results to de Montjoye et al., Mønsted et al. show that by using a more rigorous methodology with the same data, only extraversion can be inferred

(with an accuracy significantly better than the baseline) from (smart)phone data. They obtained an accuracy improvement of +36% (wrt the baseline) for that specific trait. Therefore, we cannot compare our work with their results, except for those of Mønsted et al. who used a (rigorous) methodology similar to ours. Hence, we can assert that personality inference models using WAT data outperform those using CDR as they achieve a higher accuracy for extraversion as well as accuracies significantly higher than the baseline for neuroticism and openness.

More recently, Stachl et al. inferred personality traits from richer smartphone data [287] using smartphone data of 624 participants collected over 30 days. Their features were more diverse and richer than those used in the other studies. The features were derived from call detail records, music consumption, application usage, mobility, overall phone activities, and daily activities. They show that it is possible to infer openness, extraversion, and conscientiousness from these data.

In summary, we are the first to demonstrate that WAT data brings valuable information to classify users according to their personality traits. Moreover, regarding related work that used similar methodological approaches (ternary classification), we show WAT data is more helpful for such classification than phone data. Also, by using a rigorous evaluation methodology, and thus, in comparison with most of the previous works, fairly evaluating the inference performance, we are the first to show how users can be classified according to their neuroticism level with an accuracy significantly higher than the baseline. Finally, we show that WAT data are correlated to openness, which was not the case with the data considered in prior work (e.g., CDR).

## 4.8 Discussion

Our experimental results demonstrate that processed data from WATs bring valuable information about at least three of the Big Five personality traits. Indeed, WAT data correlates with at least three of the five personality traits, which is consistent with multiple previous findings showing that behavior indicators are particularly informative for some traits (especially for extraversion) [288, 316], that WAT data can help assess stress resilience [70], or that it can be used to infer someone's mood [89]. One can argue that, as openness and extraversion tend to be positive traits (or are perceived as such), inferring them does not represent a particularly serious threat. However, that would overlook two important points. On the one hand, even if these traits are (per-



ceived as) positive, inferring that an individual has a low score in these traits would therefore be rather (perceived as) negative, and, on the other hand, the goal of the adversary is not to point out the positive/negative aspects of their targets' personality, but rather to gather information about their personality to better influence them afterward. Another important point to raise is that even if the inference accuracy showed in this work might be considered rather low, as we used only WAT *processed* data collected from a limited number of individuals during a limited amount of time, our results constitute a lower bound of what data brokers can do. On the one hand, they can access training data from many more individuals and thus can build stronger models. On the other hand, they can easily link WAT data with other types of data to improve the inference models. In their research, Aimeur et al. [291] showed how easy it is to link data of the same individual through different data broker databases. They voiced concerns about how easy it is to collect personal data about given individuals in general. Furthermore, it is known that few individuals read privacy policies and that among those who do, one-third claim to have no (or very little) understanding of what they read [319]. Considering this, and that most WAT users tend to forget about the (not always honest [53, 52, 54]) third-party apps they share their data with and highly underestimate their number, as shown in Chapter 3, it is likely that many data brokers have access to individuals' WAT data along with other types of personal data that can be used together to accurately infer personality profiles. Moreover, as Google recently acquired Fitbit [27] and plans to force Fitbit users to migrate their Fitbit account into their Google accounts [29], they will be in position to build the strongest possible inference models. Furthermore, the magnitude of this threat can only increase as the technology improves with the addition of new sensors (e.g., ECG), better sensor accuracy, and more efficient machine-learning algorithms. Furthermore, whereas rather low accuracy may not be considered a serious threat for particular individuals, the case is different when we consider an entire population (or a large part of it). Even if an accuracy increase of 15 points above the random baseline is not particularly impressive when about a single individual, on a large scale, it may help an adversary to better target thousands of people. This raises obvious privacy and societal issues, especially in light of the recent scandals related to personality-based influence campaigns. Additionally, our results may be considered rather low compared to other data sources (e.g., smartphone and app behavior, online social networks), however, WATs are still an emergent technology and companies regularly implement more and more (and more efficient) sensors and/or

functionalities and therefore collect more and more data (and more accurate data). Furthermore, unlike online social networks, for example, it is particularly astonishing to be able to infer certain personality traits simply based on acceleration, orientation, and light sensors (for movements and heart rate). This “unsuspected feel” may exacerbate the threat to the general population. Indeed, almost half of our participants answered that personality traits inference would be “Not at all precise” or “Slightly precise”. Nevertheless, WATs are indeed activity-centered and will probably therefore never collect data that can be used to infer only-social related traits such as agreeableness. As for conscientiousness, although we collected the battery level with the intention of inferring this trait (with the hypothesis that the most conscientious people recharged their devices more regularly), it is important to note that the battery of the Fitbit device that we used for this study lasts approximately five days and that, consequently, most of our battery data series are about 20 points long, unlike the step-count data series, for example, which has tens of thousands of points.

To address this threat, a first step is to raise awareness of it. This thesis chapter makes a contribution by providing concrete evidence of this threat based on a rigorous risk assessment. Based on this assessment, privacy protection techniques should be designed. A first protection technique would be to limit the amount of data shared with the service provider, keeping as much data as possible on the users’ devices. As all Fitbit users collected data are stored on Fitbit’s servers, a simple solution would be to let the user choose whether to store each type of data on Fitbit’s servers or to only store them on a personal synchronized smartphone/tablet. Except for some specific data, the raw sensor signal-processing is directly computed either on the WAT or on the smartphone. This means that as long as the user does not need to share personal data and the smartphone’s storage capacity is sufficient, they could increase their privacy while keeping the same level of utility. Furthermore, if a given piece of information needs more computing power than provided by the user’s smartphone, so it has to be processed on Fitbit’s servers, it can simply be deleted from the servers once transferred back to the user. This will leave the data inaccessible to most of the potential adversaries and reduce the data-leakage risks. Additionally, the data shared could be obfuscated to further enhance users’ privacy. A commonly used solution is to add noise to the data, which should be done in a controlled way in order to provide formal guarantees, such as differential privacy. However, we decided to evaluate a different, simpler (and so more understandable by users), technique which consists in

aggregating data over some period of time. For instance, only the daily step count or the daily average heart rate could be shared with the service provider. We showed the efficacy of such an obfuscation technique in Section 4.6. By doing so, an adversary loses substantial information about when the data has been collected, which is particularly useful as seen in Section 4.6 (e.g., steps at night). Indeed, our results suggest that only intra-day data brings information about personality traits. Therefore, an adversary whose goal is to infer individuals' personality would probably not obtain significant results using aggregated WAT data. Furthermore, in the case of the adversary being the service provider, it would still be able to store their users' (aggregated) data, and to provide them with attractive services. Indeed, recent works, including Chapter 3 of this Thesis manuscript, show that most users view this obfuscation technique as having little impact on their utility [46], and are inclined to use it when sharing their data.

Another possible solution would be to empower users by letting them choose which sensors to enable or disable and which data to keep on the device or share with the servers of the service provider.<sup>7</sup>

An important lead for future work is to evaluate the acceptability of such protection techniques by end users. Would users be interested in disabling some of their WAT sensors (and which ones)? Do users need to synchronize their data with the service provider (which data)? Do users need to synchronize their step counts for every minute and with a one-step precision? Indeed, research has shown that users usually do risk-benefit analysis or so-called privacy calculus when using wearable devices [321]. For example, when purchasing healthcare wearable devices, users trade-off receiving relevant and personalized health information, the sensitivity of this information, and the existence of legislative data-protection mechanisms [37]. However, some users are not fully aware of the potential privacy risks of WATs [132] and, as shown in Chapter 3, they also are not fully aware of their own data-sharing behavior and lack knowledge on the data-sharing ecosystem, which might negatively affect their utility-privacy trade-offs, and ultimately lead them to take wrong privacy decisions [102]. Some individuals are willing to decrease their privacy for an increase in utility, especially when they consider that the device provides them considerable benefits [113], whereas other individuals are willing to accept lower benefits to gain more privacy [322]. The latter users probably

---

<sup>7</sup>Note that Fitbit already enables their users to deactivate some sensors directly on some of its devices [320]. However, this option is not particularly highlighted on the user interface and is limited to a binary choice.

prefer to use WATs that implement protection mechanisms, even if the activation of such mechanisms decreases their utility. They could then trade off utility and privacy directly when using the device and fine-tune the parameters with respect to their concerns. This could be studied through the lens of privacy calculus [321, 37, 323].

## 4.9 Limitations and Generalization of the Results

Our work has some limitations, beyond those related to the use of Fitbit, as mentioned above. In particular, we only show that, for three of the five traits, WAT data can be used to reach significantly higher inference accuracies compared to the random baseline. Thus, future studies are needed to optimize the model and show that WAT data can be used to develop highly effective models for personality inference. Also, while we can assume that our ground truth is particularly accurate given the detailed questionnaire we relied upon, we want to highlight that the participants' answer quality could be degraded due to the well-documented respondent fatigue [324], as well as the social desirability bias [325]. There is clearly a trade-off between the details of the psychological profiles and the quality of the collected survey data. Furthermore, the participants' responses about their privacy concerns may have been biased as they were aware of the study's purpose. Additionally, while the study participants are somewhat representative of the local student population, they are not representative of the general adult population. Finally, a larger duration and a larger number of participants would have increased the significance of our results.

However, even if we study a particular type of population in this work, it is highly likely that our results can be generalized to a more global type of population. Although our model (trained on data collected from a specific type of population) can probably not be generalized to other types of populations, as for example, older individuals tend to not have the same activities as students (e.g., going out on Friday night), our methodology can. In the case of a different age population, for example, whereas specific behaviors change over time [326], different specific remarkable patterns could still be used by an inference model. Therefore, another inference model can be trained with data from another type of population and will probably reach a similar performance, but using features related to different life patterns. It is also important to note

that, even if our model is trained with data collected from young individuals, as personality traits are generally stable over time [292], data collected now from young individuals will still be useful to infer their personality in the future.

## 4.10 Conclusion and Future Work

In this chapter, we showed that WAT data can help classify users according to their personality traits, especially openness, extraversion, and neuroticism. We demonstrated that the use of WATs can create privacy risks that an adversary can potentially exploit. Our study is based on the WAT data of 204 individuals collected over a period of four months. We defined classes for each of the Big Five traits according to the dataset terciles and used different features extracted from information as step count, heart rate, sleep time, activities, battery level, and gender to train a ternary classifier for each of these traits. We conducted ternary classification and used accuracy as the evaluation metric and obtained results significantly higher than the baseline for openness, extraversion, and neuroticism. Also, we showed that, regarding prior work, using WAT data outperforms the use of call detail records (CDR) for inferring individuals regarding their personality traits. Furthermore, we studied the impact of data source removal on inference accuracy and pointed out that the model could reach even higher performance if trained on a larger dataset. Moreover, we analyzed the selected features and highlighted the most informative ones for each personality trait. We also showed that aggregating step count and heart rate by day is an effective obfuscation technique. Finally, we drew links with related studies and compared our results with theirs.

For future work, as noted in Section 4.6, we consider that it would be interesting to optimize inference models by exploring more feature combinations and by training and evaluating such models on larger datasets. To this end, additional data collection may be useful. For example, knowing that some WATs provide logging functionalities (e.g., meals and food intake), those data may be used to build features to improve the inference model (prior studies state that personality and dietary habits are correlated [327]). Also, profile information or device-usage data, as the number of “Fitbit friends” or the number of times where a user taps on the device’s screen, could be helpful to increase the inference accuracy. Furthermore, as highlighted in the introduction as well as the adversarial model of this chapter, personality inference as defined in this chapter is often considered as a privacy issue.

It would also be interesting to design and evaluate other obfuscation techniques. Indeed, it might be relevant to develop obfuscation techniques that result in less data loss, and thus, would have an even better acceptability than the one that we evaluated.

In this study, we focus on a particular adversary who has full access to user data. However, it could be interesting to consider adversaries who would have only partial access to the data and study what methods they might use to obtain these data. Furthermore, we focus on only one given type of device. It would be interesting to extend our study to multiple kinds of devices and evaluate, for instance, how the quality/quantity of sensors affects the inference accuracy. Moreover, in our study, we used data collected on a very specific population. Conducting a similar experiment on a more diverse population would be useful for studying whether our results can be extended to all categories of the population.

In Chapter 3 we have shown that, due to their lack of knowledge of the WAT data-sharing ecosystem and of awareness of their own behavior, users may adopt risky practices, and in this chapter, we have shown that WAT data can be used to conduct inference attacks. Therefore, it is crucial to develop privacy/transparency technologies to help them better manage their data-sharing, which is what we study in the next chapter.



## Chapter 5

# Our Data, Our Solutions: A Participatory Approach for Enhancing Privacy in Wearable Activity Tracker Third-Party Apps

**Abstract.** Users of wearable activity trackers (WATs) lack knowledge about data sharing. Most of them are not fully aware of their own data-sharing behavior. Therefore, it is crucial to design privacy-enhancing technologies (PETs) and transparency-enhancing technologies (TETs) to help them better manage their data-sharing hence to protect their privacy. In this chapter, we take a participatory design approach to design PETs/TETs, together with WAT users. We conducted three design sessions with 8-9 users in each session. During these sessions, the participants were able to propose and evaluate new PETs/TETs related to WAT-data sharing. The outcome of these sessions was 19 different designs that we then categorized into seven categories of functionality (design features). Multiple proposed designs can be compared to designs existing in other fields (e.g., social networks, mobile permission) as they offer similar functionalities. We then evaluated these different functionalities regarding their feasibility, effectiveness, adoption, and usability as PETs. Then, to propose a general solution, we selected three identified design features. Such a solution should implement functionalities



related to partial sharing, reminders, and revocation assistance. These functionalities were evaluated overall as highly feasible and effective; Finally, the participants found them very usable and to have a high adoption potential.

## 5.1 Introduction

In the previous chapters, we discussed the fact that though WAT data are generally kept on the user's device or on the service provider's cloud, the data can also be shared voluntarily by users with other individuals (e.g., family, friends, co-workers, healthcare professionals) and entities (e.g., employers, insurance companies, third-party service providers), typically through third-party applications (TPAs) or social network profiles. Users do so for increased social or financial benefits (e.g., better projection of the self, decreased insurance premiums) and/or for additional features not offered by the original services or application. For example, users might want to share some of their fitness data to take advantage of functionalities that are not natively supported by the service provider's applications or get financial rewards (e.g., WeWard [47], Fitcoin [49]). However, they could lose track of their TPAs [56], or some TPAs could collect more data than they need to provide their services [52] then share them with other parties and/or use them against the users' consent.

In Chapter 2, we analyzed users' awareness, understanding, attitudes, and behaviors toward fitness data sharing with TPAs and individuals. We explored users' practices and actual behaviors toward fitness data sharing and their mental models. Our empirical results showed 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 have forgotten that 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, demonstrated 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 can actually do this.

Therefore, it is crucial to set up privacy-enhancing technologies (PETs), as well as transparency-enhancing technologies (TETs [275]), to help the users better manage and keep track of their multiple applications and better understand how the fitness-data sharing ecosystem works. Such PETs/TETs

could indeed help them to avoid risky behaviors for privacy, such as sharing more data than is actually required or not regularly checking the previously granted permissions to revoke them if necessary. Few studies, including the work described in this thesis, evaluate the potential for adoption of such PETs (i.e., related to TPAs) [46], and others developed PETs in the context of WAT data sharing [56, 181, 184, 186, 187, 118, 190, 191, 192]. However, for these studies, the tools are designed by the researchers (and sometimes tested by users afterward). Although such tools could be evaluated by users afterward, involving users upstream in the design process would often highlight problems and solutions that developers and researchers would not have thought of, as they do not represent the core target. Furthermore, none of these studies is focusing on data sharing. There is no study, to our knowledge, that focuses on the design of PETs for WAT-data sharing that includes users in the design process, which constitutes an important gap in the related literature.

In this chapter, we report the results obtained by conducting a participatory design study with WAT users (N=26). In this study, we answer the following research questions:

- **RQ1:** What solutions will be suggested by WAT users to help them better manage data sharing to avoid risky behaviors for privacy?
- **RQ2:** What solutions will be suggested by WAT users to help them better understand the data-sharing process?
- **RQ3:** What solutions will be suggested by WAT users to obfuscate/aggregate their data in order to improve their privacy while keeping decent utility?

In this chapter, we report the designs proposed and evaluated by 26 WAT users during three participatory design sessions (8-9 users for each session). We collected and analyzed 19 different designs that we then categorized into seven categories of functionality (design features). Multiple proposed designs can be compared to designs existing in other fields (e.g., social networks, mobile permission) as they offer similar functionalities. We then conducted an expert session with two information security & privacy experts to evaluate these different functionalities regarding their feasibility, effectiveness, adoption, and usability as PETs (the same criteria were used by the participants for their evaluation) and propose a general solution that implements multiple design features proposed by the participants. This general solution should implement functionalities related to partial sharing, reminder notifications, and revocation assistance. We evaluate these functionalities as being overall highly

feasible and effective, and the participants evaluated them as being very easy to use (i.e., high usability) and as having a high adoption potential.

### Outcome

The chapter is organized as follows. We first describe our methodology in detail in Section 5.2. We present and analyze the results in Section 5.3. We then discuss the results and propose our general solution in Section 5.4 before briefly concluding in Section 5.6.

## 5.2 Methodology

In this study, we highlight the solutions WAT users propose to help them maximize their privacy when they share their data. To this end, we conducted participatory design sessions [328] with WAT users who share data with TPAs. Similarly to our proposition in Section 3.4.4, there are few studies about privacy-enhancing technologies (PETs) as new functionalities for data sharing [46, 181] or as tools to better understand privacy policies [151]. However, to our knowledge, all published works related to PETs for WAT-data sharing were about solutions designed by developers or researchers. Whereas, we think that users themselves could bring us particularly relevant perspectives and insight, as they are the first concerned by the usage of WATs and data sharing. Participatory design is a user-centric design approach that is used by designers to include the end users in the process of the design [329]. Such an approach has been used in multiple studies related to utility, including WAT utility [330, 331, 332] and privacy [333, 334]. Participatory design is particularly useful to develop solutions related to usable security and privacy [335, 336]. As the main problems that we want to solve in this work are related to the end-users behaviors and understanding, we need to develop solutions that are particularly adapted to their needs. By directly asking users to propose solutions, we, therefore, gather information from the individuals who are the most affected by privacy issues as well as the usage of the related technology (i.e., WATs). Whereas developers may have many biases related to their particular position, and fail to see potential problems with the usability of their solutions, directly asking end users to propose solutions may be a great help in highlighting new ideas, with a form of guarantee that they will be relevant and adapted to their usage and understanding of the technology. Furthermore, previous research has shown that a participatory design approach

may actually help researchers to develop relevant and efficient solutions to help users [333, 337]. In the framework of our participatory design sessions, we performed different participatory design activities in order to make the participants aware of the risks related to data sharing and promote participants' creativity, and generate effective solutions.

We designed our study according to the participatory design approach of earlier studies [328, 338, 333]. Most of the content presented to the participants during the design sessions was adapted from the findings of Chapter 3. LABEX, a dedicated structure of the University of Lausanne (UNIL) helped us to recruit the participants, as it manages a pool of around 8'000 students from two universities (a technical one, i.e., EPFL, and a general one, i.e., UNIL itself covering a broad range of disciplines). This is the same structure that we used in Chapter 4. The students who were interested in our experiment completed a screener survey that we used to verify their eligibility for participating. We scheduled three participatory design sessions.

### 5.2.1 Recruitment

We used an online screener survey to recruit participants (5 minutes to complete). This survey was as short as possible and was composed of only questions that are necessary to filter the participants regarding our criteria (see below), basic demographics to ensure recruiting a balanced sample (e.g., with respect to gender), WAT usage, and data-sharing behavior.

831 individuals answered the screener questionnaire and 54 were compatible with the experiment criteria. The recruitment criteria are the following: to regularly use a WAT device (at least 3 days a week) for more than six hours a day, to have used their WAT for at least six months (medium-term use [42]), to share data with at least one third-party application, and to speak French (i.e., the local language at the universities). For each session, 11 individuals were invited in order to finally obtain 9 participants, as we expected a few “no shows”. When more than 9 individuals attended a given session, the last ones to arrive were sent back with 10 CHF ( $\sim$  10 USD) in compensation. The latecomers were not compensated. If an invited individual withdrew before the session began, we invited someone else. In total, we invited 40 individuals, nine of them withdrew before the session began, 2 did not attend their session, and 3 extra participants were sent back (including one who was sent back without being compensated because they were late). Finally, 26 individuals were present for the sessions.

## Ethical Considerations

Before each design session, the participants had to sign a consent form that described the conditions of participation, the data being collected (and the associated data-management plan), the procedure for withdrawing from the study, and the information about the financial incentive. The institutional review board at our university validated the consent form and approved the study. Participants were paid 70 CHF ( $\sim$  75 USD) at the end of the session.

### 5.2.2 Session Procedure

We separated the participants into design sessions (i.e., nine people in each session) in order to conduct several focused design sessions. Each group attended one session. We conducted all three sessions in two days without overlap (one on the first day in the afternoon, and the others on the second day, respectively, in the morning and in the afternoon).

Although we conducted sessions with nine participants, we set up both general (all nine participants together) and group (groups of three participants) activities. To facilitate participatory design, the study consisted of six main parts: *pre-study (screeener) survey, introduction, setting up the situation, upgrading knowledge, sketching, and value ranking* (see Figure 5.1).

The participatory design sessions were conducted by three investigators: the thesis author who led the sessions as the session moderator and two assistants. We audio-recorded all sessions and took photos of the artifacts (after collecting consent from the participants). In the following paragraphs, we explain the procedure and our rationale for each part.

Figure 5.1 summarizes the timeline of a session. During each session, after welcoming the participants, we briefly introduced the concept of data sharing (what can be shared and with whom) and asked them some thought-provoking questions about privacy. We briefly presented, based on academic research and newspaper articles, the potential threats to privacy caused by WAT-data sharing. Once they were aware of the threats, we reconstructed, with them, the WAT-data sharing ecosystem. We presented them with the current literature knowledge about users' behavior and the understanding of data sharing with third parties. This presentation was based mostly on the findings of Chapter 3. Next, after briefly giving them a few tips about design, we set up discussions (in small groups) on how to improve users' understanding of the whole data-sharing ecosystem, their awareness of their own behavior, and the user experience, and on how to develop multiple solutions. The outcome of these

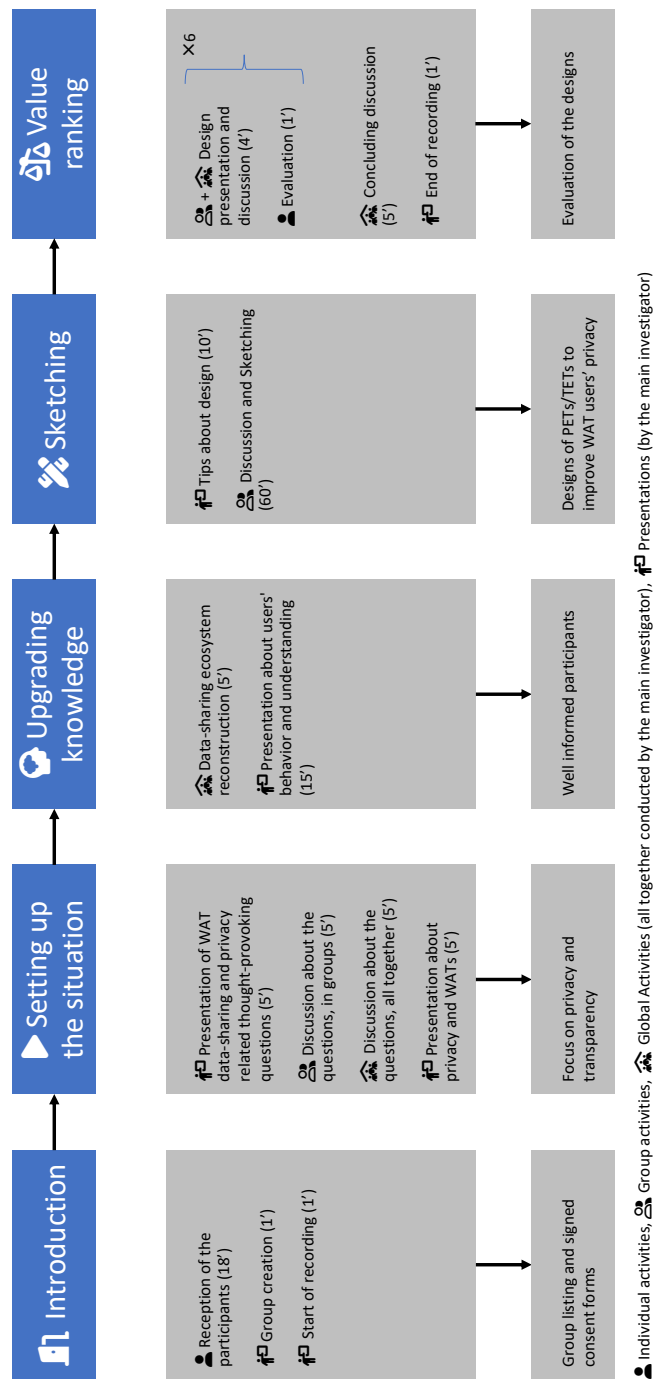


Figure 5.1: Session timeline. This figure summarizes the different steps of one participatory design session, shows the different activities and their expected outcomes.

sessions is PETs/TETs that assist WAT users in the data-sharing process, hence that increase their privacy. The form of solutions were storyboards or low-fidelity paper prototypes.

Prior to running the study, in order to refine the protocol, we ran the session with two researchers who are from our university, have expertise, respectively, in distributed systems and information security & privacy, and who did not take part in this research.

The following sections summarize the different parts of each session. During all parts involving the participants discussing with each other (generally in small groups of three individuals), we encouraged them to share their own user experience and to raise the positive and negative points of their own experience with data sharing.

### **Introduction (20 min.)**

The participants were invited to attend participatory design sessions. They were asked to wear their WATs and to bring their phones. Two researchers were present to welcome them. The first checked their identity and brought them into the room, while the second asked them to fill out and sign the consent and payment forms. When every participant was welcomed and had signed their forms, they sat around a table, they were free to choose their seats. Once everyone was seated, we asked some participants to switch places in order to form a gender-balanced group, we began recording the session (audio and video), and the participatory design session officially started. We followed the methodology from earlier participatory design studies [328, 338, 333]. Before commencing the first activities, we described the schedule and reminded the participants about the main goals of this study (i.e., designing tools to help users better manage their data-sharing and/or better understand the data-sharing process).

### **Setting Up the Situation (20 min.)**

We began each session by briefly presenting how WAT-data sharing can impact users' privacy. One of the investigators (this thesis author) first showed them the different ways, for a WAT user to share their data, and he displayed a short video showing them how to grant and revoke access to the data to a given TPA (Strava). Then, we asked a few thought-provoking questions about data privacy:

- Who do you think might be interested in accessing your fitness data, and why?
- What do you think it is possible to do with or learn from your fitness data?

Participants discussed these questions with their group mates for five minutes. By asking them to discuss these questions in small groups, we encouraged everyone to participate and think deeply about these questions, as it provides more opportunities to express their opinions and thoughts to others. Indeed, pedagogical research has shown that, in comparison to simple lectures, asking people to discuss specific questions in small groups increases their engagement and retention of knowledge [339, 340]. Then, all participants (from all groups) shared and debated their answers to these questions; they raised additional related questions and answers. This discussion was supervised by the main investigator. The goal of this activity was to ensure that the participants were aware of the problem and that they share their personal concerns and experience with data sharing. After the discussion, the main investigator briefly presented the potential threats to privacy caused by WAT-data sharing. This presentation was based on academic research [4, 8, 9, 12, 13] and newspaper [30, 26, 341] articles.

### **Upgrading Knowledge (20 min)**

We discussed the process of WAT-data sharing with TPAs and how the data-sharing environment works [52, 61]. Together with the participants, we reconstructed the data flow by asking them (and correcting them if they are wrong) what the different entities are, what their relations are, and how the data are shared between them. Our purpose at this step is to reconstruct a correct drawing of the process on a flip chart. As active learning increases knowledge acquisition and performance [339], by involving everyone in this process, we increase their engagement thus ensuring that they indeed acquire a correct mental model of the ecosystem. However, as we wanted to provide them with only the correct model and needed the activity to be reasonably brief, we directly conducted this activity all together. We also briefly presented the current literature knowledge about users' behavior and their understanding toward data sharing with third parties and the related threats; in particular, the findings discussed in Chapter 3. This was done by showing them a short presentation with slides. Thus, we helped the participants to be at the same






Figure 5.2: Photo of one of the tables during the sketching phase.

knowledge level and to be aware of the problems we wanted to address during the session. It is important for participants to have the correct mental model in order for them to be able to design solutions that indeed apply to the described problems and that can be implemented in a way that corresponds to a specific and existing layer of the system. In order to not influence them with regard to the solutions they could have proposed, we discussed only our findings on users' behavior and understanding and did not present the results about our own PET designs (i.e., the countermeasures presented in Chapter 3 and Chapter 4).

### Sketching (70 min)

We conducted a short presentation to give the participants tips about designs and how to sketch, or storyboard [342, 343]. By doing this, we ensure that they have all the necessary information to create designs that we will be able to correctly interpret and classify. We then asked each group to design and propose at least two solutions to improve users' privacy related to data sharing (more if they have more ideas). We asked them to do this in three steps: (1) Determine one or two specific problems (challenges) that they want to solve. (2) Imagine at least two new functionalities/solutions to improve the problems (i.e., either two solutions for one problem or two solutions for two



Figure 5.3: Photo of one of the groups presenting one of their design.

problems). And (3) draw sketches to visualize how it could be implemented in the data sharing (WAT device, companion app, servers, ...). Figure 5.2 shows one of the tables during the sketching phase as well as the drawings produced by the corresponding group. Time permitting, we welcomed more solutions from each group. We provided the participants with large paper sheets (A3), sticky-note papers, colored pens, and markers. Each group worked separately from each other. There were no interactions between groups at this stage. The investigators went from time to time to the different tables to observe the progress of the activity. To do this, the investigators asked a few questions, without too much priming, to understand where participants were in the definition of their problem and/or the design of their solution and also to check that participants understood the process and had no questions about what they were doing.

### **Value Ranking (30 min)**

In order to compare the reactions the WAT users had about the proposed designs with our own evaluation as security & privacy experts, we asked the participants to evaluate them. To do so, each group (either one or multiple persons by group) presented their sketches and discussed them with the other session participants. Each presentation (5 min.) was in three phases: presentation, questions and answers, and evaluation. Figure 5.3 shows participants during the presentation of one of their designs. After each presentation, each

participant (they had to indicate if they were one of the designers or not) was asked to evaluate the proposed solution regarding two points: usability [344] (i.e., if the solution is easy to use) and adoption [345] (i.e., if they would use the solution in everyday life). For each of these points, they attributed a grade on a five-point Likert Scale. Their grades were collected using an online form that the participants could access with their phones. The participants graded the solutions before we proceeded to the coding (see Section 5.2.5) in order for us to identify specific design features from the proposed solutions, as grading the identified features after the sessions would have required us to contact each participant once again after the coding process. Therefore, the evaluation by the participants is directly related to the proposed solutions and not to the design features that we identified in these solutions and that we discuss in the results section. However, we consider that this evaluation still provides us with insightful information about how the participants perceived the different proposed functionalities. After grading, one of the investigators collected all the material (text and drawing) related to the presented design. Furthermore, the evaluation was anonymous (we know only if the evaluation of a given design was done by one of its designers or not). Hence, we minimized all social biases that could have influenced their evaluation. We then asked them if there were any comments or questions about the session, or information security & privacy in general, and we discussed them if necessary. After the session, each participant was paid in cash upon leaving.

### 5.2.3 Room Layout

As shown in Figure 5.4, during the sessions, we arranged the room as follows: We placed three tables in the middle of the room, along with three chairs each, each one can accommodate three participants. At one end of each table, standing on a stool, we placed an audio recorder to record the discussion at the table. At the beginning of the session, the main investigator (i.e., the author of this thesis) was at the back of the room, or in front of the participants when sitting. The main investigator was equipped with a laptop, a beamer, and a flipchart in order to supervise the general discussions and the presentations. Two video cameras are also placed in the room: One, fixed on a tripod stand, was in one of the corners of the room, thus enabling a global vision. The other camera was placed on a long table and turned towards the presentations and was regularly moved by a second investigator in order to film pieces of the discussion taking place at certain tables during the session.




Figure 5.4: Layout of the room where all the participatory design sessions were conducted. To conduct the sessions, we used three tables along with three chairs each as well as drawing material (e.g., paper sheets, pens), a flipchart, and a video projector (and a connected laptop). To record the sessions we used three audio recorders (one per table) as well as two video cameras.

### 5.2.4 Participants & Groups Composition

The session participants and the groups were composed in such a way as to obtain a gender balance. Table 5.1 in the appendix shows the details about the sessions and group composition. We conducted three different sessions with three groups of three people, except for Group 2 of Session 1, which was composed of only two participants. Among the participants, 42% were women (11 participants), and 58% were men (15 participants). They were 21.1 years old on average, with a standard deviation of 2.5 years. The participants wore their WAT 5.9 days a week on average, with a standard deviation of 1.4 days. The days they wore the device, 35% of them wore it during 7 to 12 hours, 27% during 13 to 18 hours, and 38% during 19 to 24 hours. The participants were composed of 65% of Apple users, 12% of Fitbit users, 19% of Garmin users, and only one of them (4%) had another type of device. Half of them (50%) share their data with only one TPA, 42% with two to five TPAs, only one of them (4%) share their data with six to nine TPAs, and also only one of them

| Session | Group | Gender | Age   | Days | Hours | Nb TPAs | Device |        |
|---------|-------|--------|-------|------|-------|---------|--------|--------|
| 1       | 1     | Woman  | 23    | 6    | 13-18 | 1       | Garmin |        |
|         |       | Man    | 30    | 7    | 13-18 | 10+     | Apple  |        |
|         |       | Man    | 20    | 7    | 13-18 | 1       | Apple  |        |
|         | 2     | Woman  | 22    | 7    | 19-24 | 2-5     | Fitbit |        |
|         |       | Man    | 22    | 4    | 7-12  | 2-5     | Garmin |        |
|         | 3     | Woman  | 19    | 7    | 7-12  | 2-5     | Apple  |        |
|         |       | Man    | 22    | 7    | 19-24 | 1       | Apple  |        |
|         |       | Man    | 20    | 7    | 19-24 | 6-9     | Apple  |        |
|         | 2     | 4      | Man   | 25   | 4     | 7-12    | 2-5    | Garmin |
| Woman   |       |        | 22    | 7    | 19-24 | 1       | Fitbit |        |
| Woman   |       |        | 19    | 5    | 19-24 | 2-5     | Garmin |        |
| 5       |       | Woman  | 24    | 6    | 13-18 | 2-5     | Apple  |        |
|         |       | Man    | 19    | 5    | 13-18 | 2-5     | Apple  |        |
|         |       | Man    | 20    | 5    | 19-24 | 1       | Garmin |        |
| 6       |       | Woman  | 23    | 5    | 19-24 | 1       | Other  |        |
|         |       | Man    | 21    | 7    | 13-18 | 1       | Apple  |        |
|         |       | Woman  | 20    | 7    | 13-18 | 2-5     | Apple  |        |
| 3       |       | 7      | Man   | 21   | 7     | 7-12    | 2-5    | Apple  |
|         |       |        | Woman | 20   | 5     | 19-24   | 1      | Apple  |
|         |       |        | Woman | 21   | 5     | 7-12    | 1      | Apple  |
|         | 8     | Woman  | 18    | 7    | 7-12  | 1       | Apple  |        |
|         |       | Man    | 20    | 7    | 7-12  | 1       | Apple  |        |
|         |       | Man    | 20    | 3    | 7-12  | 2-5     | Apple  |        |
|         | 9     | Man    | 20    | 7    | 19-24 | 1       | Apple  |        |
|         |       | Man    | 19    | 7    | 19-24 | 2-5     | Fitbit |        |
|         |       | Woman  | 19    | 3    | 7-12  | 1       | Apple  |        |

Table 5.1: Details of participants for each session and group.

share their data with 10 or more TPAs.

### 5.2.5 Coding Process

After all the sessions, we collected 19 drawings that represent the participants' designs (all groups submitted two designs except for one group that submitted three). Then, we used open coding [269] to categorize the multiple functionalities (i.e., design features) included in the different designs. Two of the researchers working on this study (the coders) independently developed a codebook on their own before pooling and discussing their respective re-

sults. By doing so, we noted a high rate of overlap between the codes and design feature categories defined in the two codebooks. After comparing both codebooks, one of the coders (i.e., the author of this thesis) built a new final codebook by merging overlapping codes and designing feature categories; then, the second coder reviewed and provided feedback. Finally, both coders reached an agreement on the coding. In total, as shown in Table A.1, we identified 16 distinct codes classified into seven categories (or themes). As they describe design features, and one specific design could implement multiple features, these codes and categories are non-exclusive. As a result, regarding their functionalities, each design could correspond to multiple design-feature categories.

### 5.2.6 Expert Review Meeting

After having coded all the proposed designs, we were able to classify them into seven design-feature categories, each category corresponding to one specific type of PET/TET. In addition to the evaluation of the design provided by the participants themselves, and in order to provide a more informed perspective on the way these solutions could be implemented and used, we evaluated these different types of technologies (i.e., the pre-defined categories) during an expert evaluation session with two information security & privacy experts, i.e., the two other researchers who work on the study and who did not take part in the coding process. The two experts are professors in our institution and their specializations are Information Security & Privacy (Expert 1), and Cybersecurity (Expert 2). The coders (and in particular one of them, the author of this thesis) first welcomed the experts and presented the protocol of the review meeting (10 min.). Then, we proceeded as follows, individually for each design feature category (10 min. each): The author of this thesis presented the design feature to the others (2 min.). During the presentation, the author of this thesis showed one slide that included the name of the design feature, a brief description, and a few examples of drawings that are the most representative of the feature. Then, the experts discussed it (5 min.). During the discussion, they could ask questions to the coders. Finally, they provided feedback (3 min.). This feedback consisted of a graded evaluation on a 5-point Likert scale of the feasibility [346] (i.e., if it is feasible to develop) and effectiveness [347] (i.e., if it is effective to protect users' privacy) of the feature, and of a free discussion including (1) comments about their graded evaluation, (2) any suggestions on how to improve it, and (3) any example

that such design is implemented in a different context. During the evaluation, the coders also graded the different designs, together with the experts. However, in order to not be influenced by the experts, they did this before hearing the experts' comments. The author of this thesis allocated sufficient speaking time to both experts. During each step of this session, the experts had pens and paper at their disposal to take notes and write down their ideas and notes when required. The session was audio recorded. The session lasted approximately 90 minutes.

## 5.3 Results

In this section, we present seven features extracted through the coding process, from the 19 designs collected during the participatory design sessions. For each feature, we begin with a complete description of its functionalities, then, if possible, we provide examples of similar features that already exist in another context (e.g., mobile phone permissions, social networks), that usually help users to easily monitor what type of data is shared with which application. We present qualitative evaluations related to the feature, before presenting the quantitative evaluation by the participants and the evaluation by the experts which are also summarized, respectively, in Table 5.2.

Table A.1 in the appendix summarizes the results of the coding with all the features (categories) and codes that were identified during the coding process.

### 5.3.1 Feature 1 - Partial Sharing

Partial sharing enables WAT users to share only part of their data according to a specific time frame or a given context. In fact, granting access to WAT data permits the TPAs to access every data of a specified type regardless of when the data has been collected by the device. In other words, a TPA can access WAT data that was collected before a user granted access. Using this feature, the user would be able to choose a specific data-collection time frame that they want to share (excluding the others). Feature 1 was present in three different designs proposed by the participants. Whereas one of the designs allows users to select a time frame by indicating dates, another one simply enables the user to choose between sharing all the data or only the data that has been collected since the access was granted. In a different approach, the third design, shown in Figure 5.5, is context-aware, allowing users to indicate the data type and the activity type they would want to share (e.g., sharing



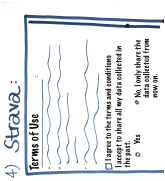
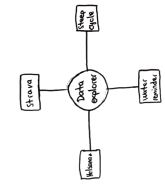
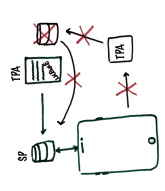
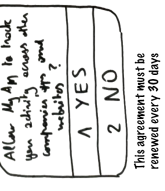
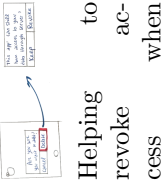
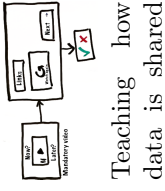

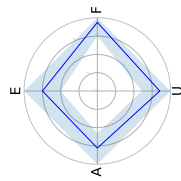
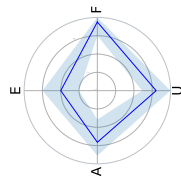
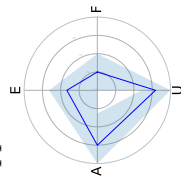
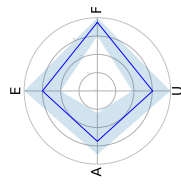
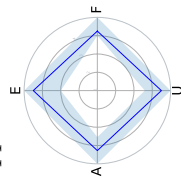
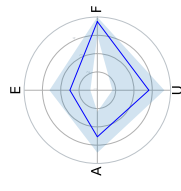
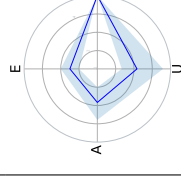
|  |   |  |  |   |   |   |
|--|---|--|--|---|---|---|
| <p><b>Partial sharing</b></p>  <p>4) Strava:<br/>Terms of Use<br/>I agree to the terms and conditions<br/>I agree to share all my data collected in<br/>the app with other users.<br/>No. I do not agree to the<br/>terms and conditions.</p> | <p><b>Visualisation</b></p>  <p>Other ways to visualize which data is shared with which TPA.</p> | <p><b>Centralization</b></p>  <p>Specific TPA's app store or plugins directly implemented in SP's mobile app.</p> | <p><b>Reminders</b></p>  <p>Reminding that data is shared. Periodic "opt-in" or "opt-out" renewal.</p> | <p><b>Revocation Assistance</b></p>  <p>Helping to revoke access when uninstalling a TPA's mobile app.</p> | <p><b>Sensitization, Education</b></p>  <p>Teaching how data is shared and what are the risks.</p> | <p><b>TPAs limit</b></p>  <p>Limit the number of TPAs users can share their data with.</p> |
| <p><b>&amp; Android permissions</b></p>  <p>SP</p>  | <p><b>SP</b></p>  <p>SP</p>  | <p><b>SP, TP</b></p>  <p>Facebook privacy<br/>checkup</p> <p>SP, TP</p>   | <p><b>SP</b></p>  <p>SP</p>  | <p><b>SP, OS (low involvement)</b></p>  <p>Trading apps</p> <p>SP, OS (low involvement)</p>                | <p><b>SP</b></p>  <p>WhatsApp limit for device pairing</p> <p>SP</p>                               | <p><b>SP</b></p>  <p>SP</p>  |

Table 5.2: Presentation of the seven identified design feature categories. For each of them (top to bottom), we provide (1) the title, (2) a translated example, (3) a short definition, (4) the result of the evaluation by the participants and the experts – The scores are from 1 to 5 for feasibility (F), effectiveness (E), adoption (A), and usability (U), graded by the participants – (5) examples of similar existing designs in other contexts if any, and (6) the stakeholders who should take action to implement the design who are either the service provider (SP), the third-party (TP), or the company providing the smartphone's operating system paired with the WAT (OS).



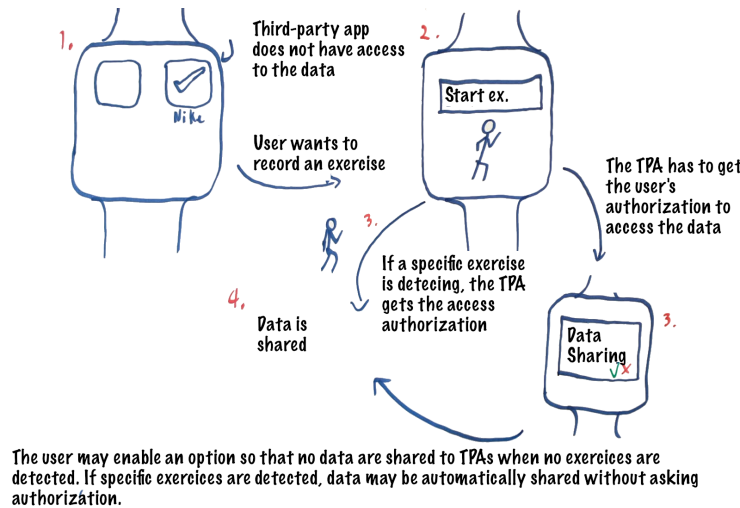


Figure 5.5: Translated version of the design that proposes partial sharing (Feature 1) regarding specific contexts (exercises) and/or time frames (when the user enables data sharing). In this version of the design, we replaced all the text (written in French) with an English translation.

only the heart rate data that were collected while running). This last design also proposes a “start sharing” feature that the user could enable and disable to select the time frame during which the data is shared (and only the data collected during this specific time frame). Such features should be implemented by the service providers (e.g., Apple or Fitbit).

This feature received positive feedback from the experts. Expert 1 mentioned that “[...] it’s also pretty good addressing one of the issues detected in previous work, that is people misunderstand that when you grant access you also grant access to data that was collected in the past.” [61]. The same expert added “[...] find it a bit restrictive in a way, they [the participants] could have gone further other than just basing their access control on time [...] I think they could have imagined other mechanisms like the granularity of the data and so on [...]”.

Regarding the scores, this feature received the second-highest score given by the experts, for feasibility (4.75) and effectiveness (4.00). As for the evaluations by the participants, this feature received the second-highest score given by the participants, for adoption (4.12) and usability (4.41). We can, therefore, affirm that, with all scores above 4, this feature is particularly appreciated.




Figure 5.6: Screenshot of the current Garmin companion app showing the TPAs list. When the user taps on one of the TPAs, it opens a panel showing all types of data that is shared with that TPA.

### 5.3.2 Feature 2 - Visualization

Solutions under this category aim to help users have a better overview of their data-sharing behavior by designing new visualization tools. These tools can help the users explore the shared data and the different TPAs with whom they share their data by classifying them either by data type or by TPAs. Some proposed visualization features also allow users to keep track of all shared data through a logging system and by displaying an accurate data-sharing history. Finally, such a feature can also help users monitor their own behavior toward data sharing by presenting them with specific statistics about their usage of the different TPA services that are installed on their phones. Feature 2 was present in five different designs. Currently, most platforms allow users to check a list of connected TPAs (e.g., see Garmin interface in Figure 5.6). However, no companion app provides a list of TPAs classified based on the type of shared data. Such features should be implemented by service providers. An example of Feature 2, in a different context, is available on iOS and Android for access management of mobile applications, as shown in Figure 5.7. The behavioral and log statistics feature is also similar to the macOS screen time.<sup>1</sup>

For the experts, the weakest aspect of Feature 2 is effectiveness. Expert 2 mentioned that *“The main drawback of this approach is that it’s not very effective, maybe it can lead to a change or increase awareness of the user [...] but as a mechanism itself it is not directly protecting privacy.”* This feature received a low score for effectiveness (3.00). However, feasibility received the second-highest score given by the experts (4.75, tied with three other features). As for the evaluation by the participants, this feature did not receive a high score for adoption (3.82); however, it received a decent score for usability (4.21). We think that such a solution might be perceived as useful and multi-

<sup>1</sup><https://support.apple.com/en-us/HT210387>

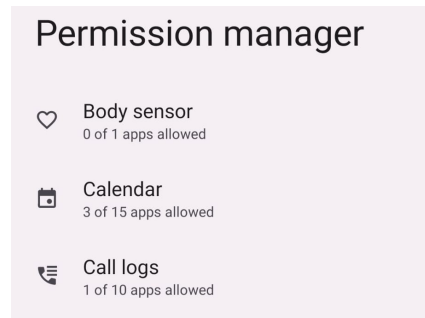


Figure 5.7: Screenshot of the permission management panel on Android. When the user taps on one of the data types, it opens a panel showing all applications having access to that data/sensor.

ple WAT users might be interested in accessing information about their data sharing. However, most of the users will probably not use it, as they have to actively check a dedicated section in the service provider’s mobile app, which is already quite complex. Indeed, previous research on online social networks and Android permission has shown that most of the users never update or even check their privacy settings [57, 247].

### 5.3.3 Feature 3 - Centralization

Centralization is not a new feature but rather a solution that guarantees secure data sharing among users. Two different solutions were proposed. The first solution suggests that the main service provider should have its own TPA app store. Any TPA interested in offering services in the app store would first need to obtain approval from the main service provider. This approval would act as a guarantee to users that the TPA will confidentially and securely process their data and that their privacy will not be compromised. In a slightly different context, a known example of this feature is Google Play’s privacy labels<sup>2</sup> (or data safety section), allowing developers to disclose information about their app’s data collection, sharing, and security measures. The second proposed method is to eliminate the possibility of sharing users’ data with TPAs and replace it with a plugin system directly integrated into the main application. This solution would guarantee that the users’ data would not be stored on the TPA’s server at any moment, as the main service provider would still be the data-processing entity. Feature 3 was present in three different designs.

<sup>2</sup><https://blog.google/products/google-play/data-safety/>

This solution would involve the active participation of service providers and the TPAs' companies.

This feature received multiple criticisms from the experts in general. [Expert 1]: *“The only positive aspect I see is that a dedicated store would force TPAs to be more transparent about what they really do with the data, but regarding the plugin solution, I don't think that such solution can be put in place.”*. [Expert 2]: *“In terms of feasibility, I don't even know how it could be done.”*. This feature received by far the lowest evaluation for feasibility (2.00), and the second lowest score for effectiveness (2.67). As for adoption and usability, it received average scores (respectively, 4.00, and 4.18).

### 5.3.4 Feature 4 - Reminders

Feature 4 is designed to address the well-known problem of users forgetting to revoke access [56, 61] by proposing notification reminders. Such a system could simply remind users periodically that they are sharing their data with TPAs. Multiple designs propose further engaging features by directly asking the users to renew the previously granted access (i.e., to opt-in again) or by asking them if they want to revoke it (i.e., to opt-out). Similar reminder mechanisms were implemented in other contexts. For example, Facebook implemented a privacy checkup system [348] to periodically remind users about the TPAs they share their data with and ask them if they want to revise the access authorizations. Such features could be implemented by service providers. Feature 4 was present in seven solutions. Figure 5.8 depicts one of the examples.

This feature was generally well-perceived by the experts. [Expert 1]: *“I don't think it's gonna solve the privacy issues all together [...] but will it solve an existing problem? I think yes, absolutely, it solves the problem of forgetting. [...] and I think it would be used.”*. [Expert 2]: *“I gave pretty much the same scores as for [Feature 2], except for the effectiveness because [...] at least it prompts the user to take some action [...] so it's a bit more effective than just being transparent [...]”*. This feature received the second highest score for feasibility (4.75) and effectiveness (4.00), and even if the mean score for usability (4.03) is not one of the highest, it is greater than 4, which is a decent score. As the score of adoption (3.75) is slightly lower than 4, we would recommend implementing that feature with an option to disable it or choosing the reminder frequency to avoid bothering users who do not want to use it.



Figure 5.8: Translated version of a design that implements a reminder notification feature (Feature 4).

### 5.3.5 Feature 5 - Revocation Assistance

This feature assists users in revoking data access. Two of the related proposed solutions include features for directly asking the user if they want to revoke access to their data when they *uninstall* a TPA’s mobile app from their phone. This feature is relevant because some users would be concerned about their data being deleted after uninstalling TPAs [61]. Feature 5 was present in three different designs. One of these designs also includes an automatic data-revocation option for when a TPA’s mobile app is not used for a while. A similar technique was implemented by Google on Android phones called “Remove permissions for unused apps”<sup>3</sup> to automatically remove permissions for apps than you did not use for a certain amount of time. The third design implements an option for directly sending a message to the TPA’s company to ask them to delete any related data that are stored on their servers. This feature is supported by Article 17 of the General Data Protection Regulation (GDPR) about the “right to be forgotten”.<sup>4</sup> Figure 5.9 is one of the designs implementing a feature that would enable revoking access while uninstalling a TPA’s mobile app on the phone. Service providers should implement such features, and depending on the specific version of the feature, it may also require the involvement of the company that provides the OS of the phone (e.g.,

<sup>3</sup><https://support.google.com/android/answer/9431959?hl=en#zippy=%2CAutomatically-remove-permissions-for-unused-apps>

<sup>4</sup><https://gdpr.eu/right-to-be-forgotten/#:~:text=In%20Article%2017%2C%20the%20GDPR,originally%20collected%20or%20processed%20it>

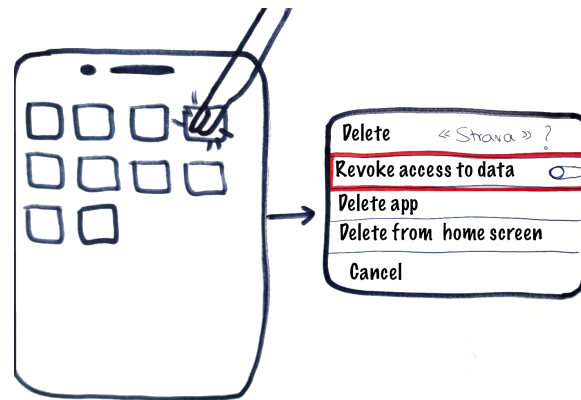


Figure 5.9: A translated example of design implementing Feature 5 enabling revoking access while uninstalling a TPA’s mobile app on the phone.

sending a revoking request when uninstalling the app). Furthermore, the option to automatically send a data removal request to the TPA company could also be imposed by law, as it corresponds to an article of the GDPR.

This feature received one of the most positive feedback. [Expert 2]: *“I have a strong opinion on this option; I think it would increase privacy overall without decreasing utility [...] I guess it’s a good option.”*. This feature received the highest score for effectiveness (4.50) and a decent score for feasibility (4.50). It also received the highest mean scores for adoption (4.28) and usability (4.50). As we can see, except for feasibility (for which it still received a decent mean score), this feature is the best-rated one.

### 5.3.6 Feature 6 - Education & Sensitization

Participants proposed adding a tutorial or awareness-raising video during the data-sharing process. Such a video would serve as educational design friction to encourage users to be mindful and considerate about the consequences of WAT data sharing, hoping it can be more effective than the typical text-based “terms of services”. Earlier literature showed that users usually would skip reading such text-based privacy notices [126]. Feature 6 was present in four different designs. One of the designs, shown in Figure 5.10, proposes to show a short video to the user in order to explain to them how data-sharing works and what are the multiple related risks to their privacy. This design also specifies that after watching the video, the users would have to answer a short quiz, and if they fail, they could not share their data. The fourth design aims to




Figure 5.10: Translated version of a design implementing a sensitization video (Feature 6) and requiring to pass a test to share data with TPAs.

implement an informative and interactive consent form, enabling the users to click on different links to obtain more information about how their data is processed. Such features should be implemented by the service provider or by the TPA’s company. The use of educational videos as privacy-preserving interventions has been proposed in various contexts, such as for multiparty privacy conflicts on social media [349]. Also, trading apps usually offer brief training when users create an account.<sup>5</sup>

This feature did not receive much positive feedback from the experts (except for feasibility). They found forcing users to watch a video challenging because they could be doing something else while the video played. Besides the possibility of refraining from watching enforced videos, Expert 1 also thought that such interventions harm the sense of gratification that users would perceive when using a new technology. [Expert 1]: “[...] *This is not promising [...] you just installed the Strava app, you want to test it immediately, your interest in such things is modest.*”. Despite a decent score in feasibility (4.75), Feature 6 received the lowest score for effectiveness (2.50). Furthermore, it received the second-lowest score for adoption (3.55) and usability (3.82).

### 5.3.7 Feature 7 - TPAs Limit

This feature aims to limit the number of TPAs the users can share their data with. If a user wants to share their data with a new TPA and this number is already reached, they will first have to revoke a previously granted access. Only one design implements this feature. Such features should be implemented by the service provider. Such a limitation is implemented, for example, in the messaging app WhatsApp that permits linking an account to only four different

<sup>5</sup><https://www.degiro.ch/helpdesk/en/trading-possibilities/why-do-i-have-complete-test-i-can-trade-product>

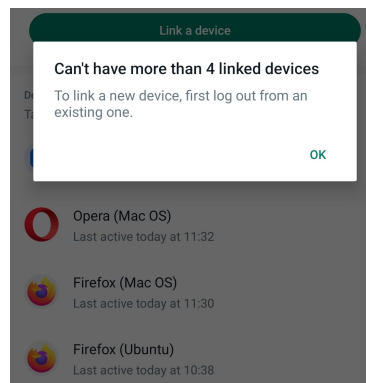


Figure 5.11: Screenshot of the WhatsApp linked devices panel. The user cannot link their account with more than four different devices at a time.

devices at the same time, as shown in Figure 5.11.<sup>6</sup>

Except for feasibility, this feature received mostly negative feedback from the experts. [Expert 1]: “[...] if that feature could be enabled or disabled, I’m pretty sure everyone would disable it on the first time they get prevented from installing something [...]”. As for the participants, they mostly seem to dislike it. [Man, 20 y.o., Apple]: “It’s kind of annoying to have a limit, let’s say someone needs a lot of apps.”. In this case, the experts (and the coders) gave a low score for effectiveness (2.50). Indeed, even if the feature of having a limited number of TPAs with which users can share their data would certainly increase users’ privacy, users would not like it and would not want such a feature to be implemented. Furthermore, users could simply revoke/grant access multiple times, which does not help them. Considering the simplicity of this feature, it received the highest mean score for feasibility (5.00). As for the evaluation by the participants, this feature also received the lowest mean score for adoption (2.83) and usability (3.17).

## 5.4 Discussion

We investigated the widespread problem of user data-sharing in the context of third-party applications (TPAs) and wearable activity trackers (WATs). Through participatory design sessions, our participants provided us with multiple designs in order to help them better manage their WAT-data sharing and

<sup>6</sup>[https://faq.whatsapp.com/378279804439436/?helpref=uf\\_share](https://faq.whatsapp.com/378279804439436/?helpref=uf_share)



protect their privacy. These proposed solutions offer novel insights into the future design and development of privacy-enhancing technologies for WAT-data sharing with TPAs. In the following parts, we further discuss these findings, including their limitations and technical feasibility, and envision possible combinations of these solutions to build effective PETs.

After having classified and evaluated the various proposed features by experts, we found that a general solution combining Features 1, 4, and 5 would be a promising tool to help WAT users effectively protect their privacy. Next, we revisit these three features.

Enabling the users to **share selectively based on context, or specific timeframes** (i.e., which data and activity type they want to share regarding the time it was collected or the corresponding activity) could address one major misunderstanding regarding data sharing as users tend for example to think that they only share the data that was collected from the moment the granted an access authorization, which is not the case as once a TPA has access to a user's given type of data (e.g., step-count, heart-rate), they can access all data corresponding to this type, regardless of when it was collected [61]. Furthermore, it could likely increase user privacy by substantially reducing the amount of personal data<sup>7</sup> that a potential adversary would have access to. As suggested during the evaluation by the experts, it could be particularly interesting to also limit the amount of shared data by allowing users to share aggregated data. Indeed, previous research already discussed options to share data aggregated over time (e.g., aggregating the data series by the day) [187] and showed that it is an effective technique for mitigating inference attacks [63] and is likely to be adopted by a large number of WAT users [61].

Mechanisms such as **reminder notifications** and “opt-out” or “opt-in” access-authorization renewal were also evaluated as having high usability and effectiveness (especially according to the evaluation by the experts). An advantage of such solutions is their feasibility to develop them without many technical challenges. A similar feature was also proposed and evaluated in previous research [61], showing that WAT users are particularly inclined to use reminder notifications. However, we recommend implementing only “opt-out” renewal, as “opt-in” could cause utility issues because such a feature would revoke the access if the user ignores the message. Furthermore, the user should be able to choose the frequency of such notifications or disable them, for example, by checking a box that appears with the notification (e.g., “don't

---

<sup>7</sup>following the concept of data minimization <https://edps.europa.eu/data-protection/data-protection/glossary/d.en>

ask me again”).

The feature allowing users **revoke data access** when uninstalling a TPA’s mobile app or asking a TPA’s company to remove data from their servers received the most positive feedback from the participants and the experts. Therefore, we find that such a protection mechanism should be implemented. Indeed, as multiple WAT service providers implement data access for TPAs by using API keys (e.g., using services as OAuth [43]), it might not be evident for users that the access authorization is not necessarily revoked when they delete a TPA’s mobile app from their phone and that the TPA can still access their data from server to server. Solutions in Feature 5 could not only remind the users to revoke the access but also teach them that they must do it if they want to stop sharing their data with a given TPA. Furthermore, a feature to help WAT users ask a TPA’s company to remove data from their servers is not only a particularly good feature for increasing privacy but is also conformed with Article 17 of the GDPR: “*The data subject shall have the right to obtain from the controller the erasure of personal data concerning him or her without undue delay and the controller shall have an obligation to erase personal data without undue delay [...]*” [350]. However, simple notifications, as suggested in Feature 4, would be preferable to automatic revocation, as the former could cause utility issues (e.g., an access authorization being removed without the user noticing).

Therefore, we propose a meta-solution called **RePaRe**, which stands for REminder, PArtial sharing, and REvocation assistance. **RePaRe** is a comprehensive approach comprising the previously mentioned design features to help WAT users better manage data-sharing. It implements partial sharing (i.e., timeframe, context, and temporal aggregation), periodical reminders with “opt-out” renewal (i.e., the user has to revoke the access actively) as well as a disabling option, and an option to revoke access authorization when uninstalling the TPA’s mobile app from the phone as well as the option to send an automatic data removal request to the corresponding company.

Figure 5.12 shows the workflow of WAT-data sharing and the different features of **RePaRe**. The workflow, informed by earlier literature on WATs, consists of three main steps in the usage of TPAs (and so the data sharing with them): (1) adoption (i.e., the moment when users start to use WATs) [351, 352], (2) adherence (i.e., the period when users continue to use WATs) [353], and (3) abandonment (i.e., the moment when users stop using WATs) [354, 355]. Next, we explain **RePaRe** according to the different stages of the workflow. During the adoption step, a given user contemplates using a TPA, usually

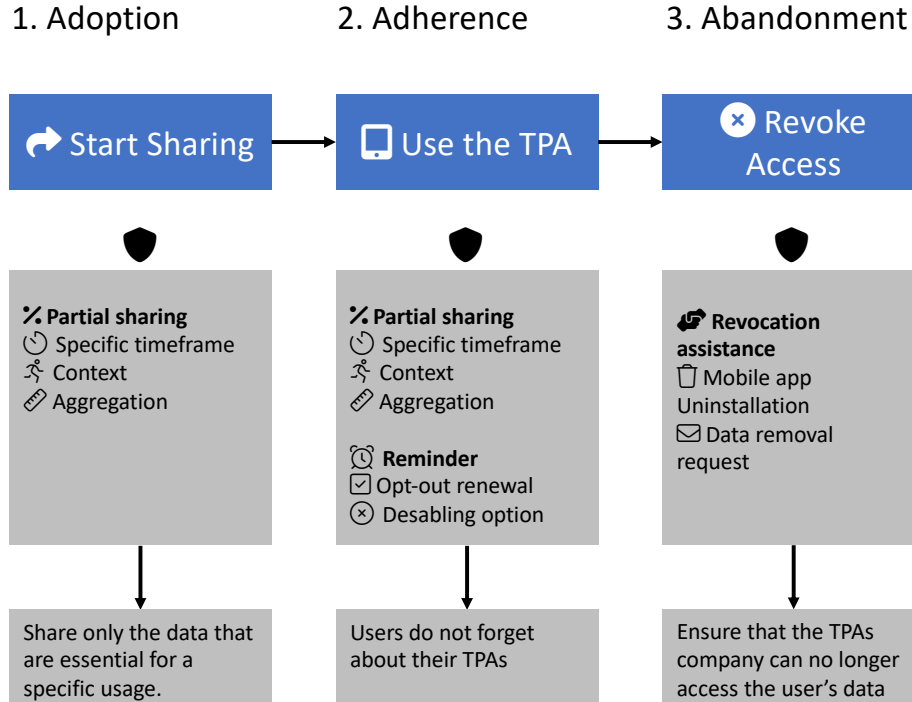


Figure 5.12: (Top) Workflow of WAT-data sharing process with TPAs. (Bottom) RePaRe: an example meta-solution we propose for the workflow.

installing the corresponding app on their phone, and sharing their data with the TPA. To do so, the user usually has the possibility to select the type of data they want to share (e.g., step count, heart rate, activities) and have to agree to share their data, generally by tapping/clicking on an “accept” button. In that step, RePaRe proposes partial sharing (i.e., Feature 1), offering the user different options to share more specific data regarding context and timeframe, and to only share aggregated data. Implementing such a feature early in the process is important as the TPA will have access to all data for a given type as soon as the user accepts to share. The user interface of such a multi-aspect partial sharing feature is out of the scope of our work and should be investigated by future studies. During the adherence step of the data-sharing process, the user passively shares their data to a given TPA by just wearing the TPA and potentially using the corresponding mobile app. RePaRe still offers partial sharing, in this step, as the user may want to modify them, either to share more or less data and adjust the privacy risks. Additionally,

**RePaRe** reminds the user, helping them not forget about the previously granted access to their data, and the fact that they have the possibility to revoke this authorization. In the abandonment step, **RePaRe** assists the user in revoking previously granted access when they uninstall or remove the TPA mobile app from their phone. Lastly, **RePaRe** reminds the user that they can request the corresponding company to delete their data, including personally identifiable information and WAT data, according to GDPR.

Such a meta-solution requires (almost exclusively) the service provider (e.g., Fitbit) to be involved, which makes them the primary stakeholder, allowing users to increase their privacy. This is an advantage for users or any other party (e.g., a legal authority), as they would not need to request new features from multiple third parties, but only from a single entity (the service provider), requiring less effort on their part. Recognizing that the presented meta-solution is just one manifestation of the diverse set of potential design configurations is essential. The participatory approach we took uncovered a range of innovative design features, each with the potential to enhance user privacy in distinct ways. Thus, other combinations and arrangements of these features have the potential to produce equally effective solutions. Having said that, some of the proposed ideas require more consideration. For example, a particularly drastic one would be to limit data storage directly on the WAT or the smartphone. However, this could lead to severe drawbacks for utility, as these devices (especially the WAT) have very limited computing and storage capacities.

## 5.5 Limitations

This work has a few limitations. First, a participants' sample including more Fitbit users and fewer Apple users would have been more representative of the population of WAT users. Indeed, as seen in Chapter 3, whereas Fitbit is one of the most permissive companies regarding TPAs, Apple is more restrictive and its users are more used to a particularly closed environment. Having more participants that use devices from other companies would probably have helped to generate more diverse designs. Another limitation is that the participants evaluated each design as a whole and not the design features or each proposed functionality, individually. Indeed, multiple designs included several different functionalities but were evaluated as a global solution. Therefore, even if this helps us draw general conclusions, the participants' evaluations are not strictly representative of how they would have evaluated each category, whereas, with

the experts, the evaluations were made for each category, individually. As we conducted multiple participatory design sessions, all the designs were not evaluated by the same participants. Moreover, we (the investigators) also decided to evaluate the different categories by regarding the same criteria as the participants did, including adoption, for which an evaluation by WAT users would be more relevant. Therefore, as we are not necessarily representative of the WAT-user population, our evaluation for adoption could be biased. However, the evaluation by the experts of the three other discussed criteria (i.e., feasibility, effectiveness, and usability) was still highly relevant. Finally, we lack qualitative feedback from the participants. Indeed, although we encouraged them to engage in discussion after each presentation, they mostly asked questions to be sure they correctly understood the design, but very few of them emitted remarks about their appreciation of the presented design.

## 5.6 Conclusion

In this chapter, we have described a participatory design study that was conducted with 26 WAT users in order to design new functionalities that could help WAT users better manage data sharing, thereby increasing their privacy.

We have classified the 19 different designs that were proposed by the participants into 7 different design feature categories. Then, we have described and evaluated these categories. We have also compared our experts' evaluations with the participants' evaluations. We have used this information, as well as other protection mechanism ideas that we already proposed in the previous chapters, to develop a general (perfect) solution that would be, in our opinion, highly effective for increasing WAT-users privacy while keeping a decent level of feasibility, usability, and adoption. This "meta-solution" combines three of the seven previously defined features, to which we have added one particular PET that we already proposed and tested in the previous chapters (Chapters 3 and 4).

For future work, we plan to present our previously proposed "meta-solution", called **RePaRe**, in detail to WAT users in order to have precise evaluations and feedback. We also intend to implement and deploy a tool that would enable WAT users to use similar functionalities in order to conduct a longitudinal study on how users can adopt and use such a solution. Such a study would help us improve this solution and to better understand to what extent it would be useful to protect WAT-users privacy. Another interesting study could be conducted on WAT companion-app developers and/or companies in order to

better understand their motivations and to what extent they would consider implementing such functionalities. Finally, we should actively follow the new data-sharing trends in the WAT market, as well as the corresponding functionalities, in order to observe if new privacy-enhancing technologies are indeed implemented and to study potential new data-sharing behavior that could be harmful to privacy.



# Chapter 6

## Conclusion

In this final chapter, we conclude this thesis in two parts. First, we summarize our contributions to the wearable activity tracker security & privacy research field. Then, we discuss the future of this research field by highlighting various ways to complement the research presented in this thesis and to further explore our findings.

### 6.1 Contributions

In this thesis, we have provided key findings about the quantification of the privacy of users of wearable activity trackers by assessing the risks and threats associated with the use of such devices and by proposing countermeasures. After reviewing the current literature about WAT security & privacy in Chapter 2, we highlight, in Chapter 3, the risks (i.e., “the possibility of something bad happening” [356]) related to the usage of WAT devices by assessing the users’ behavior and understanding toward data sharing. Our findings show that a large number of users have poor knowledge of the WAT-data sharing ecosystem and are often are not aware of all the data they share and with whom. Such findings suggest that they could make decisions that, without realizing it, are harmful to their privacy and that they are not fully aware of the issues at stake in their privacy-utility trade-off [37, 134]. We have analyzed, in Chapter 4, a particular threat to the privacy of WAT users by assessing the extent to which an adversary can infer individuals’ personality from their WAT data. In this chapter, we show a significant correlation between this data and personality traits. We have discussed how such inferences can be harmful (e.g., targeted advertising, discrimination) to the users and/or to society



as a whole (e.g., massive political manipulation). Finally, in Chapter 5, we have described the outcome of three participatory design sessions about privacy/transparency enhancing technologies conducted with WAT users, and we have presented and commented on multiple draft designs for improving WAT-data sharing, in the sense that they could help users to better understand it, be more aware of their behavior, and share their data with fine-grained access control; these designs could give them a means to improve their privacy. We then discussed and proposed a meta-solution, called **RePaRe**, composed of the best-evaluated features extracted from the different designs. This solution would, in our opinion, highly improve WAT users' privacy by providing them with new features to help them better manage their data sharing and share data more selectively.

## 6.2 Future Work and Perspectives

In Chapter 3, we have explored the *actual* behavior of WAT users toward data sharing, as well as their understanding of the WAT-data sharing ecosystem. However, we explore their behavior regarding only the entities (e.g., other users, TPAs) to which WAT users share their data. For future work, it could be interesting to further explore the types of data that are shared with TPAs, as well as with other entities (e.g., friends, physicians, insurance, and social networks). Indeed, as shown in Chapter 4, the amount of information that an adversary can infer from WAT data varies depending on the type of data to which they have access. And as attitudes and concerns about data sharing regarding the different types of data and categories of entities WAT users can share their data with were explored in previous work [19], this is not the case of their *actual* behavior.

In Chapter 3, we have presented the results related only to how many TPAs users share their data *at a specific point in time* (i.e., when they answered the questionnaire). However, many of them could have shared their data with far more TPAs before revoking their access. Some users regularly try, for example, new TPAs before revoking the previously granted access. Such behavior can also be harmful, as these TPAs can generally access all the data collected in the past. As a result, it is crucial to study the frequency of such behavior in the population of WAT users by investigating, through longitudinal studies, how users grant and revoke access to their TPAs over time and by putting this into perspective with their actual use of the TPAs.

We could also explore, from a more technical point of view, other types of

inferences. Indeed, in this thesis, we studied how WAT data can be used to infer a user's personality. However, many other personal attributes, such as religious beliefs, political views, and marital status, can probably be inferred from such behavioral and contextual data. This could be either done using a methodology similar to the one described in Chapter 4 or by collecting data on a larger scale. For example, by collecting previously collected WAT-data from users (after obtaining their consent) and ground truth about the personal information, we could try to infer more. Indeed, as the main adversarial model described in this thesis corresponds to companies that can access vastly larger amounts of data than the data used in this thesis, it is crucial to understand that the extent to which the accuracy of the inference of the personal attributes of WAT users could indeed vary on a much larger scale and subsequently to develop adapted mitigation techniques.

Another aspect of information privacy that we did not explore in this dissertation and that could bring valuable knowledge about the privacy of WAT users is to study the extent to which WAT data can be de-anonymized in a large dataset. Indeed, health-data breaches (e.g., hospital databases) are particularly frequent [357, 358] and, even if such data could be anonymized or pseudonymized, it is generally not the case for WAT data that have been shared by the users (e.g., with TPAs). As WAT data are highly related to certain types of health data (e.g., heart rate), it is crucial to understand how an adversary can cross-reference the data of a user to create user profiles containing even more information.

Regarding mitigation techniques, we could design more secure communication protocols for WAT data, from either the WAT to the companion app (i.e., the phone) or from the app to the cloud; as we have seen in Chapter 2 these aspects have multiple vulnerabilities. However, as described in this thesis, we consider that the way WAT users share their data constitutes a risk greater than these security vulnerabilities. Hence, it is important to further study PETs/TETs, as we have in Chapter 5. For example, we could implement and evaluate, on a large scale, the final general solution that we proposed in this manuscript, as well as others. It could also be interesting to open a discussion with WAT developers/companies to evaluate to which extent they would be inclined to develop such functionality in order to help their users to protect their privacy. Although they are one of the main stakeholders related to this topic hence are involved in the multiple issues about privacy, there are particularly few that are studied/solicited by researchers.

Although, in this thesis, we focus on a specific type of wearable device

(i.e., WATs), there exist a large variety of other devices, with similar types of sensors that should be analyzed from the angle of information security & privacy. Indeed, such devices are increasingly part of the daily lives of numerous individuals and are equipped with ever more accurate and diverse sensors. A large number of emerging technologies, such as the recent Apple Vision Pro [359] could be widely adopted in the near future. Therefore, it is crucial to not only conduct “technical” studies related to security & privacy on this type of device (e.g., analyzing communication protocols, highlighting security breaches, personal information inferences) but also to conduct studies with user-centric approaches. This is essential to better understanding all the risks for privacy and for proposing the necessary mitigation methods adapted to the usage that individuals make of these technologies.

# Bibliography

- [1] Statista Research Department. Fitness/activity tracking wrist-wear users worldwide 2018-2027, 2023. URL <https://www.statista.com/forecasts/1314613/worldwide-fitness-or-activity-tracking-wrist-wear-users>.
- [2] Number of connected wearable devices worldwide from 2016 to 2022, 2021. URL <https://www.statista.com/statistics/487291/global-connected-wearable-devices/>.
- [3] Eun Kyoung Choe, Nicole B. Lee, Bongshin Lee, Wanda Pratt, and Julie A. Kientz. Understanding quantified-selfers' practices in collecting and exploring personal data. In *Proc. of the conf. on Human Factors in Computing Systems (CHI)*, 2014. Association for Computing Machinery. doi: 10.1145/2556288.2557372.
- [4] Simon Eberz, Giulio Lovisotto, Andrea Patane, Marta Kwiatkowska, Vincent Lenders, and Ivan Martinovic. When Your Fitness Tracker Betrays You: Quantifying the Predictability of Biometric Features Across Contexts. In *S&P*, 2018. IEEE. doi: 10.1109/SP.2018.00053.
- [5] Anindya Maiti, Murtuza Jadliwala, Jibo He, and Igor Bilogrevic. (Smart)Watch Your Taps: Side-channel Keystroke Inference Attacks Using Smartwatches. In *Proc. of the ACM Int'l Symp. on Wearable Computers (ISWC)*, 2015. ACM. doi: 10.1145/2802083.2808397.
- [6] Anindya Maiti, Oscar Armbruster, Murtuza Jadliwala, and Jibo He. Smartwatch-Based Keystroke Inference Attacks and Context-Aware Protection Mechanisms. In *Proc. of the ACM Asia Conf. on Computer and Communications Security (ASIA CCS)*, 2016. ACM. doi: 10.1145/2897845.2897905.
- [7] Anindya Maiti, Ryan Heard, Mohd Sabra, and Murtuza Jadliwala. Towards Inferring Mechanical Lock Combinations Using Wrist-Wearables As a Side-Channel. In *Proc. of the ACM Conf. on Security and Privacy in Wireless and Mobile Networks (WiSec)*, 2018. ACM. doi: 10.1145/3212480.3212498.
- [8] Hyeokhyen Kwon, Gregory D. Abowd, and Thomas Plötz. Adding structural characteristics to distribution-based accelerometer representations for activity recognition using wearables. In *Ubicomp*, 2018. ACM. doi: 10.1145/3267242.3267258.

- [9] Vishvak S. Murahari and Thomas Plötz. On attention models for human activity recognition. In *Proc. of the ACM Int'l Symp. on Wearable Computers (ISWC)*, 2018. ACM. doi: 10.1145/3267242.3267287.
- [10] Reem Abdel-Salam, Rana Mostafa, and Mayada Hadhood. Human Activity Recognition Using Wearable Sensors: Review, Challenges, Evaluation Benchmark. *Deep Learning for Human Activity Recognition*, 2021. doi: 10.1007/978-981-16-0575-8\_1.
- [11] Emre Ertin, Nathan Stohs, Santosh Kumar, Andrew Raij, Mustafa al'Absi, and Siddharth Shah. AutoSense: unobtrusively wearable sensor suite for inferring the onset, causality, and consequences of stress in the field. In *Proc. of the ACM Conf. on Embedded Networked Sensor Systems (SenSys)*, 2011. ACM Press. doi: 10.1145/2070942.2070970.
- [12] Annamalai Natarajan, Abhinav Parate, Edward Gaiser, Gustavo Angarita, Robert Malison, Benjamin Marlin, and Deepak Ganesan. Detecting cocaine use with wearable electrocardiogram sensors. In *Proc. of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2013. ACM. doi: 10.1145/2493432.2493496.
- [13] Robert P Hirten, Matteo Danieletto, Lewis Tomalin, Katie Hyewon Choi, Eddy 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. 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.
- [14] Wajih Ul Hassan, Saad Hussain, and Adam Bates. Analysis of Privacy Protections in Fitness Tracking Social Networks -or- You can run, but can you hide? In *Proc. of the USENIX Security Symp.*, 2018. USENIX Association.
- [15] Jaron Mink, Amanda Rose Yuile, Uma Pal, Adam J Aviv, and Adam Bates. Users Can Deduce Sensitive Locations Protected by Privacy Zones on Fitness Tracking Apps. In *Proc. of the conf. on Human Factors in Computing Systems (CHI)*, 2022. Association for Computing Machinery. doi: 10.1145/3491102.3502136.
- [16] Karel Dhondt, Victor Le Pochat, and Alexios Voulimeneas. A Run a Day Won't Keep the Hacker Away: Inference Attacks on Endpoint Privacy Zones in Fitness Tracking Social Networks. *Los Angeles*, 2022.
- [17] Alex Hern. Fitness tracking app Strava gives away location of secret US army bases. *The Guardian*, 2018.
- [18] Abdulmajeed Alqhatani and Heather Richter Lipford. "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)*, 2019. USENIX Association.

- [19] Sandra Gabriele and Sonia Chiasson. Understanding Fitness Tracker Users' Security and Privacy Knowledge, Attitudes and Behaviours. In *Proc. of the CHI Conf. on Human Factors in Computing Systems (CHI)*, 2020. Association for Computing Machinery. doi: 10.1145/3313831.3376651.
- [20] Yuan Zhong, Nicholas Jing Yuan, Wen Zhong, Fuzheng Zhang, and Xing Xie. You Are Where You Go: Inferring Demographic Attributes from Location Check-ins. In *Proc. of the ACM Int'l Conf. on Web search and data mining (WSDM)*, 2015. doi: 10.1145/2684822.2685287.
- [21] Alan Mislove, Bimal Viswanath, Krishna P. Gummadi, and Peter Druschel. You are who you know: inferring user profiles in online social networks. In *Proc. of the ACM Int'l Conf. on Web search and data mining (WSDM)*, 2010. doi: 10.1145/1718487.1718519.
- [22] Antoine Boutet and Sébastien Gambs. Inspect What Your Location History Reveals About You: Raising user awareness on privacy threats associated with disclosing his location data. In *Proc. of the ACM Int'l Conf. on Information and Knowledge Management (CIKM)*, 2019. doi: 10.1145/3357384.3357837.
- [23] Fitness trackers chase after the corporate market, 2014. URL <http://www.washingtonpost.com/blogs/on-leadership/wp/2014/12/18/fitness-trackers-chase-after-the-corporate-market/>.
- [24] Ramon Llamas, Jitesh Ubrani, and Michael Shirer. Xiaomi and Apple Tie for the Top Position as the Wearables Market Swells 17.9% During the First Quarter, According to IDC, 2017. URL <https://www.businesswire.com/news/home/20170605005391/en/Xiaomi-and-Apple-Tie-for-the-Top-Position-as-the-Wearables-Market-Swells-17.9-During-the-First-Quarter-According-to-IDC>.
- [25] André Henriksen, Martin Haugen Mikalsen, Ashenafi Zebene Woldaregay, Miroslav Muzny, Gunnar Hartvigsen, Laila Arnesdatter Hopstock, and Sameline Grimsgaard. Using Fitness Trackers and Smartwatches to Measure Physical Activity in Research: Analysis of Consumer Wrist-Worn Wearables. *Journal of Medical Internet Research*, 2018. doi: 10.2196/jmir.9157.
- [26] Christopher Rowl. With fitness trackers in the workplace, bosses can monitor your every step - and possibly more, 2019. URL [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](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).
- [27] Jon Porter. Google completes purchase of Fitbit, 2021. URL <https://www.theverge.com/2021/1/14/22188428/google-fitbit-acquisition-completed-approved>.
- [28] Giles Bruce. Google parent Alphabet's health insurance company grew nearly sixfold in '22, 2023. URL <https://www.beckershospitalreview.com/disruptors/>

- google-parent-alphabets-health-insurance-company-grew-nearly-sixfold-in-22.html.
- [29] Jess Weatherbed. All Fitbit users will require a Google account by 2025, 2022. URL <https://www.theverge.com/2022/9/26/23372438/fitbit-changes-update-google-account-new-2025>.
- [30] Melanie Ehrenkranz. The Plan to Use Fitbit Data to Stop Mass Shootings Is One of the Scariest Proposals Yet, 2019. URL <https://gizmodo.com/the-plan-to-use-fitbit-data-to-stop-mass-shootings-is-o-1837710691>.
- [31] Office of the United Nations High Commissioner for Human Rights (OHCHR). Privacy and surveillance. URL <https://www.ohchr.org/en/taxonomy/term/767>.
- [32] Alan Furman Westin. *Privacy and freedom*. IG Publishing, new edition edition, 1967.
- [33] The European Parliament and the Council of The European Union. REGULATION (EU) 2016/679 of the European Parliament and the Council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official J. Eur. Union*, 2016.
- [34] Privacy Impact Assessment, . URL <https://gdpr-info.eu/issues/privacy-impact-assessment/>.
- [35] Aaron Yi Ding, Gianluca Limon De Jesus, and Marijn Janssen. Ethical hacking for boosting IoT vulnerability management: a first look into bug bounty programs and responsible disclosure. In *Proceedings of the Eighth International Conference on Telecommunications and Remote Sensing*, 2019. Association for Computing Machinery. doi: 10.1145/3357767.3357774.
- [36] Marleen Weulen Kranenbarg, Thomas J. Holt, and Jeroen van der Ham. Don't shoot the messenger! A criminological and computer science perspective on coordinated vulnerability disclosure. *Crime Science*, 2018. doi: 10.1186/s40163-018-0090-8.
- [37] He Li, Jing Wu, Yiwen Gao, and Yao Shi. Examining individuals' adoption of health-care wearable devices: An empirical study from privacy calculus perspective. *International Journal of Medical Informatics*, 2016. doi: 10.1016/j.ijmedinf.2015.12.010.
- [38] Teresa Robertson Ishii and Philip Atkins. Essential vs. Accidental Properties. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2020 edition, 2020.
- [39] Moritz Becker, Andreas Kolbeck, Christian Matt, and Thomas Hess. Understanding the Continuous Use of Fitness Trackers: A Thematic Analysis. In *PACIS Proc.*, 2017.
- [40] Matthew B. Hoy. Personal Activity Trackers and the Quantified Self. *Medical Reference Services Quarterly*, 2016. doi: 10.1080/02763869.2016.1117300.

- [41] Zablón Pingo and Bhuvá Narayan. “My smartwatch told me to see a sleep doctor”: a study of activity tracker use. *Online Information Review*, 2019. doi: 10.1108/OIR-04-2018-0115.
- [42] Kavous Salehzadeh Niksirat, Lev Velykoivanenko, Noé Zufferey, Mauro Cherubini, Kévin Huguenin, and Mathias Humbert. *Wearable Activity Trackers: A Survey on Utility, Privacy, and Security*. 2023.
- [43] Blaine Cook and Chris Messina. OAuth 2.0 — OAuth, 2012. URL <https://oauth.net/2/>.
- [44] apple. HealthKit | Apple Developer Documentation. URL <https://developer.apple.com/documentation/healthkit>.
- [45] Mohammad Hossein Jarrahi, Nicci Gafinowitz, and Grace Shin. Activity trackers, prior motivation, and perceived informational and motivational affordances. *Personal and Ubiquitous Computing*, 2018. doi: 10.1007/s00779-017-1099-9.
- [46] Lev Velykoivanenko, Kavous Salehzadeh Niksirat, Noé Zufferey, Mathias Humbert, Kévin Huguenin, and Mauro Cherubini. 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*, 2021. doi: 10.1145/3494960.
- [47] WeWard - The mobile app that motivates you to walk, 2023. URL <https://en.weward.fr/>.
- [48] Actifit - Rewarding Your Everyday Activity, . URL <https://actifit.io>.
- [49] fitcoin. Fitcoin | Much more than just a fitness cryptocurrency. URL <https://fitcoin.io/>.
- [50] Andrei Kazlouski, Thomas Marchioro, and Evangelos Markatos. I just wanted to track my steps! Blocking unwanted traffic of Fitbit devices. In *Proceedings of the International Conference on the Internet of Things*, 2023. Association for Computing Machinery. doi: 10.1145/3567445.3567457.
- [51] David Curry. Fitbit Revenue and Usage Statistics (2020), 2023. URL <https://www.businessofapps.com/data/fitbit-statistics/>.
- [52] Mehdi Nobakht, Yulei Sui, Aruna Seneviratne, and Wen Hu. PGFit: Static permission analysis of health and fitness apps in IoT programming frameworks. *Journal of Network and Computer Applications*, 2020. doi: 10.1016/j.jnca.2019.102509.
- [53] William Enck, Peter Gilbert, Byung-Gon Chun, Landon P. Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N. Sheth. TaintDroid: an information flow tracking system for real-time privacy monitoring on smartphones. *Communications of the ACM*, 2014. doi: 10.1145/2494522.



- [54] Anna Maria Mandalari, Daniel J. Dubois, Roman Kolcun, Muhammad Talha Paracha, Hamed Haddadi, and David Choffnes. Blocking Without Breaking: Identification and Mitigation of Non-Essential IoT Traffic. *Proc. on Privacy Enhancing Technologies (PoPETs)*, 2021. doi: 10.2478/popets-2021-0075.
- [55] Georgios Kontaxis, Michalis Polychronakis, and Evangelos P. Markatos. Minimizing information disclosure to third parties in social login platforms. *International Journal of Information Security*, 2012. doi: 10.1007/s10207-012-0173-6.
- [56] Ilaria Torre, Odnan Ref Sanchez, Frosina Koceva, and Giovanni Adorni. Supporting users to take informed decisions on privacy settings of personal devices. *Personal and Ubiquitous Computing*, 2018. doi: 10.1007/s00779-017-1068-3.
- [57] Shumin Guo and Keke Chen. Mining Privacy Settings to Find Optimal Privacy-Utility Tradeoffs for Social Network Services. In *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing*, 2012. IEEE. doi: 10.1109/SocialCom-PASSAT.2012.22.
- [58] Rachna Dhamija, J. D. Tygar, and Marti Hearst. Why phishing works. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*, 2006. ACM. doi: 10.1145/1124772.1124861.
- [59] Yuting Liao. Sharing Personal Health Information on Social Media: Balancing Self-presentation and Privacy. In *Proc. of the Int'l Conf. on Social Media and Society (SM-Society)*, 2019. Association for Computing Machinery. doi: 10.1145/3328529.3328560.
- [60] Stefan Schneegass, Romina Poguntke, and Tonja Machulla. Understanding the Impact of Information Representation on Willingness to Share Information. In *Proc. of the conf. on Human Factors in Computing Systems (CHI)*, 2019. Association for Computing Machinery. doi: 10.1145/3290605.3300753.
- [61] Noé Zufferey, Kavous Salehzadeh Niksirat, Mathias Humbert, and Kévin Huguenin. "Revoked just now!" Users' Behaviors Toward Fitness-Data Sharing with Third-Party Applications. *Proc. on Privacy Enhancing Technologies (PoPETs)*, 2023. doi: 10.56553/popets-2023-0004.
- [62] Robert R. McCrae, Paul T. Costa, Jr., and Thomas A. Martin. The NEO-PI-3: A More Readable Revised NEO Personality Inventory. *Journal of Personality Assessment*, 2005. doi: 10.1207/s15327752jpa8403\_05.
- [63] Noé Zufferey, Mathias Humbert, and Kévin Huguenin. Watch your Watch: Inferring Personality Traits from Wearable Activity Trackers. *Proc. of the USENIX Security Symposium*, 2023.
- [64] Barbara Kitchenham, O. Pearl Brereton, David Budgen, Mark Turner, John Bailey, and Stephen Linkman. Systematic literature reviews in software engineering – A systematic literature review. *Information and Software Technology*, 2009. doi: 10.1016/j.infsof.2008.09.009.

- [65] EM Aromataris and Z Munn. JBI Manual for Evidence Synthesis, 2020. URL <https://jbi-global-wiki.refined.site/space/MANUAL>.
- [66] Yetong Cao, Fan Li, Huijie Chen, Xiaochen liu, Li Zhang, and Yu Wang. Guard Your Heart Silently: Continuous Electrocardiogram Waveform Monitoring with Wrist-Worn Motion Sensor. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2022. doi: 10.1145/3550307.
- [67] Jingwen Zhang, Dingwen Li, Ruixuan Dai, Heidy Cos, Gregory A. Williams, Lacey Raper, Chet W. Hammill, and Chenyang Lu. Predicting Post-Operative Complications with Wearables: A Case Study with Patients Undergoing Pancreatic Surgery. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2022. doi: 10.1145/3534578.
- [68] Gabriel Guo, Hanbin Zhang, Liuyi Yao, Huining Li, Chenhan Xu, Zhengxiong Li, and Wenyao Xu. 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*, 2021. doi: 10.1145/3494970.
- [69] Robert P Hirten, Matteo Danieletto, Lewis Tomalin, Katie Hyewon Choi, Micol Zweig, Eddy Golden, Sparshdeep Kaur, Drew Helmus, Anthony Biello, Renata Pyzik, Alexander Charney, Riccardo Miotto, Benjamin S Glicksberg, Matthew Levin, Ismail Nabeel, Judith Aberg, David Reich, Dennis Charney, Erwin P Bottinger, Laurie Keefer, Mayte Suarez-Farinas, Girish N Nadkarni, and Zahi A Fayad. Use of Physiological Data From a Wearable Device to Identify SARS-CoV-2 Infection and Symptoms and Predict COVID-19 Diagnosis: Observational Study. *Journal of Medical Internet Research*, 2021. doi: 10.2196/26107.
- [70] Daniel A. Adler, Vincent W.-S. Tseng, Gengmo Qi, Joseph Scarpa, Srijan Sen, and Tanzeem Choudhury. Identifying Mobile Sensing Indicators of Stress-Resilience. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2021. doi: 10.1145/3463528.
- [71] Ruixuan Dai, Thomas Kannampallil, Jingwen Zhang, Nan Lv, Jun Ma, and Chenyang Lu. Multi-Task Learning for Randomized Controlled Trials: A Case Study on Predicting Depression with Wearable Data. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2022. doi: 10.1145/3534591.
- [72] Sangwon Bae, Anind K. Dey, and Carissa A. Low. Using passively collected sedentary behavior to predict hospital readmission. In *UbiComp*, 2016. Association for Computing Machinery. doi: 10.1145/2971648.2971750.
- [73] Laleh Jalali and Ramesh Jain. Building health persona from personal data streams. In *Proc. of the ACM Int'l workshop on Personal data meets distributed multimedia (PDM)*, 2013. Association for Computing Machinery. doi: 10.1145/2509352.2509400.

- [74] Jamie A Ward, Daniel Richardson, Guido Orgs, Kelly Hunter, and Antonia Hamilton. Sensing interpersonal synchrony between actors and autistic children in theatre using wrist-worn accelerometers. In *Proc. of the ACM Int'l Symp. on Wearable Computers*, 2018. ACM. doi: 10.1145/3267242.3267263.
- [75] Mohammad Tahaei, Julia Bernd, and Awais Rashid. Privacy, Permissions, and the Health App Ecosystem: A Stack Overflow Exploration. In *Proceedings of the 2022 European Symposium on Usable Security*, 2022. Association for Computing Machinery. doi: 10.1145/3549015.3555669.
- [76] Alireza Abedin, Mahsa Ehsanpour, Qinfeng Shi, Hamid RezaTofighi, and Damith C. Ranasinghe. Attend and Discriminate: Beyond the State-of-the-Art for Human Activity Recognition Using Wearable Sensors. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2021. doi: 10.1145/3448083.
- [77] Cihang Liu, Lan Zhang, Zongqian Liu, Kebin Liu, Xiangyang Li, and Yunhao Liu. Lasagna: towards deep hierarchical understanding and searching over mobile sensing data. In *Proc. of the Annual Int'l Conf. on Mobile Computing and Networking*, 2016. ACM. doi: 10.1145/2973750.2973752.
- [78] Manuel Dietrich and Kristof van Laerhoven. A typology of wearable activity recognition and interaction. In *Proc. of the Int'l Workshop on Sensor (based Activity Recognition and Interaction)*, 2015. Association for Computing Machinery. doi: 10.1145/2790044.2790048.
- [79] Yasmin F van Kasteren, Lua Perimal-Lewis, and Anthony Maeder. Detecting short-duration ambulatory episodes in Fitbit data. In *Proc. of the Australasian Computer Science Week MultiConf. (ACSW)*, 2018. Association for Computing Machinery. doi: 10.1145/3167918.3167954.
- [80] Igor Lopes de Faria and Vaninha Vieira. A Comparative Study on Fitness Activity Recognition. In *Proc. of the Brazilian Symp. on Multimedia and the Web (WebMedia)*, 2018. Association for Computing Machinery. doi: 10.1145/3243082.3267452.
- [81] Edison Thomaz, Irfan Essa, and Gregory D. Abowd. A practical approach for recognizing eating moments with wrist-mounted inertial sensing. In *UbiComp*, 2015. ACM Press. doi: 10.1145/2750858.2807545.
- [82] Gary M. Weiss, Jessica L. Timko, Catherine M. Gallagher, Kenichi Yoneda, and Andrew J. Schreiber. Smartwatch-based activity recognition: A machine learning approach. In *Proc. of the IEEE-EMBS Int'l Conf. on Biomedical and Health Informatics (BHI)*, 2016. IEEE. doi: 10.1109/BHI.2016.7455925.
- [83] Joan-Isaac Biel, Nathalie Martin, David Labbe, and Daniel Gatica-Perez. Bites'n'Bits: Inferring Eating Behavior from Contextual Mobile Data. *Proc. of the Conf. on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2018. doi: 10.1145/3161161.

- [84] Mario A. Gutierrez, Michelle L. Fast, Anne H. Ngu, and Byron J. Gao. Real-Time Prediction of Blood Alcohol Content Using Smartwatch Sensor Data. In Xiaolong Zheng, Daniel Dajun Zeng, Hsinchun Chen, and Scott J. Leischow, editors, *Smart Health (Lecture Notes in Computer Science)*, 2016.
- [85] Muhammad Shoailb, Hans Scholten, Paul J. M. Havinga, and Ozlem Durmaz Incel. A hierarchical lazy smoking detection algorithm using smartwatch sensors. In *Proc. of the IEEE Int'l Conf. on e-Health Networking, Applications and Services (Healthcom)*, 2016. IEEE. doi: 10.1109/HealthCom.2016.7749439.
- [86] Sheng Shen, He Wang, and Romit Roy Choudhury. I Am a Smartwatch and I Can Track My User's Arm. In *Proc. of the Int'l Conf. on Mobile Systems, Applications, and Services*, 2016. ACM. doi: 10.1145/2906388.2906407.
- [87] Juhi Ranjan and Kamin Whitehouse. Object hallmarks: identifying object users using wearable wrist sensors. In *UbiComp*, 2015. Association for Computing Machinery. doi: 10.1145/2750858.2804263.
- [88] N. Vinayaga-Sureshkanth, A. Maiti, M. Jadliwala, K. Crager, J. He, and H. Rathore. A Practical Framework for Preventing Distracted Pedestrian-Related Incidents Using Wrist Wearables. *IEEE Access*, 2018. doi: 10.1109/ACCESS.2018.2884669.
- [89] Jiayu Li, Zhiyu He, Yumeng Cui, Chenyang Wang, Chong Chen, Chun Yu, Min Zhang, Yiqun Liu, and Shaoping Ma. Towards Ubiquitous Personalized Music Recommendation with Smart Bracelets. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2022. doi: 10.1145/3550333.
- [90] Jack Sturgess, Simon Eberz, Ivo Sluganovic, and Ivan Martinovic. Inferring User Height and Improving Impersonation Attacks in Mobile Payments using a Smartwatch. In *Proc. of the IEEE Int'l Conf. on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*, 2022. IEEE. doi: 10.1109/PerComWorkshops53856.2022.9767287.
- [91] Yafei Wang, Ingmar Weber, and Prasenjit Mitra. Quantified Self Meets Social Media: Sharing of Weight Updates on Twitter. In *Proc. of the Int'l Conf. on Digital Health Conf. (DH)*, 2016. Association for Computing Machinery. doi: 10.1145/2896338.2896363.
- [92] Ulku Meteriz, Necip Fazil Yildiran, Joongheon Kim, and David Mohaisen. Understanding the Potential Risks of Sharing Elevation Information on Fitness Applications. In *Proc. of the IEEE Int'l Conf. on Distributed Computing Systems (ICDCS)*, 2020. doi: 10.1109/ICDCS47774.2020.00063.
- [93] Christoph Amma, Marcus Georgi, and Tanja Schultz. Airwriting: Hands-Free Mobile Text Input by Spotting and Continuous Recognition of 3d-Space Handwriting with Inertial Sensors. In *Int'l Symp. on Wearable Computers*, 2012. doi: 10.1109/ISWC.2012.21.

- [94] Chao Xu, Parth H. Pathak, and Prasant Mohapatra. Finger-writing with Smartwatch: A Case for Finger and Hand Gesture Recognition using Smartwatch. In *Proc. of the Int'l Workshop on Mobile Computing Systems and Applications (HotMobile)*, 2015. Association for Computing Machinery. doi: 10.1145/2699343.2699350.
- [95] L. Ardüser, P. Bissig, P. Brandes, and R. Wattenhofer. Recognizing text using motion data from a smartwatch. In *Proc. Pervasive Computing and Communication Workshops (PerCom Workshops)*, 2016. doi: 10.1109/PERCOMW.2016.7457172.
- [96] Qingxin Xia, Feng Hong, Yuan Feng, and Zhongwen Guo. MotionHacker: Motion sensor based eavesdropping on handwriting via smartwatch. In *Proc. of the IEEE INFOCOM - IEEE Conf. on Computer Communications Workshops (INFOCOM WK-SHPS)*, 2018. doi: 10.1109/INFOCOMW.2018.8406879.
- [97] Raveen Wijewickrama, Anindya Maiti, and Murtuza Jadliwala. deWristified: handwriting inference using wrist-based motion sensors revisited. In *Proc. of the ACM Conf. on Security and Privacy in Wireless and Mobile Networks (WiSec)*, 2019. Association for Computing Machinery. doi: 10.1145/3317549.3319722.
- [98] Jessica Vitak, Yuting Liao, Priya Kumar, Michael Zimmer, and Katherine Kritikos. Privacy Attitudes and Data Valuation Among Fitness Tracker Users. In Gobinda Chowdhury, Julie McLeod, Val Gillet, and Peter Willett, editors, *Transforming Digital Worlds (Lecture Notes in Computer Science)*, 2018. Springer International Publishing. doi: 10.1007/978-3-319-78105-1\_27.
- [99] Liezel Cilliers. Wearable devices in healthcare: Privacy and information security issues. *Health Information Management Journal*, 2020. doi: 10.1177/1833358319851684.
- [100] Philip Dahlstrøm, Erlend Fauchald, Benjamin Fimreite', and Miriam Lillebo. Users knowledge and attitudes towards data collection in activity trackers.
- [101] Alexander Wieneke, Christiane Lehrer, Raphael Zeder, and Reinhard Jung. Privacy-related Decision-Making in the Context of Wearable Use. *PACIS 2016 Proceedings*, 2016.
- [102] B. Lowens, V. G. Motti, and K. Caine. Wearable Privacy: Skeletons in The Data Closet. In *Proc. of the IEEE Int'l Conf. on Healthcare Informatics (ICHI)*, 2017. doi: 10.1109/ICHI.2017.29.
- [103] Laura Burbach, Chantal Lidynia, Philipp Brauner, and Martina Ziefle. Data protectors, benefit maximizers, or facts enthusiasts: Identifying user profiles for life-logging technologies. *Computers in Human Behavior*, 2019. doi: 10.1016/j.chb.2019.05.004.
- [104] Chris Xiaoxuan Lu, Bowen Du, Hongkai Wen, Sen Wang, Andrew Markham, Ivan Martinovic, Yiran Shen, and Niki Trigoni. Snoopy: Sniffing Your Smartwatch Passwords via Deep Sequence Learning. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2018. doi: 10.1145/3161196.

- [105] Jamie Pinchot and Donna Cellante. Privacy Concerns and Data Sharing Habits of Personal Fitness Information Collected via Activity Trackers. *Journal of Information Systems Applied Research*, 2021.
- [106] Sara Boysen, Barbara Hewitt, David Gibbs, and Alexander McLeod. Refining the Threat Calculus of Technology Threat Avoidance Theory. *Communications of the Association for Information Systems*, 2019. doi: 10.17705/1CAIS.04505.
- [107] Nina Gerber, Benjamin Reinheimer, and Melanie Volkamer. Investigating People’s Privacy Risk Perception. *Proc. on Privacy Enhancing Technologies (PoPETs)*, 2019. doi: 10.2478/popets-2019-0047.
- [108] Deborah Lupton. “Sharing Is Caring:” Australian Self-Trackers’ Concepts and Practices of Personal Data Sharing and Privacy. *Frontiers in Digital Health*, 2021. doi: 10.3389/fdgth.2021.649275.
- [109] Emilee Rader and Janine Slaker. The importance of visibility for folk theories of sensor data. 2017.
- [110] Anubha Mishra, Lori Baker-Eveleth, Prachi Gala, and Julia Stachofsky. Factors influencing actual usage of fitness tracking devices: Empirical evidence from the UTAUT model. *Health Marketing Quarterly*, 2023. doi: 10.1080/07359683.2021.1994170.
- [111] Krutheeka Baskaran, Vijayan Sugumaran, and Saji K Mathew. Are You Coping or Coping Out? Wearable Users’ Information Privacy Perspective. 2020.
- [112] Riccardo Reith, Christoph Buck, Torsten Eymann, and Bettina Lis. Integrating Privacy Concerns Into the Unified Theory of Acceptance and Use of Technology to Explain the Adoption of Fitness Trackers. *International Journal of Innovation and Technology Management*, 2020. doi: 10.1142/S0219877020500492.
- [113] Michael Zimmer, Priya Kumar, Jessica Vitak, Yuting Liao, and Katie Chamberlain Kritikos. ‘There’s nothing really they can do with this information’: unpacking how users manage privacy boundaries for personal fitness information. *Information, Communication & Society*, 2020. doi: 10.1080/1369118X.2018.1543442.
- [114] Seoyoung Kim, Arti Thakur, and Juho Kim. Understanding Users’ Perception Towards Automated Personality Detection with Group-specific Behavioral Data. In *Proc. of the conf. on Human Factors in Computing Systems (CHI)*, 2020. Association for Computing Machinery. doi: 10.1145/3313831.3376250.
- [115] Miikael Lehto and Martti Miikael. Health Information Privacy of Activity Trackers. In *ECCWS European Conf. on Cyber Warfare and Security*, 2017.
- [116] Chantal Lidynia, Philipp Brauner, and Martina Ziefle. A Step in the Right Direction – Understanding Privacy Concerns and Perceived Sensitivity of Fitness Trackers. In Tareq Ahram and Christianne Falcão, editors, *Advances in Human Factors in Wearable Technologies and Game Design (Advances Intelligent Systems and Computing)*, 2018. Springer International Publishing. doi: 10.1007/978-3-319-60639-2\_5.

- [117] Hyunsoo Lee, Soowon Kang, and Uichin Lee. Understanding Privacy Risks and Perceived Benefits in Open Dataset Collection for Mobile Affective Computing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2022. doi: 10.1145/3534623.
- [118] Angeliki Aktypi, Jason R.C. Nurse, and Michael Goldsmith. Unwinding Ariadne’s Identity Thread: Privacy Risks with Fitness Trackers and Online Social Networks. In *Proc. of the Multimedia Privacy and Security (MPS)*, 2017. Association for Computing Machinery. doi: 10.1145/3137616.3137617.
- [119] Moritz Becker. Understanding Users’ Health Information Privacy Concerns for Health Wearables. 2018. doi: 10.24251/HICSS.2018.413.
- [120] Nanna Gorm and Irina Shklovski. Sharing Steps in the Workplace: Changing Privacy Concerns Over Time. In *Proc. of the conf. on Human Factors in Computing Systems (CHI)*, 2016. Association for Computing Machinery. doi: 10.1145/2858036.2858352.
- [121] Nina Gerber, Paul Gerber, and Melanie Volkamer. Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior. *Computers & Security*, 2018. doi: 10.1016/j.cose.2018.04.002.
- [122] Christian Matt, Moritz Becker, Andreas Kolbeck, and Thomas Hess. Continuously Healthy, Continuously Used? – A Thematic Analysis of User Perceptions on Consumer Health Wearables. *Pacific Asia Journal of the Association for Information Systems*, 2019. doi: 10.17705/1pais.11105.
- [123] Jason Orlosky, Onyeka Ezenwoye, Heather Yates, and Gina Besenyi. A Look at the Security and Privacy of Fitbit as a Health Activity Tracker. In *Proc. of the ACM Southeast Conf. (ACM SE)*, 2019. Association for Computing Machinery. doi: 10.1145/3299815.3314468.
- [124] Xinru Page, Paritosh Bahirat, Muhammad I. Safi, Bart P. Knijnenburg, and Pamela Wisniewski. The Internet of What? Understanding Differences in Perceptions and Adoption for the Internet of Things. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2018. doi: 10.1145/3287061.
- [125] Marco Furini, Silvia Mirri, Manuela Montangero, and Catia Prandi. Can IoT Wearable Devices Feed Frugal Innovation? In *Proc. of the Workshop on Experiences with the Design and Implementation of Frugal Smart Objects (FRUGALTHINGS)*, 2020. Association for Computing Machinery. doi: 10.1145/3410670.3410861.
- [126] Zablon Pingo and Bhuvan Narayan. Users’ Responses to Privacy Issues with the Connected Information Ecologies Created by Fitness Trackers. In Milena Dobрева, Anika Hinze, and Maja Žumer, editors, *Maturity and Innovation in Digital Libraries (Lecture Notes in Computer Science)*, 2018. Springer International Publishing. doi: 10.1007/978-3-030-04257-8\_25.

- [127] Aylin Ilhan and Kaja J. Fietkiewicz. Data privacy-related behavior and concerns of activity tracking technology users from Germany and the USA. *Aslib Journal of Information Management*, 2020. doi: 10.1108/AJIM-03-2020-0067.
- [128] Kaja Fietkiewicz and Aylin Ilhan. Fitness Tracking Technologies: Data Privacy Doesn't Matter? The (Un)Concerns of Users, Former Users, and Non-Users. 2020. doi: 10.24251/HICSS.2020.421.
- [129] Nico Ebert, Kurt Alexander Ackermann, and Peter Heinrich. Does Context in Privacy Communication Really Matter? A Survey on Consumer Concerns and Preferences. In *Proc. of the conf. on Human Factors in Computing Systems (CHI)*, 2020. Association for Computing Machinery. doi: 10.1145/3313831.3376575.
- [130] France Bélanger, Robert E. Crossler, and John Correia. Privacy Maintenance in Self-Digitization: The Effect of Information Disclosure on Continuance Intentions. *ACM SIGMIS Database: the DATABASE for Advances in Information Systems*, 2021. doi: 10.1145/3462766.3462769.
- [131] Krutheeka Baskaran and Saji K. Mathew. Understanding Coping Intentions of Fitness Tracker Users: An Empirical Investigation Using Fear Appeals. *International Journal of Human-Computer Interaction*, 2022. doi: 10.1080/10447318.2022.2124358.
- [132] Krutheeka Baskaran and Saji K. Mathew. Danger vs Fear: An Empirical Study on Wearable Users' Privacy Coping. In *Proc. of the Computers and People Research Conf. (SIGMIS (CPR))*, 2020. Association for Computing Machinery. doi: 10.1145/3378539.3393856.
- [133] Krutheeka Baskaran, Vijayan Sugumaran, and Saji K. Mathew. What Do I Do? Uncovering Fitness Tracker Users' Privacy Coping Strategy. *AMCIS 2021 Proceedings*, 2021.
- [134] Susan B. Barnes. A privacy paradox: Social networking in the United States. *First Monday*, 2006. doi: 10.5210/fm.v11i9.1394.
- [135] E. Schomakers, C. Lidynia, and M. Ziefle. Listen to My Heart? How Privacy Concerns Shape Users' Acceptance of e-Health Technologies. In *Int'l Conf. on Wireless and Mobile Computing, Networking and Communications (WiMob)*, 2019. doi: 10.1109/WiMOB.2019.8923448.
- [136] Karthik S. Bhat and Neha Kumar. Sociocultural Dimensions of Tracking Health and Taking Care. *Proceedings of the ACM on Human-Computer Interaction*, 2020. doi: 10.1145/3415200.
- [137] Moritz Becker, Christian Matt, and Thomas Hess. It's Not Just About the Product: How Persuasive Communication Affects the Disclosure of Personal Health Information. *ACM SIGMIS Database: the DATABASE for Advances in Information Systems*, 2020. doi: 10.1145/3380799.3380804.



- [138] Humira Ehrari, Frank Ulrich, and Henning Andersen. Concerns and Trade-offs in Information Technology Acceptance: The Balance between the Requirement for Privacy and the Desire for Safety. *Communications of the Association for Information Systems*, 2020. doi: 10.17705/1CAIS.04711.
- [139] Helen Nissenbaum. Privacy as contextual integrity. *HeinOnline*, 2004.
- [140] Laura Calloway, Hilda Hadan, Shakthidhar Gopavaram, Shrirang Mare, and L. Jean Camp. Privacy in Crisis: Participants' Privacy Preferences for Health and Marketing Data during a Pandemic. In *Proc. of the Workshop on Privacy in the Electronic Society (WPES)*, 2020. Association for Computing Machinery. doi: 10.1145/3411497.3420223.
- [141] Anastasia Kuzminykh and Edward Lank. How Much is Too Much? Understanding the Information Needs of Parents of Young Children. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2019. doi: 10.1145/3328923.
- [142] Mikkel S. Jørgensen, Frederik K. Nissen, Jeni Paay, Jesper Kjeldskov, and Mikael B. Skov. Monitoring children's physical activity and sleep: a study of surveillance and information disclosure. In *Proc. of the Australian Conf. on Computer (Human Interaction)*, 2016. Association for Computing Machinery. doi: 10.1145/3010915.3010936.
- [143] Qingyang Li, Clara Caldeira, Daniel A. Epstein, and Yunan Chen. Supporting Caring among Intergenerational Family Members through Family Fitness Tracking. In *Proc. of the Conf. on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, 2020. Association for Computing Machinery. doi: 10.1145/3421937.3422018.
- [144] Kyrill Potapov and Paul Marshall. LifeMosaic: co-design of a personal informatics tool for youth. In *Proc. of the Interaction Design and Children Conf. (IDC)*, 2020. Association for Computing Machinery. doi: 10.1145/3392063.3394429.
- [145] Roxanne Leitão. Anticipating Smart Home Security and Privacy Threats with Survivors of Intimate Partner Abuse. In *Proc. of the Designing Interactive Systems Conf. (DIS)*, 2019. Association for Computing Machinery. doi: 10.1145/3322276.3322366.
- [146] C. Braghin, S. Cimato, and A. Della Libera. Are mHealth Apps Secure? A Case Study. In *Proc. of the IEEE Annual Computer Software and Applications Conf. (COMPSAC)*, 2018. doi: 10.1109/COMPSAC.2018.10253.
- [147] Greig Paul and James Irvine. Privacy Implications of Wearable Health Devices. In *Proc. of the Int'l Conf. on Security of Information and Networks (SIN)*, 2014. Association for Computing Machinery. doi: 10.1145/2659651.2659683.
- [148] Munyaradzi Katurura and Liezel Cilliers. Privacy in wearable health devices: How does POPIA measure up? In *Kalpa Publications in Computing*, 2019. doi: 10.29007/qsp7.

- [149] Luke Hutton, Blaine A Price, Ryan Kelly, Ciaran McCormick, Arosha K Bandara, Tally Hatzakis, Maureen Meadows, and Bashar Nuseibeh. Assessing the Privacy of mHealth Apps for Self-Tracking: Heuristic Evaluation Approach. *JMIR mHealth and uHealth*, 2018. doi: 10.2196/mhealth.9217.
- [150] Joshua Gluck, Florian Schaub, Amy Friedman, Hana Habib, Norman Sadeh, Lorie Faith Cranor, and Yuvraj Agarwal. How Short Is Too Short? Implications of Length and Framing on the Effectiveness of Privacy Notices. 2016.
- [151] Wentao Guo, Jay Rodolitz, and Eleanor Birrell. Poli-see: An Interactive Tool for Visualizing Privacy Policies. In *Proc. of the Workshop on Privacy in the Electronic Society (WPES)*, 2020. Association for Computing Machinery. doi: 10.1145/3411497.3420221.
- [152] Olha Drozd and Sabrina Kirrane. Privacy CURE: Consent Comprehension Made Easy. In Marko Hölbl, Kai Rannenberg, and Tatjana Welzer, editors, *ICT Systems Security and Privacy Protection (IFIP Advances Information and Communication Technology)*, 2020. Springer International Publishing. doi: 10.1007/978-3-030-58201-2\_9.
- [153] Patrick Murmann, Matthias Beckerle, Simone Fischer-Hübner, and Delphine Reinhardt. Reconciling the what, when and how of privacy notifications in fitness tracking scenarios. *Pervasive and Mobile Computing*, 2021. doi: 10.1016/j.pmcj.2021.101480.
- [154] Kristin Masuch, Maike Greve, and Simon Trang. *Fitness First or Safety First? Examining Adverse Consequences of Privacy Seals in the Event of a Data Breach*. 2021. doi: 10.24251/HICSS.2021.469.
- [155] Bhavani Thuraisingham, Murat Kantarcioglu, Elisa Bertino, Jonathan Z. Bakkdash, and Maribel Fernandez. Towards a Privacy-Aware Quantified Self Data Management Framework. In *Proc. of the ACM on Symp. on Access Control Models and Technologies (SACMAT)*, 2018. Association for Computing Machinery. doi: 10.1145/3205977.3205997.
- [156] Elizabeth A. Brown. The Fitbit Fault Line: Two Proposals to Protect Health and Fitness Data at Work. *Yale Journal of Health Policy, Law and Ethics*, 2016.
- [157] Nancy H. Brinson and Danielle N. Rutherford. Privacy and the quantified self: A review of U.S. health information policy limitations related to wearable technologies. *Journal of Consumer Affairs*, 2020. doi: <https://doi.org/10.1111/joca.12320>.
- [158] Vishakha Kumari and Sara Anne Hook. The Privacy, Security and Discoverability of Data on Wearable Health Devices: Fitness or Folly? In Margherita Antona and Constantine Stephanidis, editors, *Universal Access in Human-Computer Interaction. Human and Technological Environments (Lecture Notes in Computer Science)*, 2017. Springer International Publishing. doi: 10.1007/978-3-319-58700-4\_5.
- [159] Dominik Leibenger, Frederik Möllers, Anna Petrlc, Ronald Petrlc, and Christoph Sorge. Privacy Challenges in the Quantified Self Movement – An EU Perspective. *Proc. on Privacy Enhancing Technologies (PoPETs)*, 2016. doi: 10.1515/popets-2016-0042.

- [160] Cristina M. Mares. To Cover or Not to Cover: The Relationship between the Apple Watch and the Health Insurance Portability and Accountability Act. *DePaul Journal of Health Care Law*, 2016.
- [161] Kaja J. Fietkiewicz and Maria Henkel. Privacy Protecting Fitness Trackers: An Oxymoron or Soon to Be Reality? In Gabriele Meiselwitz, editor, *Social Computing and Social Media. User Experience and Behavior (Lecture Notes in Computer Science)*, 2018. Springer International Publishing. doi: 10.1007/978-3-319-91521-0\_31.
- [162] L. Tuovinen and A. F. Smeaton. Unlocking the Black Box of Wearable Intelligence: Ethical Considerations and Social Impact. In *Proc. of the IEEE Congress on Evolutionary Computation (CEC)*, 2019. doi: 10.1109/CEC.2019.8790173.
- [163] Angela Daly. The Law and Ethics of 'Self Quantified' Health Information: An Australian Perspective. SSRN Scholarly Paper ID 2559068, Social Science Research Network, 2015. URL <https://papers.ssrn.com/abstract=2559068>.
- [164] Stefania Marassi and Philippa Collins. Is That Lawful? Data Privacy and Fitness Trackers in the Workplace. *International Journal of Comparative Labour Law and Industrial Relations*, 2021.
- [165] David Vandervort. Medical Device Data Goes to Court. In *Proc. of the Int'l Conf. on Digital Health Conf. (DH)*, 2016. Association for Computing Machinery. doi: 10.1145/2896338.2896341.
- [166] Kristen R. Moore, Natasha Jones, Bailey S. Cundiff, and Leah Heilig. Contested sites of health risks: using wearable technologies to intervene in racial oppression. *Communication Design Quarterly*, 2018. doi: 10.1145/3188387.3188392.
- [167] M. Siddiqi, S. T. Ali, and V. Sivaraman. Forensic Verification of Health Data From Wearable Devices Using Anonymous Witnesses. *IEEE Internet of Things Journal*, 2020. doi: 10.1109/JIOT.2020.2982958.
- [168] F. Hantke and A. Dewald. How can data from fitness trackers be obtained and analyzed with a forensic approach? In *Proc. of the IEEE European Symp. on Security and Privacy Workshops (EuroS PW)*, 2020. doi: 10.1109/EuroSPW51379.2020.00073.
- [169] A. Almogbil, A. Alghofaili, C. Deane, T. Leschke, A. Almogbil, and A. Alghofaili. Digital Forensic Analysis of Fitbit Wearable Technology: An Investigator's Guide. In *Proc. of the IEEE Int'l Conf. on Cyber Security and Cloud Computing (CSCloud)/IEEE Int'l Conf. on Edge Computing and Scalable Cloud (EdgeCom)*, 2020. doi: 10.1109/CSCloud-EdgeCom49738.2020.00017.
- [170] Yung Han Yoon and Umit Karabiyik. Forensic Analysis of Fitbit Versa 2 Data on Android. *Electronics*, 2020. doi: 10.3390/electronics9091431.

- [171] Sarah McNary and Aaron Hunter. Wearable Device Data for Criminal Investigation. In Guojun Wang, Jinjun Chen, and Laurence T. Yang, editors, *Security, Privacy, and Anonymity in Computation, Communication, and Storage (Lecture Notes in Computer Science)*, 2018. Springer International Publishing. doi: 10.1007/978-3-030-05345-1\_5.
- [172] Mathias Humbert, Benjamin Trubert, and Kévin Huguenin. A Survey on Interdependent Privacy. *ACM Computing Surveys*, 2019. doi: 10.1145/3360498.
- [173] Courtney Hassenfeldt, Shabana Baig, Ibrahim Baggili, and Xiaolu Zhang. Map My Murder: A Digital Forensic Study of Mobile Health and Fitness Applications. In *Proc. of the Int'l Conf. on Availability, Reliability and Security (ARES)*, 2019. Association for Computing Machinery. doi: 10.1145/3339252.3340515.
- [174] Deborah Lupton. The diverse domains of quantified selves: self-tracking modes and dataveillance. *Economy and Society*, 2016. doi: 10.1080/03085147.2016.1143726.
- [175] Etye Steinberg. Run for Your Life: The Ethics of Behavioral Tracking in Insurance. *Journal of Business Ethics*, 2021. doi: 10.1007/s10551-021-04863-8.
- [176] Valentina Di Pasquale, Valentina De Simone, Martina Radano, and Salvatore Miranda. Wearable devices for health and safety in production systems: a literature review. *IFAC-PapersOnLine*, 2022. doi: 10.1016/j.ifacol.2022.09.410.
- [177] Ekaterina Svertoka, Salwa Saafi, Alexandru Rusu-Casandra, Radim Burget, Ion Marghescu, Jiri Hosek, and Aleksandr Ometov. Wearables for Industrial Work Safety: A Survey. *Sensors*, 2021. doi: 10.3390/s21113844.
- [178] Chia-Fang Chung, Nanna Gorm, Irina A. Shklovski, and Sean Munson. Finding the Right Fit: Understanding Health Tracking in Workplace Wellness Programs. In *Proc. of the conf. on Human Factors in Computing Systems (CHI)*, 2017. Association for Computing Machinery. doi: 10.1145/3025453.3025510.
- [179] Nanna Gorm and Irina Shklovski. Steps, Choices and Moral Accounting: Observations from a Step-Counting Campaign in the Workplace. In *Proc. of the ACM Conf. on Computer-Supported Cooperative Work & Social Computing*, 2016. ACM. doi: 10.1145/2818048.2819944.
- [180] Casetext. Sunbelt Rentals, Inc., Plaintiff, v. Santiago Victor, Defendant, 2014. URL <https://casetext.com/case/sunbelt-rentals-inc-v-victor>.
- [181] Abdulmajeed Alqhatani and Heather R. Lipford. Exploring The Design Space of Sharing and Privacy Mechanisms in Wearable Fitness Platforms. In *Workshop on Usable Security and Privacy (USEC)*, 2021.

- [182] Liangyuan Na, Cong Yang, Chi-Cheng Lo, Fangyuan Zhao, Yoshimi Fukuoka, and Anil Aswani. Feasibility of Reidentifying Individuals in Large National Physical Activity Data Sets From Which Protected Health Information Has Been Removed With Use of Machine Learning. *JAMA Network Open*, 2018. doi: 10.1001/jamanetworkopen.2018.6040.
- [183] Pooja Parameshwarappa, Zhiyuan Chen, and Gunes Koru. An Effective and Computationally Efficient Approach for Anonymizing Large-Scale Physical Activity Data: Multi-Level Clustering-Based Anonymization. *International Journal of Information Security and Privacy (IJISP)*, 2020. doi: 10.4018/IJISP.2020070105.
- [184] Yanmin Gong, Yuguang Fang, and Yuanxiong Guo. Private data analytics on biomedical sensing data via distributed computation. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2016. doi: 10.1109/TCBB.2016.2515610.
- [185] Andrew Garbett, David Chatting, Gerard Wilkinson, Clement Lee, and Ahmed Kharufa. ThinkActive: Designing for Pseudonymous Activity Tracking in the Classroom. In *Proc. of the Annual ACM Conf. on Human Factors in Computing Systems (CHI)*, 2018. Association for Computing Machinery. doi: 10.1145/3173574.3173581.
- [186] Jing Wang, Na Wang, and Hongxia Jin. Context Matters? How Adding the Obfuscation Option Affects End Users' Data Disclosure Decisions. In *Proc. of the Int'l Conf. on Intelligent User Interfaces (IUI)*, 2016. Association for Computing Machinery. doi: 10.1145/2856767.2856817.
- [187] Daniel A. Epstein, Alan Borning, and James Fogarty. Fine-grained sharing of sensed physical activity: a value sensitive approach. In *UbiComp*, 2013. Association for Computing Machinery. doi: 10.1145/2493432.2493433.
- [188] Ayanga Imesha Kumari Kalupahana, Ananta Narayanan Balaji, Xiaokui Xiao, and Li-Shiuan Peh. SeRaNDiP - Leveraging Inherent Sensor Random Noise for Differential Privacy Preservation in Wearable Community Sensing Applications. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2023.
- [189] Sarah L. Alvarez, Stephanie L. Baller, and Anthony Walton. Who Owns Your Health Data? Two Interventions Addressing Data of Wearable Health Devices among Young Adults and Future Health Clinicians. *Journal of Consumer Health on the Internet*, 2021. doi: 10.1080/15398285.2020.1852386.
- [190] Odnan Ref Sanchez, Ilaria Torre, Yangyang He, and Bart P. Knijnenburg. A recommendation approach for user privacy preferences in the fitness domain. *User Modeling and User-Adapted Interaction*, 2020. doi: 10.1007/s11257-019-09246-3.
- [191] Jaco du Toit. PAUDIT: A Distributed Data Architecture for Fitness Data. In Hein Venter, Marianne Loock, Marijke Coetzee, Mariki Eloff, and Jan Eloff, editors, *Information and Cyber Security (Communications in Computer and Information Science)*, 2020. Springer International Publishing. doi: 10.1007/978-3-030-43276-8\_4.

- [192] Kambiz Ghazinour, Emil Shirima, Vijayasimha Reddy Parne, and Abhilash Bhoom-Reddy. A Model to Protect Sharing Sensitive Information in Smart Watches. *Procedia Computer Science*, 2017. doi: 10.1016/j.procs.2017.08.322.
- [193] Xiao Liu, Bonan Gao, Basem Suleiman, Han You, Zisu Ma, Yu Liu, and Ali Anaissi. Privacy-Preserving Personalized Fitness Recommender System P3FitRec: A Multi-level Deep Learning Approach. *ACM Transactions on Knowledge Discovery from Data*, 2023. doi: 10.1145/3572899.
- [194] Aavek K. Das, Parth H. Pathak, Chen-Nee Chuah, and Prasant Mohapatra. Uncovering Privacy Leakage in BLE Network Traffic of Wearable Fitness Trackers. In *Proc. of the Int'l Workshop on Mobile Computing Systems and Applications (HotMobile)*, 2016. Association for Computing Machinery. doi: 10.1145/2873587.2873594.
- [195] K. Lotfy and M. L. Hale. Assessing Pairing and Data Exchange Mechanism Security in the Wearable Internet of Things. In *Proc. of the IEEE Int'l Conf. on Mobile Services (MS)*, 2016. doi: 10.1109/MobServ.2016.15.
- [196] Rohit Goyal, Nicola Dragoni, and Angelo Spognardi. Mind the tracker you wear: a security analysis of wearable health trackers. In *Proc. of the Annual ACM Symp. on Applied Computing (SAC)*, 2016. Association for Computing Machinery. doi: 10.1145/2851613.2851685.
- [197] M. Rahman, B. Carbutar, and U. Topkara. Secure Management of Low Power Fitness Trackers. *IEEE Transactions on Mobile Computing*, 2016. doi: 10.1109/TMC.2015.2418774.
- [198] Q. Zhang and Z. Liang. Security analysis of bluetooth low energy based smart wristbands. In *Int'l Conf. on Frontiers of Sensors Technologies (ICFST)*, 2017. doi: 10.1109/ICFST.2017.8210548.
- [199] X. Fafoutis, L. Marchegiani, G. Z. Papadopoulos, R. Piechocki, T. Tryfonas, and G. Oikonomou. Privacy Leakage of Physical Activity Levels in Wireless Embedded Wearable Systems. *IEEE Signal Processing Letters*, 2017. doi: 10.1109/LSP.2016.2642300.
- [200] Jaewoo Shim, Kyeonghwan Lim, Jaemin Jeong, Seong-je Cho, Minkyu Park, and Sangchul Han. A Case Study on Vulnerability Analysis and Firmware Modification Attack for a Wearable Fitness Tracker. 2017.
- [201] Jiska Classen, Daniel Wegemer, Paul Patras, Tom Spink, and Matthias Hollick. Anatomy of a Vulnerable Fitness Tracking System: Dissecting the Fitbit Cloud, App, and Firmware. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2018. doi: 10.1145/3191737.
- [202] Florina Almenares Mendoza, Lucía Alonso, Andrés Marín López, , and Daniel Díaz Sánchez Patricia Arias Cabarcos. Assessment of Fitness Tracker Security: A Case of Study. *Proceedings*, 2018. doi: 10.3390/proceedings2191235.

- [203] Guillaume Celosia and Mathieu Cunche. Fingerprinting Bluetooth-Low-Energy Devices Based on the Generic Attribute Profile. In *Proc. of the Int'l ACM Workshop on Security and Privacy for the Internet (of-Things)*, 2019. Association for Computing Machinery. doi: 10.1145/3338507.3358617.
- [204] Chaoshun Zuo, Haohuang Wen, Zhiqiang Lin, and Yinqian Zhang. Automatic Fingerprinting of Vulnerable BLE IoT Devices with Static UUIDs from Mobile Apps. In *Proc. of the ACM SIGSAC Conf. on Computer and Communications Security (CCS)*, 2019. Association for Computing Machinery. doi: 10.1145/3319535.3354240.
- [205] Johannes K Becker, David Li, and David Starobinski. Tracking Anonymized Bluetooth Devices. *Proc. on Privacy Enhancing Technologies (PoPETs)*, 2019. doi: 10.2478/popets-2019-0036.
- [206] Matthew L. Hale, Kerolos Lotfy, Rose F. Gamble, Charles Walter, and Jessica Lin. Developing a platform to evaluate and assess the security of wearable devices. *Digital Communications and Networks*, 2019. doi: 10.1016/j.dcan.2018.10.009.
- [207] Jiliang Wang, Feng Hu, Ye Zhou, Yunhao Liu, Hanyi Zhang, and Zhe Liu. BlueDoor: breaking the secure information flow via BLE vulnerability. In *Proc. of the Int'l Conf. on Mobile Systems, Applications, and Services (MobiSys)*, 2020. Association for Computing Machinery. doi: 10.1145/3386901.3389025.
- [208] O. M. Gouda, D. J. Hejji, and M. S. Obaidat. Privacy Assessment of Fitness Tracker Devices. In *Int'l Conf. on Computer, Information and Telecommunication Systems (CITS)*, 2020. doi: 10.1109/CITS49457.2020.9232503.
- [209] Ludovic Barman, Alexandre Dumur, Apostolos Pyrgelis, and Jean-Pierre Hubaux. Every Byte Matters: Traffic Analysis of Bluetooth Wearable Devices. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2021. doi: 10.1145/3463512.
- [210] Jaime Fúster, Sonia Solera-Cotanilla, Jaime Pérez, Mario Vega-Barbas, Rafael Palacios, Manuel Álvarez Campana, and Gregorio Lopez. Analysis of security and privacy issues in wearables for minors. *Wireless Networks*, 2023. doi: 10.1007/s11276-022-03211-6.
- [211] Hossein Fereidooni, Tommaso Frassetto, Markus Miettinen, Ahmad-Reza Sadeghi, and Mauro Conti. Fitness trackers: fit for health but unfit for security and privacy. In *Proc. of the IEEE/ACM Int'l Conf. on Connected Health: Applications, Systems and Engineering Technologies (CHASE)*, 2017. IEEE Press. doi: 10.1109/CHASE.2017.54.
- [212] Hossein Fereidooni, Jiska Classen, Tom Spink, Paul Patras, Markus Miettinen, Ahmad-Reza Sadeghi, Matthias Hollick, and Mauro Conti. Breaking Fitness Records Without Moving: Reverse Engineering and Spoofing Fitbit. In Marc Dacier, Michael Bailey, Michalis Polychronakis, and Manos Antonakakis, editors, *Research in Attacks, Intrusions, and Defenses (Lecture Notes in Computer Science)*, 2017. Springer International Publishing. doi: 10.1007/978-3-319-66332-6.3.

- [213] Andrei Kazlouski, Thomas Marchioro, Harry Manifavas, and Evangelos Markatos. I still See You! Inferring Fitness Data from Encrypted Traffic of Wearables:. In *Proc. of the Int'l Joint Conf. on Biomedical Engineering Systems and Technologies*, 2021. SCITEPRESS - Science and Technology Publications. doi: 10.5220/0010233103690376.
- [214] A. Maiti, M. Jadliwala, J. He, and I. Bilogrevic. Side-Channel Inference Attacks on Mobile Key pads Using Smartwatches. *IEEE Transactions on Mobile Computing*, 2018. doi: 10.1109/TMC.2018.2794984.
- [215] Mohd Sabra, Anindya Maiti, and Murtuza Jadliwala. Keystroke inference using ambient light sensor on wrist-wearables: a feasibility study. In *Proc. of the ACM Workshop on Wearable Systems and Applications*, 2018. ACM. doi: 10.1145/3211960.3211973.
- [216] C. Wang, X. Guo, Y. Chen, Y. Wang, and B. Liu. Personal PIN Leakage from Wearable Devices. *IEEE Transactions on Mobile Computing*, 2018. doi: 10.1109/TMC.2017.2737533.
- [217] Guglielmo Cola, Marco Avvenuti, Fabio Musso, and Alessio Vecchio. Gait-based authentication using a wrist-worn device. In *Proc. of the Int'l Conf. on Mobile and Ubiquitous Systems: Computing, Networking and Services (MOBIQUITOUS)*, 2016. Association for Computing Machinery. doi: 10.1145/2994374.2994393.
- [218] Andrew H. Johnston and Gary M. Weiss. Smartwatch-based biometric gait recognition. In *Int'l Conf. on Biometrics Theory, Applications and Systems (BTAS)*, 2015. IEEE. doi: 10.1109/BTAS.2015.7358794.
- [219] S. Vhaduri and C. Poellabauer. Wearable device user authentication using physiological and behavioral metrics. In *Proc. of the IEEE Annual Int'l Symp. on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, 2017. doi: 10.1109/PIMRC.2017.8292272.
- [220] F. Tehranipoor, N. Karimian, P. A. Wortman, and J. A. Chandy. Low-cost authentication paradigm for consumer electronics within the internet of wearable fitness tracking applications. In *Proc. of the IEEE Int'l Conf. on Consumer Electronics (ICCE)*, 2018. doi: 10.1109/ICCE.2018.8326233.
- [221] Wenqiang Chen, Lin Chen, Yandao Huang, Xinyu Zhang, Lu Wang, Rukhsana Ruby, and Kaishun Wu. Taprint: Secure Text Input for Commodity Smart Wristbands. In *The Annual Int'l Conf. on Mobile Computing and Networking (MobiCom)*, 2019. Association for Computing Machinery. doi: 10.1145/3300061.3300124.
- [222] Xiaopeng Li, Fengyao Yan, Fei Zuo, Qiang Zeng, and Lannan Luo. Touch Well Before Use: Intuitive and Secure Authentication for IoT Devices. In *The Annual Int'l Conf. on Mobile Computing and Networking (MobiCom)*, 2019. Association for Computing Machinery. doi: 10.1145/3300061.3345434.



- [223] Yiran Shen, Bowen Du, Weitao Xu, Chengwen Luo, Bo Wei, Lizhen Cui, and Hongkai Wen. Securing Cyber-Physical Social Interactions on Wrist-Worn Devices. *ACM Transactions on Sensor Networks*, 2020. doi: 10.1145/3378669.
- [224] Prakash Shrestha and Nitesh Saxena. Hacksaw: biometric-free non-stop web authentication in an emerging world of wearables. In *Proc. of the ACM Conf. on Security and Privacy in Wireless and Mobile Networks (WiSec)*, 2020. Association for Computing Machinery. doi: 10.1145/3395351.3399366.
- [225] Jack Sturgess, Simon Eberz, Ivo Sluganovic, and Ivan Martinovic. WatchAuth: User Authentication and Intent Recognition in Mobile Payments using a Smartwatch. 2022.
- [226] Christian Skalka, John Ring, David Darias, Minseok Kwon, Sahil Gupta, Kyle Diller, Steffen Smolka, and Nate Foster. Proof-Carrying Network Code. In *Proc. of the ACM SIGSAC Conf. on Computer and Communications Security (CCS)*, 2019. Association for Computing Machinery. doi: 10.1145/3319535.3363214.
- [227] Wenqing Yan, Sam Hylamia, Thiemo Voigt, and Christian Rohner. PHY-IDS: a physical-layer spoofing attack detection system for wearable devices. In *Proc. of the ACM Workshop on Wearable Systems and Applications (WearSys)*, 2020. Association for Computing Machinery. doi: 10.1145/3396870.3400010.
- [228] Jingyu Xin, Vir V. Phoha, and Asif Salekin. Combating False Data Injection Attacks on Human-Centric Sensing Applications. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2022. doi: 10.1145/3534577.
- [229] J. Mnjama, G. Foster, and B. Irwin. A privacy and security threat assessment framework for consumer health wearables. In *Information Security for South Africa (ISSA)*, 2017. doi: 10.1109/ISSA.2017.8251776.
- [230] Sophia Moganedi and Dalenca Pottas. Identification of Information Security Controls for Fitness Wearable Manufacturers. In Hein Venter, Marianne Loock, Marijke Coetzee, Mariki Eloff, Jan Eloff, and Reinhardt Botha, editors, *Information and Cyber Security (Communications in Computer and Information Science)*, 2020. Springer International Publishing. doi: 10.1007/978-3-030-66039-0\_8.
- [231] Finn Kensing and Jeanette Blomberg. Participatory Design: Issues and Concerns. *Computer Supported Cooperative Work (CSCW)*, 1998. doi: 10.1023/A:1008689307411.
- [232] Angela Chen. What happens when life insurance companies track fitness data?, 2018. URL <https://www.theverge.com/2018/9/26/17905390/john-hancock-life-insurance-fitness-tracker-wearables-science-health>.
- [233] Jennifer King, Airi Lampinen, and Alex Smolen. Privacy: is there an app for that? In *Proc. of the USENIX Symp. on Usable Privacy and Security (SOUPS)*, 2011. ACM Press. doi: 10.1145/2078827.2078843.

- [234] Hanna Krasnova, Nicole Eling, Oleg Schneider, Helena Wenninger, Thomas Widjaja, and Peter Buxmann. Does This App Ask For Too Much Data? The Role Of Privacy Perceptions In User Behavior Towards Facebook Applications And Permission Dialogs. *ECIS*, 2013.
- [235] Pamela Wisniewski, Heng Xu, Heather Lipford, and Emmanuel Bello-Ogunu. 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*, 2015. doi: 10.1002/asi.23299.
- [236] Jennifer Golbeck and Matthew Mauriello. User Perception of Facebook App Data Access: A Comparison of Methods and Privacy Concerns. *Future Internet*, 2016. doi: 10.3390/fi8020009.
- [237] Wanpeng Li and Chris J. Mitchell. Security Issues in OAuth 2.0 SSO Implementations. In Sherman S. M. Chow, Jan Camenisch, Lucas C. K. Hui, and Siu Ming Yiu, editors, *Information Security (Lecture Notes in Computer Science)*, 2014. Springer International Publishing. doi: 10.1007/978-3-319-13257-0\_34.
- [238] X. Li, J. Xu, Z. Zhang, X. Lan, and Y. Wang. Modular Security Analysis of OAuth 2.0 in the Three-Party Setting. In *Euro S&P*, 2020. doi: 10.1109/EuroSP48549.2020.00025.
- [239] K. I. Manktelow and Man Cheung Chung, editors. *Psychology of reasoning: theoretical and historical perspectives*. Psychology Press, 1 edition, 2004.
- [240] Na Wang, Heng Xu, and Jens Grossklags. 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)*, 2011. ACM Press. doi: 10.1145/2076444.2076448.
- [241] Patricia Arias-Cabarcos, Saina Khalili, and Thorsten Strufe. 'Surprised, Shocked, Worried': User Reactions to Facebook Data Collection from Third Parties, 2022. URL <http://arxiv.org/abs/2209.08048>.
- [242] Jaime Delgado, Eva Rodríguez, and Silvia Llorente. User's privacy in applications provided through social networks. In *Proc. of the ACM SIGMM workshop on Social media (WSM)*, 2010. ACM Press. doi: 10.1145/1878151.1878163.
- [243] Mohamed Shehab, Said Marouf, and Christopher Hudel. ROAuth: recommendation based open authorization. In *Proc. of the USENIX Symp. on Usable Privacy and Security (SOUPS)*, 2011. ACM Press. doi: 10.1145/2078827.2078842.
- [244] Pauline Anthonysamy, Awais Rashid, James Walkerdine, Phil Greenwood, and Georgios Larkou. 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)*, 2012. ACM Press. doi: 10.1145/2185354.2185359.

- [245] Yuan Cheng, Jaehong Park, and Ravi Sandhu. Preserving user privacy from third-party applications in online social networks. In *Proc. of the Int'l Conf. on World Wide Web - WWW Companion*, 2013. ACM Press. doi: 10.1145/2487788.2488032.
- [246] Seyed Hossein Ahmadijad, Philip W.L. Fong, and Reihaneh Safavi-Naini. 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)*, 2016. ACM. doi: 10.1145/2897845.2897878.
- [247] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. Android permissions: user attention, comprehension, and behavior. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, 2012. ACM. doi: 10.1145/2335356.2335360.
- [248] Lydia Kraus, Ina Wechsung, and Sebastian Moller. Using Statistical Information to Communicate Android Permission Risks to Users. In *2014 Workshop on Socio-Technical Aspects in Security and Trust*, 2014. IEEE. doi: 10.1109/STAST.2014.15.
- [249] Irina Shklovski, Scott D. Mainwaring, Halla Hrunn Skúladóttir, and Höskuldur Borgthorsson. Leakiness and creepiness in app space: perceptions of privacy and mobile app use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2014. ACM. doi: 10.1145/2556288.2557421.
- [250] Panagiotis Andriotis, Martina Angela Sasse, and Gianluca Stringhini. Permissions snapshots: Assessing users' adaptation to the Android runtime permission model. In *2016 IEEE International Workshop on Information Forensics and Security (WIFS)*, 2016. IEEE. doi: 10.1109/WIFS.2016.7823922.
- [251] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhammedi, Shikun (Aerin) Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions. 2016.
- [252] K. Olejnik, I. Dacosta, J. S. Machado, K. Huguenin, M. E. Khan, and J. P. Hubaux. SmarPer: Context-Aware and Automatic Runtime-Permissions for Mobile Devices. In *S&P*, 2017. doi: 10.1109/SP.2017.25.
- [253] Primal Wijesekera, Arjun Baokar, Lynn Tsai, Joel Reardon, Serge Egelman, David Wagner, and Konstantin Beznosov. The Feasibility of Dynamically Granted Permissions: Aligning Mobile Privacy with User Preferences. In *2017 IEEE Symposium on Security and Privacy (SP)*, 2017. IEEE. doi: 10.1109/SP.2017.51.
- [254] Lynn Tsai, Primal Wijesekera, Joel Reardon, Irwin Reyes, Serge Egelman, David Wagner, Nathan Good, and Jung-Wei Chen. Turtle Guard: Helping Android Users Apply Contextual Privacy Preferences. 2017.
- [255] Weicheng Cao, Chunqiu Xia, Sai Teja Peddinti, David Lie, Nina Taft, and Lisa M. Austin. A Large Scale Study of User Behavior, Expectations and Engagement with Android Permissions. 2021.

- [256] Stefan Palan and Christian Schitter. Prolific.ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance*, 2018. doi: 10.1016/j.jbef.2017.12.004.
- [257] Statista. Wearable band market share in North America by vendor 2018-2020, 2022. URL <https://www.statista.com/statistics/1042044/north-america-quarterly-wearable-band-market-share-by-vendor/>.
- [258] Christina Farr. Fitbit has a new health tracker, but you can only get it through your employer or insurer, 2019. URL <https://www.cnbc.com/2019/02/08/fitbit-releases-inspire-for-employers.html>.
- [259] Rajesh Pandey. Android 11 will automatically revoke permissions from unused apps, 2020. URL <https://www.neowin.net/news/android-11-will-automatically-revoke-permissions-from-unused-apps/>.
- [260] David Jonassen and Young Hoan Cho. Externalizing Mental Models with Mindtools. In Dirk Ifenthaler, Pablo Pirnay-Dummer, and J. Michael Spector, editors, *Understanding Models for Learning and Instruction*, pages 145–159. Springer US, 2008. doi: 10.1007/978-0-387-76898-4.7.
- [261] Alexandra Mai, Katharina Pfeffer, Matthias Gusenbauer, Edgar Weippl, and Katharina Krombholz. User Mental Models of Cryptocurrency Systems - A Grounded Theory Approach. In *Proc. of the USENIX Symp. on Usable Privacy and Security (SOUPS)*, 2020. USENIX Association.
- [262] Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara Kiesler. “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)*, 2015. USENIX Association.
- [263] Katharina Krombholz, Karoline Busse, Katharina Pfeffer, Matthew Smith, and Emanuel von Zezschwitz. “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 (SP)*, 2019. IEEE. doi: 10.1109/SP.2019.00060.
- [264] Maggie Oates, Yama Ahmadullah, Abigail Marsh, Chelse Swoopes, Shikun Zhang, Rebecca Balebako, and Lorrie Faith Cranor. Turtles, Locks, and Bathrooms: Understanding Mental Models of Privacy Through Illustration. *Proc. on Privacy Enhancing Technologies (PoPETs)*, 2018. doi: 10.1515/popets-2018-0029.
- [265] Naresh K. Malhotra, Sung S. Kim, and James Agarwal. Internet Users’ Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model. *Information Systems Research*, 2004. doi: 10.1287/isre.1040.0032.
- [266] Tenga Matsuura, Ayako A. Hasegawa, Mitsuaki Akiyama, and Tatsuya Mori. Careless Participants Are Essential for Our Phishing Study: Understanding the Impact of Screening Methods. In *European Symp. on Usable Security*, 2021. ACM. doi: 10.1145/3481357.3481515.

- [267] Ruba Abu-Salma, M. Angela Sasse, Joseph Bonneau, Anastasia Danilova, Alena Naiakshina, and Matthew Smith. Obstacles to the Adoption of Secure Communication Tools. In *Proc. of the IEEE Symp. on Security and Privacy (SP)*, 2017. doi: 10.1109/SP.2017.65.
- [268] Rick Wash and Emilee Rader. Influencing mental models of security: a research agenda. In *Proc. of the Conf. New Security Paradigms Workshop (NSPW)*, 2011. Association for Computing Machinery. doi: 10.1145/2073276.2073283.
- [269] Johnny Saldana. *The Coding Manual for Qualitative Researchers*. SAGE Publishing, 4th ed edition, 2021.
- [270] Yuvin Ha, Maria Karyda, and Andrés Lucero. Exploring Virtual Rewards in Real Life: A Gimmick or a Motivational Tool for Promoting Physical Activity? In *Proc. of the ACM Designing Interactive Systems Conf. (DIS)*, 2020. Association for Computing Machinery. doi: 10.1145/3357236.3395477.
- [271] Rocket Fuel. 'Quantified Self' Digital Tools: A CPG Marketing Opportunity. Technical report, Rocket Fuel, 2014.
- [272] Hamza Harkous, Kassem Fawaz, Rémi Lebret, Florian Schaub, Kang G. Shin, and Karl Aberer. Polisis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning. 2018.
- [273] Fitbit. Fitbit SDK, 2020. URL <https://dev.fitbit.com/>.
- [274] garmin. Overview | Garmin Connect Developer Program | Garmin Developers. URL <https://developer.garmin.com/gc-developer-program/overview/>.
- [275] Milena Janic, Jan Pieter Wijbenga, and Thijs Veugen. Transparency Enhancing Tools (TETs): An Overview. In *Workshop on Socio-Technical Aspects in Security and Trust*, 2013. IEEE. doi: 10.1109/STAST.2013.11.
- [276] Apple. Legal - Data & Privacy - Apple. URL <https://www.apple.com/legal/privacy/data/en/health-app/>.
- [277] James Heckman. Varieties of Selection Bias. *American Economic Review*, 1990.
- [278] Carole Cadwalladr and Emma Graham-Harrison. Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach. *The Guardian*, 2018.
- [279] Sonia Roccas, Lilach Sagiv, Shalom H. Schwartz, and Ariel Knafo. The Big Five Personality Factors and Personal Values. *Personality and Social Psychology Bulletin*, 2002. doi: 10.1177/0146167202289008.
- [280] Carlos A. Rivera. The Big-Five Personality Test and Cambridge Analytica, 2019. URL <https://galindes.wordpress.com/2019/05/03/the-big-five-personality-test-and-cambridge-analytica/>.

- [281] Elizabeth Gibney. The scant science behind Cambridge Analytica's controversial marketing techniques. *Nature*, 2018.
- [282] Charles Duhigg. What Does Your Credit-Card Company Know About You? *The New York Times*, 2009.
- [283] Rodrigo de Oliveira, Alexandros Karatzoglou, Pedro Concejero Cerezo, Ana Armenta Lopez de Vicuña, and Nuria Oliver. Towards a psychographic user model from mobile phone usage. In *CHI · Work-in-Progress*, 2011. ACM. doi: 10.1145/1979742.1979920.
- [284] Gokul Chittaranjan, Jan Blom, and Daniel Gatica-Perez. Who's Who with Big-Five: Analyzing and Classifying Personality Traits with Smartphones. In *ISWC*, 2011. IEEE. doi: 10.1109/ISWC.2011.29.
- [285] Yves-Alexandre de Montjoye, Jordi Quoidbach, Florent Robic, and Alex Pentland. Predicting Personality Using Novel Mobile Phone-Based Metrics. *Social Computing, Behavioral-Cultural Modeling and Prediction*, 2013. doi: 10.1007/978-3-642-37210-0\_6.
- [286] Bjarke Mønsted, Anders Mollgaard, and Joachim Mathiesen. Phone-based metric as a predictor for basic personality traits. *Journal of Research in Personality*, 2018. doi: 10.1016/j.jrp.2017.12.004.
- [287] Clemens Stachl, Quay Au, Ramona Schoedel, Samuel D. Gosling, Gabriella M. Harari, Daniel Buschek, Sarah Theres Völkel, Tobias Schuwerk, Michelle Oldemeier, Theresa Ullmann, Heinrich Hussmann, Bernd Bischl, and Markus Bühner. Predicting personality from patterns of behavior collected with smartphones. *Proceedings of the National Academy of Sciences*, 2020. doi: 10.1073/pnas.1920484117.
- [288] Lisa M Pytlik Zillig, Scott H Hemenover, and Richard A Dienstbier. What Do We Assess when We Assess a Big 5 Trait? A Content Analysis of the Affective, Behavioral, and Cognitive Processes Represented in Big 5 Personality Inventories. 2002. doi: 10.1177/0146167202289013.
- [289] R E Rhodes and N E I Smith. Personality correlates of physical activity: a review and meta-analysis. *British Journal of Sports Medicine*, 2006. doi: 10.1136/bjism.2006.028860.
- [290] Mirka Hintsanen, Sampsa Puttonen, Kylie Smith, Maria Törnroos, Markus Jokela, Laura Pulkki-Råback, Taina Hintsala, Päivi Merjonen, Terence Dwyer, Olli T. Raitakari, Alison Venn, and Liisa Keltikangas-Järvinen. Five-factor personality traits and sleep: Evidence from two population-based cohort studies. *Health Psychology*, 2014. doi: 10.1037/hea0000105.
- [291] Esma Aimeur, Gilles Brassard, and Muxue Guo. How data brokers endanger privacy. *Transactions on Data Privacy*, 2022.
- [292] Deborah A. Cobb-Clark and Stefanie Schurer. The stability of big-five personality traits. *Economics Letters*, 2012. doi: 10.1016/j.econlet.2011.11.015.

- [293] Paul T. Costa and Robert R. McCrae. Four ways five factors are basic. *Personality and Individual Differences*, 1992. doi: 10.1016/0191-8869(92)90236-I.
- [294] Ana Carolina E.S. Lima and Leandro Nunes de Castro. A multi-label, semi-supervised classification approach applied to personality prediction in social media. *Neural Networks*, 2014. doi: 10.1016/j.neunet.2014.05.020.
- [295] M. Kosinski, D. Stillwell, and T. Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 2013. doi: 10.1073/pnas.1218772110.
- [296] Michal Kosinski, Yoram Bachrach, Pushmeet Kohli, David Stillwell, and Thore Graepel. Manifestations of user personality in website choice and behaviour on online social networks. *Machine Learning*, 2014. doi: 10.1007/s10994-013-5415-y.
- [297] Naa Amponsah Dodoo and Cynthia Morton Padovano. Personality-Based Engagement: An Examination of Personality and Message Factors on Consumer Responses to Social Media Advertisements. *Journal of Promotion Management*, 2020. doi: 10.1080/10496491.2020.1719954.
- [298] Brahim Zarouali, Tom Dobber, Guy De Pauw, and Claes de Vreese. Using a Personality-Profiling Algorithm to Investigate Political Microtargeting: Assessing the Persuasion Effects of Personality-Tailored Ads on Social Media. *Communication Research*, 2020. doi: 10.1177/0093650220961965.
- [299] Fitbit Development: Fitbit Web API Basics, 2021. URL <https://dev.fitbit.com/build/reference/web-api/basics/>.
- [300] Salvatore Tedesco, Marco Sica, Andrea Ancillao, Suzanne Timmons, John Barton, and Brendan O’Flynn. Accuracy of consumer-level and research-grade activity trackers in ambulatory settings in older adults. *PLOS ONE*, 2019. doi: 10.1371/journal.pone.0216891.
- [301] Fitbit Counts on Women as Device Buyers, Just Not Board Members, 2015. URL <https://www.businessoffashion.com/articles/technology/fitbit-counts-on-women-as-device-buyers-just-not-board-members>.
- [302] Isabel Wagner and David Eckhoff. Technical Privacy Metrics: A Systematic Survey. *ACM Computing Surveys*, 2019. doi: 10.1145/3168389.
- [303] Paul T. Costa Jr. and Robert R. McCrae. Domains and Facets: Hierarchical Personality Assessment Using the Revised NEO Personality Inventory. *Journal of Personality Assessment*, 1995. doi: 10.1207/s15327752jpa6401\_2.
- [304] scikit learn. scikit-learn: machine learning in Python — scikit-learn 0.24.1 documentation. URL <https://scikit-learn.org/stable/>.
- [305] Randy J. Larsen. Individual differences in circadian activity rhythm and personality. *Personality and Individual Differences*, 1985. doi: 10.1016/0191-8869(85)90054-6.

- [306] Colin G. DeYoung, Lynn Hasher, Maja Djikic, Brock Criger, and Jordan B. Peterson. Morning people are stable people: Circadian rhythm and the higher-order factors of the Big Five. *Personality and Individual Differences*, 2007. doi: 10.1016/j.paid.2006.11.030.
- [307] Gokul Chittaranjan, Jan Blom, and Daniel Gatica-Perez. Mining large-scale smartphone data for personality studies. *Personal and Ubiquitous Computing*, 2013. doi: 10.1007/s00779-011-0490-1.
- [308] Clemens Stachl, Florian Pargent, Sven Hilbert, Gabriella M. Harari, Ramona Schoedel, Sumer Vaid, Samuel D. Gosling, and Markus Bühner. Personality Research and Assessment in the Era of Machine Learning. *European Journal of Personality*, 2020. doi: 10.1002/per.2257.
- [309] Ada H. Zohar, C. Robert Cloninger, and Rollin McCraty. Personality and Heart Rate Variability: Exploring Pathways from Personality to Cardiac Coherence and Health. *Open Journal of Social Sciences*, 2013. doi: 10.4236/jss.2013.16007.
- [310] A. Sano, A. J. Phillips, A. Z. Yu, A. W. McHill, S. Taylor, N. Jaques, C. A. Czeisler, E. B. Klerman, and R. W. Picard. Recognizing academic performance, sleep quality, stress level, and mental health using personality traits, wearable sensors and mobile phones. In *Proc. of the IEEE Int'l Conf. on Wearable and Implantable Body Sensor Networks (BSN)*, 2015. doi: 10.1109/BSN.2015.7299420.
- [311] Christoph Randler. Morningness–eveningness, sleep–wake variables and big five personality factors. *Personality and Individual Differences*, 2008. doi: 10.1016/j.paid.2008.03.007.
- [312] Yanna J. Weisberg, Colin G. DeYoung, and Jacob B. Hirsh. Gender Differences in Personality across the Ten Aspects of the Big Five. *Frontiers in Psychology*, 2011. doi: 10.3389/fpsyg.2011.00178.
- [313] Stephen Bright, Eyal Gringart, Emily Blatchford, and Samantha Bettinson. A quantitative exploration of the relationships between regular yoga practice, microdosing psychedelics, wellbeing and personality variables. *Australian Journal of Psychology*, 2021. doi: 10.1080/00049530.2021.1882266.
- [314] Páraic S. Ó Súilleabháin, Siobhán Howard, and Brian M. Hughes. Openness to experience and adapting to change: Cardiovascular stress habituation to change in acute stress exposure. *Psychophysiology*, 2018. doi: 10.1111/psyp.13023.
- [315] Iva Čukić and Timothy C. Bates. Openness to experience and aesthetic chills: Links to heart rate sympathetic activity. *Personality and Individual Differences*, 2014. doi: 10.1016/j.paid.2014.02.012.
- [316] John A. Johnson. Units of Analysis for the Description and Explanation of Personality. In *Handbook of Personality Psychology*, pages 73–93. Elsevier, 1997. doi: 10.1016/B978-012134645-4/50004-4.



- [317] Lakmal Meegahapola, Marios Constantinides, Zoran Radivojevic, Hongwei Li, Daniele Quercia, and Michael S Eggleston. Quantified Canine: Inferring Dog Personality From Wearables. 2023.
- [318] Daniel Arp, Erwin Quiring, Feargus Pendlebury, Alexander Warnecke, Fabio Pierazzi, Christian Wressnegger, Lorenzo Cavallaro, and Konrad Rieck. Dos and Don'ts of Machine Learning in Computer Security. In *Proc. of the USENIX Security Symp.*, 2020. doi: <https://doi.org/10.48550/arXiv.2010.09470>.
- [319] Brooke Auxier, Lee Rainie, Monica Anderson, Andrew Perrin, Madhu Kumar, and Erica Turner. Americans and Privacy: Concerned, Confused and Feeling Lack of Control Over Their Personal Information, 2019. URL <https://www.pewresearch.org/internet/2019/11/15/americans-and-privacy-concerned-confused-and-feeling-lack-of-control-over-their-personal-information/>.
- [320] Fitbit. Fitbit Inspire HR User Manual, 2019. URL [https://staticcs.fitbit.com/content/assets/help/manuals/manual\\_inspire\\_hr\\_en\\_US.pdf](https://staticcs.fitbit.com/content/assets/help/manuals/manual_inspire_hr_en_US.pdf).
- [321] Cory Hallam and Gianluca Zanella. Wearable Device Data and Privacy: A study of Perception and Behavior. *World Journal of Management*, 2016. doi: 10.21102/wjm.2016.03.71.06.
- [322] Igor Bilogrevic, Kévin Huguenin, Stefan Mihaila, Reza Shokri, and Jean-Pierre Hubaux. Predicting Users' Motivations behind Location Check-Ins and Utility Implications of Privacy Protection Mechanisms. In *Proc. of the Network and Distributed System Security Symp. (NDSS)*, 2015. Internet Society. doi: 10.14722/ndss.2015.23032.
- [323] Mary J. Culnan and Pamela K. Armstrong. Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation. *Organization Science*, 1999. doi: 10.1287/orsc.10.1.104.
- [324] Stephen R. Porter, Michael E. Whitcomb, and William H. Weitzer. Multiple surveys of students and survey fatigue. *New Directions for Institutional Research*, 2004. doi: 10.1002/ir.101.
- [325] Anton J Nederhof. Methods of coping with social desirability bias: A review. 1985. doi: 10.1002/ejsp.2420150303.
- [326] L. M. Verbrugge, A. L. Gruber-Baldini, and J. L. Fozard. Age Differences and Age Changes in Activities: Baltimore Longitudinal Study of Aging. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 1996. doi: 10.1093/geronb/51B.1.S30.
- [327] Sara J. Weston, Grant W. Edmonds, and Patrick L. Hill. Personality traits predict dietary habits in middle-to-older adults. *Psychology, Health & Medicine*, 2020. doi: 10.1080/13548506.2019.1687918.

- [328] Susanne Weber, Marian Harbach, and Matthew Smith. Participatory Design for Security-Related User Interfaces. In *Proc. Workshop on Usable Security*, 2015. Internet Society. doi: 10.14722/usec.2015.23011.
- [329] Finn Kensing and Andreas Munk-Madsen. PD: structure in the toolbox. *Communications of the ACM*, 1993. doi: 10.1145/153571.163278.
- [330] SR Davis, D Peters, RA Calvo, SM Sawyer, JM Foster, and L Smith. “Kiss myAsthma”: Using a participatory design approach to develop a self-management app with young people with asthma. *Journal of Asthma*, 2018. doi: 10.1080/02770903.2017.1388391.
- [331] Stephen Lindsay, Daniel Jackson, Guy Schofield, and Patrick Olivier. Engaging older people using participatory design. In *Proc. of the conf. on Human Factors in Computing Systems (CHI)*, 2012. Association for Computing Machinery. doi: 10.1145/2207676.2208570.
- [332] Meethu Malu and Leah Findlater. Toward Accessible Health and Fitness Tracking for People with Mobility Impairments. In *Proc. of the EAI Int’l Conf. on Pervasive Computing Technologies for Healthcare*, 2016. ACM. doi: 10.4108/eai.16-5-2016.2263329.
- [333] Kavous Salehzadeh Niksirat, Evanne Anthoine-Milhomme, Samuel Randin, Kévin Huguenin, and Mauro Cherubini. “I thought you were okay”: Participatory Design with Young Adults to Fight Multiparty Privacy Conflicts in Online Social Networks. In *Designing Interactive Systems Conf.*, 2021. ACM. doi: 10.1145/3461778.3462040.
- [334] Alethia Hume, Nicolás Ferreira, and Luca Cernuzzi. The design of a privacy dashboard for an academic environment based on participatory design. In *XLVII Latin American Computing Conf. (CLEI)*, 2021. doi: 10.1109/CLEI53233.2021.9640155.
- [335] Keld Bødker, Finn Kensing, and Jesper Simonsen. Participatory Design in Information Systems Development. In Hannakaisa Isomäki and Samuli Pekkola, editors, *Reframing Humans in Information Systems Development*, Computer Supported Cooperative Work, pages 115–134. Springer, 2011. doi: 10.1007/978-1-84996-347-3\_7.
- [336] Douglas Zytka, Pamela J. Wisniewski, Shion Guha, Eric P. S. Baumer, and Min Kyung Lee. Participatory Design of AI Systems: Opportunities and Challenges Across Diverse Users, Relationships, and Application Domains. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*, 2022. Association for Computing Machinery. doi: 10.1145/3491101.3516506.
- [337] Kavous Salehzadeh Niksirat, Diana Korka, Hamza Harkous, Kévin Huguenin, and Mauro Cherubini. On the Potential of Mediation Chatbots for Mitigating Multiparty Privacy Conflicts - A Wizard-of-Oz Study. *Proceedings of the ACM on Human-Computer Interaction*, 2023. doi: 10.1145/3579618.

- [338] Peter Leo Gorski, Yasemin Acar, Luigi Lo Iacono, and Sascha Fahl. Listen to Developers! A Participatory Design Study on Security Warnings for Cryptographic APIs. In *Proc. of the CHI Conf. on Human Factors in Computing Systems*, 2020. ACM. doi: 10.1145/3313831.3376142.
- [339] Scott Freeman, Sarah L. Eddy, Miles McDonough, Michelle K. Smith, Nnadozie Okoroafor, Hannah Jordt, and Mary Pat Wenderoth. Active learning increases student performance in science, engineering, and mathematics. *Proceedings of the National Academy of Sciences of the United States of America*, 2014. doi: 10.1073/pnas.1319030111.
- [340] Christie van Diggele, Annette Burgess, and Craig Mellis. Planning, preparing and structuring a small group teaching session. *BMC Medical Education*, 2020. doi: 10.1186/s12909-020-02281-4.
- [341] Rina Torchinsky. How period tracking apps and data privacy fit into a post-Roe v. Wade climate. *NPR*, 2022.
- [342] Saul Greenberg, Sheelagh Carpendale, Nicolai Marquardt, and Bill Buxton. The narrative storyboard: telling a story about use and context over time. *Interactions*, 2012. doi: 10.1145/2065327.2065340.
- [343] Lahari Goswami, Thibault Estier, Pegah Sadat Zeinoddin, and Mauro Cherubini. Supporting Collaboration in Introductory Programming Classes Taught in Hybrid Mode: A Participatory Design Study. 2023.
- [344] Raquel Benbunan-Fich. Usability of Wearables without Affordances. *Americas Conference on Information Systems, Boston*, 2017.
- [345] Karen Renaud and Judy van Biljon. Predicting Technology Acceptance and Adoption by the Elderly: A Qualitative study. *South African Institute of Computer Scientists & Information Technologists (SAICSIT)*, 2008.
- [346] feasibility, 2023. URL <https://dictionary.cambridge.org/dictionary/english/feasibility>.
- [347] effectiveness, 2023. URL <https://dictionary.cambridge.org/dictionary/english/effectiveness>.
- [348] Facebook Privacy Checkup | Facebook Help Centre, 2023. URL <https://www.facebook.com/help/443357099140264>.
- [349] Renita Washburn, Tangila Islam Tanni, Yan Solihin, Apu Kapadia, and Mary Jean Amon. Bottom-up psychosocial interventions for interdependent privacy: Effectiveness based on individual and content differences. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023. Association for Computing Machinery. doi: 10.1145/3544548.3581117.

- [350] Art. 17 GDPR – Right to erasure (‘right to be forgotten’), . URL <https://gdpr-info.eu/art-17-gdpr/>.
- [351] Jourdan Owen, Delano Archibald, and Damith Wickramanayake. The Willingness to Adopt Fitness Wearables in Jamaica: A Study on Wearable Fitness Trackers in Kingston and St. Andrew. *International Journal of Internet of Things*, 2019. doi: 10.5923/j.ijit.20190802.02.
- [352] Daniel A. Epstein, An Ping, James Fogarty, and Sean A. Munson. A lived informatics model of personal informatics. In *UbiComp*, 2015. Association for Computing Machinery. doi: 10.1145/2750858.2804250.
- [353] Lie Ming Tang, Jochen Meyer, Daniel A. Epstein, Kevin Bragg, Lina Engelen, Adrian Bauman, and Judy Kay. Defining Adherence: Making Sense of Physical Activity Tracker Data. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2018. doi: 10.1145/3191769.
- [354] Louis Faust, Priscilla Jiménez-Pazmino, James K. Holland, Omar Lizardo, David Hachen, and Nitesh V. Chawla. What 30 Days Tells Us About 3 Years: Identifying Early Signs of User Abandonment and Non-Adherence. In *Proc. of the Conf. on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, 2019. Association for Computing Machinery. doi: 10.1145/3329189.3329196.
- [355] James Clawson, Jessica A. Pater, Andrew D. Miller, Elizabeth D. Mynatt, and Lena Mamykina. No longer wearing: investigating the abandonment of personal health-tracking technologies on craigslist. In *UbiComp*, 2015. Association for Computing Machinery. doi: 10.1145/2750858.2807554.
- [356] risk, 2023. URL <https://dictionary.cambridge.org/dictionary/english/risk>.
- [357] Aarti Shahani. The Black Market For Stolen Health Care Data. *NPR*, 2015.
- [358] Thomas H. McCoy and Roy H. Perlis. Temporal Trends and Characteristics of Reportable Health Data Breaches, 2010-2017. *JAMA*, 2018. doi: 10.1001/jama.2018.9222.
- [359] Apple. Apple Vision Pro, 2023. URL <https://www.apple.com/apple-vision-pro/>.



# Appendix A

## A.1 Questionnaire of Chapter 3

### STUDY

You are invited to participate in a research survey about **fitness-data sharing**. The survey takes approximately 25 minutes to complete. For an optimal experience, **we recommend you take the survey on a computer with your smartphone close to you.**

### PARTICIPATION CRITERIA

To be eligible for this study you must:

- **regularly** use an **activity tracker** (i.e., a wrist-worn device that collects personal fitness data like step counts, activities, and/or heart rate) from **Apple, Fitbit, or Garmin,**
- have your activity tracker paired with an **iOS or Android smartphone,**
- use the **official** companion **app** (i.e., Apple Health, Fitbit, Garmin Connect) - preferably in English, and,
- have **granted access to** your **fitness data** to a **third-party app** (i.e., an app that is **not** provided by the company that manufactured your device).

### YOUR RIGHTS

You will be paid **\$5** for your complete participation in the study. You may choose to terminate your participation in this study at any time and for any

reason. In this case, however, you will not be compensated and your data will be deleted. If you participate, your answers will be kept **confidential**. Also, we do not collect personally identifying information such as your name and e-mail address. All data will be **stored** on a **secured** server and only researchers participating in this study will have access to it. The results of this research study might be published in scientific journals or conferences. Any published information will be **aggregated** and/or **anonymized**.

### CONSENT

If you wish to participate in this research study, please select the “**Agree**” option to continue. It will indicate that you are eligible for this study, that you will answer all questions truthfully, and that you consent that we use the collected data under the conditions stated above. If you select “**Disagree**” you will not participate in this research survey and will not be paid.

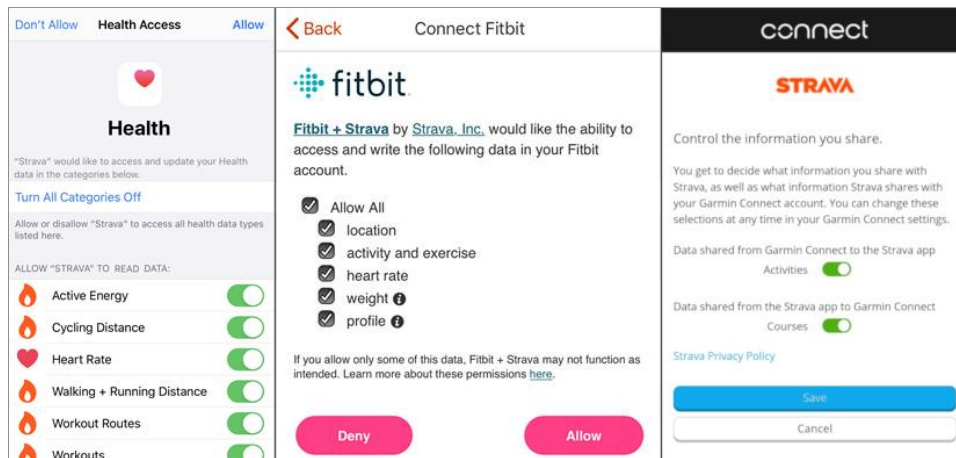
- Agree
- Disagree

### Screening - sharing

Have you ever **granted access** to your fitness data to any **third-party app**?

A **third-party app** is an app that is **not** provided by the company that manufactured your device.

For example, the **Strava app qualifies** as a third-party app: you can grant it reading access to your data collected by devices from Apple, Fitbit, and Garmin. The picture below shows granting Strava access to your device's companion app's data. **Apple Health/Fitbit/Garmin app does not qualify** as a third-party app.



- Yes  
 No

### device infos

**How many days** per week on average do you wear your main activity-tracker?

- 0  
 1  
 2  
 3  
 4  
 5  
 6  
 7

What is the **brand** of your main activity tracker?



- Garmin
- Apple
- Fitbit

What type of **smartphone** have you **paired** with your main activity tracker?

- iOS (Apple)
- Android

#### **Block B - Device Information (device type and usage)**

From now on please keep your **smartphone** close to yourself (the one with which your fitness tracker **is paired**).

Some questions will explicitly ask you to check your **app settings**.

From now on please keep your **smartphone** close to yourself (the one with which your fitness tracker **is paired**).

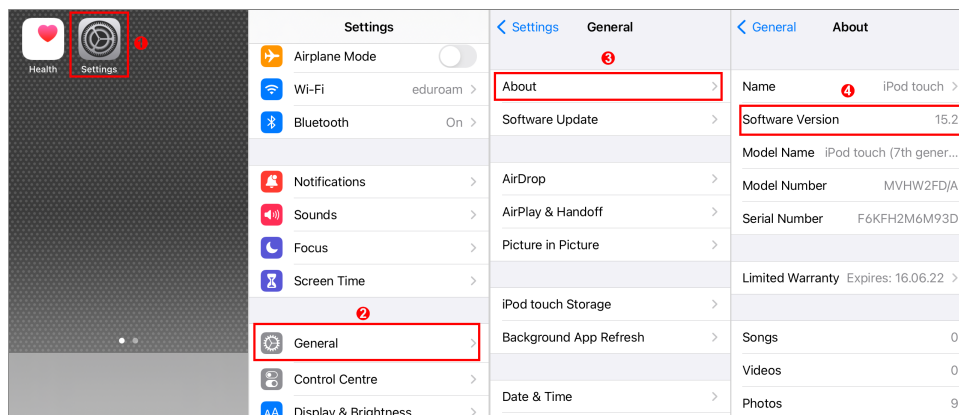
Prepare your Fitbit **account credentials** (we will not ask you the credentials, only to connect to your Fitbit account).

Some questions will explicitly ask you to check your **app settings**.

What **version** of **iOS** is your device using?

Follow the instructions below:

1. Tap on "Settings"
2. Tap on "General" (scroll down if needed)
3. Tap on "About"
4. Look at the current "Software Version"

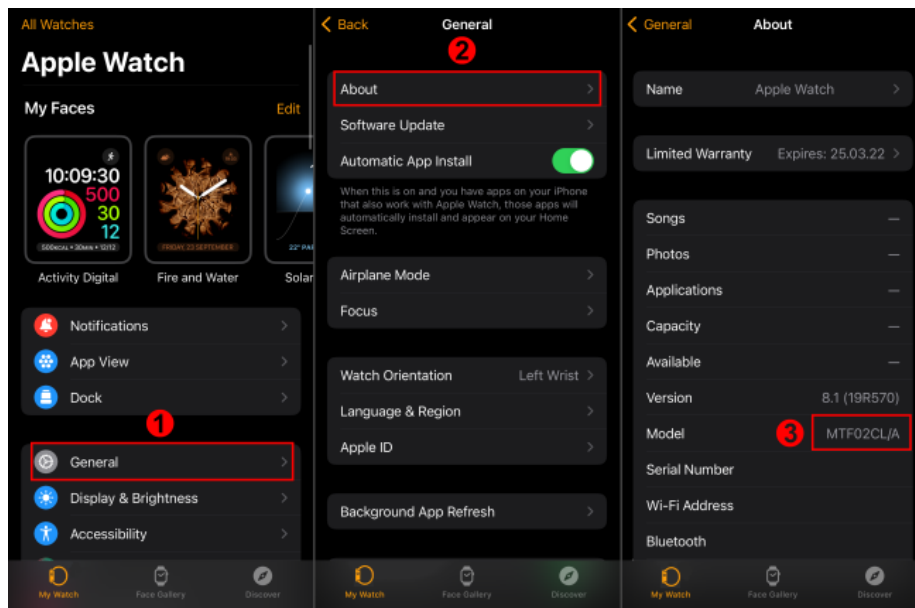


- v.15 or newer (e.g., v.15.2)
- v.14.9 or older (e.g., v.13)

What is the **model** of your **Apple Watch**, please indicate the model series number?

Open Apple Watch app and follow the instructions below:

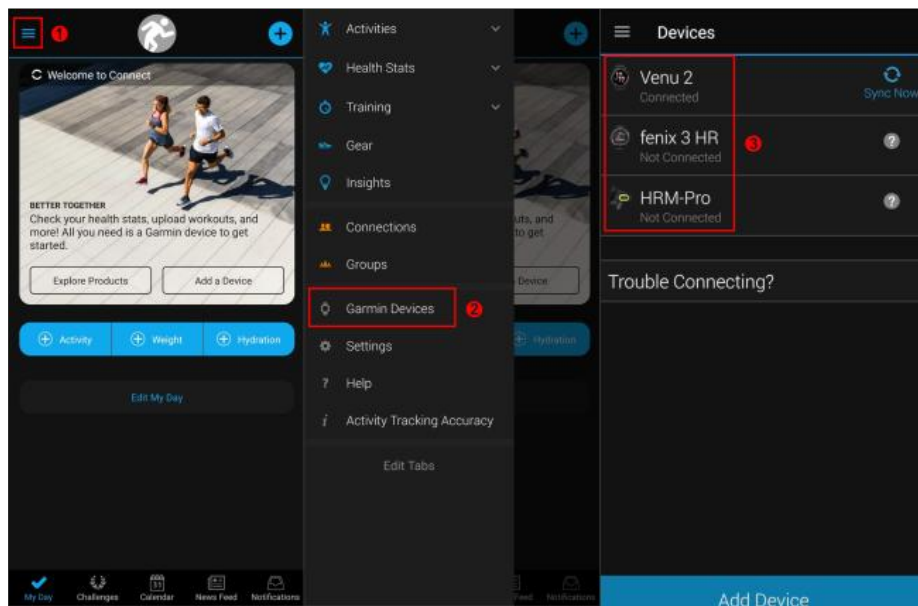
1. Tap on "General"
2. Tap on "About"
3. The series number is indicated in the "Model" section



What is the **model** of your **main activity tracker**?

Open the Garmin Connect app and follow the instructions below:

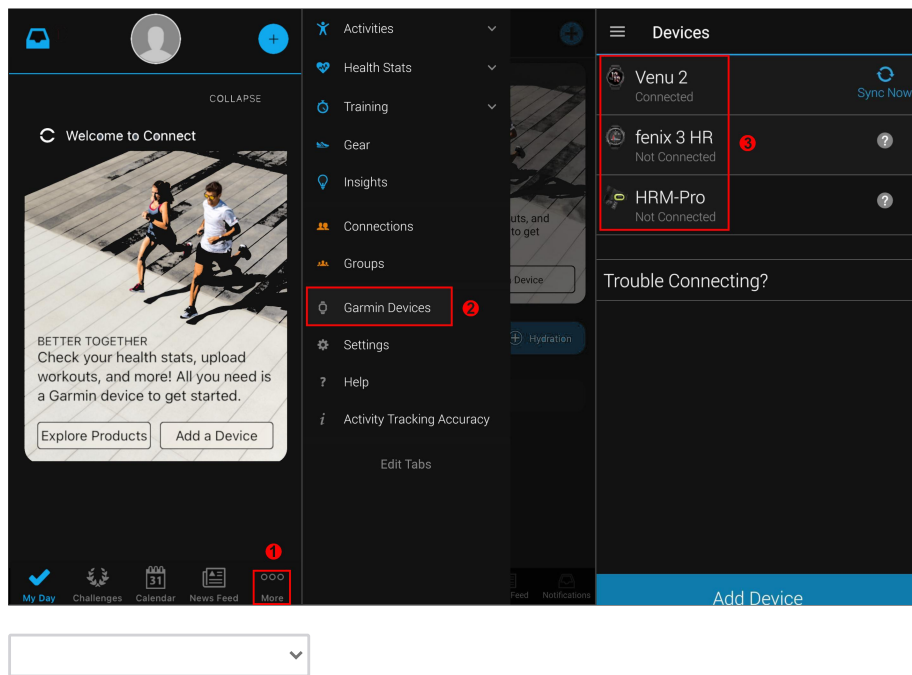
1. Tap on the left top of the screen (the three lines)
2. Tap on "Garmin Devices"
3. Please indicate the device you **use the most**



What is the **model** of your **main activity tracker**?

Open the Garmin Connect app and follow the instructions below:

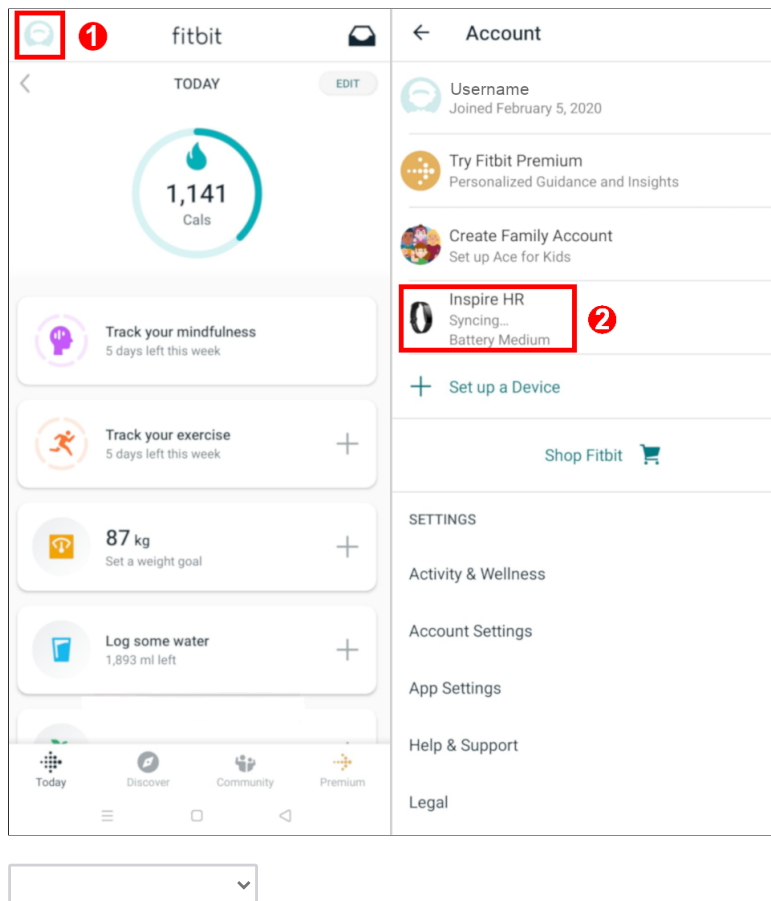
1. Tap on the left top of the screen (the three lines)
2. Tap on "Garmin Devices"
3. Please indicate the device you **use the most**



What is the model of your **main activity tracker**?

Open the Fitbit app and follow the instructions below:

1. Tap on your profile picture (top left)
2. Please indicate the device you use **the most**



Which of the following **functionalities** of your fitness tracker do you use?  
*Select all that apply.*

- Steps tracking
- Heart-rate tracking
- Sleep tracking
- Activity tracking
- Calorie tracking
- Stress monitoring
- None

**Block D - Access via Public Profile**

**Off the top of your head (i.e., without checking on your smartphone),** which of the following data is visible on your Fitbit **profile**?

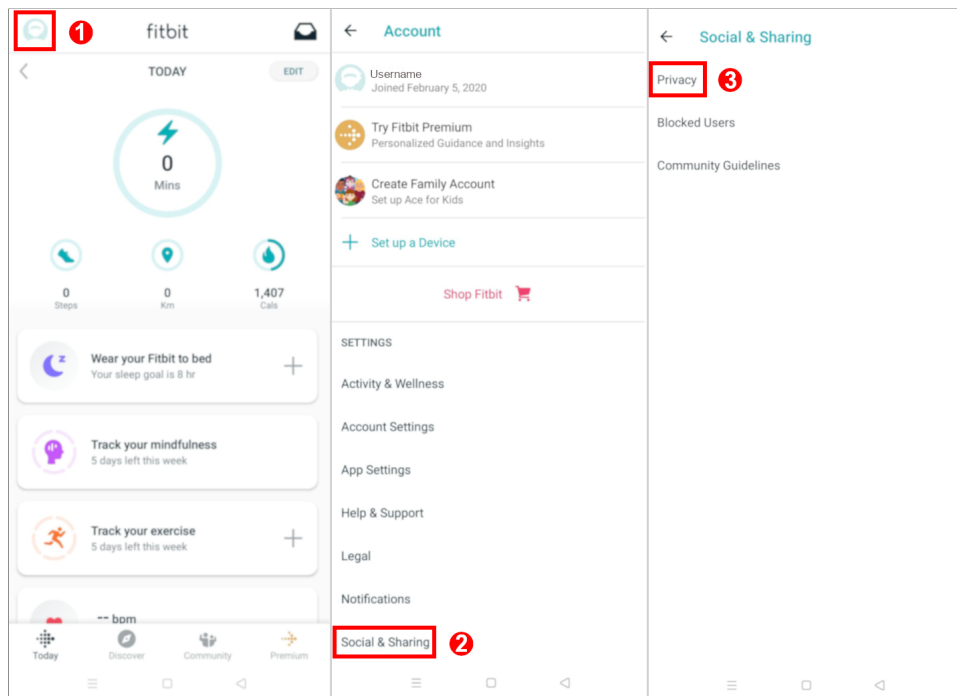
*Select all that apply.*


- Birthday
- Sex
- Height
- Weight
- Location
- My Friends
- Badges & Trophies
- Lifetime Steps, Distance, and Floors
- Average Daily Step Count
- None

Now please **check in your Fitbit app** (see instructions below) and indicate the **privacy settings** for each type of data.

Open the Fitbit app and follow the instructions below:

1. Tap on the top left corner (where the profile picture is)
2. Tap on "Social & Sharing"
3. Tap on "Privacy"
4. Tap on each data type and check privacy settings (Private, Friends, or Public)



|   |  Private |  Friends |  Public |
|---|---|---|--|
| Birthday                                | <input type="radio"/>   | <input type="radio"/>   | <input type="radio"/>  |
| Sex                                     | <input type="radio"/>   | <input type="radio"/>   | <input type="radio"/>  |
| Height                                  | <input type="radio"/>   | <input type="radio"/>   | <input type="radio"/>  |
| Weight                                  | <input type="radio"/>   | <input type="radio"/>   | <input type="radio"/>  |
| Location                                | <input type="radio"/>   | <input type="radio"/>   | <input type="radio"/>  |
| My Friends                              | <input type="radio"/>   | <input type="radio"/>   | <input type="radio"/>  |
| Badges & Trophies                       | <input type="radio"/>   | <input type="radio"/>   | <input type="radio"/>  |
| Lifetime Steps,<br>Distance, and Floors | <input type="radio"/>   | <input type="radio"/>   | <input type="radio"/>  |
| Average Daily Step<br>Count             | <input type="radio"/>   | <input type="radio"/>   | <input type="radio"/>  |



**Off the top of your head (i.e., without checking on your smartphone),** which one of the following data is publicly visible on your Garmin Connect **profile?**

*Select all that apply.*

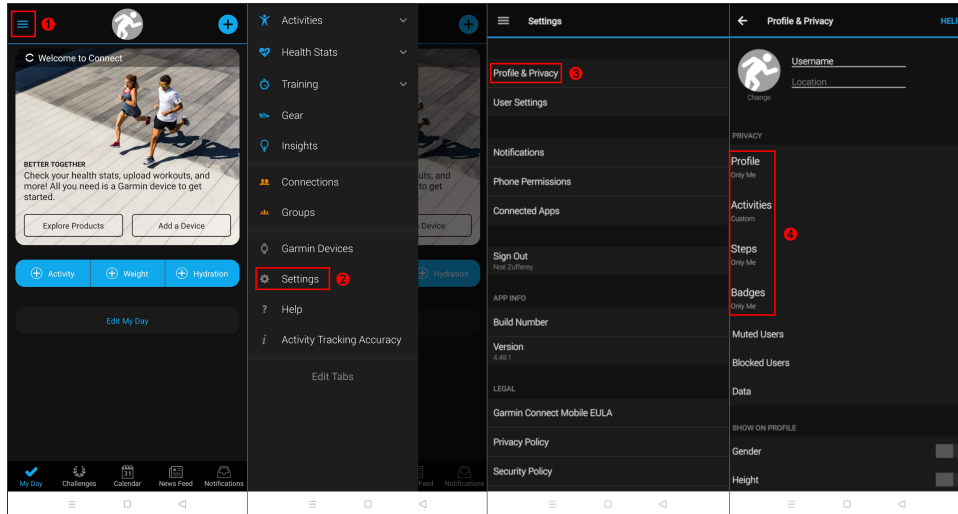
- Gender
- Height
- Weight
- Age
- VO<sub>2</sub> max
- Personal records
- Lifetime Totals
- Last 12 months
- Garmin device
- Segments leaderboard
- Running
- Cycling
- Walking
- Swimming
- Gym & fitness equipment
- Multisport
- Diving
- Winter sports
- Hiking
- Other (sports)
- None

Now please **check in your Garmin Connect app** (see instructions below) and indicate the **privacy levels** for each type of data.

Open the Garmin Connect app and follow the instruction below:

1. Tap on the left top of the screen (the three lines)

2. Tap on "Settings"
3. Tap on "Profile & Privacy"
4. Check which data type is shared with which level (e.g., "Everyone")

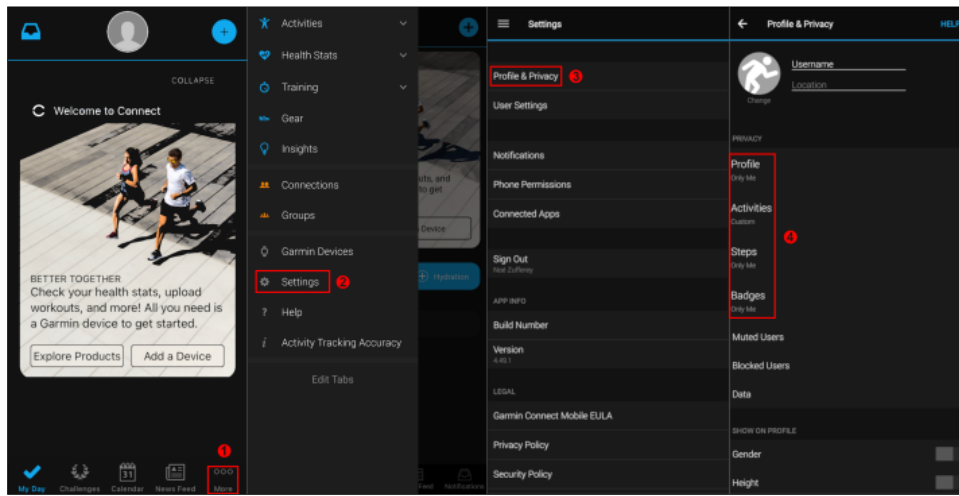


|            | Only me               | My connections        | My connections & groups | Everyone              | Custom                |
|------------|-----------------------|-----------------------|-------------------------|-----------------------|-----------------------|
| Profile    | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>   | <input type="radio"/> | <input type="radio"/> |
| Activities | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>   | <input type="radio"/> | <input type="radio"/> |
| Steps      | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>   | <input type="radio"/> | <input type="radio"/> |
| Badges     | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>   | <input type="radio"/> | <input type="radio"/> |

Now please **check in your Garmin Connect app** (see instructions below) and indicate the **privacy levels** of each type of data.

Open the Garmin Connect app and follow the instruction below:

1. Tap on "More" at the bottom right of the screen (the three dots)
2. Tap on "Settings"
3. Tap on "Profile & Privacy"
4. Check which data type is shared with which setting (e.g., "Everyone")



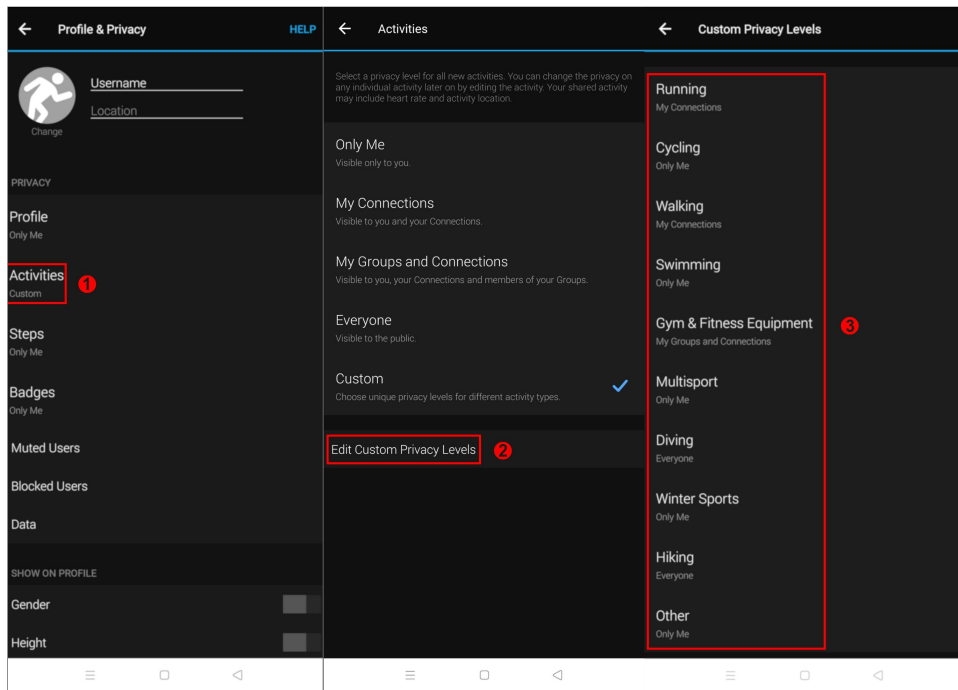
|            | Only me               | My connections        | My connections & groups | Everyone              | Custom                |
|------------|-----------------------|-----------------------|-------------------------|-----------------------|-----------------------|
| Profile    | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>   | <input type="radio"/> | <input type="radio"/> |
| Activities | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>   | <input type="radio"/> | <input type="radio"/> |
| Steps      | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>   | <input type="radio"/> | <input type="radio"/> |
| Badges     | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>   | <input type="radio"/> | <input type="radio"/> |

In the previous question you said that you defined **custom** privacy settings for activities.

Please check your "Activities" privacy settings and select the **privacy levels** for each activity type.

Follow the instructions below:

1. In the same "Profile & Privacy" settings, tap on "Activities" (second option from the top)
2. Tap on "Edit Custom Privacy Levels"
3. Check which type data is shared with which setting (e.g., "Everyone")

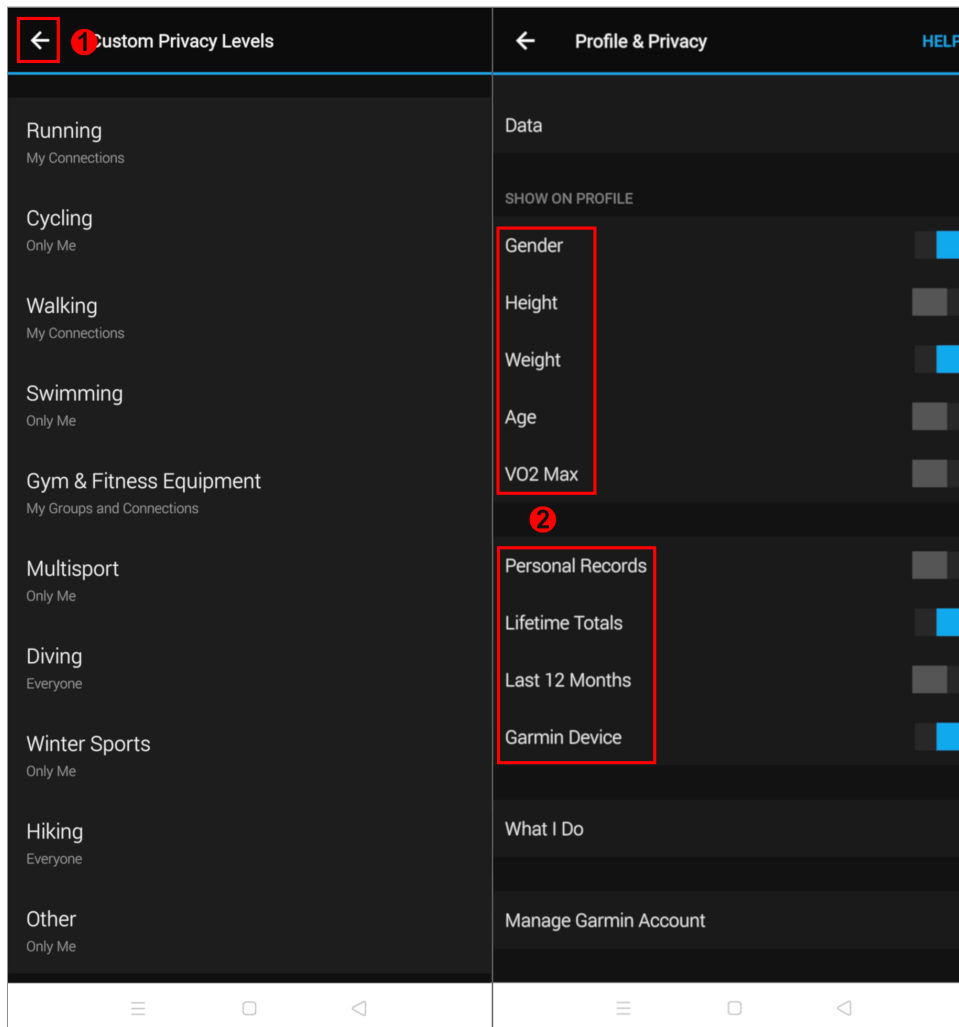


|                         | Only me               | My connections        | My connections and groups | Everyone              |
|-------------------------|-----------------------|-----------------------|---------------------------|-----------------------|
| Running                 | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>     | <input type="radio"/> |
| Cycling                 | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>     | <input type="radio"/> |
| Walking                 | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>     | <input type="radio"/> |
| Swimming                | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>     | <input type="radio"/> |
| Gym & fitness equipment | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>     | <input type="radio"/> |
| Multisport              | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>     | <input type="radio"/> |
| Diving                  | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>     | <input type="radio"/> |
| Winter sports           | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>     | <input type="radio"/> |
| Hiking                  | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>     | <input type="radio"/> |
| Other                   | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>     | <input type="radio"/> |

Please check your "Profile & Privacy" settings and select among the following types of data the ones that are **shown** on your profile.

Follow the instructions below:

1. Go back to "Profile & Privacy" settings, slide down to the "Show on profile" section
2. Select all the following data that are available according to your settings



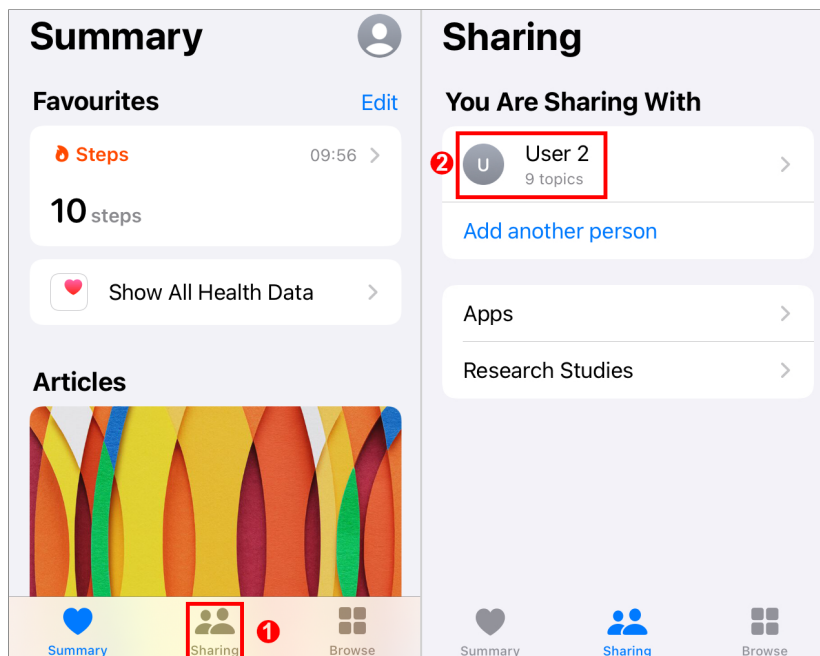
Gender

- Height
- Weight
- Age
- VO2 Max
- Personal Records
- Lifetime Totals
- Last 12 Months
- Garmin Device
- None

With which types of **individuals** do you generally share the following type of data?

Open the Apple Health app and follow the instructions below:

1. Tap the "Sharing" tab (bottom)
2. You can check the sharing settings for each of your "sharing contacts"



Please select the **types of individuals** that you share your data with.  
 Select if you share with **at least one** corresponding individual.  
**Do not** select anything if you do not share your data with anyone.

|                          | Friends                  | Family                   | Co-workers               | Health Professional      | Employers                | Strangers                |
|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| Activity related data    | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Heart related data       | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Mindfulness related data | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Mobility related data    | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

With **how many** individuals of each of the following types do you share **at least one** type of data?

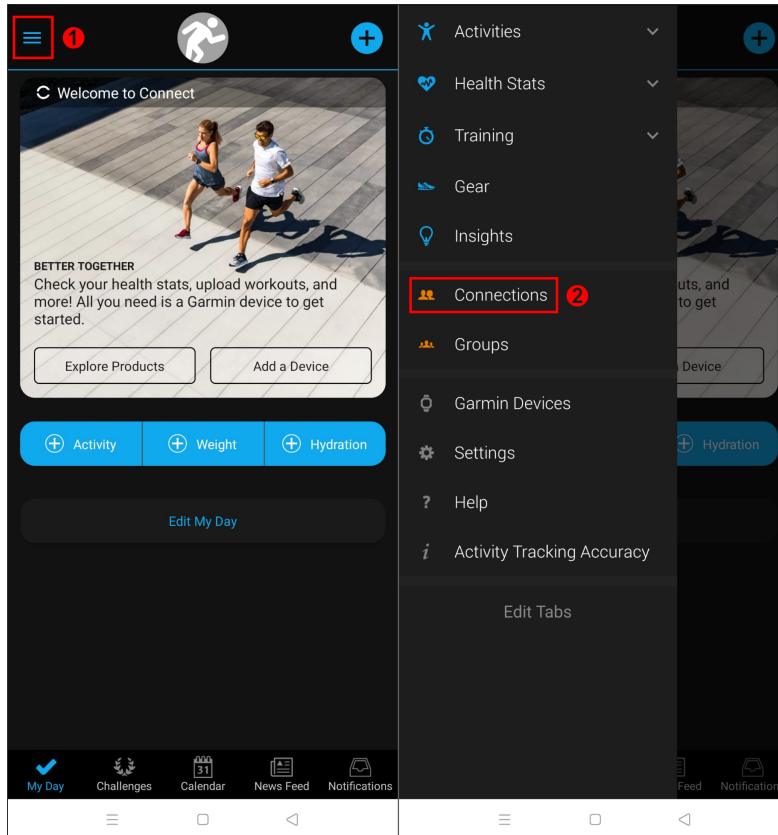
Please select "0" when **no one** corresponds to the concerned type of individuals.

|                     | 0                     | 1-2                   | 3-4                   | 5-7                   | More than 7           |
|---------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| Friends             | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Family              | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Co-workers          | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Health Professional | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Employers           | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Strangers           | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

**How many** individuals of each of the following types do you have as Garmin "Connections"?

Open the Garmin Connect app and follow the instructions below:

1. Tap on the left top of the screen (the three lines)
2. Tap on "Connections"



Please select "0" when **no one** corresponds to the concerned type of individuals.

|                     | 0                     | 1-2                   | 3-4                   | 5-7                   | More than<br>7        |
|---------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| Friends             | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Family              | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Co-workers          | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Health Professional | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

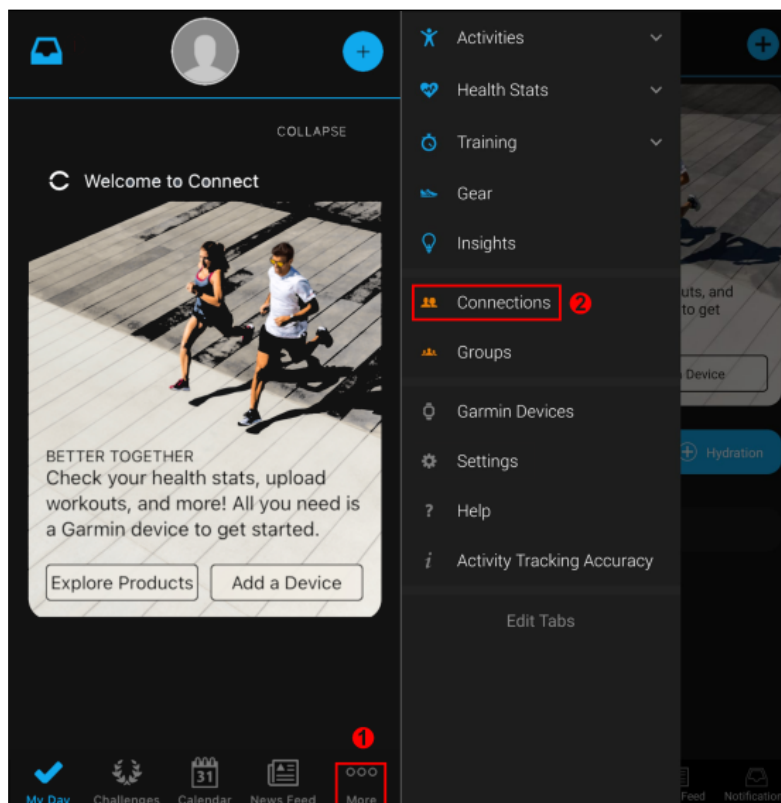


|           | 0                     | 1-2                   | 3-4                   | 5-7                   | More than 7           |
|-----------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| Employers | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Strangers | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

**How many** individuals of each of the following types do you count in your Garmin "**connections**"?

Open the Garmin Connect app and follow the instructions below:

1. Tap on "More" at the bottom right of the screen (the three dots)
2. Tap on "Connections"



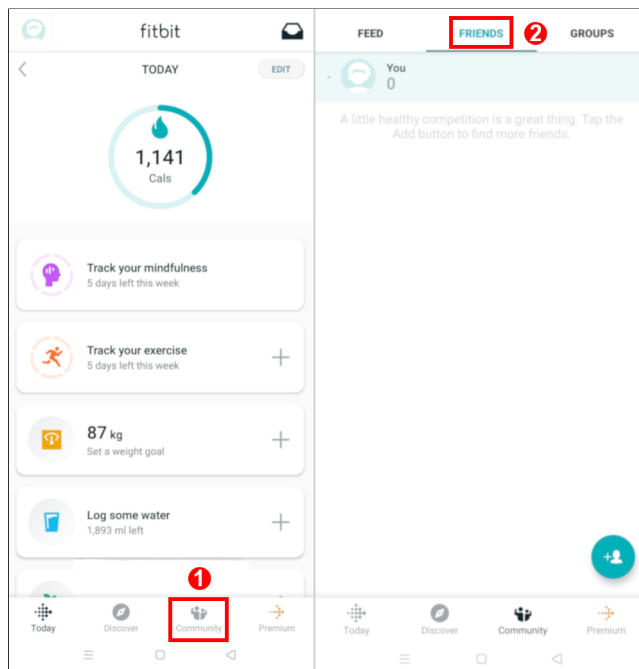
Please select "0" when **no one** corresponds to the concerned type of individuals.

|                     | 0                     | 1-2                   | 3-4                   | 5-7                   | More than<br>7        |
|---------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| Friends             | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Family              | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Co-workers          | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Health Professional | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Employers           | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Strangers           | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

**How many** individuals of each of the following types do you count in your Fitbit "**friends**"?

Open the Fitbit app and follow the instructions below:

1. Tap on the bottom "Community"
2. Tap on the top "Friends"
3. Specify your relationship with the people on the list



Please select "0" when **no one** corresponds to the concerned type of individuals.

|                     | 0                     | 1-2                   | 3-4                   | 5-7                   | More than 7           |
|---------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| Friends             | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Family              | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Co-workers          | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Health Professional | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Employers           | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Strangers           | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

**Do you share your fitness data as part of a health program (e.g., with your employer and/or health insurance company)?**

- Yes  
 No

Have you ever modified the default **privacy settings** to change the availability of some of the data on your **profile**?

- Yes. I **increased** the availability of some of the data
- Yes. I **decreased** the availability of some of the data
- No

#### Bloc 5

**Off the top of your head (i.e., without checking on your smartphone), how many third-party apps** currently have access to your fitness data?

- 0
- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10+

#### Bloc 6

**Off the top of your head (i.e., without checking on your smartphone),** please select the **third-party** apps that currently have access to your fitness data.

- Samsung Health
- MyFitnessPal
- Google Fit
- Strava Running and Cycling
- Weight Watchers Mobile
- Lose It!
- Planet Fitness
- Noom Weight
- SleepIQ
- RENPHO
- ...others (one app per line)

**Off the top of your head (i.e., without checking on your smartphone),** please select the **third-party** apps that currently have access to your fitness data.

- MyFitnessPal
- Sweatcoin
- Achievement
- Flo period & Ovulation Tracker
- Weight Watchers Mobile
- Sleep Watch
- Lose It!
- Planet Fitness
- Noom Weight

- Strava Running and Cycling
- ...others (one app per line)

### Block E - Given Authorizations and Control

Do you think it is possible to **revoke** (cancel) previously granted third-party app access?

- No
- Yes

Imagine that you **granted access** to a **third-party app** and you agreed to share **all the data that it is possible to share**.

Select which of the following data you think the third party app **has access** to.

- Steps
- Sleep data
- Stress level
- E-mail for your activity tracker account
- Username for your activity tracker account
- Password for your activity tracker account
- Birthdate
- Weight and height
- Physical activities

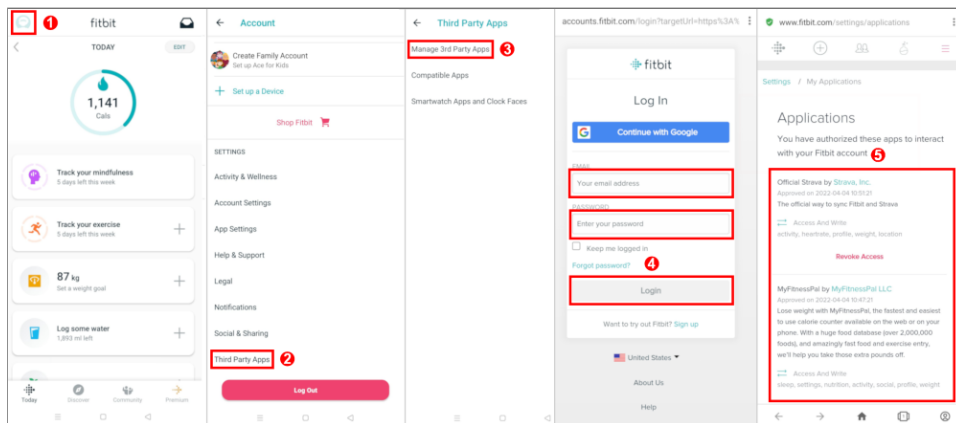
- Location
- Gender
- Menstrual cycle
- None

Please open your Fitbit app and check **how many third-party apps** currently **have access to your fitness data**.

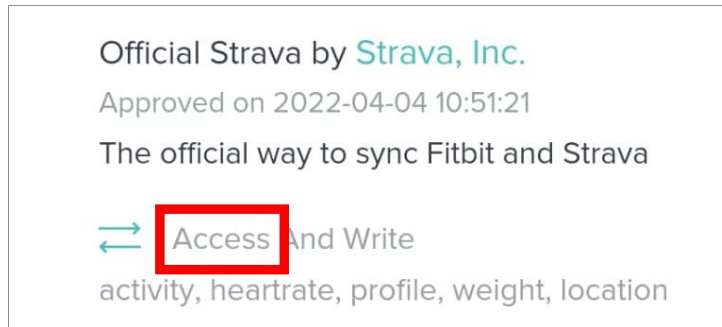
**For this question, you will need your Fitbit credentials.**

Open the Fitbit app and follow the instructions below:

1. Tap on the top left corner (where the profile picture is)
2. Tap on "Third-Party Apps" (Last option)
3. Tap on "Manage 3rd party apps"
4. If requested, login to your profile
5. Count the number of 3rd party apps with access, this number can be 0.



Please only take into account the third-party apps that **have access to your data**, like in the example below (do **not** take into account apps that **only** have writing authorization).




- 0
- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10+

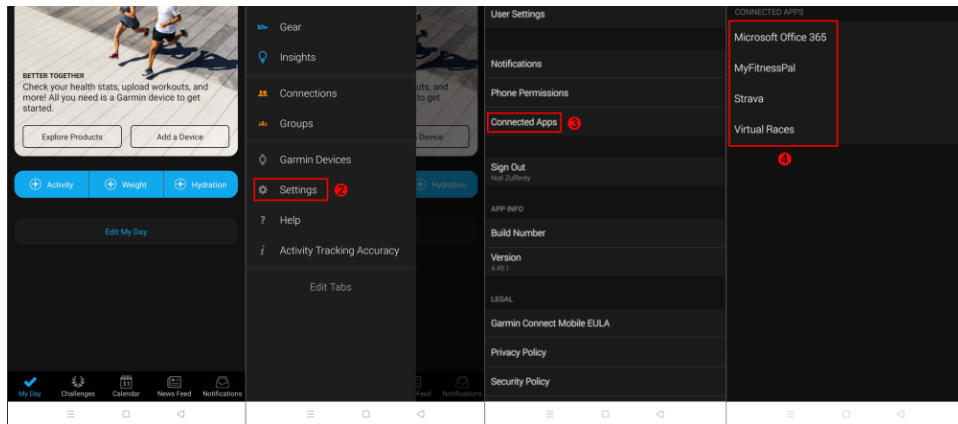
Please open your Garmin Connect app and check **how many third-party apps** currently have access to your fitness data.

Open the Garmin Connect app and follow the instructions below:

1. Tap on the top left of the screen (the three lines)
2. Tap on "Settings"
3. Tap on "Connected Apps"
4. Count the number of 3rd party apps with access (connected apps), this number can be 0.

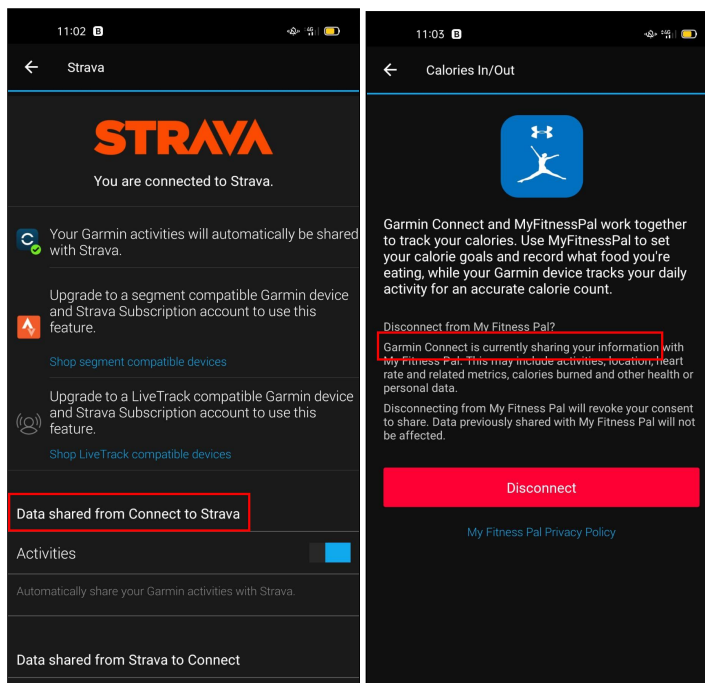






Please only take into account the third-party apps **with which Garmin Connect shares data**.

To verify if Garmin Connect indeed share data with a given third-party app, tap on this app in the list and it w






- 0
- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10+

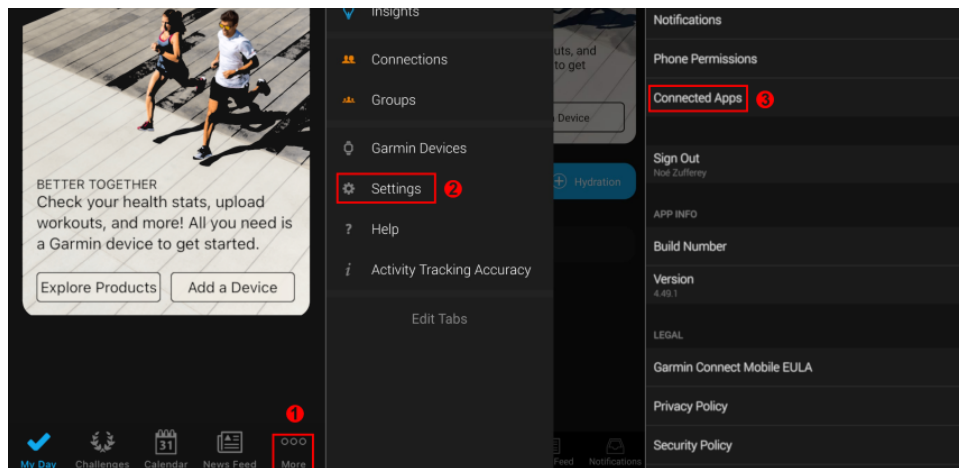
---

Please open your Garmin Connect app and check **how many third-party apps** currently have access to your fitness data.

Open the Garmin Connect app and follow the instructions below:

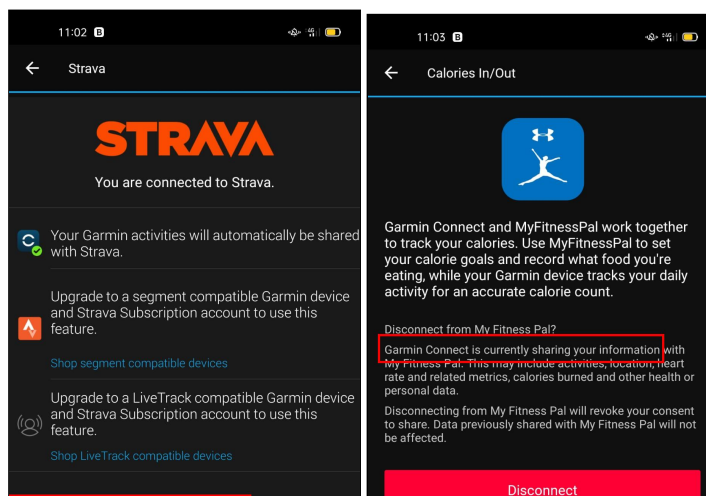
1. Tap on "More" at the bottom right of the screen (the three dots)
2. Tap on "Settings"
3. Tap on "Connected Apps"
4. Count the number of 3rd party apps with access (connected apps), this number can be 0.

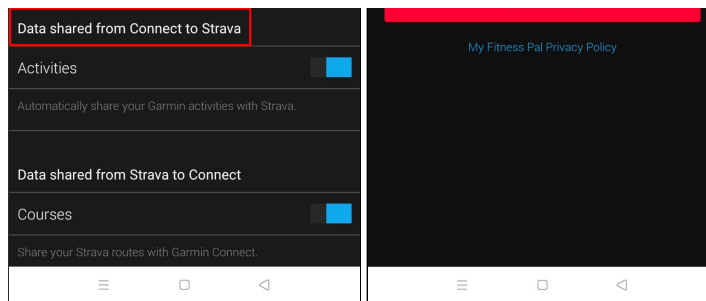




Please only take into account the third-party apps with which **Garmin Connect** shares data.

To verify if Garmin Connect indeed share data with a given third-party app, tap on this app in the list and it w





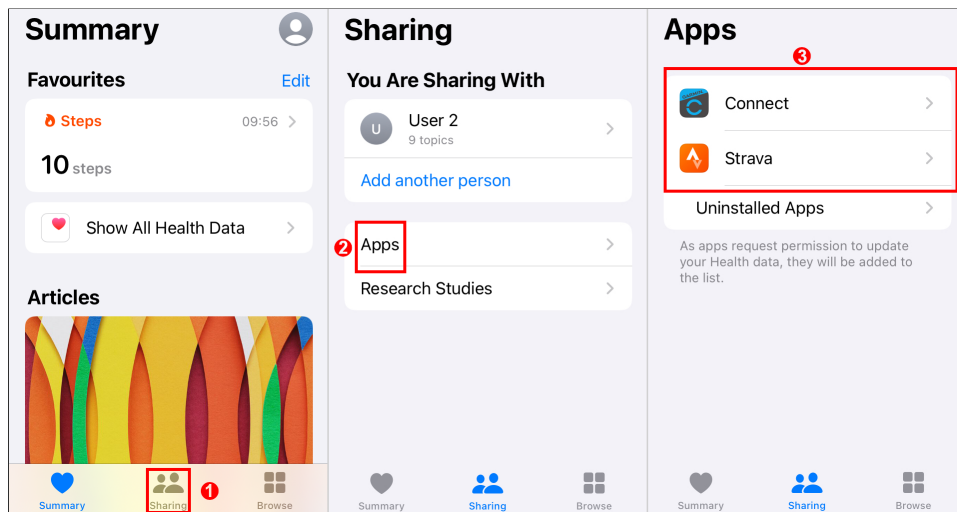
- 0
- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10+

---

Please open the Health app and check **how many third-party apps** currently have access to your fitness data.

Open the Apple Health app and follow the instructions below:

1. Open your profile (top right corner)
2. Tap on "Apps" or "Apps AND Uninstalled Apps" (depending on the version)
3. Count the number of apps with access, this number can be 0.

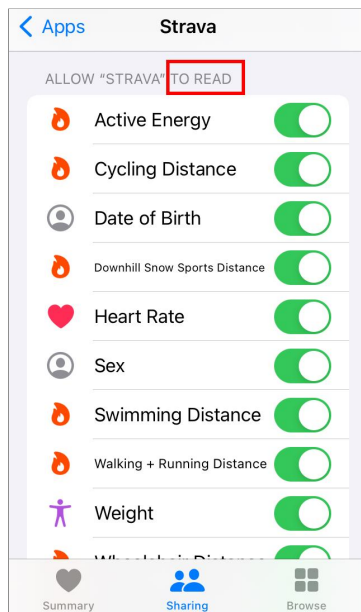


Please **do not consider** apps that are classified as "**Uninstalled Apps**".

Please only take into account the third-party apps **with which Apple Health shares data**.

To verify if Apple Health indeed share data with a given third-party app, **tap on this app in the list** and it will specify it (see the examples below).

You may have to scroll down to find it.



- 0
- 1
- 2
- 3
- 4
- 5
- 6
- 7
- 8
- 9
- 10+

**After checking your  $\{q://QID7/ChoiceGroup/SelectedChoices\}$  app,** please list the names of the **third-party apps** that currently have access to your fitness data.

- >> Samsung Health
- >> MyFitnessPal

- >> Google Fit
- >> Strava Running and Cycling
- >> Weight Watchers Mobile
- >> Lose It!
- >> Planet Fitness
- >> Noom Weight
- >> SleepIQ
- >> RENPHO
- >> ...others (one app per line)

**After checking your `{q://QID7/ChoiceGroup/SelectedChoices}` app,** please list the names of the **third-party apps** that currently have access to your fitness data.

- >> MyFitnessPal
- >> Sweatcoin
- >> Achievement
- >> Flo period & Ovulation Tracker
- >> Weight Watchers Mobile
- >> Sleep Watch
- >> Lose It!
- >> Planet Fitness
- >> Noom Weight
- >> Strava Running and Cycling

- >> ...others (one app per line)

**After checking your Apple Health app**, please list the names of the **third-party apps** that currently have access to your fitness data.

- >> MyFitnessPal
- >> Sweatcoin
- >> Achievement
- >> Flo period & Ovulation Tracker
- >> Weight Watchers Mobile
- >> Sleep Watch
- >> Lose It!
- >> Planet Fitness
- >> Noom Weight
- >> Strava Running and Cycling
- >> ...others (one app per line)



Are you still **actively** using all these apps?

- Yes, I actively use all of these apps.
- I actively use most of these apps.
- I actively use only some of these apps.
- No, I actively use none of these apps.

Please select the **third-party apps** that currently **have access** to your fitness data and that you use **actively**.

- >> Samsung Health
- >> MyFitnessPal
- >> Google Fit
- >> Strava Running and Cycling
- >> Weight Watchers Mobile
- >> Lose It!
- >> Planet Fitness
- >> Noom Weight
- >> SleepIQ
- >> RENPHO
- >> ...others (one app per line)

Please select the **third-party apps** that currently **have access** to your fitness data and that you use **actively**.

- >> MyFitnessPal
- >> Sweatcoin
- >> Achievement
- >> Flo period & Ovulation Tracker
- >> Weight Watchers Mobile
- >> Sleep Watch
- >> Lose It!
- >> Planet Fitness
- >> Noom Weight
- >> Strava Running and Cycling
- >> ...others (one app per line)

How often have you **revoked** third-party app access to your fitness data?

- Never
- Only once
- 2-5 times
- 6-10 times
- More than 10 times

Please select the **third-party apps** whose access you **revoked**.

- >> Samsung Health
- >> MyFitnessPal
- >> Google Fit

- >> Strava Running and Cycling
- >> Weight Watchers Mobile
- >> Lose It!
- >> Planet Fitness
- >> Noom Weight
- >> SleepIQ
- >> RENPHO
- >> ...others (one app per line)

Please select the **third-party apps** whose access you **revoked**.

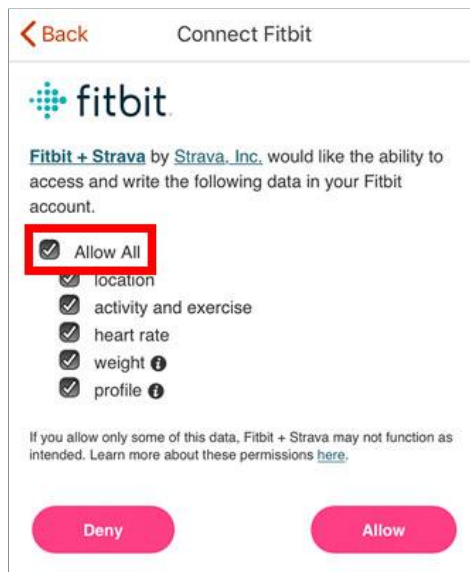
- >> MyFitnessPal
- >> Sweatcoin
- >> Achievement
- >> Flo period & Ovulation Tracker
- >> Weight Watchers Mobile
- >> Sleep Watch
- >> Lose It!
- >> Planet Fitness
- >> Noom Weight
- >> Strava Running and Cycling

>> ...others (one app per line)

Why did you **revoke** access to those third-party apps?

You mentioned that you are **not actively** using one or several **third-party apps** that currently **have access** to your fitness data. Please explain why you **did not revoke** their access.

Usually, when you **grant access** to **third-party apps**, how do you select the **types** of data that you want to share?

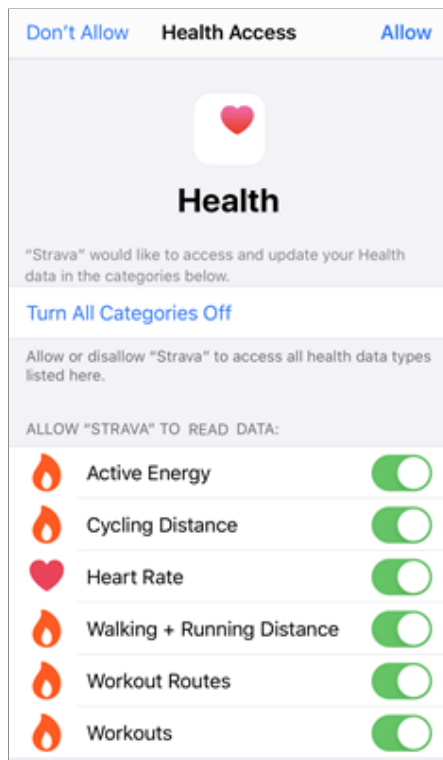


*In the above example, you can tap on "Allow All" to grant access to all requested data or select only the data you want to share.*

*Do not focus on the data types in the example, please provide your answer regarding all types of data that are collected by your device.*

- I share everything
- I share everything only when it is necessary to use the app otherwise I share selectively
- I share selectively

Usually, when you **grant access** to **third-party apps**, how do you select the **types** of data that you want to share?

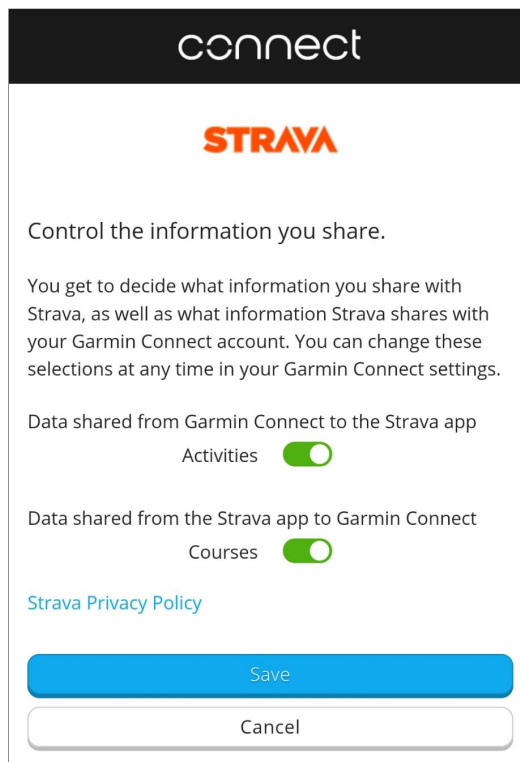


*In the above example, all types of data are selected "by default" and you can unselect the data you do not want to share.*

*Do not focus on the data types in the example, please provide your answer regarding all types of data that are collected by your device.*

- I share everything
- I share everything only when it is necessary to use the app otherwise I share selectively
- I share selectively

Usually, when you **grant access** to **third-party apps**, how do you select the **types** of data that you want to share?



*In the above example, all types of data are selected "by default" and you can unselect the data you do not want to share.*

*Do not focus on the data types in the example, please provide your answer regarding all types of data that are collected by your device.*

- I share everything
- I share everything only when it is necessary to use the app otherwise I share selectively
- I share selectively

During the access-granting process, **how likely** would you use additional sharing options related to the **precision** of the data, as in the example below?

< Back Data sharing

Strava by Strava, Inc. would like the ability to access the following data in your account.

Allow All

Steps

- Not rounded (e.g., 6243 steps)
- Rounded to the tens (e.g., 6240 steps)
- Rounded to the hundreds (e.g., 6200 steps)
- Rounded to the thousands (e.g., 6000 steps)

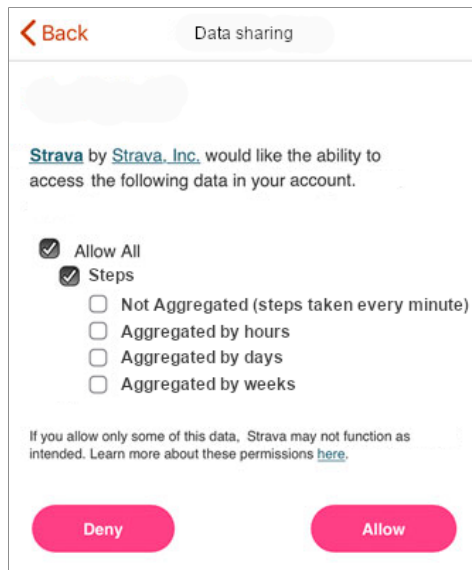
If you allow only some of this data, Strava may not function as intended. Learn more about these permissions [here](#).

Deny Allow

- Extremely likely
- Likely
- Slightly likely
- Neutral
- Slightly unlikely
- Unlikely
- Extremely unlikely

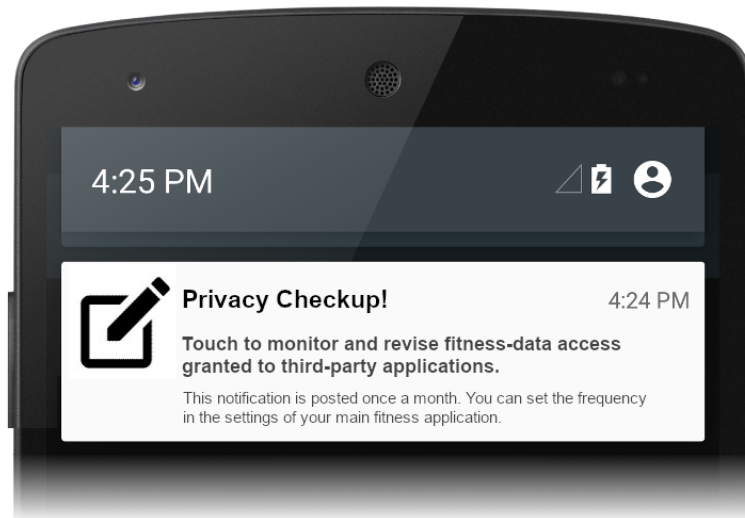
During the access-granting process, **how likely** would you use additional sharing options related to **time granularity** of the data, as in the example below?





- Extremely likely
- Likely
- Slightly likely
- Neutral
- Slightly unlikely
- Unlikely
- Extremely unlikely

**How likely** would you use notification functionalities to periodically (e.g., every three months) **remind you to monitor** access granted to third-party apps, as shown in the example below?



- Extremely likely
- Likely
- Slightly likely
- Neutral
- Slightly unlikely
- Unlikely
- Extremely unlikely

### Bloc F

Assume that you **granted** access to your fitness data to a **third-party app**. For each of the following statements, please answer if you think that they are **true** or **false**.

|  | True                  | False                 |
|--|-----------------------|-----------------------|
| The 3rd party is able to access the fitness data that was <b>collected</b> (by my fitness tracker) <b>before</b> I granted access. | <input type="radio"/> | <input type="radio"/> |

|  | True                  | False                 |
|--|-----------------------|-----------------------|
| The 3rd party is able to access the fitness data that was <b>collected</b> (by my fitness tracker) <b>after</b> I granted access.  | <input type="radio"/> | <input type="radio"/> |
| The 3rd party is able to store on their own servers any data they have access to.  | <input type="radio"/> | <input type="radio"/> |
| The 3rd party app is <b>legally</b> allowed - according to the <b>federal laws</b> in force in the United States - to <b>store</b> any data they have access to on their own servers.  | <input type="radio"/> | <input type="radio"/> |
| The 3rd party app is <b>legally</b> allowed - according to <code>\$(q://QID7/ChoiceGroup/SelectedChoices)</code> 's <b>terms of service</b> - to <b>store</b> any data they (the 3rd party app) have access to on their own servers. | <input type="radio"/> | <input type="radio"/> |

Assume that you **revoked** the access previously granted to a **third-party app**.

For each of the following statements, please answer if you think that they are **true** or **false**.

|   | True                  | False                 |
|---|-----------------------|-----------------------|
| The 3rd party will be able to access the data <b>collected after revoking</b> , using the previously granted authorization.   | <input type="radio"/> | <input type="radio"/> |
| The 3rd party will be able to access the data <b>collected before revoking</b> , if they stored it on their own servers.  | <input type="radio"/> | <input type="radio"/> |
| The 3rd party will still be able to access the data <b>collected before revoking</b> , using the previously granted authorization (without storing it on their own server). | <input type="radio"/> | <input type="radio"/> |

How difficult do you find it to **monitor** or **revoke** access granted to third-party apps?

- Very easy
- Easy
- Moderately easy
- Neutral
- Moderately difficult
- Difficult
- Very difficult

What are your suggestions to **facilitate** the process of **monitoring**, **granting**, or **revoking** the access to **third-party apps**?

### IUIPC

Please indicate to what extent you **agree** with each of the the following statements.

|                      |          |                        |                                     |                     |       |                   |
|----------------------|----------|------------------------|-------------------------------------|---------------------|-------|-------------------|
| Strongly<br>disagree | Disagree | Moderately<br>disagree | Neither<br>agree<br>nor<br>disagree | Moderately<br>agree | Agree | Strongly<br>Agree |
|----------------------|----------|------------------------|-------------------------------------|---------------------|-------|-------------------|



|   | Strongly disagree     | Disagree              | Moderately disagree   | Neither agree nor disagree | Moderately agree      | Agree                 | Strongly Agree        |
|---|-----------------------|-----------------------|-----------------------|----------------------------|-----------------------|-----------------------|-----------------------|
| It bothers me to give personal information to so many online companies.                     | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>      | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| I am concerned that online companies are collecting too much personal information about me. | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/>      | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

### Last demographics

**For how long** have you been using your main activity tracker?

- Less than 1 month
- 1 to 12 months
- 1 to 3 years
- 3 to 5 years
- More than 5 years

On average, **how many hours** per day do you wear your main activity tracker?

- 1-6
- 7-12
- 13-18

19-24

With which gender do you identify the most?

- Woman
- Man
- Non-binary
- Prefer to self describe
- Prefer not to answer

### Mental Model

Last, we have an **optional** question for you. If you answer this question, you will have the opportunity to participate in a **draw to win a bonus payment of \$10 (in addition to the initial 5\$)**. **One in five people (i.e., 20%) will be chosen as winners** by a random draw.

You will need to **draw a picture** representing how you think the **access granting** to third-party apps is processed, and how your **fitness data** is transferred between different entities.

Please choose one of these options:

- I will draw the picture
- I prefer not to send any drawings and to skip this question

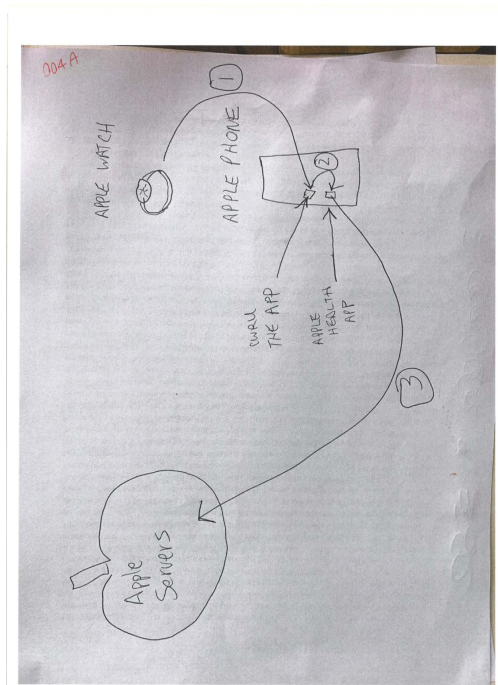
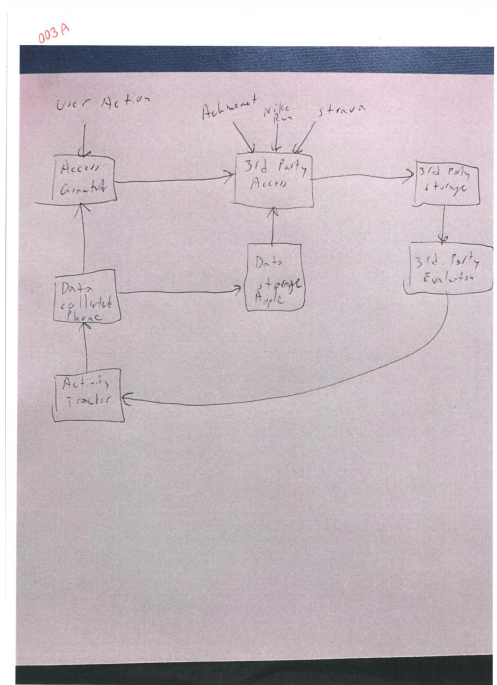


Before drawing the picture, please read **carefully** the following instructions:



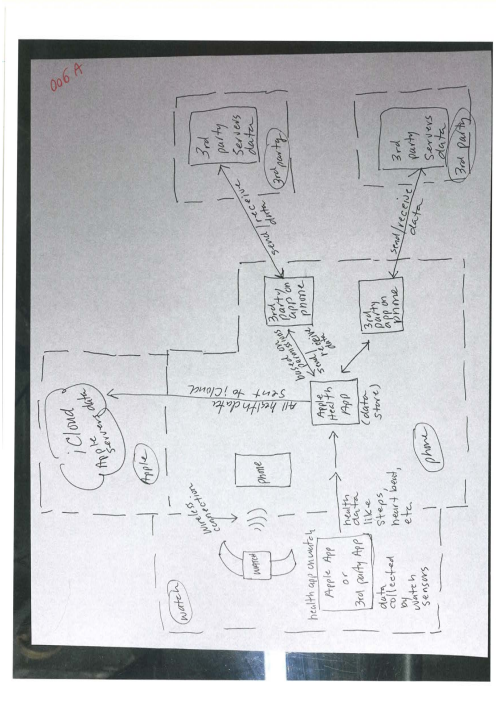

1. Do not spend more than 5 minutes on the drawing.


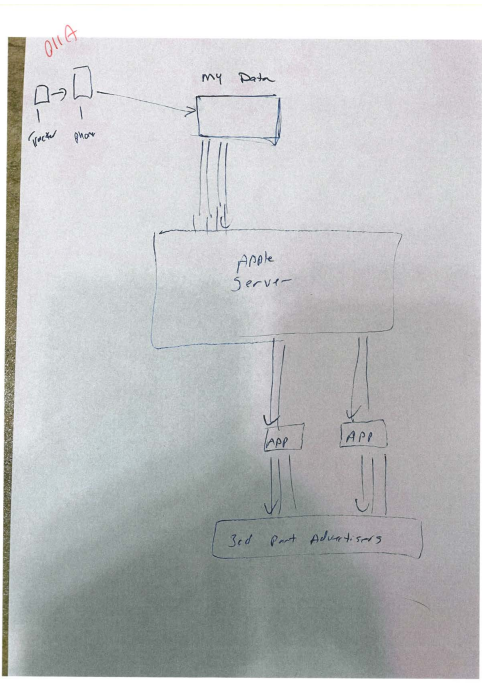
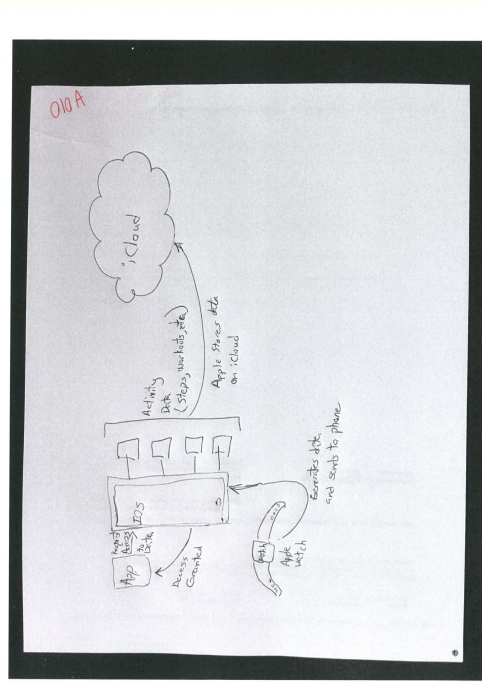

2. Please take a clean and **white sheet of paper**.
3. Use a pencil, pen, or ideally color pens for your drawing.
4. When drawing consider these two questions: **(i)** how do you think **access granting** to third-party apps is processed? **(ii)** how do you think your **fitness data** is transferred between different entities?
5. Include all relevant elements in your drawing, including **(i)** your activity tracker, **(ii)** your smartphone, **(iii)** `q://QID7/ChoiceGroup/SelectedChoices`'s servers, **(iv)** the apps (i.e., your service provider's app and third-party apps that have access to your fitness data), **(v)** any other elements you think are relevant to be illustrated.
6. You can use **arrows** and **lines** to connect these entities to each other.
7. Please use text to **label** the entities and their relationships with each other.
8. Use your smartphone to **take a picture** of your drawing. Please try your best to take a good-quality photo such that you can clearly read the labels and see drawn entities.
9. Please note that we will **not judge** your drawing skills and technical understanding.
10. Upload your picture here.




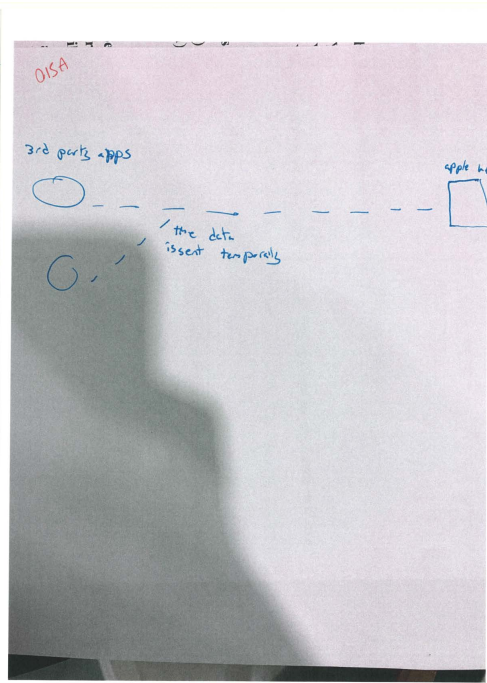

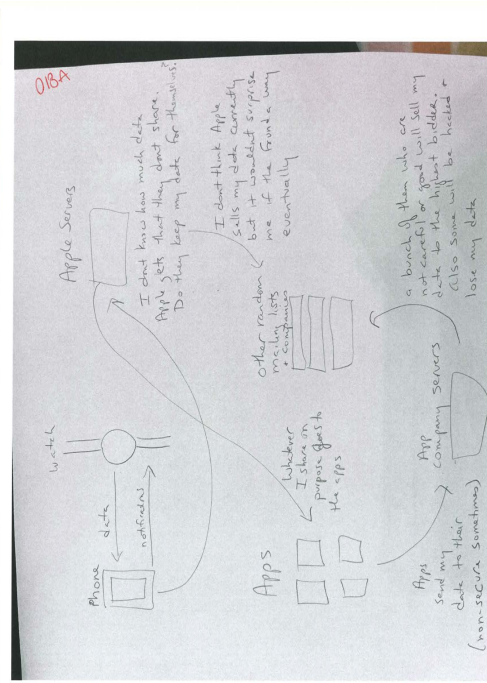
# A.2 All Mental Models

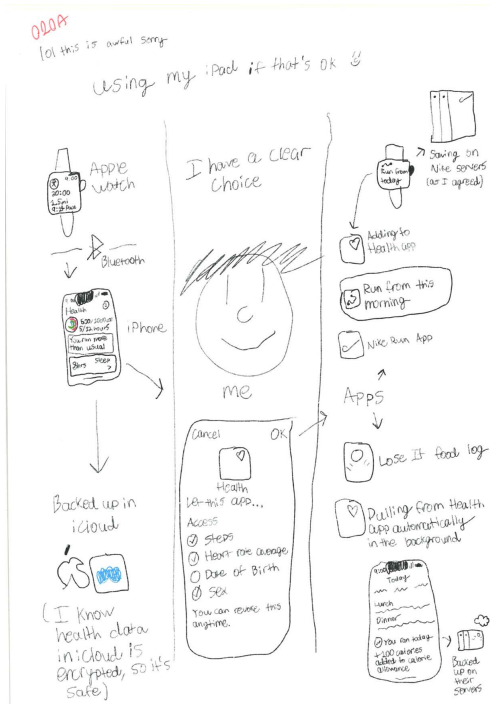

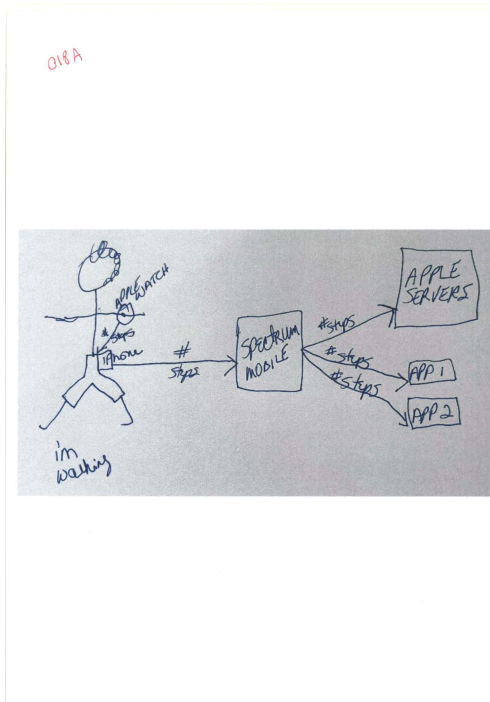
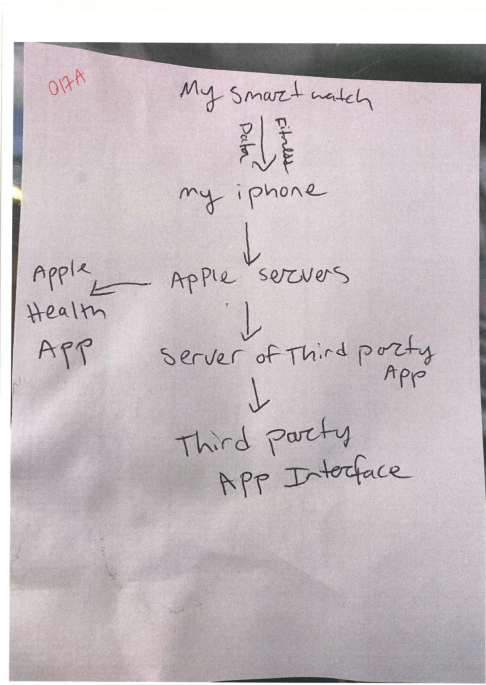


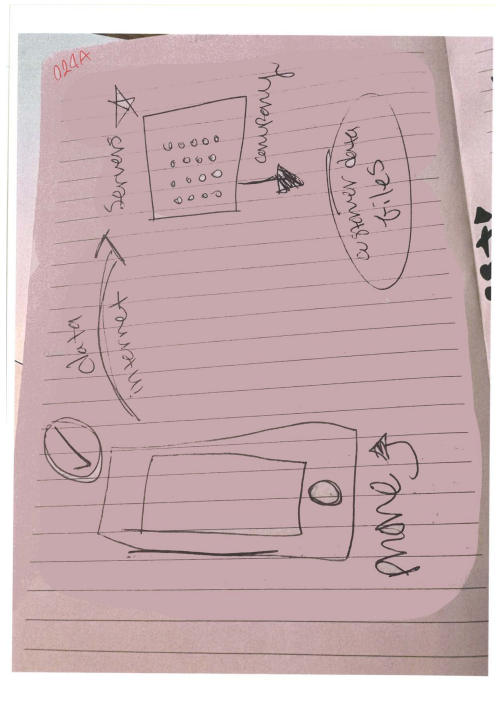
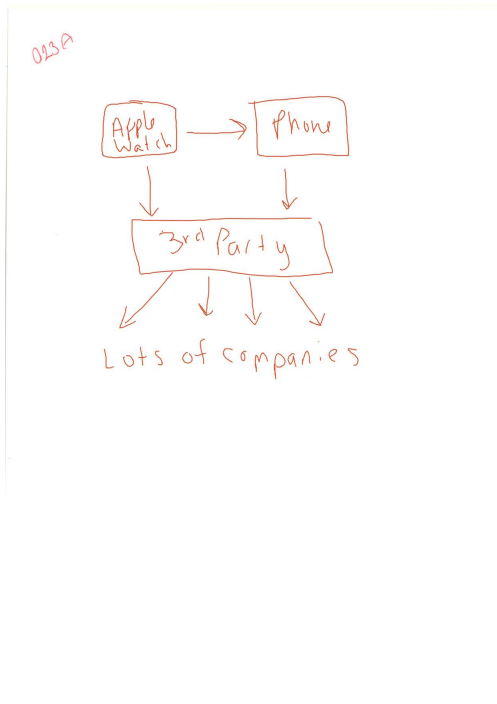

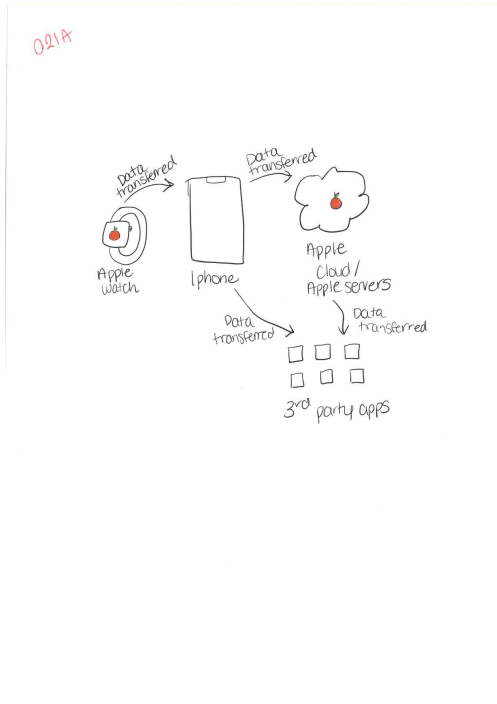




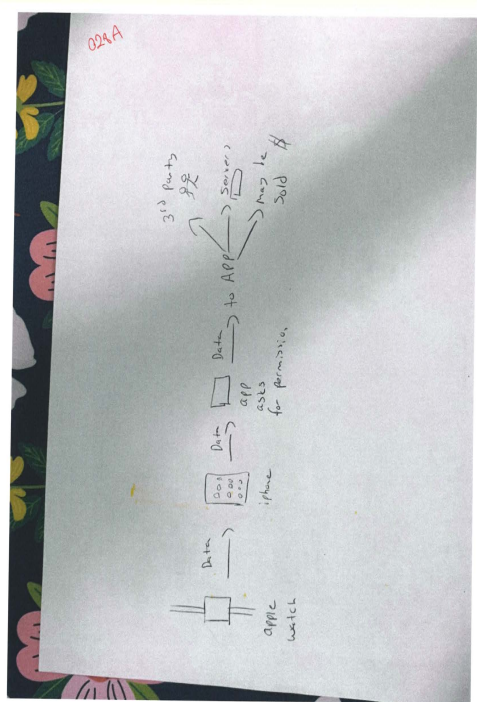
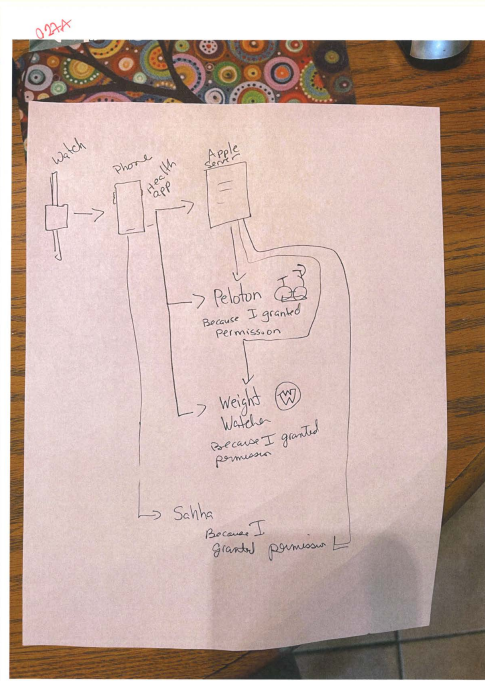
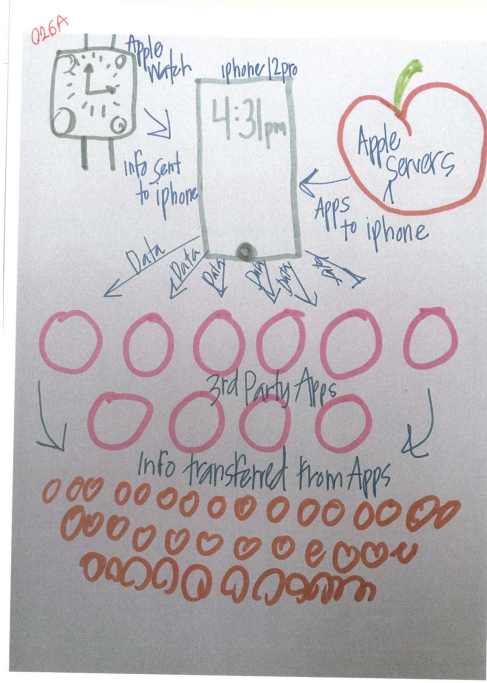
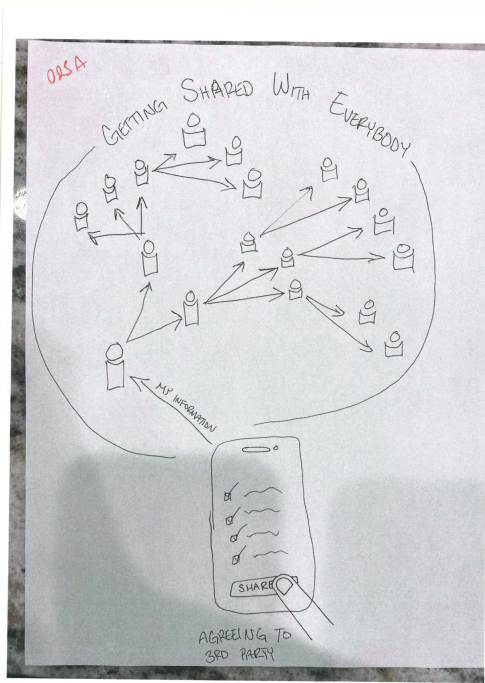



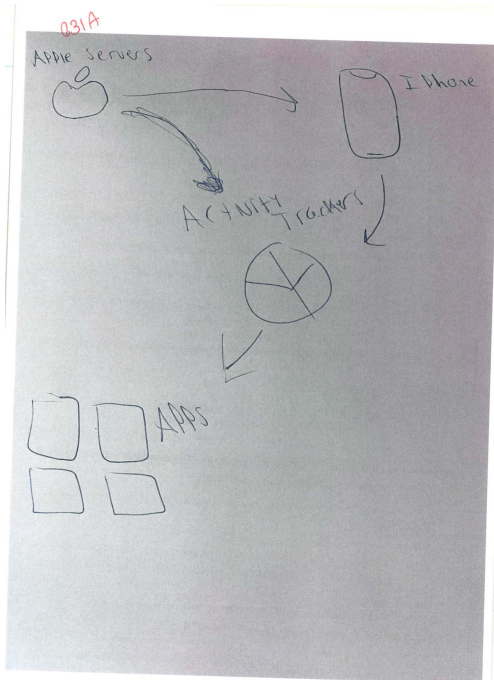







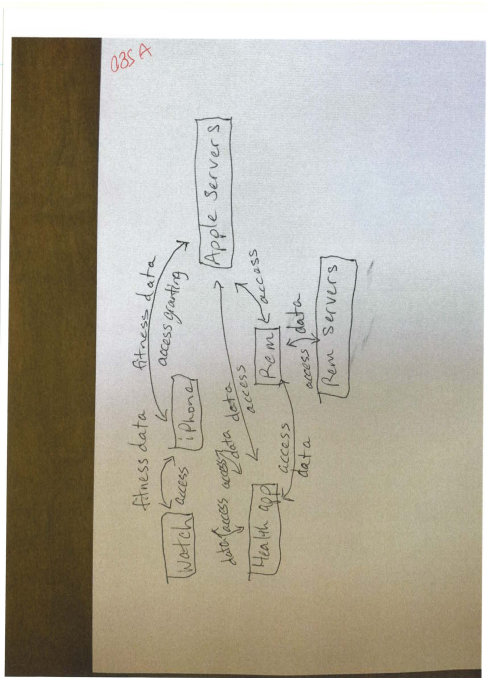
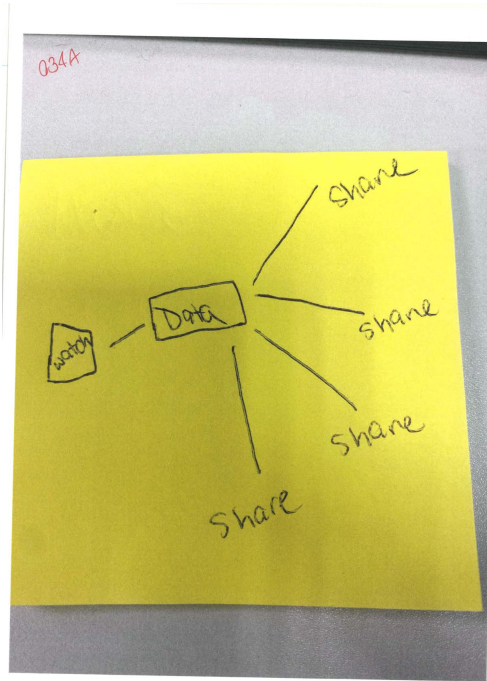
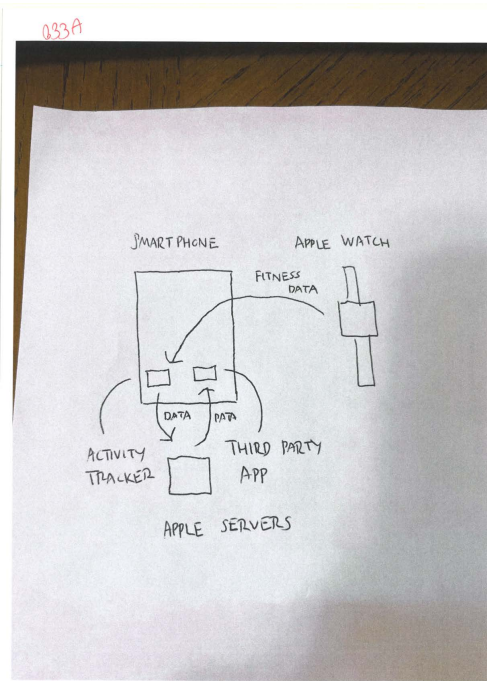






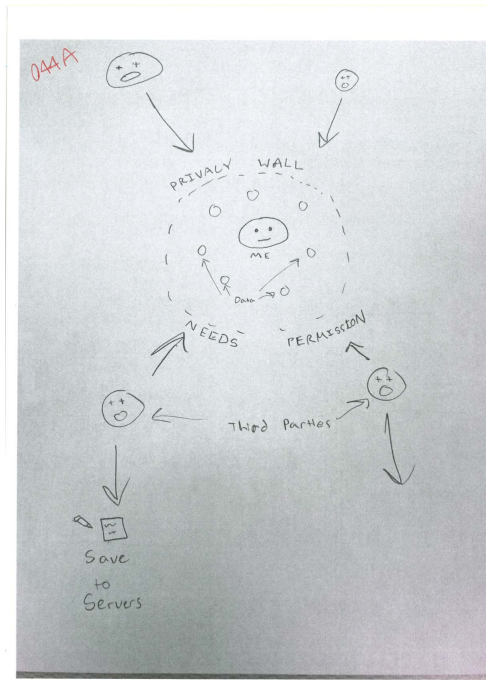
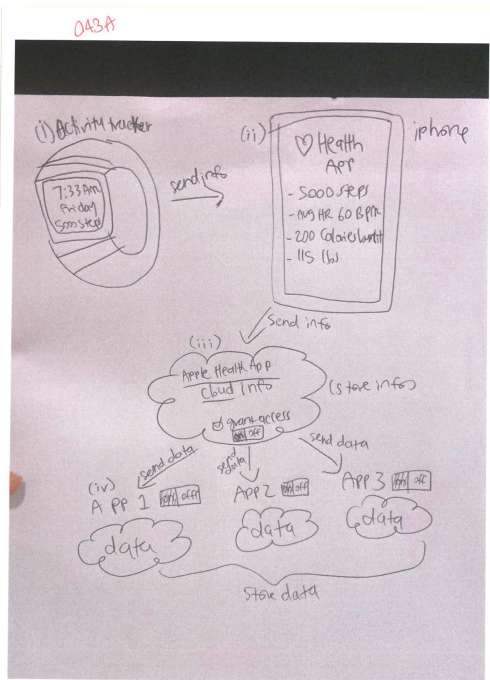
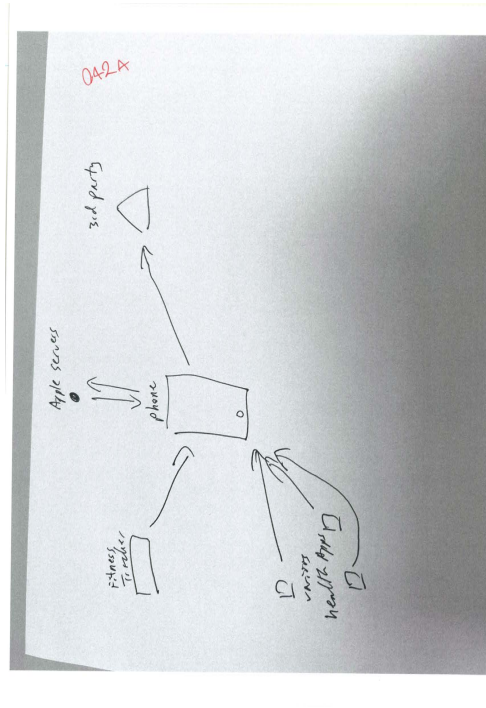
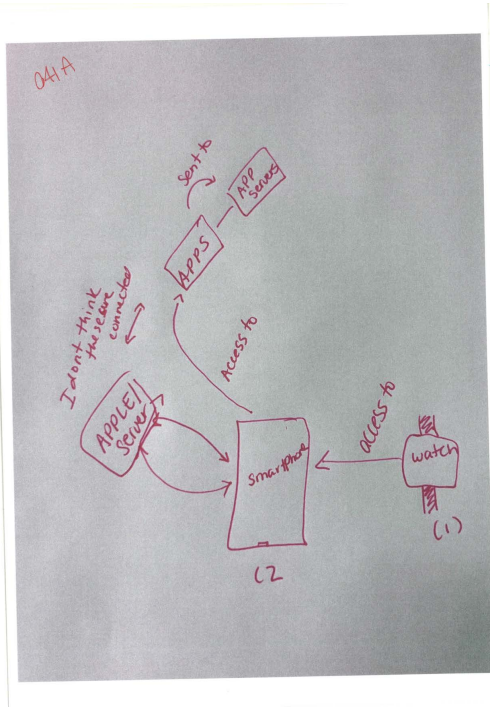


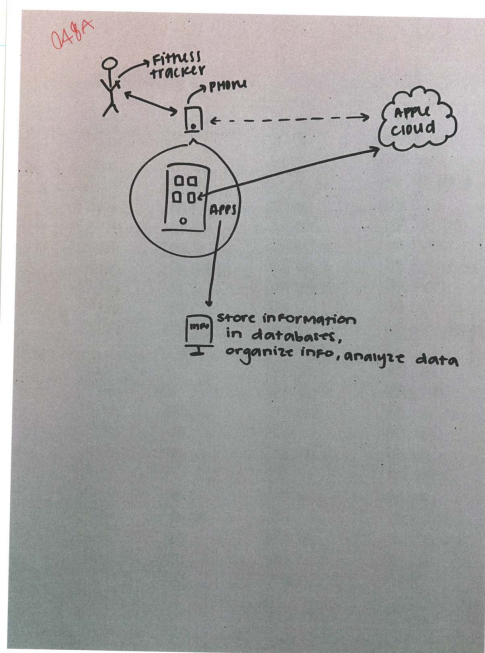






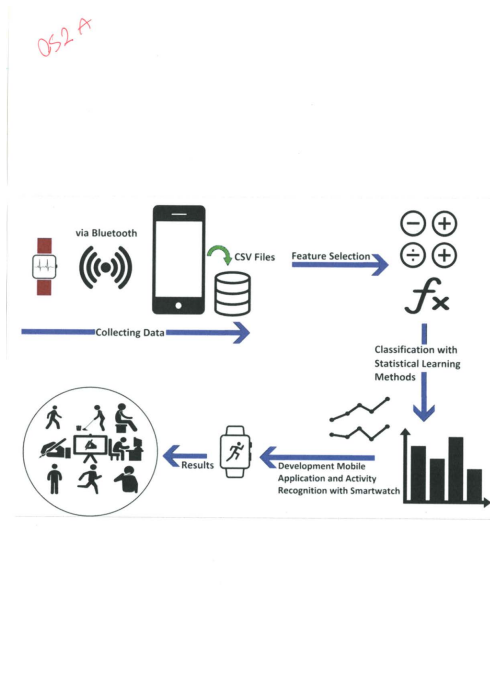
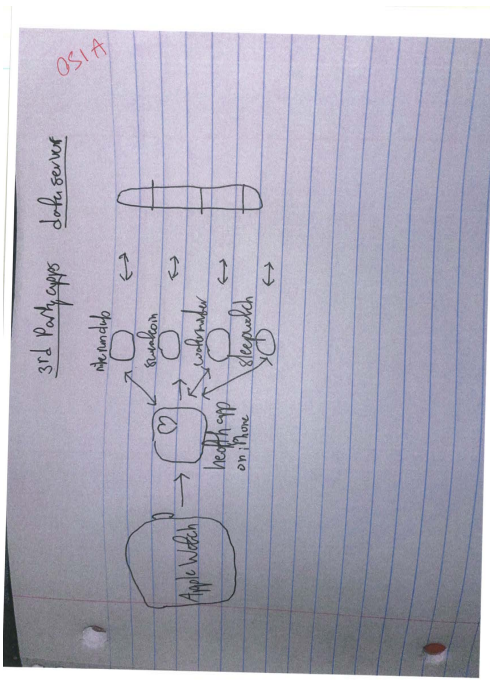
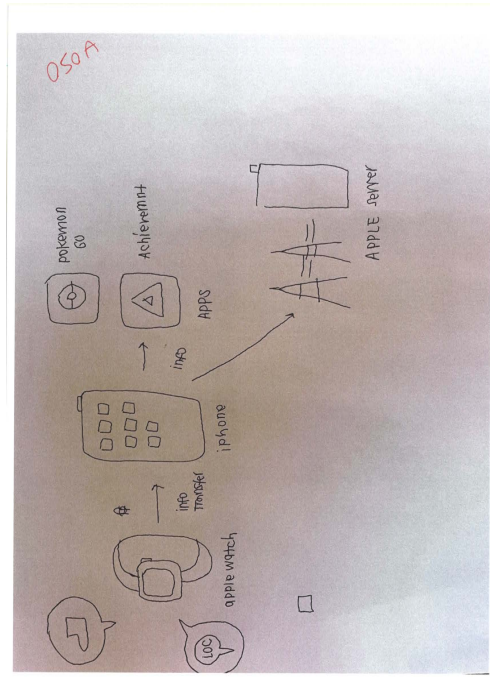
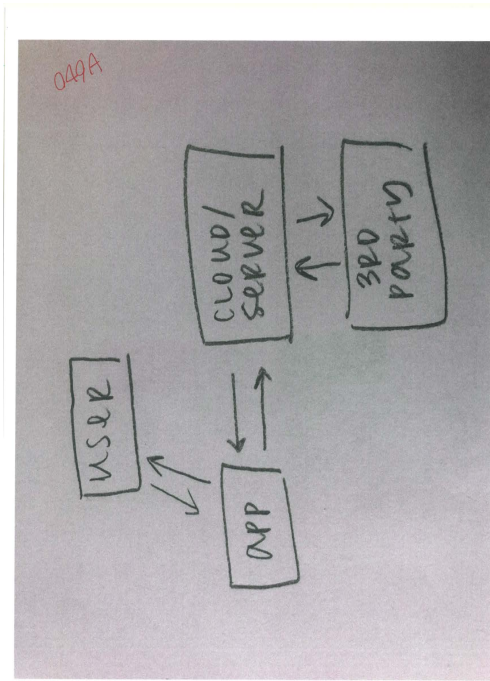



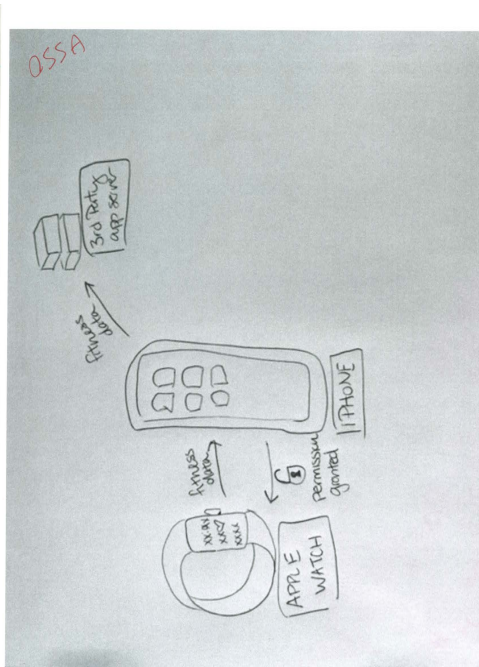
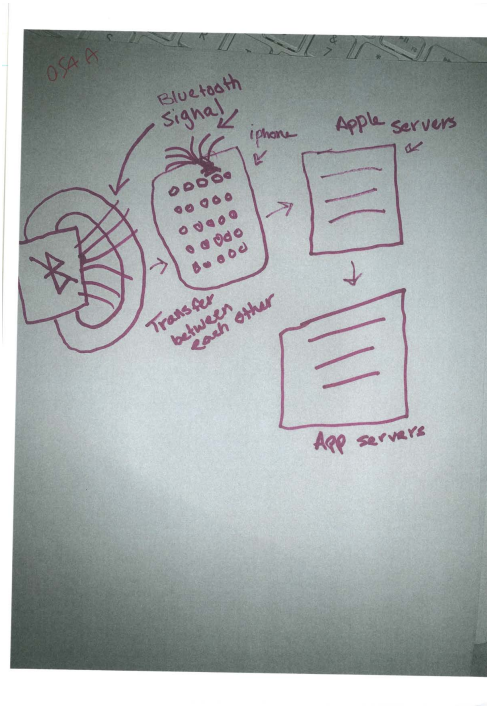








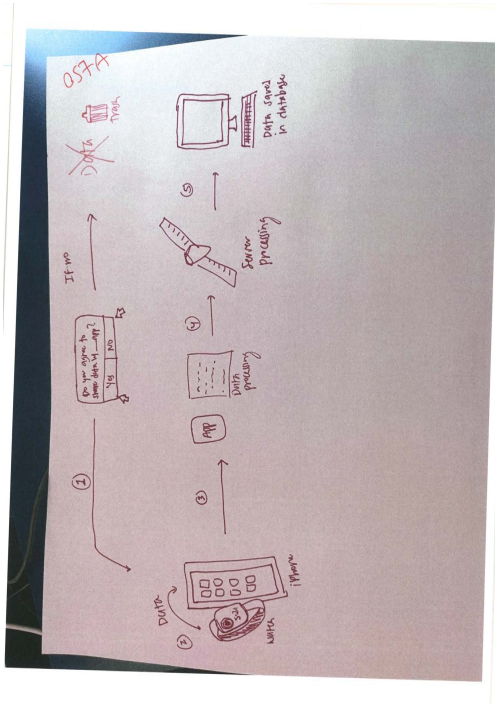



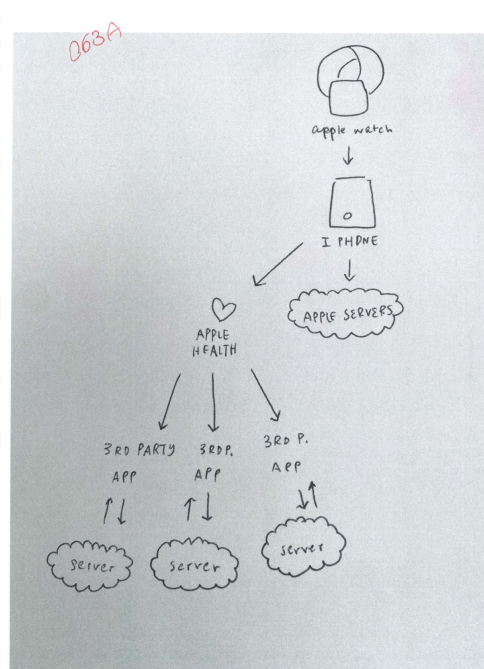
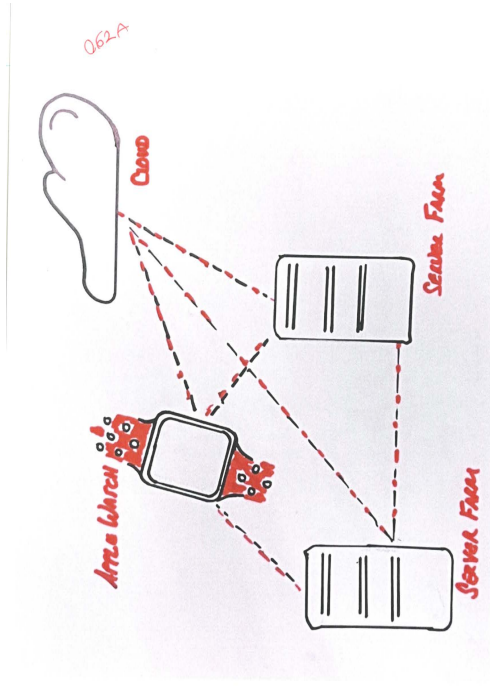






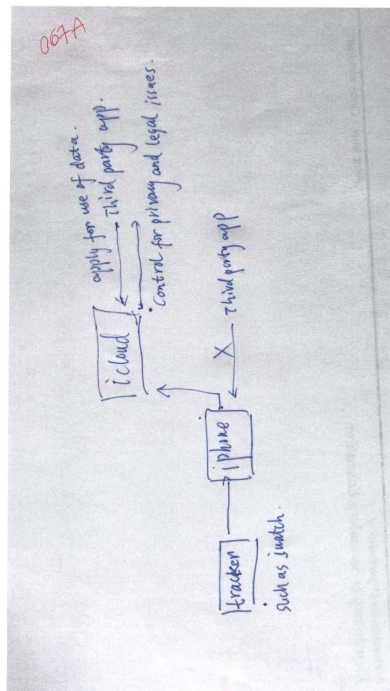






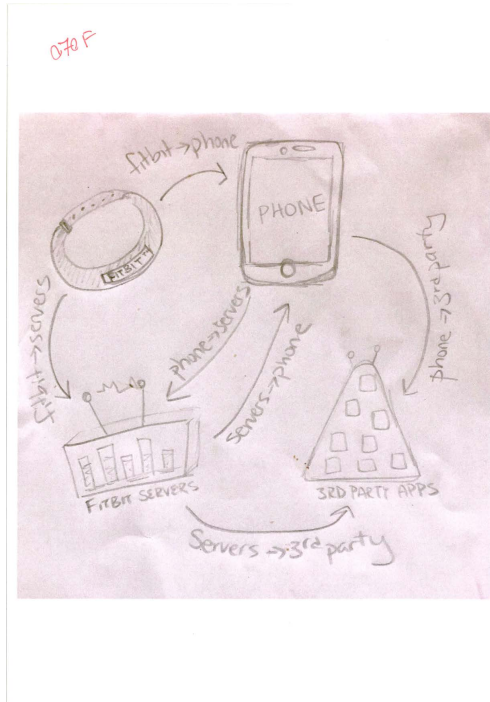





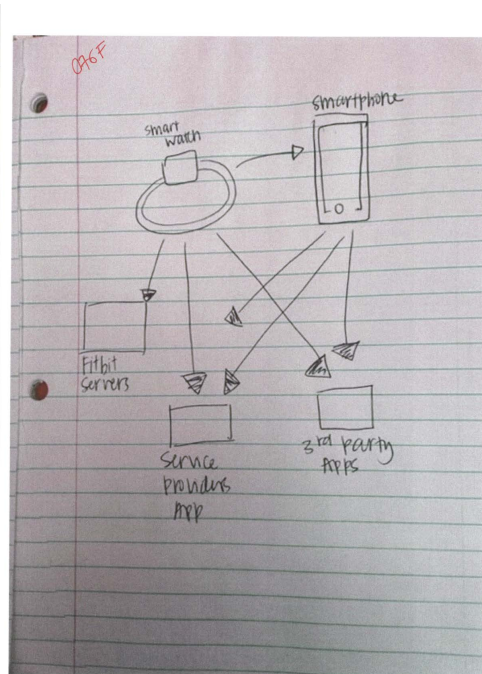
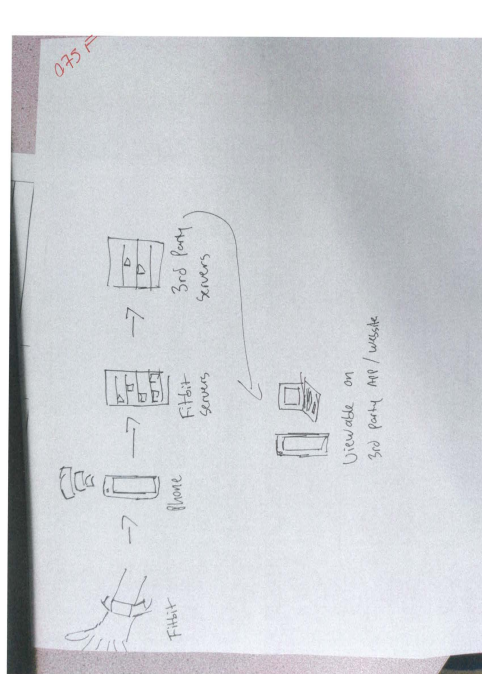
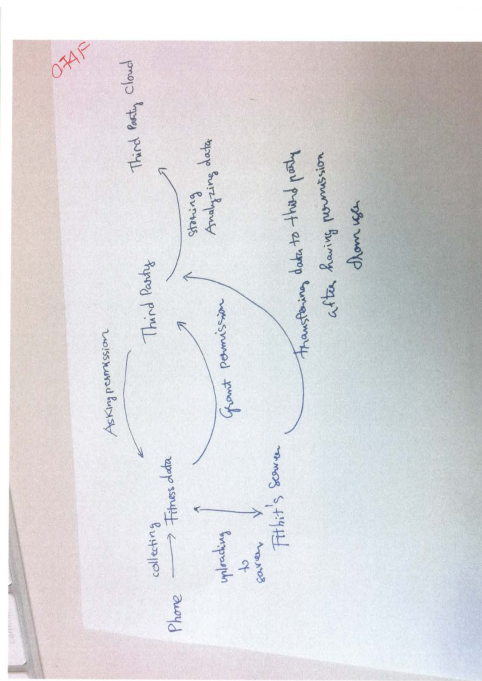
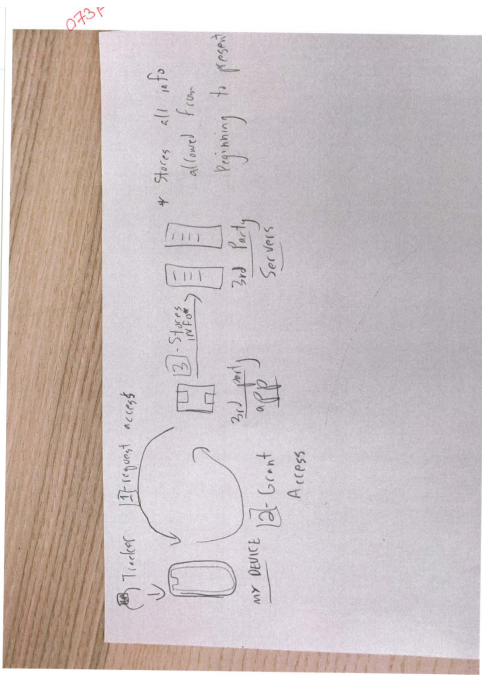



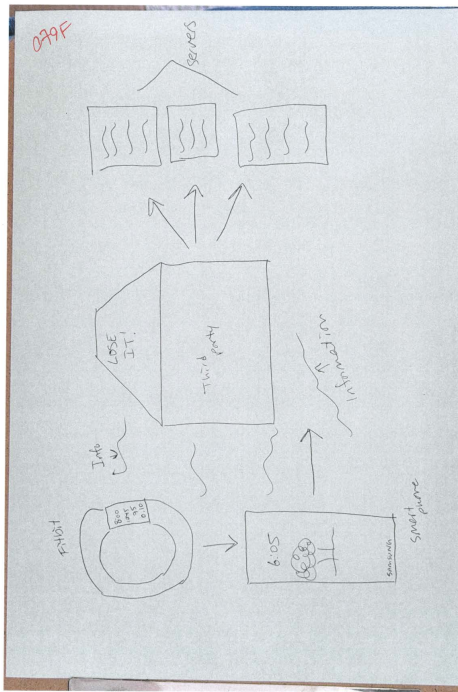
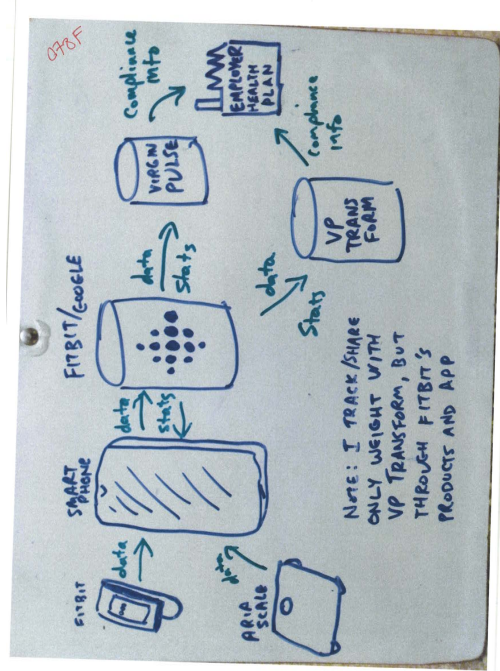
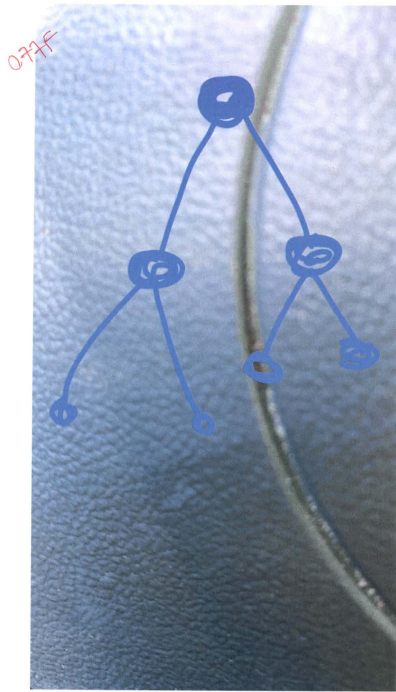




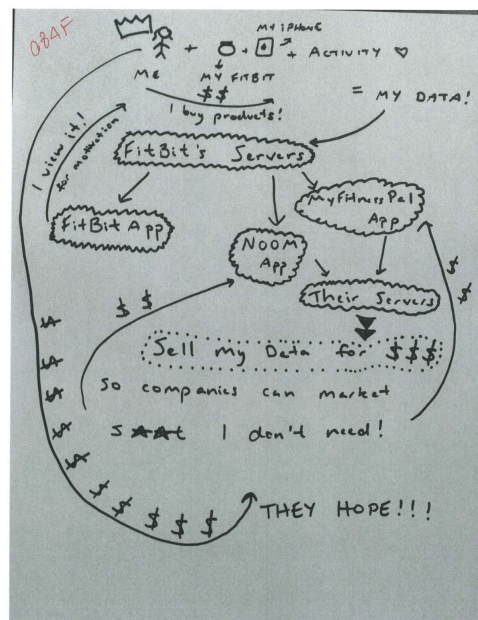


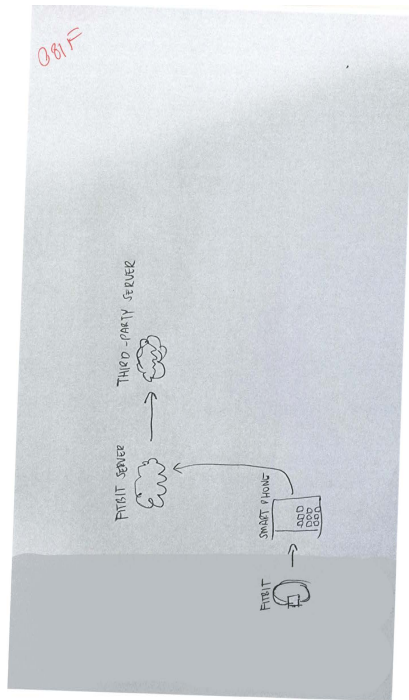


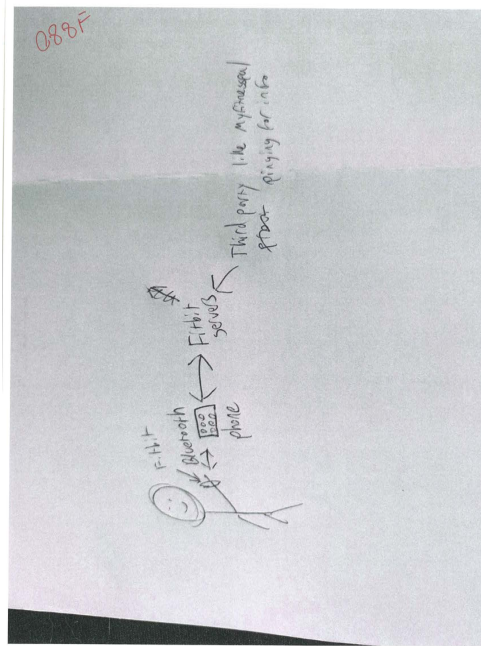

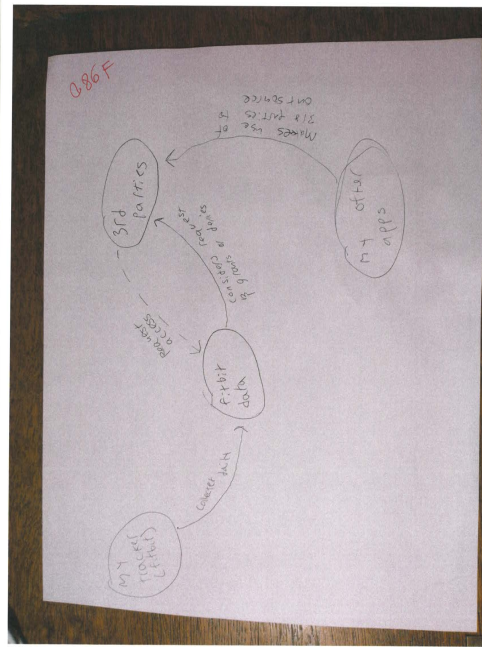
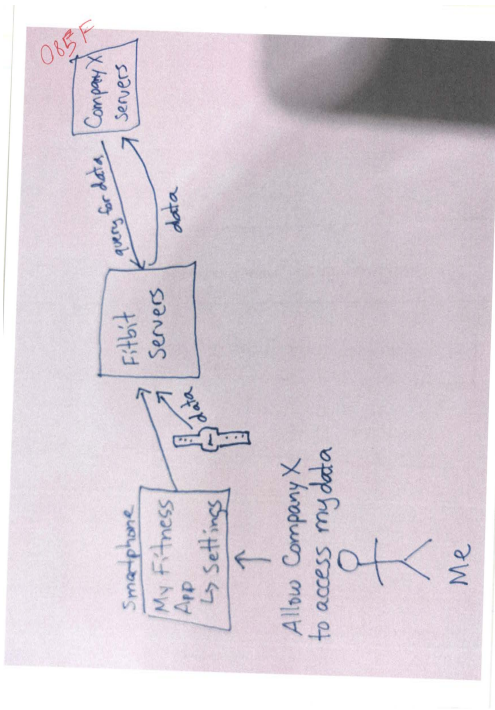





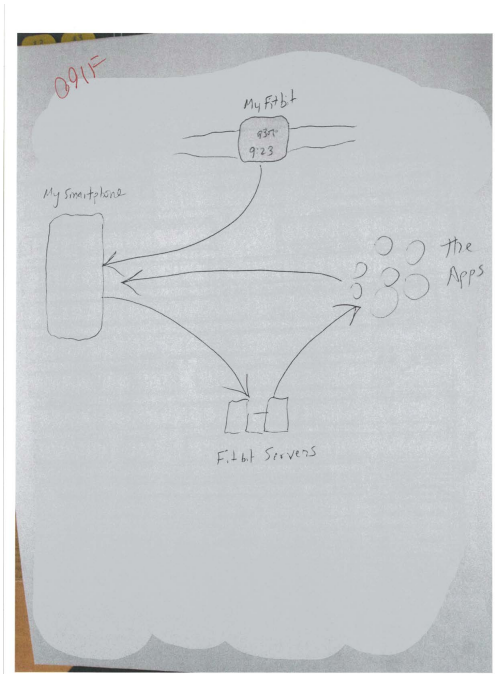
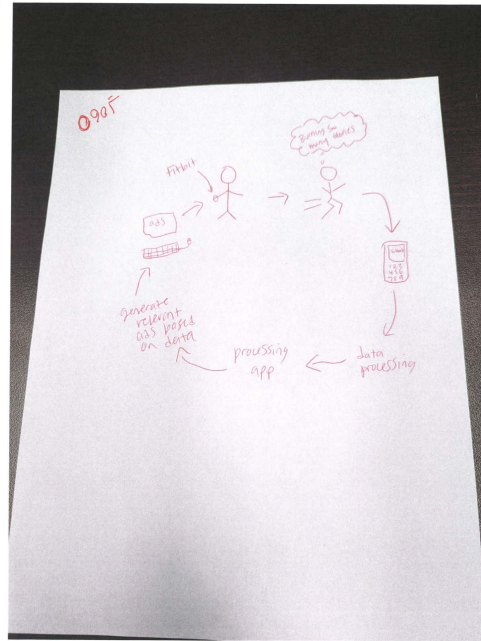



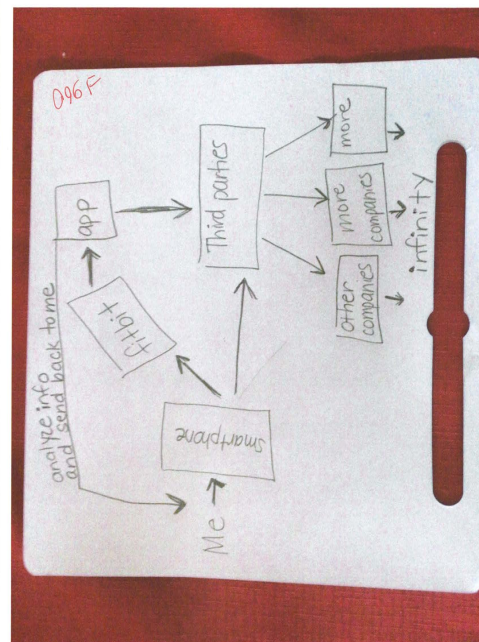
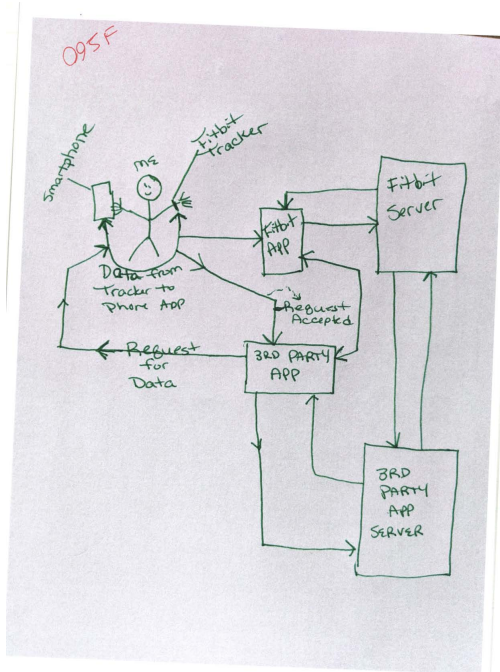
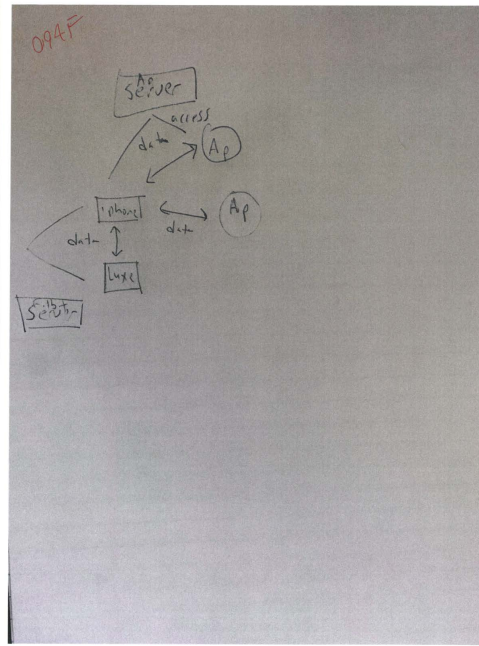






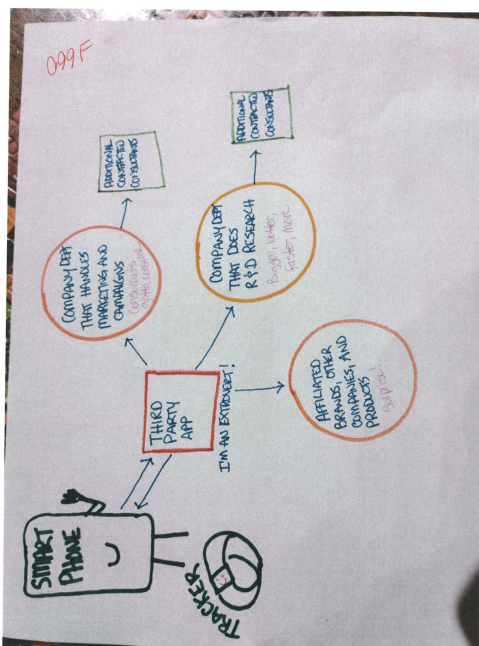
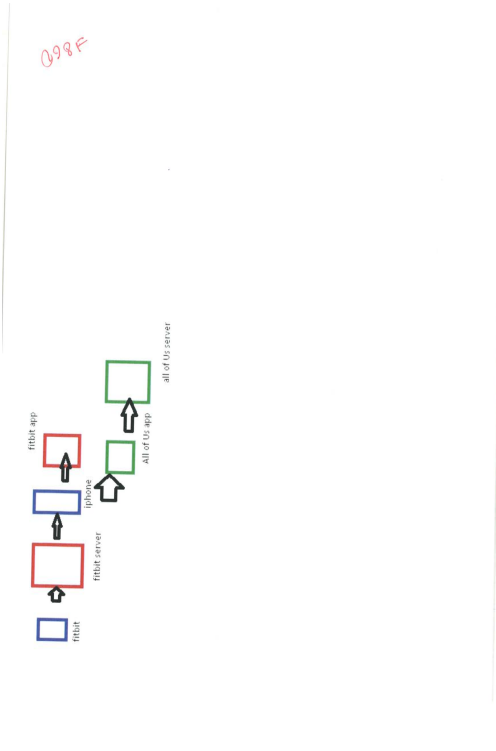



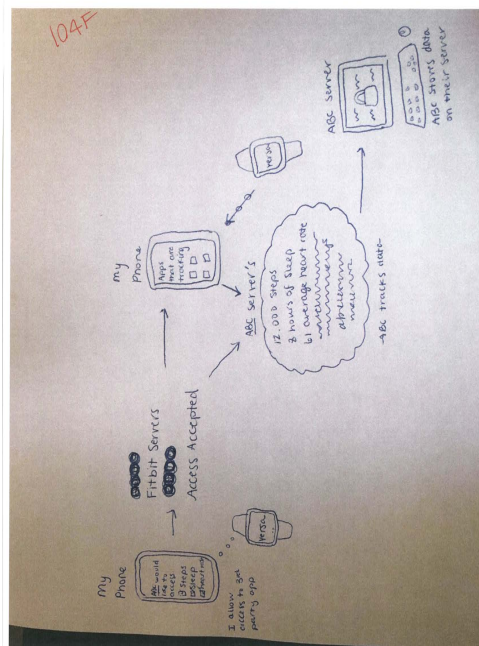
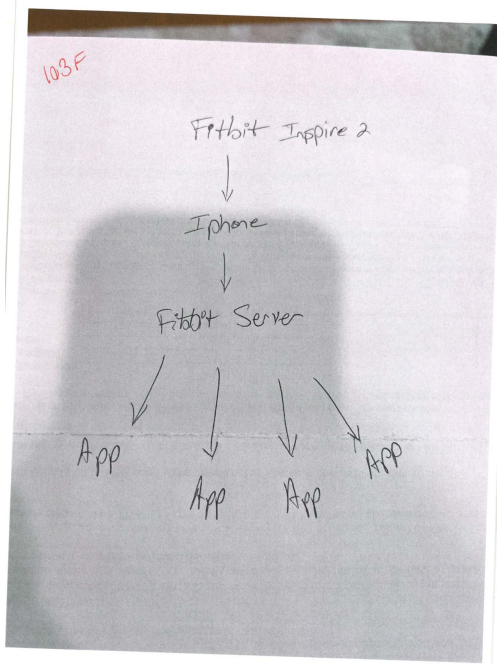

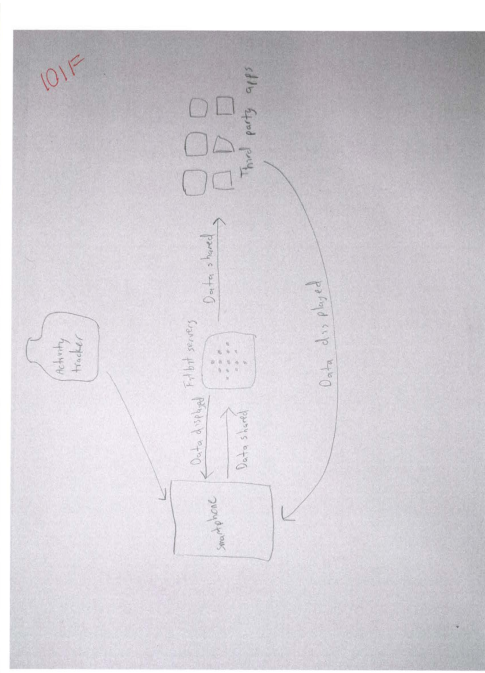




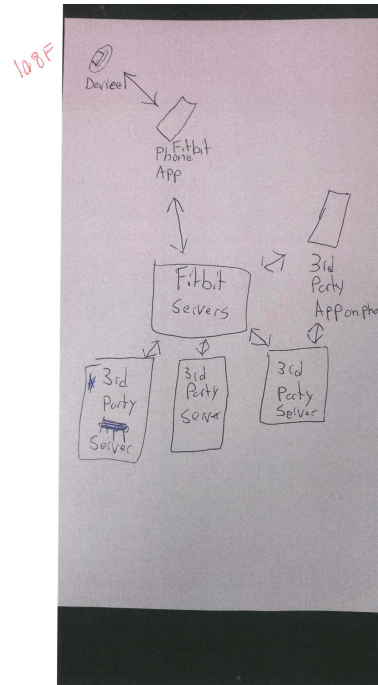

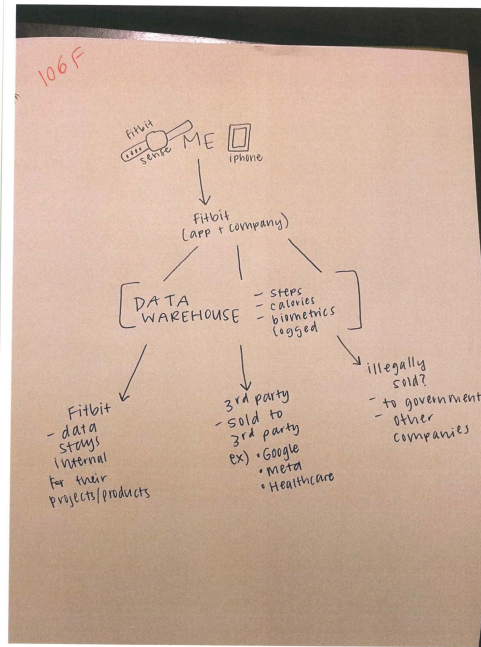
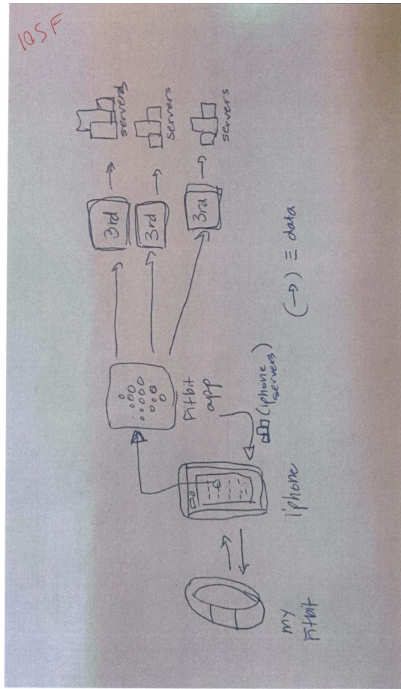



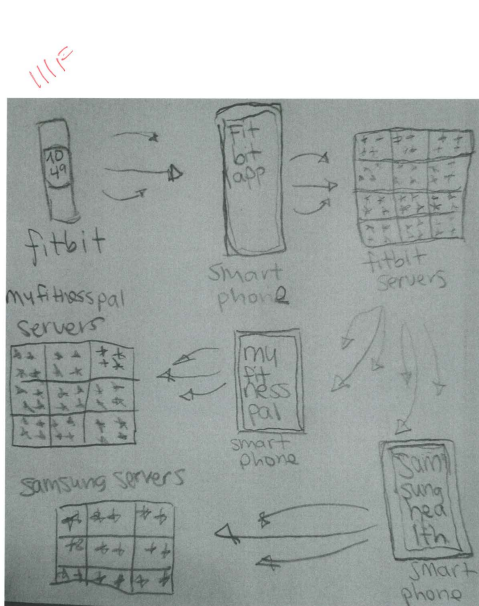
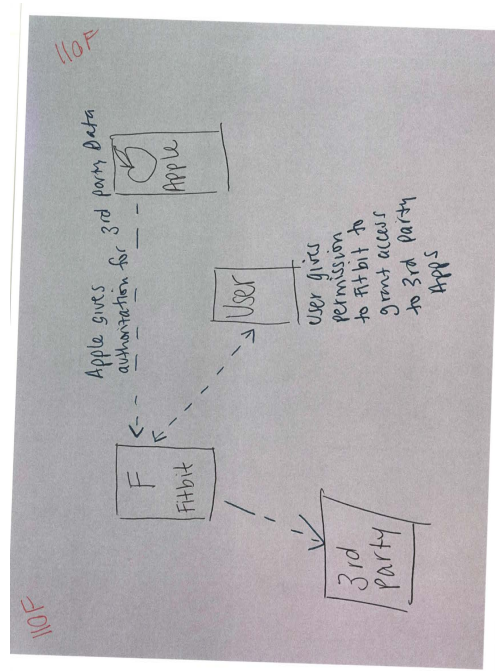





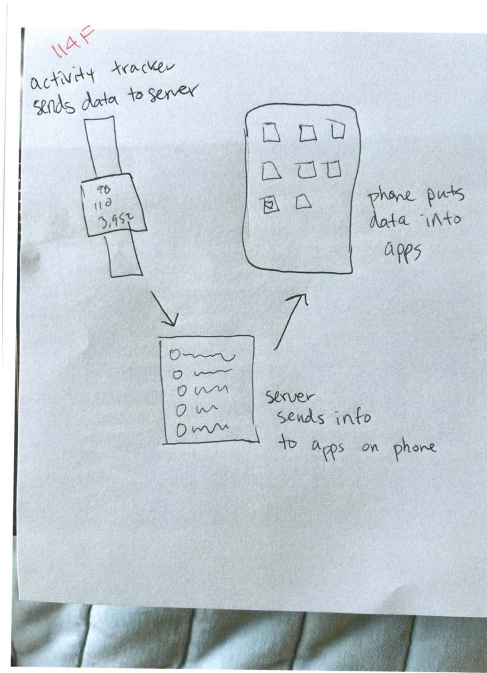
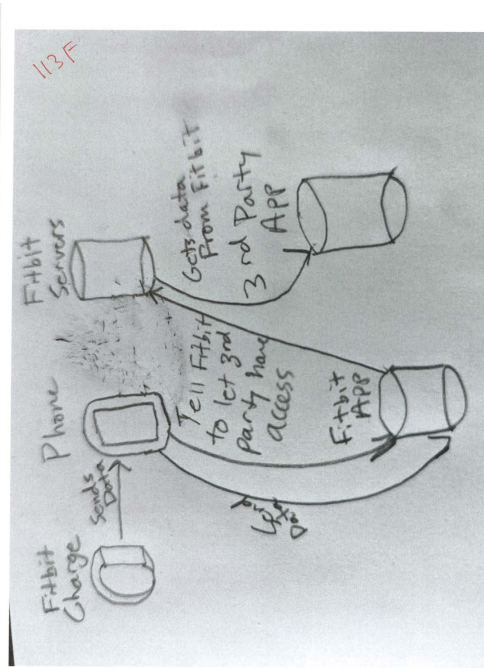















117F



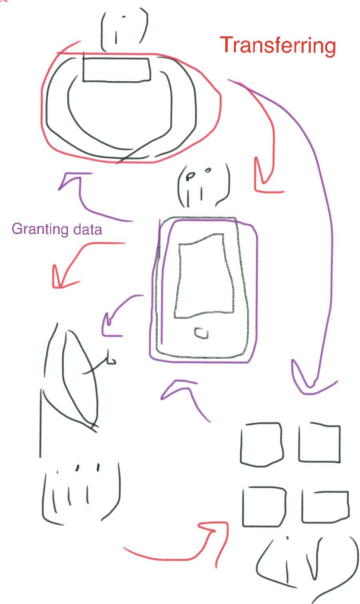
118F




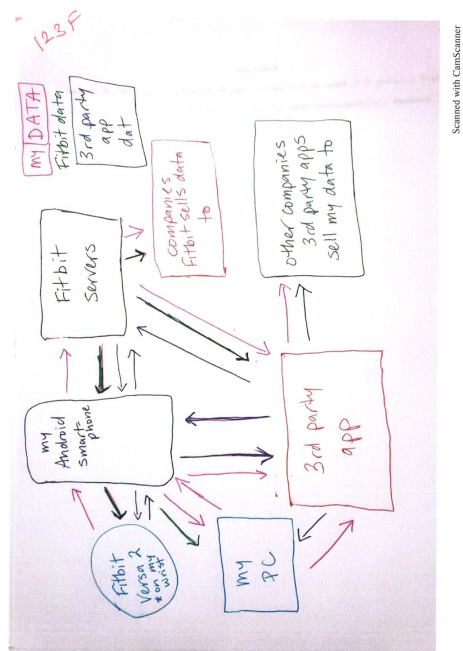
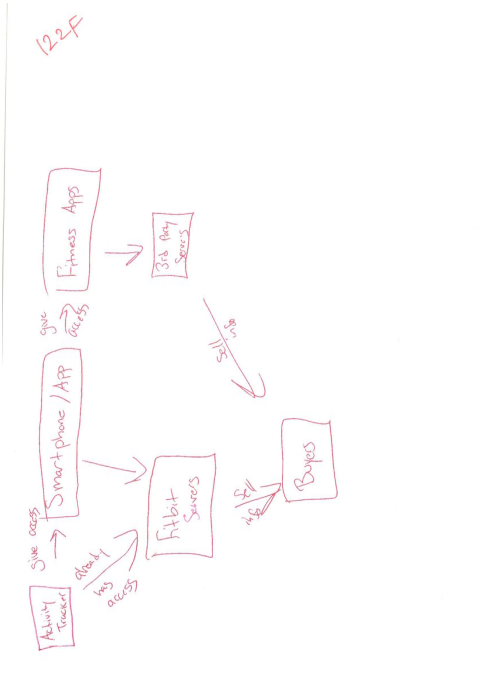
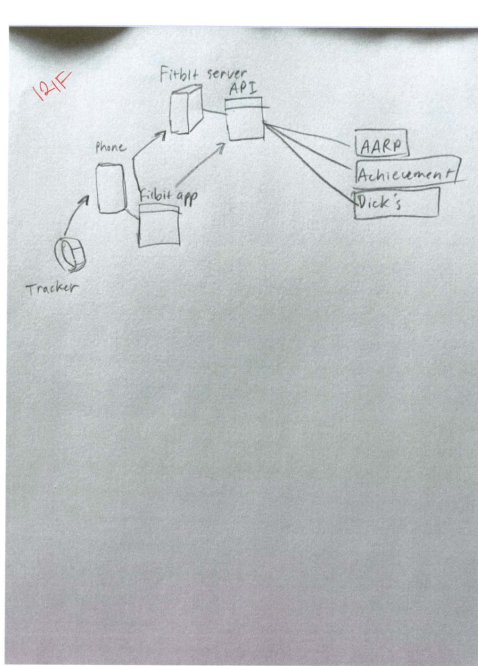
119F


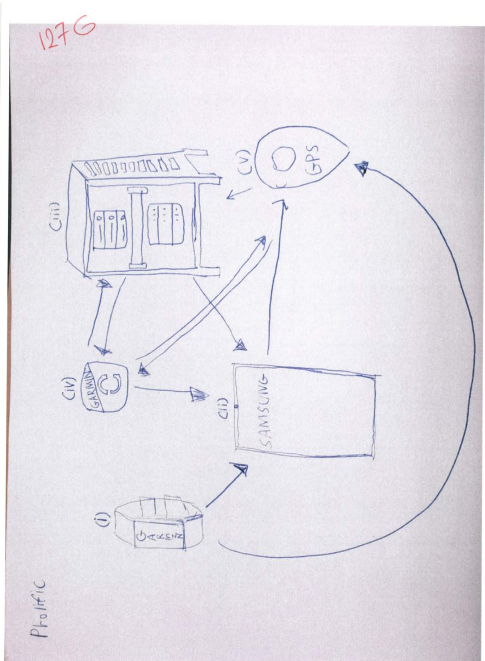
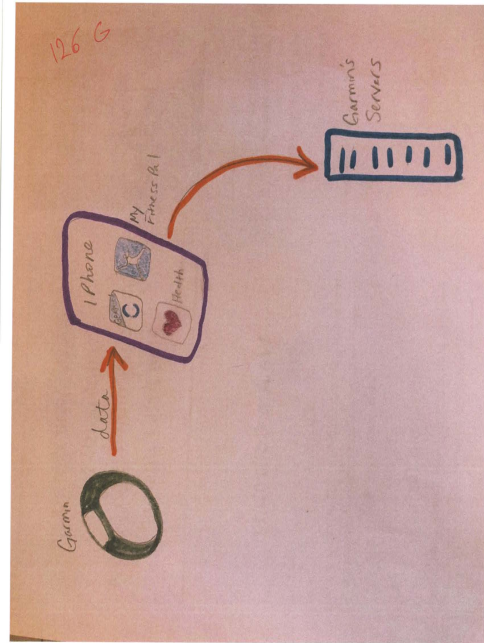
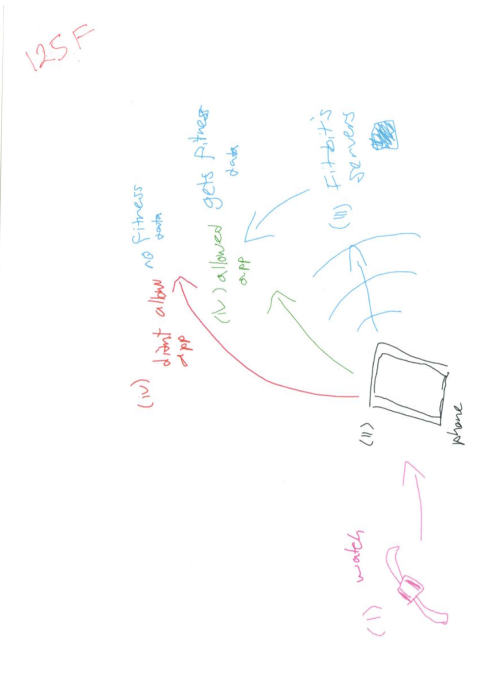


120F

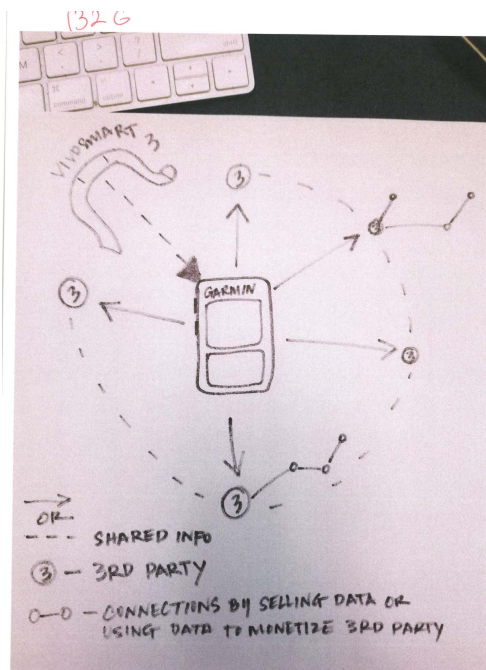


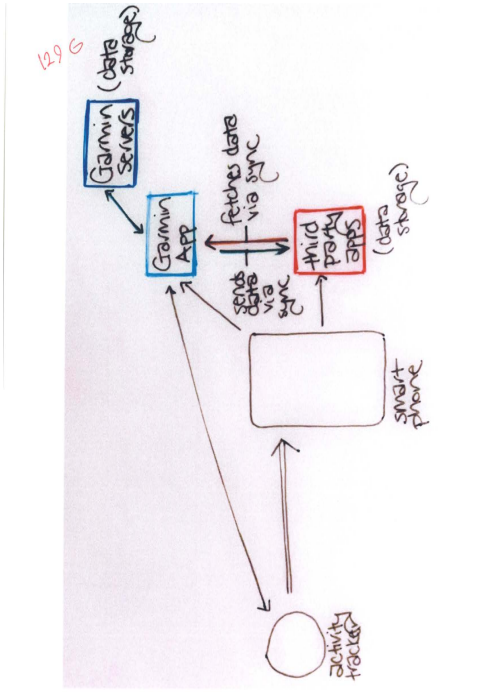








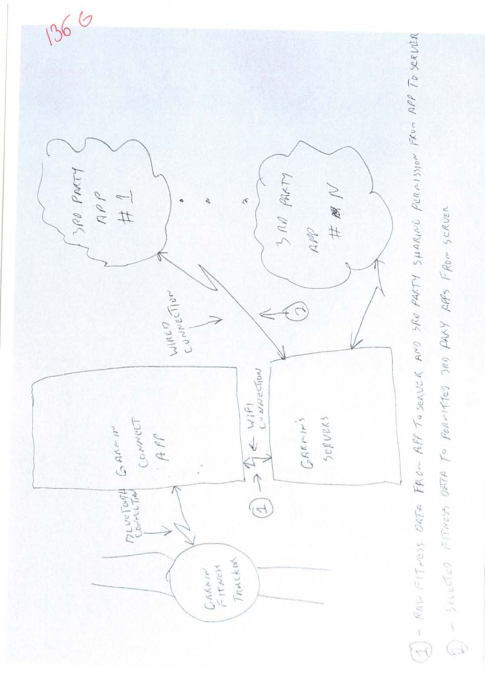

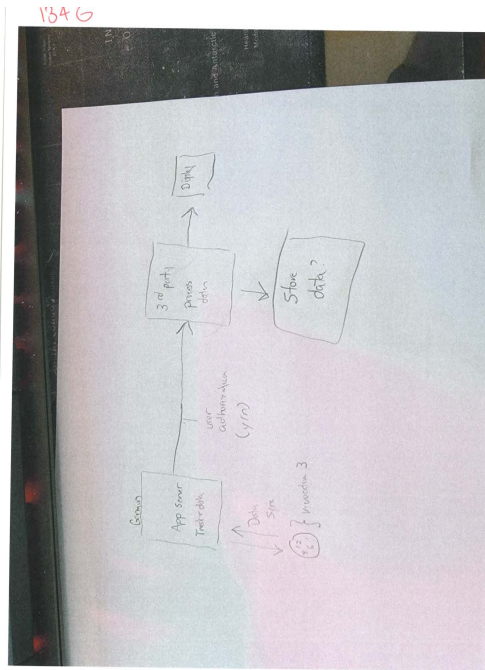


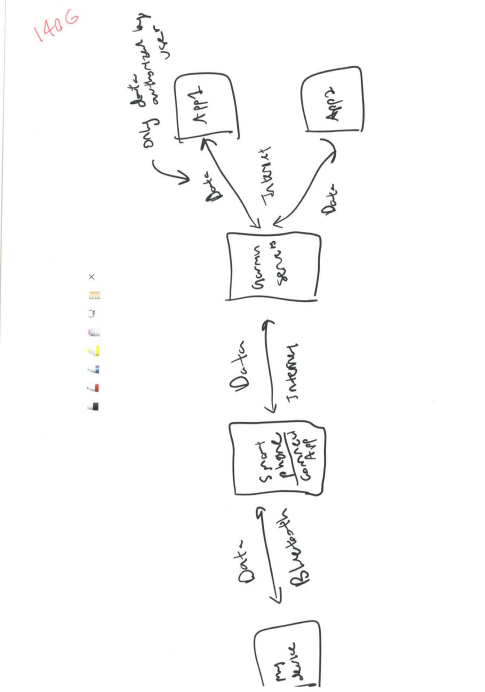
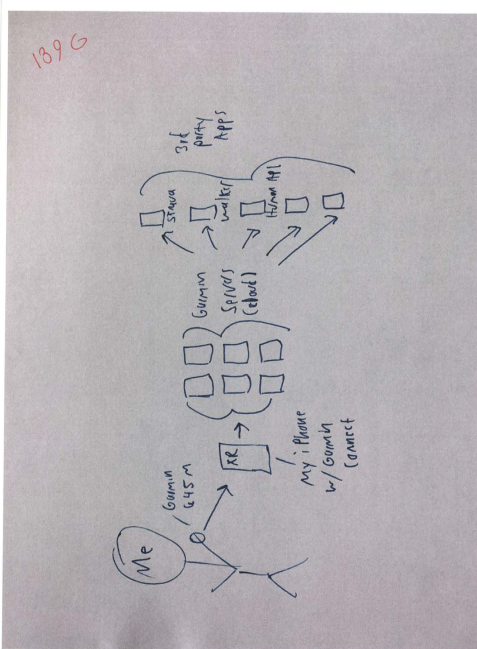
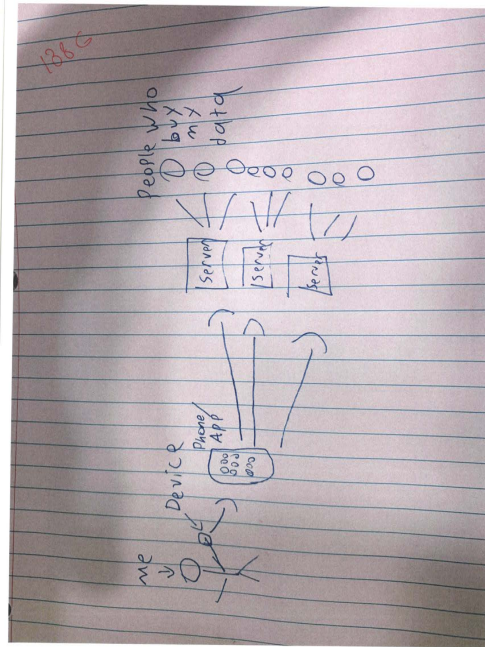
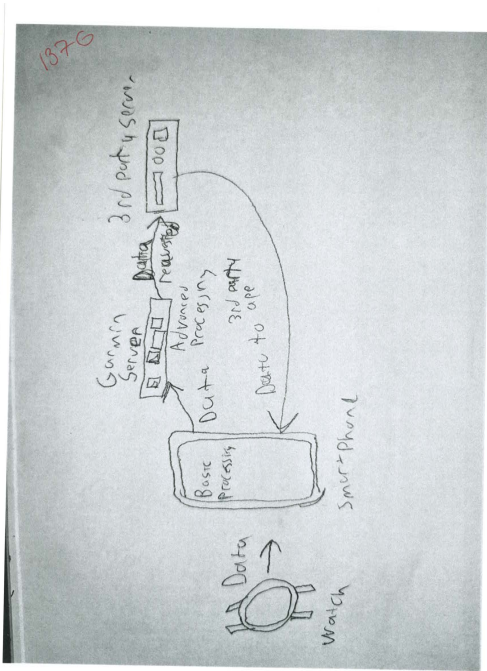


133 G

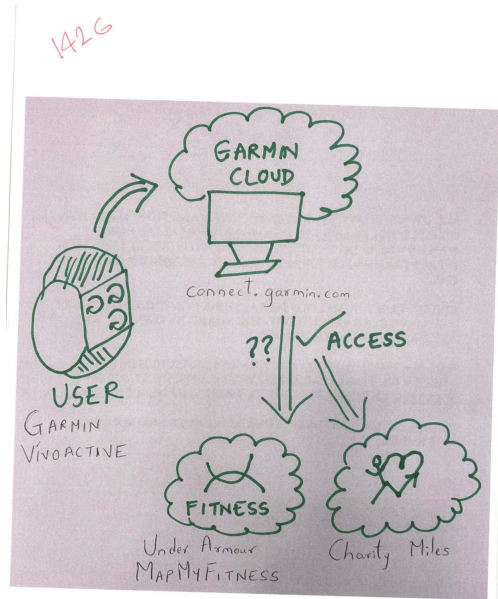
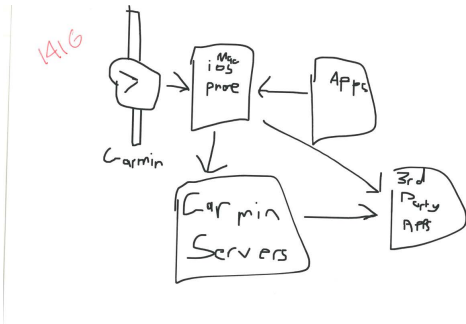
Request  
 Vivoactive 3 music → smartphone → Garmin servers → Strava app  
 → Strava servers

Approval  
 Strava servers → Garmin app → Strava app → smartphone  
 → Vivoactive 3 music









## A.3 Technical Codebook

| Brand   | Apple                  |              |             | Fitbit    |              |             | Garmin    |              |             | TOTAL      |            |
|---|------------------------|--------------|-------------|-----------|--------------|-------------|-----------|--------------|-------------|------------|------------|
|   | N                      | ratio        | correctness | N         | ratio        | correctness | N         | ratio        | correctness | N          | ratio      |
| Mental Models   |                        |              |             |           |              |             |           |              |             |            |            |
| mm1   | 18                     | 29.03        | incorrect   | 24        | 45.28        | correct     | 6         | 35.29        | correct     | 48         | 36.36      |
| mm2   | 2                      | 3.23         | incorrect   | 6         | 11.32        | incorrect   | 3         | 17.65        | incorrect   | 11         | 8.33       |
| mm3   | 23                     | 37.1         | correct     | 10        | 18.87        | incorrect   | 2         | 11.76        | incorrect   | 35         | 26.52      |
| mm1 & mm3   | 3                      | 4.84         | inaccurate  | 1         | 1.89         | inaccurate  | 1         | 5.88         | inaccurate  | 5          | 3.79       |
| others - mm4  | 16                     | 25.81        | incorrect   | 12        | 22.64        | incorrect   | 5         | 29.41        | incorrect   | 33         | 25         |
| <b>TOTAL</b>  | <b>62</b>              | <b>46.97</b> |             | <b>53</b> | <b>40.15</b> |             | <b>17</b> | <b>12.88</b> |             | <b>132</b> | <b>100</b> |
| 36.36% of the respondents have mm1. 26.52% of the respondents have mm3. 8.33% of the respondents have mm2. 3.79% of the respondents share both mm1 and mm3.   |                        |              |             |           |              |             |           |              |             |            |            |
| Finally, 25% of the respondents have different (other) models.  |                        |              |             |           |              |             |           |              |             |            |            |
|   | <b>correct</b>         | 23           | 37.1        | 24        | 45.28        |             | 6         | 35.29        |             | 53         | 40.15      |
|   | <b>incorrect</b>       | 36           | 58.06       | 28        | 52.83        |             | 10        | 58.82        |             | 74         | 56.06      |
|   | <b>inaccurate</b>      | 3            | 4.84        | 1         | 1.89         |             | 1         | 5.88         |             | 5          | 3.79       |
|   | <b>TOTAL</b>           | <b>62</b>    | <b>100</b>  | <b>53</b> | <b>100</b>   |             | <b>17</b> | <b>100</b>   |             | <b>132</b> | <b>100</b> |
| 40.15% of the respondents had a correct mental model and 56.06% of them had an incorrect mental model. 3.79% of the respondents mixed correct and incorrect models (i.e., inaccurate model)                             |                        |              |             |           |              |             |           |              |             |            |            |
| Fitbit users have more correct mental models with 45.28% than Apple (37.1%) and Garmin (35.29%) users.  |                        |              |             |           |              |             |           |              |             |            |            |
|   | <b>correct</b>         | 23           | 37.1        | 24        | 45.28        |             | 6         | 35.29        |             | 53         | 40.15      |
|   | <b>incorrect</b>       | 18           | 29.03       | 18        | 33.96        |             | 8         | 47.06        |             | 44         | 33.33      |
|   | <b>different model</b> | 21           | 33.87       | 11        | 20.75        |             | 3         | 17.65        |             | 35         | 26.52      |
|   | <b>TOTAL</b>           | <b>62</b>    | <b>100</b>  | <b>53</b> | <b>99.99</b> |             | <b>17</b> | <b>100</b>   |             | <b>132</b> | <b>100</b> |
| The mental model of 26.52% of the users is related to a different brand than their own device. In particular, one-third (33.78%) of the Apple users think their device can sync online with servers which are not true. |                        |              |             |           |              |             |           |              |             |            |            |

## A.4 Contextual Codebook

|                   |  | Count (n) | Count (%) (out of 73 drawings) |
|-------------------|--|-----------|--------------------------------|
| Code 1.1          | being concerned TPA share/sell data                                | 28        | 38.36                          |
| Code 1.2          | being concerned about data being stored                            | 14        | 19.18                          |
| Code 1.3          | concerned about information analysis                               | 5         | 6.85                           |
| Code 1.4          | being concerned about user profiling by TPA                        | 4         | 5.48                           |
| Code 1.5          | being concerned about TPA network security                         | 6         | 8.22                           |
| <b>Category 1</b> | <b>Lack of trust in TPAs</b>                                       | <b>47</b> | <b>64.38</b>                   |
| Code 2.1          | privacy concern about WAT provider                                 | 7         | 9.59                           |
| Code 2.2          | generic privacy concern about the backstage                        | 1         | 1.37                           |
| Code 2.3          | not happy with TPA/WAT privacy                                     | 1         | 1.37                           |
| Code 2.4          | being concerned about the phone company                            | 1         | 1.37                           |
| Code 2.5          | operating systems such as iOS and Android should protect the users | 1         | 1.37                           |
| Code 2.6          | being concerned about TPA access                                   | 1         | 1.37                           |
| <b>Category 2</b> | <b>General Privacy concerns about WAT/TPAs/operating systems</b>   | <b>12</b> | <b>16.43</b>                   |
| Code 3.1          | knowledgable about access management                               | 31        | 42.47                          |
| Code 3.2          | knowledgable about selective sharing                               | 6         | 8.22                           |
| <b>Category 3</b> | <b>Being knowledgeable about access management</b>                 | <b>35</b> | <b>47.94</b>                   |
| Code 4.1          | believing that data will be deleted later                          | 1         | 1.37                           |
| Code 4.2          | confident with WAT provider S&P                                    | 1         | 1.37                           |
| Code 4.3          | collected data can be use to improve WAT services                  | 2         | 2.74                           |
| <b>Category 4</b> | <b>Being positive about the data sharing process</b>               | <b>4</b>  | <b>5.48</b>                    |

## A.5 Why Revoking Access Codebook

| Question: Why do respondents revoke access? |            |                         |           |           |             |   |
|---|------------|-------------------------|-----------|-----------|-------------|---|
|   | Count (n)  | Valid Collected Answers |           |           | CODEBOOK    | Example Quote   |
|   |            | Apple                   | Fitbit    | Garmin    |             |   |
|   | 171        | 93                      | 65        | 13        |             |   |
|   |            | CODEBOOK                |           |           |             |   |
| Code 1.1                                    | 23         | 8                       | 9         | 6         | -           | "The app kept disconnecting from my data and I had to constantly re-link it. I decided to stop using the app altogether."                                   |
| Code 1.2                                    | 111        | 63                      | 42        | 6         | -           | "No longer was interested in using that app, so I wanted to revoke permissions prior to uninstalling."  |
| Code 1.3                                    | 4          | 3                       | 1         | 0         | -           | "I needed to be able to look at the data from the apps but I did not want them to continue collecting data. For this app, I found others I preferred more." |
| <b>Category 1</b>                           | <b>138</b> | <b>74</b>               | <b>52</b> | <b>12</b> | <b>80.7</b> | -   |
| <b>Category 2</b>                           | <b>27</b>  | <b>17</b>               | <b>9</b>  | <b>1</b>  | <b>15.8</b> | "I was nervous about the data they were accessing."   |
| Code 3.1                                    | 2          | 1                       | 1         | 0         | -           | "To be honest, I do not remember. I used to use it, and earlier I thought it was still on my phone."  |
| Code 3.2                                    | 1          | 0                       | 1         | 0         | -           | "Just did and best to do!"  |
| Code 3.3                                    | 1          | 0                       | 1         | 0         | -           | "The option to revoke access popped up on my phone as an option so I did it."   |
| Code 3.4                                    | 2          | 1                       | 1         | 0         | -           | -   |
| <b>Category 3</b>                           | <b>6</b>   | <b>2</b>                | <b>4</b>  | <b>0</b>  | <b>3.5</b>  | -   |
| <b>others</b>                               |            |                         |           |           |             |   |

## A.6 Why Not Revoking Access Codebook

Question: Why do respondents not revoke access?

|                                |  | Count<br>(n)         | Apple        | Fitbit        | Garmin        | Count (%)        | Example Quote   |
|--------------------------------|--|----------------------|--------------|---------------|---------------|------------------|---|
| <b>Total Answers Collected</b> |  | <b>394</b>           | <b>221</b>   | <b>141</b>    | <b>32</b>     |                  |   |
| <b>Valid Answers</b>           |  | <b>384</b>           | <b>206</b>   | <b>129</b>    | <b>29</b>     |                  |   |
| <b>CODEBOOK</b>                |  |                      |              |               |               |                  |   |
|                                |  | <b>Count<br/>(n)</b> | <b>Apple</b> | <b>Fitbit</b> | <b>Garmin</b> | <b>Count (%)</b> |   |
| Code 1.1                       | did not think about it or thought not necessary to revoke access       | 67                   | 45           | 18            | 4             | -                | "I just never think about it and do not think it is an issue to leave them on."   |
| Code 1.2                       | because they don't care as fitness data is not sensitive               | 41                   | 29           | 12            | 0             | -                | "Honestly, I'm not concerned of these third-party apps that's why I kept them."   |
| <b>Category 1</b>              | <b>comfortable to share data (not interested in access management)</b> | <b>108</b>           | <b>74</b>    | <b>30</b>     | <b>4</b>      | <b>29.7</b>      |   |
| Code 2.1                       | forgo/did not notice   | 93                   | 46           | 33            | 14            | -                | "I forgot and didn't realize the apps had access until completing this survey."   |
| Code 2.2                       | revoked or will revoke after answering the survey                      | 14                   | 7            | 5             | 2             | -                | "I was not thinking about this app having access to my data until I was prompted to think of the third-party apps I may have connected to my Fitbit...I plan to revoke access to it now." |
| <b>Category 2</b>              | <b>forgot about installed TPAs (might revoke later)</b>                | <b>107</b>           | <b>53</b>    | <b>38</b>     | <b>16</b>     | <b>29.4</b>      |   |
| <b>Category 3</b>              | <b>contemplate using the TPA (actively) again in the (near) future</b> | <b>97</b>            | <b>50</b>    | <b>41</b>     | <b>6</b>      | <b>26.7</b>      |   |
| Code 4.1                       | didn't know I can do or how to do it                                   | 20                   | 13           | 7             | 0             | -                | "I didn't realize I could. I didn't even know where to look for this information until today."  |
| Code 4.2                       | uninstalled and thought it will revoke automatically                   | 6                    | 3            | 2             | 1             | -                | "I uninstalled the apps and assumed they would no longer be able to access my app information."   |
| Code 4.3                       | didn't realize sharing my data   | 42                   | 24           | 13            | 5             | -                | "I didn't realize that after I stopped using the app, they would still have access to my data."   |
| <b>Category 4</b>              | <b>not familiar with data sharing and access management</b>            | <b>68</b>            | <b>40</b>    | <b>22</b>     | <b>6</b>      | <b>18.7</b>      |   |
| <b>Category 5</b>              | <b>perceive access management as complex / difficult (hassle)</b>      | <b>14</b>            | <b>9</b>     | <b>5</b>      | <b>0</b>      | <b>3.9</b>       |   |
| Code 6.1                       | to keep receiving financial benefits                                   | 2                    | 0            | 2             | 0             | -                | "Third party app needs to be saved for my annual insurance incentive."  |
| Code 6.2                       | keep saving data for future use  | 2                    | 2            | 0             | 0             | -                | "I want to keep my data there for a historical record in case I was to refer back."   |
| <b>Category 6</b>              | <b>want to get more benefits (health or monetary)</b>                  | <b>4</b>             | <b>2</b>     | <b>2</b>      | <b>0</b>      | <b>1.1</b>       |   |
| Code 7.1                       | trust the app  | 2                    | 1            | 1             | 0             | -                | "I trust that there isn't really anything they would do with my data."  |
| Code 7.2                       | to help the app  | 1                    | 0            | 1             | 0             | -                | "I thought it would help with the app."   |
| <b>Category 7</b>              | <b>trust TPAs</b>  | <b>3</b>             | <b>1</b>     | <b>2</b>      | <b>0</b>      | <b>0.8</b>       |   |
| Code 8.1                       | only log info and doesn't let tracking                                 | 3                    | 2            | 1             | 0             | -                | "Because I never let them have access in the first place. I only used 3rd parties like my fitness pal to log the information."  |
| Code 8.2                       | Can't remember   | 1                    | 1            | 0             | 0             | -                | "Didn't remember!"  |
| Code 8.3                       | someone else used the app  | 1                    | 1            | 0             | 0             | -                | "My little cousin uses the app, I don't!"   |
| Code 8.4                       | did not have time as I just stopped using it                           | 1                    | 0            | 1             | 0             | -                | "I only recently stopped using the app."  |
| Code 8.5                       | I limited the access   | 1                    | 0            | 1             | 0             | -                | "I did not revoke access entirely but limited access."  |
| Code 8.6                       | My WAT doesn't support it  | 1                    | 0            | 1             | 0             | -                | "Not compatible with my version of Fitbit device!"  |
| Code 8.7                       | Not meaningful responses [removed later]                               | 2                    | 0            | 2             | 0             | -                |   |
| <b>Category 8</b>              | <b>others</b>  | <b>10</b>            | <b>4</b>     | <b>6</b>      | <b>0</b>      | <b>2.7</b>       |   |

## A.7 Suggestions Codebook

| Question: suggestions on how to facilitate the TPA access management process |  |            |              |   |
|--|--|------------|--------------|---|
| Total Answers Collected  |  | Count (n)  |              |   |
|  |  | 614        |              |   |
| Valid Answers  |  | Count (n)  |              | Count (%)   |
|  |  | 480        |              |   |
| CODEBOOK   |  |            |              |   |
|  |  | Count (n)  | Count (%)    | Example Quote   |
| Code 1.1   | notification/reminder  | 191        | 39.79        | "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."  |
| Code 1.2   | privacy checkups (auto-turn off unused permissions)                        | 17         | 3.54         | "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."   |
| Code 1.3   | specialized app/feature for managing granted access                        | 19         | 3.96         | "An app that can track which 3rd-party applications have access to my data and help me choose which to revoke. It is too much to go through every individual app to see what has access to what!"   |
| Code 1.4   | the phone should clearly show for each TPA what data data collect or store | 30         | 6.25         | "Place the permissions in a consolidated location, rather than skipping around to apps that may or may not be reading data."  |
| <b>Category 1</b>  | <b>access monitoring systems</b>   | <b>257</b> | <b>53.54</b> | -   |
| Code 2.1   | clear, transparent, and easy to comprehend privacy policies                | 41         | 8.54         | "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." |
| Code 2.2   | partial (selective) sharing  | 10         | 2.08         | "Most apps make you share everything regardless of what their purpose is, they should only request access to what is necessary."  |
| Code 2.3   | making temporary access only   | 5          | 1.04         | "I would do temporary access instead of permanent access in the beginning."   |
| Code 2.4   | not sharing (maximum privacy) should be default                            | 3          | 0.63         | "Make privacy selections default to the minimum authorization levels. User must actively select information sharing."   |
| Code 2.5   | asking users to double check what they share                               | 3          | 0.63         | "Apple should be more straightforward about what allowing access actually means. Like when I select "allow" a popup should say, "Are you sure you want to allow access?"  |
| <b>Category 2</b>  | <b>granting access</b>   | <b>62</b>  | <b>12.92</b> | -   |
| Code 3.1   | educating users about TPAs   | 19         | 3.96         | "I think that having a guide like this would be helpful for anyone. I never really considered how much access 3rd parties have to all of my health data."   |
| Code 3.2   | facilitate access managing procedure (interface) and make it visible       | 100        | 20.83        | "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."   |
| Code 3.3   | new legislation or law enforcement   | 12         | 2.5          | "We need laws in place to protect people."  |
| Code 3.4   | deleting previously stored data after revoking access                      | 12         | 2.5          | "Entering into the agreement that certain apps can only access data from the date of authorization moving forward and once access is revoked all data will be deleted from any saved data hubs."  |
| Code 3.5   | allowing users to delete their data from the servers                       | 2          | 0.42         | "Allow the users the option to 'wipe' their data."  |
| Code 3.6   | uninstalling TPA should revoke it  | 7          | 1.46         | "As soon as the app is deleted it should be auto revoked."  |
| <b>Category 3</b>  | <b>generic solutions</b>   | <b>152</b> | <b>31.67</b> | -   |
| Code 4.1   | users should avoid using many TPAs at the same time                        | 3          | 0.63         | "Just try out one app at a time until you are familiar with each other."  |
| Code 4.2   | user should be more careful  | 34         | 7.08         | "Make sure you take care to only grant access to apps that you will actually use and when you stop using them, remember to revoke access."  |
| Code 4.3   | users should do regular check  | 3          | 0.63         | "Regularly check on which apps have access."  |
| <b>Category 4</b>  | <b>Users should take more responsibility</b>                               | <b>40</b>  | <b>8.34</b>  | -   |
| Code 5.1   | keeping account private  | 1          | 0.21         | "I prefer to keep my account of stats private with only trusted access."  |
| Code 5.2   | adjusting privacy settings   | 1          | 0.21         | "Adjust the privacy settings of each app."  |
| Code 5.3   | tangible consent collection  | 1          | 0.21         | "Put something in writing as well as electronically and follow up to confirm receipt and desired action was taken."   |
| Code 5.4   | authentication   | 2          | 0.42         | "I suggest that you receive an email confirmation."   |
| Code 5.5   | warning users about implications   | 1          | 0.21         | "An app that warns you of implications."  |
| Code 5.6   | manage it via phone  | 1          | 0.21         | "I think it would be easier for all users if they can monitor/revoke it through the app itself."  |
| <b>Category 5</b>  | <b>others</b>  | <b>7</b>   | <b>1.47</b>  | -   |



## A.8 WAT Data Sharing


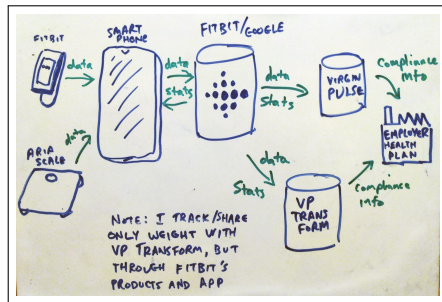
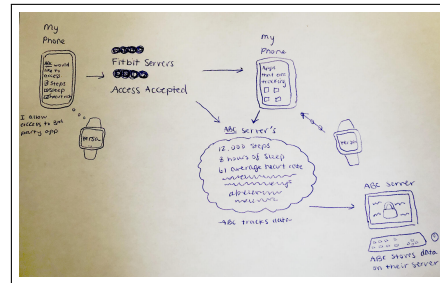


Figure A.1: Privacy level of different profile information for Fitbit and Garmin users. Here, we decided to refer to similar concepts of both service providers using the Garmin’s labels (e.g., “Badges” and “Badges and Trophies”), and to use Fitbit’s privacy labels and to refer to both Garmin’s “My Connections” and “My Groups and Connection” as “Friends”.

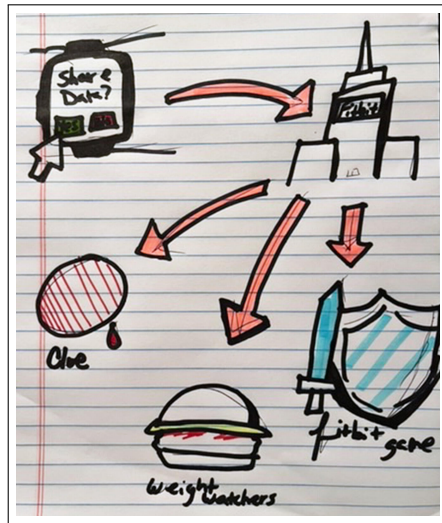
## A.9 Mental Models



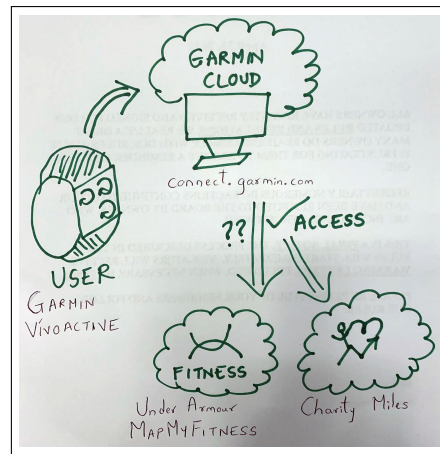
(a) The first type of the mental models ( $mm_1$ ): The fitness data is transmitted to TPAs via a connected device and the WAT server.



(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.

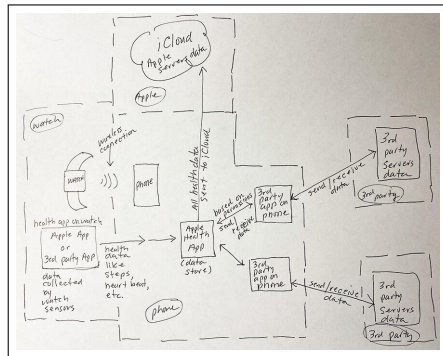


(c) The second type of the mental models ( $mm_2$ ): The fitness data is transmitted without passing via a connected device.




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


Figure A.2: 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.

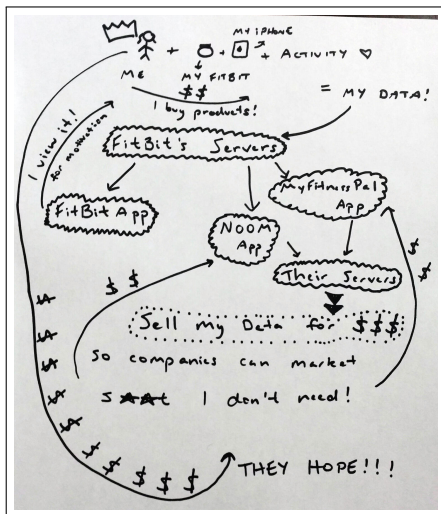
Figure A.3: Examples of users' mental models - 2




(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 A.4: 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 A.5: Examples of users' mental models - 4

## A.10 Design Feature Coding

| Code / Group  | 1 | 1 | 2 | 2 | 3 | 3 | 4 | 4 | 5 | 5 | 6 | 6 | 7 | 7 | 8 | 8 | 8 | 9 | 9 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| <b>Category 1 - Partial sharing</b>   | • |   |   |   |   |   |   |   | • |   |   |   |   |   |   |   |   | • |   |
| Code 1.1 - Sharing regarding the context  | • |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Code 1.2 - Sharing regarding a specific timeframe                                 | • |   |   |   |   |   |   |   | • |   |   |   |   |   |   |   |   | • |   |
| <b>Category 2 - Visualization</b>   |   | • |   |   |   |   |   |   | • | • | • |   | • |   |   |   |   | • |   |
| Code 2.1- Interactive tool for exploring shared data / data flow                  |   | • |   |   |   |   |   |   | • | • | • |   | • |   |   |   |   | • |   |
| Code 2.2- Data sharing logs (history)   |   |   | • |   |   |   |   |   | • |   |   |   |   |   |   |   |   |   |   |
| Code 2.3 - TPA usage Statistics   |   |   |   |   |   |   |   |   |   |   |   |   | • |   |   |   |   |   |   |
| <b>Category 3 - Centralization</b>  |   |   | • |   |   |   |   |   | • |   |   |   |   |   |   |   |   |   |   |
| Code 3.1 - Specific app store   |   |   | • |   |   |   |   |   | • |   |   |   |   |   |   |   |   |   |   |
| Code 3.2 - Plugins  |   |   |   |   |   |   |   |   |   | • |   |   |   |   |   |   |   |   |   |
| <b>Category 4 - Reminders</b>   |   |   | • | • | • |   |   |   |   |   |   |   | • | • |   |   |   | • | • |
| Code 4.1 - "Opt-in" data access renewal   |   |   | • | • | • |   |   |   |   |   |   |   | • | • |   |   |   | • | • |
| Code 4.2 - "Opt-out" data access renewal  |   |   |   |   |   |   |   |   |   |   |   |   |   |   | • |   |   |   |   |
| Code 4.3 - Only information   |   |   | • |   | • | • |   |   |   |   |   |   | • | • |   |   |   |   |   |
| <b>Category 5 - Revocation Assistance</b>   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   | • | • | • | • |
| Code 5.1 - Assistance for access revocation when uninstalling TPA's service       |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   | • | • | • | • |
| Code 5.2 - Assistance for asking TPA to remove the user's data from their servers |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   | • |   |
| Code 5.3 - Automatic revocation   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| <b>Category 6 - Sensitization, Education</b>                                      |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |
| Code 6.1 - Video  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   | • |   |
| Code 6.2 - Informative and Interactive consent form                               |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   | • |   |
| <b>Category 7 - TPAs limit</b>  |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   |   | • |

Table A.1: Coding table. Each column corresponds to one specific design. For each design, we display the id of the group who designed it, this id corresponds to the group ids in Table 5.1.