RoboCode-Ethicists – Privacy-friendly robots, an ethical responsibility of engineers?

Christoph Lutz

Institute for Media & Communications
Management, University of St. Gallen,
Blumenbergplatz 9, CH-9000
christoph.lutz@unisg.ch

Visiting: Oxford Internet Institute,
University of Oxford, 34 St Giles
UK-Oxford OX1 3LD
christoph.lutz@oii.ox.ac.uk

Aurelia Tamò

Chair for Information and
Communication Law, University of
Zurich, Rämistrasse 74/49, CH-8001
aurelia.tamo@uzh.ch

Visiting: Institute for Pervasive
Computing, ETH Zurich,
Universitätsstrasse 6, CH-8092 Zurich
aurelia.tamo@inf.ethz.ch

ABSTRACT

This article asks why engineers building robots should consider privacy aspects when programming their gadgets. We start with a definition of robots, differentiating active, social robots from passive, non-social robots. We then discuss the related literature on the privacy implications of social robots. Two aspects are of fundamental concern in this context: the pervasiveness and intrusiveness of robots on the one hand and a general lack of awareness and knowledge about how robots work, collect and process sensitive data on the other hand. We explain how the existing literature on robot ethics provides a suitable framework to address these two issues. In particular, robot ethics are useful to point out how engineers' and regulators' mindset towards privacy protection differs. The paper argues that different – at first sight incommensurable – rationalities exist when it comes to robotic privacy. As a contribution to the emerging field of robotic privacy, we propose an interdisciplinary and collaborative approach that bridges the two rationalities. This approach considers the role of code as the central governing element of robots. RoboCode-Ethicists, trans-disciplinary experts trained in the technical/computational, legal and social aspects of robotics, should lead the way in the discussion on robotic privacy. They could mediate between different stakeholders – mainly regulators, users and engineers – and address emerging privacy issues as early as possible.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

WebSci '15, June 28 - July 01, 2015, Oxford, United Kingdom

© 2015 ACM. ISBN 978-1-4503-3672-7/15/06...\$15.00

<http://dx.doi.org/10.1145/2786451.2786465>

Keywords

Privacy, Robots, Ethical Coding, Applied Ethics.

Categories and Subject Descriptors

I.2.9. [Artificial Intelligence]: Robotics – *commercial robots and applications*. K.7.4. [The Computing Profession]: Professional Ethics – *ethical dilemmas*. K.4.1. [Computers and Society]: Public Policy Issues – *privacy*.

General Terms

Algorithms, Performance, Design, Security, Human Factors, Legal Aspects.

1. INTRODUCTION

Recent media coverage about robots abounds. A PEW report on the Internet of Things [52], for instance, predicted that the next revolution in digital technology will be led by increased embedded and wearable computing. Previously known only by computer scientists, the terminologies Ambient Intelligence (AmI), Artificial Intelligence (AI), or Internet of Things (IoT) are taking form and their potential is recognized by the wide population [53]. Important characteristics of “smart” environments are the ability to capture and take into account the context of particular scenarios as well as the ability to adapt to individual users [72].

Currently, robots are used in many professional and social contexts, such as labor and services, military and security, research and education, healthcare, as personal companions or as toys [42]. The world's robot population goes in the millions and the numbers are rapidly increasing. A current estimation assumes that between 2013 and 2016, 22 million robots will be sold [36].

Yet this “robot revolution” also raises broader ethical questions. Already in 2004 scientists formulated their expectations for the next-generation robots in the Fukuoka World Robot Declaration [31, as cited in 69:29]. Those include that (1) robots will be partners that coexist with human beings, that (2) they will assist human beings both physically and psychologically, and that (3)

they will contribute to the realization of a safe and peaceful society. In particular with respect to informational privacy, concerns about the seamless and imperceptible data collection of robots were raised. The aim of this article is to ponder on the ethical question of why engineers¹ should consider designing privacy-friendly RoboCode. The article contributes by showing ways how this can be done.

The paper contains three main parts. After the introduction, we discuss the topic of privacy in the context of robots (Chapter 2). We define the concept of robot and distinguish different forms (2.1). Moreover, we show how and why privacy in the context of robots matters (2.2). We propose that two characteristics of robots make them especially susceptible for privacy violations: their constant collecting and processing of data in the background (calmness) and the black box problem, i.e. users' unawareness of and missing knowledge about how robots work and how the algorithms they apply function. In Chapter 3 we focus on ethical issues in robotics. We present a brief outline of robot ethics (3.1), and present two different rationalities in this context, applying them to the issue of robot privacy: the practical, pragmatical engineer's rationality (3.2) and the normative regulator's approach (3.3). We show how these two views – that are represented by two main lines of thoughts of robot ethics – conflict (3.4). Finally, in Chapter 4, we argue for a middle ground between these perspectives. In doing so, we lay a focus on the role of code. We discuss how code affects privacy protection in robots and why we should consider such issues from a coding perspective.

2. ROBOTS & PRIVACY

2.1 Characteristics of Robots

The science of robotics originated from technical research fields such as mechanics, physics, computer science, cybernetics, automation and control. Recently, it has drawn upon several other disciplines such as logic, linguistics, neuroscience, psychology, physiology, anthropology, art, design and others [69]. With the amplified spectrum of involvements, the complexity of issues around machines “capable of carrying out complex series of actions automatically” (Oxford Dic.) has increased too. The intricacy already starts with defining the characteristics of a robot, let alone evaluating the consequences of the employment of robots in daily settings. The following subchapters define robots, delineate them from “traditional” machines and elaborate on the typology of robots before dwelling upon the privacy implications of their use.

2.1.1 Sense, Process, Weigh, Act

Bekey [6:18] provides a useful working definition of the term, defining robots as “a machine, situated in the world, that senses, thinks, and acts. Thus, a robot must have sensors, processing ability that emulates some aspects of cognition and actuators.” This “sense-think-act” paradigm has been acknowledged in the literature [13, 22, 62:67].

Yet, the term “think” should be treated with some reservation. Indeed, Bekey [6:18] puts the term consistently in brackets.

¹ In this paper the terminologies (software, robotics) engineer and (software, robotics) developer are used as synonyms. Engineers have a broader, more strategic influence over the construction of robots, while developers are more focused on the actual implementation of code.

Thinking implies having a “particular belief or idea” (Oxford Dic.), i.e., having formed an opinion or particular thought in one's mind. Whether this attribute is a distinctively human characteristic is not central to the argument in this paper. However, it is important to acknowledge that (so far) robots process incoming information and weigh their potential reactions to this collected data according to a pre-determined set of rules programmed into their RoboCode. In other words, robots sense information, process it, weigh potential outcomes out of possible actions and act upon those computations as programmed to do. Therefore, we propose to use the term “*sense-process-weigh-act*” paradigm.

Furthermore, the attribute “situated in the real world” distinguishes robots from software bots [6]. In essence, robots are complex, programmed devices able to intake, filter and act upon real-world information.

Notwithstanding, the nature of complex technologies makes a clear-cut definition difficult. Therefore, it seems useful to distinguish robots from “simpler” technologies by defining the attributes that other technologies do not have [13], i.e. elaborating on the paradigm change robots bring along. By this logic, for example, a traditional car equipped with cameras has some degree of ability to sense the outside world, yet relies – so far – usually on the driver to act upon the information sensed. A self-driving car on the other hand, equipped with cameras and able to sense, “think” and act upon the information of the world around it, would be considered a robot under the “sense, process, weigh, act” paradigm.

2.1.2 Towards Greater Autonomy

Another characteristic of robots rests upon the degree of autonomy and self-governance. Darling [19:fn. 8] specifies that the term autonomy in robotics can mean “as little as the ability to *perform tasks without continuous human input* or control.” The criterion here is whether individuals remain “in the loop” when robots operate [10:252]. Robots would have gained full autonomy if they were able to decide upon what actions to execute. Rather than merely following a predetermined action sequence, they base their decision on their own perception of the environment surrounding them [61]. Yet, even if robots can achieve a certain degree of autonomy, they remain, first and foremost, human instruments [13]. Thus, to some extent they are never fully autonomous, as they are programmed by humans.

Especially, the transition from “semi-autonomous” to “fully autonomous” is not always clear-cut. According to Del Campo et. al. [21], a fully autonomous robot is able to gain information about its surroundings, to work without human intervention for an extended period of time, to move through its environment without human support. It might be able to learn from and adjust itself to a changed environment. For example, self-driving cars as currently developed by Google are an example for almost fully autonomous robots, while parking assist systems or auto-pilots are semi-autonomous, since the driver or pilot can decide when to switch off the application and control the car him/herself.

2.1.3 Social & Non-Social Robots

Within the field of robots, various subtypes can be distinguished. One broad distinction is between social and non-social robots. Darling [19:4] defines a social robot as a “physically embodied, autonomous agent that communicates and interacts with humans on an emotional level.” Thus in this sub-class of robots the focus lies on the *interaction with humans*. For instance, robots that act as personal companions for elderly people (robot caregivers),

entertainment robots such as Pleo, or therapy robots such as Paro are social robots [13]. These social robots often imitate human behaviors, such as talking or showing emotions. They sometimes resemble humans in their looks and are constituted of clearly visible (human or animal) body parts. Individuals tend to anthropomorphize robots in general and social robots in particular [13, 19, 30, 60].

Unlike non-social robots, social robots are able to make (limited) decisions about their actions and behavior they exhibit. In a certain sense, their “process-weigh” trait is more elaborated than in industrial settings where actions of robots are more monotonous. Social robots base their decisions on their internal states and perceptions, while robots in industrial settings perform actions based on very specific preprogrammed commands [60]. The application of non-social robots, by contrast, is usually restricted to *industrial processes*. Their appearance is more machine-like [60].

While the dependability on non-social robots for industrial processes is unquestionable, the trend towards the adaptation of social robots is seen as more disruptive, as the immediate contact of social robots with individuals’ challenges and generates emotions and feelings. In particular, the emotional dependence on social robots [60] as well as the increased reliance and personalized adaptation raise issues with respect to privacy.

2.2 Privacy Implications of Robots

Technological advancements have led to controversies and fears around privacy long before the “robot revolution” [63]. Automated data processing machines have had a disruptive impact on the way information is collected, analyzed, employed or shared. The rupture embodied by seamless, dehumanized and sometimes invasive collection structures as well as intransparent processing patterns has raised privacy concerns – not only in social science but also in computer science [38]. The described abilities of robots to sense, process, weigh, and act upon the world around them likewise stir up privacy concerns. Privacy is a human, social phenomenon and privacy implications are thus social implications. Therefore, concerns arise in particular with social robots, which interact with humans, and less so with industrial, non-social robots, which interact mainly with machines or things.

In this article, our understanding of privacy refers mainly to *informational privacy*. In contrast to physical privacy, which encompasses the “access to an individual and/or the individual’s surroundings and private space” [63:990], informational privacy describes access and control of personally identifiable information (ibid.). Floridi [25:52] defines informational privacy as the “freedom from epistemic interference”. Robots challenge both, the access to and control over such personal information. This paper elaborates on the major privacy challenges of the “robot revolution”. Concerns in a democratic, modern society rest, among others, on the customized information collection and processing. In particular, the consequences of (1) opaque collection practices and issues resulting from the (2) ignorance of evaluation practices shall be discussed in more detail.

2.2.1 Pervasiveness & Acclimatization

To a certain extent, social robots fulfill Weiser’s [72:19] prophecy:

“The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it.”

While robots are not literally “weaved into the fabric” they merge into their social environment, accomplish certain needed activities and are recognized as part of our *daily entourage*, similarly as cellphones or wearable health trackers nowadays. They will merge into daily routines, collect and evaluate (personal) data on the go. The trend of so-called “co-inhabitant robots”, i.e., robots that are present “in our homes, assisting us in cleaning, housekeeping, child care, secretarial duties” [6:25] makes robots increasingly omnipresent, more tangible and embedded into our daily life.

The merging of technology into everyday life increases undetectable, subtle or *imperceptible data collection*. The seamless collection of data renders the process ordinary. Already today, users are accustomed to accept various service agreements, privacy policies, or cookie notices on the Web, when downloading mobile applications, or when signing in to accounts to retrieve the evaluated data of mobile devices. The notion of (informed) consent is the dominant legal mechanism for transactions involving data processing [12]. Users are given notice prior to the data gathering and consent, with a check mark, to those practices. By merging newer technologies into daily life users will get accustomed to agree to newer data collection practices that come with the increased employment of social robots. Users will consent to being surrounded by robots at home, and agree to the service contracts and privacy policies of the manufacturers, the software providers (running the RoboCode), remote operators, or companies selling or leasing the robot. As currently seen with mobile applications, consenting to such practices is easy and we get accustomed that app developers require us to grant access to our location, contacts and other personal information. In this sense, consent has been criticized being the main lever to undo legal restrictions and protections [15, 43:173, 65]. The latter point has implications for social robot developers as elaborated in 4.2.

The described developments are central as social robots increasingly impact privacy in three ways [13]:

(1) They facilitate *surveillance*. Robots, such as drones, often remain undetected when gathering information. Especially the military employs such machines to increase their capacity of surveillance. Surveillance triggers chilling effects, i.e., drones might inhibit individuals from engaging in certain activities [13, 59].

(2) They increase the *access* to (sensitive) personal information. In particular, social robots, which have access to our homes, or daily, private activities and life collect more detailed personal data than traditional website providers do. Examples of more intense data gathering devices include home robots or caregiver robots. Depending on the outcome of the processing and usage of the evaluated data as well as its dissemination, individual harms such as reputational or emotional damage could be triggered [13, 64].

(3) The pervasiveness of technology is amplified by the *social* bonding of humans and robots. The third issue deals with the “social meaning” of robots. Social robots decrease the opportunities for solitude as they become abundant and are mobile [67]. They are thus more likely to extract sensitive information from users. Combined with robots perfect memory and ability to link events and data, the increased aggregation of data catalyzes the potential of misuse of data [13].

These three aspects of robots' privacy implications may overlap and sometimes even reinforce each other. An example is a social robot with *access* to private rooms, such as an individual's home. Because of his proximity to the individuals in the home, the robot seamlessly *monitors* his surroundings, senses the individuals' whereabouts, actions, emotions and needs [67]. A scenario in which a robot enters an individual's bedroom to screen the environment while the latter is asleep is thinkable. Yet, some people would feel that thereby their privacy is infringed upon as the machine gathers data while the human is in a state of unconsciousness. In addition, the *social bonding* between individuals and robots increases the likelihood of sharing personal information with a robot, such as characteristic daily routines or emotions.

2.2.2 Black Box & Predictive Analysis

While a gradual acclimatization to robotic data gathering and evaluation is not *per se* unacceptable, such a changed perspective should be publicly discussed and the risks and benefits of current data processing practices analyzed before users become dull towards potential privacy issues – especially as the acclimatization to tracking, processing and targeting practices might reinforce the gap between users' awareness over how robots process data and the potential privacy implications of such processes. Therewith, the so-called "privacy paradox" might simultaneously be strengthened too. The privacy paradox refers to the tendency of users to be concerned about their privacy and fearful of a loss of control, but simultaneously not adapting their behavior, by, for example, disclosing less or choosing encrypted online services and other privacy-protecting techniques [66]. Therefore, the tendency to be ignorant vis-à-vis data collection and processing practices of robots should be carefully studied. Especially, because robots are more invasive than most mainstream Internet applications, the privacy repercussions might be more pronounced and far-reaching.

The problem of *ignorance* vis-à-vis data collection and evaluation practices has been referred to as the black box problem [47, 70]. It describes the lack of knowledge and understanding on how algorithms work. The distress rests upon the illiteracy and incapability of individuals to rationally understand the computations and resulting outputs of big data analysis [70]. Since users do not understand how their data are being analyzed, profiled, and used, the potential of misuse increases. This potential of misuse goes hand in hand with a *lack of check-and-balances*, i.e., the lack of being able to supervise the use of the data [34]. This issue has been termed the "awareness challenge" in the context of filtering and selection as well as data collection mechanisms on the Internet (such as search engines, social media algorithms such as Facebook's Edge Rank or recommender systems seen on online-shopping platforms, like Amazon). A large part of users is simply not aware that and how information about them is being collected, analyzed and traded to third parties, and how filtering algorithms work.

The argument here is obviously related to the opaque data collection mentioned above. The black box problem is further amplified when it comes to predictive analysis scenarios. Predictive analysis are conclusions based on the evaluation of an abundant amount of different data centered on an individual [18]. Privacy harms result from the data jigsaw employed to attain certain deductions and building assumptions on how an algorithm respectively a robot should react. Thereby, the individual and regulator lose control over the processing of data and data

protection legislations can hardly be enforced. Yet, the inferences of predictive analysis might damage the individuals, e.g., by leading to incorrect conclusions or discrimination [8, 18].

These privacy implications caused by the widespread adoption of (social) robots are complex and their impact is difficult to predict. Nevertheless, they touch upon a problem area, which has already been genuinely discussed in the field of information or computer ethics. The research on robot ethics constructively sheds light on the social impact of robots (and potential ways to handle them) - also in terms of privacy - and we therefore chose to build upon already existing literature in this field.

3. AN ETHICAL CLASH?

3.1 Robot Ethics

What should robots do? What laws and rules should they obey? What morals should they follow? Such questions touch the subject of robot ethics, a field concerned with the application of ethical principles to robotic behavior. Robot ethics are important in every area where robots are applied: military, households, social care, industry etc. Privacy is frequently described as a branch or topic of ethics [56, 63:991]. Thus, in this understanding, robotic privacy is a topic that can be analyzed within the wider margin of robot ethics (as we do here).

3.1.1 Asimov's Laws of Robotics as a Background

Science-fiction writer Isaac Asimov developed one of the earliest and probably the most widely known ethical code for robots. His so-called "Three Laws" state:

- (1) *A robot may not injure a human being or, through inactivity, allow a human being to come to harm.*
- (2) *A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.*
- (3) *A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.*

Clarke [17] discusses some problems and dilemmas with the Three Laws. One big issue is the vagueness of human communication and the potential for misunderstanding. Other problems arise when two or more humans give the robot conflicting orders or when humans are involved that want to harm other people, e.g., terrorists. By the First Law, the robot is not allowed to injure a terrorist. At the same time, his inactivity would lead to injuries or deaths among the attacked people.

To address some of the issues Asimov added the "Zeroth Law":

- (4) *A robot may not injure humanity, or, through inaction, allow humanity to come to harm.*

Despite their impetus and preeminent role in the field, the Three Laws are insufficient as a foundation for robot ethics. Improved ethical frameworks strive to come up with principles that can inspire moral learning and thus lead to self-reflection [4].

3.1.2 Ethical Frameworks & Information Ethics

Gips [33] distinguishes a range of different robot or machine ethical frameworks: consequentialist theories, deontological theories, virtue-based theories, and, AI approaches. The distinction between consequentialist and deontological theories with respect to privacy protection will be reflected in more depth in the subsequent parts (3.2 and 3.3).

In essence, the *consequentialist approach* considers the consequences of actions and recommends the actions with the best

possible outcomes. The most accepted consequence maximized is the personal luck, benefit or good that emerges from the action. Such a utilitarian view would suggest to maximize $\sum p_i w_i$, where p_i is the pleasure arising from the action for each individual i and w_i is the weight assigned for each individual i [33:245]. In a similar vein, Hospers [35:3] argues that one is “morally obliged to choose that action which maximizes total happiness (summed over all affected persons) according to utilitarian ethical theory.” Applied to robotic privacy, a consequentialist approach would design robots that decide what privacy-related action to take by considering the weight and pleasure of the involved individuals (this corresponds with the privacy calculus perspective).

Deontological theories, on the other hand, evaluate actions in themselves and not in terms of their consequences. Actions are judged by their imminent morality or immorality. The Ten Commandments are an example of a deontological moral system. Applied to robotic privacy, a deontological approach would come up with maxims that safeguard the value of privacy, such as “A robot may never violate a user’s privacy by involuntarily publishing sensitive information about him in the wrong context.”

Virtue-based theories focus on the being instead of doing. Instead of asking “What shall I do?” they ask “What shall I be?” [33:250, 3]. Schopenhauer’s two cardinal virtues (benevolence, justice) are a prominent example of a virtue-based ethics. Virtue-based theories seem difficult to apply to robotic privacy because privacy in itself is not a virtue. At the same time, an always privacy-respecting robot might be seen as a virtuous one at first sight. However, given that privacy is a continuum with an optimum, i.e., there can be both *too little* or *too much* privacy [46], such a robot might give the user too much privacy and be too careful and considerate.

Finally, the *AI approach* takes an explorative stance and wants to investigate and develop suitable ethical principles on-the-go, i.e., while concretely developing and investigating technological solutions. “The hope is as we try to implement ethical systems on the computer we will learn much more about the knowledge and assumptions built into the ethical theories themselves.” [33:251] Such a bottom-up perspective begins with the engineers and their perspectives. Applied to robotic privacy, such an approach would not treat ethical values as given in a new context, but as something to be explored and experimented with. The value-sensitive design paradigm is an example [27]. Here, good privacy-balancing mechanisms should be developed iteratively.

Robot ethics can be seen as part of a wider ethical framework, namely *information ethics*, which has been defined as „the branch of ethics that focuses on the relationship between the creation, organization, dissemination, and use of information, and the ethical standards and moral codes governing human conduct in society.”² Floridi [26] gives a concise overview of information ethics. In contrast to bio-centric ethics, which take living entities as the basis for moral principles and worthy of protecting, information ethics go a step further. They consider information as the principal entity of analysis, worthy of protection. In this sense, inanimate objects or machines – such as robots – and even abstract ideas have basic moral value because they contribute to the *infosphere* (the information environment). Several interesting

sub-domains can be derived from and analyzed within the broader framework of information ethics: computer ethics, AI or even robot ethics. Such sub-domains tend to be more applied and bottom-up, driven by concrete developments and emerging moral questions. The following subsection outlines a practical, applied perspective to robot ethics.

3.1.3 Applied Ethics & Domains of Application

Riek and Howard [58:6] follow a bottom-up and practical approach. They have established guiding principles for a Human Robot Interaction (HRI) code of ethics and distinguish four categories of consideration, where the principles apply: human dignity, design, legal, and social. The human dignity considerations are the most elementary ones and touch upon aspects such as the consideration of the emotional needs of humans and their right to privacy. Design considerations encompass, among others, transparency in the design process, predictability and the provision of opt-out mechanisms (kill switches). The legal aspects point to the necessity of robots abiding to current legislation. Finally, the social aspects entail considerations such as the anthropomorphic potential of many robots during the design process and the avoidance of racist or sexist appearance of robots.

Several of the principles encompass aspects of privacy [58:6]:

- (1) *The humans’ right to privacy shall always be respected to the greatest extent consistent with reasonable design objectives.*
- (2) *Maximal reasonable transparency in the programming of robotic systems is required.*
- (3) *Trustworthy system design principles are required across all aspects of a robot’s operation, for both hardware and software design, and for any data processing on or off the platform.*
- (4) *All relevant laws and regulations concerning individuals’ rights and protections are to be respected.*
- (5) *The tendency for humans to form attachments to and anthropomorphize robots should be carefully considered during design.*

While the first principle embodies a human dignity consideration (cf. 3.3), the second, third and fifth principles point towards design considerations for engineers. In other words, engineers should keep in mind the need for transparency and trust as well as the potential issues resulting from the human tendency to be attached to objects, when developing (social) robots.

In addition, the importance of privacy has been highlighted in the Code of Ethics of the Association for Computer Machinery (ACM), listing in paragraph 1.7 the *respect of the privacy of others* as a fundamental principle:

“This imperative implies that only the necessary amount of personal information be collected in a system, that retention and disposal periods for that information be clearly defined and enforced, and that personal information gathered for a specific purpose not be used for other purposes without consent of the individual.”

However, it is crucial to distinguish different *domains of application*. Ethically designed and privacy respecting robots must meet different requirements depending on the ethical and privacy norms at play in a given situation. Nissenbaum’s [49] notion of *contextual integrity* is useful to account for such context-specific privacy, also when it comes to robots. Contextual

² Joan, Reitz M. "Information Ethics." Online Dictionary For Library And Information Science. http://www.abc-clio.com/ODLIS/odlis_i.aspx

integrity describes an approach where data collecting entities respect the privacy norms in a given context – instead of collecting data on a catch-it-all basis. Thus, privacy is secured as long as data collectors fulfill the *norms of appropriateness* (what constitutes private information in a given situation) and *distribution* (how and to whom should information be given in a certain context). In practical robot ethics, this entails an understanding of these two norms in the application domain of the robot. Household robots, for example, are used in private spaces, where appropriate shared information between household members is quite encompassing and can include religious beliefs, sexual orientations and strongly emotional secrets (depending of course on the composition of the household). Medical and care robots, by contrast, are confronted with a context where appropriate information is much more restricted, e.g., to health indicators – blood pressure, weight, height – and to relatively insensitive personal data, such as age, gender and marital status. In terms of information distribution or flow, ethical robots should consider that in the household setting it might be more appropriate to act in a bidirectional, conversational way, while in the medical context, such a behavior could be understood as a privacy breach (since the norm of information distribution or flow are more unidirectional in the latter scenario).

3.2 Engineer's Rationality: Make it Work

Engineers design robots and when doing so they have certain goals in mind. Their main objective is to make it work, i.e., the robot should be functioning and flawlessly fulfill specific tasks it was designed to carry out. Such a pragmatic and functional approach (necessarily) leaves aside certain unintended negative consequences that can arise. Among others, because of their real-world agency robots can cause physical injury, emotional harm or threats to privacy.

Weber's concept of *means-end* or *instrumental rationality* captures some of the underlying foundations of the engineer's rationality. Weber [71] distinguishes four types of rationality: means-end/instrumental, value-/belief-oriented, traditional, and emotional rationality. In this contribution, we focus on the first and second type. Means-end or instrumental rationality describes a social action that is rationally pursued in terms of the outcomes. Thus, an individual carries out a certain action if it serves his purposes and meets her expectations. Such a scheme comes close to the notion of the *homo oeconomicus*. Someone who acts strictly within the logic of instrumental rationality optimizes the cost-benefit relation.

In our context, the instrumental rationality means evaluating a new technology in terms of its costs and benefits. Developers envision clear benefits from their innovations, such as simplifying tedious work tasks (the 3 Ds – dangerous, dirty, dull) [43] or providing emotional and physical support to solitary or handicapped people. They are certainly aware of some of the direct costs, such as development cycles, material costs or possible malfunctions. Privacy protection, however, does not form part of the (development) costs because the calculation and quantification of the costs in case of privacy breaches is complex, but also because engineers do generally not view privacy protection as a task within their domain of responsibility [29, 38]. According to Langheinrich [38:14], the main reasons why engineers do not consider privacy implications are: not feeling morally responsible; not seeing it a necessity yet or anymore; not considering privacy as a problem when building only prototypes;

or arguing that privacy is too abstract of a problem and/ or not part of the deliverables.

The slogan “make it work” captures the developer's rationality. When designing new products, engineers have a concrete outcome in mind. “Tool makers and users regularly evaluate how well tools accomplish the purposes for which they were designed.” [48:14] Central dimensions of evaluation include usability, correctness and reliability but not necessarily human values [29]. Every mean to achieve these goals – except for cost and resource restrictions – is legitimate [50]. In other words, the goal of the project justifies the processing of data even if privacy implications could arise.

3.3 Regulator's Rationality: Respect Privacy

The regulator's rationality differs from the engineer's rationality. It represents users' interests and points of view rather than the ones of engineers and designers. Instead of “making it work”, it puts “respect privacy” at the center of its actions. Hence, it caters much more to the value-/belief-oriented rationality than to the means-end or instrumental rationality [71]. *Value-/belief-oriented rationality* focuses on values as the main drivers of actions. Hence, individuals who act strictly according to value-rationality do not compromise their firmly held principles, even if a better outcome would be possible in case these principles are forgone. An environmentalist, who uses public transport instead of a car to go work, even if she spends double the time on the road, is an example of value-/belief-oriented rationality. This rationality corresponds with the *deontological* perspective described above and is close to a Kantian philosophy.

Regulators strive to protect individual rights and liberties. The right of privacy, enshrined in various human rights charters and national constitutions, is one fundamental right which regulators safeguard. Privacy is linked to fundamental societal interests such as the preservation of individual freedom and dignity or empowering thriving pluralism and democracy. Therefore, privacy is given a superior weight or value to be protected by normative means [12]. The duty of regulators rests in the establishment of national legislations and frameworks that defend the right to privacy. They also need to balance such rights with other interests, e.g., corporate innovation and economic growth – a non-trivial challenge [1].

In the case of privacy, regulators need to make sure that individuals' privacy is not involuntarily invaded. Because data collectors have an information and knowledge advantage vis-à-vis the end users, regulators primary focus is (or should be) on empowering the consumers, not the data collectors. Recent case law and legislative attempts in the European Union (Right to be Forgotten) show this tendency. In this contribution, we assign a somewhat ideal-typical user interest to the regulator in order to make the distinction between engineers and regulators – and thus our point – clearer.

Furthermore, a broad stream of research on online privacy shows that a large proportion of people in different countries are concerned about their privacy on the Internet [24, 44, 54, 66]. For example, 91 percent of respondents in a representative US survey agree or strongly agree with the statement: “Consumers have lost control over how personal information is collected and used by companies.” [54] While such concerns are not directly transferable to robots (and we are not aware of representative surveys assessing users' privacy concerns when it comes to robots), we suspect similar or even stronger privacy inhibitions in the case of robotic applications. From a consumer protection point

of view, regulators feel compelled to empower users. Privacy concerns of citizens and consumers are answered with political actions. In particular in Europe, regulators have stressed the importance of privacy protection via data protection legislation. In this vein, developer's instrumental rationality ("make it work") should factor in privacy considerations because of their importance for the end user (in terms of reputation gains) and regulators (in terms of compliance with the legal frameworks).

3.4 Make it Work vs. Respect Privacy – Conflicting Rationalities

A useful dichotomy to capture the different logics of engineers and regulators is the distinction of *efficiency* and *effectivity*. While efficiency asks whether the (pre-defined) things are done right (and remains silent about the goodness or rightness of the things to be done), efficacy/effectivity describes whether the right things are done. This can be in terms of economic benefits but also in terms of moral values. Hence, the developer's perspective functions more in terms of efficiency, while the regulator operates more in terms of effectivity. While regulators prefer to think in "terms of abstraction", engineers "like to think in terms of buildable designs" [2:59]. "Bridging these two cultures is not a trivial task" [2:59].

For robots, such conflicts are not yet in the foreground of attention because mass adoption in private households has yet to occur. However, some applications have sparked controversy between the two positions. The next chapter provides examples which illustrate the conflicting positions and look at why and how privacy – as an ethical value – can (or cannot) be encoded into robots. In the scholarly community, debates about how to regulate – if at all – autonomous technologies and robots are only emerging.

4. ENCODING PRIVACY IN ROBOCODE

4.1 A Path Worth Going Down?

The literature on how much robot engineers should encode ethical and legal standards when building their devices is divided – also in terms of privacy. On the one hand, proponents of a more constructivist approach to technology argue more in favor of a Laissez-Faire approach [23]. On the other hand, Calo [14] argues in favor of new regulation with respect to the robotic industry.

These clashing opinions on whether the production and encoding of ethical principles into technology is possible and desirable can be traced along existing discussion in philosophy of technology. One extreme position claims that technology itself is inherently neutral and only becomes value-laden when it is used in certain ways. This position corresponds largely with the social construction of technology (SCoT) paradigm [37, 55]. In contrast to the deterministic view of technology, this more socio-technical view states that technology cannot be understood outside the social and political context [45]. Under this perspective, technology is not an independent force, i.e., once set in motion it does not follow its own, independent course (in contrast to [73]). It is therefore impossible or even contra productive to regulate the production of technology via hard law, as the human involvement will considerably affect implementations.

"A robot is basically a computer that causes some physical change in the world. We can and do regulate machines, from cars to drills to implanted defibrillators. But the thing that distinguishes a power-drill from a robot-drill is that the robot-drill has a driver: a computer that operates it. Regulating that

computer in the way that we regulate other machines – by mandating the characteristics of their manufacture – will be no more effective at preventing undesirable robotic outcomes than the copyright mandates of the past 20 years have been effective at preventing copyright infringement (that is, not at all)." [23]

Instead, certain unwanted uses should be legislated (e.g., it should not be prohibited to produce knives but it should be prohibited to stab people with a knife) and self-regulating mechanisms applied, such as code reviews or rigorous user tests [23].

The contrasting position argues that technology is not neutral but often value-laden. According to Garfinkel, the notion that technology is neutral is a "comforting idea but it's wrong" [32:150]. Technological artifacts afford certain built-in consequences from their use – many of which are not neutral [9]. Applied to privacy Garfinkel [32:150] comes to the conclusion: "Although it's possible to use technology to protect or enhance privacy, the tendency of technological advances is to do the reverse. It is harder, and frequently more expensive, to build devices and construct services that protect people's privacy than to destroy it." In a similar vein, Oosterlaken sees technology as a "capability expansion" [51:94] and acknowledges the important role of developers and engineers when designing new technologies such as robots. She argues for a development of technology that considers and incorporates moral values in the design process.

Depending on the context-specificity and autonomy of a technology, these built-in consequences can be uniform and central: a spectacle case has a clearly-defined built-in consequence of protecting eyeglasses; or they can be varied and diverse: a smartphone has a range of built-in consequences, from being in touch with friends to helping people kill time at the bus station via app-based games [9]. Robots are one of the most autonomous technologies today. Thus, they have a multitude of built-consequences, which makes it difficult to judge them in terms of their embedded values. For example, caregiver robots main built-in consequence is assistance to elderly or handicapped people by carrying out tasks that are difficult for these people to do themselves. This consequence promotes the value "quality of life" – a positive outcome. Thus, it makes sense to develop caregiver robots and would be contra productive to ban their development. At the same time, caregiver robots, as an unintended side-consequence, can violate users' physical and informational privacy and become very invasive, just by their mere presence. In this paradigm, technology should/could be prohibited from being developed if it promotes disvalues rather than values.

Since privacy protection is a less tangible topic than the protection of a human life, a second example shall outline the intricacies of code regulation in the context of robots: the RoboGun. Such a robot consists of a mobile gun tied to a mechanical vehicle to move the gun, equipped with sensors to detect movements. It could be employed in private households as a guardian in conjunction with an alarm system, in the sense that – when the house owner switches the RoboGun on – it automatically detects and shoots at intruders. Is it possible and desirable to ban the manufacture of such RoboGuns (including the code and software for auto-detection of intruders as well as the shooting mechanism)? From a value-driven perspective – including the deontological frameworks as well as the regulator's point of view – the answer would be "yes"; from a SCoT, consequentialist and technology-neutral position the answer would be "no".

We posit that if there is a middle ground, where engineers and regulators come together and their rationalities are reconciled (cf. 3.4). In this sense, we aver that extreme forms of disvalue-producing technologies should be regulated by law, while more ambiguous and abstract technological artifacts should be dealt with other means, such as industry-wide agreements, rigorous use testing, scenario analysis or code reviews. In the following section we discuss the legitimacy of one specific encoding – namely privacy protection.

4.2 Privacy and Code in the Context of Robots

While the question of code or software is less dominant in classical robot ethics, it seems important to steer the discourse back towards the “infinitely reproducible nugget at the core of the system” [23], i.e. code. It is justified to link theories of robot ethics and the analysis on the regulation via code, since code, or RoboCode, is the foundation for every ethical discourse.

“If you accept that robots are just machines (...) and that the thing that makes them “robot” is the software that runs on a general-purpose computer that controls them, then all the legislative and regulatory and normative problems of robots start to become a subset of the problems of networks and computers.” [23]

In addition, code constrains the possible behaviors and interactions *ex ante*. Lessig [40, 41] and Reidenberg [57] have highlighted this rule-making phenomenon of code or technical architecture early on. In particular, Lessig’s oeuvre influenced the academic discourse around code and law. Lessig argues that computer code regulates – together with law, markets, and norms – the way we interact and employ technology [40, 41]. If code is believed to regulate, then this triggers fundamental questions about the (ethical) accountability of engineers designing such technical standards. It also raises questions about the political or social participation in the decision-making that produces such technical solutions or instruments [7].

Is it legitimate for regulators to ask engineers to build privacy protections into RoboCode? What are arguments, taking the more philosophical discourse outlined above (4.1) into consideration, in favor of regulatory oversight over built-in privacy protections in robots? The *theory of human rights law* [5, 11] remains most dominant when pondering the question of the legitimacy of code-based regulation. The fact that the Universal Declaration of Human Rights or the European Convention of Human Rights considers privacy as a fundamental human right presents an argument in favor of regulation. This argument will also be susceptible to different ethical views, such as the utilitarian view, advocating for maximizing human welfare, the human right view, focusing on the rights of individuals, and the more duty-driven rooted dignitarian perspective claiming that human dignity is an uncompromisable duty [11]. Those (ethical) paradigms influence the weight given to the notion of consent (being the dominant notion in liberal societies). The consent notion assumes that people are free to choose and enter “contracts” with companies, e.g., companies selling robots (see 2.2.1.). Depending on the ethical views, one can accept consent and thereby privacy infringement by consent or not. The utilitarian view for instance places high stakes on consent since there can be utility for a person to consent to something. By contrast, the human right and dignity perspective would not accept “consented” privacy infringements, as it places dignity as a higher good compared to

the freedom of contract [11]. Regulation should be considered whenever not regulating might cause harm. By considering the potential harm of a regulation, one takes a more pragmatic approach. Yet, numerating the “potential harm” will more often than not be a difficult undertaking [11]. Agreement might be easier to find with human dignity than privacy argumentation. Most people would certainly agree that research involving human beings with deadly outcomes must be regulated by legislation. Yet, the development of robots which interfere with our right to privacy – having less tangible consequences – are less unanimously subject of regulation.

Nevertheless, it seems reasonable, when acknowledging human rights as a fundamental value worth of protection, to ask for privacy protection in the context of robots. Also because such a value-driven perspective, compared to the means-end rationality of developers, is more inclusive and taking into account a broader range of perspectives. Negative externalities and long-term consequences of robots are not taken into account in a strict instrumental paradigm. The NSA scandal shows how the aim of engineers to ameliorate their data processing has simultaneously enabled the seamless monitoring by intelligence agencies. In addition, the potential of decreasing serendipity [34] and the consequences of profiling and being “fed” only tailored information is a negative externality not considered in the calculus of developers. Therefore, privacy protection should prevail even if it requires working on the “nugget of the system”, namely RoboCode. However, such a call for privacy should use realistic, applicable and feasible approaches. In the end privacy must be *realistically* encodable. In order to provide for a realistic foundation, both, engineers and regulators must take a step towards the “middle ground” (cf. 4.3).

4.3 Reconciling Clashing Rationalities

This part aims at reconciling the two contrasting perspectives outlined in 3.4. In applied ethics, confronting the practical ethical principles – e.g., the ACM Code of Ethics – with Nissenbaum’s [49] notion of contextual integrity means defining more fine-grained principles that take into account specific application contexts. This requires both: bottom-up mechanisms, breaking up the engineer’s rationality (4.3.1), and top-down ones, breaking down the regulator’s rationality (4.3.2)

4.3.1 Bottom-Up: Breaking Up the Engineer’s Rationality

Providing developers/engineers with clear-cut principles that account for users’ privacy considerations is a good way to respect their instrumental rationality. Lederer and colleagues [39] do that by elaborating five pitfalls to avoid when designing interactive systems: “(1) obscuring potential information flow, (2) obscuring actual information flow, (3) emphasizing configuration over action, (4) lacking coarse-grained control, and (5) inhibiting existing practice.”

The first two pitfalls affect users’ *understanding*, while the last three touch on their *action*. Applied to robotics, designers’ devices should (1) clearly show the users what kind of data can be collected, to whom they might be transferred, how long and where they are stored, and how much risk for unintended disclosure there is. They should also indicate clearly (2) what information is processed. “The disclosure should be obvious to the user as it occurs; if this is impractical, notice should be provided within a reasonable delay. Feedback should sufficiently inform but not overwhelm the user.” [39:446]

When it comes to concrete actions, privacy settings should be easy and intuitive to configure (3). Especially, they should not overwhelm the user. Applied to robots, designers should make sure that its privacy configurations are easy to grasp. A toy or conversation robot, for example, should refrain from asking the user about its desired level of privacy in difficult, abstract or legal jargon. Instead, it should talk in plain language, while being as precise as possible (“May I know about your favorite music or is this too personal for you?”). Similarly, value-sensitive design (VSD) aims at taking human values into account in an inclusive way when designing computer systems and code [9, 28]. The focus lies on “human values with ethical import”, including privacy, autonomy, accountability, human welfare, informed consent, and trust [28]. The VSD approach strives to design technologies that respect and balance different stakeholders’ values. It encompasses a set of tools and methods for designers to guide them with regards to different values in the design process [9:54].

Also, designers should offer visible mechanisms to switch off a device or some of its privacy-related functions (4). This is especially important for robots, since – due to their mobility and ability to act and move – they can be much more intrusive than other technologies. Finally, designers should be aware of existing practices and well-rounded frameworks (5). Here, considering established findings in Human Robot Interaction – especially in terms of privacy management – is a good tactic.

4.3.2 Top-Down: Breaking Down the Regulator’s Rationality

The value-based rationality, which is more inclusive than the instrumental one taking different perspectives into account, must be broken up into more manageable guidelines in order to meet developers’ needs. One example how that can be achieved is the *privacy by design* movement [16]. Cavoukian argues that “privacy must be incorporated into networked data systems and technologies, by default. Privacy must become integral to organizational priorities, project objectives, design processes, and planning operations. Privacy must be embedded into every standard, protocol and process that touches our lives.” [16]

Privacy by design is a holistic concept including the employment of privacy-friendly information technology (such as the use of privacy-enhancing technology), measures on an organizational level (e.g., management support for privacy) and on a physical level (e.g., access controls). The concept aims to break up the vague philosophical principles into smaller bricks or patterns. Transparency of the data evaluation, for example, is a central cornerstone for privacy-friendly devices. Transparency can be broken down into elements such as defining and articulating which data is being processed, when the data is erased from the server, how users can control the data, or object to it. Thus, similarly to the five pitfalls mentioned above, the privacy by design approach breaks the regulator’s rationality into more manageable modules. Regulators need to split the general data protection or privacy principles into sub-elements, use terminologies employed in other disciplines such as computer science, find the ontologies and taxonomies within them, and provide for more flexible regulations. Once they have been broken up, more tangible sub-rules can be formulated and implemented on a case-by-case basis.

Designing for privacy implies creating a plan or blueprint. When such designs are complex, the term *system engineering* is

employed [68:62]. In order to provide a certain level of abstraction design patterns (or plan patterns) are used. “A design pattern is an abstraction of a design, in the sense that it is not concerned with implementation details.” [68:64]. These patterns generate best-practice rules throughout the entire life cycle of technologies. At the very beginning, the focus lies on privacy requirements and establishing patterns of privacy needs, privacy risks, vulnerability or threat assessments, and selecting privacy controls. Privacy Impact Assessments (PIAs) already provide regulators with helpful guidelines in this respect. As a second step, more technical design patterns can be employed. They relate to potential solution mechanisms, such as anonymization, minimization of data, transparency and informed consent. Different tools or patterns to address such issues technically are available [68]. Herewith, the loop to the engineer’s rationality is closed.

4.4 Call for RoboCode-Ethicists

Both, the bottom-up and top-down approach, call for an alignment and a holistic view on the topic. Simultaneously there is a need for specialization in terms of the training/education of robot scholars. Philosophers, ethicists, legal scholars and social scientists working on the topic should be adequately trained in the technological aspects (especially programming and code), while engineers should possess a basic understanding of the privacy implications and theories currently discussed in the study of information systems in general and robots in particular. Robotics is a complex, interdisciplinary research field. It calls for greater *specialization* and experts among others in the areas of computer science, mechanics, and psychology. Like the need for algorithmists for big data analysis (a term coined by [47]), who act as “reviewers of big-data analysis and predictions” [47:180], robotics needs RoboCode-Ethicists. Such independent individuals or entities could monitor developers work, evaluate the data processing practices of robots, the choice of analytical tools, the bounding between robots and humans, the pervasiveness of data collection, and determine whether privacy implications have been deliberated about before the development of a prototype [69].

Such RoboCode-Ethicists should deal with emerging *value conflicts* that occur with the wider diffusion of robotics. The main question here centers on how to program for value conflicts. The role of RoboCode-Ethicists in this context could be a consulting and expert function, trying to resolve these value conflicts and balance different rationalities. For example, in terms of privacy, such value conflicts arise between users and designers on questions of data retention and re-use. RoboCode-Ethicists should address these conflicts and elaborate recommendations that contribute to a privacy-respecting eco-system, serving the needs of both parties.

5. CONCLUSION

This article addressed the issue of privacy in the context of robots. Our focus lay on autonomous and social robots. We discussed how such robots – as a technology that will likely see massive diffusion in the years to come – might present privacy threats, such as surveillance, access and social meaning [13]. We then added two additional issues that are already discussed in other contexts (IoT, big data, digital culture): the opacity of robotic technology, i.e., the fact that robots will become a taken for granted part of our everyday lives and “melt” into our environments; and the black box problem. The latter describes our unawareness of what robots do and how they function – especially

how the algorithms work that they apply. We then surveyed some of the literature in robot, machine, computer, and information ethics – a field with a long tradition that presents a useful set of concepts to productively approach the privacy implications of robots. In this context, we contrasted two clashing rationalities that vastly correspond with two dominant perspectives in ethics in general and robot ethics in particular as well as Weber’s distinction of means-end rationality vs. value-based rationality: the engineer’s rationality, which largely follows a consequentialist – means-end-rational – approach (“make it work”); and the regulator’s rationality, which largely follows a deontologist – value-based – principle (“respect privacy”). The ensuing tension between these perspectives was illustrated. We then presented a bottom-up and a top-down approach to reconcile the tension. The bottom-up approach takes the robot engineer as the starting point and presents her/him with clear-cut, feasible principles to implement privacy during the development stage [39]. The top-down perspective starts from the regulator’s perspective. However, instead of offering abstract notions, it operationalizes privacy protection with a clear set of implementable rules. Privacy by design [16] served as our example for the top-down approach.

The analysis has several *implications*. First of all, it is a call for *interdisciplinary collaboration* in the research and development of robots. Engineers, legal scholars, sociologists, HRI scholars and philosophers/ethicists should work together to think of privacy-friendly solutions, because privacy is a multi-disciplinary phenomenon and touches on a diverse set of issues in the context of robots. Second, the study points to the necessity of “thinking privacy” *ex ante* and not *ex post*, i.e., addressing the emerging privacy problems as early as possible. Similar to addressing other societal issues, it is arguably on the long term more efficient to address the causes and not only the symptoms of an issue. Third, the article calls for a *holistic view* on the topic and *specialization* in terms of the training/education of robot scholars. Because of the complexity of the research field, *RoboCode-Ethicists*, i.e. experts with knowledge of both the developers and regulators’ rationality are needed.

Our analysis has several *limitations* pointing to fruitful avenues for future research. First, this is a conceptual paper without an empirical contribution. Future research should use qualitative and quantitative methods to explore the privacy implications of social robot technology in depth. A number of studies have provided good efforts [20, 59] but more research is needed to access the topic from a concisely conceptualized perspective. For example, Calo’s [13] three privacy implications could be operationalized within a survey design to assess people’s evaluation of them. Also, qualitative interviews or focus groups with a wide range of experts on the topic could inspire additional insights in how to face the privacy issues as early as possible. Second, given the space constraints, we could not give full justice to previous research. This is especially true for the section on the ethical questions around robots. A rich body of literature from long-standing philosophical traditions has developed around the topic and we had to severely simplify the literature, to an extent that important lines of thought (e.g., virtue-based ethics) could only be tapped into superficially. Future research with more space at hand could expand on the existing ethical frameworks and dig deeper. Third, we did not account for the cultural contingency of robots and privacy. The adoption and acceptance of robots varies strongly depending on the cultural context – and the same is true for understandings of privacy and privacy concerns [1]. Unfortunately, we could not do justice to these nuances. Our

analysis is thus very generalizing and abstract. Future research could analyze the role of robotic privacy (implications) in different contexts, using comparative methodology.

6. REFERENCES

- [1] Aeschlimann, L., Harasgama, R., Kehr, F., Lutz, C., Milanova, V., Müller, S., Strathoff, P., and Tamò, A. 2015. Re-Setting the Stage for Privacy: A Multi-Layered Privacy Interaction Framework and Its Application. In S. Brändli, R. Harasgama, R. Schister, and A. Tamò (eds.), *Mensch und Maschine – Symbiose oder Parasitismus*, Stämpfli, Berne, 1–43.
- [2] Allen, C. and Wallach, W. 2012. Moral Machines: Contradiction in Terms or Abdication of Human Responsibility? In P. Lin, G. Bekey, and K. Abney (eds.), *Robot Ethics: The Ethical and Social Implications of Robotics*, MIT Press, Cambridge (MA), 55–68.
- [3] Anderson, S. L. 2011. Philosophical Concerns with Machine Ethics. In M. Anderson and S. L. Anderson (eds.), *Machine Ethics*, Cambridge University Press, Cambridge (UK), 162–167.
- [4] Anderson, S. L. 2011. The Unacceptability of Asimov’s Three Laws of Robotics as a Basis for Machine Ethics. In M. Anderson and S. L. Anderson (eds.), *Machine Ethics*, Cambridge University Press, Cambridge (UK), 285–296.
- [5] Asscher, L. 2006. ‘Code’ as Law. Using Fuller to Assess Code Rules. In E. Dommering and L. Asscher (eds.), *Coding Regulation: Essays on the Normative Role of Information Technology*, Asser Press, The Hague, 61–90.
- [6] Bekey, G. 2012. Current Trends in Robotics: Technology and Ethics. In P. Lin, G. Bekey, and K. Abney (eds.), *Robot Ethics: The Ethical and Social Implications of Robotics*, MIT Press, Cambridge (MA), 17–34.
- [7] Bennet, C. J. and Raab, C., D. 2006. *The governance of Privacy – Policy Instruments in Global Perspective*. MIT Press 2006, Cambridge (MA).
- [8] boyd, d. and Crawford, K. 2012. Critical Questions for Big Data. *Information, Communication & Society* 15, 5, 662–679.
- [9] Brey, P. 2010. Values in technology and disclosive computer ethics. In L. Floridi (ed.), *The Cambridge Handbook of Information and Computer Ethics*, Cambridge University Press, Cambridge (UK), 41–58.
- [10] Borenstein, J. and Pearson, Y. 2012. Robot Caregivers: Ethical Issues across the Human Lifespan. In P. Lin, G. Bekey, and K. Abney (Eds.), *Robot Ethics: The Ethical and Social Implications of Robotics*, MIT Press, Cambridge (MA), 251–265.
- [11] Brownsword, R. 2008. *Rights, Regulation, and the Technological Revolution*. Oxford University Press, Oxford (UK).
- [12] Bygrave, L. A. (2014). *Data Privacy Law*. Oxford University Press, Oxford (UK).
- [13] Calo, R. 2012. Robots and Privacy. In P. Lin, G. Bekey, and K. Abney (Eds.), *Robot Ethics: The Ethical and Social*

Implications of Robotics, MIT Press, Cambridge (MA), 187–202.

- [14] Calo, R. 2014. Robotics and the New Cyberlaw. *SSRN Electronic Journal*, 101–146. Online: <http://robots.law.miami.edu/2014/wp-content/uploads/2013/06/Calo-Robotics-and-the-New-Cyberlaw.pdf>
- [15] Cate, F. H. and Mayer-Schönberger, V. 2013. Notice and consent in a world of Big Data. *International Data Privacy Law* 3, 2, 67–73.
- [16] Cavoukian, A. 2009. *Privacy by Design – The 7 Foundational Principles*. Information and Privacy Commissioner of Ontario. Online: https://www.iab.org/wp-content/IAB-uploads/2011/03/fred_carter.pdf
- [17] Clarke, R. 2011. Asimov’s Laws of Robotics: Implications for Information Technology. In M. Anderson and S. L. Anderson (eds.), *Machine Ethics*, Cambridge University Press, Cambridge (UK), 254–284.
- [18] Crawford, K. and Schultz, J. 2014. Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms, *Boston College Law Review* 55, 93, 2014.
- [19] Darling, K. 2012. Extending Legal Rights to Social Robots. *SSRN Online Journal*, 1–18. Online: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2044797
- [20] De Graaf, M. M. A., Ben Allouch, S., and Klamer, T. 2015. Sharing a life with Harvey: Exploring the acceptance of and relationship-building with a social robot. *Computers in Human Behavior* 43, 1–14.
- [21] Del Campo, M., Fure, A., McGee, W., Manninger, S., and Flexer, A. 2013. *Autonomous Tectonics – A Research into Emergent Robotics Construction Methods*. In F. Scheurer, J. Nembrini, A. Kilian, and C. Gengnagel (eds.), *Rethinking Prototyping: Proceedings of the Design Modelling Symposium Berlin 2013*, 1–13.
- [22] Denning, T., Matuszek, C., Koscher, K., Smith, J. R., Kohno, T., and Allen, P. G. 2009. A Spotlight on Security and Privacy Risks with Future Household Robots: Attacks and Lessons. in *Proceedings of the 11th International Conference on Ubiquitous Computing (UbiComp '09)*, (Orlando, FL, September 30–October 03 2009), 1–10.
- [23] Doctorow, C. 2014. *Why it is not possible to regulate robots here*. The Guardian Technology Blog, Robots, April 2 2014. <http://www.theguardian.com/technology/blog/2014/apr/02/why-it-is-not-possible-to-regulate-robots>
- [24] EU Eurobarometer (2011). *Attitudes on Data Protection and Electronic Identity in the European Union*. Research Report. Online: http://ec.europa.eu/public_opinion/archives/ebs/ebs_359_en.pdf
- [25] Floridi, L. 1999. Information ethics: on the philosophical foundations of computer ethics. *Ethics and Information Technology* 1, 1, 37–56.
- [26] Floridi, L. 2010. *Information ethics*. In L. Floridi (ed.), *The Cambridge Handbook of Information and Computer Ethics*, Cambridge University Press, Cambridge (UK), 77–100.
- [27] Friedman, B. and Kahn, P. H. 2002. The Ethics of Systems Design. In M. D. Ermann and M. S. Shauf (eds.), *Computers, Ethics, and Society*, Oxford University Press, Oxford (UK), 55–63.
- [28] Friedman B. and Kahn P. 2008. Human Values, Ethics and Design. In J. Jacko and A. Sears (eds.), *The Human Computer Interaction Handbook* (2nd edition), Lawrence Erlbaum Associates, Mahwah (NJ), 1241–1266
- [29] Friedman, B., Kahn, P. H., and Borning, A. 2002. *Value sensitive design: Theory and methods*. University of Washington technical report, 02-12. Online: <http://www.urbanism.org/pub/Research/ResearchPapers/vsd-theory-methods-tr.pdf>
- [30] Fussell, S. R., Kiesler, S., Setlock, L.D., and Yew, V. 2008. How people anthropomorphize robots. In *Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction*, (Amsterdam, The Netherlands, March 12–15 2008), 145–152.
- [31] Fukuoka World Robot Declaration 2004. Online: <http://www.prnewswire.co.uk/news-releases/world-robot-declaration-from-international-robot-fair-2004-organizing-office-154289895.html>
- [32] Garfinkel, S. 2003. *Privacy in a Database Nation*. In M. D. Ermann and M. S. Shauf (eds.), *Computers, Ethics, and Society*, Oxford University Press, Oxford (UK), 137–152.
- [33] Gips, J. 2011. *Towards the Ethical Robot*. In M. Anderson and S. L. Anderson (eds.), *Machine Ethics*, Cambridge University Press, Cambridge (UK), 244–253.
- [34] Hoffmann, C. P., Lutz, C., Ranzini, G., and Meckel, M. 2015. Diversity by Choice: Applying a Social Cognitive Perspective to the Role of Public Service Media in the Digital Age. *International Journal of Communication* 9, 1–20 (in press).
- [35] Hospers, J. 2002. *The Best Action Is One with the Best Consequences*. In M. D. Ermann, M. N. Williams, and M. S. Shauf (eds.), *Computers, Ethics, and Society*. Oxford University Press, Oxford (UK), 3–11.
- [36] International Federation of Robotics 2013. *World Robotics 2013 Report*.
- [37] Klein, H. K. and Kleinman, D. L. 2002. The Social Construction of Technology: Structural Considerations. *Science Technology Human Values* 27, 1, 28–52.
- [38] Langheinrich, M. 2005. *Personal Privacy in Ubiquitous Computing: Tools and System Support*. Dissertation submitted to the Swiss Federal Institute of Technology Zurich, 2005. Online: <http://e-collection.library.ethz.ch/eserv/eth:28011/eth-28011-01.pdf>.
- [39] Lederer, S., Hong, J. I., Key, A. D., and Landay, J. A., 2004. Personal Privacy through Understanding and Action: Five Pitfalls for Designers. *Journal of Personal and Ubiquitous Computing* 8, 6, 440–454.
- [40] Lessig, L. 1999. *Code and Other Laws of Cyberspace*. Basic Books, New York (NY).
- [41] Lessig, L. 2006. *Code version 2.0*. Basic Books, New York (NY).
- [42] Lin, P. 2012. Introduction to Robot Ethics. In P. Lin, K. Abney, and G. A. Bekey (eds.), *Robot Ethics: The Ethical*

- and *Social Implications of Robotics*, MIT Press, Cambridge (MA), 3–16.
- [43] Lin, P., Abney, K., and Bekey, G. 2011. Robot ethics: Mapping the issues for a mechanized world. *Artificial Intelligence* 175, 5-6, 942–949.
- [44] Lutz, C. and Strathoff, P. 2013. Privacy Concerns and Online Behavior – Not so Paradoxical After All? Viewing the Privacy Paradox through Different Theoretical Lenses. In S. Brändli, R. Schister, and A. Tamò (eds.), *Multinationale Unternehmen und Institutionen im Wandel – Herausforderungen für Wirtschaft, Recht und Gesellschaft*, Stämpfli, Berne, 81–99.
- [45] MacKenzie, D. and Wajcman, J. 1999. *The Social Shaping of Technology*, Open University Press, Buckingham (UK).
- [46] Margulis, S. T. 2011. Three Theories of Privacy – An Overview. In S. Trepte and L. Reinecke, L. (eds.), *Privacy Online – Perspectives on Privacy and Self-Disclosure in the Social Web*, Springer, Berlin/Heidelberg, 9–18.
- [47] Mayer-Schönberger, V. and Cukier, K. 2013. *Big Data: A Revolution That Will Transform How We Live, Work and Think*. HMH Books, Boston (MA) & New York (NY).
- [48] Moor, J. H. 2011. The nature, importance, and difficulty of machine ethics. In M. Anderson and S. L. Anderson (eds.), *Machine Ethics*, Cambridge University Press, Cambridge (UK), 13–20.
- [49] Nissenbaum, H. 2004. Privacy as Contextual Integrity. *Washington Law Review* 79, 101–139.
- [50] Omohundro, S. 2014. Autonomous technology and the greater human good. *Journal of Experimental & Theoretical Artificial Intelligence* 26, 3, 303–315.
- [51] Oosterlaken, I. 2009. Design for development: A capability approach. *Design Issues* 25, 4, 91–102.
- [52] PEW Research 2014. *Internet of things Report*. Online: http://www.pewinternet.org/files/2014/05/PIP_Internet-of-things_0514142.pdf
- [53] PEW Research 2014. *AI, Robotics, and the Future of Jobs*. Online: <http://www.pewinternet.org/files/2014/08/Future-of-AI-Robotics-and-Jobs.pdf>
- [54] PEW Research 2014. *Public Perceptions of Privacy and Security in the Post-Snowden Era*. Online: http://www.pewinternet.org/files/2014/11/PI_PublicPerceptionsOfPrivacy_111214.pdf
- [55] Pinch, T. J. and Bijker, W. E. 1984. The social construction of facts and artefacts: or how the sociology of science and the sociology of technology might benefit each other. *Social Studies of Science* 14, 3, 399–441.
- [56] Pfleeger, C. P. and Pfleeger, S. H. 2007. *Security in Computing* (4th edition). Prentice Hall, New York (NY).
- [57] Reidenberg, J.R. 1998. Lex Informatica: The Formulation of Information Policy Rules Through Technology. *Texas Law Review* 76, 3, 554–593.
- [58] Riek, L. D., and Howard, D. 2014. A Code of Ethics for the Human-Robot Interaction Profession. In *We Robot Conference*, (Coral Gables, FL, April 04-05 2014), 1–10.
- [59] Rosenthal-von der Pütten, A. M. and Krämer, N. 2014. How design characteristics of robots determine evaluation and uncanny valley related responses. *Computers in Human Behavior* 36, 422–439.
- [60] Scheutz, M. 2012. The Inherent Dangers of Unidirectional Emotional Bonds between Humans and Social Robots. In P. Lin, G. Bekey and K. Abney (eds.), *Robot Ethics: The Ethical and Social Implications of Robotics*, MIT Press, Cambridge (MA), 205–222.
- [61] Scheutz, M., Crowell, C. 2007. The Burden of Embodied Autonomy: Some Reflections on the Social and Ethical Implications of Autonomous Robots. In *Workshop on Roboethics at the 2007 IEEE International Conference on Robotics and Automation (IRCA)*, (Rome, Italy, 10-14 April 2007), 1–7. <http://www.roboethics.org/icra2007/contributions/SCHEUTZ%20The%20Burden%20of%20Embodied%20Autonomy.pdf>
- [62] Singer, P. 2009. *Wired for War*. Penguin Press, New York (NY).
- [63] Smith, H. J., Dinev, T., and Xu, H. 2011. Information privacy research: an interdisciplinary review. *MIS Quarterly* 35, 4, 989–1016.
- [64] Solove, D. J. 2008. *Understanding Privacy*. Harvard University Press, Cambridge (MA).
- [65] Solove, D. J. 2012. Introduction: Privacy Self-Management and the Consent Dilemma. *Harvard Law Review*, 126, 1880–1881.
- [66] Tufekci, Z. 2008. *Can you see me now? Audience and disclosure regulation in online social network sites*. *Bulletin of Science, Technology & Society* 28, 1, 20–36.
- [67] Turkle, S. 2011. Authenticity in the age of digital companions. In M. Anderson and S. L. Anderson (eds.), *Machine Ethics*, Cambridge University Press, Cambridge (UK), 62–76.
- [68] Van Rest, J., Boonstra, D., Everts, M., van Rijn, M., van Paassen, R. 2014. Designing Privacy-by-Design. In B. Preneel and D. Ikonoumou (eds.), *Privacy Technologies and Policy*, Springer, Berlin & Heidelberg, 55–72.
- [69] Veruggio, G. 2007. *EURON Robotics Roadmap*. Online: http://www.roboethics.org/index_file/Roboethics%20Roadmap%20Rel.1.2.pdf
- [70] Watson, S. 2014. If Customers Knew How You Use Their Data, Would They Call It Creepy? *Harvard Business Review*, April 29 2014. Online: <https://hbr.org/2014/04/if-customers-knew-how-you-use-their-data-would-they-call-it-creepy/>.
- [71] Weber, M. 1978. *Economy and Society*. University of California Press, Oakland (CA).
- [72] Weiser, M. 1991. *The computer for the 21st century*. *Scientific American* 265, 3, 66–75. Reprinted in *IEEE Pervasive Computing* 1, 1, Jan.-Mar. 2002, 19–25.
- [73] Winner, L. 1977. *Autonomous Technology: Technics-out-of-Control as a Theme in Political Thought*, MIT Press, Cambridge (MA).