

Document Version

Final published version

Licence

CC BY-NC-ND

Citation (APA)

Allhoff, F., & Henschke, A. (2018). The Internet of Things: Foundational ethical issues. *Internet of Things (Netherlands)*, 1-2, 55-66. <https://doi.org/10.1016/j.iot.2018.08.005>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



The Internet of Things: Foundational ethical issues

Fritz Allhoff^{a,*}, Adam Henschke^{b,c}

^a Department of Philosophy, Western Michigan University, United States

^b National Security College, Australian National University, Australia

^c Faculty of Technology, Policy and Management, Delft University of Technology, The Netherlands

ARTICLE INFO

Article history:

Received 15 August 2018

Accepted 18 August 2018

Available online 22 August 2018

Keywords:

Internet of Things

Ethics

Informed consent

Privacy

Information security

Physical safety

Trust

ABSTRACT

This paper surveys foundational ethical issues that attach to the Internet of Things (IoT). In [Section 1](#), we provide an overview of the technology, indicating both current and future applications. Subsequent sections consider particular ethical issues, including: informed consent ([Section 2](#)), privacy ([Section 3](#)), information security ([Section 4](#)), physical safety ([Section 5](#)), and trust ([Section 6](#)). [Section 7](#) emphasizes that these ethical issues do not exist in isolation, but converge and intersect in myriad ways. And that these issues are not comprehensive, but rather are foundational starting points that stand to be expanded and further elucidated through future research.

© 2018 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license.

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

1. Internet of Things

The Internet of Things (IoT) is poised to be the next step in the information revolution, portending societal change that will rival that of the internet itself.¹ The size of the IoT is expected to be immense: by 2020, 20–50 billion things are estimated to be connected as part of the IoT [64], leading some to predict an investment of US\$1.7 trillion by 2020 ([16]). It thus has the attention of companies, governments, and citizens worldwide, giving rise to research in industry and academia alike [50,85].

The IoT refers to a complex network of interactive and technical components clustered around three key elements: sensors [51], informational processors [2,105], and actuators [90]. It is this “ability [of] objects to communicate that delivers the power of the IoT” ([30], p. 20). The IoT’s networked communication will radically alter the way that we interact with technologies, particularly as our physical relationships to those technologies transform. For example, the IoT allows for increased automation or action-at-a-distance. But the communicative capacities enabled by the IoT are not simply for the sake of themselves. Rather they are meant to achieve some useful objective, specifically through intervention in the physical world. To illustrate what that means in practice, consider that:

a sensing and actuation utility will not only exist in public spaces, but also extend into the home, apartments, and condominiums. Here, people will be able to run health, energy, security, and entertainment apps on the infrastructure. Installing and running new apps will be as easy as plugging in a new toaster into the electric utility. One app may help monitor and control heart rate, another may perform financial and investments services, and another automatically

* Corresponding author.

E-mail address: fritz.allhoff@wmich.edu (F. Allhoff).

¹ This section draws in part from Henschke [38].

ordering food and wine, or even predicting an impending medical problem that should be addressed early to mitigate or even avoid the problem [90], p. 4).

In this regard, the IoT's informational and communicative functions have direct physical impacts.² Worryingly, this duality manifests novel pathways for malicious attacks. Indeed, a physical target accessible through cyberspace is more preferable for attackers than one that must be physically accessed, since the attacker would not need to incur the expense or risk of physical access [43]. As we discuss in Sections 4 and 5, this informational insecurity has implications for physical safety as well.

While it seems uncontroversial that we should follow the U.S. Department of Commerce's recommendation "to support a stable, secure and trustworthy IoT environment", the word "ethics" is conspicuously absent from their analysis [67]. This is especially noteworthy given the general consensus that the security of the IoT is worryingly underdeveloped [78]. In other words, we simultaneously have an informationally insecure IoT, without a corresponding emphasis on what that means from an ethical perspective. Unfortunately, this is not a terribly novel situation in which to find ourselves with regards to emerging technologies: the ethics *often* lags behind the technological innovation [3]. And that is true even if "ethics" is construed broadly, comprising not just what might be the philosophical dimensions, but also the policy and legal components; call this approach more broadly "normative" if "ethics" seems too narrow.

To be sure, there has been significant work in the ethics of emerging technologies more generally [3,15]. But this is too broad as it does little to identify the novel features of the IoT in particular. And much of that work just attaches to fairly abstract discussions, say as relates to the precautionary principle; this principle invites us to consider broad targets, like risk and uncertainty, without a particular operational context [3]. So, while acknowledging that things like risk and uncertainty matter, it is hard to gain traction on some specific domain unless we take a deeper dive. There has also been substantial work on the ethics of nanotechnology [6–8], which provides a useful model for thinking about the IoT—insofar as it is a fairly new, technologically-empowered discipline. But, as we will see, the IoT raises a host of its own issues and therefore deserves a dedicated forum.

And so this paper aims to articulate what those issues are, as well as some ways to begin to think about them. As with all futures-related work, a healthy dose of humility is needed: some of the details will be wrong, and the technological evolution will meander in ways that had not been predicated. But even if we try to separate what we know from what we don't know, the latter comes in two important flavors, "known unknowns" and "unknown unknowns":

[A]s we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns—the ones we don't know we don't know. And if one looks throughout the history of our country and other free countries, it is the latter category that tend to be the difficult ones [24].

However, the existence of "unknown unknowns" need not paralyze us, either. Rather, we can focus on near- and mid-term issues, which are more predictively tractable than long-term issues. And many of these do not have to be uncritically hypothesized: they are already in front of us. In the remainder of this paper, we will therefore try to survey those that are most readily applicable to the IoT. It is important to note from the outset, though, that these issues do not separate as neatly as would be convenient: they intersect with each other in all sorts of ways, ranging from the simple to the complicated.

The issues that we discuss in the ensuing sections were chosen for two principal reasons. First, we consider them to foundational in the sense that they underlie many other ethical issues that are likely to develop. Second, they have been explored in other contexts, and so there is a wealth of other thinking on which to draw for our analyses. For example, informed consent (Section 2) and privacy (Section 3) have a long history in medical ethics, as well as other domains that trade heavily on shared data. Informational security (Section 4) and safety (Section 5), by contrast, are less at home in the medical ethics literature, but have received extended treatments in literatures on internet ethics and robotics. Trust (Section 6) frequently appears in literatures relating to governance, regulation, and oversight. With regards to organization, we have structured the discussion to start with informational concerns, particularly those derived from sensors and data gathering (cf., informed consent and privacy), to then move to communications (cf., informational security), actuators (cf., physical safety) and decision-making (cf., trust).

Finally, we want to emphasize that this approach is meant to be pluralistic and open. By "pluralistic", we mean that there are a range of ethical issues relevant to the IoT. Furthermore, these issues should not be taken in isolation. While we have organized the paper around the aforementioned five issues, we want to stress that they overlap and intersect in myriad ways. Trust, for example, is as relevant to privacy as privacy is to trust. By "open", we mean two complementary things. First, we are open to the idea that this ethical inquiry can—and indeed should—be carved up differently by different people pursuing different purposes or explanatory priorities [37], pp. 174–75.³ Second, we intend for this paper to *begin*—or, in some regards, to *continue*—a critical exploration of the IoT's ethical dimensions. This will hardly be the last word, nor

² Moreover, there is a tension between the informational and physical "spaces". On the one hand, informational accessibility is important to *reduce* physical risk—if the device has problems or needs updates, physical safety could be promoted through, for example, remote access. However, informational security is important to *increase* physical safety—in order to prevent malicious hacking, informational access needs to be restricted and controlled. See Henschke [38] for more discussion.

³ For other approaches, see, Popescu and Georgescu [71], AboBakr and Azer [1], and [38].

should it be. To the contrary, we embrace the sentiment that the “*first step* is to pursue the discussions, studies, task forces, commissions, and pilots that will help develop governance for an empowering and enabling IoT” ([13], p. 7); emphasis added). And so the discussion offered in the following sections is meant to be a first step in this regard, fully aware that many more steps both should and will ensue.

2. Informed consent

In this section, we will discuss informed consent directly, though it intersects with the discussions in the other sections quite extensively. To motivate the discussion, consider the case in 2016 where a smart sex toy was discovered to be collecting and communicating user data to the product’s maker. A U.S. woman sued the company upon learning that the maker was “collecting information about her and other users’ preferred vibration settings, the dates and times the device is used, [and] the email addresses of [device] owners who had registered their devices . . . [obtaining] all this data without the permission of its users” [75]. This case demonstrates how the IoT intersects with existing moral concepts like informed consent: a device gathers personal information through sensors and communicates that information to some receiver, all without the user necessarily being actively involved. While this case trades on highly personal information, even more innocuous information—like times and dates of use—are nevertheless both commonly transmitted and morally significant.

Informed consent has been most robustly considered in the context of medical ethics; seminal work has come in the past few decades, but the concept goes back to Ancient medicine and was substantially developed in the 19th and 20th centuries [28,46]. Much of that work is more technical than we need for present purposes, but can be distilled into the following analysis. Specifically, the American Medical Association suggests that “[i]nformed consent . . . is fundamental in both ethics and law” and that stakeholders “have the right to receive information and ask questions . . . so that they can make well-considered opinions” [10], Section 2.1.1].

While it is common to conflate “consent” with “informed consent”, the latter means something more robust. Specifically, “consent” should not mean simply “assent”—e.g., merely an affirmative response. The reason is that stakeholders need to know what they are assenting to, and the “informed” part adds that component. Informational asymmetry presents an obstacle here: for example, say that the stakeholder does not know—nor would even be able to understand—the relevant pharmacological mechanisms behind how some drug works. But it would be specious to think that informed consent requires a thoroughgoing understanding of, say, the relevant molecular interactions. Rather, informed consent requires something short of that, like knowing that some drug may cause nausea—exactly *why* the drug causes nausea is not necessarily so important for a decision about whether that treatment is, all things considered, advisable. Or, to the extent that it is—e.g., how it affects a patient’s pre-existing conditions or sensitivities—it *would* be relevant for informed consent. The point is simply that informed consent is generally an intermediate threshold between mere assent and technical expertise.

Informed consent has been extended from medicine to digital contexts, starting in the early 2000s [27]. This extension has translated from the patient or research subject to the digital consumer. The principal application is vis-à-vis personal data, often as would be collected by or shared with a third party, such as a website or smartphone application [68]. The applications to the IoT are initially straightforward insofar as any IoT application will similarly seek informed consent from the consumer. Therefore, there needs to be a standard that establishes informed consent for the acquisition, retention, and sharing of personal data. For example, the European Union has recently implemented a definition of consent as: “any freely given, specific, informed and unambiguous indication of the data subject’s wishes by which he or she, by a statement or by a clear affirmative action, signifies agreement to the processing of personal data relating to him or her” [26].

The challenges start when we consider how a standard like this intersect with those whose data is at stake. The standard solution is an “end user license agreement” (EULA), which are meant to specify exactly how this all works. Two immediate problems arise with EULAs. On the one hand, they tend to be too detailed, often running several pages replete with technical and legal jargon that is inaccessible to many users. Furthermore, the sheer length and density of these documents virtually ensures that users do not read them, instead clicking anything to make them go away. On the other hand, they tend to be not detailed enough, comprising fairly generic boilerplate. For example, “they are the same regardless of the type of user, they do not consider different contexts (e.g., office or home), and they address . . . geo-political boundaries [in a limited way]” [68], p. 789].

Another issue with EULAs is that users often lack meaningful alternatives to click-through practice. Imagine, for example, that a user accesses a news site that says something like “by using this site, you accept the terms of service.” But insofar as most users don’t even know what those terms of service are—e.g., because they can’t be bothered to click on some document and read it—it is hard to seriously interpret their continued use of the site as manifesting informed consent. And given, for example, the implementation of GDPR—which triggers click-to-accept on every newly-visited website—it is unreasonable to expect that users are doing full diligence. And so we can separate the question as to whether, in principle, EULAs are a useful way to track informed consent with the manifest practice that they are not read.

Remember also that, from the medical context, informed consent evolved as a standard under which patients could engage in a conversation with their providers in a clinical or medical context as to available treatments, risks and benefits, and so on. A EULA lacks any of these interactive features. Rather, it is quite literally a “click here to advance”, completely bereft of any interactive context or effective recourse. Say a user had a legitimate question about a EULA, then what? If the service provider is iTunes, is someone supposed to simply dial up Apple in Cupertino? It is extraordinarily unlikely that, in any reasonable timeframe, a consumer would be able to elucidate the salient contractual terms or bargain on any of them.

In fact, these sorts of concerns threaten to undermine the legitimacy of EULAs altogether. In contract law, for example, we speak of “contracts of adhesion”, which may be unenforceable. These are ubiquitous in commercial practice, but are unconscionable in certain cases, because of their propensity to lead to “unfair surprise”; principally because they are rarely read, and so users do not know what they include ([42], p. 28]; citing [9], Section 2-302] cmt. 1). For example, in *Step-Saver* [93]—a seminal EULA case—the court held that a vendor could not enforce a box-top license under which simply opening software packaging rendered the consumer liable to contractual terms printed on the box.⁴ But other courts have held the opposite, leaving this an open, context-dependent inquiry.⁵

As a way to deal with these issues, various commentators have proposed regulatory frameworks [11,70,102]. Interestingly, this might come full circle to the medical context discussed from the outset as the sensitivity of personal medical data intersects with IoT applications, like Fitbits, Apple watches, and so on. Moreover, as one of us has argued elsewhere [37], the aggregation of personal information enabled by convergent technologies like those composing the IoT creates virtual identities that can become as revealing as the sorts of information gathered and produced in a medical context. And to the extent that IoT devices capture and produce this sort of personal information, existing medical data regulation may need to transpose or be adapted to the IoT context [11]. Nevertheless, an ongoing obstacle is likely to be that boilerplate EULAs are comparatively cheap and easy to deploy as against many of the alternatives, either of which is a welcome outcome for industry. Regulation will impose further burdens on industry, many of which may be then passed on to consumers. And so the “informed consent gold standard”—whatever that may be—will need to be weighed against the increased costs that such a gold standard introduces.

3. Privacy

As a philosophical concept, “privacy” can refer to a suite of related, but separable, concepts [66], tracking both descriptive and normative components ([37], pp. 28–55). Aristotle’s account contrasted the public sphere of politics and political activity (*πóλις*) with the private sphere of the family (*οἶκος*) ([12], *Politics* 1253b1–14, *Politics* 1260b8–27, *Eudemian Ethics* 1242a40–b2; [83]). John Locke approached privacy through the lens of property rights, contrasting property held in common with that acquired through the mixing of one’s labor [55], Section 5]. John Stuart Mill’s thinking traded on the distinction between the appropriate extent of government authority, as against the realm of self-regulation; for him, the regulative power of the government lacked application in the private domain [23,63].⁶

The intersection between privacy to the IoT starts with the observation that devices connected to the IoT collect vast amounts of user data and that data can be analyzed, shared, and so on [103].⁷ For example, in a widely publicized case, Target mined a client’s purchasing habits, predicted that she was pregnant, and send a mailer promoting baby items to her home. As it turns out, she was still in high school and, while she was in fact pregnant, her family did not know; they literally found out because of the mailer. As one of Target’s statisticians reported:

[Target] ran test after test, analyzing the data, and before long some useful patterns emerged. Lotions, for example. Lots of people buy lotion, but [an analyst] noticed that women on the baby registry were buying larger quantities of unscented lotion around the beginning of their second trimester. Another analyst noted that sometime in the first 20 weeks, pregnant women loaded up on supplements like calcium, magnesium and zinc. Many shoppers purchase soap and cotton balls, but when someone suddenly starts buying lots of scent-free soap and extra-bag bags of cotton balls, in addition to hand sanitizers and washcloths, it signals they could be getting close to their delivery date [39].

So here the issues are at least two-fold. First, Target was collecting significant user data, likely tied to the shopper’s credit card or rewards number. And then it used this data to profile her, correctly predicting that she was pregnant, based on her purchasing. There could well be ways to opt out of this sort of profiling (e.g., paying cash), but this would simply never occur to most consumers. Second, Target then collated the information in such a way and sent targeted mailings—to

⁴ As we are writing for an international audience, the point in mentioning U.S. cases is simply to map conceptual possibilities, not to drill down into a single country’s jurisprudence.

⁵ See, for example, *ProCD, Inc. v. Zeidenberg*, 86 F.3d 1447 (7th Cir. [73]). See Hillman and Rachlinski [42] for a comprehensive survey of case law.

⁶ In addition to the philosophical approaches, courts have ruled on the legal status of privacy as well. These rulings tend to be too jurisdictionally-dependent to survey easily, but we will briefly mention the U.S. context—which, if not necessarily representative, is at least informative. *Griswold v. Connecticut* [35] affirmed the right of married couples to possess oral contraception, but more generally discussed the right of privacy as pertaining to marriage and sexual relations of married couples. *Griswold* was soon extended to allow for interracial marriage, possession of obscene materials at home, and possession of contraception by non-married individuals as well [23]. The high-water mark for this analysis came in *Roe v. Wade* [80], in which the Court upheld a right to abortion. In that holding, the Court ruled that, while there was no explicit assignment of a privacy right in the U.S. constitution, such a right nevertheless could be derived from the “penumbra” (i.e., shadow) of other explicitly granted rights which, taken in aggregate, can be used to infer a protected, private sphere.

Article 8 of the European Convention on Human Rights also provides “a right to respect for private and family life” and continues that this right shall not be interfered with except “such as in accordance with the law and is necessary in a democratic society in the interests of national security, public safety or the economic well-being of the country” Relevant to present purposes, the European Court of Human Rights held in *Rotaru v. Romania* [81] that public information collected by the state is protected by Article 8. The European right is generally broader than the American one; for example, *Pretty v. United Kingdom* [72] established a right to die that is not universal across American jurisdictions—and does not exist at all federally. The European so-called “right to be forgotten” does also not exist in the United States. See, for example, *Google v. Costeja* [34]. For more general discussion, see Rosen [82].

⁷ For more detailed discussion, see Henschke [37], pp. 126–98].

a minor no less—that were noticed by her parent. Certainly profiling can well be good business practice, but it can still be invasive—or, at a minimum, relevantly intersect with the informed consent discussion offered in Section 2.

And so one worry is what sorts of information becomes available, and to whom, once we integrate our lives more thoroughly with the IoT. Simply encrypting the information is not likely to solve all the problems, either. As discussed above, a privacy suit was lodged against the company that produced the IoT-connected smart sex toy [48]. We can imagine such data being used to enhance a user's experience, but it could just as easily be used to provide detailed demographic information that could go to marketing or product development, neither of which might be the user's preference.⁸ Another worry is that even if the data is encrypted, the *metadata* can nevertheless reveal pertinent details. For example, consider the case of General David Petreus, former director of the U.S. Central Intelligence Agency. In the summer of 2012, General Petreus was having an affair with his biographer, Paula Broadwell. The two communicated through a Gmail account, by leaving messages in a shared draft folder. The point of doing this was to create less data than actually sending and receiving messages; in fact, this technique has been used by terrorists for exactly this reason [29]. But investigators were able to match metadata (e.g., IP addresses, time stamps) to information they already knew about Broadwell (e.g., which hotel she was staying in on a given night) to connect the dots [32].

Once we start talking about fully-integrated smart homes, these issues compound. Consider something like Amazon Echo which, by design, is always listening—how else would it know when you said “Alexa” [22]? And, therefore, it is listening whether someone is using the product or simply talking to their family. Where that information goes and how it is recorded, processed, and stored raise important ethical questions. Smart homes are routinely communicating information back to manufacturers, not all of which is even encrypted [40]. But, per above, even if it were, that hardly solves all the issues. For example, suppose an IoT coffee machine fires up twice on Friday morning, when it otherwise only fires up once. Suppose the EvaDrop smart shower also displays more activity than usual. Did our resident meet someone at a bar the night before? What if it happens every Friday—does that reveal even more about the resident's personal life or practices? The point here is that seemingly innocuous information, like the time a coffee machine is turned on, can become highly revealing when integrated and aggregated with other information like shower duration.⁹ Not only could hackers potentially access this information, but the manufacturers themselves would start to have more information about their customer's habits, and the personal information that is derived from their habits can reveal deeply intimate aspects of a person's life like being pregnant.

There have started to be trends to protect consumers' privacy, but much of the IoT is still unregulated. The European Union's GDPR just went into effect in mid-2018 and makes substantial strides toward increased data protection. Courts have also started to recognize rights to digital privacy. For example, in *Carpenter v. United States* [19], law enforcement had used metadata from cell phone towers to place a suspect at the location of several bank robberies. But the Court held that the Fourth Amendment—which protects against unreasonable searches and seizures—proscribed this practice without the prior acquisition of a warrant. And so, moving forward, regulatory or legal regimes may further secure rights to digital privacy, but, at least at present, risks and pitfalls abound [102]. The overall point to recognize here is that the IoT, with its wealth of sensors and integrated communications, creates the perfect opportunity for gathering incredible amounts of personal information; that information can significantly bear on an individual's privacy.

4. Information security

Many devices connected to the IoT have limited or no effective information security [20,33,98,104]. Failure to ensure effective information security has important ramifications for privacy and physical safety, and can effectively invalidate informed consent and undermine trust. In other words, informational security plays a key role, not just as an intrinsic value, but also instrumentally insofar as it helps secure other moral values as well. However, in order to meet the promises of a free and frictionless, fully-integrated, and user-friendly IoT, informational flow between various components is also important. The challenge here is not simply to identify the ethical risk, but also to recognize that there are likely to be tradeoffs between important ethical values.

The underlying problem with information security is that many IoT components have sensors coupled to communicators. For example, a camera, microphone or other sensor picks up data from the environment and is coupled communicator that relays that data to a remote location, like the cloud or some proprietary server. It is not so much that smartphones with cameras cause concern merely in virtue of this functionality. Rather, the worry is that that the smartphones are potentially in communication with remote and opaque data receivers. And these can process the data and potentially distribute the resulting information without the user's knowledge or consent [31]. There have been a series of cases already where smart televisions and other smart devices have been sent data picked up by cameras and microphones from people's homes back to the producer's servers for analysis, without clear or obvious consent from the users [22,61].

The potential violations of privacy and informed consent should be obvious here. However, there are also implications for physical safety (see Section 5 below), as well as wider security implications of limited information security. Not only can remote operators potentially observe users in their homes or other private spaces, but the lack of information security in

⁸ Again, informed consent will also figure in centrally here—emphasizing the point that these foundational ethical issues are interdependent and cross cutting.

⁹ For a sustained development and presentation of this argument, see Henschke [37].

the communications themselves has led to IoT-connected devices being used as part of distributed denial of service (DDoS) attacks. In 2016, for example, IoT devices like video recorders and smart refrigerators were used in DDoS attacks that brought down the website of cybersecurity expert Bryan Krebs and internet services on the east coast of the U.S. [49,97].

To further highlight that this is not simply an issue of privacy or informed consent, certain IoT devices have implications for national security. For example, a wearable device connected to the internet and uploaded data to a publicly-accessible website:

Strava, a fitness-tracking app, is revealing potentially sensitive information about military bases and supply routes via its global heatmap website. The data map shows 1 billion activities and 3 trillion points of latitude and longitude from Strava's global network of athletes Using satellite imagery, you can see base buildings, for example. But on the heatmap, you can see which buildings are most used, or the jogging routes of soldiers ([14]; internal quotation marks omitted).

In this case, we are quite literally shown how a network of informationally insecure IoT devices place individuals at physical risk, as well as relay sensitive national security information—in this case, that of soldiers and site infrastructure.

A large part of the problem with IoT security is that, while many of the components have some security features, such as passwords to limit access to users, those passwords are typically set to a factory default and never changed [45]. Furthermore, the default passwords are widely available online [41]. For these reasons, someone—whether curious or malicious—could hack into these devices, effecting anything from privacy violations all the way up to national security risks. So why are these devices left “open to the world”? In other words, if this lack of information security presents such risks, then why are the devices and components designed and released with such informationally porous features? One answer is cost: adding more comprehensive security features increases the cost of the individual components [47]. And in a marketplace and consumer setting unburdened with government regulation and oversight, the cheapest cost point will be widely pursued.¹⁰ There is, though, a role for governments to set and enforce *minimum* information security standards for a product to be available in the market [11,70,102].

This would require us to first identify what the minimum set of necessary information security features are. This is a complex undertaking. For instance, the Strava application mentioned above is only a national security issue when used by military and intelligence employees in a potential conflict zone. Does it therefore make sense to have stringent information security for all users of the technology, simply because some of those users might, sometimes, be in conflict zones? Minimum information security standards are likely to be variable, depending in part on the user and the context of use. The role of an informed and active public debate here is essential both to recognize what the minimum levels ought to be and to enable a broad range of stakeholders to articulate how their particular circumstances bear on those minimum levels.

A second issue concerns the responsibility of the user. For example, if changing passwords were simple enough, we might think this is a core user responsibility [101]. In this regard, it would be analogous to seat belts in cars—it is the driver's responsibility to put them on. But that assumes the changes are simple; in other words, that users have the requisite time and cyberliteracy to manage secure passwords. In terms of information security, it would be ideal for every device to have a 32-character password requiring upper and lower case, numbers and symbols, that needs to be changed daily. But this is patently unreasonable in terms of human capacities and workflow; security demands must be balanced against human capacities and workflow in order to achieve functional benefits.

A third issue related to usability. The more information security limits the ease of communication between devices, the less effective interoperability will be. Ultimately, heightened information security can lead to deleterious impacts on the very thing that the IoT is supposed to provide: seamless and invisible integration into our working and personal lives [47]. And so security cuts against some aspirational elements of the IoT, namely effective and easy use. In this regard, it—like much else—requires balancing of competing values. To put it a different way, a commitment to information security runs the risks of systems that are stove-piped, unconnected, and disintegrated (i.e., because the converses carry security risks). This sets the scene for particular producers to create information ecosystems that are closed to other company's products. Rather than a collective, global, IoT, we would have a balkanized, localized “splinternet of things”. The end result of simply pursuing constraints on communications, even if motivated by the sound concerns about information security, is that we lose the interoperability and ease of use. This is not to caricature a binary decision—e.g., complete informational openness versus an absolute lock down on communications—since there are surely settings that reduce many of the security concerns while maximizing ease of use and interoperability. However, as we have indicated, operationalizing this framework is non-trivial.

5. Physical safety

As mentioned in Section 1, one principal difference between the IoT and the “traditional internet” is that the IoT has the potential to be active in the physical realm. In addition to sensors and communication, many elements of the IoT include

¹⁰ We note here that Microsoft has recently publicly called for government regulation around facial recognition technologies, arguing that “[w]hile we appreciate that some people today are calling for tech companies to make these decisions—and we recognize a clear need for our own exercise of responsibility, as discussed further below—we believe this is an inadequate substitute for decision making by the public and its representatives in a democratic republic. We live in a nation of laws, and the government needs to play an important role in regulating facial recognition technology. As a general principle, it seems more sensible to ask an elected government to regulate companies than to ask unelected companies to regulate such a government” [88].

actuators; i.e., components that move or change in the physical realm. For example, consider a smart home with a door lock that is activated when the user—or, more likely, the user's smartphone—is within five meters of the door. In this example, there are sensors (e.g., one in the smartphone and one in the door's locking mechanism) and communicators (e.g., the smartphone and the door's locking mechanism confirming handshakes). Critically, though, the sensors and communications only serve their purpose when they cause some physical change in the world, through the actuator.

As another example, consider a person in a smart driverless vehicle that steers, accelerates, brakes, and so on in response to other vehicles (i.e., driverless and driven), traffic lights, and other road users like cyclists and pedestrians. In contrast to an internet largely constrained to cyberspace, this technology has physical presence and capacities, some of which might threaten physical safety. Should the driverless vehicle malfunction, its passengers, those in other vehicles, and the aforementioned cyclists and pedestrians may all be at risk [76,95]. Again, this draws a significant distinction between the IoT and cyberspace, which, as a virtual space, does not pose risks of direct physical harm.¹¹

The immediate ethical importance of this is that we have a responsibility mitigate or avoid physical harms. “We have learned by now that new technologies, first and foremost, need to be safe” [52], p. 7]. Following existing literatures in both robot ethics and autonomous vehicles, we can expect the IoT to be significantly concerned with physical safety moving forward [53,54]. “Safety in IoT means being able to reason about the behavior of IoT devices, especially actuators, and being able to detect and prevent unintended or unexpected behavior” [2], p. 41]. Any components with actuators that pose a physical risk must take those risks seriously. While this is but a moral platitude, things become interesting when we seek detail about what processes are required to *ensure* physical safety and what processes are needed to *assure* the users and public at large that the physical components are safe [79].

A key mechanism here is government oversight and enforcement of minimum safety standards. Given the range of applications and contexts of use, consider cars and children's toys here as examples of the ways that such minimum safety standards would be developed. For cars, there are range of minimum features that a car must meet in order for it to be commercially viable. Given the physical risks posed to the driver, passengers and other road users, not only are those standards required at the time of sale, but regular car safety checks are required to ensure that things like tires, brakes, windshields, and so on meet minimum standards for road-worthiness. Children's toys are typically expected to be reasonably safe for children of a given age range to play with: if the toy has parts that pose a choking hazard to an infant or young child, then it is not safe for that child. Toys also typically undergo testing to establish that, for example, they are non-toxic and inflammable. The parameters of that testing understandably vary based on the specific products and context of use, including the users themselves (e.g., children) [60]. This same principle should apply to IoT devices as well: an IoT kitchen blender with sharp blades for instance, would likely require different safety features than a piece of IoT-enabled sporting equipment. And this applies to non-physical risks as well: an IoT-enabled sex toy should likely have greater constraints on what information is transmitted to the manufacturer than a smart blender because of the relative sensitivities of that information.

Before moving on, though, let us consider some objections to what otherwise sounds like a simple proposal. First, the claims about government oversight and enforcement of regulations assume that these roles fall within the government's purview; not everyone finds this attractive. For example, libertarians espouse limited governmental intervention and suppose individuals are well-suited to act in accordance with their own risk tolerances [69,100]. But it bears emphasis that even industry partners may support some role for the government as against a completely *laissez faire* free-market model. For example, Microsoft is actively calling for government engagement on controversial technology like facial recognition technology [88], and Facebook's CEO Mark Zuckerberg has not ruled out an appropriate role for government regulation following the Cambridge Analytica scandal [92].

Second, any enforcement mechanisms—particularly ones that seek to apply civil or criminal liability as a result of physical harms—generally need to be able to correctly identify the responsible parties [96].¹² However, with the IoT, the causal networks are complex, and determinations of liability can be quite complex. For example, suppose multiple parties bear a causal relationship to some injury, such that their respective causal contributions have to be partitioned [74]. For example, suppose that a driver of a smart vehicle negligently fails to update his firmware, thus leading to a preventable accident with some other motorist who negligently ran a red light. Or suppose that a smart home is robbed because some security patch for the smart lock did not install—of course the robber is (at least partially) liable, but the manufacturer might be (at least partially) liable as well. Or map onto either of these possibilities the government's role in failing to appropriately regulate the field. Or even a lightning strike that—against all probabilities—destroyed the door lock. Could the government now be liable? Does the storm mitigate the robber's liability? What if the homeowner failed to repair the lock on some expedient timeline? Causal analysis can get complicated, quickly.

And so it again bears emphasis that what distinguishes the IoT from the traditional internet is the former's ability to act in the physical world, thus opening the possibilities of physical risk. This opens up significantly different ethical analyses, not just in terms of the overall risk, but also in terms of the associative liabilities. And any of this intersects with other ethical

¹¹ This is not to say that there are no risks from actions in cyberspace, nor even that events in cyberspace are necessarily sealed from the physical realm: for example, people suffer significant psychological harm from activity online [18], and certain actions in cyberspace can have significant geopolitical impacts [36]. However, the physical impacts from cyberspace are generally mediated through humans, human activity, or some longer causal network. In contrast, the IoT can have proximately unmediated physical impact.

¹² But see *Summers v. Tice* [94] and *Sindell v. Abbott Laboratories* [87], which allow for recovery absent established causation. These holdings have not gained much traction beyond particular sets of facts in a particular jurisdiction, though raise a host of interesting issues.

values already discussed, like informed consent and privacy: in the case of the failed door lock, any of these has applications. For example, informed consent would go to whether the user fully understood the risks of failure before installation. Privacy here could be compromised, not just in the virtual senses discussed in Section 3, but also in the physical sense that a user's home is insecure. But certainly physical safety matters here, both directly and indirectly in terms of its bearing on other ethical values.

6. Trust

Trust is essential for the complex network of systems to bring about its desired outcomes, yet often flies off the radar. “[T]rust tends not to be talked about very much. Most of the time, it is an invisible assumption” [86], p. 550]. When working efficiently, seamlessly, and safely, our trust in the IoT will be implicit, as opposed to explicit; individual components work reliably, the networks of components and operators run as expected, and we feel comfortable that the overall system is promoting the desired ends. It is only when any of these elements break down that we notice the myriad ways in which trust bears on the IoT. In this section, we look at three different conceptions of trust, their relationships to the IoT, and the ways in which artificial intelligence (AI)—an increasingly central feature of the IoT—also intersects with trust.

Trust refers to a number of different concepts. “There is a strong *prima facie* case for supposing that there is no single phenomenon that ‘trust’ refers to, nor that our folk concept has determinate rules of use. Nonetheless, this does not mean that there is no philosophical understanding of the concept to be had” [86], p. 551]. A common conception of trust is technically focused and interchangeable with the notion of “reliance” [62], pp. 18–19]. “Reliance” is a term typically used for the function of inanimate objects that does not take into account human factors. For example, while we *rely* on our vehicle's brakes to work, we *trust* our companions to drive carefully. Another conception of trust, sometimes called the “strategic” model, involves humans and sees trust as simply predictive [99], pp. 20–26]. In other words, it is concerned with what we might reasonably expect when interacting with other people. The question “can we trust our friends?” is interchangeable with the question “given past experience, how can we expect our friends to act in the future?” A third conception of trust considers the affective aspect of human relations [44], p. 5]. On this conception, trust is not simply risk analysis, but rather leads “one to anticipate that the other will have and display goodwill” [44], p. 6]. In relation to the IoT, trust can apply to the IoT in terms of whether the components will reliably serve their function, if we can expect other users and developers to act in predictable ways, and if we believe that other people are treating each other with goodwill.

Think again of the example of a driverless vehicle. Insofar as we *rely* on the driverless vehicle to work, we trust that the brakes are working properly and that the car will stop whenever that is needed. If this does not happen because the brakes malfunction, then we lose trust in that we can no longer *rely* on the car to operate properly—obviously intersecting with the discussion of physical safety in Section 5. Here, the failure of a component to function effectively points to a lack of reliability in the component and can lead to a loss of trust in the IoT more generally: if we cannot rely on the brakes in driverless vehicles to function, we either lack or lose trust in driverless vehicles more generally. Contrast this scenario with one in which we learn that a given model of car injured someone because its brakes were deactivated by a careless or malicious remote operator. Rather than mere reliance, we lose trust in that service when we do not *expect* the remote operator to drive the vehicle safely. This then goes to the discussion of information security in Section 4, where vulnerabilities in the security of communications, particularly in this case commands, has implications for our trust in the network. That is, if the information security is weak such that uncaring or maliciously motivated people can take control of components, we lose trust in the relations between user, component and operator. This manifests a reduced expectation that the network is trustworthy.

Alternatively, we can lose trust when we do not believe that people inside the company are *motivated* by our best interests. The minimum responsibility for the company is a level of openness regarding the security breach or the possibility of some nefarious intervention. Trust “is not a narrow target: it is the trust always found in friendship, often found between professionals and their clients, sometimes found between strangers, and sometimes, even, between people and their governments” [44], p. 5]. Distinct from reliability and expectation, the actions of humans within the system potentially lead to an overall affective judgment that the system itself is not trustworthy. The revelations that people working for Volkswagen had falsified and covered up emissions data:

led to a disaster in confidence. The news of the scandal received broad global coverage in the press and social media, and was the subject of political and civic discussion. In the month following the EPA's press release Volkswagen was held responsible for weakening the German economy . . . a loss of public trust in corporations generally. . . and for the possible overturning of the entire motor industry [77], p. 1509].

Should such events occur with the IoT, a similar loss of trust would occur.

In the IoT, trust also bears on the role of artificial intelligence (AI). Continuing with the example of driverless vehicles, in order for these individual vehicles to operate effectively, there will need to be coordination *between* the vehicles. The complexity of such coordination will likely require AI not just for the decisions made by discrete vehicles, but also for the complex system as a whole [25]. And in order for us to have trust in the components and the system as whole, we will need some assurance that the decisions trade on our best interests. For example, humans jaywalk: if the driverless vehicle is simply programmed to follow road laws, it might not “look” for pedestrians outside of crosswalks, thus eventuating in fatalities. Similarly, if the system-level decisions are only directed at speed, efficiency, or coordination—as opposed to also

integrating additional moral values—then we might be unhappy with the outcomes. The point here is that the combination of sensors, communications, and actuators that make up the IoT will generally require AI for coordination, and that the decision-making processes of that AI need to be trustworthy in the third sense of trust, namely that it takes our best interests seriously.

This leads us to contemplate a range of ethical issues in AI. Given the production and communication of revealing personal information, as well as the potential for physical risk, we need to ask whether AI should make morally important decisions. For example, should AI be tasked to handle intimate personal information or decisions in which people's physical safety is involved?¹³ We might think that decisions that involve significant physical risk ought to remain the province of humans; AI should not make these decisions. For example, the Campaign to Stop Killer Robots argues that “[g]iving machines the power to decide who lives and dies on the battlefield is an unacceptable application of technology” [17]. The underpinning principle is that robots and AI lack moral agency and, therefore, decisions of significant moral weight ought to remain humans'. While attractive, as James Moor presciently wrote several decades ago:

[W]hen computers drive cars, not only are human needs carried out more efficiently but there is a substantial reduction in deaths, injuries, and property damage. Further suppose in those cases in which humans override computer driving decisions, the accident rate soars. Under such circumstances there is a persuasive moral and prudential argument to have computers do the decision making and not to allow humans to override the decisions [[65], pp. 226–227].

It therefore might be that such important decisions should be exclusively the province of AI; if and when the IoT involves significant risk to people's safety, and it can be shown that humans make worse decisions than AI, then we have a *prima facie* argument that AI should be making these decisions. Furthermore, if we follow a more controversial argument that personal information only becomes a privacy concern when a human operator accesses it, then “[a]utomated surveillance . . . brings with it a degree of anonymity and privacy” [[58], p. 163]. The upshot of this view is that the more that AI can access and handle intimate personal information, the better for privacy.¹⁴

To summarize, trust is directly related to the other ethically relevant features we have discussed. For example, if we are confident that privacy will be respected, then we have established trust. Similarly, if informed consent is respected, then we will know how our data is being used or what the risks of some IoT-enabled device are, thus furthering trust. In information security were compromised, then we would retain trust by expecting we be notified of breaches and that they be expeditiously patched. Finally, if there are failures that eventuate in risks of physical safety, then we lose trust in both the component itself and in the oversight mechanisms. The point here is that trust is dependent upon the other ethically relevant features, as well as potentially indicative of when those features have not been respected.

7. Future research

In closing, we want to emphasize a point made throughout: the ethical issues that we have surveyed do not exist in isolation, but instead converge and intersect in myriad ways. For example, if informed consent is not properly tended to, risks abound with regards to privacy or information security. If privacy and information security are not properly tended to, risks abound with regards to physical safety. If *anything* is not properly tended to, risks abound with regards to trust. Much work in technology ethics is oversimplified in this regard, focusing on a single issue, while excluding the ways in which it connects to others. Dialectically, this is harder to pursue: surely it is easier to separate issues than to treat them holistically. And so our approach has been designed as much as for expedient presentation as for any other reason. That is all fine, though, so long as we recognize the limits of this approach and keep in mind that convergence and intersectionality still apply; future research can continue to elucidate these relationships.

Also, as indicated at the outset, we consider this paper to be an opening foray into an ongoing discussion regarding ethics and the IoT. The list of ethical issues we have covered, while not arbitrary, is surely incomplete. For example, we have not opined on larger societal issues, such as distributive justice or the ways in which the IoT either exacerbates or mitigates the digital divide. We have also not discussed gender or race, nor the ways that design principles could bear on discrimination, availability, or access to essential services and products. Nor have we discussed intellectual property, including trade-offs between development and regulation. Indeed, there is much work still to be done, and we encourage others to continue and to expand the conversation. Our aim has not been to be comprehensive, but rather to indicate some ways in which the IoT triggers ethical analysis. In this regard, our hope is that informed consent, privacy, information security, physical safety, and trust is foundational—in the sense that they are useful starting points that adhere to short- and mid-term applications and have established literatures, both in general and as pertains to technology ethics.

¹³ For the purposes of this paper, we use “decision” to refer some input-weighting process that results in a change in the world. This conception is tailored to be inclusive of something like AI, as well as the ways in which human beliefs and desires act as input-weighting factors that lead to actions. It fits with a “minimal definition of autonomy vis-à-vis a [system that is] capable of some significant operation without direct human oversight” [[89], p. 94]. This description of decision making may be controversial, but the very notion of “decision making” is already challenging since it must articulate both what a *decision* is, as well as what *decision making* is. For more discussion, see Steele and Stefansoon [91] and Charland [21].

¹⁴ This is not to say that we endorse this view; Henschke [[37], pp. 257–61] rejects the notion that non-semantic information is free of privacy concerns. Our point is broader in that discussions of ethics and the IoT need to recognize and respond to developing trends in AI ethics. See also Macnish [59].

But while we recognize that our focus has been limited for space constraints, we also want to suggest that at least some of this discussion generalizes. In other words, while we have explored various ethical issues through an IoT lens, those explorations can similarly elucidate discussions beyond simply the IoT context. To give one example, the military applications of the IoT are likely to be immense. While there has been increasing attention to cybernorms, specifically as relates to cyberwarfare [5,56,57,84], until recently, the cyberwarfare discussion has tended to focus on state-state conflict (e.g., Stuxnet), or state-substate conflict (e.g., North Korea hacking Sony) [57]. Substate-substate conflict has largely been ignored in this debate, and that is where much of the IoT action will fall. Furthermore—but relatedly—the cyberwarfare debate tends to be about different issues altogether, such as sovereignty, or commensurability (e.g., between cyber- and kinetic attacks). We anticipate that the evolution of the IoT will significantly impact military practices, and the practical issues that will be faced in the military context, including but not limited to the theatres of conflict ought to play a role in shaping the ethical concepts underpinning the just war tradition, ethical espionage, covert operations and the like. And so IoT matters for this context, but so do completely other literatures relating to military ethics and security studies [4].

Acknowledgment

The authors thank Patrick Lin, Jonathan Milgrim, and Fatos Xhafa for helpful discussions regarding this paper.

References

- [1] A. AboBakr, M. Azer, IoT ethics challenges and legal issues, in: *Proceedings of the Twelfth International Conference on Computer Engineering and Systems*, 2017.
- [2] Y. Agarwal, A.K. Dey, Toward building a safe, secure, and easy-to-use internet of things infrastructure, *IEEE Comput.* 49.4 (2016) 88–91.
- [3] F. Allhoff, Risk, precaution, and emerging technologies, *Stud. Ethic. Law Soc.* 3.2 (2009) 1–27.
- [4] F. Allhoff, G. Nicholas, Evans, and Adam Henschke, *The Routledge Handbook of Ethics and War: Just War Theory in the Twenty-First Century*, Routledge, London, 2013.
- [5] F. Allhoff, A. Henschke, B. Jay Strawser, *Binary Bullets: The Ethics of Cyberwarfare*, Oxford University Press, Oxford, 2016.
- [6] F. Allhoff, P. Lin, *Nanotechnology & Society: Current and Emerging Ethical Issues*, Springer, Dordrecht, 2008.
- [7] F. Allhoff, P. Lin, J. Moor, J. Weckert, *Nanoethics: The Social & Ethical Implications of Nanotechnology*, John Wiley & Sons, Hoboken, NJ, 2007.
- [8] F. Allhoff, P. Lin, D. Moore, *What Is Nanotechnology and Why Does It Matter: From Science to Ethics*, Wiley-Blackwell, Oxford, 2010.
- [9] American Law Institute, *Uniform Commercial Code*, American Law Institute, 2002.
- [10] American Medical Association, *Code of Medical Ethics*, American Medical Association, Chicago, 2016.
- [11] S. Banerjee, T. Hemphill, P. Longstreet, "Is IoT a Threat to Consumer Consent?" (unpublished, 2017).
- [12] J. Barnes, *The Complete Works of Aristotle*, 2, Princeton University Press, Princeton, 1984.
- [13] F. Berman, V. Cerf, *Social behavior in the Internet of Things*, *Commun. Assoc. Comput. Mach.* 60.2 (2017) 6–7.
- [14] A. Bogle, "Strava has published details about secret military bases, and an Australian was the first to know," *ABC News* (2018, 30 January). Available at <http://www.abc.net.au/news/science/2018-01-29/strava-heat-map-shows-military-bases-and-supply-routes/9369490>.
- [15] P. Brey, *Anticipatory ethics for emerging technologies*, *Nanoethics* 6.1 (2012) 1–13.
- [16] *Business Wire*, "Explosive internet of things spending to reach \$1.7 Trillion in 2020, According to IDC" (June 2, 2015). Available at <https://www.businesswire.com/news/home/20150602005329/en/Explosive-Internet-Things-Spending-Reach-1.7-Trillion>.
- [17] Campaign to Stop Killer Robots, "Learn" (2018). Available at <https://www.stopkillerrobots.org/learn/>.
- [18] D. Canetti, M.L. Gross, I. Waismel-Manor, "Immune from cyberfire?: The psychological and physiological effects of cyberwarfare." In Allhoff et al. (2016), pp. 157–76.
- [19] *Carpenter v. United States*, 585U.S. (2018).
- [20] E. Chapman, T. Uren, *The Internet of Insecure Things*, Australian Strategic Policy Institute, Canberra, 2018.
- [21] L.C. Charland, "Decision-making capacity," *The Stanford Encyclopedia of Philosophy* (2015). Available at <https://plato.stanford.edu/archives/fall2015/entries/decision-capacity>.
- [22] N. Chokshi, "Is Alexa listening?: Amazon echo sent out recording of couple's conversation," *New York Times* (May 25, 2018). Available at <https://www.nytimes.com/2018/05/25/business/amazon-alexa-conversation-shared-echo.html>.
- [23] J. DeCew, "Privacy," *Stanford Encyclopedia of Philosophy* (2018). Available at <https://plato.stanford.edu/entries/privacy/>.
- [24] Department of Defense, "DoD News Briefing: Secretary Rumsfeld and Gen. Myers" (February 12, 2002).
- [25] A. Dreves, M. Gerdt, *A generalized nash equilibrium approach for optimal control problems of autonomous cars*, *Optim. Control Appl. Methods* 39 (1) (2017) 326–342.
- [26] European Parliament and Council of the European Union, *Regulation (EU) 2016/679 (General Data Protection Regulation)* (2016).
- [27] G. Eysenbach, J.E. Till, *Ethical issues in qualitative research on internet communities*, *Br. Med. J.* 323 (2001) 1103–1105.
- [28] R.R. Faden, T.L. Beauchamp, *A History and Theory of Informed Consent*, Oxford University Press, Oxford, 1986.
- [29] M. Fisher, "Here's the e-mail trick Petreus and broadwell used to communicate," *Washington Post* (November 12, 2012). Available at https://www.washingtonpost.com/news/worldviews/wp/2012/11/12/heres-the-e-mail-trick-petraeus-and-broadwell-used-to-communicate/?utm_term=.53d9e006c076.
- [30] D. Fletcher, *Internet of Things*, in: M. Blowers (Ed.), *Evolution of Cyber Technologies and Operations to 2035*, Springer, Dordrecht, 2015, pp. 19–32.
- [31] L. Franceschi-Bicchieri, "Internet of things teddy bear leaked 2 million parent and kids message recordings," *Motherboard* (February 27, 2017). Available at: https://motherboard.vice.com/en_us/article/internet-of-things-teddy-bear-leaked-2-million-parent-and-kids-message-recordings.
- [32] M. Garber, "Email location data led to the discovery of the Petraeus affair," *The Atlantic* (November 12, 2012). Available at <https://www.theatlantic.com/technology/archive/2012/11/email-location-data-led-to-the-discovery-of-the-petraeus-affair/265093/>.
- [33] M. Goodman, "The Internet of Things will turn our machines against us," *Wired* (January 15, 2016). Available at <https://www.wired.co.uk/article/internet-of-hackable-things>.
- [34] *Google v. Costeja*, C-131/12 (2014).
- [35] *Griswold v. Connecticut*, 381U.S. 479 (1965).
- [36] S. Haggard, J.R. Lindsay, *North Korea and the Sony Hack: Exporting Instability through Cyberspace*, *AsiaPacific Issues* 117 (2015) 1–8.
- [37] A. Henschke, *Ethics in an Age of Surveillance*, Cambridge University Press, Cambridge, 2017.
- [38] A. Henschke, *The Internet of Things and dual layers of ethical concern*, in: P. Lin, K. Abney, R. Jenkins (Eds.), *Robot Ethics 2.0*, Oxford University Press, New York, 2017, pp. 229–243.
- [39] K. Hill, "How target figured out a teen girl was pregnant before her father did," *Forbes* (February 16, 2012). Available at <https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/#7f6e25046668>.
- [40] K. Hill, S. Mattu, "The house that spied on me," *Gizmodo* (February 7, 2018). Available at <https://gizmodo.com/the-house-that-spied-on-me-1822429852>.

- [41] J. Hiner, "New research: most IoT devices can be hacked into botnets," TechRepublic (March 7, 2018). Available at <https://www.techrepublic.com/article/new-research-most-iot-devices-can-be-hacked-into-botnets/>.
- [42] R.A. Hillman, J.J. Rachlinski, Standard-form contracting in the electronic age, *N.Y.U. Law Rev.* 77 (2002) 429–495.
- [43] R. Jenkins, "Cyberwarfare as Ideal War" in Allhoff et al. (2016), pp. 89–114.
- [44] K. Jones, "Trust as an Affective Attitude," *Ethics* 107 (1) (1996) 4–25.
- [45] M. Kan, IoT botnet highlights the dangers of default passwords," *InfoWorld* (October 4, 2016). Available at <https://www.infoworld.com/article/3127167/password-security/iot-botnet-highlights-the-dangers-of-default-passwords.html>.
- [46] J. Katz, Informed consent—must it remain a fairy tale, *J. Contemp. Health Law Policy* 10.69 (1994) 69–91.
- [47] N. Kobie, "The Internet of Things: convenience at a price," *The Guardian* (March 30, 2015). Available at: <https://www.theguardian.com/technology/2015/mar/30/internet-of-things-convenience-price-privacy-security>.
- [48] D. Kravets, "Makes of internet of things-connected vibrator will settle privacy suit," *Ars Technica* (December 8, 2016). Available at <https://arstechnica.com/tech-policy/2016/12/maker-of-internet-of-things-connected-vibrator-will-settle-privacy-suit/>.
- [49] M. Lazarescu, "Hacked by your fridge: the internet of things could spark a new wave of cyber attacks," *The Conversation* (October 6, 2016). Available at <https://theconversation.com/hacked-by-your-fridge-the-internet-of-things-could-spark-a-new-wave-of-cyber-attacks-66493>.
- [50] A. Levy, "Cisco Wagers \$1.4 billion on 'Internet of Things,'" *CNBC* (February 3, 2016). Available at <https://www.cnbc.com/2016/02/03/cisco-wagers-14-billion-on-internet-of-things.html>.
- [51] S. Li, L.D. Xu, S. Zhao, The Internet of Things: a survey, *Inf. Syst. Front.* 17 (2) (2015) 243–259.
- [52] P. Lin, "Introduction to Robot Ethics" in Lin et al. (2012), pp. 3–16.
- [53] P. Lin, K. Abney, G. Bekey, *Robot Ethics: The Ethical and Social Implications of Robotics*, MIT Press, Cambridge, MA, 2012.
- [54] P. Lin, K. Abney, R. Jenkins, *Robot Ethics 2.0*, Oxford University Press, Oxford, 2017.
- [55] J. Locke, (Indianapolis, IN: Hackett Publishing, 1980).
- [56] G. Lucas, "Emerging norms for cyberwarfare," in Allhoff et al. (2016), pp. 13–33.
- [57] J. Katz, *Ethics and Cyber Warfare*, Oxford University Press, Oxford, 2016.
- [58] K. Macnish, Unblinking eyes: the ethics of automating surveillance, *Ethics and Information Technology* 14 (2) (2012) 151–167.
- [59] K. Macnish, *The Ethics of Surveillance: An Introduction*, Routledge, London, 2017.
- [60] S.L. Mak, H.K. Lau, An implementation of toy safety assessment model, in: *Proceedings of the IEEE Symposium on Product Compliance Engineering*, 2014, pp. 12–16.
- [61] C. Matyszczyk, "Samsung's warning: our smart TVs record your living room chatter," *CNET* (February 8, 2015). Available at <https://www.cnet.com/news/samsungs-warning-our-smart-tvs-record-your-living-room-chatter/>.
- [62] C. McLeod, "Trust," *Stanford Encyclopedia of Philosophy* (2015). Available at <https://plato.stanford.edu/entries/trust/>.
- [63] J.S. Mill, *On Liberty*, in: E. Rapaport (Ed.), Hackett Publishing, Indianapolis, IN, 1978.
- [64] J. Mohammed, "5 Predictions for the Internet of Things in 2016," *World Economic Forum* (December 15, 2015). Available at <https://www.weforum.org/agenda/2015/12/5-predictions-for-the-internet-of-things-in-2016/>.
- [65] J. Moor, Are there decisions computers should never make? *Nat. Syst.* 1 (4) (1979) 217–229.
- [66] A. Moore, Defining privacy, *J. Soc. Philos.* 39 (3) (2008) 411–428.
- [67] National Telecommunications & Information Administration [NTIA], "Fostering the advancement of the Internet of Things" (January 2017). Available at https://www.ntia.doc.gov/files/ntia/publications/iot_green_paper_01122017.pdf.
- [68] R. Neisse, G. Baldini, G. Steri, Y. Miyake, S. Kiyomoto, A.R. Biswas, An agent-based framework for informed consent in the Internet of Things, in: *Proceedings of the 2015 IEEE Second World Forum on Internet of Things*, 2015, pp. 789–794.
- [69] R. Nozick, *Anarchy, States, and Utopia*, Basic Books, New York, 1974.
- [70] S.R. Peppet, Regulating the Internet of Things: first steps toward managing discrimination, privacy, security, and consent, *Texas Law Rev.* 93 (2017) 85–178.
- [71] D. Popescu, M. Georgescu, "Internet of Things: some ethical issues," *USV Annal. Econ. Publ. Administr.* 13 (2) (2013) 208–214.
- [72] *Pretty v. United Kingdom*, 2346/02 (2002).
- [73] *ProCD, Inc. v. Zeidenberg*, 86 F.3d 1447 (7th Cir. 1996).
- [74] W.L. Prosser, *Comparative negligence*, *Michig. Law Rev.* 51 (4) (1953) 465–508.
- [75] M. Redden, "Tech company accused of collecting details of how customers use sex toys," *The Guardian* (September 14, 2016). Available at <https://www.theguardian.com/us-news/2016/sep/14/wevibe-sex-toy-data-collection-chicago-lawsuit>.
- [76] Reuters, "Uber settles with family of woman killed by self-driving car," *The Guardian* (March 29, 2018). Available at <https://www.theguardian.com/technology/2018/mar/29/uber-settles-with-family-of-woman-killed-by-self-driving-car>.
- [77] C. Rhodes, "Democratic business ethics: Volkswagen's emissions scandal and the disruption of corporate sovereignty," *Org. Stud.* 37 (10) (2016) 1501–1518.
- [78] P. Roberts, "Pretty much all consumer internet of things vulnerabilities are avoidable," *The Security Ledger* (September 13, 2016). Available at <https://securityledger.com/2016/09/pretty-much-all-consumer-internet-of-things-vulnerabilities-are-avoidable/>.
- [79] S. Robbins, A. Henschke, Designing for democracy: bulk data and authoritarianism, *Surveil. Soc.* 15.3 (2017) 582–589.
- [80] *Roe v. Wade*, 410 U.S. 113 (1973).
- [81] *Rotaru v. Romania*, 28341/95 (2000).
- [82] J. Rosen, The right to be forgotten, *Stanford Law Rev.* 64 (2012) 88–92.
- [83] J. Roy, 'Polis' and 'Oikos' in classical Athens, *Greece Rome* 46 (1) (1999) 1–18.
- [84] M. Schmidt, L. Vihul, "The emergence of international legal norms for cyber-conflict" in Allhoff et al. (2016), pp. 34–55.
- [85] A. Scroton, "UK IoT research hub opens with support from academic world," *ComputerWeekly.com* (January 6, 2016). Available at <https://www.computerweekly.com/news/4500270034/UK-IoT-research-hub-opens-with-support-from-academic-world>.
- [86] T.W. Simpson, What is trust? *Pacif. Philosop. Q.* 93 (4) (2012) 550–569.
- [87] *Sindell v. Abbott Laboratories*, 607 P.2d 924 (Cal. 1980).
- [88] B. Smith, "Facial recognition technology: the need for public regulation and corporate responsibility," *Microsoft Issues* (July 13, 2018). Available at <https://blogs.microsoft.com/on-the-issues/2018/07/13/facial-recognition-technology-the-need-for-public-regulation-and-corporate-responsibility/>.
- [89] R. Sparrow, Robots and respect: assessing the case against autonomous weapon systems, *Ethics Int. Affairs* 30.1 (2016) 93–116.
- [90] J.A. Stankovic, Research directions for the Internet of Things, *IEEE Int. Things J.* 1.1 (2014) 3–9.
- [91] K. Steele, H. Orri Stefánsson, "Decision theory", *The Stanford Encyclopedia of Philosophy* (2016). <https://plato.stanford.edu/archives/win2016/entries/decision-theory/>.
- [92] E. Stewart, "What the government could actually do about Facebook," *Vox* (April 10, 2018). Available at <https://www.vox.com/policy-and-politics/2018/4/10/17208322/facebook-mark-zuckerberg-congress-testimony-regulation>.
- [93] *Step-Saver Data Systems, Inc. v. Wyse Technology* 939 F.2d 91 (3rd Cir. 1991).
- [94] *Summers v. Tice*, 199 P.2d 1 (Cal. 1947).
- [95] T.S., "Why Uber's self-driving car killed a pedestrian," *The Economist* (May 29, 2018). Available at <https://www.economist.com/the-economist-explains/2018/05/29/why-ubers-self-driving-car-killed-a-pedestrian>.
- [96] A. Twerski, A.S. Weinstein, H.R. Piehler, W.A. Donaher, Product liability: a study of the interaction of law and technology, *Duquesne Law Rev.* 12.3 (1974) 425–464.

- [97] H. Ungureanu, Massive dyn DDOS attack: experts blame smart fridges, dvrs, and other iot devices why your internet went down, *Tech Times* (October 24, 2016). Available at <https://www.techtimes.com/articles/183339/20161024/massive-dyn-ddos-attack-experts-blame-smart-fridges-dvrs-and-other-iot-devices-why-your-internet-went-down.htm>.
- [98] L. Urquhart, D. McAuley, Avoiding the internet of insecure industrial things, *Comput. Law Secur. Rev.* 34.3 (2018) 450–466.
- [99] E.M. Uslaner, *The Moral Foundations of Trust*, Cambridge University Press, Cambridge, 2010.
- [100] P. Vallentyne, B. van derVossen, "Libertarianism," *Stanford Encyclopedia of Philosophy* (2014). Available at <https://plato.stanford.edu/entries/libertarianism/>.
- [101] I. van de Poel, Z. Robaey, Safe-by-design: from safety to responsibility, *NanoEthics* 11.3 (2017) 297–306.
- [102] R.H. Weber, Internet of Things—governance *Quo Vadis*, *Comput. Law Secur. Rev.* 29 (2013) 341–347.
- [103] R.H. Weber, Internet of Things: privacy issues revisited, *Comput. Law Secur. Rev.* 31 (2015) 618–627.
- [104] M. White, Internet of Hackable Things: wired world wide open to new age of cyber crime, *Sydney Morning Herald* (May 22, 2015). Available at <https://www.smh.com.au/technology/internet-of-hackable-things-wired-world-wide-open-to-new-age-of-cyber-crime-20150522-gh7c3q.html>.
- [105] A. Whitmore, A. Agarwal, L.D. Xu, The Internet of Things—a survey of topics and trends, *Inf. Syst. Front.* 17 (2) (2015) 261–274.

Fritz Allhoff, J.D., Ph.D. is a Professor in the Department of Philosophy at Western Michigan University. Parts of this paper were written while as a Fellow in the Center for Law and the Biosciences at Stanford University and as a Fulbright Specialist in the Faculty of Political Science at the University of Iceland; he thanks those institutions for their support. Finally, he thanks the United States National Science Foundation, which has provided generous support under award # 1317798.

Adam Henschke, Ph.D. is a Senior Lecturer in the National Security College at the Australian National University and a Senior Research Fellow at the Delft University of Technology (Netherlands). Preliminary research on this topic was conducted as a Visiting Researcher at the Brocher Foundation (Switzerland) and as a Visiting Fellow at the Uehiro Centre for Practical Ethics at the University of Oxford; he thanks those institutions for their support. Finally, he thanks the European Research Council and the Australian Research Council, which have provided generous support under awards Advanced Grant project on Collective Responsibility and Counterterrorism and # DP180103439, respectively.