The State of Responsible IoT 2018

Experts from the ThingsCon community explore current challenges and opportunities of responsible IoT. Learn more at thingscon.com.

THINGS

Version: v1.03 / 30 August 2018
First published: 24 August 2018

ThingsCon Report: The State of Responsible Internet of Things 2018

Published by ThingsCon e.V., Berlin
August 2018

Other available formats

The report is available as a publication on Medium as well as a PDF download. An overview including all links can be found at: bit.ly/riot-report

Table of Contents

Content listed in alphabetical order, by first names.

Prolog

A lot has happened since we published the first ThingsCon State of Responsible IoT report in 2017: Responsibility and ethics in tech have begun to enter mainstream conversations, and these conversations are having an effect. The media, tech companies, and policy makers all are rethinking the effect of technology on society.

The lines between the Internet of Things (IoT), algorithmic decision-making, Artificial Intelligence/Machine Learning (AI/ML), and data-driven services are all ever-more blurry. We can't discuss one without considering the others. That's not a bad thing, it just adds complexity. The 21st century one for black and white thinking: It's messy, complex, quickly evolving, and a time where simple answers won't do.

It is all the more important to consider the implications, to make sure that all the new data-driven systems we'll see deployed across our physical and digital environments work well—not just for the users but for all who are impacted.

Things have evolved and matured in big strides since our last State of Responsible IoT. This year's report reflects that evolution, as well as the enormous breadth and depth of the debate. We couldn't be happier with the result.

Solid advice for you IoT gunslingers and quacksters out there: Time to change your business model

By Christian Villum

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

Imagine yourself being in some place called Tombstone Creek or Deadwood Gorge, Wyoming or Colorado, around 1858 or so - or any other frontier town in the wild west around that time, the way they are often shown in Western movies: Gunslingers, quack doctors, prostitutes, gold diggers, fraudsters and racketeers in this place are feeding on all the fortune seekers rushing in from the old world out east. At the height of the gold rush, the American 18th century frontier of the heartland and west coast outback presented ample opportunity for short-term small racketeering business success with its vague and hastily crafted laws and only a few easily corruptable sheriffs to keep some degree of order. Shady entrepreneurs had their heyday back then, no doubt about that.

In many ways this scenario - the good ol' Wild West, which of course really wasn't all that good - resembles that of the global landscape of the IoT industry today. Quick bucks are made by hasty entrepreneurs and nimble hardware startups, whose ship-now-and-deal-with-trouble-later strategies makes up for a very fragmented and perplexing market with little to no regulation and the Internet to easily propel all kinds of shady products to every corner of the world. Will it stay this way? If not, what comes next? History and capitalism will tell us.

Let's start out by looking at the business model of one of those opportunistic merry men and women of the wild west - the quack doctor: Upfront payment for a product such as a balm or wonder liquid which, according to the sales pitch given impromptu from the back of the prairie wagon, could [insert some amazing and

unique benefit], but which most likely didn't yield anything, and in worse cases might have left the gullible customer with a rash. By the time the lack of quality became apparent, the quack would have hauled their prairie wagon off to the next town and maybe taken a new business name. It was quick money for sure, and maybe good short term business for the seller. For the customer? Not so much.

Does this sound vaguely familiar in today's settings? Well, buy an IoT product nowadays from one of the present day digital prairie wagons and you might miss the days when a wasting a little money was all that happened to you. More likely the damage done by the IoT-balm these days would be some horrible privacy violation, lacking cybersecurity, or no interoperability with the other cheap IoT-balms you've acquired. You might end up not just having wasted money, but potentially gotten hacked, had your credit card misused, or find your new acquisition has become part of an evil bot army without your consent or knowledge. Worst case scenarios include you having your identity stolen or house burned down from a short-circuiting connected toaster that went haywire. It's a wild west out there in IoT-land these days.

Historically in America, at some point there was no longer any outlawed frontier land to flee to for the quacks, and society started to become more and more organized. The shanty gold digger towns turned into cities, the prairie turned into farmland and business became organized and taxed. The arm of the law became longer, and the offering of pharmaceuticals became more restricted and regulated. (I am not sure they got rid of *all* the quacks, but you get the picture.)

Will this happen in the global IoT space? I am sure it will, although maybe not only through the same kind of regulatory processes as we saw helped shape modern American society. There will be no cavalry coming to the rescue in the darkest hour on the global Internet: We don't have one, and that is probably for the best. Instead I would argue that the IoT industry actors will start to clean up their act to stay competitive and to be able to build business in the 2020s and beyond. Offering responsible tech will simply be the best business model. Customers will flock to the suppliers of trustworthy, secure, interoperable and open products and services.

Why? The IoT global market will not remain an industry driven mainly by selling hardware, if it even ever was. Once most homes, workplaces, cars, watches and implants become fully connected, the market for hardware - while still large - will be much more saturated, and entry will be harder. As will scaling and growth. The real driver of good and scalable IoT-business will be data. Data and subscription models, where businesses try to get the customer to obtain the actual hardware as

easily and cheaply as possible (maybe even freely) in order to then establish the foundation for real business: A mutually beneficial relationship to ensure a long-term and much more profitable business opportunity.

All this will be based on trust, which as a business you will have to earn from your customer. This will not happen if you cut corners on hardware that violates users rights, or is easy to hack, or does not give customers the freedom to own their own data or move it elsewhere in a click or two. You will simply not be able to build the foundation for good long-term business that way.

That is not to say that we will not need some form of regulation. Countries will have to instate their own protections to shelter citizens from predatory IoT quacks. We need some international standards that customers can use to navigate the field and to help them distinguish between proper businesses and the quacks. A good example is the IoT Trustmark, which has the potential to create a bar that it will be hard for businesses to refrain from staying above in order to stay profitable.

Moving from where we are now to a future global market dominated by a more credible IoT-industry will not happen overnight. I think it's fair to say that in comparison with the timeline of the American frontier (as it moved from east to west) we are maybe somewhere around Utah. California is still some way out on the horizon, but if there is anything we've learned from modern capitalism it is that it is often not the best strategy to hang on to the old cash cows for too long instead of adapting towards what comes next. The industry needs to start changing their business models now if they want to save their hides (pun intended) and stay relevant. Let's learn from history. Yeehaw!

Driven by a keen interest in exploring new boundaries for strategic design, **Christian Villum**'s work as Director of Digital & Future Thinking at the Danish Design Centre examines new ideas in the span between technology and design thinking.

With a background in maker technology, new business models, sharing cultures, open data and open design, internet culture and hacktivism, he enjoys developing communities and bringing people together to share new ideas and generate change. His work explores future currents in technology from a design perspective, and includes, among other things, programs for new open source business models

for manufacturing, establishment of global Fab Cities and human-centric approaches to technology. He is a frequent public speaker, blogger and was the editor and co-writer of the book 'Open Source City' (2016).

Christian's previous work includes co-founding and heading the experimental Platform4 Art & Technology hub, being a frontrunner in the use of Creative Commons content licenses, building global communities for the UK-based non-profit organisation Open Knowledge Foundation and initiating a wide range of companies and projects.

IoT and Value. A dangerous game.

By Dries de Roeck

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

"It's a sin with no name
Like a hand in a flame
And our senses proclaim
It's a dangerous game."
— Jekyll and Hyde.

As part of last year's thingscon RIOT report, I wrote about IoT design processes [1]. The central conclusion made was the lack of (conscious) human centred design approaches in IoT startups, where a technology-first approach is still very dominant. This year, I would like to dig a little deeper and touch upon the somewhat ambiguous term 'value', which is - in its many forms and nuanced appearances - always part of a design process at some point. I believe thinking about value can help a lot in becoming more conscious about how a more humane internet of things can be framed and understood.

Your value isn't mine

The ambiguity around 'value' is mostly built upon the very diverse interpretations given to it. Suppose that you'd put a sociologist, a marketeer, a psychologist, a computer scientists and an economist in the same room it should be no surprise that they'd each give a very different definition of 'value' from their perspective. In

general, I believe there are three larger clusters of value to be introduced. I very much approach these from a product or service perspective, which generates a flurry of different types of value.

Image adapted from Aneeque Ahmed - the noun project

Human values

This is all about what the stakeholders involved in using a product or service perceive. These values are very emotionally driven and are hard to quantify. It is about the 'feeling' something gives you when using it. In a publication on this topic, Irene Ng[2] talks about 'Phenomenal' (P) versus 'Access' (A) value, where P-type value focusses on conscious and measurable experiences and A-type value is about the 'heightened awareness' or our personal perception of something we experience.

When it comes to designing IoT products and services, it obviously is very hard to design for something that every person can perceive or experience differently. Nevertheless, elements such as perceived quality (material quality and/or service quality), transparency about data collection and usage, agency and openness in product usage and personalisation can all have a significant impact on the values delivered to all involved stakeholders.

To illustrate this with an example, PLEQ is an internet connected sensor used for predictive maintenance systems. PLEQ allows 'upgrading' older machinery by adding a sensor box to it which monitors anomalies in machine behaviours (primarily using vibration detection). It is the type of hardware that is hidden and is monitored mainly via software systems. However, once these sensors were introduced to the market, the users asked for more visible sensor boxes (brightly coloured, perhaps not in 'just' a square box). The reason behind this was that the companies using the product wanted to be able to show to their clients visiting their warehouses or production halls that their machines were being monitored by this IoT system. However blunt, this example does show that user value should be taken seriously from the start.

Societal value

The second type of value is societal, which is again harder to quantify and sits more on the human values side. It is, however, different because societal value relates to a group of people (or a culture). Related to IoT, a product might not impact you directly - but might substantially impact society. An example of this is the UK based flood.network, which is a service focused on reporting floods throughout the country. A distributed network of privately owned and maintained sensors keeps track of the water level around the country. When aggregating all data points, trends in data can be spotted and local communities can be warned about imminent floods or risks thereof. Interestingly, this only works because multiple people collaborate. One sensor on its own doesn't really do much, which means that by investing in the flood network there are societal values at play. Let's not forget that societal value can also be impacted negatively, and IoT might (unfortunately) be very good at doing so. An inherent characteristic of IoT is that it uses a digital medium by default, implying that having access to this medium is essential. In many cases, a large slice of society is left out of IoT products and services because of not being able to access it in the first place. Examples such as the internet connected urban participation pavilion do bridge this gap, by providing a way to interact with a digital system for everyone.

Company value

The third type of value is company value, which should be regarded as a more quantifiable type of value. It very much ties in to the ‘hard factors’ like cost, revenue and overall business model related to a product or service.

Another aspect to company value, specific to internet of things products, are less tangible ‘assets’ which can be gained. Typically this is about gathering data coming from a device (through sensors) or some kind of user monitoring. Gathering this data has its potential, but can very quickly end up being abused or misused. Sometimes this abuse is conscious, in other cases it may not be intentional or even truly clear, but collecting data can have a nasty side. On the other hand, the premise of data gathering about the usage of a product or service is to develop a stronger relationship between companies and clients and build a two way conversation using the product or service as a mediator. One of the true strengths of an IoT product is that companies can change a product’s behaviour after it has been launched. The schoolbook example thereof is the Tesla car, which constantly updates its GUI based on user interaction and feedback. Doing so, Tesla can push novel features and other updates to the car based on (amongst others) gathered user data. In this case, the gathered user data increases company value. How it impacts user value depends on the case.

Value interplay

All of these types of value are in constant interplay with each other. Interestingly, there is no one-to-one relationship between the different types of value involved: Focussing on generating company value doesn't necessarily lead to negative impact on societal or human values or vice versa. It is essential to note that the before-mentioned types of value and values are not at all exclusive to design and development in an IoT context. What is, however, very specific to IoT development is that technology is thrown into the mix. Technology has the capability to impact value in a hidden way. For instance, data can be gathered from a device and used by a company to increase a product’s company value. In many cases, the people using that product are either not aware how their data is used, or they have no (direct) access to this data at all. If they ‘found out’ or if a data breach occurred it would negatively impact the perceived human value that this product delivers.

On the other hand, if a company included a sensor to enhance the end user experience (thus increasing perceived human values) but neglects the impact on the company's cost structure, one could argue that the focus on human values impacts company value negatively: The company might end up with a happy customer but doesn't necessarily have a viable product.

While there is nothing *really* new to this thinking, technology is making it more difficult for people in design and development to be fully aware of the impact design choices have on the 'invisible', technological, parts of a product or service.

Impacting value through IoT

Knowing that these values are constantly at play, not only during the design process but also when an IoT product is used, is one thing. But how to impact them, or at least take informed decisions during the design process? To do so, there are three aspects to the internet of things that can be helpful and should be considered when defining an IoT concept.

Identifying opportunities

IoT allows to combine data across diverse sectors. This has the premise to open up a design space for radically novel ideas. When defining these ideas, it is important to constantly have a very broad view on how the product or service being created relates to other sectors or other technologies. In a design and development team, this means that organisations should be aware and offer flexibility to reach out to industries which might seem very unrelated in the beginning. In literature, these activities are referred to as 'boundary spanning', which is by definition a very open and unstructured activity [3]. Boundary spanning related work is not necessarily required for conceptualising an IoT product, but by considering unrelated market sectors it becomes possible to have a much larger impact on different types of value.

Digital connectivity

The premise of designing an IoT product is by adding a digital component to a product, its perceived human value will go up as well as its measurable company value. In reality, we should be very wary about this. Firstly, just by adding a digital component to a product it won't automatically become an IoT product. There are specific characteristics to take into account which need to be consciously considered during the design and development process. An example of this is Alexandra Deschamp-Sonsino's Litmus Test for the IoT which hints towards several of such categories. Secondly, in current practice, digital technology is too often considered the driving force of a digitally connected product or service. Adding more technology shouldn't be the goal. Instead we should strive to find ways to get inspired by the 'value opportunities' digital components can offer from a human and company perspective. In both cases, responsibility and awareness play a central role.

Future implications

A last element very specific to IoT design is to consciously consider future product or service changes over time and the implications on human, company and societal value thereof. An IoT design and development team should be prepared and open for ever-changing functions and be willing to revisit the underlying value propositions of the product. When working on future implications of a product, reflections are made on how an internet connected product can evolve over time. This type of design work is typically done in future forecasting projects (i.e. by setting out scenarios in the short or long term), and has the intention to be divergent instead of convergent. However challenging, this activity is important because it allows and forces a design team to point out the spectrum of possibilities a product or service might hold over time. This 'future implications' work should be done during the ideation process, as part of a concept definition and not as something which is done after a product launch. Defining future implications could eventually become part of a product launch strategy, where not all product functions are implemented or included - deliberately - leaving flexibility for the organisation to figure out how their offering can be adapted to better match the market.

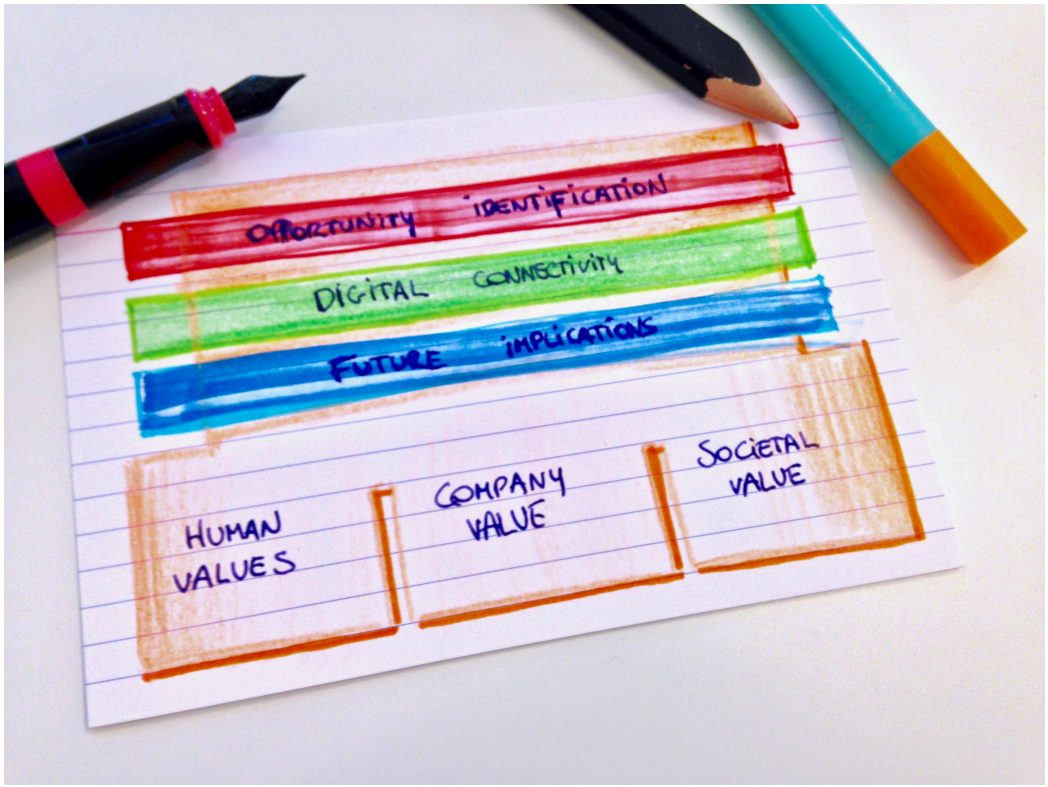

Work in progress: a visual consolidation of this value framework for IoT.

tl;dr: Be explicit

In order to understand value better, it is important to understand and be clear about the intended actions of a product. A helpful way to be explicit is by at least knowing which elements play part in the designed system:

- Which people interact with objects, which objects interact with other objects
- What type(s) of interaction are used? Are they hardware based, or do they solely rely on data?
- Are there objects that interact between each other by sharing data or aggregating data from linked sources?
- In which context or environment does all of this take place, what is the role of this environment?

Getting insight into this interlinked, underlying system of interactions was the spark that led up to creating the IoT ideation card deck. It helps in structuring and communicating about network connected product service systems by offering a personalisable deck of cards to build system maps. This tool is for sure not the holy grail, but it does support diverse design and development teams in taking more conscious design decisions.

A reaction often encountered when presenting a tool like the IoT ideation cards is that it takes long and merely states the obvious. The case I would argue for is that **taking the time and stating the obvious** might as well be what our industry needs in order to consciously design the responsible IoT we're all trying to contribute to.

Shoutouts

This thinking luckily didn't just sprout randomly out of my own brain. Thanks Pieter, Alexis, Ingrid, Karin, Albrecht, Alex, Iskander, Laura, Nathalie, Nik, Peter, Simon, Simone, Sören, Elisa and Iohanna for challenging and helping me in understanding all of this better.

Footnotes

1. Dries De Roeck, (2017), On IoT Design Processes, Thingscon RIOT report
2. Irene CL Ng, Laura A Smith, Stephen L Vargo, (2012), An integrative framework of value
3. Susan E. Reid and Ulrike de Brentani, (2004), The Fuzzy Front End of New Product Development for Discontinuous Innovations: A Theoretical Model.

Dries de Roeck is a designer, researcher and leads all things research at the creative agency Studio Dott (Belgium). In his research work, he questions how design processes change when digital and physical products become increasingly intertwined. He is the creator of the IOT ideation cards and sporadically hosts local Thingscon events.

Governance of Internet of Things and Ethics of Intelligent Algorithms

By Prof. Dr. Eduardo Magrani & Dr. Ronaldo Lemos

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

New technical artifacts connected to the Internet constantly share, process, and storage a huge amount of data. This practice is what unifies the concept of Internet of Things to the concept of Big Data[4]. A wide range of data's variances can significantly change the way we live[5]. With the growing dissemination of Big Data and computing techniques, technological evolution and economic pressure spread rapidly, and algorithms have become a great resource for innovation and business models. This rapid diffusion of algorithms and their increasing influence, however, have consequences for the market and for society, consequences which include questions of ethics and governance[[6]].

Given that algorithms can permeate countless branches of our lives, as they become more sophisticated, useful, and autonomous, there is a risk that they will make important decisions, replacing human beings. To foment the integration of algorithms into social and economic processes, algorithms governance tools are needed[7].

The governance of algorithms can vary from the strictly legal and regulatory point of view, to the purely technical point of view. This depends on some factors, such as the nature of the algorithm, the context or its risks. Market-oriented solutions or grassroots government mechanisms can occur, as seen, at multiple levels. In the first case, there is the possibility, for example, of regulation by private companies, through internal organization, and self-regulation of the entire industry. In both

cases, the standards adopted should be based on the public interest. However, in case of government regulation, the standards should focus on requirements such as the level of transparency or quality of service.

Among the regulation points are transparency, responsibility - which is linked to notions of justice and due process - and technical guarantees, as well as the development of ethical principles regarding the use of personal data. It should be noted that algorithms are constantly working and facing unplanned and unprecedented situations frequently, so that their monitoring must be constant.

One of the main themes raised by doctrine when it comes to governance is the *opacity of the algorithms*. The problem of opacity is related to the difficulty of decoding the result generated by the algorithm. Thus, there has been talk of the need for greater transparency, which could be achieved by regulating.

Researchers at the University of Zurich[8] argue that algorithm governance must be based on identified threats and suggest a *risk-based approach*, highlighting those related to manipulation, bias, censorship, social discrimination, privacy breaches, property rights and abuse of market power. To prevent these risks from materializing, it is necessary to resort to governance.

Considering these complex systems of these non-human agents, the debate on liability and ethics - already raised when presenting technical artifacts - returns. Issues such as the liability of developers and the existence of morality in nonhuman actors - with a focus here on technological objects - need a response or, at least, reflections that contribute to the debate in the public sphere.

For this analysis, we will focus on advanced algorithms with *machine learning*, and on robots equipped with artificial intelligence, considering that they are technical artifacts (Things) attached to sociotechnical systems with a greater potential for autonomy (based largely on the processing of Big Data) and unpredictability.

The implementation of programs capable of "learning" to perform functions that relate to people creates new ethical and regulatory challenges, since it increases the possibility of obtaining results other than those intended or even totally unexpected. This is because, as previously argued, these mechanisms also act as agents in society, and end up influencing the environment around them, even though they are non-human elements. It is not, therefore, a matter of thinking only about the "use" and "repair" of new technologies, but mainly about the proper ethical orientation for their development[9].

In addition, the more adaptable the artificial intelligence programs become, the more unpredictable are their actions, bringing new risks. This makes it necessary for developers of this type of program to be more aware of the ethical responsibilities involved in this activity. The Code of Ethics of the Association for Computing Machinery indicates that professionals in the field should develop "comprehensive and thorough assessments of computer systems and their impacts, including the analysis of possible risks".

The ability to amass experiences and learn from massive data processing, coupled with the ability to act independently and make choices autonomously can be considered preconditions for damages liability. However, since Artificial Intelligence is not recognized today as a subject of law, it cannot be held individually liable for the potential damage it may cause. In this sense, according to Article 12 of the United Nations Convention on the Use of Electronic Communications in International Contracts, a person (natural or an entity) on behalf of whom a program was created must, ultimately, be liable for any action generated by the machine. This reasoning is based on the notion that a tool has no will of its own.

On the other hand, in the case of damage caused by acts of an artificial intelligence, another type of responsibility is the one that makes an analogy with the responsibility attributed to the parents by the actions of their children (*strict vicarious liability*). Thus, adopting the theory of "robots as tools", the responsibility for the acts of an AI could fall on its producer, users or their programmers, responsible for their "training".

Should an act of an Artificial Intelligence cause damages by reason of deceit or negligence, manufacturing defect or design failure as a result of poor programming, existing liability rules would most often indicate the "fault" of its creators.

However, it is often not easy to know how these programs come to their conclusions or even lead to unexpected and possibly unpleasant consequences. This harmful potential is especially dangerous in the use of Artificial Intelligence programs that rely on *machine learning* mechanisms, in which the very nature of the *software* involves the intention of developing an action that is not predictable, and which will only be determined from the data and events with which the program comes into contact.

As the behavior of an AI is not totally predictable, and its behavior is the result of the interaction between several human and nonhuman agents that make up the sociotechnical system and even of *self-learning* processes, it can be extremely

difficult to determine the *causal nexus*[10] between the damage caused and the action of a human being or legal entity.

According to the legal framework we have today, this can lead to a situation of "distributed irresponsibility" (the name attributed in the present work to refer to the possible effect resulting from the lack of identification of the causal nexus between the agent's conduct and the damage caused) among the different actors involved in the process. This will occur mainly when the damage transpires within a complex sociotechnical system, in which the liability of the intelligent Thing itself, or of a natural or legal person, will not be obvious[11].

The ideal regulatory scenario would guide the development of the technical artifacts and manage it from a perspective of fundamental rights protection. But no reliable answers have yet been found on how to deal with the potential damages that may arise due to programming errors, or even due to *machine learning* processes that end up incorporating undesired conducts into the behavior of the machine that were not predicted by developers. Therefore, establishing minimum ethical foundations for regulating purposes is just as important as developing these new technologies.

When dealing with Artificial Intelligence, it is essential to promote an extensive debate about the ethical guidelines that should guide the construction of these machines. After all, there is a strong growth of this segment of scientific research, regulatory scenario included. However, clear parameters of how to conduct this study, from the point of view of ethics, has yet to be defined. The need to establish a *regulatory framework* for this type of technology has been highlighted by some initiatives.

The General Data Protection Regulation in Europe (GDPR) already established important guidelines concerning, for example, data collection storage and privacy, setting key principles, such as: *Purpose Limitation, Data Minimisation, Storage Limitation, Integrity and Confidentiality (security) and Accountability.*

On the other hand, a conference held in January 2017 in Asilomar[12], CA, aimed to establish the definitions of a series of principles so that the development of Artificial Intelligence programs can be beneficial. Twenty three principles were indicated, the most notable among them are:

1. Race Avoidance: Teams developing AI systems should actively cooperate to avoid corner-cutting on safety standards;
2. Safety: AI systems should be safe and secure throughout their operational

lifetime, and verifiably so where applicable and feasible;
3. Failure Transparency: If an AI system causes harm, it should be possible to ascertain why;
4. Responsibility: Designers and builders of advanced AI systems are stakeholders in the moral implications of their use, misuse, and actions, with a responsibility and opportunity to shape those implications;
5. Risks: Risks posed by AI systems, especially catastrophic or existential risks, must be subject to planning and mitigation efforts commensurate with their expected impact.

Designers and builders of advanced AI systems are considered stakeholders in the moral implications of their use, misuse, and actions of the Thing and its damaging autonomous actions, with a responsibility and opportunity to shape these implications.

Additionally, there should also be considered responsibility/liability of the designer the concern in guaranteeing values such as privacy, safety and ethics in the design of the artifacts. This aims to avoid problems to *a posteriori*, always taking into account what is within the sphere of control and influence of the designer. Hence, the challenge of thinking, therefore, of a "value-sensitive design". As an example, we can mention the commands of: "*privacy by design*", "*security by design*" and, "*ethics by design*".

From a legal standpoint, it is fundamental to keep in mind the new nature of a control and diffuse liability, potentially dispersed in space, time and agency of the various actants in the public sphere. We need to think about the context in which assumptions on liability are made. The question that is presented to us is not only how to make computational agents liable, but how to reasonably apply the mentioned liability. We must, therefore, think of a "shared liability" between the different actors working in the sociotechnical network and their

spheres of control and influence over the presented situations and the other agents, which makes necessary a whole new interpretation for the the role of law in this context.

Footnotes

1. Big Data is an evolving term that presents any amount of accumulated, semi-structured or unstructured data that has the potential to be exploited for information.
2. The projections for the impact of this scenario of hyperconnection in the economy are also impressive. There are researches that estimate that in 2020,

the number of interconnected objects will reach 25 billion, and could reach 50 billion. Available at: http://www.zdnet.com/article/25-billion-connected-devices-by-2020-to-build-the-Internet-of-things/.
3. SAURWEIN, Florian; JUST, Natascha; LATZER, Michael. Governance of algorithms: options and limitations. Info, v. 17, n. 6, p. 35-49, 2015.
4. DONEDA, Danilo; ALMEIDA, Virgilio A. F. What Is Algorithm Governance? IEEE Internet Computing, v. 20, p. 60, 2016.
5. SAURWEIN, Florian; JUST, Natascha; LATZER, Michael. Governance of algorithms: options and limitations. Info, v. 17, n. 6, p. 37, 2015.
6. WOLF, Marty, et al. Why We Should Have Seen That Coming: Comments on Microsoft's tay "Experiment", and Wider Implications. 2017. Available at: http://digitalcommons.sacredheart.edu/computersci_fac/102/. Accessed on 27 September 2017.
7. 'Causal nexus' is the link between the agent's conduct and the result produced by it. "Examining the causal nexus determines what were the conducts, be them positive or negative, gave rise to the result provided by law. Thus, to say that someone has caused a certain fact, it is necessary to establish a connection between the conduct and the result generated, that is, to verify if the action or omission stemmed from the result caused. Available at: https://www.jusbrasil.com.br/topicos/291656/nexo-causal. Accessed on 27 September 2017.
8. Available at: http://unesdoc.unesco.org/images/0025/002539/253952E.pdf. Accessed on: 27 September 2017.
9. Available at: https://futureoflife.org/ai-principles/. Accessed on: 25 May 2017.

Ronaldo Lemos (@lemos_ronaldo) is one of the Co-founders and a Director of the Institute for Technology & Society of Rio de Janeiro (ITS Rio). PhD in Law from the University of São Paulo (USP), Master's Degree in Law from Harvard University, and studied Law in an Undergraduate level also at USP. He is a Law professor at the Rio de Janeiro State University (UERJ) and a visiting researcher at MIT Lab. He was a visiting professor at Princeton University, affiliated with the Information Technology Policy Center. Ronaldo was also a visiting professor at the Oxford University (Michaelmans term, 2005). He us the Director of the Creative Commons project in Brazil and Co-founder of the project Overmundo, which won the Golden Nica in the Digital Communications category. He is a member of the Social Communications Council, created by article 224 of the Brazilian Constitution, with headquarters in the Federal Senate. He is a Liaison Officer at MIT Media Lab for Brazil and member of the Administration Council of Mozilla Foundation.

Eduardo Magrani (@eduardomagrani) is Coordinator of the Institute for Internet and Society of Rio de Janeiro (ITS Rio). PhD. Senior Fellow at the Alexander von Humboldt Institute for Internet and Society in Berlin. Eduardo Magrani has been working with public policy, Internet regulation and Intellectual Property since 2008. He is Professor of Law and Technology and Intellectual Property at FGV Law School, UERJ, IBMEC and PUC-Rio. Researcher and Project Leader at FGV in the Center for Technology & Society (2010-2017). Author of the books "Digital Rights: Latin America and the Caribbean" (2017), "Internet of Things" (2017) and "Connected Democracy" (2014) in which he discusses the ways and challenges to improve the democratic system through technology. Associated Researcher at the Law Schools Global League and Member of the Global Network of Internet & Society Research Centers. Ph.D. and Master of Philosophy (M.Phil.) in Constitutional Law at Pontifical Catholic University of Rio de Janeiro with a thesis on Internet of Things and Artificial Intelligence Regulation through the lenses of Privacy Protection and Ethics. Lawyer, acting actively on Digital Rights, Corporate Law and Intellectual Property fields. Magrani has been strongly engaged in the discussions about Internet regulation and was one of the developers of Brazil's first comprehensive Internet legislation: the Brazilian Civil Rights Framework for the Internet ("Marco Civil da Internet"). He is coordinator of Creative Commons Brazil and the Digital Rights: Latin America and the Caribbean Project, alongside with prestigious Latin American organizations.

Learning to avoid users infantilization

by Prof. Dr. Gaia Scagnetti

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

Twelve weeks ago we had a baby; by the day she was born my partner and I had already determined that we would not post any pictures on social media to protect her privacy. For us it was an easy decision, like disregarding the wipe warmer, choosing only gender-neutral colors, and eating more organic food during breastfeeding. Our friends believed that being parents would transform us into obsessive oversharers of baby pictures and eventually change our relationship with social media. Instead, having a baby turned out to be a reflection on the smart home, voice commanded apps, and data.

I need a robotic smart home

I have always preferred to interact with my environment in a quiet and reserved way: I favor text messages over phone calls, I choose to read a map rather than ask for directions. I did not develop an effective relationship with Siri or Cortana or any dictation app: initially because my Italian accent confused any voice recognition software, but ultimately because I would rather not vocalize my activities; I like to keep them for myself.

When the baby arrived, something changed. I found myself spending hours nursing, rocking, changing diapers, and expressing milk. What all these activities have in common is that they required the use of both hands.

For those of you who had never fed a baby, I'll give you a quick rundown of the principal modern peculiarities of this activity. My house is quiet; the baby is sleeping. All of a sudden the baby is hungry and starts fussing. There are only about sixty seconds before she starts screaming at a pitch that can break windows (metaphorically) and wake up the entire neighborhood (literally). In that minute window, I pick up the baby, sit, and start nursing. A newborn can eat for five to twenty minutes. While sitting in the quiet again, I suddenly realize my needs; they emerge as the itching feeling during a meditation session. I would love to have a pillow to rest on under my elbow, maybe one behind my back, I want to search the internet, reach my phone, text a friend, do anything that keeps me awake. But both of my hands are occupied supporting the baby. The phone is twenty inches away but I can't grab it, the computer might be just in front of me but I cannot type.

For the first time in my life, I want to be able to command everything by voice. While I am stuck supporting a hungry infant trying to grasp my mobile phone with my foot, I regret not having trained the Google assistant to understand me and perform activities when I speak. During the hours spent staring at something out of reach, I dream of a smart home where every single thing responds to my voice command and where objects communicate with each other. A house where the sensor of my baby's onesie would turn on the bottle warmer when she starts fussing for food, where the bottle warmer would select the right container based on the date and time that my breast pump recorded, where my phone did not need my touch to record how long the baby slept, nursed, and when her diaper needed a change.

For weeks the Internet of Things seemed like a great idea. It might increase our quality of life!

This period did not last long, and already at ten weeks the baby could support her neck decently and I mastered the one-hand-football-hold. I also realized that even if my house was a high tech responsive environment as the one I helped design as a researcher in MIT a long time ago[13], I would have probably just been able to make an order on Amazon but not get it out of the box, start the laundry but not loading the washing machine, order food but not open the door. What I needed was more of a robotic house than a smart one. With a voice controlled device as Alexa or Google home I could have played Jeopardy[14], booked a restaurant[15], stayed mindful with meditation[16] and ordered, bought, and purchased as much as I want to. All the data I could record about me and my baby would be shared with companies to customize my ads, rather than to other devices to facilitate my activities.

The motivation behind the present development of our home technologies is deeply rooted in capitalist objectives more than in the desire of increasing the quality of life.

Technology should serve its users rather than the interest of manufacturing companies; in reality, the user is more often an element of the system rather than its beneficiary, the system is designed to persuade the user to perform specific actions, almost entirely of consumption.

The fundamental conversation we should have is about trust in the companies who are designing these devices. If Alexa was a real person working for Amazon – helping you around the house but also reporting back to her employer all you do and say – would you trust her? How would being listened continuously by someone who works and reports to a company feel? Would we be more comfortable with an independent OS for IoT? Could we DIY an autonomous device with no capitalist purpose?

I need full surveillance

Furthermore, this parenting experience changed my relationship with data. Since the day I delivered at the hospital we were encouraged to keep a log of the activities of our baby. How much she ate (ml per bottle) for how long she nursed (minutes at the breast), how long she slept (hours), how many wet and dirty diapers, and which shade of color their content. We measured height, weight and head circumference. This log helps new parents learn their baby's behavior and pattern match it with the average. It allows spotting when something is wrong and reassures that everything is in the norm. The early days of parenting are made of data.

For the first time, I thought that a total surveillance through data collection was a great idea. It might lead to significant discoveries in human behaviors!

This period did not last long, because the data we collect are useful only when secondary to the parents' intuition. An over-reliance on them disrupts the parents' ability to listen to a baby needs. The normative average baby is anyway an illusion: "every baby is different" we are told over and over again by wise nurses. When using a paper log we are also keeping all the information for ourselves, the same way our simple fridge is not communicating with any other device. The data we produce stays with us. Collecting detailed data about my baby does not feel particularly problematic: it does not pose the same challenges of collecting data

about an adult. Data collection can be a serviceable activity when the object of observation it is naturally subjected to the power of who owns the data. The subjection of the infant to the parents is proper, the subjection of an adult to a company is not. An infant is not yet able to survive autonomously and does not have freedom, the survival of an adult depends on her autonomy and freedom.

Do not treat me like an infant

Infantilization through technology is a thought-provoking framework to discuss the design of a responsible IoT. The concept of infantilization has been described as a treat of the postmodern adult[17], [18] by many[19], [20], [21], [22]: Baudrillard[23] describes Disneyland as the archetype of this world, a metaphor of an American society where the cult of youth is used by capitalism to "infantilize the consumer as a means of non-aggressive control."[24]

A responsible Internet of Things avoids infantilization. An infant is always at the center of the world and needs to be continuously heard and monitored by her carers. The infant gets fed, washed and changed, gets put to sleep and dressed: all trivial tasks are handled by someone else so she can fully dedicate her energy to growth. Consumption is the primary activity of the infant. The infant does not understand the system around her. She is not aware of how and why her world works and does not need to know: stories are designed to explain reality.

When we design artifacts with the assumption that we need a superior entity to make the right decision for us, we are infantilizing our determination. A technology that creates the illusion that we are of central importance treats us as children unable to understand the vast scale of our society. When we develop digital services and use data as a currency to access them, we are infantilizing and objectifying our users. Technological innovations delegating trivial tasks to free our time to for personal growth are trivializing our identities. When we hide the complexity of our systems behind over-simplified interfaces, we are paternalistic.

A responsible Internet of Things should be human-centric, where human refers to an adult with fully autonomous will, identities and rights.

...

Footnotes

1. Scagnetti, Gaia, and Federico Casalegno. "Social Sustainability in Design: The Window as an Interface for Social Interaction." In Cross-Cultural Design, 321–

30. Lecture Notes in Computer Science. Springer, Cham, 2014. doi:10.1007/978-3-319-07308-8_31.
2. Sony Pictures Television. "Jeopardy!" Amazon.com. Accessed July 19, 2018. https://www.amazon.com/Sony-Pictures-Television-Jeopardy/dp/B019G0M2WS.
3. Google Assistant. "OpenTable." Assistant.google.com. Accessed July 19, 2018. https://assistant.google.com/services/a/uid/000000e768ed4de9?hl=en-US.
4. Stoked Skills LLC. "Amazon.com: Mindful Meditation: Alexa Skills." Amazon.com. Accessed July 19, 2018. https://www.amazon.com/Stoked-Skills-LLC-Mindful-Meditation/dp/B0784GXF2M/ref=lp_14284837011_1_5?s=digital-skills&ie=UTF8&qid=1531937136&sr=8-5.
5. Bernardini, Jacopo. "The Infantilization of the Postmodern Adult and the Figure of Kidult." Postmodern Openings/Deschideri Postmoderne 5, no. 2 (2014).
6. Bernardini, Jacopo. "The Role of Marketing in the Infantilization of the Postmodern Adult." Fast Capitalism, Search Results The University of Texas at Arlington, 10, no. 1 (2013). https://www.uta.edu/huma/agger/fastcapitalism/10_1/bernardini10_1.html.
7. McHugh, Molly. "How Technology Is Creating a Generation of Adult Babies." The Ringer, January 25, 2018, sec. Tech. https://www.theringer.com/tech/2018/1/25/16933668/2018-ces-rocking-bed-infantalization.
8. Cain, Benjamin. "The Ironies of Modern Progress and Infantilization (by Ben Cain)." Three Pound Brain, February 4, 2014. https://rsbakker.wordpress.com/2014/02/04/the-ironies-of-modern-progress-and-infantilization-by-ben-cain/.
9. Elkus, Adam. "The Infantilizing Nature of Technophobia: A Matter of Will." Medium, April 20, 2015. https://medium.com/strategies-of-the-artificial/the-infantilizing-nature-of-technophobia-a-matter-of-will-813ef97efab9.
10. Singer, Natasha. "Technology That Prods You to Take Action, Not Just Collect Data." The New York Times, December 21, 2017, sec. Technology. https://www.nytimes.com/2015/04/19/technology/technology-that-prods-you-to-take-action-not-just-collect-data.html.
11. Baudrillard, Jean. America. Verso, 1989.
12. Silva, Erick da. "The Infantilization of Society and the Cult of Youth." The Ivory Tower, September 4, 2015. https://the-ivory-tower.com/the-infantilization-of-society-and-the-cult-of-youth/.

Dr. Gaia Scagnetti is program co-coordinator and full-time Assistant Professor at the Pratt Institute's Graduate Communications Design department in New York. Her current research projects focus on new pedagogies and decolonization for higher education in design. In 2010 Gaia Scagnetti completed a Post Doctoral research at the Design Lab at the Massachusetts Institute of Technology. In 2009 she obtained a PhD degree cum Meritus in Industrial Design and Multimedia Communication at the Politecnico di Milano. During the doctoral research Gaia has worked as a designer and researcher at the Density Design Lab in Milan, where she carried out research, design projects and teaching activities on Information Visualization and Mapping. Her works have been featured in several conferences and exhibitions and publications and showcases. Her complete portfolio can be found at www.namedgaia.com.

Where Does the Responsibility Lie: A Case Study

By Holly Robbins

on behalf of the Just Things Foundation

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

There is a growing appetite for data-intensive and internet-connected technologies, often referred to as Internet of Things (IoT) technologies, to be cast as responsible in terms of how they relate to us as individuals and in terms of their larger role within society. There is a growing discomfort for what *could* go wrong, and what has gone wrong, with technologies with these capabilities. But what is to be done?

There are a few routes to approach framing these technologies as responsible. First is to address how they are regulated, on the level of legislation and policy. The second concerns how they are designed: what exactly are the capabilities that these technologies are designed to have; are the intentions of the designers and entrepreneurs responsible? Relatedly, the third approach considers how people interface or work together with these technologies, and visa versa: how these technologies interface or work together with people; do people make use of the technology in responsible ways; is it possible that the artificial intelligence of the technology develops irresponsible tendencies? While each route is critical to address and opens up very large areas for discussion, this essay is concerned with the ways that these last two routes are entangled with one another. Specifically, **how can design support the ways that people understand what these technologies do, and the role that they, as users or consumers, play in helping them do what they were designed to do.**

How can we step beyond the conventions and expectations of how these technologies are designed to contribute to reframing these technologies as responsible? This essay's scope targets the relation between design and the users of that technology as pivotal; and likewise turns to challenge design conventions by exploring their alternatives with design itself. This essay will discuss and unpack one particular conceptual design to consider the role that design can have in this effort of framing these technologies as responsible. This design, of an IoT charging station for eclectic cars known as the "Transparent Charging Station," helps to guide our questioning of where does the responsibility reside when framing IoT technologies as responsible. In this case, it appears to be in how the design makes the work that the technology does insightful for the people using it, as well as the opportunities that people have to act upon that insight.

Haunted by the IoT

IoT products and services work in ways that are not always apparent to us. Their sleek exteriors (if they even have any) do not reveal the inner workings of these technologies, specifically how they connect, communicate, exchange, and evaluate data within a network of other data-intensive and connected technologies. Conventional design practices favor hiding the complexity that comes with connectivity; and on one hand, this is sensible. It would be overwhelming to be troubled with all the particularities and technicalities of how these technologies work. It would be far easier if the technology functioned the way it was designed to and we as users of these technologies could just enjoy and consume the technology. Technology is intended to make our life easier, more efficient, to limit or overcome barriers. Yet on the other hand, with IoT technologies, we may be unwittingly participating in a network that can take some liberties in terms of our relation with these technologies. For example: a chair that we once merely sit upon could now, as an IoT-enabled chair, detect a change in our weight and start broadcasting to advertisers that we are a prime audience to target diet plans at, resulting in a web browser inundated with ads reflecting the chair's assessment of how we occupy it. *The audacity of that chair!*

Instead, what if it was apparent how the technology responded to our use, how it connected, communicated, exchanged, and evaluated data within a network of other data-intensive and connected technologies?

Instead of being haunted by IoT technologies, where they determine how to engage with us in ways that are invisible and inscrutable, what if we could see the ways that we work with them and how they likewise work with us?

In this essay, we turn to the conceptual design of the Transparent Charging Station to provoke the very norms and conventions of design that promote this inscrutable haunting.

The Transparent Charging Station: Making the IoT Insightful

The Transparent Charging Station (TSC), designed by The Incredible Machine, speaks to this very provocation of making the way that people and IoT technologies work together more insightful (see left figure). The TSC addresses one truism of IoT technologies, which is that there is a complex network of other technologies being tapped into to contribute to how the technology works, and that our use of the technology impacts that network's very functioning. The TSC embraces this truism and is dedicated to making it legible and insightful to the consumer or user of its services, as opposed to the convention to make such operations invisible. Commissioned by a Dutch energy company and sister company specializing in "smart charging" infrastructure (Alliander and Elaad respectively), this IoT electric car charging station offers an alternative to the convention of inscrutable haunting.

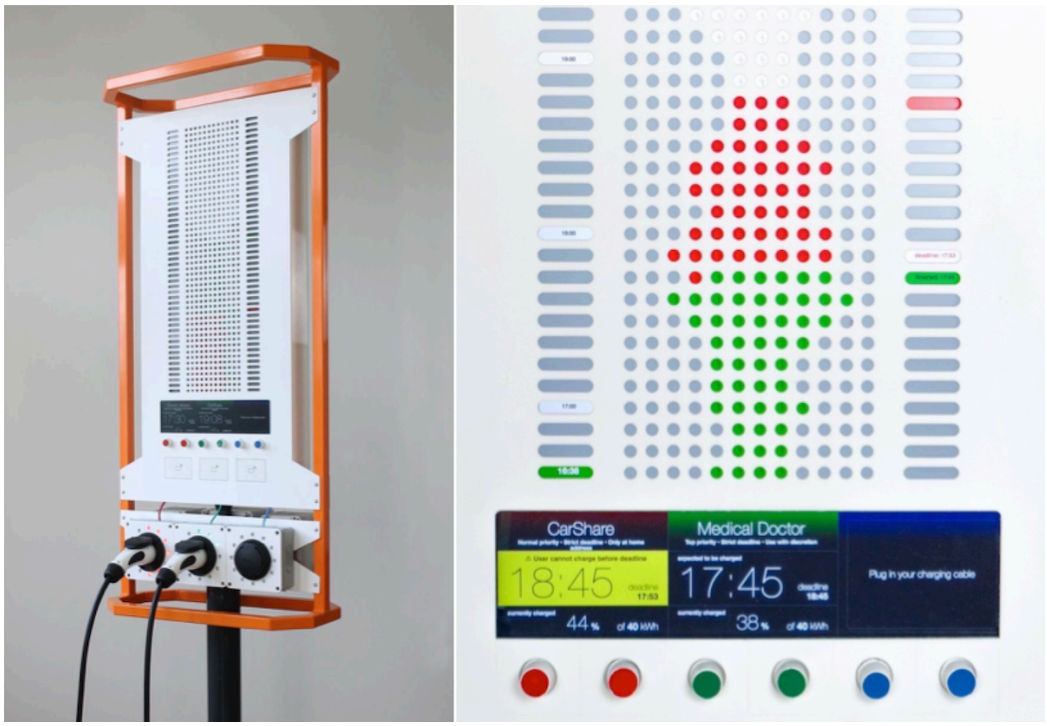

(Left) The Transparent Charging Station is an electric car charging station that allows people to negotiate how much of their battery is to be charged and within what time period according to the networked constraints of an electric energy grid.

(Right) The interface of the Transparent Charging Station demonstrates the constraints and demands on the electric grid with a Tetris-like screen. The figure in the middle illustrate what energy is available on the grid over the course of the day. By turning the dials of your port (red, green, or blue), a driver can negotiate within the constraints of what is available on the grid, and the demands other patrons are making of that station. Design and images: The Incredible Machine.

The TCS is a speculative design developed in anticipation of the widespread prevalence of electric cars. The reality behind the infrastructure of electric charging will radically challenge existing practices and expectations surrounding how we fuel cars. Car batteries require more time to recharge than it takes to fill a car with petrol at the station. We've also grown accustomed to the fact that, under normal circumstances, a petrol station will always have a reservoir of fuel available on demand. However, with an electric fueling infrastructure, the availability of electric energy at the "pump" will fluctuate in response to a number of factors: what is the demand on the electric grid at any particular moment; the availability of renewable resources; the weather; and what's already currently stored. There are a network of factors that will influence how and when the car can possibly be charged.

As a result of the constraints of this electric infrastructure, namely the fluidity of the resource and the network it relies upon, there are two significant design hurdles. First how to prioritize fueling protocols. The system will not be able to accommodate the influx of cars being charged after rush hour, and therefore decisions will have to be made about how to prioritize requests. The second design challenge considers how to conceptualize and make insightful to users the fluctuating availability and networked qualities behind this charging station and resource.

To address these particular design challenges, the TSC features a design and interaction where people can negotiate with the algorithms behind the charging infrastructure. In their interaction, people determine how much of their battery needs to be filled and within what timeframe. How, or even if, these requests will be fulfilled will vary based upon how the request impacts the other the constraints posed by the other nodes in the network. For example, a request for a full battery charge within an hour may not be possible if there is a high demand on the system on a cloudy day. Perhaps if there was more solar energy being contributed to the grid that day it may be able to accommodate that request.

Each station has three ports from which energy is dispensed for three different vehicles. To communicate how the fluidity of the resource is impacted by our energy demand, the charging station features a large interface that resembles a

Tetris game (figure on right). This interface contains an outline of an irregular shape that represents the boundaries of what is available from the electric grid at that particular charging station, and what it predicts to be available over time. Within that outline there are colored blocks that represent the request being made from the other ports of that station.

The person who has come to charge their car has two dials which they can rotate to indicate how much charge they will need (15%, 85%, 100%, etc.) and within what time frame (1 hour, 8 hours, 24 hours, etc.). The board will modify the Tetris-like configuration in accordance with these requests as they relate to what's available to that station and what are the demands of the others currently utilizing that station. The demand I make on one of the station's ports will impact that of the other user who is also charging from that station. If I ask for a full and rapid charge from the station, it will draw energy from the other car's charging arrangement.

In this dynamic interactional exchange that one has with the station, we negotiate with the algorithmic constraints that govern this system, as well as with the network itself as a whole.

The TCS makes the dynamic among the energy grid, the station, and the person insightful. It does so by making explicit what factors are taking into account by the charging algorithm, and making experiential how they affect the projected distribution of resources. In this case, the TCS looks at the available energy and the stress (amount of energy needed before deadline) of each of the patrons, and their privilege (some car owners require a full charge at all times, such as emergency responders). The ability to play with the parameters empowers citizen to scrutinize the algorithm and be better informed about how they are being treated, as well as to decide how they want to navigate the system itself.

Where Does Responsibility Reside?

The TCS has been well received at various industry and research venues. Recently it was even awarded the very prestigious Dutch Design Award under the product category. These accolades attest to the fact that this design offers us a very plausible future for IoT technologies. So plausible in fact that even a city is exploring how to deploy this conceptual design within their municipality.

The station provides us with a lens to address the question: from what position should we frame the responsibility to reside within the IoT? This essay opened by proposing that there are three routes to addressing this:

legislatively, in design, and in the exchanges between people and the technology. The scope of this essay sought to examine how can design support the ways that people understand what IoT technologies do, and the role that they, as users or consumers, play in helping them do what they were designed to do. **With the TSC we find that responsibility is framed in terms of making the way the technology works "transparent," and in the opportunities that people have to be autonomous agents within this system.**

The TCS makes an interaction out of something that was formally made invisible, to make it insightful for the people using it. Algorithms tend to be designed and trained to promote predefined optimal outcomes. In this case, we find that the TCS does not address the intention behind the algorithm, but the design language proposed by TCS could be used to make the difference in different intentions transparent to the people using the technology. With this design, responsibility is being framed in terms of having available for scrutiny how our interactions with, and demands of, the technology draw on and likewise impact the network behind it. It isn't just a question of making these particularities of the technology available for scrutiny, but in a form that is insightful for people to be able to be able to develop a perspective on. Further, this insight and perspective can be put to use with the opportunities that are made possible for the user to exercise their autonomy and agency, such as with negotiating the charging request with the station.

We also find some nuance with the TCS regarding how the concept of "transparency" relates to responsibility. Insight into the dynamics of the station aren't literal or explicit. If the station was explicit about all the work it was doing and the connections it was making, it would likely be difficult for the layperson to process and make sense of. This would hinder our ability to decipher and scrutinize the technology.

The TCS cannot be the perfect or complete solution for all our questions regarding how to create a culture of responsible IoT technologies; but this design does offer us some direction in provoking potentially problematic design conventions. We can develop a vocabulary through examples such as the TCS to advance our agenda of responsible IoT design. Let us remember, this particular project started as a conceptual design, was then recognized as a product, and will potentially be implemented in a municipality. This is a radical trajectory for a conceptual design. There are more questions that the TCS will surface the more it is developed and made use of, but they are likely to be exactly the questions we need to be asking at this time.

Part of what makes the TCS so special is its acknowledgment of what about the infrastructure of electric energy is unique. Rather than attempting to create a design or system that assimilates those qualities into design convention (of hiding the complexity), these qualities are explored to consider what opportunities lie within them to frame them as more responsible. We need more work to follow this suit.

Holly Robbins is a postdoc at Delft University of Technology (Netherlands) where her research is in values and ethics of technology. She focuses on internet-connected and data-intensive technologies. She is also a co-author of the IoT Manifesto and a co-founding board member of the Just Things Foundation.

Full disclosure: Marcel Schouwenaar and Harm van Beek are both partners in The Incredible Machine as well as co-founders and board members of the Just Things Foundation.

More-than-Human Design for the Future of AI in the Home

By Iohanna Nicenboim, Prof. Dr. Elisa Giaccardi, Dr. James Pierce

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

Inhabited by smart devices and intelligent assistants, our future home will be certainly more-than-human. Expanding to almost every fabric of our everyday future, the Internet of Things (IoT) and Artificial Intelligence (AI) promise new exciting possibilities for designers. But they also surface new anxieties (see for example the projects Network Anxieties and Objects of Research).

Existing anxieties such as privacy and democracy become even more prominent when IoT is combined with AI. AI is seemingly everywhere, but the term actually means different things in different contexts. Beyond the humanized and romanticized 'he' or 'she' that we see in science fiction movies and television, AI comes in many forms, and it is already underlying many of the products and services we use today, from social media to public services. Even though we interact with AI every day, its complexity and opacity makes extremely difficult for people to 'see' it and grasp its benefits and pitfalls. This situation leaves us uncertain on whether we need to protect ourselves against the 'entity' AI or rather against the people building, training, and operating it.

Understanding and critically evaluating AI is difficult also for designers and researchers. Complex AI agents often exhibit emergent behaviors that are impossible to predict with precision, even by their own programmers. MIT researchers explain that to evaluate AI algorithms it is not enough to simply look at their source code or internal architecture. For this reason, they recently proposed a new field called Machine Behaviour, the scientific study of machines not as

engineering artifacts, but as a new class of actors with their unique behavioral patterns and ecology. This new field overlaps with computer science and robotics, but it is different, because it treats machine behavior observationally and experimentally.

When we start looking at algorithms from an anthropological perspective, we begin to see that they are "unstable objects that are enacted through the varied practices in which we engage with them" (Seaver 2017; Giaccardi et al. 2016). This is interesting in IoT, because intelligent algorithms are performed by everyday objects, which might be fundamentally different than other enactments of AI. The way AI is enacted or performed by everyday objects can foreground certain issues while occluding others. For this reason, Nick Seaver proposes that critical researchers should research algorithms ethnographically, seeing them as heterogeneous and diffuse sociotechnical systems, rather than rigidly constrained and procedural formulas. To do so, he suggests thinking of algorithms as part of broad patterns of meaning and practice that can be engaged with empirically.

So when it comes to design, we too can begin to observe smart objects ethnographically to evaluate the future of IoT + AI. This means to research AI not just as code or behaviour, but as performed by everyday objects in the context of mundane practices, within the messy ecologies of our homes. But design can do more than understanding algorithms ethnographically. Design can help us imagine our more-than-human home, and the role algorithms will play in that future.

For this type of inquiry, traditional design methods such as human-centered design might be insufficient. In the new domestic landscape of IoT+AI, not only people will interact with objects, but also objects with each other. To better understand these complex ecologies, we need to include also the perspective of things, and actively enlist them as partners in the design process (Giaccardi 2018). Thing-centered design is a novel design approach that gives designers access to fields and trajectories normally unattainable to human observation. This ethnographic engagement is called thing ethnography, and it is usually applied to existing things.

To research ethnographically future things instead, and help us imagine the future of AI in the home, we have combined thing ethnography and future-oriented design techniques. Future-oriented techniques, such as design fiction, make future paradigms of technology more tangible, and develop critical discourse on the impact technologies might have on individual lives and society at large. Designers

often use fictional techniques to project and evaluate the encounters people may have with a technology, and the actions and decisions that people may take in response.

Thing-Centered Design meets Design Fiction

When taking a thing-centered approach in design fiction work, we not only access future perspectives of humans, we also gain access to the nonhuman perspective of things. These new perspectives can "enhance, complicate, and sometimes even challenge the perspective of humans" (Giaccardi 2018).

We have explored these possibilities in two projects: "Affective Things: Entanglements of the Connected Home"; and "Unpredictable Things: Objects that Withdraw".

Affective Things: Entanglements of the Connected Home

This project is a series of design fictions in the form of videos that explore the complex interactions of things with things in the more-than-human home of the future. By positioning everyday objects within complex ecologies, these works show how things may become entangled with us and with each other, and how they might co-perform tasks. In this work, smartness is explored as fluid and things have the possibility to become 'other things' as this smartness is shared and gains meaning through interaction.

Unpredictable Things: Objects that Withdraw

This project investigates the boundary of what algorithms can see (and recognize) and what they cannot see, as a productive design space for resilience against surveillance at home. The work explores how things could hide from different cameras by altering their materials and shapes. It does so by proposing different strategies to make objects unique, and thus impossible to be captured by object recognition: from a virus for digital fabrication codes to a home lab to create diversity. The design process in this project was done by co-designing with things

themselves, for example, by letting living organisms reshape everyday objects in unique and unpredictable ways, or by crafting together with a machine by looking from its view to make design decisions.

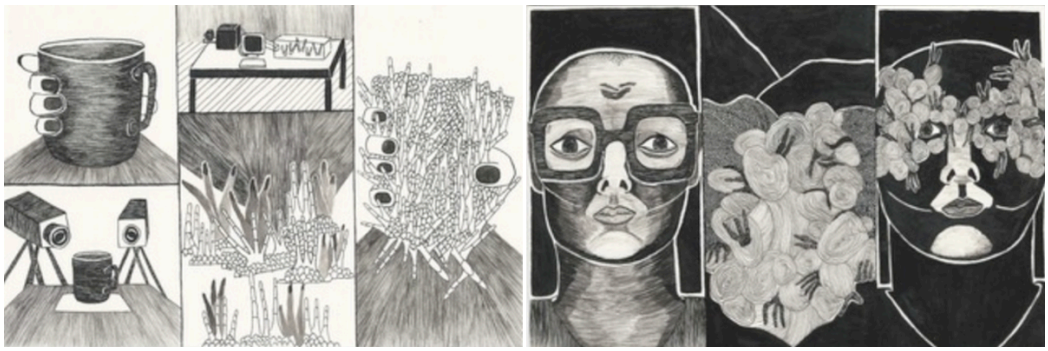

Implications for IoT design

Taking a thing-centered approach in design fiction work can help designers explore a future everyday from a novel perspective and gain unique insights that might be difficult to obtain with traditional design methods. Conducting a thing ethnography by means of a speculative prototype can shed light on the ecologies that may configure around a future thing, including what meaningful data it might collect and how people would react to it. It can also help designers to finally imagine interactions that are less animistic, avoiding the tendency to excessively anthropomorphize or zoomorphize machines.

A thing perspective offers a fundamentally different role for things in design beyond their functional use. As argued in Giaccardi 2018, it helps us "cast things in design as partners, overcome our human biases, problematize our design space and

possibly be more humble in our worldview". For example in a more recent project, Connected Resources, a thing-centered approach helped us to reframe and imagine smartness as something shared between people and things, instead of an exclusive property of artefacts or humans.

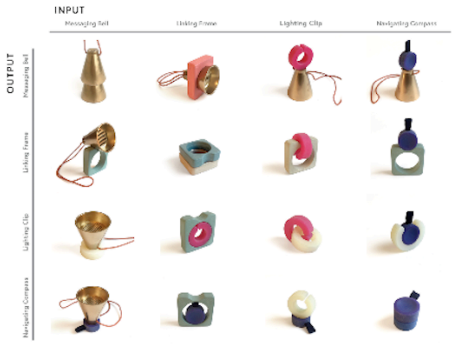

Connected Resources

A family of recombinant sensors for older people, designed to emulate in physical form and digital functionality the material affordance of the mundane objects used by older people in their everyday strategies of resourcefulness.

In summary, combining future-oriented techniques and a thing-centered approach can contribute to understand how algorithms will be enacted by smart objects in everyday futures, positioning things as agents within complex socio-cultural ecologies. This can help us reframe the design space and inspire more ethical and resourceful approaches, where smartness is not exclusive but shared.

Credits

Affective Things: In collaboration with The Incredible Machine. Photos Andreas Dhollandere. Thanks to Design United.

Unpredictable Things: In collaboration with Daniel Suarez. Photos by Bart van Overbeeke. Drawings by Alexandra Sebag. Thanks to Everyday Futures Network.

Connected Resources: Graduation Project by Masako Kitazaki. Photos by Andreas Dhollandere. Part of 'Resourceful Ageing' funded by STW under the Research through Design program (2015/16734/STW).

References

Giaccardi, Elisa. 2018. "Things Making Things How Things Will Help Us Design the Internet of Reinvented Things." IEEE Pervasive Computing, 2018.

Giaccardi, Elisa, Chris Speed, Nazli Cila, and Melissa L. Caldwell. 2016. "Things As Co-Ethnographers: Implications of a Thing Perspective for Design and Anthropology." In Design Anthropological Futures, edited by Rachel Charlotte Smith, Kasper Tang Vangkilde, Mette Gislev Kjaersgaard, Ton Otto, Joachim Halse, and Thomas Binder. Bloomsbury Academic.

Seaver, Nick. 2017. "Algorithms as Culture: Some Tactics for the Ethnography of Algorithmic Systems." Big Data & Society 4 (2): 2053951717738104.

Iohanna Nicenboim is a design researcher at the Connected Everyday Lab, TU Delft. Her research focuses on designing for the IoT as part of complex socio-technical systems in everyday futures. Her background brings together industrial design with digital media, and currently focuses on IoT and AI. Following a thing centered speculative approach, her designs often use Design Fiction to provoke reflections on more desirable futures. She received the Internet of Things Award for the Best Design Fiction project in 2015-16, and is a Thingscon IoT fellow since 2017. She participated in residency programs, gave talks and exhibited her work in different international exhibitions and conferences, like CHI, DIS, FutureEverything, Transmediale, Milan Design Week and Dutch Design Week.

Elisa Giaccardi, Ph.D. is Professor and Chair at Delft University of Technology, where she is director of the Connected Everyday Lab. After conducting pioneering work in metadesign, collaborative and open design processes, Elisa has during the last years focused on the challenges that a permeating digitalization means for the field of design. Her recent research engages with 'things' in new ways, with the starting point that these now hold both perception and possible agency (e.g., AI), and thus 'participate' in design and use in ways that previous industrially produced objects could not. Her online course "Thing-centered design" can be found on the TU Delft Online Learning platform.

James Pierce is an Assistant Professor at California College of the Arts and research scientist at the University of California Berkeley. As a design researcher, he practices and reflects upon design as a mode of inquiry, critical engagement, and speculative exploration. His recent research focuses on issues of digital privacy, security, and surveillance.

Responsibility in IoT: What does it mean to "do good"?

By Prof. Dr. Irina Shklovksi

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

> *"The door refused to open. It said, "Five cents, please."*
>
> *He searched his pockets. No more coins; nothing.*
>
> *"I'll pay you tomorrow," he told the door.*
>
> *Again it remained locked tight. "What I pay you," he informed it, "is in the nature of a gratuity; I don't have to pay you."*
>
> *"I think otherwise," the door said. "Look in the purchase contract you signed when you bought this conapt."*
>
> *...he found the contract. Sure enough; payment to his door for opening and shutting constituted a mandatory fee. Not a tip.*
>
> *"You discover I'm right," the door said. It sounded smug."*
>
> *— From Ubik, by Philip K. Dick. Published by Doubleday in 1969*

Philip K. Dick had an uncanny sense of the possibility of technologies and of their potential impact. I find this excerpt eery in how uncomfortably close to our current technological realities it is, even if it was written nearly 50 years ago. There are so many new devices that are rapidly coming on the market that are smart, connected and wanting to help. You may have heard about the Amazon Echo personal

assistant named Alexa that constantly listens to its environment and can play music, report weather or even order products directly from Amazon on command. This device has been in the news relatively frequently as it first took to ordering doll houses at children's command and cookies on its own initiative, then proceeded to laugh at random, startling its owners and even recorded snippets of conversation and sent these to recipients randomly selected from the contact list. There is the Google Home personal assistant that can do much the same thing as Amazon's Echo but when two such devices were put together they got into some deep arguments about the nature of the universe. There are smaller much more specific objects as well - internet connected toothbrushes, TV's, hairbrushes, mirrors and, of course, there are many smart locks. These do not yet require a payment every time they open the door, but they can be hacked by enterprising hackers, broken by a simple 'dumb' screwdriver or suddenly made inoperable by an errant firmware update.

Complex consumer-oriented IoT devices such as Google and Amazon home assistants or Samsung mobile phones are suddenly implicated in sending unexpected types of data to unexpected recipients with increasing frequency. Whether clever hacks or results of unexpectedly buggy software, such discoveries are invariably troubling and creepy, prompting efforts to reverse-engineer ways to check what our own devices might "have on us". Of course, the smart printers, toothbrushes and TVs are not far behind in constantly phoning home and reporting on their users. The problem is that when physical environments become instrumented with all manner of smart sensors and devices, the flows of data and decisions about its collection and use become ever-more invisible to the end-user. In this situation the onus of morality shifts further towards the developer and designer of the technology in question because they get to unilaterally decide the rights and wrongs of device behavior. The use of these devices assumes and requires increasing amounts of trust from the end-user. Perhaps it is possible to hold the end-user responsible for the decision to put an Echo device into their kitchen or for purchasing a smart tooth-brush, but such arguments can only go so far. After all, in some places it is practically impossible to buy a "dumb" TV these days and who has the time to really pay attention and read all of the interminable end- user license agreements and privacy policies? Such expectations are intractable.

Connected devices are entering our homes, our lives and ever more intimate spaces of bedrooms, bathrooms and boudoirs, collecting data about incredibly private moments and consumers are asked to trust these opaque systems with the data they collect. Their usefulness is sometimes questionable and sometimes undeniable, but the crucial question is: Who is responsible for the behavior of

these devices? There is no playbook, no rules of conduct – technology developers and designers at large companies and those working to disrupt and innovate, entrepreneurs, makers, hackers – are charged with making moral choices and are expected to get it "right". The EU GDPR has set out a number of guidelines and each member country has a complicated web of laws and policies with regards to data, but as we all know laws are complex and yet limited. Besides can we really expect entrepreneurs, developers, designers or innovators working in maker and hacker spaces to be able to navigate this complex web where even the lawyers get tangled?

Many designers and developers in startups as well as in mature companies are struggling because, as one developer explained to me during an IoT meet up in Copenhagen: "we don't yet have much of an idea of what is ok and not ok." In other words, the decision making about what is "good and responsible behavior" does not yet have real precedents or pre-existing experience to guide it. Even if IoT developers are attempting to be responsible, what constitutes responsibility is not yet agreed upon. Clearly communities of IoT developers must come together in collaboration with their stakeholders to develop ideas about what responsibility means in this context. What relations must be considered, what obligations must be taken on and enacted are important decisions precisely because building new systems requires acknowledgment and renegotiation of interrelations of responsibilities. At the same time the shifting standards and new regulations continuously shape and structure what sorts of decisions might be made. Who gets to make these decisions and whose values might guide these are also pertinent questions. In a globalized economy, the notion of "good" does not work as a local concept and yet "good" is always contextual, so who is responsible for moments when "good" pivots and takes on negative consequences? If nobody can predict the future, is it actually worth trying? This stuff is complicated and what constitutes responsible action does not have a clear answer.

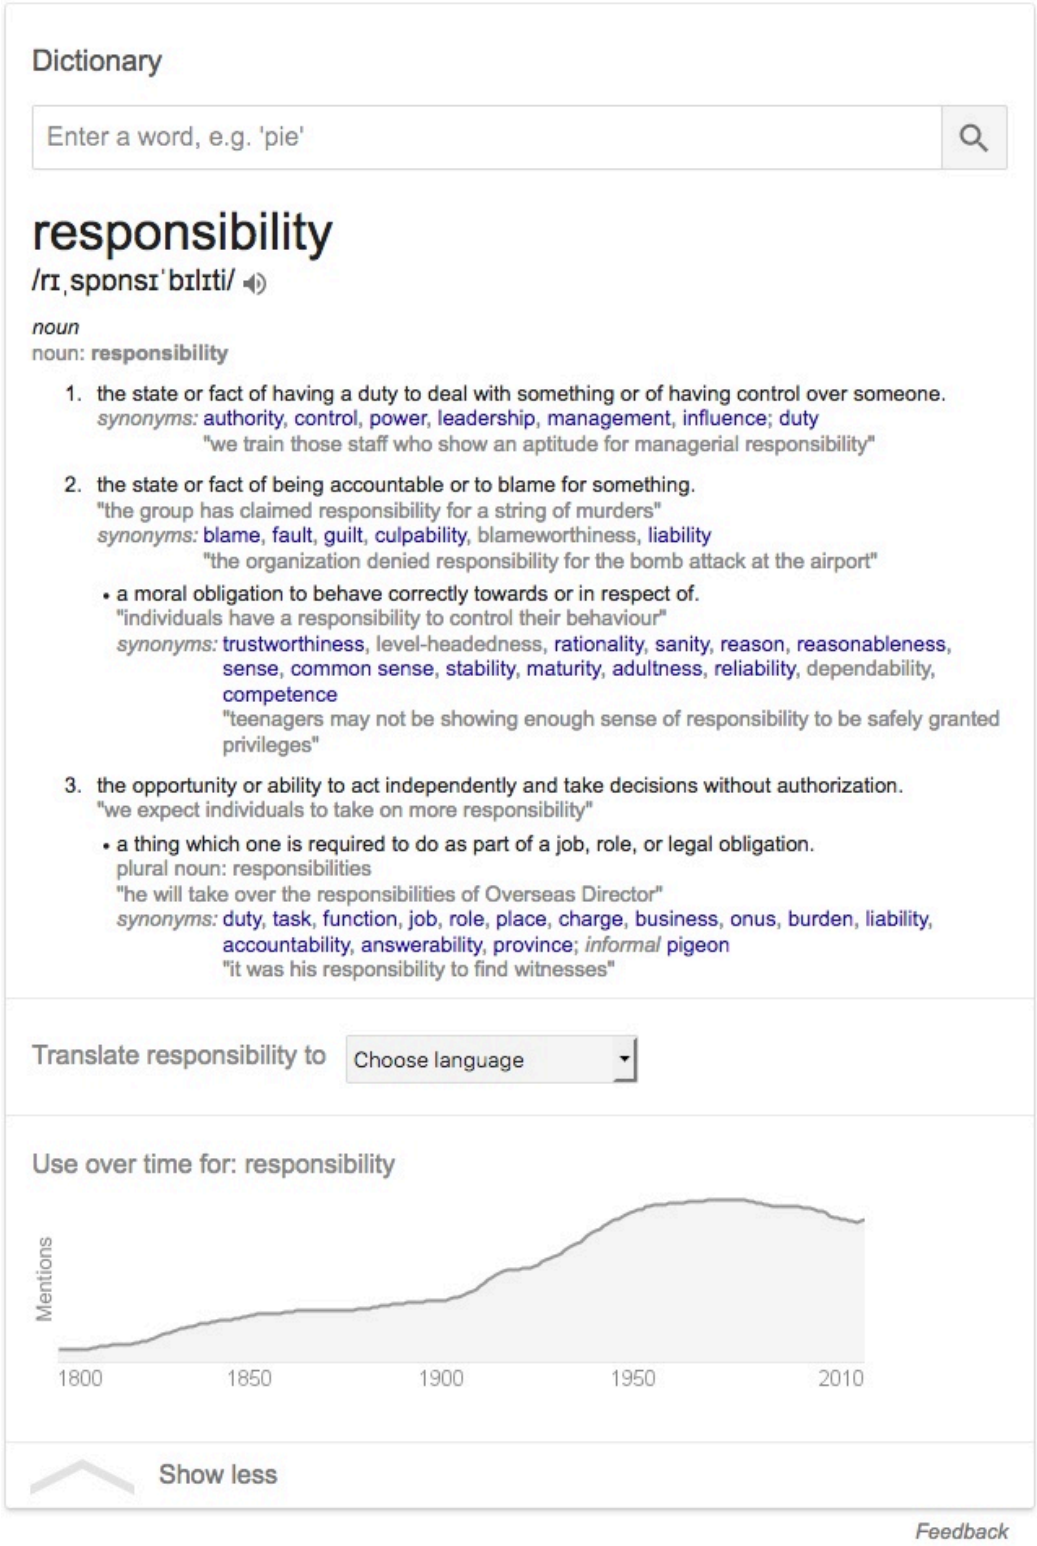

Google offers many definitions of the term responsibility.

When asked to define the term responsibility, Google produces many definitions. This variety of definitions makes one thing clear - as individuals, all of us are enmeshed in a variety of different interdependencies - responsibilities to many others that are sometimes complimentary and sometimes at odds and must be negotiated. We are responsible in different ways and for different things to our families, friends, neighbors, workplaces, institutional arrangements in which we take part, the state and even global communities of many kinds. I am responsible to the editors of this RIOT collection for producing this essay by the deadline. At the same time, I owe my family time and attention in these summer months, I owe my friends some thought in absence, I am obligated to my employer to respond to email and I must also submit reports and deliverables to the European Union. I have made promises that I must fulfill and more often than not these obligations clash in their demands on my rather finite time and other resources. All promises and obligations are inter-related and at times competing; their fulfillment is a balancing act. What's more, many of these varying types of responsibilities are reciprocal - my commitment to spend time with friends or family is moot unless they make that commitment as well.

Individuals are always entangled in a diversity of relationships that hold contradictory values and conflicting demands. For example, collaboration is seen and acknowledged as an important value among the IoT community (e.g. open access software, off-the-shelf hardware). At the same time, for many startups, the pressures of ‘making it' in the ever more competitive IoT market push people to focus on ‘survival' thus privileging some collaborative relationships over others and perhaps even eschewing relationships that previously held significant sway. So how might these notions of responsibility be translated with respect to IoT? Who must be responsible, for what, how and why?

Technologies in general and IoT technologies in particular are, of course, not neutral. They embody and reflect their designers' values and ideas of what counts as "good" or "responsible." After all, if a smart lock or a digital home assistant is intended to improve people's lives, the design of these technologies is driven by someone's idea of what counts as "improvement." Among the calls for building technologies responsibly and for doing good, what does it mean to "do good"? The Dutch ethnographer and philosopher Annemarie Mol says that "It is important to do good, to make life better than it would otherwise have been. But what it is to do good, what leads to a better life, is not given before the act. It has to be established along the way." Importantly, it does not mean that every developer or designer must focus on figuring this out for themselves, separately from others. No matter the emphasis on personal improvement, perhaps it is time to acknowledge that we are never separate individuals, but are instead composed of our many

memberships, relationships and social entanglements that span our lives. We might want to hold those responsible for the design choices as accountable for their positive and negative outcomes. Maybe the engineer responsible for the "like" button on Facebook is worried about addictive behaviors or the designer who developed "pull to refresh" behavior is appalled at how it has been used and feels personally responsible. I would like to propose, however, that feeling guilty for these outcomes is not going to get us anywhere useful.

I have no real use for guilt. Instead, let's acknowledge the problems and try again. This, I think, is a way around the paralyzing realizations of downright apocalyptic possibilities of IoT that my colleagues and I have previously observed in our analysis of IoT manifestos. If calling for being responsible, let's reflect of what we mean by responsibility and consider who ought to be responsible for what, how and why. Being responsible individually is often lauded as an ideal, but that's one lonely mountain-top and I think responsibility ought to be taken on together as groups and communities. If what constitutes "good" needs to be established along the way, then it needs to be established together. One way to allow for this deliberation along the way is to design with legal scholar Julie Cohen's idea of "semantic discontinuity - the opposite of seamlessness" - a call for strategically under-designing technologies in order to allow spaces for experimentation and play. Such intentional building in of flexibility may be one way to offer possibilities for alternatives, for seeking out what a "good life" ought to look like with IoT.

Irina Shklovski is an Associate Professor at the IT University of Copenhagen. Although for her primary field as human computer interaction, her work spans a lot of other fields from computer science to sociology and science & technology studies. Irina's research focuses on big data, information privacy, social networks and relational practice. Her projects address online information disclosure, data leakage on mobile devices and the sense of powerlessness people experience in the face of massive personal data collection. She is very much concerned with how everyday technologies are becoming increasingly "creepy" and how people come to normalize and ignore those feelings of discomfort. To that end she recently launched a "Daily Creepy" Tumblr to monitor the latest in creepy technology. She leads an EU-funded collaborative project VIRT-EU, examining how IoT developers enact ethics in practice in order to co-design interventions into the IoT development process to support ethical reflection on data and privacy in the EU context.

Things As Citizens in the future Cities of Things

By Iskander Smit

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

Arizona law gives delivery robots same rights as pedestrians – but they must abide by same rules; we read on the website of Fox News end of May. It triggers my attention as it seems to be an example of the new relation we will have with things in our city. Things that are more intelligent, have agency, things that live on the streets of our cities as a new kind of creatures, even like citizens.

There is a lot to say on these new relations with things. With non-human artifacts, or should we say post-human as more and more often is mentioned. This is the subject of new research program we started in 2017 at Delft University of Technology, called PACT: Partnerships in Cities of Things. The research program combines different topics, like the development of Things that will become intelligent and can interact with each other based on algorithms. An example is to illustrate is this movie Affective Things based on the research of Iohanna Nicenboim and Elisa Giaccardi.

Watch on Vimeo.

The delivery bots that are mentioned in the example of Arizona are rather serious. There are several companies developing variants, the one of Starship Technologies seems the most market-ready. Started in 2014 by the Skype co-founders they just received a serious 25 million investment round.

The third wave of the smart city; Cities of Things

The evolution of the smart city has more layers. For the kick-off of PACT Elisa Giaccardi distinguished three waves of the smart city.

The first wave - that is still ongoing - is the city as a dashboard. Sensors are added to the city to capture data of the all kinds and that data is served back to the user of the city. First to the makers of applications, the governments and if we are lucky also to the citizens. This is what the current debate is all about mainly. It is what Martijn de Waal divides into three types of Smart Cities; The Control Room with a focus on economic values, The Creative City as a innovation lens, and the The Smart Citizen where the city becomes a political and civic community (De Waal, 2017). In this wave we improve the city on the angle of the citizen mainly.

The second wave focuses on the smart city as an intelligent infrastructure. Starting as a sensing city that is used to nudge citizens in a certain behaviour. Next is the infrastructure that will become more and more adaptive to the collective and individual behaviour. Examples are lighting systems that adapts. But also a public

transport systems based on autonomous moving vehicles that personalise the routes based on the travellers. MIT is doing interesting research here in the Senseable City Lab. The angle is the city as infrastructure. In projects like Minimum Fleet where they look how the number of taxies can be reduced in Manhattan by adding autonomous vehicles, and with the Roboat project in Amsterdam, where the boats can form pop-up bridges if needed, autonomous vehicles become part of the intelligent city.

The work of Saskia Sassen on the intelligent city is also inspiring. The city as an autonomous entity, not so much reacting on citizens, but also acting and interacting. The network of actors, both human and non-human getting meaning while acting together, as in best tradition of Latour.

The third wave has still to start. What if the Things in the city we live with become more than Things, and will take a role as citizens. What a city is will not only be defined by the citizens or the representing governments, humans and non-humans will live in concert and shape the city together. That is what we call the City of Things. Autonomous moving objects, things as social entities in the city. Things that cannot be controlled, but like humans, Things can be governed. Humans and non-humans have pact/social contracts to live together in these cities.

We focus in the PACT research on this last phase. We look at Things as social entities, data-enabled artifacts with performing capabilities. These Things connect to existing networks for the necessary data, and combine that with the real-time data it senses. More than now, these things act proactively and behave socially. The PACT research is linked to the research work on co-performance of things and humans, and during the last period we looked especially to the role of Things As Citizens in different (on-going) research.

Co-performance for more-than-human partnerships

As Kuijer and Giaccardi show in their paper is co-performance (Kuijer, 2018) between human and non-human actors crucial. Artifacts have artificial body/minds capable of performing social practices next to people. In the research co-performance important aspect is the 'appropriateness' of human and artificial performances under situated circumstances.

People and things have different capabilities. People are better in judging, machines are better in optimizing and generalizing. Humans are better in improvising. That is why appropriateness is so key in the distributions of tasks. The

way to treat the appropriateness of the interaction between human and artificial things is looking to changing divisions of roles and responsibilities between human and artificial things. Design plays a key role here to delegate task and judgements. Where the appropriateness is defined in the the margins designers leave to the interplay of humans and things in the everyday performance.

Design qualities for Things as Citizens

In her master research on Things As Citizens Louise Hugen found that a model for democratic citizenship could function well to indicate the different aspects divided in obligations and rights, that are important in the relations of citizens in cities (Butts, 1988). She found in her research the values authority, truth, participation, equality and privacy as most significant.

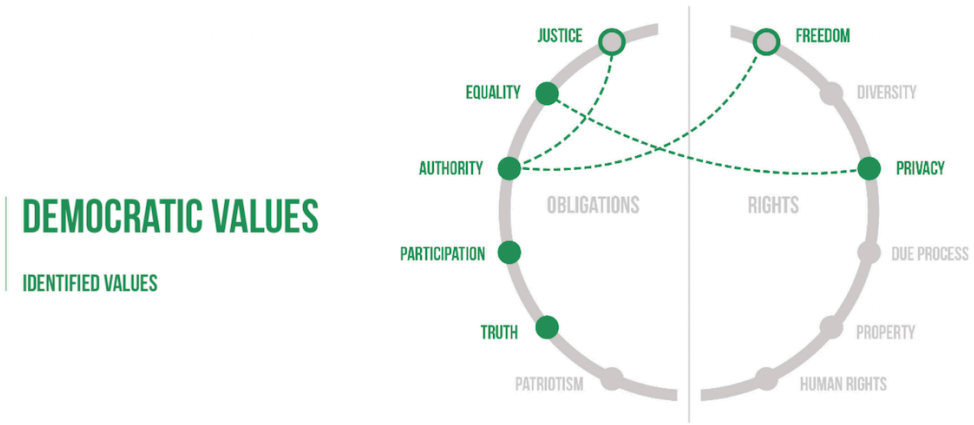

These values are transformed in design qualities for the Things as Citizens. Hugen created a model where thing capabilities and citizen requirements meet. The design qualities can be used in the creation of the new city things. In her study she focused on one demonstrator, more explorations are needed to test the design qualities. It seems that these six are the most in the core:

Based on the value of Truth - a thing is able to promote its sensitivities and is able to show the lack of ethical sensitivities - a citizen is able to understand decisions made by Things - a thing and a citizen can bridge differences in perceptions and intentions.

Based on the value of Authority - a thing is able to react in different ways to be partially directable in its actions - a citizen has provisioned space for negotiability - a thing can enter a continuous agreement

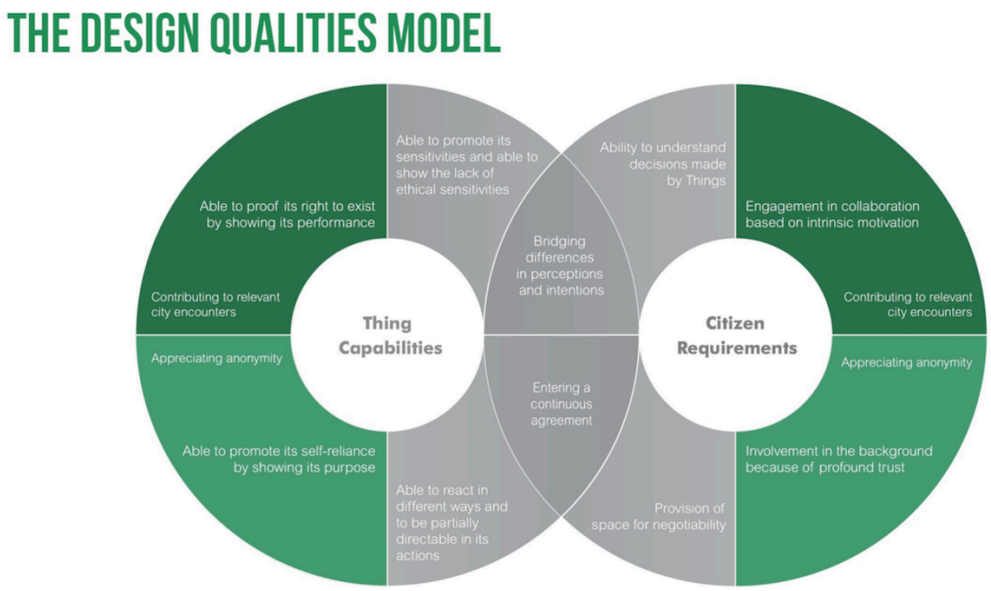

Hugen used the design qualities in her project to design a autonomous delivery pod combined with air purification service, and designed specific thing-thing and thing-human interactions to stimulate communication between things and citizens. How city things behave is a key element in having humans and non-humans living together in the cities of things.

Using ideation to learn on Things as Citizens

To research how people think on the new relations with things as citizens Maria Luce Lupetti, postdoc researcher in PACT, created TaCIT; the Things as Citizens Ideation Toolkit. The aim, as Lupetti formulates it: identifying, collaboratively, a set of topics representing the main opportunities and challenges related to the design and development of autonomous things for the city.

We tested the toolkit in sessions at TU Delft, and at public conferences like Border Sessions and The Next Web. One of the key elements in the toolkit is to look at five dilemmas; - Responsibility; private to public, who is responsible for the behaviour of the autonomous thing - Priority; who rules? The human or the system, or the

thing? - Relationship; is the aim to behave social or antisocial - Adaptation; aspects of adapting behaviour, presence, from the standpoint of the human or the thing. - Delegation; are tasks partially or totally delegated to the system?

In the workshop different briefs are made for delivery pods and last mile vehicles that vary in the usage scenario and regulation scenario. This triggers ideation. Reflections are than done from different roles: government, industry and citizen. What does that make a difference. Via a debate and cluster session the challenges and opportunities are filtered.

The ideation process during a PACT Workshop on autonomous things and government regulation at the Border Sessions Festival in June. Picture: Ashlee Valdes.

The workshops taught the participants to take different perspectives, especially the thing-perspective and their role in the city.

We will continue the PACT research as there are still open questions to address. How much do we prone to accept and adapt to things that behave out of our control? And how can we design the appropriate interplays between human and things. What are the morphologies, the non-verbal behaviour, the interaction schemas for the Things of Citizens?

Looking back at the case of the delivery bots in Arizona. Not long after the news on the rights it turned out that people start kicking the delivery pods. It is a form of antipathy Starship Technology (the manufacturer) thinks. It is also a process however to learn to live the new citizens in the city. New rituals will grow on us. And on the Things, in the form of the design rules.

Illustrative is the way people like to game self-driving cars by jumping in front of them. Humble as the cars are they will stop for humans. But too humble might lead to a standstill of society: We need new manners that respect humans and things in the future living together in the Cities of Things.

In the design we need a certain character, but also a language to communicate intentions. And you can imagine that the interactions are not always the same. In the design of self-driving cars like Daimler is research conducted to try to recognise the intentions of pedestrians in the visual recognition system. Is the intention to pass the street or just walking to the side of the pavement. If these conclusions are combined with a learning systems the car might start to recognise types of behaviour.

Which triggers the question what that might deliver. Are all cars rating the pedestrians based on predictive models, categorising future behaviour? Will the car combine older data with newer data to make better judgements? Is this a bias we like to tolerate?

PACT is still in progress and we will continue the research. The dilemmas sketched in the TaCIT workshop are helping us to discuss how much we accept and adept to things that behave out of our control. We hope to develop insights to design for appropriate interplays between human and things.

Iskander Smit is educated as Industrial Design Engineer and worked as creative and strategist in digital services since 1994. He has a longtime track record for designing and thinking on the internet of things. Since 2009 Iskander is member of Council Internet of Things and in 2013 co-founder of the Behavior Design AMS meetup and co-organising Tech Solidarity NL.

In 2014 he initiated and co-organised the Amsterdam edition of Berlin conference ThingsCon, a leading conference on the design of the internet of things, that will be organised for the fifth time in 2018.

Iskander Smit is now innovation director at agency Info.nl in Amsterdam, that crafts connected digital products and services. Iskander is responsible for R&D and leading labs.info.nl. Since 2017 he is appointed as visiting professor at Connected Everyday Lab at Delft University of Technology, where he coordinates the research program PACT (Partnerships in Cities of Things) and the City of Things Delft Design Lab.

Responsible and trustworthy IOT

By Dr. Laura James

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

Why do we want responsible IOT?

For the same reasons we want responsible technology in general – so that we get technology that is useful, has benefits definitely outweighing harms, that we can rely on. We look to the developers and operators of technology to act responsibly, so that we can have confidence in their products and services. Whilst it's easy to get carried away thinking that this is just an issue for the big silicon valley companies, or about personal information, the issues we face with irresponsible technology go beyond these things. This is not a new concern, but one I have been reflecting and working on more recently.

As I wrote last year:

> *Ten years ago I was at AlertMe, architecting and developing an internet of things system. We were within 6 months of a shipping product — in January 2008 people we didn't know were buying AlertMe systems online, and receiving boxed kits ready to install. AlertMe set out to create broadband home security, bringing burglar alarms into the internet age, and redesigning them to be rather more useful to householders. We made a hub to connect to your router, and used a secure ZigBee mesh network to link up a mixture of detectors and buttons, both mobile and static devices. We thought about how the consumer would buy and set up the kit, did low power radio for key exchange, designed for hardware sale on the second hand market, considered the different people in a household and their data and privacy needs. We did*

user research and user testing and field testing and got independent security experts to review our architecture, and we followed standards where they existed (and contributed to development where they didn't). It was quite a lot of work, but seemed like the least we should do for a connected home product, especially as we were thinking about the security market (we also thought about future extensions which would use the AlertMe platform, such as energy monitoring and control, which also has security and privacy requirements). It took two years to go from a three-word brief to a shipping product (with customer support and sales channels and all the rest of it), which felt simultaneously very fast and frustratingly slow, as it seemed as if the market for such things was about to take off.

Ten years on, there's nothing much like that available for mainstream consumers. The IOT products we see for the home and for individuals more generally are mostly simple things — perhaps one device, connecting via wifi to the internet or via bluetooth to some other device such as a phone. Mesh networks are rare, security holes seem disturbingly common, embedded systems hold personal data and forget to get rid of it when they are sold, connected home systems (with very few exceptions) seem to have forgotten that houses have many different people in them. Not only are the products made much simpler, but they often don't seem to have thought through what even ten years ago were reasonably obvious privacy and security basics. What happened? Did we oversell the "anyone can make hardware" idea? (When I've said "hardware is hard," I didn't just mean the manufacturing bit :)

Even though all this was possible a decade ago, somehow there is still work to do in driving up standards, and making it easier and more valuable to make secure and trustworthy systems.

Only a few months after I wrote this, the AlertMe servers (now owned by British Gas) were turned off, disabling the systems of those of us still using it for home monitoring. Requests for the service to be passed to the community of users were not responded to.

I think we did a good job of producing trustworthy technology at AlertMe – by the standards of development 10-12 years ago. Undoubtedly we would do things differently now – different technologies, different values, different regulations. I'd like to think that contributed to the longevity of the service, which was exceptional. And yet how odd it is to celebrate ten years of service, when the vision and functionality of so many IOT devices is designed to form part of the fabric of our homes, our workplaces, our public spaces. Things which are designed for many

decades, not one. We still have a lot to learn about architecting and creating digital systems that can last, still operational and secure, and be maintained and adapted for the lifetimes we might reasonably expect. The challenges here are partly technical, and partly about business model – having organisation, capability and capacity to tend to and repair these systems. Even if we opted for a more dynamic society, where our environments change more often, we would need to think about the energy and materials used in replacing our connected infrastructures and things.

(I'd like to learn of other examples of consumer IOT which has operated for over a decade, and which is of greater complexity than a Bluetooth device that connects to a phone.)

Aside from that personal example, the last year has seen some interesting bits of news about the state of responsible IOT. Tesla's attitude to information after an accident; the black market of smart agricultural equipment repair; fitness trackers and privacy as a feminist issue. In all of the recent focus on individual user experience design, we seem to have forgotten the need for adversarial thinking (although not all IOT devices will be at risk of unlikely hacks, such as sound waves to fool accelerometers), or to consider devices which are used and interacted with by more than one person. Perhaps it's just because there's more IOT in the field now, or more visibility of issues when they arise, but it feels like the level of responsibility in practice is getting worse, not better.

Luckily there are more people working on this critical issue now.

A year ago, I wrote an outline of ten aspects of responsible technology for Doteveryone. It includes a mix of areas very specific to digital with broader business responsibilities related to tech. Responsible practice isn't just about the technology itself—it's about the people who develop, manage and invest in tech, about the users, individually and collectively, and the wider context. Responsibility is more than simply applying an encryption standard, or complying with GDPR – it is integrated through everything an organisation making a tech product.

In responsible technology:

1. The **business model, ownership and control of the organisation is responsible** and appropriate for what the organisation does and the products/services made
2. Every worker, including suppliers and subcontractors, is **paid fairly and has good working conditions** in an inclusive environment

3. The people, communities, projects and businesses contributing effort or information to the organisation are **rewarded fairly**
4. The organisation's products and services make a **positive (or neutral) contribution to public and societal value**
5. Risks, systems effects, side effects, **potential harms and unintended consequences have been considered**, both for the organisation overall and for the products/services – including for potential acquisition, or market dominance
6. **Plans for maintenance and support of products/services into the future**, including clarity on how long support and updates will be available for and what happens when they stop, have been considered and published
7. People can **easily find out and understand** how the product or service works, how it is supported, what costs there may be, and what happens with data.
8. The product/service **follows relevant standards and best practices** - in design, architecture, developing, testing, deploying, maintaining and supporting technology
9. The product or service is **usable and accessible for the range of users** who may need to use it, and appropriate support is provided for them
10. **The wider context around the product/service has been considered** and addressed appropriately, including thinking about the people who may encounter the service and their lives, the environment and sustainability in terms of energy and materials.

All of these apply well to IOT, as well as to broader technologies in general. It's not always helpful to silo specific technologies when thinking about ethics or responsibility – people encounter holistic products and services and their effects on the wider world. IOT is both a buzzword (perhaps an aging one now, but still alive), and a big tent, encompassing a huge range of different kinds of tools, toys, infrastructure, and more. Coming up with actionable principles that apply across this range is not easy.

The IOTMark project has been working on codifying responsible concepts over the last year too, and has made some good progress. Many of the above ten things, especially the technical ones, are echoed in the IOTMark principles. Still, the challenge comes when ideals hit reality. Are we setting a gold standard few products will reach, but which articulates our ideals and where we should be aiming? Or a pragmatic one, which will build momentum in the industry and raise the bar a little (perhaps allowing for future adjustment upward later)?

What – or perhaps whose - values do we want to capture anyway?

Is responsible IOT about European cultural values? Or North American, Silicon Valley ones? Or Chinese ones? Or something else, or something more specific? If we are to have practical guidelines for people developing IOT, we need to be able to answer this. There will not be a single global solution. We may think that GDPR sets out how the personal data aspects of IOT should be, and that these European values are good for people everywhere. Disagreement comes more quickly if we consider what fair value exchange for information might look like, or whether it's essential to publish the provenance of all the hardware components, or which software should be open source and under what circumstances, or what acceptable documentation would look like for a mass consumer product some of whose customers may not be highly literate. A well designed, 21st-century mark to help consumers choose IOT products could be valuable, but the IOTMark community have yet to nail down what values it will embody.

Today we hear calls for ethical design in tech and in IOT more than ever. There are many IOT manifestos, and lists of principles (as well as ethical tech ones in general). There are varied technical privacy and security standards, relating to different parts of the IOT or different application areas. These initiatives often very siloed, when IOT is always a cross-cutting endeavour, with decisions about hardware, software, data, application area and users intertwined. We need approaches to responsibility that reflect this, and which support collaborative discussion across the teams making and maintaining products.

A year on, Doteveryone has streamlined the 10 aspects to 3 key ones, which we now champion and are building into an initial toolkit which fits in an agile design process. These reflect our values as a think tank focussed on responsible technology, which we believe will be better for everyone in society. Assuming that the business is responsible already (and there are tools and support systems available to enable a business to sort out its employment, governance, and practices in this way), what are the really critical components of responsibility for digital technologies and IOT?

Context - looking beyond the individual user and taking into account the technology's potential impact and consequences on society

Technology that understands and respects the greater contexts and ecosystems it operates within and the potential impacts - positive, negative or a bit of both - it could have on the institutions, communities and relationships that make up society. This is about deciding on tradeoffs and explaining these to not only the direct stakeholders of your technology but those who might be affected.

Contribution - sharing how value is created in a transparent and understandable way

Determining all of the ways different parties contribute value to a technology product/service; this can include information, formal or informal labour. Then sharing publicly these value flows - about who is involved in them, what is being exchanged - in a clear way that is easy to understand.

Continuity – ensuring ongoing best practice in technology and design, that accounts for real human lives

We should be creating and supporting products and services that are safe, secure and reliable for real, messy human lives and situations. Ensuring people with different needs and abilities who might reasonably use a system are accounted for with inclusive design, and that the technology is suitably supported and maintained. Following appropriate best practices for the specific hardware and software elements of a product, and anticipating and adapting to new needs and threats as they emerge.

We think these are practical and reasonable for today's tech sector to work towards – but they are still principles, requiring thought and effort to put them into practice, not a simple recipe to follow.

From conversations with some tech businesses, simple recipes are wanted. But ethics and responsibility aren't a free lunch – they take hard work. Consider an organisation struggling between good treatment of user information and the business model their investors think will give highest growth, or an individual concerned about practice in their work: knowing what is right is not always obvious, and knowing what to do about it even less so.

It's somewhat easier for organisations who set out to do the right thing from the start; their values, the people they recruit, the customers and investors they target are more aligned. It's not an easy ride, even for them – there's a lot to think about, and practical tools can help even these organisations to make good choices throughout their work.

Still being responsible doesn't demand diving into deep philosophy in most cases. It's often basic common sense, thinking through risks and planning sensibly for whatever you are doing. (If you are designing a lock, think about how malicious people could open it!) Incidents of bad design affect the perception of IOT as a

whole. We need to do better at calling out silly mistakes early, helping each other to build better products, learning together, and encouraging others to be part of the responsible IOT movement.

Because consumer trust in IOT is starting to be threatened by incidents like this, and by fears and misunderstandings about the internet companies who are so pervasive in our lives. This is worrying for those of us who would like to see connected technologies delivering valuable services and benefitting society – if trust is lost, the potential benefits will be diminished.

At an event last year I heard many people in an educated, thoughtful audience express genuine concern that Amazon Echo devices are listening to them all the time, and that Amazon gets that information. I am reasonably confident that this is not the case; but it was a real belief, and one which it is very hard to counter. Part of my confidence comes from knowing some of the developers personally; that does not scale as a route to trust.

Trust is not a wholly rational response to the world. It is necessarily only possible when the trusted thing or person cannot be perfectly known (if we know something entirely, we do not trust it – we have well-founded confidence in it instead.) An IOT system is opaque, complex, impossible for a person (even a deeply technical developer of it) to know completely.

We can't engineer people's trust. That would be manipulative – and sometimes people are right not to trust some technologies.

We *can* engineer trustworthy IOT products, services and systems, which are competently made, reliable, and honest about what they do and how they do it.

> *"Those who want others' trust have to do two things. First, they have to be trustworthy, which requires competence, honesty and reliability. Second, they have to provide intelligible evidence that they are trustworthy, enabling others to judge intelligently where they should place or refuse their trust."*
>
> *— Onora O'Neill*

It is an individual responsibility on each IOT developer, designer, leader.

More than that, changing the landscape – through customers, investors and employees all demanding better IOT development practice – will drive change. It may not be rapid, though.

Creating, using (and for those who commission or buy tech, requiring) frameworks that strongly encourage good practice is helpful too. It doesn't have to be regulation, with all the heavyweight process and slowness that people expect. We can have governance through open technical standards; improved practice sector-wide through templates for better IOT design, tools to evaluate impact and think through risks. We can make it easier to build responsible IOT, and we can research and showcase how being responsible can realise business value, too.

I'm excited by the new technology concepts, such as distributed machine learning which can keep data on the device but still create powerful, actionable IOT insights, and platforms like DataBox or the Hub Of All Things, which change the data dynamics. GDPR is likely to change the IOT data space in some ways – although it's not clear how, yet. More people will produce toolkits which help organisations to think about ethical issues, and to use good design patterns. Conferences and online communities Existing responsible business tools and certification systems like BCorp and Responsible100 are looking to enhance their technology cover, and are being joined by the Zebra movement, platform co-ops and tech co-ops more generally. There are new ways of doing things across business and technology that give me hope for greater responsibility in many IOT systems in coming years.

There's no perfectly responsible, ethical and trustworthy IOT project. There will be compromises and tradeoffs, especially in the tough competitive landscape of consumer 'things'. The landscape will always include shining examples of good practice, and shocking mistakes and malicious products. But individually, and together, we can shift the balance.

I wonder where we'll be ten years from now.

Dr Laura James is Entrepreneur in Residence at the University of Cambridge Computer Laboratory, catalysing multidisciplinary research and activities around trust and technology, and Technology Principal at Doteveryone. She has spent nearly twenty years exploring cutting edge technologies and turning them into useful products and systems, in technology and leadership roles in diverse contexts. Laura has been the first employee at a connected home startup, on the management team of an AI startup, scaled a rapidly growing civic tech nonprofit,

ran mission critical open source systems for a whole university, and cofounded a community workshop, a startup humanitarian NGO, and most recently a member owned co-operative.

IoT - upcoming challenges for digital ethics

by Luca van der Heide

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

Technologies related to and today grouped under the name "Internet of Things" have many times been associated with words like "pervasive" and "invasive". Not particularly appealing words, that suggest a hostile quality of the technology, almost as if we expect it to revolt and take over. If you think about the idea of innovation, that entails the introducing of new ideas as well as new technologies, you may confidently say that it is a very fundamental aspect of this idea to replace the old with the new – meaning that the less innovative cannot stand in the way of the more innovative. One may argue that this feeling of inevitability does not necessarily apply to every technology, as it is true that many innovations remain, to this day, optional. But can the same be said of a technology which has been called pervasive and even ubiquitous? These words certainly recall the idea of inevitability. And if nowadays these words are thought of as a little obsolete, then the name "Internet of Things" itself contains one of the most inevitable aspects of modern life: of course, the Internet.

The fact that these words, "pervasive" and "invasive", are heard so often in this context tells us that there is some degree of uneasiness connected to the omnipresence of this technology in our lives. An uneasiness that cannot but increase when the already inevitable Internet is about to take bold steps out of the digital world and into the realm of things; something that seems to add an extra dimension to the wide range of privacy and autonomy issues that the immense quantity of data produced by the Internet have already brought forward. Certainly if digital ethics is struggling so much to keep the pace with this ever-expanding

world right now, making sure that a newer, more pervasive version of the technology develops responsibly and in a more or less controlled way feels like a rather daunting task.

From the point of view of digital ethics, the challenges we can expect to be facing are all somehow related to this concern of maintaining the human in control of the system or, more precisely, to safeguard the autonomy of individuals in a highly complex, intelligent and autonomous system of objects. So taking into consideration the main peculiarity of the IoT, pervasiveness, we could assume that the most IoT-specific ethical issue is exactly this impossibility to opt out. With this I intend not only a physical impossibility to abstain from engaging with the system, but also and perhaps most importantly a social pressure to conform to certain standards the user might not be willing to conform to. As the technology becomes more and more entangled in our daily lives, we will be required to use it and be familiar with it, whether we want to or not. Failing to do so could mean having less chances than others; with other words, it would create social injustice. Like it happened with Facebook, WhatsApp, or LinkedIn for job opportunities, the person not conforming is left behind or even coerced, to further his personal interests, to participate. Apply this to real, physical environment, and it seems like there won't be much choice for anyone to simply let innovation do its course, and conform to it.

Apart from being pervasive, one of the major characteristic of the IoT – and one of its major appeals as well – is to be unobtrusive, invisible. That is to say, acting in the background of our attention, relieving the user from wasting precious energy and time. Anyone can see the appeal of it: invisible servants demanding none of your attention and working together with true mechanic precision to deliver you from the nuisances of day-to-day chores. What might make someone feel uneasy, however, is the fact of not being part of this process. The idea of invisibility, for how attractive it is in this context, presupposes some degree of unawareness, and a lack of awareness may be easily linked back to a lack of control. For the simple fact that interaction amongst different agents going about their own agenda always presents the possibility of misunderstanding. And how can we expect every user interacting with invisible IoT systems in public spaces to be aware at all times of every implication and consequence of this interaction?

Now, so far we have seen the IoT is described as pervasive, invisible and autonomous. As I mentioned above, another widespread word in the field is invasive – or intrusive. Given that the same technology is expected to be unobtrusive, the use of these concepts might create some confusion. What is intended when using the word "invasive" is, in the vast majority of cases, an

invasion of the private sphere of the individual, a breach of a right to privacy. In the case of pervasive technology, we are mainly interested in a specific kind of privacy that can be called spatial privacy.

IoT systems – as connected to the Internet – will be present both in public and private spaces. And as the Internet will be implemented into physical objects, the issue of protecting our private sphere won't refer only to protecting our personal information in the digital world, but also in the physical world where we act and live. The massive collection of data now happening largely digitally will happen in real spaces, public or private, by means of systems that are made to remain undetected and to constantly communicate with one another.
With spatial privacy is here intended also a privacy of the relation of the individual in space, that is, his location and movements. It is clear that people moving in and across IoT environments will be exposed to tracking by third parties. And when considering unwanted action being taken by someone or something, we also have to consider how and when would consent be given for such actions that might be unwanted – such as data collection. Users need to be reassured of having some kind of control, because the feeling of having no or limited control would most likely entail reticence in giving consent. That is to say, if users cannot always be aware of what the system is doing at the present moment, they must at least be aware of the motivations and purpose of the system, and be sure that such action is in line with previously given consent and respectful of his right to privacy.

If these fundamental rights are to be preserved, we must take a distance from such unforgiving concepts as intrusiveness, inevitability, impossibility to opt out. What I mean is, there must be some kind of leeway for users to keep exercising their autonomy in the system; if not to be able to opt out of it altogether, at least to be given the possibility to choose one's degree of involvement. The system should then be adaptable to the user – we might say, user-friendly. The choice whether to engage or not in the interaction should be of the user, primarily. For this reason it is clear there will have to be ways for the system to "show itself" when needed, so as to allow the user to make an informed decision whether or not to engage with the system, and to what extent.

Key steps for this to be possible is to turn "invisible" into "transparent" and "invasive" into "inclusive". Systems can remain unobtrusive while their intentions being clear and accessible at all times. There has to be an openness regarding the uses of IoT systems and sufficient information has to be provided for users to be aware of what the system is doing while operating "in the background". That consent will be present and informed should never be assumed; the engagement of the user with the technology needs to be active and fully conscious.

At the same time, this information should be of common knowledge, so that users are less likely to feel the threat of misbehaviour or even deceit of the technology. The IoT doesn't need to be thought of as invading the private space of users; it should rather be designed with the purpose of including the user into its processes. Given appropriately transparent systems, users would also feel more in control and have make more conscious decision on their degree of involvement with the technology. The possibility of someone choosing to have less or a different kind of engagement with the technology than others must therefore be taken into account; and if the development of the IoT is accompanied by values of acceptance and promoted from the start as an inclusive technology, social injustice due to different degrees of involvement can be largely limited or even prevented.

A whole range of adequate safety conditions are to be observed when designing such a complex technology like the IoT. To avoid fear and unwanted consequences, it is paramount for possible misbehaviour of interconnected systems to be predictable and preventable, and for the user to be and feel in control at all times. Even more than that, the introduction and development of the IoT needs to be embedded in the right set of values, always aiming at preserving the autonomy of the individual over that of the system and at banishing concepts like "pervasiveness" and "intrusiveness" in favour of "user-friendliness", "transparency" and "inclusivity".

Luca van der Heide is a writer and teacher graduated in Applied Ethics with a focus on contemporary issues in digital ethics. During his MA he has worked for the Rathenau Instituut, contributing to a project for the foundation of an ethics committee for emerging technologies in the Netherlands. Said project was submitted to the Ministry of Internal Affairs and subsequently approved. Working at the Rathenau he had to research in depth current ethical issues in modern technology and specialized on the Internet of Things. He has written his thesis on the problem of moral agency for users and devices in interconnected systems. Luca is currently travelling, writing and teaching English around the world. You can reach him at lucavdheide@gmail.com

A-words: Accountability, Automation, Agency, AI

By Maya Indira Ganesh

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

In this essay I discuss approaches to accountability in human and non-human systems in contexts of error, breakdown and disaster. Machine learning and AI-related technologies are often applied to automated decision-making without an established review or audit process. In many situations they may be applied by people who do not have adequate knowledge of how they work. There are already instances of how these applications fail, resulting in discrimination and bias.

Accountability is a set of practices to understand how disasters, accidents and breakdowns have occurred in technical systems. Opening up a system to see how it works and identifying causes of error and breakdown can also feed into a productive forward-looking process of better design of the system. Accountability practices and approaches require an odd combination of skills: bureaucracy, investigation, and a deep and broad knowledge of how a system works and how it connects to other systems.

My interest in accountability in the AI context is part of an ongoing research project that investigates autonomy: Can you hold something that is not human and is autonomous accountable for something? What is autonomy then? How do various states of autonomy in non-humans and humans in a large, complex technical system result in errors and breakdowns?

In this essay I draw on cultural critiques of algorithmic culture, Science and Technology Studies (STS) ethnographies of infrastructure and disaster, and art and design. I end by suggesting that where accountability seeks clarity about

constituent actors in a system and their interactions to understand how something occurred, not everything can be mapped and known.

Autonomous accounting

Consider the elaborate socio-technical architectures of a semi-autonomous car, a biometric border, a credit scoring algorithm or a lethal autonomous weapon. Each has similar components: data sets, programming architectures, commercial proprietors, workflows, front and back ends, middles, contracts, global trade flows, geopolitics, risk assessments, project managers, clients, suppliers, interfaces and dashboards, lawyers, engineers, local and global regulations, and specific industrial practices and their legacies.

When something goes wrong in such large technical systems, accountability cannot rest with a single individual. "Complex systems are rarely, if ever, the product of single authorship; nor do humans and machines operate in autonomous realms" (Schuppli 2014, 5). Shared and distributed accountability for errors in a complex technical systems is accepted in industries such as aviation (Galison 2014).

Yet, AI is imagined as somehow different. The popular imagination of AI conjures up a machine system that is somehow 'autonomous', atomised, singular, and capable of accounting for its actions, making moral decisions, or decisions in changing and uncertain circumstances. Conveying autonomy in a sense that "fetishizes individuality" (Fisch 2017, 122), AI systems are calibrated as autonomous through constructed measures such as 'ethics' or 'intelligence'.

There is the ambitious imaginary of the fully autonomous vehicle that makes decisions for itself. Martha Poon refers to it as "the perfect neoliberal subject that tootles along making decisions for itself" (Ganesh 2017). This is the kind of object that James Moor refers to in his discussion of 'explicit ethical agents' (2006). This is also the vision we're handed down through cinema and literature: the robotic, autonomous, 'awesome thinking machine' (Natale and Balatore 2017) modelled on humans that can be programmed as a force for good or evil, and makes decisions accordingly. A recent version of this is Eva in Ex Machina that models human cunning, deception and violence in order to survive. While current AI technologies are not at the Eva stage, it is important to acknowledge that the anxieties and drama associated with this new technology are part of its emergence (Bell 2018)

This explicit, self-accounting autonomous machine relies on the notion of 'ethics' which is leveraged variously as a measure, test or outcome: does the machine 'have' ethics? Can it 'do' ethics? The quest for software that makes decisions according to ethical principles has been in the works for some time. Referred to as 'machine ethics' its goal is

> *"to create a machine that's guided by an acceptable ethical principle or set of principles in the decisions it makes about possible courses of action it could take. The behavior of more fully autonomous machines, guided by such an ethical dimension, is likely to be more acceptable in real-world environments than that of machines without such a dimension." (Anderson and Anderson 2006, page 10)*

A 'machine guided by ethical principles' in its decision-making is epitomised by the Trolley Problem as applied to future driverless cars. Anyone listening to a tech podcast or watching a TED Talk in the past few years has probably come across this thought experiment. It has entered mainstream awareness as a prescriptive suggestion for programming an autonomous vehicle to make a complex moral choice (known as 'ethics') about the value of life.

MIT's Moral Machine project is an academic project based on an iteration of the Trolley Problem (Rahwan 2016). In this, the problem is gamified into scenarios involving a driverless car with failed brakes and a series of different human and non-human actors—legally or illegally—crossing at a crosswalk ahead. In some instances the driverless car has passengers. The question is the same: should the driverless car risk the life of someone or something on the crosswalk, or bring damage to itself or its occupants by avoiding them? The online game has generated 40 million responses from 3 million people in 200 countries and territories (Rahwan and Awad 2018). They believe this dataset could be the path to a "universal machine ethics" (ibid).

I read the application of the Trolley Problem in Computer Science projects as a "calculative device" (Amoore and Piotukh 2016) that transforms values about killing and dying into quantifiable metrics; and this is constructed as 'ethics'. I believe that this accrues power to computation and invokes a kind of 'cybernetic control fantasy' that manages risk and produces a futurity in which the outcomes of a crash are foreseen, and then perfectly managed (Coleman 2016). This is a kind of perfect accountability, perhaps.

US 20170285585A1

(19) **United States**

(12) **Patent Application Publication** (10) **Pub. No.: US 2017/0285585 A1**

Weast et al. (43) **Pub. Date: Oct. 5, 2017**

(54) **TECHNOLOGIES FOR RESOLVING MORAL CONFLICTS DURING AUTONOMOUS OPERATION OF A MACHINE**

(71) Applicants: **John C. Weast**, Portland, OR (US); **Tobias M. Kohlenberg**, Portland, OR (US); **Brian D. Johnson**, Portland, OR (US)

(72) Inventors: **John C. Weast**, Portland, OR (US); **Tobias M. Kohlenberg**, Portland, OR (US); **Brian D. Johnson**, Portland, OR (US)

(21) Appl. No.: **15/089,541**

(22) Filed: **Apr. 2, 2016**

Publication Classification

(51) **Int. Cl.**

G05B 13/02	(2006.01)
G06N 99/00	(2006.01)
G05D 1/00	(2006.01)
G06N 5/04	(2006.01)

(52) **U.S. Cl.**

CPC ***G05B 13/028*** (2013.01); ***G06N 5/045*** (2013.01); ***G06N 5/046*** (2013.01); ***G06N 99/005*** (2013.01); ***G05D 1/0088*** (2013.01)

(57) **ABSTRACT**

Technologies for controlling a machine include a compute system configured to control operation of the machine. The compute system is configured to detect a moral conflict related to the operation of the machine and determine operational choices for operation of the machine to resolve the moral conflict. The compute system also determines a moral agent likely to be affected by each operational choice and one or more moral rules applicable to the moral conflict. The moral agents may be weighted based on a set of weighting rules, which may vary based on geographical location and/or other criteria. Each moral rule defines a goal to be achieved by the operation of the machine. The compute system is further configured to select one of the operational choices to resolve the conflict based on the determined moral agents and the moral rules and control the machine to perform the selected operational choice.

Patent for resolution of moral conflicts in autonomous operation of a machine. Weast et al 2017

People are Infrastructure (also)

It is going to be a while till we arrive at a fully autonomous vehicle; two decades possibly. Till then, what is really autonomous? Which human, or machine, is not embedded in a complex chain of actors and relations? Even a future fully autonomous vehicle will be entangled in a dense network of computer vision databases tagged and annotated by humans (something we are already doing by filling in CAPTCHAs), internet infrastructure to connect to the cloud, other vehicles, laws and regulations.

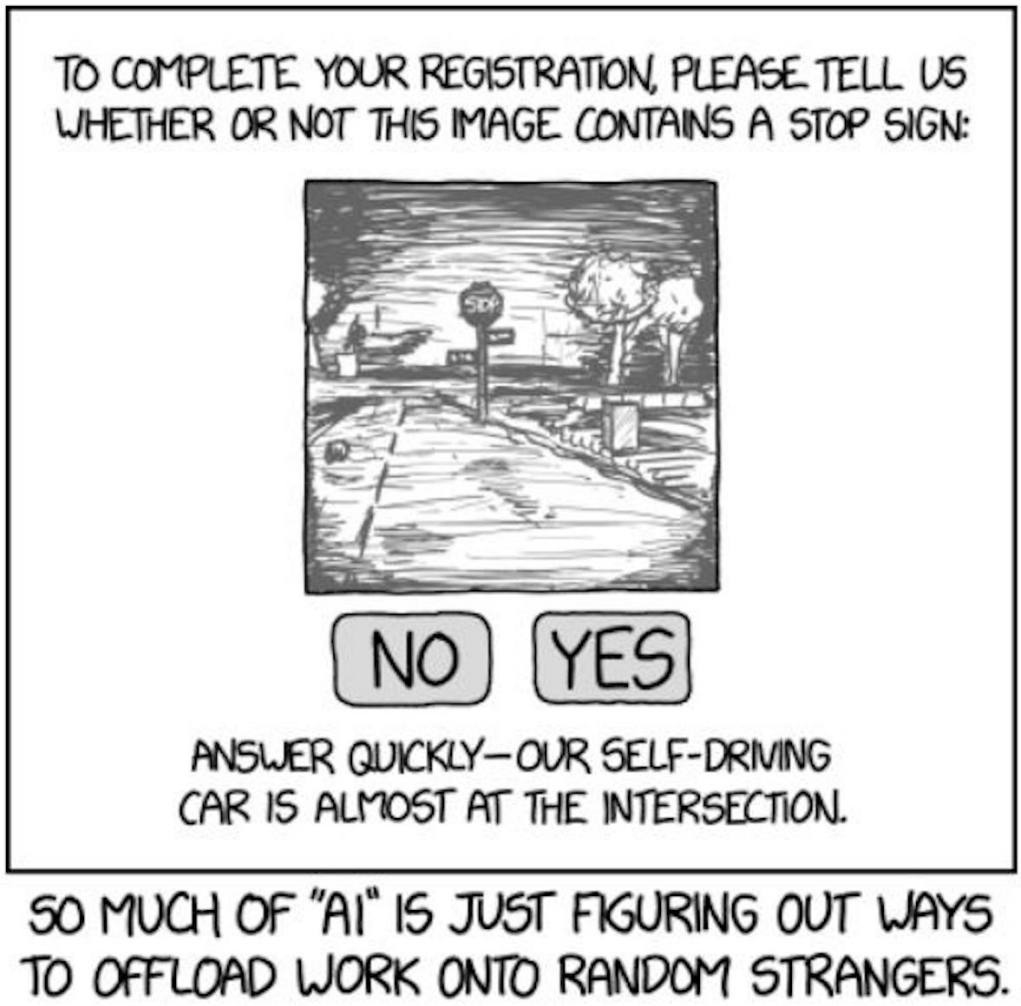

XKCD comic grabbed from the internet

In 2015 Shanghai-based designers Mathieu Cherubini and Simone Rebaudengo made a speculative object called the Ethical Fan. The Ethical Fan is a portable electric fan swivel-mounted on an input dashboard with dials and connected to the internet. When the fan is placed between two people, input buttons record 'ethical' information about the individuals such as their gender, education level and religion. The fan can compute who it should turn toward depending on the input. If it cannot decide because of a 'conflict', then this information is sent to a Mechanical Turk worker to resolve the question. The faraway Turker is also expected to offer a short justification for their choice. The results can be hilarious and nonsensical. For example, in one case the Turker says that the heavier of the two people should be fanned because fat people sweat more.

Watch on Vimeo.

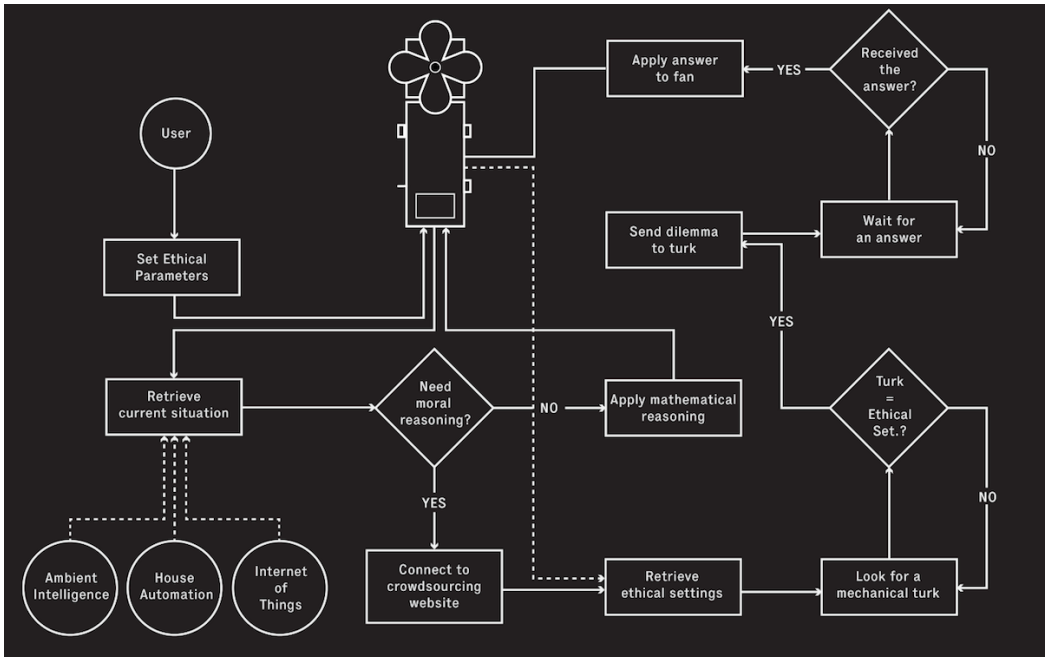

Flow chart accompanying the video of the Ethical Fan offers a blueprint for how decision-making is constructed in the system.

The designers seem to want the results from the Ethical Fan to be ridiculous in order to raise questions about how, or if at all, decisions are made by machines. Like the original medieval Turk, there is a human inside the machine making a decision. But the Fan gives the impression of arriving at the decision itself.

Complex socio-technical systems need to be pulled apart. In understanding how they work, and who and what they are comprised of, we can identify how power flows through the system. This is relevant in cases of error and breakdown: what interactions between which agents resulted in decisions that led to collapse? But it isn't always just that something breaks down because someone flipped the wrong switch. Technical disasters are often social, and incubate for long periods of time.

Diane Vaughan's book (1997) about the 1986 Challenger explosion shows that the potential for a disaster matures through poor communication, organisational culture and social and political pressure. She identifies "scripts" - a way of talking about technical knowledge, essentially - that NASA engineers came to believe about the faulty design of O-Rings on the shuttle that led to the explosion. Vaughan found that technical information about what was risky in the O-Ring design was re-classified as not-risky as assessments of the design flaw traveled between engineers, bureaucrats and managers at NASA, amidst incredible pressure to win the Space Race. Infrastructure is people, and maintaining and managing complex technical infrastructure requires attentiveness to the interaction of human and non-human.

A cluster of initiatives around 'algorithmic accountability' are opening up the black box of algorithmic infrastructures. These include projects such as algorithm audits, protocols to standardise training data sets (Gebru et al 2018), public impact accountability practices (Reisman et al 2018), a government-convened algorithm review task force in New York City (Powles 2017), the Fairness Accountability and Transparency conference, and the 2018 Global Data Protection Regulation (GDPR).

The mental model of a computational system is usually:

input – black box / process – output.

Emphasising the inscrutability of the black box and opening it up is important work. But it is equally important to ask how accountability initiatives re-inscribe particular approaches to what exactly the problem is. It is possible that imagining the black box as where the problem lies feeds our assumptions about exactly what fails when an algorithm is discriminatory or biased - computational, sociological, legal, political or cultural, or some combination of these? There is a risk that algorithmic accountability remains a computational fix. Like the self-driving car's morality algorithm, machine learning could well be programmed to regulate itself. Thus mechanisms of accountability need to themselves be interrogates: are they accountable too?

Recent organising and resistance among workers in Silicon Valley is an interesting development that complements initiatives for algorithmic accountability. Human workers are petitioning their employers – Google, Microsoft and Amazon - to not sell technology or expertise to government programs that are being used in the criminal justice, defence and immigration control. In the case of Google, the company withdrew from its discussions with the US Department of Defence on its drone program, Project Maven.

Accounting for the irregular

Mapping the vast technical systems of human and non-human agents has its limits. There is hubris at the heart of map-making: it is never possible to complete the map. Also, the map reveals the values and social position of the map-maker: what is considered worth mapping? What is left out?

Accountability mechanisms and approaches can be proactive: By understanding how a system works, and fails, there can be measures to improve its design. An ethnographer of infrastructure, Michael Fisch, and the artist and designer Ilona Gaynor push us to think about disasters in complex, large scale systems in terms of that which cannot be mapped: the irregular and the uncertain.

Nuclear reactors are "absolutely determined technologies" meaning that every part of the operation must be carefully mapped and regulated in order to foresee and manage errors. Any nuclear accident is a significant disaster. A nuclear reactor is inert and 'finished' once it is complete because it cannot remain open; changes can introduce instability that might affect the fragile chemical processes at the core of the reactor. In his exploration of accounts of the 2011 TEPCO (Tokyo Electrics) Fukushima nuclear reactor disaster, Michael Fisch finds that the word *soteigai* was invoked by TEPCO as the cause of the disaster. "Soteigai translates loosely as referring to something that is beyond expectations. Accordingly, it is commonly understood to denote something that can not be anticipated via existing risk management models and technologies." (p 1)

But, Fisch pushes past this explanation, showing that negligence and corporate mismanagement aside, eventually it was the closed nature of the system that led to the failure. Unlike an organic system that can evolve, a nuclear technology is a closed and tightly coupled system. Anything irregular – in this case, a tsunami that exceeded existing data about the effects of tsunamis – cannot be accounted for within the system. He concludes that *soteigai* was never about limits in thinking

about the possible causes and contingencies of failure of the system, but that "it has always been about the impossibility of thinking the consequences of the nuclear crisis." (p 6)

Everything Ends in Chaos (Gaynor 2011) touches on similar themes of the limits of what can be known or imagined. This finely detailed work emerged in the aftermath of the 2008 financial crisis so it spans economics, finance, global markets, risk management, insurance and mathematics. Gaynor reverse-engineers fictional global catastrophes through various scenarios: one starts with the kidnapping of a wife of a wealthy senator; a second is about a bomb in the boardroom of an insurance firm. Reminiscent of 1950s scenario planning (Galison 2014) she traces the 'what if' path to understand how imaginable and unimaginable events might be predicted, managed and reversed. Inspired by the idea of a 'Black Swan' event, Gaynor asks how we might know and manage risk and disaster through instruments of precision which themselves may not be necessarily precise.

She says that designing exaggerated and hypothetical scenarios reveals how certain systems work, as well as how future economic and financial systems might be re-designed. Gaynor says:

> *"I do think that as its complexity continues to grow and get increasingly denser, it starts to tangle and knots occur. It's becoming more and more difficult to control such a living organism and I don't think we can continue down a pathway that's so obviously treacherous. The critical discourse lies in my aim to celebrate such a system. It's a non-human entity with non-human goals, and it's deliciously destructive." (deBatty 2011)*

In thinking about complex and large architectures in systems like AI, Fisch and Gaynor ask that we identify the limits imposed by technical system themselves, and in our own thinking, about causes and consequences of breakdowns and errors. This may take us to a beginning, not an end, of where we might articulate ethics:

> *"An account of oneself is always given to another, whether conjured or existing, and this other establishes a scene of address as a more primary ethical relation than a reflexive effort to give an account of oneself. Moreover, the very terms by which we give an account, by which we make ourselves intelligible to ourselves and to others, are not of our making. It may be possible to show that the question of ethics emerges precisely at the limits of our schemes of intelligibility, the site where we ask ourselves what it might*

mean to continue a dialogue where no common ground can be assumed, where one is, as it were, at the limits of what one knows yet still under the demand to offer and receive acknowledgment" (Butler 2005, p 20-21).

References

Amoore, L. and Piotukh, V. (2016).Eds: *Algorithmic Life: Calculative Devices In The Age Of Big Data*. London and New York: Routledge

Anderson, M. and Anderson, S.L. (2007). The status of machine ethics: a report from the AAAI symposium. Minds and Machines 17: 1-10.

Butler, J. (2005) *Giving an account of the self*. New York: Fordham University Press.

Coleman, R (2016) 'Calculating Obesity, Pre-Emptive Power And The Politics Of Futurity' in Amoore and Piotukh (eds) Algorithmic Life: Calculative Devices In The Age Of Big Data. London and New York: Routledge. p 185

deBatty, R. (2011) Everything ends in Chaos: Interview with Ilona Gaynor. *We Make Money Not Art*. Retrieved 1 August 2018 http://we-make-money-not-art.com/everything_ends_in_chaos/

Fisch, M. (n.d) Meditations on the Unthinkable (soteigai). In Erez Golani Solomon, Editor. *The Space of Disaster*. Tel-Aviv, Resling Publishing. Retrieved 1 August 2018 https://anthropology.uchicago.edu/people/faculty_member/michael_fisch/

Fisch, M. (2017) Remediating infrastructure: Tokyo's commuter train network and the new autonomy in Infrastructures and Social Complexity: A companion by Penny Harvey, Casper Bruun Jensen and Atsuro Morita (eds). London and New York: Routledge.

Ganesh, M. I (2017) In a personal conversation with Martha Poon in Brussels, January 2017.

Galison, P. (2000). "An Accident of History." Ed. Peter Galison and Alex Roland. Atmospheric Flight in the Twentieth Century. Springer Science and Business Media: 3-43.

Vaughan, D. (1997). The Challenger Launch Decision: Risky Technology, Culture and Deviance at NASA. Chicago: University of Chicago.

Galison, P (2014). The Future of Scenarios: State Science Fiction In *The Subject of Rosi Braidotti: Politics and Concepts*, edited by Bolette Blaagaard and Iris van der Tuin, 38-46. London and New York: Bloomsbury Academic.

Gebru, T., Morgenstern, J., Vecchione, B., Wortman, J., Wallach, H., Daumeé , H, and Kate Crawford. 2018. Datasheets for datasets. Retrieved 1 August 2018: https://arxiv.org/abs/1803.09010

Johnson D.G. (2011) Software Agents, Anticipatory Ethics, and Accountability. In: Marchant G., Allenby B., Herkert J. (eds) The Growing Gap Between Emerging Technologies and Legal-Ethical Oversight. The International Library of Ethics, Law and Technology, vol 7. Springer, Dordrecht

Moor, J. (2006) The Nature, Importance, and Difficulty of Machine Ethics. *Machine Ethics: IEEE Intelligent Systems*.

Natale, S. and Balatore, A. (2017) Imagining the thinking machine: Technological myths and the rise of artificial intelligence. Convergence: The International Journal of Research into New Media Technologies. 1–16 Retrieved DOI: 10.1177/1354856517715164. 22 January 2018

Powles, J. (2017) New York City's Bold, Flawed Attempt to Make Algorithms Accountable. New Yorker December 20, 2017. Accessed online https://www.newyorker.com/tech/elements/new-york-citys-bold-flawed-attempt-to-make-algorithms-accountable 5 January 2018

Rahwan, I and Awad, E (2018). The Moral Machine Experiment: 40 Million Decisions and the Path to Universal Machine Ethics. Invited talk at the Artificial Intelligence and Ethics Conference, AIES, New Orleans, LA, February 1-3, 2018. Retrieved 2 August 2018 http://www.aies-conference.com/invited-talks/

Rahwan, I (2016) The Social Dilemma of Driverless Cars. Tedx Cambridge. Retrieved https://www.youtube.com/watch?v=nhCh1pBsS80 14 November 2017.

Reisman, D., Schultz, J., Crawford, K and Meredith Whittaker (2018) Algorithmic Impact Assessments. A Practical Framework For Public Agency and Accountability. AI Now Institute. https://ainowinstitute.org/aiareport2018.pdf

Schuppli, S. (2014) Deadly Algorithms: Can Legal Codes hold Software accountable for Code that Kills? Radical Philosophy, Issue 187 UK, (2014): 2-8. Accessed online http://susanschuppli.com/writing/deadly-algorithms-can-legal-codes-hold-

software-accountable-for-code-that-kills/ 12 February 2018

Vaughan, D. (1996) *The Challenger Launch Decision: Risky Technology, Culture and Deviance at NASA*. Chicago: University of Chicago Press. ISBN 0-226-85176-1.

Weast, J.C., Kohlenberg, T.M., Johnson, B.D (2017) Technologies for resolving moral conflicts during autonomous operation of a machine. US Patent application publication number US2017/0285585 A1 Oct 5, 2017.

Maya Ganesh is a technology researcher, educator and writer who works with industry, arts and culture organisations, academia and NGOs. She is working on a PhD about autonomy, ethics and AI. Her other research interests include: design; financial technologies; post-humanism; and the contested term 'Anthropocene'. She has worked with Tactical Technology Collective, Point of View Bombay, UNICEF India, and the APC Women's Rights Program. Her writing and publications are here. She tweets @mayameme.

Tech, Trust, Transparency: The Trustable Tech Mark

By Peter Bihr

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

Tech has a trust issue. It's the year 2018. We live in a world *post-Snowden, post-Cambridge Analytica, post-scandal after scandal of security and data and privacy breaches*. The overly optimistic, *gung-ho*, maybe even naive tech optimism that reigned supreme until the mid-2000s has served its time for now. This is particularly palpable in the space of IoT: Adding microphones, cameras and an internet connection to everyday objects has a way of making people think just that little bit harder about their privacy. I like to think this isn't necessarily a bad thing but a great opportunity.

> *"Let's be clear: none of our instincts will guide us in our approach to the next normal."*
>
> ***— Adam Greenfield, Radical Technologies***

In his book *Radical Technologies*, Adam Greenfield points out that networked technology changes the way we interact with our world, and that it does so in ways that are often pervasive, invisible, unintuitive. If our instincts cannot guide us, that makes it all the more important that connected devices are designed and built responsibly: They need to ship and function with respect for users, and their rights, privacy, and everyday context. They need to be *better* and more *responsibly* than most connected products are today.

Within the ThingsCon community, we've been advocating for better design and data practices since day one. If ever there was a time to put our thinking into action, it is now.

Enter the Trustable Technology mark.

The Trustable Technology Mark

The Trustable Technology mark is our attempt to establish a consumer trustmark for the Internet of Things (IoT). It's one of ThingsCon's core initiatives this year (and hopefully for some years to come), and made possible with support from the Mozilla Foundation, who invited me to join as a Mozilla Fellow for the year.

Consumers don't currently have the tools to find out which connected products are trustworthy. What's more, there are companies out there who go out of their way to build responsible products that respect their users's privacy and rights, but they don't have an effective way of communicating their commitment. Here's huge potential for a trustmark for IoT.

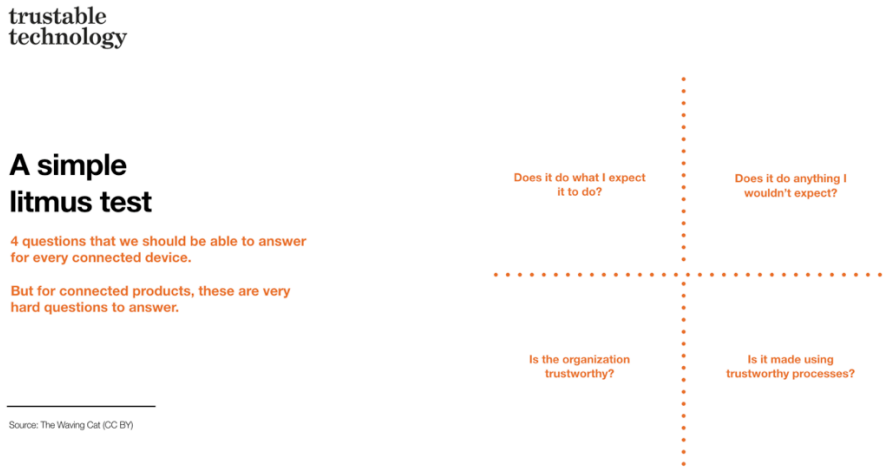

After doing extensive research, we're convinced that a trustmark can address these issues meaningfully.

We believe that trust is holistic, systemic in nature: An insular focus on *exclusively security* or *only privacy* won't do. However, because of the hybrid nature of IoT products—hardware, software (on-device) and service (often server or cloud

based) there are some tricky aspects to external audits that we haven't seen solved anywhere yet. So we went a different route.

> *Trust is always earned, never given.*
>
> *— Proverb*

We're designing an self-assessment tool that allows companies to evaluate where their product meets or doesn't meet our trust requirements. This tool can also serve as a guideline for designing better and more trustworthy products in the future. It's all openly licensed and free to use, forever.

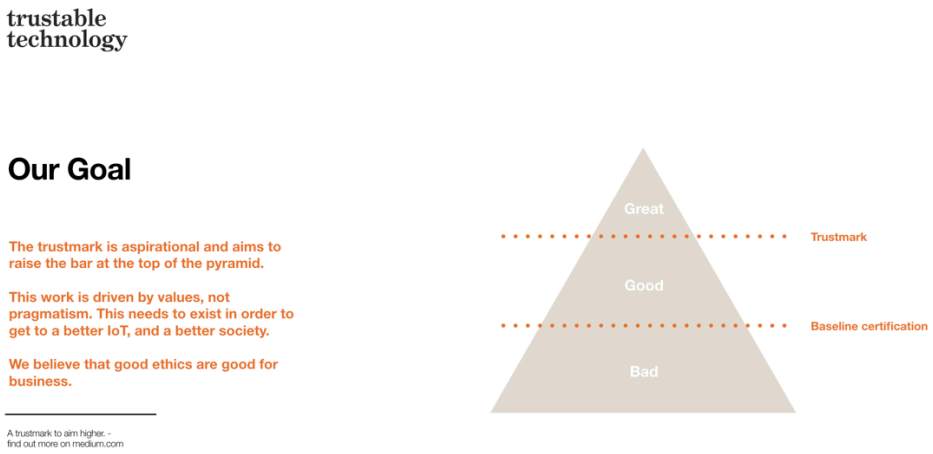

The self-assessment tool - which also doubles as the application form - guides the company through a series of questions. The company ultimately decides if they are confident to clear the bar, and if they do, they submit their assessment as an application. Our experts review the answers and check for obvious gaps, or follow up for clarification. Once both the company and our reviewers give their go, the self-assessment is published in full: It's a public commitment that their practices match these answers. It's part of the certification requirements that the full assessment is published under an open license for everyone to peruse.

What does the self-assessment look like? Concretely, imagine a questionnaire that consists of simple but tough YES/NO questions. A YES counts towards the success, a NO counts against. A NOT APPLICABLE (N/A) won't be scored. For every answer we highly encourage an explainer paragraph what this means in this specific context. This sounds more complicated than it is; it is a really simple, straightforward process.

Our initial testing shows that those companies who already put in the effort find it easy and quick to answer these questions, and the ones who don't tend to struggle. So we're confident it works fairly robustly.

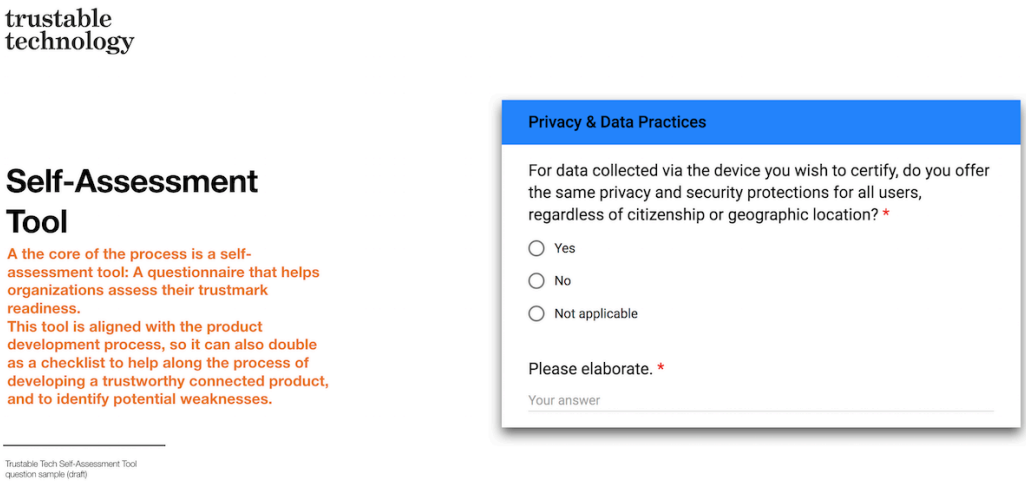

A prototype of the self-assessment tool that also doubles as the application form. At the time of writing, this prototype is available online as a Google Form.

What are we looking at?

We identified 5 dimensions that we believe are relevant to anyone inviting a connected device into their lives:

- Security
- Privacy & Data Practices
- Openness
- Transparency
- Stability

The first four are largely self-explanatory, the fifth requires a bit of context. Think of *stability* as an indicator of robustness and longevity: Will the product still work if the company goes belly-up or switches off their servers? Is there an exit strategy to keep the products working after an acquisition? Does the company commit to software and security updates to make sure the device can be safely used for a few years after the initial sale?

Our initial research has shown that this approach can be quite powerful: While the questions are simple, they do cut deep. Answering them requires a level of commitment to openness and transparency that comes easy only to those companies who do the right thing anyway, and will be nearly impossible who those who don't.

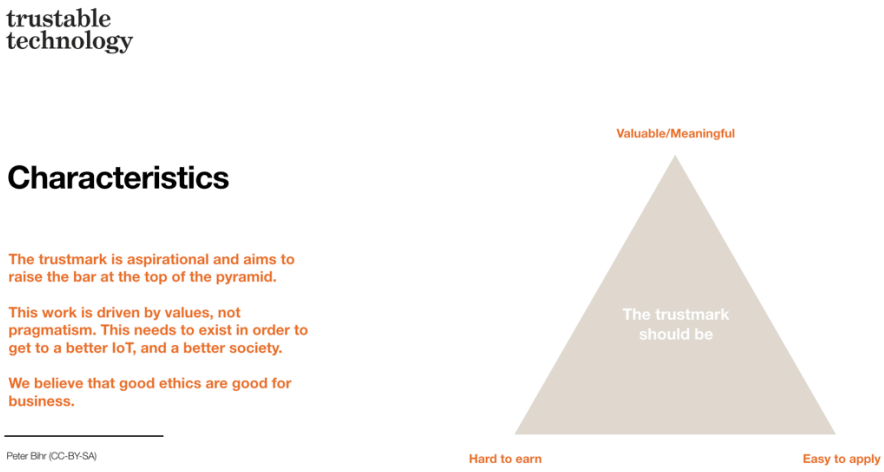

Consider this example from the *security* section of the self-assessment tool:

- What is the core functionality of this product? *Please explain the core functionality of your product.*
- Are there any other features or functionalities outside the core functionality? *Please explain why the choice was made to include this feature or functionality.*
- Which non-core (non-essential) features could be enabled in the future? *Please explain why the choice was made to potentially enable this feature or functionality in the future.*

This series of three question aims to determine the risk to security through *feature creep*, because any non-essential feature might open new security holes. These questions all need to be considered during the design process; we've heard multiple times that product owners and designers might in fact find it useful to be forced to be explicit about these decisions internally. The questions are also simple enough to answer, but also incredibly hard. Not only do you have to be clear on the one thing your product tries to do well, but also give a clear glimpse into your decision making. Finally, the question about the product development roadmap exists to ensure transparency about the potential for features (and security risks) that could be enabled through the next software update: If there's a microphone in the device that isn't required as part of the core feature set, this is the time to disclose it.

The beginnings of an ecosystem

The trustmark can stand by itself, and it would provide significant value. However, we hope that there will be more to it. We envision a whole ecosystem to grow around the trustmark. We're building the self-assessment and the trustmark that derives from it. When designing them we did so with third party services in mind.

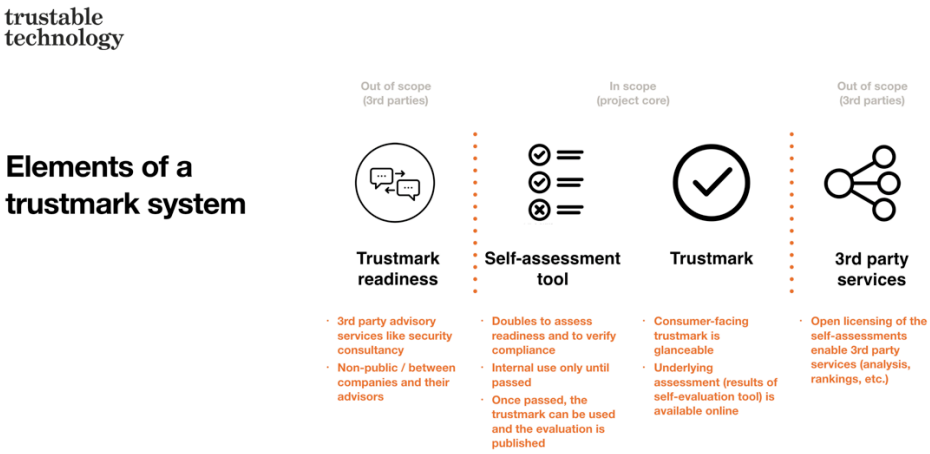

It's all build on open licenses so that the published documentation (the results of the self-assessments) can be aggregated, analyzed, made accessible or sliced and diced in any kind of way we can't even think of now. How about an app? A shopping guide? A ranking of the most privacy-respecting toys?

We also think there's potentially a whole small but healthy opportunity for advisors to get companies "trustmark ready". This could happen commercially or through volunteers. The fantastic network that our friends of the UK-based OpenIoTMark (whose excellent design principles we've also integrated into the Trustable Technology mark) has been building for that purpose seems like an obvious great fit. If you're a startup in need of security advice, this is where you can find an expert who's willing to engage.

Open questions & what's next?

In an issue as complex as IoT certification, the devil is in the details. So we're in the nitty-gritty of testing and further prototyping the self-assessment tool. We do that by talking to companies who help us test the trustmark process, and by hosting workshops with experts in the field. We're figuring out the best way to make the trustmark legally binding, and of figuring out questions around governance. We're lining up commercial and academic partners so that we have a strong alliance once we're ready to launch.

We've started to speaking about the trustmark more publicly at meetups and conferences to expose the idea to more eyes, ears, and minds: The quality of feedback has been astounding, and the level of interest shows just how needed

(and timely!) this initiative is.

I'm convinced that we can make a significant contribution towards a more trustworthy Internet of Things. One trustmark and product at a time.

You can learn more about ThingsCon on thingscon.com and about the Trustable Technology mark at trustabletech.com.

If you or your organization would like to be involved in the trustmark initiative in some way, please get in touch.

I'd like to thank the ThingsCon community for all the input. Mozilla Foundation for all of their support through my fellowship. And I'd particularly like to thank Pete Thomas and Jason Schultz. Pete (of University of Dundee) has taken the lead in the branding and design of the mark, and been an excellent sparring partner in strategic questions. Jason (of NYU Law) has been exploring legal and policy implications of the Trustable Tech mark.

Peter Bihr co-founded ThingsCon, a global community & event platform that fosters the creation of a responsible Internet of Things. In 2018, Peter is a Mozilla Fellow. He also is the founder and managing director of The Waving Cat, a boutique digital strategy, research & foresight company. We explore the impact of emerging technologies — like Internet of Things (IoT) and artificial intelligence — and how your organization can harness them effectively. Interested in working together? Let's have a chat.

Full disclosure: Peter's partner works for the Mozilla Foundation.

Responsible IoT after techno-solutionism

by Prof. Dr. Seyram Avle, David Li & Prof. Dr. Silvia Lindtner

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

In the first two weeks of July 2018, the world was captivated by news coverage of a rescue mission for a team of young footballers in Thailand who had been stranded in a cave. Towards the end of their ordeal and the news spectacle, Elon Musk, the CEO of Tesla and SpaceX announced that his team had built a minisub using cutting-edge space technology to aid the rescue mission. They had tested it in a pristine swimming pool in LA and then Musk flew to deliver it in the murky waters outside the cave where the boys were stuck. While Musk had been busy building hype for the minisub, a team of local and international experts and volunteers had already started getting the boys out.[25]

While often not attracting the same broad media coverage, stories of such failed techno-solutionism, the belief that technological innovations in their own right can solve complex societal challenges, abound. Across regions, hackathons, pitch contests, and design competitions are positioned as ideal to help cultivate a mindset of entrepreneurial agility and innovation thinking in turn portrayed as crucial to facilitate societal and economic change ([26] Irani 2015, [27] Lindtner 2015, [28] 2017).With taglines such as "Design for Good," "AI for Good," "HCI for Good," civic hacking, and so on, they garner investment from foundations, government, and corporations eager to signal social responsibility. Common to these events is the gathering of technologists, engineers, and designers charged with the task to create solutions for the disadvantaged, often without their close involvement. The proposed solutions are often more attuned to investors' interests in value accumulation than to the realities they promise to intervene in.

All of these projects represent a much broader and by now several decades old approach of designing technology "for good", for "development" (also often referred to as ICT4D), and for others[29], [30], [31]. While often well-meaning and motivated by commitments to designing technology that serves rather than undermines humanity and economic opportunity, research has shown for decades now the continuous failures of such techno-solutionist ideals and projects centered around designing or developing for others[32], [33], [34].

In this short article, we explore an alternative view of what counts as responsible technology innovation and responsible IoT in particular. Specifically, we tell a story of IoT innovation that starts from a position of designing with and within rather than for, and from an attitude of partnership rather than a rhetoric of "do good" or "feel good". We urge for the importance to locate responsibility not only in the technical product (e.g. a product or service that serves low-income populations or that enables a more equitable life for minorities), but also within the social and intercultural processes of design and production.

In the spring of 2018, two of us traveled to Accra, Ghana, where we met Kamal Yakub, the CEO of two Ghana-based companies, Farmable[35] and Trotro Tractor[36]. When we met Kamal, he was preparing for a trip to Shenzhen in the South of China to solidify a deal with a Chinese company, ThinkRace, to design a new tracking product with elderly travellers in mind. Two months earlier, Kamal's team had found ThinkRace online and reached out to them. The engineering and design team took their prototypes, got on a plane to Shenzhen and worked with ThinkRace for three weeks to turn their prototypes, tested in the fields and with farmers in Ghana, into a product.

Farmable is a crowd farming platform that connects small holder cattle farmers in Ghana with a global market. The company's suites of IoT devices helps farmers track cattle, monitor their health and connect with potential buyers. Buyers can invest in cows, and also follow their progress till they are ready "for harvest". Trotro Tractor targets small holder farmers and allows them to rent tractors through a sharing model, thereby reducing their equipment cost and increasing their farm productivity. The IoT devices designed by the company provide real time data that connects farmers to the nearest and available tractor, cutting down wait times, in largely rural areas where internet access is quite low.

We followed Kamal to Shenzhen and joined his visit with his Chinese counterpart, ThinkRace's CEO Rick Tang. What struck us about the interactions between Kamal and Rick, and what we heard from their team member's prior interactions, was that they were enacted and articulated as mutual learning from one another, as a form

of partnership and co-dependence even, with each side gaining and contributing equally. They were both equally invested, even if the type of investment differed. While Kamal was invested in creating low-cost IoT devices for farmers in Ghana, Rick was invested in positioning his company as a trustworthy and ethical Chinese company serving an international market. Both businessmen were also upfront about and openly discussed the necessity and importance of making money to sustain and advance their respective companies. In techno-solutionist discourse, making money is often de-emphasized as it is rendered to signal a lack of creativity, or greed, or in some cases as not representing a hacker or maker ethic. Truly innovative products, so the story often goes, emerge from creative play and tinkering rather than economic interests or concerns.

For Kamal, who had participated in a myriad of pitch competitions and start-up events and gatherings over the last years, this partnership was different. After one of his meetings with another company, he remarked that "doing this here is different than in Silicon Valley because in Silicon Valley there are no social problems attached to the technology. [...] Yes, of course, I do this for money, I'm a businessman... but I also need to solve problems... In Shenzhen, when you are here, you just think of innovation. You think about how you can build it. It's not about hype, it's about implementing change." Rick put it another way "talk is cheap — put some money and some hope on it". Both men felt energized by their collaboration and saw a mutually beneficial opportunity to learn and craft new futures for their companies and the people that worked in them.

The story of Kamal and Rick makes visible the limitations and consequences of more typical portrayals of responsible technology as emerging from designing for others less privileged. This 'othering' of people in need "out there"^5 undermines and reduces their agency to recipients of aid. Once a project is constructed as aid for a presumed disadvantaged other, it runs the risk of exacerbating the asymmetry between the giver's desires and the intended recipients. It also risks undermining the build-up of trust and recognition of mutual dependencies as we saw in Kamal and Rick's partnership.

Technological designs and systems framed as *for development* are often legitimized in the name of objectivity and what Donna Haraway calls the god trick; i.e. claiming one's own partial view of the world as objective and universally applicable[10].We argue that responsible IoT can fall into this trap, when tools, techniques, and devices conceived in elite IT laboratories are framed as universally applicable and beneficial. This language of universality masks the corporate and financial interests of already powerful tech companies invested in expanding their business and reach to the very places being construed elsewhere as in need of aid.

Our aim is not to naively celebrate Kamal or Rick as somewhat more "authentic" innovators or "native" entrepreneurs. On the contrary, our goal is to work against such overly simplistic stories of success and against the imagery of the hero/savior (as the Musk story we opened with suggests). Let's tell other stories and shift our orientation from techno-solutionism to mutual accountability and co-dependencies — or as feminist scholar Sara Ahmed (2010)^9 put it: "Orientations matter. To say that orientations matter affects how we think 'matter'. Orientations might shape how matter 'matters.'" This means also orienting towards the continuous exclusions and discriminations along gender, class, and racial lines that lurk underneath a myriad of seemingly successful and mutually beneficial partnerships.

Rather than constructing responsibility for or on behalf of an 'other' out there, one path might be to interrogate for whom we are claiming responsibility and to listen and work with those already finding solutions in the places we want to help, rather than defaulting to doing things for them, usually without their input. Rather than starting from the technology, in this case IoT, responsibility could start from a place where we take seriously the aspirations and commitments to futures and worlds that are otherwise, as enacted partially by Kamal and Rick.

Our argument, simply put, is that responsible IoT is when their design and production starts from a place of mutual care, respect, and equality. We are inspired by a long lineage of feminist technoscience scholarship that provides techniques and methods for imagining and implementing technology (e.g. among others Anna Tsing, Shaowen Bardzell, Lisa Nakamura, Lucy Suchman). Their work has long shown how exclusions along the lines of gender, race, and class persist despite and perhaps even because of well-meaning "do good" initiatives in IT industry and research communities. Our approach begins with an understanding that responsibility in technology innovation starts from a commitment by the technology designer, engineer and/or researcher to remain accountable for the actions and values taken. It also includes accounting for one's partial view, rather than claiming objective expertise and universal knowledge.

The tools for IoT may have 'democratized' to some extent but the socioeconomic systems behind them have not. We therefore need to start being critical at the onset of more pervasive IoT rather than after the fact when more harm than good has taken place.Ultimately, responsible IoT may not need advanced technological solutions but commitments to accountability and responsibility.

References

1. See https://slate.com/technology/2018/07/elon-musk-is-trying-to-aid-the-thai-

cave-rescue-by-sending-engineers-and-brainstorming-on-twitter.html and https://slate.com/technology/2018/07/elon-musk-thai-soccer-team-cave-rescue-fruitless-attempt.html for an example of the media covered following these events.
2. Lilly Irani. 2015. Hackathons and the Making of Entrepreneurial Citizenship. Science, Technology & Human Values (Sage), Vol. 40, No. 5.
3. Silvia Lindtner. 2015. Hacking with Chinese Characteristics: The Promises of the Maker Movement against China's Manufacturing Culture. Science, Technology & Human Values (Sage), Vol. 40, No. 5, pp. 854-879.
4. Silvia Lindtner. 2017. Laboratory of the Precarious. Methods of the Precarious: The Seductive Draw of Entrepreneurial Living. Women's Studies Quarterly, Vol. 45, Nr. 3&4, pp. 287-305.
5. Lilly Irani, Janet Vertesi, Paul Dourish, Kavita Philip, and Rebecca Grinter. 2010. Postcolonial Computing: A Lens on Design and Development. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 1311-1320.
6. Alex Taylor. 2011. Out there. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2011), 685-694.
7. Seyram Avle and Silvia Lindtner. 2016. Design(ing) "here" and "there": Tech Entrepreneurs, Global Markets, and Reflexivity in Design Processes. In Proceedings of SIGCHI Conference on Human Factors in Computing Systems (CHI '16), 2233 - 2245.
8. James Ferguson. 1990. The anti-politics machine: 'development', depoliticization and bureaucratic power in Lesotho. CUP Archive.
9. Sara Ahmed. 2010. Ahmed, S., 2010. Orientations matter. In Coole D. and Frost S. (eds)New materialisms: Ontology, agency, and politics, pp.234-257. Duke University Press.
10. Christo Sims. 2017. Disruptive fixation: School reform and the pitfalls of techno-idealism. Princeton University Press.
11. http://www.farmable.me
12. http://www.trotrotractor.com
13. Donna Haraway. 1988. Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. Feminist Studies, Vol. 14, No. 3 (Autumn, 1988), pp. 575-599.

Seyram Avle is an assistant professor at the University of Massachusetts, Amherst. Her research is on digital technology cultures and innovation in the global south.

David Li is the executive director of Shenzhen Open Innovation Lab which facilitate the collaboration between global smart hardware entrepreneurs and Shenzhen Open Innovation ecosystem. Prior to SZOIL, he co-founded XinCheJian in 2010, the first hackerspace in China to promote hacker/maker culture and open source hardware. In 2011, he co-founded Hacked Matter, a research hub on maker movement and open innovation. In 2015, he co-founded Maker Collider, a platform to develop next generation IoT from Maker community.

Silvia Lindtner is an assistant professor at the University of Michigan in the School of Information, with a courtesy appointment in the Penny W. Stamps School of Art and Design. Lindtner's research and teaching interests include innovation and technology entrepreneurship, making and hacking cultures, shifts in digital work, labor, and industry. This work unfolds through a deep engagement with issues of gender, inequality, and enactments of masculinity in engineering, politics of design, contemporary political economy, and processes of economization. Lindtner draws from ten years of multi-sited ethnographic research on China's shifting role in global tech production, including research in urban China, Europe, the United States, Taiwan, and Africa. Her research has been awarded support from the US National Science Foundation, IMLS, Intel Labs, Google Anita Borg, and the Chinese National Natural Science Foundation. Together with Professor Anna Greenspan and David Li, Lindtner co-founded the Research Initiative Hacked Matter, dedicated to critically investigating processes of technology innovation, urban redesign, and maker-manufacturing cultures in China.

Observing Things - Responsibility and the Internet of Things

By Simon Höher

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

Before we call for a responsible IoT, let's make sure we know we actually mean by that. While the meaning of IoT appears to be relatively obvious, a precise understanding of responsibility is more elusive. One possible answer is that acting responsibly means to offer a choice in face of multiple options – management decisions at the work place, personal and social preferences at home, administrative or creative decisions in the city. What such a choice might consist of, and what makes it responsible, is the topic of this essay. In search for answers we will turn to cybernetics and Systems Theory, in particular Luhmann's notion of social systems [37] and Heinz von Foerster's proposal of the role of recursive computation for the observation of the world at hand.[38] We will close with an attempt to apply this cybernetic notion of responsibility to our own expectations toward the IoT with the help of Bruno Latour (culture) and our claims to its structure, reviewing Xanadu, a proposal for an alternative network structure from the 1990s (design).

Computing Environments

While the term *IoT* itself is somewhat bewildering – what else should the internet consist of, if not computers, that is things? – todays' notion of the IoT links back to the first conjuring of the term "ubiquitous computing", seminally proposed by Mark Weiser and his colleagues at Xerox PARC in the early 1990s[39]. Together with the idea of "ambient intelligence", coined by Eli Zelkha and his team at Philipps, they

envisioned a digital network that represents physical things by virtual components with respective virtual addresses – a hybrid network of physical and virtual references. Interestingly, the term of ubiquitous *computing* already hints at an active performativity of things that now become interconnected and ever intentional. Kevin Ashton's (proclaimed, but widely accepted) coinage of the term "Internet of Things" in 1999 then brought little new to the table, but offered a catchy description of the concept.

Aside from history and semantics, let's take a closer look at such a global network of connected things: A computational network where each physical node (thing) is represented by a virtual node and performs the triumvirate of (1) monitoring their surroundings by sensors that *encode* analog input (sound, heat, light, touch, etc.) into digital data, (2) *processing* this data by applying specific programs or algorithms, and (3) eventually generating some sort of *output* that - and this is important - likely functions as an input for the *next* triumvirate system, not rarely itself.[40] The re-introductiong of previous output as the input for the next operation makes the IoT a "non-trivial machine" that refers to its current state in order to derive its future processes. We might even refer to it as some sort of decentralized supercomputer itself, or as other have put it, a "world-sized robot"[41]. Interesting for us is the extent to which such a robot appears to be a neutral, or at least predictable machine to its observers (us), and to what extent it seems to be a new, equally purposeful[42] observer itself. We want to ask what is going on in the IoT – and what does it mean for society?

The Internet of Invisible Things

That the increasing presence of connected technologies has a severe impact on all sorts of social settings, and is equally adopted into their own respective logic, is apparent: examples reach from public management and administrations that translate the IoT into a notion of a connected and "smart city" to businesses and industrial corporations that perpetuate the idea of an "Industry 4.0" in the factory and the "Future of Work" in the office all the way to a new sense of intimacy and privacy in the setting of the family and their now allegedly "smart home".

The IoT is elusive, as it is always both, physical and virtual

This embrace and translation of ubiquitous computing into specific contexts of specific organizations is as telling to the attentive observer of such systems, as it tends to blur the technology at hand, making it to some extent invisible to the observing organization itself: **The IoT is elusive, as it is always both, physical**

and virtual, and this duality runs the implicit risk of leading us to mistake it for its other half: The physical object for its virtual representation, the virtual element for the actual thing.

It is this opaque invisibility is a quality that Weiser and his colleagues marked the ultimate triumph of any "profound technology" – quite enthusiastically, one might add. However, the implied benefit of ubiquitous computers that "weave themselves into the fabric of everyday life until they are indistinguishable from it"[43] holds the risk of being simply overlooked as what it is. A trait that might hold far-ranging implications for the way we should describe them, as they obviously are more than a trivial tool that helps making life manageable, and rather marked by a certain transformative, a non-trivial quality, an enchanted and enchanting effect, that demands to be observed just as suspiciously and wary as its sensors are monitoring its surroundings itself.[44] But how to do so, if it remains unclear what (or where) the thing really is?

Elusive Control

Once we look at issues like control and neutrality, this becomes obviously problematic, as both terms become crucial: They contradict and enable choice, depending on who controls and who sees clearly. On an economic level, some have pointed out that the sense of ostensible neutrality paired with a pervasive intangibility is indeed often not what it seems to those employing it. In his rather gloomy account of "The Epic Struggle for the Internet of Things", Bruce Sterling paints a disheartening picture of a technological development that is merely an economic one in disguise.[45] Depicting IoT as a pervasive instrument of oligopolistic organizations (most notably Google, Apple, Facebook, Amazon, and Microsoft, short: GAFAM), that "wrangle" for a dominant worldview of ever persistent sales, delivered through the connected devices in the homes, cities, and office buildings of the world, the once sanguine aspirations of Weiser and the likes seem deserted. Rather than a neutral - let alone helpful - technology, the IoT becomes an extended, massive communications channel to cater to the interest of a few giant economic organizations.

That said, we can take a closer upon a highly interesting observation of Sterling, namely the deception of the consumer, in thinking the Internet of Things would be about things (rather than sales). The elusive quality of IoT once more becomes obvious here, although embedded in an economic context, as well as its manipulative element of control. Our question thus becomes: **How does IoT manage to hide behind itself** – and how does it exert control while doing so?

We should recall here, that a ubiquitously computing environment is, indeed, *computing*. That means all references and information offered up for referral by ongoing communication are themselves the result (output) of computational processes yielding them. Sterling's assumed manipulative role of the IoT thus becomes indeed evident, as it can directly influence communication by selecting or repeating information: re-actualizing otherwise forgotten memories, and withholding potential information that might alternatively be referred to. The two sides of this selection then are *memory* and *oblivion* – a powerful lever to shift both, a system's focus of observation and its elusive blind spot.

It would stand to reason to describe this quasi-memory function of computing environments as a form of control, as, according to Ashby, control is a sort of memory of past events that constantly scans and adopts to diverging observations in the present, in order to (re-) formulate respective expectations in the future[46]. **With memory comes control**, and Sterling's hunch was right, if only not exactly (or rather: exclusively) in the way he laid it out, for this holds true any social system, depending on its respective "code". To speak with Spencer-Brown[47], this code is the initial and concurrent distinction that allows for control in the first place - and it can have many forms: buying | not-buying for economics, familiar | strange for the intimate home, public | private for the city.

As we have seen, the fact that the IoT exerts control by offering and withholding information, is not necessarily something to be worried about. Instead of a depletion of communication and an ensuing entropy of decisions (one might ask: what is there left to administer in a smart city), of intimacy (what is there left to share in an ever-watching intelligent home), or of innovation (what to manage in a predictively optimized business), it is more likely, that those observing systems notice they are observing and observed by technology - and subsequently learn to evolve and adapt, modeling strategies to ensure their own continuity, by building upon the observations of their surroundings. After all, the assumption of a computing environment is nothing new for any social system.

Choice and Truth

If we review IoT this way, we can discern its performative and controlling role in communicative contexts as making certain selection preferences more likely than others, without an explicit opportunity for objection. A computing environment is, in fact, an observer. Just like any observer, by observing and computing, it creates the reality it is observing in the first place – in this case by decoding it into data, by computing it, and by encoding it into some kind of result or output. It is this

constructed reality that any other system is then observing on their behalf. And it is through this selective construction of reality, that a computing systems *controls* how other observers might deal with it, refer to it, act upon its outputs, or address it deliberately.

This notion of control get us closer to a refined understanding of how computing environments might affect present communicative processes in society: by establishing elusive control, computing and presenting (via outputs) a specifically served plate for selection of past and future communications. We might also describe this selection proposition as a forestalled distinction, a *de facto* pre-selection of potential and perceived truths – as unrecognizable, and thus indisputable observations. We should remember one of the core introductions of cybernetics here, namely the observer and its subsequent annulment of any absolute truth. To the contrary, any unconditional adoption of constrained distinctions would run exactly diametrical to the realization that "everything said is said by an observer"[48] – including any statement, observation, or reasoning for truth. Even if there was a truth, once it enters communication it is thus doomed by its observers.

The result is not loss of reference or orientation, but rather a discovery of the autonomy of the observer, and of its "responsibility" and liability for the truths selected.[49] **Responsibility is both a prerequisite and result of choice** – and granting and marking options for choice would be responsible, as it allows others to choose deliberately.

A responsible IoT

For social systems like organizations and even society as a whole this would mean a different form of expectation formulation, a different culture when dealing with the IoT. Reviewing the aforementioned notion of responsibility on these grounds would mean an unconditional need for choice: The display and transparent presentation of information as a demand to connected systems translates into two imperatives that Heinz von Foerster lays out elegantly.[50] Building on the insight that reality is a collectively co-constructed reference, affected by both social and technological systems, by Thou and I, where both sides are mutually constitutive and form a shared reference of identity, he states:

> *reality = community.*
>
> *What are the consequences of all this in ethics and aesthetics?*

The ethical imperative: *Act always as to increase the number of choices.*

The aesthetical imperative: *If you desire to see, learn how to act."*[51]

For us this means the active exploration of and with computing environments that allows for comparison, leeriness, and refusal of observed observations, driven by active learning, critique, comparison and a closer look. Building a responsible IoT would cater to both imperatives by allowing for choice (autonomy) and inspiring to see (decision-making). It maintains the observer's autonomy by marking itself un-neutral, intransparent, and subjective. It inspires to act, but offering decisions to be made, by expanding the field of view, by re-introducing the new and inviting to act upon it. This implies both an adjusted perspective byits observers (culture) and an adjusted structure of itself (design).

A new perspective

Starting with the former, as a first proposal, we would need an extended notion of the *thing* at hand for organizations, in order to measure up to its implications for communication. It ought to be treated with reservation and its output handled as a subject for discourse, thereby itself shifting into the role of a disputable offering of potential observations, a matter of debate.

An interesting take of such a debatable thing can be found in Bruno Latour's notion of the "ding"[52]: Building upon a draft for "object-oriented-democracy", as opposed to a discourse, that is stuck with the fallacy of assuming a "real" and factual ground for decisions, Latour calls for ousting such allegedly objective discourse in favor of a focus on (and acceptance of) a ubiquitous prevalence of subjective interest and prioritization. He calls for **a shift from the matter-of-fact toward the matter-of-concern**. An organization equipped with this lens would treat *things* just like it would treat all their observations: as preliminary truths, that either call for trust to be reliable, or are rejected once another, more trustworthy one is found[53]. It would be aware of the computational nature of its environment, and treat connected objects accordingly: not as objects but as subjects, not mystically enchanted but technically equipped with the capability of filtering, focusing, and proposing observable information. It would review a thing in reference of itself and other, similar things, and formulate a counter-computational assumption on where presented information is coming from and whether or not to trust it. It would thus account for more implied potentiality and interconnectedness, and recursive selection of relevance to provide "plentiful" options.

On the other hand, the IoT's equivalence of social culture can be described as structural design. Applying von Foerster's imperatives to it would call for an equally open, recursive, and *traceable* processing of information. Traceable to the observing eyes of staff, family members, citizens, and tenants. Interestingly, looking back at the history of the Internet of Things, such an opportunity seems to have once been more palpable: Ted Nelson's account of a mutually connected internet, with bi-directional links and a clear awareness of what is actual and what is potential marked such a visionary description of treating things for what they were, or rather could be.[54] [55] In his project Xanadu, he laid out a set of principles that would not only allow for comparison with self and with the other, but along the way, introduces a structural implementation of recursion. It's an ambitious project with an equally arduous history, that does, however, shed some light on the potential of an alternatively designed Internet of Things.

While judging from to today, the opportunity of a Xanadu-internet seems to have passed for now. It is, however, well up to future research if and in how far such system might shed new light on von Foerster's call for **trust over truth**, and a vision of a collectively opened discourse on the *things* around us. A contemporary opportunity that technology itself holds might be the emergence of the blockchain: By once more re-introducing the implicit relevance of trust on a technical (that is structural) level, a blockchain driven IoT might allow for reliable traceability of information, by incorporating entrained historical data into a decentralized system of shared memory. We should thus expect to find more hints toward a re-claim and granted responsibility in von Foerster's sense when dealing with computing environments – for observation the IoT is a mutual one.

Footnotes

1. Niklas Luhmann, *Die Gesellschaft Der Gesellschaft*, 1. Aufl. (Frankfurt am Main: Suhrkamp, 1997).
2. Heinz von Foerster, On Constructing Reality. p. 211-227. In: *Understanding Understanding. Essay on Cybernetics and Cognition*, (New York: Springer, 2003).
3. Mark Weiser, "The Computer for the 21st Century," *Scientific American* 265, no. 3 (September 1, 1991): 66–75, https://doi.org/10.1038/scientificamerican0991-94.
4. Dirk Baecker, "Digitalisierung als Kontrollüberschuss von Sinn," in *Digitale Erleuchtung: Alles wird gut*, ed. Zukunftsinstitut (Frankfurt a M.: Zukunftsinstitut, 2016).
5. Bruce Schneier, "Click Here to Kill Everyone - With the Internet of Things, We're Building a World-Size Robot. How Are We Going to Control It?," New York Magazine, January 27, 2017, http://nymag.com/selectall/2017/01/the-internet-of-

things-dangerous-future-bruce-schneier.html.

6. Purpose in the sense of higher-level predictability as laid out by Arturo Rosenblueth, Norbert Wiener, and Julian Bigelow, "Behavior, Purpose and Teleology," *Philosophy of Science* 10, no. 1 (January 1, 1943): 18–24, https://doi.org/10.1086/286788.
7. Weiser, "The Computer for the 21st Century," 82.
8. The metaphor of the enchanted here is not as far-fetched as it might seem: Not only since formerly passive things now seem to proactive act and behave out of themselves, but also as a reference to the mystic connotation of any new technology, that remains inexplicable, but all the more momentous to its observers – withdrawing from control and invoking a sense of elusiveness. Both such traits, control and elusion, can in fact put us on the trail of the IoT.
9. Bruce Sterling, *The Epic Struggle of the Internet of Things* (New York: Strelka Press, 2014).
10. W. Ross Ashby, "Requisite Variety and Its Implications for the Control of Complex Systems," *Cybernetica* 1, no. 2 (1958): 83–99.
11. Brown, G. S. *Laws of Form*, Julian Press, New York, 1972, p. 2
12. Humberto R. Maturana, "Everything Said Is Said by an Observer," in *Gaia, a Way of Knowing: Polit. Implications of the New Biology*, ed. William Irwin Thompson (Hudson, NY: Lindisfarne Press, 1987), 65–82.
13. Heinz von Foerster and Bernhard Pörksen, *Wahrheit ist die Erfindung eines Lügners: Gespräche für Skeptiker* (Heidelberg: Carl-Auer Verlag GmbH, 1998).
14. von Foerster, *On constructing Reality*.
15. von Foerster, *On Constructing Reality*. p, 227
16. Bruno Latour, "From Realpolitik to Dingpolitik or How to Make Things Public," in *Making Things Public: Atmospheres of Democracy*, ed. Bruno Latour and Peter Weibel (Cambridge, Mass.; [Karlsruhe, Germany: MIT Press; ZKM/Center for Art and Media in Karlsruhe, 2005).
17. cf Dirk Baecker, *Studien zur nächsten Gesellschaft* (Frankfurt am Main: Suhrkamp, 2011).
18. Theodor H Nelson, *Computer Lib; Dream Machines* (Redmond, Wash.: Tempus Books of Microsoft Press, 1987).
19. For a reflective take on project XANADU and its potential implications for today's IoT, see Usman Haque's contribution to the first edition of the RIOT Report. See: Haque, "How Might We Grow Diverse Internets of Things? Learning from Project Xanadu & the WWW", 2017, ThingsCon - The State of Responsible IoT, at https://medium.com/the-state-of-responsible-internet-of-things-iot/how-might-we-grow-diverse-internets-of-things-learning-from-project-xanadu-the-www-47a64497750c

Simon Höher is an entrepreneur and strategy consultant who works with public and private organizations alike. With his clients and partners, he explores and maps mutually desirable futures, and develops human-centered strategies to get there. He has a background in Cultural and Political Science and currently studies Philosophy, Politics, and Economics at Witten-Herdecke University. He co-heads CURRENT COLLECTIVE, a research and strategic design agency, and chairs ThingsCon, an initiative to promote a responsible Internet of Things. Earlier he co-founded two companies around digital collaboration and critical design and worked with various organizations in the field of technology and international development throughout Africa and Europe. In his work and studies he explores the interplay of technology, culture and society in a global context, and is particularly interested in the impact and function of collective utopias from a sociological and cybernetic point of view. His research evolves around ethical, political, and economic perspectives on society, as well as concepts of robust systems design and critical innovation. Simon mentors at Seedcamp – and shares his insights and questions as a speaker and coach. He is based in Cologne.

The Manifesto Moment in IoT

by Ester Fritsch, Prof. Dr. Irina Shklovski and Prof. Dr. Rachel Douglas-Jones

The ThingsCon report The State of Responsible IoT is an annual collection of essays by experts from the ThingsCon community. With the Riot Report 2018 we want to investigate the current state of responsible IoT. In this report we explore observations, questions, concerns and hopes from practitioners and researchers alike. The authors share the challenges and opportunities they perceive right now for the development of an IoT that serves us all, based on their experiences in the field. The report presents a variety of differing opinions and experiences across the technological, regional, social, philosophical domains the IoT touches upon. You can read all essays as a Medium publication and learn more at thingscon.com.

This text plays a special role in that contains a meta-analysis of the previous State of Responsible IoT (2017). It is based on a CHI paper that examined IoT manifestos and, among many others, also the contributions contained in the ThingsCon State of Responsible IoT 2017 edition.

Across Europe, designers and developers of IoT are calling for a revolution. Growing unease with the present state of IoT investment, hype and direction has brought forth reflections about what technological ubiquity means in practice, and what the role of designers and developers should be in creating our common technological futures. Concerns vary widely, but in the past few years, IoT networks, design studios and organizations have started to write down their concerns in manifestos. Framed variously as design principles, statements on ethics and responsibility, our analysis of 28 IoT manifestos shows that the manifestos mark a specific point in the discourse of ethics of IoT, a moment when the promised technological future is faltering.

Why Manifestos?

Why would you write a manifesto? Manifesto writing is polemical, it is political. Manifestos have had a role in political and design developments of the twentieth century, inviting comment and engagement. The manifesto is the transformational style chosen by designers and developers of new technologies to express their dissatisfaction with the status quo, and to imagine different futures. In our study of

28 different manifestos about The Internet of Things[56], we explored what might be appealing about this long established, rousing format, or as the literature theorist Caws called it, a "loud genre" (2001:xxix)[57]. The 28 texts we analysed (full list below) were drawn from the European IoT scene, and included the 2017 RIOT report, design manifestos, maker movement documents, and network statements. While styles vary considerably across these texts, in our study we defined a manifesto by the two major rhetorical moves it makes. The moves are recognizable: manifestos first define the present and identify the problems with it, they then define how a better future should play out.

Manifesto Moves

The first move–defining the present–is a challenge for manifesto authors. The majority of texts that we analyzed described a world of technological ubiquity, a present of past futures. Ghosts of the optimistic past haunted the texts, where visions of a future full of ubiquitous technology would produce a world where lives were made easier by near invisible computational assistance. As authors Genevieve Bell and Paul Dourish noted in 2007, while this ubiquitous computing world was meant to be 'clean and orderly' it has instead turned out to be 'messy'[58]. The messiness of the present takes multiple forms in the manifestos. IoT devices are problematically invisible, and the politics of what they record, collect and transmit is an evident concern. How can the data they collect be known, and by whom should it be known? Mechanisms of data transfer, from its security to its frequency, also appear as concerns. There is a growing realization among manifesto authors that IoT devices are embedded in, and dependent on, a range of existing infrastructures. These layered problems form the descriptive basis for an ethics revolution for IoT.

But what will this revolution look like? Having described a present full of anxiety and confusion, the next move of a manifesto is to shift into predictive mode. Here, authors become part of reshaping the problems they have identified, putting forward possibilities, commitments, and new norms. This is particularly visible for IoT manifestos, because the people to whom they are addressed are the designers of these futures. As they write, authors invite their readers in to a common future, to think through the basis of future IoT design. Beyond 'thinking through', readers are sometimes invited to 'sign' the manifesto, indicating agreement that they will seek to orient their future work by the principles and visions it contains.

IoT Manifesto Themes

Four strong themes emerge across the different texts: transparency, openness and sustainability and responsibility. These are the sites where manifesto authors direct both their concerns, and their intentions for future change.

Transparency takes on two meanings in the documents, both related to the way an IoT device should be transparent. The first meaning emphasizes the need for consumers to know how IoT devices work and the second the need for designers to be explicit in their design choices. The first argument runs that without informed knowledge of what devices do, or might do, and what kind of data practices the company that made the device has, consumers cannot make informed decisions. The second, emanating from designers themselves, is that the control of data is not (and should not be) part of what a user has to control. Instead, the manifestos argue that designers should make evident how products work in a way that is accessible to someone using the product.

> *'Not having transparency into how the technology is working, making decisions, literally moulding our perception of the world, is inherently political.'*

(Robbins (RIOT)

These calls for transparency blur with manifestos that take a stance on openness. **Openness** might be open hardware, as in the Arduino manifesto, but it might equally be activities open to the public, organized around principles of creating a community of equal users, as in the Open IoT Studio's text and the Dowse manifesto. Openness is raised across a set of manifestos as a way of democratizing control over the making process and the data collection process.

Sustainability, discussed in more than half of the texts, emerges as an ethical concern about the production of IoT: both hardware and software. Manifestos sound an alarm about the environmental sustainability of materials used in IoT production, pointing to the limits on precious metals and the environmental costs of production. At the same time, the short lives of many IoT products concern authors: are they making things that will be out of date, un-useable or unsupported within a year?

Locating Responsibilities

Where should responsibility for these concerns be located? As they define their vision of a better future, all of the manifestos try to shift the hyped conversation from what is possible in IoT to what is responsible. Questions common across the 28 manifestos include:

• Should citizens become responsible for understanding the world of IoT (through becoming more educated) or should designers take greater responsibility for designing devices that communicate better how they work?

• Should designers consider the IoT ecosystem beyond the specific device they are working on, and consider the possible future effects of their designs?

• Should communities, networks and organisations hold one another accountable? How?

Conclusion

Manifestos get attention. They have been used throughout the twentieth century to articulate clear positions and agitate for change. Unlike the modernist manifesto, however, the cautionary manifestos of IoT offer not a single better future but designs for multiple possibilities. Across the design and development space in Europe, the manifesto moment is one of uncertainty and numerous different positions on what a 'good' IoT future looks like. As a result, the manifestos reflect considerable uncertainty, one where people are still trying to 'figure things out'. It remains to be seen what will happen when priorities or values conflict, when those who author manifestos look for the life of them in practice, or when designers and developers try to use these documents to reshape their working lives. The broader research project that this paper forms part of, Values and Ethics in Innovation for Responsible Technology in Europe (VIRT-EU), aims to be central to these conversations, and we are in the midst of ongoing field research, legal and social network analysis of IoT across Europe. Manifestos provide a window into a particular moment not only in social values but also of the documentary power of declaration, as a community is built around the project of figuring out what kind of technological present – and future – we want to live in.

MANIFESTOS

The following 28 documents constitute our corpus for analysis with short-codes marked in brackets.

[RIOT]

ThingsCon. 2017. RIOT. The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin. http://thingscon.com/responsible-iot-report/

[Deschamps-Sonsino, RIOT]

Deschamps-Sonsino, Alexandra. 2017. The Whole Internet of Things. The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 10-12.

[Krajewski, RIOT]

Krajewski, Andrea. 2017. User Centred IoT-Design. The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 13-21.

[Villum, RIOT]

Villum, Christian. 2017. Designing the Digital Futures We Want. The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 22-24.

[Dietrich, RIOT]

Ayala, Dietrich. 2017. Trust, Lies and Fitness Wearables. The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 25-32.

[De Roeck, RIOT]

De Roeck, Dries. 2017. On IoT Design Processes. The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 32-38.

[Scganetti, RIOT]

Scganetti, Gaia. 2017. The here and now of dystopian scenarios. The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 39-48.

[Robbins, RIOT]

Robbins, Holly. 2017. The Path for Transparency for IoT Technologies. The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 49-60.

[Smit, RIOT]

Smit, Iskander. 2017. Touch base dialogues with things: Responsible IoT & tangible interfaces. The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 61-68.

[Jorge, RIOT]

Appiah, Jorge. 2017. IoT in Africa: Are we waiting to consume for sustainable development? The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 69-73.

[Krüger, RIOT]

Krüger, Max. 2017. Expanding the Boundaries for Caring. The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 74-78.

[Thorne, RIOT]

Thorne, Michelle. 2017. Internet Health and IoT. The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 79-82.

[Bihr, RIOT]

Bihr, Peter. 2017. We need a more transparent Internet of Things. The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 83-87.

[Kranenburg, RIOT]

Van Kranenburg, Rob. 2017. How to run a country (I know where that door is). The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 88-91.

[Burbidge, RIOT]

Burbidge, Rosie. 2017. Design and branding: what rights do you own and what pitfalls should you watch out for? The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 92-97.

[Haque, RIOT]

Haque, Usman. 2017. How Might We Grow Diverse Internets of Things? Learning from Project Xanadu & the WWW. The State of Responsible Internet of Things (IoT). Published by ThingsCon, Berlin, 98-102.

[Ethical Design]

Balkan, Aral. 2015. Ethical Design Manifesto. Retrieved July 6 from https://ind.ie/ethical-design/.

[Doteveryone]

Doteveryone, 2017. Exploring what “responsible technology means”. Retrieved September 14, 2017 from https://medium.com/doteveryone/exploring-what-responsible-technology-means-4f2a69b50a61

[Dowse]

Dowse. Retrieved May, 2017 from http://dowse.eu

[Flaws Kit]

Flaws of the Smart City Friction Kit Version 1.3. October 2016. Designed by Design Friction. Retrieved August 8 from: http://www.flawsofthesmartcity.com

[Maker Movement Manifesto]

Hatch, Mark. 2014. The Maker Movement Manifesto. Mc Graw Hill Education.

[IoT Design Manifesto]

IoT Design Manifesto. 2015. Retrieved March 14, 2017 from https://www.iotmanifesto.com

[Open IoT]

Mozilla’s Open IoT Studio. 2016. Practices for a Healthy Internet of Things. Edited by Michelle Thorne, Jon Rogers and Martin Skelly. Published by Visual Research Centre, Duncan of Jordanstone College of Art and Design, University of Dundee.

[TCM]

Oliver, Julian, Gordan Savicic and Danja Vasiliev. 2011-2017. The Critical Engineering Manifesto. Retrieved June 23 from https://criticalengineering.org

[TOPP]

Topp Studio. 2016. R.IoT. Responsible IoT. Retrieved March 30, 2017 from https://medium.com/the-conference/responsible-iot-3-essential-iot-design-features-504ce4c62e77

[Uribe]

Uribe, Félix. 2017. The classification of Internet of Things (IoT) devices Based on their impact on Living Things. Retrieved July 15 from https://www.uribe100.com/images/Documents/classificationofiotdevices.pdf

[Apps for smart cities manifesto]

The apps for smart cities manifesto. 2012. Retrieved July 6, 2017 from http://www.appsforsmartcities.com/index.html%3Fq=manifesto.html

[Human(IT)]

The Human(IT) Manifesto. 2017. Accessible manifesto from World Economic Forum 2017: BlockChain, Ethics, AI, Humans and Shift Happens. Retrieved September 11, 2017 from http://dataethicsconsulting.com/en/world-economic-forum-2017-blockchain-ethics-ai-humans-shift-happens/

[Things Network]

The Things Network Manifesto. 2017. Retrieved July 7, 2017 from https://github.com/TheThingsNetwork/Manifest

[Arduino] Arduino 2016. IoT Manifesto. Retrieved May 5 2017 from https://create.arduino.cc/iot/manifesto/ Accessible from https://www.wired.com/beyond-the-beyond/2016/04/arduino-iot-manifesto/

References

1. Ester Fritsch, Irina Shklovski, and Rachel Douglas-Jones. 2018. Calling for a Revolution: An Analysis of IoT Manifestos. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18). ACM, New York, NY, USA, Paper 302, 13 pages. DOI: https://doi.org/10.1145/3173574.3173876

2. Caws, Mary Ann. 2001. The Poetics of the Manifesto: Newness and Nowness. In Mary Ann Caws (Ed): Manifesto. A Century of Isms. University of Nebraska Press, xix–xxxiii.
3. Genevieve Bell and Paul Dourish. 2007. Yesterday's tomorrows: notes on ubiquitous computing's dominant vision. Personal and Ubiquitous Computing 11, 2: 133–143.

Ester Fritsch holds an M.A. in Anthropology from The University of Copenhagen. Her research engages with complex ethical configurations that embrace laws, policy, humans, plants, technologies, data and other influences. She is curious towards how ethics emerges through relational practices unfolding in such hazy intertwinements indicating that ethics might not solely be a human affair, but a more than human matter. For the past five years Ester has explored this through empirical and conceptual inquiries into climate change, energy and agriculture in Denmark and Italy. As a PhD fellow in VIRT-EU she now seeks to understand how ethics is cultivated and circulated in European IoT ecologies and delves into how ethics is enacted among IoT developers as ethical subjects in continuous becoming.

Irina Shklovski is an Associate Professor at the IT University of Copenhagen. Although for her primary field as human computer interaction, her work spans a lot of other fields from computer science to sociology and science & technology science. Irina's research focuses on big data, information privacy, social networks and relational practice. Her projects address online information disclosure, data leakage on mobile devices and the sense of powerlessness people experience in the face of massive personal data collection. She is very much concerned with how everyday technologies are becoming increasingly "creepy" and how people come to normalize and ignore those feelings of discomfort. To that end she recently launched a "Daily Creepy" Tumblr to monitor the latest in creepy technology. She leads an EU-funded collaborative project VIRT-EU, examining how IoT developers enact ethics in practice in order to co-design interventions into the IoT development process to support ethical reflection on data and privacy in the EU context.

Rachel Douglas-Jones is an Associate Professor at the IT University of Copenhagen, Denmark. She was trained as a social anthropologist and STS scholar at the University of Cambridge, Harvard University and Durham University. Her research interests sit at the intersection of ethics, medical anthropology and

computational techniques. At ITU Copenhagen, Rachel teaches undergraduate classes on Society and Technology and a graduate studio class called Writing Innovation Studio. She also runs ITU's ETHOS Lab.