

Original citation:

schraefel, m.c., Gomer, Richard, Alan, Alper, Gerding, Enrico and Maple, Carsten. (2017) The internet of things : interaction challenges to meaningful consent at scale. Interactions, 24 (6). pp. 26-33.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/94295>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Copyright held by authors.

Posted under ACM academic use policy.

Published version: <http://interactions.acm.org/archive/view/november-december-2017/the-internet-of-things>

A note on versions:

The version presented in WRAP is the published version or, version of record, and may be cited as it appears here.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk



The Internet of Things: Interaction Challenges to Meaningful Consent at Scale



m.c. schraefel, University of Southampton
Richard Gomer, University of Southampton
Alper Alan, University of Southampton
Enrico Gerding, University of Southampton
Carsten Maple, University of Warwick

Insights

- To make consent meaningful, we need greater “apparency” of how data is being used.
- For the scale and speed of the IoT, apparency/consent decisions will need to be automated on our behalf.
- There is rich potential to create human-centered, nuanced services that can leverage the negotiation of consent/sharing terms.

Remember the scene in Minority Report where Tom Cruise’s character walks through a mall and is met with a barrage of ads? Somehow the stores have his data and can deliver customized ads that today’s social media campaigns can only dream of. We find the bombardment hideous. This is the predicted future of the smart home, city, and vehicle—or rather, the dystopian view. There is an alternative. Imagine walking through these same areas without the ads because, like the entourage of a celebrity, you have agents working on your behalf to fend them off, while making sure your mom—or other important or personally interesting information—can still reach you. After all, you’re not against sharing some of your personal data. Your agents know you’ll be happy to share your photos of local dogs with an SPCA citizen-science project about strays in town, where every

100 images gets you a tax credit. But you’re not sharing these with pet shops—you don’t have a pet. And they don’t need to know that. Right now, many feel like personal data use is heading toward the Minority Report dystopia. HCI and AI are well placed to reimagine personal data sharing as a more equitable, negotiable, and sustainable practice than either today’s take-it-or-leave-it, agree-or-not approach or tomorrow’s dystopian disempowered bombardment. Research and design in HCI will help create that better future. Indeed, this article focuses on how, with the Internet of Things (IoT) about to explode, it is essential that HCI, along with AI, embrace this space, and design it from a human-centered, human-valued perspective. We propose several concepts and questions to help envision the IoT as such a consensual, human-centered space.

The Internet of Things, we are told, is about to achieve epic scale, with Cisco and Ericsson (Dave Evans and Hans Vestburg, respectively) having predicted that there will be 50 billion devices connected to the Internet by 2020 (though each have since revised their estimates down to 30 billion and 28 billion). Their ecosystem is wildly heterogeneous. Many devices will be capturing the same identification data over and over; others will be part of networks sharing data, such as cars moving through various jurisdictions and associated infrastructures with their own terms for data sharing. One of the key concerns of the IoT and its high-speed cousin, the Internet of Vehicles, is just how that data may be captured and shared, not only within one fixed environment like a home, but across environments, from the wired High Street (like that shopping scene in *Minority Report* mapping ads to eyeballs to associated customer profiles) to moving from one location to another. In other words, as the objects of our environment become more connected to the Net, do we simply become another thing on the Net, reducing our privacy, and our civic values, in the names of everything from convenience to counterterrorism? Likewise, given the vast scale, speed, and heterogeneity of this new ecosystem, are we creating new risks to our personal and national security, both as citizens and as societies—not even from willful hacking but just because the scale of our IoT reach will exceed our grasp of all the necessary protections that we assume are ours in a civil society?

In response to the current data status quo and in anticipation of this changing ecosystem, new rules for data sharing have been established. For example, the General Data Protection Regulation (GDPR) approved by the EU Parliament on April 14, 2016, will be enforced across the Union (and enshrined in law by the U.K.) starting May 25, 2018. *These regulations have been put*

in place well in advance of the technical means to sustain their implementation. Consequently, there is a key moment for HCI and UX research and design to influence society for good, not just to design wonderful devices for the IoT but also to consider the wicked problems of how to make apparent for developers, designers, businesses, policymakers, and citizens the mechanisms and personal-data-driven assumptions that enable the IoT. By this deliberate engagement we can help surface the implications of data sharing in order to develop models of understanding, social expectations for understanding trade-offs, and means for developers to know their designs comply with these expectations and for citizens and policymakers to be able to trust that they do. Because of our expertise in understanding both technology and especially human-centered approaches to design, we have this key role to play in informing the shape of these exchanges and to create an ecosystem that supports social, technical, *meaningful* consent at IoT scale.

To better see these opportunities for the future, let's consider the status quo of data-sharing consent in the current digital economy.

CITIZEN CONSENT IN DATA SHARING: THE MARGARINE OF CONSENT

The Internet has made liars of us all. No one has used a browser, a social media site, or a smartphone app without encountering a box that says, before continuing, that A) we have read the terms and conditions of a service and B) we agree to them. We click “agree” when we haven't a clue if we do or don't. In other words, in a world where our personal data is largely the oil that greases the wheels that keep the Internet running, we have very little meaningful say in what data is collected and how it's used, and why we may wish to limit it, or for that matter, give it away in buckets. Given the current (some might say insulting) approach

to consent, this status quo is not surprising. Regulation, unfortunately, has to date created only a veneer of consent—a legal illusion of choice and control—but design has not delivered interactions that support genuine informed choices that regulators assume data-harvesting services should follow.

The stats are very clear about how broken terms and conditions pages are: If we *were* to read these “agreements” it would be nearly a full-time job [1]. More disturbingly, these terms are generally written to require a sophisticated level of reading comprehension—beyond the norm for the population using these services [2]. So even if we did read these T&C, most of us wouldn't understand them. When we give our consent under such conditions, it isn't meaningful.

We are also asked to consider these terms at exactly the moment work in HCI on task interruption has shown is the wrong time [3]: when that request gets in the way of our primary task. We click the “agree” button because clicking it gets rid of the screen so that we can get on with posting our cat video or uploading a draft of our paper to a co-editing site or synchronizing our calendar with a cloud service.

Many of us became acclimatized to this meaningless box clicking around installing software: *Yes, yes we don't own it, uh huh we're just leasing it, and no we won't make copies of it. Sure.* In those days, however, it was rare for software to call home to the mothership to locate our particular copies. Now, it is commonplace for software to be deployed as a service that knows exactly where it is and how many copies have been authorized. But that service, especially when deployed on phones or mobile devices, gathers far more information for very amorphous reasons than just registration confirmation. More troubling, as has been shown when installing apps on phones: Few people are even aware that the app is (re)setting permissions to access personal data not needed for its operation [4], such as our contacts and text messages. There is also the belief that this personal-data capture is a trade with the developer, and that if one pays for the app then that data trade is closed. Not so. Some apps take even more liberties in the paid version. Consent is not meaningless in this context—it's nonexistent.

In the app case, research suggests that few people are aware that data-

We click the “agree” button because clicking it gets rid of the screen so that we can get on with posting our cat video or uploading a draft of our paper.

access permissions can be set per app to limit access to that data [4]. And why would a person consider whether it's OK for an app to access some of their data if they have no awareness this data is being accessed in the first place? Indeed, the situation is not much different on the Web. For people in some areas of computing, we may take it for granted that unless we use services like ad blockers or virtual private networks, we are being tracked across interlocking webs via mechanisms like fingerprinting and cookies. For instance, every time we put a URL into a social media feed and it is shortened by that service, that URL reflects its path through the network—who has used it, who has looked at it, where they've gone after visiting it, and so on [5]. Our social networks and beliefs are effectively exposed. New research-based services like TrackMeNot (<https://cs.nyu.edu/trackmenot/>) run randomized Web searches from our browsers in order to confuse this profile that is constructible from our footsteps through these pathways of the Net. Problematically, however, the few studies that have looked at how tracking is perceived show that only a small number of people in the general public are aware of the degree to which they are being followed online, or that their Internet traffic is being shared among various, mostly commercial entities. The data suggests that when people do learn of this tracking, they characterize it as creepy [6] and want to find ways to control it [7].

We see this awareness effect in other data-related transactions: Once people are aware of what is happening to their data without their consent, they demand better conditions. What of the privacy paradox, then? In other words, that people say they're concerned about privacy, but if you put a form in front of them and request personal data, they readily hand it over. As more researchers have now shown, this response is not a contradiction. We are a sense-making people: If we are asked for something—especially tied to something we want—we assume there must be a rationale for it. We assume the best. When responses are probed, however, many people who provide very personal data to a service do so without a clear model of how that data may be used by the service itself; how that data may be used by other people accessing that service; what of that data is actually

necessary for the service to function; and the risks associated with sharing that data. We are busy: It is easier to trust there is a good reason for this data request, it seems, than to stop and check if we're being scammed. Indeed, we need only consider the outrage when those who do stop and look raise a red flag about terms and conditions. Doing so, however, has required the work of what we might call social interpreters to translate the language of the revised terms and conditions, moving it out of the abstract and into concrete terms that are meaningful to people. These changes otherwise remain opaque, again making our consent socially meaningless.

TOWARD APPARENCY AND SEMANTIC/PRAGMATIC TRANSPARENCY

Just from the above scenarios, we can see numerous opportunities for interaction research and UX design to change the status quo around data consent. Fundamental to any change, however, is to see a need for it. This is what we've been calling *apparency*. One may have very well-defined terms and conditions, but if people don't even know that their contacts are being accessed by a puzzle game they downloaded, if this use is not *apparent* in the first place, transparency about the terms of an unperceived process is at best meaningless.

As designers, we can help to develop the means to make such data processes apparent in order for the terms to be meaningfully transparent. In the context of ubiquitous computing, Matthew Chalmers [8] framed making the properties of a system apparent as “seamfulness,” as opposed to seamlessness or, more particularly, sameness. For instance, rather than hiding which cellphone tower a phone may be using, it might be better to make this information available. Some people might find it useful and empowering: Being able to look under the hood of a system at various levels of detail, specifically in order to engage with it and change it, is a valuable property.

In data-driven services—like most of those on the Internet—one can point to the terms and conditions and label them as either transparent or opaque, based on the language used and the specificity of descriptions. But such transparency refers largely to only an acknowledgement that data is being

collected and that it may be used to “improve the quality of the service”—as cookie notices on websites in the EU constantly assert without explanation of what or how, exactly. Apparency would seek to make those connections clear and traceable toward meaningful transparency. For example, there are no cues to the user of a downloaded game that make it apparent that there are personal-data settings associated with this app and that changing them (or not) will have an effect on risks of burglary (GPS access), identity theft (contacts access), workplace harassment (enabling anyone online to see pictures from social occasions), job-selection discrimination (social media commentary being available), or preferential or discriminatory pricing [9]. Nor is it readily apparent that shared data is churned into use for targeting advertisements, not only on the site where the data is initiated but also from that site to other sites, and through a network of brokers and advertisers, as a person surfs the Web [5]. The simple act of touching these sites is of course itself valuable data that is both *unapparent* and *untransparent*.

Indeed, we might reframe a progressive scale from apparency to transparency, in which we have apparency, semantic transparency, and pragmatic transparency. Let's call it *apparency to s/p transparency*. Apparency reflects how an activity—in this case a data activity—is signaled. Semantic transparency addresses whether we know what the terms of the apparent activity are and mean; pragmatic transparency reflects the degree to which we know what these data actions actually do or entail.

There are already lovely examples of apparency to s/p transparency design online. One elegant, motivating example is the very simple HTTPS protocol. That S makes a transparent process unobtrusively apparent: that the connection between you and a website is secure and encrypted, that the data is not out in the clear for anyone to see. Increasingly the S is backed up by a padlock icon in the browser's address bar to indicate a secure channel. If one is unfamiliar with the padlock, clicking on it usually displays text to make more of the semantics of the process apparent: that data is being transmitted over an encrypted channel. For pragmatic transparency, these

claims can be explored and tested. There is a certificate that can be verified regarding the claims made by the S and the padlock. These are signifiers of apparency, seams that can be exposed and tested in terms of semantic and pragmatic transparency. We can decide how far we wish to probe those signifiers, but with them, the resources are there to make a more informed judgment about the channel. The padlock is an elegant, apparent expression that makes the semantic and pragmatic transparency of a binary state richly available.

Apparency for the properties that would inform a consent decision are more nuanced, more variable, and potentially more dynamic. A challenge we set ourselves as a research team is how to raise apparency about one's current appearance on the Web, in particular to online trackers that have an interest in creating a picture of who you are for various purposes, from targeted ads to offer discrimination, which includes everything from job offers to insurance pricing. These impressions are based on one's clicks from one Web resource to another. Our challenge has been to find metaphors to express what this tracking means in an apparent, semantically and pragmatically transparent way. The approach we've been testing is called the Web Mirror.

AN APPARENCY EXAMPLE: THE WEB MIRROR

There have been third-party efforts, such as Mozilla Lightbeam and Disconnect.me, to make our traces through the Web and what sites track us more apparent by using network or spring graphs of trackers. In pilot tests with participants, these often engender a "Wow, what a big graph that is!" response, but few people use them, and, interestingly, the follow-up question of how to make it stop rarely comes up.

In an effort to help schools in particular teach students about protecting themselves online, we have been piloting a project with teachers called the Web Mirror (<http://mirror.websci.net/>). Here, we show students not an abstract graph but rather a "Web reflection of you." That is, we show them what the various trackers they've touched see of their Web history. We use topic extraction to infer what the interests of someone visiting those websites could be, and prompt them to ask, "What

could my browsing history say about me?" Our goal right now is to see if this mirroring back to students of what their browsing may portray about them helps them first to perceive that their browsing history is their personal data (apparency); that others are processing that data as the students move through the Web (semantic transparency); and how that data can be used to create a variety of pictures about them (pragmatic transparency). From this awareness, we are keen to empower them to control that reflection—in other words, to action consent. This is done by connecting the students back to how those reflections can be changed using the current means for proactive personal-data management (or consent management), which means cumbersome tools like ad blockers and VPNs.

TIMING

As stated earlier, we know from HCI research on interruption that when we're asked to consider anything that takes us away from our primary task, it's simply not going to get our full attention, especially when it's something as abstract as data permissions or terms and conditions. Just get out of our way! Beyond making a reflection of ourselves from our Web travels apparent, a key insight from the Web Mirror work is that there is high apparency value in making the revelation of what personal data is desired by a site/app/service, and whether or not they should have it, its own task in its own time.

When we look, there are multiple examples of such asynchronous transactions all around us in the physical world. Consider making purchases. Each time we withdraw cash from a bank machine or use a debit or credit card, we get a receipt of the transaction—and that's about it. We are not asked to review our purchasing history at the time of the transaction. Instead, we receive a monthly statement both as a record of our spending and debts, and as a log we are encouraged to review in case of errors. That monthly statement itself is a review process, but it is a data trail of what has happened with our various assets, from cash on hand to credit lines. The statement, however, along with our receipts, fits into a larger practice of personal money management, including tasks like setting a budget, saving for a purchase, investing, and so on. Insurance

purchasing is part of a similar genre of practices where we consider the terms and conditions of a policy as best we can, well before we actually need the policy—in fact, it's required that we have a policy before we need it. We may consider the terms and prices of the policy, if these change, before we renew, and then start shopping around again.

In other words, for many kinds of transactions, we have established practices to review attributes from the transactions as well as the terms and conditions. These reviews fit into a larger mechanism that informs our quality of life, from how we manage debt with financial planning to how we manage risk with insurance provisions. A key point, however, is that even though not all citizens practice such fiscal hygiene, the data is there to enable those processes. Such is not the case for personal-data transactions on the Internet. Surely in HCI we can draw on these analogous practices to better design our engagement with the terms and conditions of data consent, and with auditing consent transactions?

NEGOTIATION/AUTOMATION

Of course, one of the reasons for reviewing our financial transactions is to see if the terms of service are fair. After all, when we agree to terms and conditions, we engage in a contract with the supplier. In the data-driven world, however, these contracts are one way and binary: We as the consumers of the services can say only *yay* or *nay*. Sometimes, saying *nay* can feel impossible: If one's whole community is making use of a service, it's hard to be the lone holdout.

Once again, if we turn to real-world examples, negotiation is a key part of just about any other agreement of exchange between parties. We negotiate everything from our contract with employers or staff to our fee for network access. Many of us can't walk out of a shop without either talking a price down or haggling for extras at no cost. Negotiation is ubiquitous—except on the Internet. Why?

We have been exploring how we might be able to automate consent in terms of negotiable data-sharing preferences using autonomous agents [10], and thus begin to create richer, non-binary terms for data exchange and service provision. In this approach, a person can say under

which conditions or for which types of services they may be willing to share their text messages but not their images, their browsing history but not anything else, and so on.

In our studies we see that people are willing to share more data on average when they can negotiate the data-sharing terms. Our studies also demonstrate that a negotiation-oriented approach to permission management better enables people to align their data-sharing practices with their actual privacy preferences. Our recent work (in submission) perhaps not surprisingly shows that permissions are not sufficiently context-sensitive for meaningful consent: Sharing photos is far too broad; sharing photos of public spaces with health services is more appropriate. Being able to trace and retract those images is also important.

We have touched on only two designs around consent: 1) our work with mirroring back a Web tracker's reflections of us, and 2) offering asynchronous opportunities to set responsive terms about sharing conditions to automate consent. There are many more mechanisms HCI designers can offer to support richer, more nuanced engagement with a data-sharing ecosystem. It's important to be clear within these design explorations that there is a distinction between privacy and sharing. People are not averse to sharing some personal data. Much to our surprise, we often found people keen to share data (sometimes their friends' contact information but not their own) in exchange for services when they understood the terms *and* they had a say in that exchange. In other words, where apparency to s/p transparency was supported, data-sharing quality has improved and often increased. Likewise, not all businesses are driven to grab out on personal data wherever possible: At our workshops with researchers, policymakers, and industry members, we have been delighted to find that some businesses would like to see how a nuanced data policy for negotiating these terms could work for new services and be a business differentiator.

THE INTERNET OF TERMS AND CONDITIONS (OF THINGS)

We already experience what has been called consent fatigue when we are regularly asked to agree to effectively



meaningless terms and conditions. Likewise, when terms and conditions change—and we see such notices—it seems gratuitous to ask us to say we agree to new terms when, what is our choice if we do not yet we wish to buy an app from a developer that is available only through this one vendor site? The number of times UK/EU citizens see “this site uses cookies”—when there are no options not to accept them—has caused more annoyance than engagement. The current state of the art for consent, therefore, is meaningless consent. But at least we might say we are asked. We see a screen. We hit a button.

In the Internet of Things (IoT), the predicted number of devices that we will encounter in our homes, on the way to work, at work, at play, and on the road is, to use a biblical term and all it entails, legion. In the IoT, every fridge will know your name, but many things will not have interfaces through which consent can be requested and given—or not.

Interactions will be handed off from one infrastructure to another. As we move between districts, our consent may be either assumed unnecessary or implied as given, yet the data terms and conditions that apply when moving from one infrastructure to another—and the guarantees of data protection—may be different. Many may recall the problems of Google Street View taking pictures with identifiable people in the images: No consent was obtained. Likewise, there are recent examples in subways, malls, and museums in the U.K., and airports in the U.S., where the MAC address of a mobile phone is tracked without any requests for consent or

any options to shut off access to this information that was never intended for these purposes, the assertion being that MAC addresses are not tied to the individual. Both legal and technical experts [11] would argue that this assertion is at best dubious, and further that such data is all too easy to combine with one or two other seemingly innocuous data bits to de-anonymize someone.

These kinds of seemingly anonymous though personal data-tracking contexts are key examples of how ICT/HCI expertise can help shape policy: We are part of the dialogue around identifying the art of the possible for interactive technology to support citizens' well-being. Ours is the community with insight into what is possible for interactive technology to do now, or in the near future, to help shape approaches to laws for individual and social interaction.

Fundamentally, if, as a civil society, we assert a right or belief or ethical principle that we have a stake in the use of the data we generate as citizens, and a right to privacy around our metadata, whether about what we read or where we sit to read it, then we need mechanisms that can negotiate our consent on our behalf at both IoT scale and IoT speed. It is eminently possible to build such infrastructures of consent. But for them to be effective, HCI has a key contribution to make to ensure that the approaches are both meaningful and sensible. Here we discuss just a few of the questions HCI research can help address.

How do users model IoT apparency and s/p transparency? The IoT is still

largely terra incognita. We mainly hear about its failures. For example, IoT devices have been hacked to create denial-of-service attacks on domain-name servers, thus cutting off Internet communications [12]. Or you may recall the TVs that track every word we speak in their presence, where the terms and conditions say this is all fair game.

Where HCI can lead is to develop scenarios of interaction and models of people's understanding of IoT interaction. Plainly, without understanding what we as citizens think is happening with these devices—and in particular with the data we enable them to capture—we cannot design safer, more usable experiences. For instance, *trust* and *risk* are concepts often discussed as putting the IoT project itself at risk. People may wrongly trust a service, such as online baby monitors, when the user experience maps to one's expectations: *Look, I can see my child*. Success. But there could be a lack of what we might call *risk apparency* in what was happening with the data being made available. What if the service managing the connection between the camera and one's phone was snooping? Here the perception of trust is inappropriately high, and of risk, inappropriately low. How interaction design can help connect with potentially autonomous agents to help users come to more informed understandings of these systems' interactions within an IoT ecosystem is a new kind of interaction design challenge.

Groups like Consentua (<http://www.consentua.com/>) have also been working on mechanisms for systems developers to be able to collect and respond to individuals' consent, so that developers do not have to reinvent the consent wheel and can take advantage of interactions that have been designed and refined to deliver a high level of "consentfulness." There is much fundamental work just to begin to map out the possible scenarios across multiple systems, and from there to

explore how real people understand these systems and think about risks. As these are made apparent, the work will extend to the options they wish to have to address them.

Where do users' and designers' models diverge, and how should we design for this? There is fundamental work to be done to engage with industries that want to deliver smart homes, cities, hospitals, and cars to connect those aspirations with citizens' understanding and expectations. In 2016, Nest "bricked" its smart hubs, devices that cost roughly \$300. It was a blunt lack of apparency toward consumers who had made a purchase in good faith—and who believed, based on experiences with other devices in their homes, that Nest would continue to function as long as it was turned on, like a lamp or a router. Would a smart car be shut off by its developers if it too were construed to be always and only the property or IP of its service provider? This is a new model of how we think about physical devices. While we may be accustomed to buying software licenses, this approach to hardware is unexpected; when we buy something physical we are used to owning it. We need to explore whether we need new design languages or at least new semantics to signal these new properties—not just to accept them but also to be able to make choices about whether we wish to invest in them or to negotiate their terms. We also need apparency around the data flows between devices in these environments in order to understand, agree or disagree with, or change them.

For instance, the majority of fitness trackers are tied to a particular vendor's software service. The vendor accesses all the data. They may have APIs to enable other services to access that data as well under certain terms, but then both services have an individual's data. The individual cannot simply buy the hardware and set up their own software server to track their own data.

Likewise, a software service cannot come along to map to the hardware in order to create an open data repository of step counting from any tracking source. This could act as a public good or a research archive, or could support citizen science to explore who does the most stepping in what age group at what time of day. If we buy the hardware, the current model allows us to talk only to its software service or cloud. One might call this a kind of consent choice, but again, when all trackers play by these rules, one is not choosing among data models but rather among colors. This example is just of one device. Apple's proprietary ecosystem with its home and health kits promises to be the infrastructure that bridges between devices, providing analytics and a common voice-enabled interface. The apparency to consumers of how data flows behind this convenience is largely occluded. We don't mean to say that the ecosystem is evil, but rather that without these flows and constraints being apparent, we cannot truly consent to our data flowing into the common pool of these ecosystems; innovation is more throttled than enabled.

Can users form adequate models of device ecosystems and their infrastructure? In order to make meaningful regulations, policymakers need high data apparency to be modeled as part of the data flows across intersecting or competing infrastructures or local ecosystem boundaries. By way of example, right now if we go from one coffee shop to another to access the Internet, we may be asked at each one to sign in and agree to the terms and conditions. Our access to sites at these locations may be faster or slower or perhaps time-limited, but the experience is largely similar. And yet, without data apparency, it is impossible to tell whether different data is being captured and what additional tracking is being added to the sites we visit.

There is in these interactions a lack of another type of apparency signaling—what can happen to our data over time. To return to the apparency of the padlock icon, it signals a steady state process: The channel is either secure or it isn't. There is only a *now* to that signal. But our captured data can be so multipurpose and can contribute to so many other ways of constructing

Without understanding what we as citizens think is happening with these devices, we cannot design safer, more usable experiences.

a pattern. Time—in particular the future—is an unmet, open challenge. For example, in our current negotiation models, we test only those conditions where the data stays within a particular time frame. But what happens to collected data when the company is sold or closes and sells off its assets? The EU GDPR tries to take these kinds of future-proofing scenarios into account, but how do we represent these decisions in terms of interactions, from apparency to semantic and pragmatic transparency, to help consumers, businesses, and policymakers make choices? With rich apparency to s/p transparency interactions in which time is one of the variables to make apparent, consent for data use can be far more meaningful.

Apparency of these conditions can also enable developers and businesses to have new markets and can create new and valuable differentiators. For instance, hardware developers may create open trackers that output open data to a health or storage device or service of their choice, where people themselves offer up their data for open studies, in which, like open software, access to the data used by any third party must remain open.

CONCLUSION

If we do not have apparency-to-transparency models of how our data is actually being used now and in potential futures, we cannot consent in a meaningful way to its use.

To consider whether or not we consent—assuming we can have a clearer sense of consent terms—we need a prime time in which to consider the terms of our consent policy, as we would our bank statements or insurance.

To have consent, we need greater apparency of how data is being used as a result of our consent. For the scale and speed of the IoT, this apparency/consent decision will need to be automated on our behalf; there is rich potential to create nuanced human-centered services with our colleagues in AI that can leverage the negotiation of consent/sharing terms.

Having strong, clear apparency to real semantic and pragmatic transparency as a backbone to meaningful consent will also help clarify risks within the data flows of large-scale,

heterogeneous IoT infrastructures, from homes to cities to national infrastructure.

Overall, by improving apparency to s/p transparency, we make meaningful consent possible. When meaningful consent becomes part of a system, entirely new kinds of services may be imagined that create value based on visible, shareable data. We can also make services more resilient. To get there, we need the design acumen of HCI researchers and UX practitioners to help design, deliver, and evaluate apparency interactions at IoT scale.

ACKNOWLEDGMENTS

The work informing this paper is supported by the Engineering and Physical Science Research Council UK projects Meaningful Consent in the Digital Economy (EP/K039989/1) and Cyber Security and the Internet of Things (EP/N02298X/1). We are grateful to Simone Barbosa, who in discussions of apparency helped tease out and name the two ways we use transparency as semantic/pragmatic.

ENDNOTES

1. McDonald, A.M. and Cranor, L.F. The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society*. 2008.
2. Moran, S., Luger, E., and Rodden, T. Literatin: Beyond awareness of readability in terms and conditions. *Proc. of UBICOMP 2014 Adjunct*, 641–646.
3. Trafton, J.G. and Monk, C.A. Task interruptions. *Reviews of Human Factors and Ergonomics* 3, 1 (2007), 111–126.
4. Shih, F., Liccardi, I., and Weitzner, D.J. Privacy tipping points in smartphones privacy preferences. *Proc of CHI 2015*; <http://dl.acm.org/citation.cfm?id=2702404>
5. Gomer, R., Mendes Rodrigues, E., Milic-Frayling, N., and schraefel, m.c. Network analysis of third party tracking: User exposure to tracking cookies through search. *Proc. of WI-IAT 2013*, 549–556.
6. Blasé, U., Leon, P.G., Cranor, L.F., Shay, R., and Wang, Y. Smart, useful, scary, creepy: Perceptions of online behavioral advertising. *Proc. of the Eighth Symposium on Usable Privacy and Security*. ACM, New York, 2012, Article 4. DOI: <http://dx.doi.org/10.1145/2335356.2335362>
7. Melicher, W., Sharif, M., Tan, J., Bauer, L., Christodorescu, M., and Leon, P.G. (Do not) track me sometimes: Users' contextual preferences for web tracking. *Proc. on Privacy Enhancing Technologies* 2016, 2, 135–154; <https://doi.org/10.1515/popets-2016-0009>

8. Chalmers, M. and Galani, A. Seamless interweaving: Heterogeneity in the design and theory of interactive systems. *Proc. of DIS 2004*, 347–356.
9. Hanmak, A., Soeller, G., et al. Measuring price discrimination and steering on e-commerce web sites. *Proc. of IMC 2014*; <http://dl.acm.org/citation.cfm?id=2663744>
10. Baarslag, T., Alan, A.T., Gomer, R., Alam, M., Perera, C., Gerding, E.H., and schraefel, m.c. An automated negotiation agent for permission management. *Proc. of Autonomous Agents and MultiAgent Systems*. 2017, 380–390.
11. McIntyre, J.J. Balancing expectations of online privacy: Why Internet Protocol (IP) addresses should be protected as personally identifiable information. *DePaul Law Review*. 2011.
12. Woolf, N. DDoS attack that disrupted internet was largest of its kind in history, experts say. *Guardian*. Oct. 26, 2016; <https://www.theguardian.com/technology/2016/oct/26/ddos-attack-dyn-mirai-botnet>

m.c. schraefel (<http://www.ecs.soton.ac.uk/~mc/>) is a professor of computer science and human performance at the University of Southampton, U.K., where she leads the WellthLab in Human-Systems Interaction. Her work focuses on human-systems design to enhance quality of life for all.
→ mc+ix@ecs.soton.ac.uk

Richard Gomer is a researcher in the Agents, Interaction and Complexity group at the University of Southampton. His main research interests lie in designing systems that support meaningful human control and agency, and reframing design praxis to treat thoughtfulness and even outright rejection as worthwhile design goals.
→ r.gomer@soton.ac.uk

Alper T. Alan is an HCI and AI researcher with an interest in human interaction with intelligent agents. He is a postdoctoral fellow in the Agents, Interaction and Complexity group at the University of Southampton.
→ a.t.alan@soton.ac.uk

Enrico H. Gerding is an associate professor in electronics and computer science at the University of Southampton. His research field is artificial intelligence and multi-agent systems with a particular focus on automated negotiation, auctions and game theory, and applications including privacy, the smart grid, and transportation systems.
→ eg@ecs.soton.ac.uk

Carsten Maple is professor of cyber systems engineering and director of cyber-security research at WMG, University of Warwick. He leads the GCHQ-EPSC Academic Centre of Excellence in Cyber Security Research and has particular interest in multidisciplinary approaches to privacy and trust in cyber-physical systems.
→ cm@warwick.ac.uk