Studies in Neuroscience, Psychology and Behavioral Economics

Harald Baumeister
Christian Montag *Editors*

Digital Phenotyping and Mobile Sensing

New Developments in Psychoinformatics

Springer

Studies in Neuroscience, Psychology and Behavioral Economics

More information about this series at http://www.springer.com/series/11218

Harald Baumeister · Christian Montag
Editors

Digital Phenotyping and Mobile Sensing

New Developments in Psychoinformatics

Springer

Editors

Harald Baumeister
Department of Clinical Psychology
and Psychotherapy
Institute of Psychology and Education
Ulm University
Ulm, Germany

Christian Montag
Department of Molecular Psychology
Institute of Psychology and Education
Ulm University
Ulm, Germany

ISSN 2196-6605 ISSN 2196-6613 (electronic)
Studies in Neuroscience, Psychology and Behavioral Economics
ISBN 978-3-030-31619-8 ISBN 978-3-030-31620-4 (eBook)
https://doi.org/10.1007/978-3-030-31620-4

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Foreword for Digital Phenotyping and Mobile Sensing

It is an axiom in the business world that you can't manage what you can't measure. This principle, usually attributed to the business guru Peter Drucker, is equally true in medicine. Imagine managing diabetes without HbA1c, hypertension without blood pressure readings, or cancer diagnosis without pathology. Perhaps the foundational measure in medicine was thermometry. The discovery that our subjective sense of being "chilled" accompanied the objective evidence of body temperature rising and our subjective sense of "hot" matched a fall in body temperature eliminated forever the idea that we could manage in medicine without objective measurement. We need objective data to understand and interpret subjective experience.

Unfortunately, the field of mental health has failed to benefit from the kinds of measurement that revolutionized business and medicine. While it is true that we have a century of psychological research on objective tests of cognition, mood, and behavior; little of this science has translated into clinical practice. Over the last half-century, biologically oriented researchers have followed the medical model, exploring blood, urine, and cerebrospinal fluid in the hope of finding the equivalent of HgA1c or some circulating marker of mental illness. EEG readings have been mined like EKG tracings. Brain scans and protocols for brain imaging have become ever more sophisticated in the hopes of finding the engram or some circuit dysregulation or a causal lesion. And more recently, genomics seemed a promising path to finding a biomarker for mental illness. In oncology, the most clinically useful genetic signals have proven to be somatic or local mutations in tumors, not germ line genetic variants found in blood cells which is the basis of psychiatric genetics. Nevertheless, we continue to seek causal signals in circulating lymphocytes assuming these will reflect the complex genomics of brain.

This half-century search for biological markers for mental states has been, for patients, a roller coaster of hype followed by disappointment. Thus far, science has not delivered for patients with mental illness the kind of measurement that has transformed care for people with diabetes, hypertension, or cancer. There are many potential reasons why we have failed to discover objective markers. The most common explanation is that brain and behavior are more complicated than glucose regulation or vascular tone or uncontrolled cell division. Finding the EEG signal for psychosis or the brain signature of depression will take longer. I accept this excuse, but there are three other explanations that are worth our consideration.

First, most clinicians rightly value the subjective reports of their patients as the most critical data for managing mental illness. They point out that the subjective experience of pain, anxiety, or despair is the hallmark of a mental disorder. They are not looking for quantitative, objective measures. Instead, clinicians hone skills of observation to translate their patients' reports into something more objective, usually defined by clinical terms if not a clinical numerical score. Master clinicians base their assessments not only on what they observe in the patient but on their own subjective experience, which they have learned to use as a barometer of paranoia or suicidal risk. While this approach, combining the subjective reports of the patient with the subjective experience of the clinician, might work for the provider, patients are increasingly expecting something better. Many patients realize, just as we learned from thermometry, that they cannot trust their subjective experience. Just as people with diabetes learn that every moment of lethargy is not hypoglycemia and people with hypertension learn that every headache does not mean elevated blood pressure, people with mental disorders are asking for something more objective to help them to manage their emotional states, distinguishing joy from the emergence of mania and disappointment from a relapse of depression.

A second reason for our failure to develop objective markers is that we lack a ground truth that can serve as the basis for qualifying a measurement as accurate. This is one reason why it took 200 years for thermometry to become a standard for managing an infectious disease—we had no simple proof of the value of body temperature, especially when the measure did not conform with subjective experience. Much of the clinical research on measuring biological features of mental illness has tried to validate the measure against a diagnosis. If only 40% of patients with Major Depressive Disorder had abnormal plasma cortisol levels, then measuring cortisol could have little value as a diagnostic test. The problem here is that Major Depressive Disorder does not represent ground truth. It is simply a consensus of master clinicians who voted that five of nine subjective symptoms constituted Major Depressive Disorder. And none of those symptoms, including sleep disturbance and activity level, are actually measured.

For me, the most important shortcoming in our approach to measurement is that we have put the cart before the horse: we are attempting to find biological correlates of cognition, mood, and behavior before we have better objective measures of cognition, mood, and behavior. Our measures, when we make them, are usually at a single point in time (generally during a crisis), captured in the artificial environment of the lab or clinic, and represent a burden to both the patient and the clinician. Ideally, objective measures would be captured continually, ecologically, and efficiently.

That ideal is the promise of mobile sensing, which has now become the foundation for digital phenotyping. As described in detail in this volume, wearables and smartphones are collecting nearly continuous, objective data on activity, location, and social interactions. Keyboard interactions (i.e., reaction times for typing and tapping) are being studied as content-free surrogates for specific cognitive domains, like executive function and working memory. Natural language processing tools are transforming speech and voice signals into measures of semantic coherence and sentiment. Of course, the rich content of social media posts, search queries, and voice assistant interactions can also provide a window into how someone is thinking, feeling, and behaving. Digital phenotyping uses any or all of these signals to quantify a person's mental state.

While most of the focus for digital phenotyping has been on acquiring these signals, there is a formidable data science challenge to converting the raw signals from a phone or wearable into valid, clinically useful insights. What aspects of activity or location are meaningful? How do we translate text meta-data into a social interaction score? And how to define which speech patterns indicate thought disorder or hopelessness? As you will see in the following chapters, machine learning has been employed to solve these questions, based on the unprecedented pool of data generated. But each of these questions requires not only abundant digital data, we need some ground truth for validation. Ground truth in academic research means a clinical rating, which we know is of limited clinical value. Ground truth in the real world of practice is functional outcome, which is difficult to measure.

It's useful to approach digital phenotyping or, as it is called in some of these chapters, psychoinformatics, as a work in three parts. First, we need to demonstrate the feasibility. Can the phone actually acquire the signals? Will people use the wearable? Will there be sufficient consistent data to analyze? Next, we have the validity challenge. Does the signal consistently correlate with a meaningful outcome? Can the measure find valid differences between subjects or is it only valid comparing changes across time within subjects? Can this approach give comparable results in different populations, different conditions, different devices? Finally, we face the acid test: is the digital measure useful? Utility requires not only that the signals are valid but that they inform diagnosis or treatment in a way that yields better outcomes. Patients will only use digital phenotyping if it solves a problem, perhaps a digital smoke alarm that can prevent a crisis. Providers will only use digital phenotyping if it fits seamlessly into their crowded workflow. As a chief medical officer at a major provider company said to me, "We don't need more data; we need more time".

Mastering feasibility, validity, and utility will also require engaging and maintaining public trust. Trust is more than ethics, but certainly the ethical use of data, consistent protection of privacy, and full informed consent about the phenotyping process are fundamental. Trust also involves providing agency to users, so that they are collecting their data for their use. There may be technical assets that can help. For instance, processing voice and speech signals internally on the phone might prove useful for protecting content privacy. The use of keyboard interaction signals, which consist of reaction times and contain no content, might be more trustworthy for some users. But it is unlikely tech solutions will be sufficient to overcome the appearance of and very real risk of surveillance. It is important, therefore, as you read the following chapters that you distinguish between the use of this new technology in a medical setting where consenting patients and families can be empowered with information versus the use of this technology in a population where monitoring for behavioral or cognitive change can be a first step down a slippery slope toward surveillance.

If we can earn public trust, there is every reason to be excited about this new field. Suddenly, studying human behavior at scale, over months and years, is feasible. Recent research is proving out the validity of this approach, already in thousands of subjects for some measures. We have yet to see clinical utility but there is every reason to expect that in the near future, digital technology will create objective, effective measures. Finally, in mental health, we may be able to measure well and manage better. Patients are waiting.

Thomas R. Insel, MD
Mindstrong Health, Mountain View, CA, USA
e-mail: tom@mindstronghealth.com

Contents

Foreword for Digital Phenotyping and Mobile Sensing v
Thomas R. Insel

Digital Phenotyping and Mobile Sensing In Psychoinformatics—A Rapidly Evolving Interdisciplinary Research Endeavor xiii
Harald Baumeister and Christian Montag

Part I Digital Phenotyping and Mobile Sensing: Privacy and Ethics

1 Privacy in Mobile Sensing 3
Frank Kargl, Rens W. van der Heijden, Benjamin Erb and Christoph Bösch

2 Ethical Considerations of Digital Phenotyping from the Perspective of a Healthcare Practitioner 13
Paul Dagum and Christian Montag

Part II Digital Phenotyping and Mobile Sensing in Psycho-Social Sciences

3 Computerized Facial Emotion Expression Recognition 31
Mattis Geiger and Oliver Wilhelm

4 An Overview on Doing Psychodiagnostics in Personality Psychology and Tracking Physical Activity via Smartphones 45
Rayna Sariyska and Christian Montag

5 Smartphones in Personal Informatics: A Framework for Self-Tracking Research with Mobile Sensing 65
Sumer S. Vaid and Gabriella M. Harari

6 Digital Brain Biomarkers of Human Cognition and Mood 93
Paul Dagum

7 Mining Facebook Data for Personality Prediction: An Overview 109
Davide Marengo and Michele Settanni

8 Orderliness of Campus Lifestyle Predicts Academic Performance: A Case Study in Chinese University 125
Yi Cao, Jian Gao and Tao Zhou

Part III Digital Phenotyping and Mobile Sensing in Health Sciences

9 Latest Advances in Computational Speech Analysis for Mobile Sensing 141
Nicholas Cummins and Björn W. Schuller

10 Passive Sensing of Affective and Cognitive Functioning in Mood Disorders by Analyzing Keystroke Kinematics and Speech Dynamics 161
Faraz Hussain, Jonathan P. Stange, Scott A. Langenecker, Melvin G. McInnis, John Zulueta, Andrea Piscitello, Bokai Cao, He Huang, Philip S. Yu, Peter Nelson, Olusola A. Ajilore and Alex Leow

11 Studying Psychopathology in Relation to Smartphone Use 185
Dmitri Rozgonjuk, Jon D. Elhai and Brian J. Hall

12 Connecting Domains—Ecological Momentary Assessment in a Mobile Sensing Framework 201
Thomas Kubiak and Joshua M. Smyth

13 Momentary Assessment of Tinnitus—How Smart Mobile Applications Advance Our Understanding of Tinnitus 209
Winfried Schlee, Robin Kraft, Johannes Schobel, Berthold Langguth, Thomas Probst, Patrick Neff, Manfred Reichert and Rüdiger Pryss

14 Mobile Crowdsensing in Healthcare Scenarios: Taxonomy, Conceptual Pillars, Smart Mobile Crowdsensing Services 221
Rüdiger Pryss

15 mHealth Applications: Potentials, Limitations, Current Quality and Future Directions 235
Eva-Maria Messner, Thomas Probst, Teresa O'Rourke, Stoyan Stoyanov and Harald Baumeister

16 Using Chatbots to Support Medical and Psychological Treatment Procedures: Challenges, Opportunities, Technologies, Reference Architecture 249
Rüdiger Pryss, Robin Kraft, Harald Baumeister, Jens Winkler, Thomas Probst, Manfred Reichert, Berthold Langguth, Myra Spiliopoulou and Winfried Schlee

17 Persuasive E-Health Design for Behavior Change 261
Harald Baumeister, Robin Kraft, Amit Baumel, Rüdiger Pryss
and Eva-Maria Messner

18 Optimizing mHealth Interventions with a Bandit 277
Mashfiqui Rabbi, Predrag Klasnja, Tanzeem Choudhury,
Ambuj Tewari and Susan Murphy

Digital Phenotyping and Mobile Sensing In Psychoinformatics—A Rapidly Evolving Interdisciplinary Research Endeavor

Harald Baumeister and Christian Montag

Many scientists are currently considering whether we are seeing a paradigm shift in the psychosocial and behavioral health sciences from narrow experimental studies to ecological research driven by big data. At the forefront of this trend is the implementation of smart device technologies in diverse research endeavors. This enables scientists to study humans in everyday life on a longitudinal level with unprecedented access to many relevant psychological, medical, and behavioral variables including communication behavior and psychophysiological data. Although the smartphone without doubt presents the most obvious “game changer” (Miller 2012), it only represents a small part of a larger development toward the Internet of Things, where everything from household machines to the car will be connected to the Internet (Montag and Diefenbach 2018). Therefore, human interaction with all these Internet-connected devices will leave digital traces to be studied by scientists in order to predict bio-psycho-social variables ranging from personality to clinical variables including states of physical and mental health (Markowetz et al. 2014; Montag and Elhai 2019).

The present volume gives an overview on current developments in this area, looking at digital phenotyping and mobile sensing as two prominent approaches in *Psychoinformatics*, i.e., the research field that combines innovative technological attempts with the psychosocial and behavioral health science traditions (Montag et al. 2016). Digital phenotyping extends the construct of phenotypes as the observable (biological) traits of organism to digital traces of people in a digital era of mankind (Jain et al. 2015; Insel 2017). Given almost omnipresent human–machine interactions, people’s digital traces might allow for diagnostic, prognostic, and intervention activities in different areas of life such as predicting product needs (e.g., by GPS tracking or bio-sensing approaches; even microtargeting (see Matz et al. 2017)), estimating personality traits, attitudes and preferences (e.g., predicting people’s political orientation by social–network interaction (Kosinski et al. 2013)) or improving patients health care (e.g., estimating disease and treatment trajectories based on ecological momentary assessment data). Mobile sensing is the most prominent driver of this new approach, given the already substantial penetration of smartphones around the world (at the time of writing 3.3 billion smartphone users have been estimated (Statista.com 2019)).

Systematizing this dynamic and fast developing field of research is challenging, as technological development in diverse and unconnected research areas might already have stimulated the next wave of innovations prior to this book being published. However, while the specific approaches might vary, a first framework for using digital phenotyping and mobile sensing in psychosocial and behavioral health sciences can be proposed, as is depicted in Fig. 1.

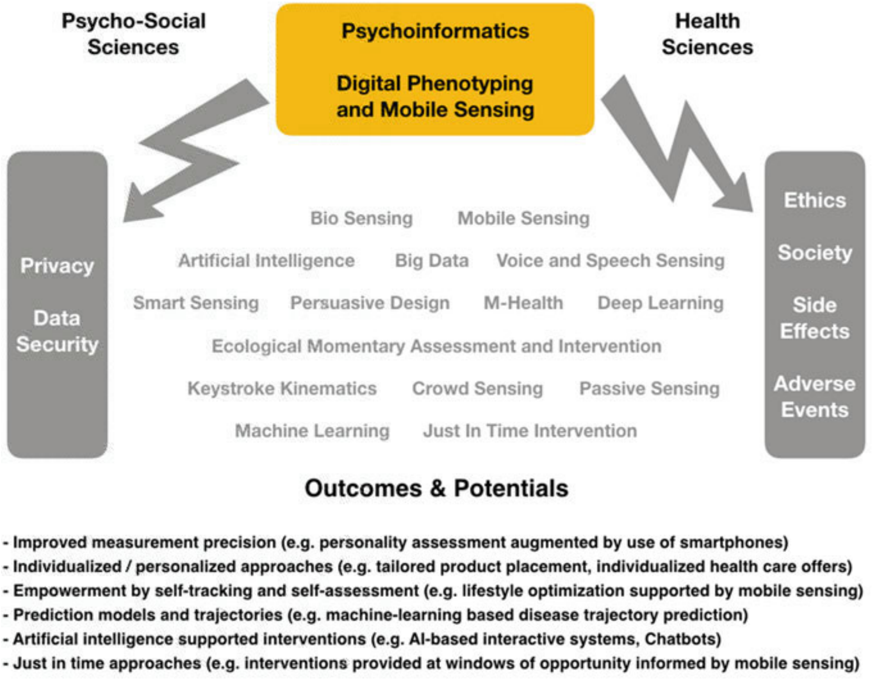

- Improved measurement precision (e.g. personality assessment augmented by use of smartphones)
- Individualized / personalized approaches (e.g. tailored product placement, individualized health care offers)
- Empowerment by self-tracking and self-assessment (e.g. lifestyle optimization supported by mobile sensing)
- Prediction models and trajectories (e.g. machine-learning based disease trajectory prediction)
- Artificial intelligence supported interventions (e.g. AI-based interactive systems, Chatbots)
- Just in time approaches (e.g. interventions provided at windows of opportunity informed by mobile sensing)

Fig. 1 Digital phenotyping and mobile sensing conceptual framework

A multitude of buzzwords such as machine and deep learning, big data, crowdsensing, bio-sensing, EMA, and EMI (ecological momentary assessment/intervention) are frequently used to express the impact of digitalization on people's lives and on societies as a whole. Note that we use the term *digitalization* here, because it describes how the use of digital technologies shapes society (or here scientific research), whereas digitization refers to the mere process of transforming analog into digital data.

Mobile sensing is often specified by the device providing the mobile sensing data, i.e., smart sensing devices such as smartphone, smartwatch, and smart-wearables, and by the data a smart sensing device is tracking, e.g., voice and speech sensing, bio-sensing, passive sensing, crowdsensing, and ecological momentary assessment or facial emotion recognition sensing. Using these terms already stimulates our imaginations regarding both risks and potentials associated with this still new technology, capable of altering peoples′ and societal life to a degree that has not yet been fully grasped.

Digital phenotyping constitutes one relevant application in the area of mobile sensing, with the potential to substantially improve our knowledge in the realm of latent constructs such as personality traits and mental disorders. A better understanding of these variables is of great relevance, because personality traits such as conscientiousness are good predictor for a healthy living style (Bogg and Roberts 2004) and mental disorders are a tremendous source of individual suffering and high costs for society (Trautmann et al. 2016). At the same time, digitally supported health care offers seemingly are capable of improving mental and behavioral health (Ebert et al. 2017; Bendig et al. 2018; Ebert et al. 2018; Paganini et al. 2018).

Our life has become digital and this digital image of our lives and persons can be ephemeral or used to provide the data basis necessary to estimate people's traits, states, attitudes, cognitions, and emotions (Montag et al. 2016; Martinez-Martin et al. 2018; Lydon-Staley et al. 2019). What does your smartphone usage pattern tell us about you and your state of mind? What does your vacation, social, and work life pictures posted in social networks tell us about your happiness and your attitude toward work, holiday time, and your spouse? Would we recognize a change in mental state when comparing voice recordings from today and 5 years ago? What will your lunch look like tomorrow? You might not know it, but maybe your bio-sensing signals will tell us now already. Two advances enable us to provide increasingly sophisticated digital phenotyping estimates: Big Data and machine/deep learning approaches.

Big Data constitutes a precondition for digital phenotyping based on a granular matrix of our digital traces, consisting of a multitude of (longitudinally) assessed variables in large cohorts coming at different degrees of velocity (speed), variety (data format), and volume. In short, Big Data can be described by varying degrees of VVVs (Markowetz et al. 2014). Once those databases have been established, we are confronted with a large set of complex data for which established statistical methods are often not the best fit.

Machine learning and deep learning (ML/DL) are the buzzwords that promise to make sense out of the big data chaos (Lane and Georgiev 2015). While some rightly argue for a more well-thought through scientific basis in the current machine learning hype (Kriston 2019), these analytical approaches are undoubtedly the key for the last puzzle of digital phenotyping and mobile sensing covered in the present volume: artificial intelligence (AI). This said, ML as an integral part of AI comes with many problems such as a lack of understanding of what kind of patterns the computer actually recognizes or learns when predicting a variable such as a tumor from an MRI scan (Mohsen et al. 2018). Aside from this, a demanding topic is that of programming ethics into deep learning algorithms (the field of machine ethics, see in, e.g., Moor and James 2006; Brundage 2015).

Artificial intelligence (AI) might become central to several fields of application, by pattern recognition using deep learning algorithms (Ghahramani 2015; Topol 2019). For instance, in a first development step, users of AI-based medical programs will be supported in interpreting diagnostic results and receiving AI-based prognostic feedback regarding the current treatment course of their patients. Predicting economic or environmental impacts of political decisions and tailoring product placement according to peoples' personality based on AI algorithms are further likely fields of application. Once developed, these artificial intelligence applications might again use mobile sensing techniques to further improve their prognostic power by means of a deep-learning-based self-improvement cycle (see Fig. 2).

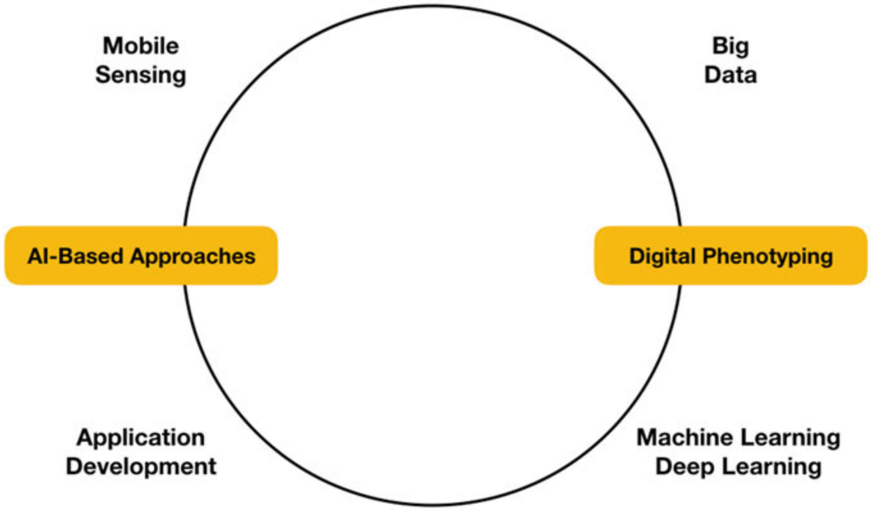

Fig. 2 Mobile sensing, digital phenotyping, and artificial intelligence life cycle

The chapters of this book provide a snapshot of what is already possible and what science might allow in the near future. Thus, most chapters not only focus on the areas of applications and the potentials that come along with these approaches but also on the risks that need to be taken into account, principally in terms of data privacy and data security issues as well as ethical and societal considerations and possible side effects of mobile sensing approaches. Regarding the risks, chapters by Kargl and colleagues (Chap. 1), who provide an overview of privacy issues as inherent aspect of mobile sensing approaches, and by Dagum and Montag (Chap. 2), who reflect on ethical implications of digital phenotyping, consider the ethical boundaries of our actions. Examples of unintended de-anonymization of data summarized by Kargl et al. (Chap. 1) shows how easily supposedly anonymized data can quickly become person identification data when combined with the almost infinite information on everything and everyone in our Internet of Things world. Answers on these ethical questions need a scientific discourse on how to exploit the potential of mobile sensing in an ethical way but also a societal discourse on what we are willing to accept in light of the conveniences mobile sensing approaches provide (e.g., accepting that Google knows where we are as a trade-off for using Google maps to navigate through traffic or find our ways in unknown places).

With these privacy and ethics boundaries in mind, readers of this book are provided with a look into the future that is already happening.

Several chapters provide exemplary research and conceptualization frameworks on mobile sensing approaches across the psychosocial and behavioral health sciences fields. Digital phenotyping, mental health prediction models, ecological momentary assessments, and academic performance estimates (Cao, Gao, and Zhou: Chap. 8; Kubiak and Smyth: Chap. 12; Marengo and Settanni: Chap. 7; Rozgonjuk, Elhai and Hall: Chap. 11; Sariyska and Montag: Chap. 4; Schlee et al.: Chap. 13; Vaid and Harari: Chap. 5) are only some of the possibilities in the realm of mobile sensing discussed in this book. The chapters range from established fields of mobile sensing that can already draw on substantial evidence (Saeb et al. 2015; Sariyska et al. 2018; Montag et al. 2019), such as predicting personality or mental and behavioral health status by means of smartphone usage patterns to less established fields such as the potential of bio-sensing, which might allow in future such things as physiologically delineated measures to improve health (e.g., cortisol implants measuring stress-related symptoms that could inform a mobile health application to provide a just-in-time intervention).

A second group of chapters focuses on the potential for machine learning and deep learning approaches. Geiger and Wilhelm (Chap. 3), for example, illustrate the research potentials of combining mobile devices with face recognition software allowing for immediate facial emotion expression recognition based on machine learning algorithms. Similarly, Hussain and colleagues (Chap. 10) reflect on the potential of machine-learning-based keyboard usage and corresponding typing kinematics and speech dynamics analysis for predicting mental health states. Regarding speech-analysis-based machine learning approaches, Cummins and Schuller (Chap. 9) present a selection of already available open-source speech analysis toolkits along with a discussion on their potentials and limitations.

Finally, four chapters provide frameworks and insights into how these fields of research can be used and combined to inform complex intervention developments in order to further improve people's health and living. While Messner and colleagues (Chap. 15) report on the current state of research regarding mHealth Apps and the potential that comes with new sensing and AI-based approaches, Pryss and colleagues (Chap. 16) present a framework for chatbots in the medical field. While most of the current chatbot approaches are based on a finite amount of answer options the chatbot can access (Bendig et al. 2019), the framework presented in this chapter looks at how the expert knowledge database necessary for complex communication situations such as psychotherapy can be generated and iteratively improved until a truly artificial intelligent chatbot therapist is at place. Baumeister et al. (Chap. 17), with a focus on persuasive-design-based intervention development, and Rabbi et al. (Chap. 18), with a model for just-in-time interventions, provide further details on how technology can be used to further improve existing interventions, enhance intervention uptake and adherence, and ultimately increase effectiveness by exploiting the full potential of mobile sensing.

Writing these few paragraphs on the content of our book fills us with excitement about the perspectives digital phenotyping and mobile sensing offer for research and practice. At the same time, however, we feel uneasy in light of the obvious risks for individuals and society at large. Researchers should not usually argue based on their emotions, but in this case these two emotions—positive and negative—might guide the next steps by a development process for future innovations that is ethical and informed on issues of privacy. It therefore seems that the development of new technical solutions will take place anyway given their potential economic value, leaving research to establish conceptual frameworks and guidelines to set the guardrail. Focusing on the example of artificial intelligence, large-scale companies will probably revolutionize the product market with ever more intelligent systems, increasing the convenience of consumers and at the same time reducing human workforce needs. However, these companies will probably not provide the urgently needed answers on how to develop and implement such innovations in a way that benefits society (Russell et al. 2015) considering all the scenarios relating to potentially malevolent AI (Pistono and Yampolskiy 2016). Broadening the focus again, we need to establish good scientific practice for mobile sensing in order to exploit its full potential. First, scientists currently discuss whether the paradigm shift postulated at the beginning of this editorial needs to undergo fine-tuning, setting the ecological correlation research approach into the context of explorative and explanatory research paradigms. Exploratively fishing for hypotheses is the beginning and not the end of methodologically sound psychosocial and behavioral health science (Kriston 2019).

This said, at the end of this short introduction we want to express our gratitude to all our authors for their important chapters in this book. They all invested a lot of their time and energy to provide insights into their different research perspectives.

We have not mentioned so far that Thomas Insel, former director of the National Institute of Mental Health (NIMH) in the USA and a prominent advocate of the digital phenotyping movement, was kind enough as to provide us with his thoughts on this relevant research area. His efforts are also much appreciated.

References

Bendig E, Bauereiß N, Ebert DD, Snoek F, Andersson G, Baumeister H (2018) Internet- and mobile based psychological interventions in people with chronic medical conditions. Dtsch Aerzteblatt Int 115:659–665

Bendig E, Erb B, Schulze-Thuesing L, Baumeister H (2019) Next generation: chatbots in clinical psychology and psychotherapy to foster mental health—a scoping review. Verhaltenstherapie 1–15. 10.1159/000499492

Bogg T, Roberts BW (2004) Conscientiousness and health-related behaviors: a meta-analysis of the leading behavioral contributors to mortality. Psychol Bull 130:887–919. 10.1037/0033-2909.130.6.887

Brundage M (2015) Limitations and risks of machine ethics. In: Risks of artificial intelligence. Chapman and Hall/CRC, pp 141–160

Ebert DD, Cuijpers P, Muñoz RF, Baumeister H (2017) Prevention of mental health disorders using internet-and mobile-based interventions: a narrative review and recommendations for future research. Front Psychiatry 8. 10.3389/fpsyt.2017.00116

Ebert DD, Van Daele T, Nordgreen T, Karekla M, Compare A, Zarbo C, Brugnera A, Øverland S, Trebbi G, Jensen KL, Kaehlke F, Baumeister H (2018) Internet-and mobile-based psychological interventions: applications, efficacy, and potential for improving mental health: a report of the EFPA E-Health taskforce. Eur Psychol 23:167–187. 10.1027/1016-9040/a000318

Ghahramani Z (2015) Probabilistic machine learning and artificial intelligence. Nature 521:452–459. 10.1038/nature14541

Insel TR (2017) Digital phenotyping. JAMA 318:1215. 10.1001/jama.2017.11295

Jain SH, Powers BW, Hawkins JB, Brownstein JS (2015) The digital phenotype. Nat Biotechnol 33:462–463. 10.1038/nbt.3223

Kosinski M, Stillwell D, Graepel T (2013) Private traits and attributes are predictable from digital records of human behavior. Proc Nat Acad Sci 110:5802–5805. 10.1073/pnas.1218772110

Kriston L (2019) Machine learning's feet of clay. J Eval Clin Pract jep.13191. 10.1111/jep.13191

Lane ND, Georgiev P (2015) Can deep learning revolutionize mobile sensing? In: Proceedings of the 16th international workshop on mobile computing systems and applications—HotMobile'15. ACM Press, New York, New York, USA, pp 117–122

Lydon-Staley DM, Barnett I, Satterthwaite TD, Bassett DS (2019) Digital phenotyping for psychiatry: accommodating data and theory with network science methodologies. Curr Opin Biomed Eng 9:8–13. 10.1016/J.COBME.2018.12.003

Markowetz A, Błaszkiewicz K, Montag C, Switala C, Schlaepfer TE (2014) Psycho-informatics: big data shaping modern psychometrics. Med Hypotheses 82:405–411. 10.1016/j.mehy.2013.11.030

Martinez-Martin N, Insel TR, Dagum P, Greely HT, Cho MK (2018) Data mining for health: staking out the ethical territory of digital phenotyping. npj Digit Med 1:68. 10.1038/s41746-018-0075-8

Matz SC, Kosinski M, Nave G, Stillwell DJ (2017) Psychological targeting as an effective approach to digital mass persuasion. Proc Natl Acad Sci U S A 114:12714–12719. 10.1073/pnas.1710966114

Miller G (2012) The smartphone psychology manifesto. Perspect Psychol Sci 7:221–237. 10.1177/1745691612441215

Mohsen H, El-Dahshan E-SA, El-Horbaty E-SM, Salem A-BM (2018) Classification using deep learning neural networks for brain tumors. Futur Comput Informatics J 3:68–71. 10.1016/J.FCIJ.2017.12.001

Montag C, Baumeister H, Kannen C, Sariyska R, Meßner E-M, Brand M, Montag C, Baumeister H, Kannen C, Sariyska R, Meßner E-M, Brand M (2019) Concept, possibilities and pilot-testing of a new smartphone application for the social and life sciences to study human behavior including validation data from personality psychology. J 2:102–115. 10.3390/j2020008

Montag C, Diefenbach S (2018) Towards homo digitalis: important research issues for psychology and the neurosciences at the dawn of the Internet of Things and the digital society. Sustainability 10:415. 10.3390/su10020415

Montag C, Duke É, Markowetz A (2016) Toward psychoinformatics: computer science meets psychology. Comput Math Methods Med 2016:1–10. 10.1155/2016/2983685

Montag C, Elhai JD (2019) A new agenda for personality psychology in the digital age?. Personality Individ Differ 147:128–134. 10.1016/j.paid.2019.03.045

Moor JH, James H (2006) The nature, importance, and difficulty of machine ethics. IEEE Intell Syst 21:18–21. 10.1109/MIS.2006.80

Paganini S, Teigelkötter W, Buntrock C, Baumeister H (2018) Economic evaluations of internet-and mobile-based interventions for the treatment and prevention of depression: a systematic review. J Affect Disord 225:733–755

Pistono F, Yampolskiy RV (2016) Unethical research: how to create a malevolent artificial intelligence. In: Proceedings of ethics for artificial intelligence workshop (AI-Ethics-2016). New York

Russell S, Dewey D, Tegmark M (2015) Research priorities for robust and beneficial artificial intelligence. AI Mag 36:105. 10.1609/aimag.v36i4.2577

Saeb S, Zhang M, Karr CJ, Schueller SM, Corden ME, Kording KP, Mohr DC (2015) Mobile phone sensor correlates of depressive symptom severity in daily-life behavior: an exploratory study. J Med Internet Res 17:e175. 10.2196/jmir.4273

Sariyska R, Rathner E-M, Baumeister H, Montag C (2018) Feasibility of linking molecular genetic markers to real-world social network size tracked on smartphones. Front Neurosci 12:945. 10.3389/fnins.2018.00945

Statista.com (2019) Number of smartphone users worldwide from 2016 to 2021 (in billions). https://www.statista.com/statistics/330695/number-of-smartphone-users-worldwide. Accessed 16 Oct 2019

Topol EJ (2019) High-performance medicine: the convergence of human and artificial intelligence. Nat Med 25:44–56. 10.1038/s41591-018-0300-7

Trautmann S, Rehm J, Wittchen H (2016) The economic costs of mental disorders. EMBO Rep 17:1245–1249. 10.15252/embr.201642951

Part I
Digital Phenotyping and Mobile Sensing: Privacy and Ethics

Chapter 1
Privacy in Mobile Sensing

Frank Kargl, Rens W. van der Heijden, Benjamin Erb and Christoph Bösch

Abstract In this chapter, we discuss the privacy implications of mobile sensing and modern psycho-social sciences. We aim to raise awareness of the multifaceted nature of privacy, describing the legal, technical and applied aspects in some detail. Not only since the European GDPR, these aspects lead to a broad spectrum of challenges of which data processors cannot be absolved by a simple consent form from their users. Instead appropriate technical and organizational measures should be put in place through a proper privacy engineering process. Throughout the chapter, we illustrate the importance of privacy protection through a set of examples and also technical approaches to address these challenges. We conclude this chapter with an outlook on privacy in mobile sensing, digital phenotyping and, psychoinformatics.

1.1 Introduction

While mobile sensing provides substantial benefits to researchers and practitioners in many fields including psychology, the data collected in the process is often sensitive. The data collected by smartphones and other devices with sensors, such as fitness trackers, is clearly related or relatable to persons. Therefore, the researcher or practitioner that collects, processes, and stores such data has moral and legal obligations to handle this data responsibly. This is especially important if the data is related to health or mental disorders of a person.

F. Kargl (✉) · R. W. van der Heijden · B. Erb · C. Bösch
Institute of Distributed Systems Ulm University, Albert-Einstein-Allee 11,
89081 Ulm, Germany
e-mail: frank.kargl@uni-ulm.de

R. W. van der Heijden
e-mail: rens.vanderheijden@uni-ulm.de

B. Erb
e-mail: benjamin.erb@uni-ulm.de

C. Bösch
e-mail: christoph.boesch@uni-ulm.de

H. Baumeister and C. Montag (eds.), *Digital Phenotyping and Mobile Sensing*,
Studies in Neuroscience, Psychology and Behavioral Economics,
https://doi.org/10.1007/978-3-030-31620-4_1

The right to protection of personal data has been recognized as a central human right and is, for example, embedded in the European Charter on Human Rights.[1] In this chapter we want to raise awareness of the importance of privacy and data protection. To this end, we introduce privacy from a legal, technical, and applied perspective, as well as discussing some of the associated challenges. In particular, we would like to dispel the myth that a consent form from a study participant relieves the researcher from legal obligations. Beyond legal obligations, we discuss some evidence that a lack of privacy may negatively affect the participants' or patients' trust in systems or procedures. Finally, we will provide a positive outlook on how both the legal and ethical obligations could be achieved by proper privacy engineering and the application of privacy-enhancing technologies (PETs).

Privacy protection has already been recognized as an important issue within the psycho-social research community, after controversial incidents such as the Tearoom Trade study by Humphreys. The name of the study refers to male-male sexual behavior in public bathrooms. In his work, Humphreys not only surveyed unwitting subjects in extremely private and intimate situations (sexual intercourse in public bathrooms) without their consent, he also collected personally identifiable data (license plates) to later de-anonymize the subjects and visit their homes under false pretenses for follow-up interviews. This study demonstrated the fatal effects when personal data is collected in studies without regard to privacy (Kelman 1977).

One of the first to investigate the anonymization of health data was Latanya Sweeney, who showed that anonymous hospital discharge records contained sufficient information to de-anonymize 87% of the US population by matching the zip code, gender, and date of birth information of the records to US census data (Sweeney 2002). Many more examples have led to the conclusion that proper anonymization becomes extremely hard if the opponent has sufficient context knowledge.

De-anonymization is also an issue for location privacy and other data collected by mobile devices such as smartphones and fitness trackers. For example, the company Strava released data from its fitness tracking service, where many people uploaded GPS traces from their daily runs. People analyzing this massive dataset quickly found out that it contained runs from soldiers from supposedly undisclosed military bases in Afghanistan and elsewhere.[2] Similarly, anonymous trip records published from New York taxis have been used to identify trips from celebrities and find out whether they tipped the driver or not.[3] The extent by which seemingly innocent data gives away our most intimate information is best illustrated by an example reported in the New York Times where shopping data from customers was analyzed to learn that a high-school girl was pregnant even before her own family knew.[4] Such examples have led to the conclusion that good anonymization is really hard and the category

[1] http://www.europarl.europa.eu/charter/pdf/text_en.pdf.

[2] https://www.theguardian.com/world/2018/jan/28/fitness-tracking-app-gives-away-location-of-secret-us-army-bases.

[3] https://research.neustar.biz/2014/09/15/riding-with-the-stars-passenger-privacy-in-the-nyc-taxicab-dataset/.

[4] https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html.

of personal identifiable information needs to be broadened up substantially. This is also reflected in the modern understanding of privacy and current lawmaking, such as the European General Data Protection Regulation (GDPR).

1.2 Privacy as a Multifaceted Concept

1.2.1 Privacy as a Legal Concept

Many countries have regulated different aspects of privacy in their laws. Legal protection of personal data is termed data protection. One of the most holistic data protection frameworks is the European General Data Protection Regulation (GDPR), which went into force throughout Europe in May 2018, unifying the different data protection regimes throughout Europe.

While it is beyond the scope of this chapter to provide a complete overview of the GDPR, we still use it here to illustrate major concepts of data protection. GDPR, or "Regulation (EU) 2016/6791" as it is officially named, regulates the *processing* of *personal data* relating to a natural person in the EU, by an individual, a company or an organization. As described in Article 4, *personal data* includes data that is indirectly related to a person, while *processing* has to be understood to include activities such as collection or storage of personal data.

The GDPR places many requirements on anyone that either conducts data processing (termed "data processor") or is responsible for data processing (termed "data controller"). In addition, "data subjects", which are the natural person(s) to whom the data relates, are given a broad set of rights, such as the right to be informed or the right to erasure of the processed data. It is important to note that any processing of personal data by default requires *informed consent* of the data subject—exceptions being, e.g., if that data is required for legal reasons or to fulfill a contract with that person. In psychological studies, such informed consent is typically the basis for processing of personal data. However, such consent does not free the data controller or processor from the broad set of obligations that come with the right to process personal data (Schaar 2017). These obligations include the information rights of data subjects, such as the right to be informed, the right to access such data in a portable format, and the right to object to data processing, even after the fact. This obviously clashes with some obligations of researchers on research data management.

In this chapter, we want to focus on yet another aspect of the GDPR: Privacy by Design (PbD) and Privacy by Default. Especially when processing sensitive data or in high-risk cases, the GDPR requires any system that processes personal data to be designed as privacy-friendly from the ground up. This is done by following a PbD design and development process based on a Data Protection Impact Assessment (DPIA) that investigates potential privacy issues right from the start. The GDPR further requires that state-of-the-art technical and organizational measures (TOMs) be integrated into the foundations of every system. However, it is currently still open how courts will decide on the definition of "state-of-the-art".

This will have substantial implications on how psychological research can be conducted. In the remainder of this chapter, we will first illustrate these consequences and then address how such privacy and data protection may be achieved. Finally, we will provide an outlook on how modern privacy engineering might also be applied to conduct research in psycho-social research in a compliant and responsible way.

1.2.2 Privacy as a Technical Concept

In the technical literature, privacy is often defined in a quantifiable way, which greatly simplifies the analysis of technical solutions employed to protect privacy. The simplest conception of privacy in the technical literature is anonymity, which means that "a subject is not identifiable within a set of subjects, referred to as the anonymity set" (Pfitzmann and Hansen 2010). The anonymity set of a subject is the set of subjects that have the same properties (e.g., age category, gender, disease characteristics) so that one cannot distinguish one particular person from the others in the group. If the anonymity set is a database published as research data, removing names or pseudonyms from the subjects may not be enough to achieve anonymity, since each subject may have other unique properties. Another related but distinct concept is that of linkability, which is defined between "two or more items of interests (e.g., subjects, messages, actions, ...) [...], the attacker can sufficiently distinguish whether these are related or not". This is a stronger requirement than that of anonymity, since knowing whether two messages originate from the same source does not (necessarily) identify the source. However, in many cases, linkability of messages implies the possibility of de-anonymization. For example, in the location privacy example of New York taxis above, linkability of locations led to de-anonymization of individuals, as explained (Douriez et al. 2016).

To design and validate practical privacy enhancing technologies (PETs), many technical articles also quantify the privacy associated with a specific system. In particular, many different metrics exist (Wagner and Eckhoff 2018) to quantify exactly how anonymous a subject is within their anonymity set. The fundamental challenge is that many de-anonymization attacks typically work using external information sources. This led to the development of many classes of metrics, the most well-known of which are data similarity (where k-anonymity is a widely-known example) and indistinguishability (where differential privacy is a prominent example). We refer interested readers to Wagner and Eckhoff (2018) for a detailed survey of these metrics. But how does this all apply to psychological research?

1.2.3 Privacy as a Concept in Research Studies and Treatments

In the context of research studies, experiments, and treatments related to healthcare and psychological well-being, privacy is particularly relevant from two perspectives. First, the protection of data from participants or patients raises strong obligations for researchers or therapists, as the ethical principles in these professions go well beyond the legal requirements. Second, the participants' or patients' perception of their privacy influences their trust in the procedure, which can potentially even negatively affect the results or the outcome of the treatment.

Organizations such as the American Psychological Association (APA) take into account privacy obligations as part of their ethical principles. For instance, section four of the APA Ethics Code specifically addresses privacy and confidentiality, requiring that the confidentiality of collected information is maintained, as well as the minimization of privacy intrusions. Best practices have been established to adhere to these principles in traditional settings (e.g., usage of pseudonymization codes). However, novel approaches such as mobile sensing and smartphone-based data collection, entail new threats to privacy and confidentiality, which cannot be addressed with existing practices alone. In fact, an increasing technologization of experiments and treatments requires equal advancements in the safeguarding and (technical) implementation of ethical principles.

Orthogonal to the ethical considerations, addressing the privacy concerns of participants and patients has a positive impact on the procedure outcome. According to a model of Serenko and Fan (2013), informational privacy (i.e., information acquisition and ownership) has the strongest influence on a patient's privacy perception in the healthcare context. The level of perceived privacy is associated with the level of trust of the patient in the treatment. Again, this trust level is associated with the behavioral intentions of the patient such as commitment, adherence, and compliance with the treatment. Furthermore, anonymity and confidentiality in studies work against the social desirability bias (Krumpal 2013)—a tendency to give answers that are considered to be favorably by other peers. Joinson (1999) found similar effects in early, web-based questionnaires.

On the other hand, recent developments in psycho-social research have shown many results were not reproducible by further studies. This has given rise to the open science movement, whose primary goal is to improve reproducibility through data availability. A wide variety of data management platforms, such as the open science framework (https://osf.io) and Zenodo (https://zenodo.org), have risen in this context, whose primary aim is to widely disseminate research data. However, sensitive personal data clearly cannot be published in this fashion under the GDPR regime; technical solutions are required to ensure that the data is used only in correspondence with the consent provided by the user, under the obligations posed by the GDPR.

As we have shown, privacy is a complex and challenging concept to data processing in psycho-social research from many perspectives. We now continue to refine these challenges in more detail before discussing possible technical solutions.

1.3 Challenges in Privacy

While mobile sensing enables new forms of participatory research and discovery by collecting almost endless amounts of sensor data, such systems create distributed and massive databases of individuals' data and maybe even of unrelated surrounding people. This collection and processing of large amounts of sensor data leads to a unique source of information about, e.g., environmental conditions and changes, user activity patterns and health, and behavioral habits. By collecting this various data about the human body as well as its direct and extended environment, it is possible to create precise personal profiles including massive amounts of sensitive information. This data can often even be used to make behavioral predictions about an individual. Thus, protecting this data from unauthorized access during the whole data lifecycle, i.e., during its collection, transfer, processing, storage, potential release or final deletion, is a major challenge.

Even anonymization of data often mentioned in this context cannot convincingly eliminate this risk. Often one can infer identities from anonymous data when linking it to other (public) data. An example of such indirect leakage is data regarding location information. Tracking a person's location inevitably implies collecting and processing data about the whereabouts—and thus the behavior—of persons. This might infer sensitive information about users' activities, inter alia, sexual habits/orientation, drinking and social behavior, physical or mental health, religious and/or political beliefs. A famous example of the analysis of driving behavior is the 2012 Uber blogpost "Rides of Glory",[5] in which Uber showed how to spot candidates for one-night stands among its riders. While the risks from location-based attacks are fairly well understood given years of previous research, our understanding of the dangers of other modalities (e.g., activity inferences, social network data) are less developed.

For some data, the impact on an individual's privacy appears insignificant at first glance, but contains sensitive information that can be derived or inferred from the data. Most users are not aware of the amount and extent of data as well as the expressiveness and significance of the collected data. Since Sweeney's work (Sweeney 2002) on de-anonymization through seemingly anonymized, innocent data, many other examples of de-anonymization attacks have been reported. This includes those of the web search queries of over half-million America Online (AOL) clients (Barbaro and Zeller 2006) and the movie reviews of a half-million Netflix subscribers (Narayanan and Shmatikov 2008).

In AOL's case, user IDs of search queries were replaced with unique identifiers per user to allow researchers using the data to access the complete list of a person's search queries. Unfortunately, the complete lists of search queries were so thorough that individuals could be de-anonymized simply based on the contents of search queries. In the Netflix case, Narayanan and Shmatikov used a different approach and re-identified several Netflix users by correlating the available research data with publicly available movie ratings data. In addition, they were able to infer more sensitive information from the available data since "[m]ovie and rating data contains

[5] https://web.archive.org/web/20140828024924/http://blog.uber.com/ridesofglory.

information of a more highly personal and sensitive nature. The member's movie data exposes a Netflix member's personal interest and/or struggles with various highly personal issues, including sexuality, mental illness, recovery from alcoholism, and victimization from incest, physical abuse, domestic violence, adultery, and rape."[6]

These examples have shown that the concept of personally identifiable information (PII) is a challenging concept that is not as straight-forward as it appears. Due to the diversity and efficiency of modern de-anonymization algorithms (Narayanan and Shmatikov 2008, 2009), it is often possible to re-identify an individual, even in the absence of personally identifiable information. While some attributes in personal data may not be uniquely identifying on their own, almost all information can be personal or identifiable when combined with enough other relevant context information (Narayanan and Shmatikov 2010).

Furthermore, an intentional release or unintended leak of personal data brings new challenges, since a data set cannot be protected to preserve privacy once it is public. A major challenge is thus to determine whether the data stored or to be released is sufficiently and adequately anonymized. There is a growing number of examples and techniques for reconstruction attacks, where data that may look safe and innocuous to an individual user may allow sensitive information to be reverse-engineered. This also means that your data may have to be regularly reconsidered for their privacy sensitivity.

Starting from a background of legal and ethical obligations to safeguard privacy, we have illustrated how difficult and complex the notion of anonymization and privacy protection is. It becomes evident how privacy and data protection need to become essential elements in mobile sensing. Motivated by data protection regulation such as the GDPR, one needs to actively consider how to limit the collection, flow, use/processing, storage, release, and deletion of personal data in own research. We will next illustrate what role Privacy Enhancing Technologies (PETs) can play in this protection and how they can be embedded in a privacy engineering process to design mobile sensing systems.

1.4 Privacy Protection

We will start our discussion on how to properly protect privacy in mobile sensing with a recap of some GDPR principles as listed in Article 5. It already foresees legitimate purposes related to scientific research and statistics, for which specific rules are outlined in Article 89. There is also an exemption that states personal data may be archived for these purposes (which would normally violate the storage limitation principle). However, these processing and archiving of personal data for scientific research is subject to the condition that technical *and* organizational measures (TOMs) are taken to protect the data. As discussed earlier, Article 25 of the GDPR specifies that data protection should be done by design and by default

[6] http://www.wired.com/threatlevel/2009/12/netflix-privacy-lawsuit.

but it leaves open what this could mean in practice. In this section, we now discuss examples how one can practically implement this requirement.

Hoepman (2014) introduced a set of *strategies* to improve privacy during the design process. They provide overarching categories to then derive more specific design patterns that can be used to finally select appropriate technical and organizational measures to protect personal data. Hoepman's strategies are called *Minimize*, *Separate*, *Aggregate*, *Hide*, *Inform*, *Control*, *Enforce*, and *Demonstrate*. Taking *Aggregate* as an example, this strategy states that "Personal data should be processed at the highest level of aggregation and with the least possible detail in which it is (still) useful". For publication and archiving of scientific data, this strategy is often used when dealing with particularly sensitive information, such as personal profiling or health-related data. On the technical side, privacy design strategies are used to classify privacy patterns, the goal of which is to provide proven best practices for some settings and enable "off-the-shelf" usage of PETs. As discussed by Hoepman (2014) and others, not all strategies have received sufficient attention from research, and the technological maturity varies greatly between PETs. However, recent years have seen significant improvements in this regard and we highly recommend the use of these strategies and review of available privacy patterns and PETs when considering processing of personal data, especially in the context of mobile sensing or big data.

Returning to the fitness tracker example from the introduction, we now describe the practical application of the *Aggregate* strategy. For the purpose of monitoring fitness activity and overall health, a fitness tracker collects information such as the average heart rate, as well as the minimum and maximum rates during a sport session. Rather than collecting each pulse measurement and sending it directly to a server where it is stored, the aggregate strategy recommends local aggregation of this information in a way that is consistent with application goals. For example, to monitor overall health and track fitness, the average, maximum and minimum heart rates over the past minute can be computed in the device and transmitted. This information reveals a lot less about activities compared to per-second measurements. In this specific example, note how the aggregate strategy has benefits beyond privacy; removing unnecessary data from the collection process reduces the bandwidth, storage, and processing time required to analyze the data, illustrating that privacy and data processing are not always a zero-sum game. Similar aggregation can be performed by discretization of continuous data (e.g., only collecting heart rate only as low, medium, and high heart); however, the information loss associated with this process may not be suitable in every application. For this reason, the choice of PETs is inherently dependent on what the data is used for; this is one reason why the GDPR requires a specification of purpose for data processing and entails that domain experts and privacy experts need to work hand in hand in such projects.

Another relevant example is the use of analysis techniques from "big data". The purpose of these techniques is to extract useful patterns from large volumes of data. In such settings, the volume of data is too large to be analyzed on a regular computer, and thus computations are performed on the data by the database or a cluster. In medical and psychological applications, the interesting patterns are typically correlations between observable behaviors and markers on the one hand and associated conditions on the other. These can be computed by the database directly, therefore removing the need for direct access to the data, although this access is often still available. Cryptographic research has dedicated significant effort (Lindell and Pinkas 2002) to the design of *privacy-preserving data mining* techniques, whose goal is to extract such patterns reliably, while being able to provably *hide* the individual inputs and thus protect them from misuse. This privacy design strategy particularly suitable for situations where aggregation or minimization are difficult to apply. In those situations, privacy-preserving data mining techniques offer a compromise between privacy and data access: a researcher can run certain analyses on encrypted data, but not retrieve an individual's data from the dataset. For example, using suitable protocols like Secure Multi-Party Computation or Secure Function Evaluation, it is possible to calculate certain averages on data without learning any individual data items anywhere in the system. However, this often requires to carefully think about the analysis one wants to conduct before collection of data, since the modes of analysis are necessarily restricted to preserve privacy. Here, careful compromises between privacy and use of the data need to be found. Solutions that limit data access to policy-compliant operations like shown by Kargl et al. (2010) and SGX-based solutions like the one from Al-Momani et al. (2018) may allow more flexible analysis but provide different and maybe weaker privacy guarantees. Another aspect to consider is the technological maturity of these approaches is still often limited and constrained by high computation overhead that some of these solutions still imply.

1.5 Conclusion and Outlook

With this chapter we aimed to raise the awareness of privacy in mobile sensing. We outlined how legal and ethical considerations require every researcher applying mobile sensing to carefully consider the privacy implications of this data collection, and to apply a privacy by design and default process to come up with the best protection possible while still keeping the data useful for the researchers.

Current legal frameworks for privacy and data protection clearly state that collecting a consent form from study participants is not sufficient in most cases. Additional technical and organizational measures should be put in place to increase privacy protection. This typically involves collaboration with privacy experts with appropriate legal and technical background. As privacy is an inherently interdisciplinary topic, privacy researchers are always eager to find interesting, challenging use cases and data to which their technologies may be applied. The appropriate and visible application of such technologies can lead to an increase in user trust, while also reducing the impacts of specific biases, such as social desirability bias.