Anticipating GDPR in Smart Homes Through Fictional Conversational Objects

"Alexa..

blue ring flashes

.. detect new devices"

Abstract

The incoming General Data Protection Regulation framework will profoundly affect the way which connected devices – the constituents of the so-called 'Internet of Things' – are designed and implemented. In this paper we introduce the metaphor 'IoT as constellation', a proposal grounded in Object Oriented Ontology, that aims to help designers create GDPR-compliant products which appreciate and respond to the data-mediated dependencies and interconnections of the Internet of Things. In particular we focus on voice-mediated interactions within IoT-enabled smart home contexts. To explore this space we apply the speculative design method Design Fiction to: (1) demonstrate how IoT as constellations may be applied in design practice; (2) promote Design Fiction as a viable prototyping mechanism for such conceptual frameworks; (3) forward proposals for how to approach GDPR-compliant, and Human Centered, design for the IoT.

Joseph Lindley; Paul Coulton; Haider Ali Akmal; Bran Knowles

{j.lindley; p.coulton; h.a.akmal; b.h.knowles1} @lancaster.ac.uk

Introduction

The relevance of Voice User Interfaces (VUIs) in Human-Computer Interaction (HCI) has become increasingly significant. This change has been contingent on computer systems which can improve their performance based on usage, an attribute that is enabled by machine learning algorithms, which themselves are fed by the large amounts of data generated by their widespread adoption. This widespread adoption has been driven by smartphone services such as Google Assistant, Apple's Siri, Microsoft's Cortana and more recently through cloud-supported hardware such as Amazon's Echo device. As these services become more widely used, the growth and maturation of voice-driven interfaces is in the midst of a positive feedback loop of ubiquity, usability and functionality.

With any new technology, the journey from technical feasibility to widespread adoption and ultimately domestication (cf. Silverstone 2006), is an unpredictable one. This unpredictability has always been a feature of new technologies, but in the post industrial age its effects cut so deep and fast that considering implications of technological adoption is now, more than ever, an important task (Joseph Lindley, Coulton & Sturdee 2017; Lanier 2013; Toffler 1990). As this process unfolds around the so-called Internet of Things – including voice interfaces - concerns around the privacy, trust, and security implications have begun to arise (e.g. Baldini et al. 2016; Carroll 2015). Some practical examples include the ingenious-but-surreptitious use of Wikipedia to advertise burgers (Kastrenakes 2017b); the potential use of accidentally gathered voice data in court cases (Sauer 2017); children's toys that are classed as surveillance devices (Oltermann 2017); and Amazon's services being detrimental of trust in personal relationships (Cecchinato & Harrison 2017).Whilst VUIs give rise to this concern in their own right, it is within the context of their relevance to the IoT that these implications seem most pertinent, primarily because voice interfaces and the hardware which house them are, oftentimes, also the way in which users access the broader milieu of IoT products and services in their home. The interconnectedness which defines the IoT massively complicates how concerns of privacy, trust and security manifest themselves in real situations. The ways we enter into agreements with the vendors of devices which produce, rely on, and collect data and services will, in Europe, change as the EU's 2016 General Data Protection Regulations (GDPR) become active. VUIs and the IoT are prime examples of ecosystems that are entirely dependent on data to function, and hence the forthcoming GDPR will set a new and unknown standards for governance in this space.

This paper is founded on the expectation that VUIs will continue to be widely adopted as they become the hub of interaction for the IoT, specifically in home environments. We proceed by unpacking the logics of this anticipated adoption, before introducing the conceptual frame of 'IoT as constellations' in order to shed light on how the ethoses of Human Centred Design (HCD) can inadvertently disempower users. With that framing in place we explore a novel perspective which philosophically challenges HCD, known as Object Oriented Ontology (OOO). OOO's thesis suggests that rather than placing humans at the centre of being, every 'thing' (including humans, animals and inanimate doodahs) is but an 'object' and these objects all coexist on a 'flat' (as opposed to hierarchical) ontology. Finally, we recount our embarkation on a speculative design process which responds to these conceptual constructs as well as the actual contents of the GDPR. These speculations employ the emerging research technique Design Fiction to probe possible futures and augment current understandings of Human-Centred Design for the IoT.

Voice User Interfaces and Information appliances

Actual figures are not published by Amazon, however according to research conducted by Parks Associates, in the United States, somewhere between 7 and 13 million Amazon Echo devices were sold in the 6 months to February 2017 (Anon 2017). Although the Echo is the market leader in consumer smart speakers, it is not alone, Google's Home product is a direct competitor (and is also available on the ~2 billion active Android devices) and Apple's forthcoming Homepod is aimed at the same market. Echo's success is not confined to Amazon's own hardware sales and is supported by integrations with other services (at the time of writing more than 10,000 'skills' are available, each offering some value-added integration with Echo) and incorporation with non-Amazon hardware. At the Consumer Electronics Show in 2017 a variety of companies announced integrations between their hardware and the Alexa Voice Service (the platform which underpins Amazon Echo) including TVs, refrigerators, robots, portable speakers, alarm clocks, cars, lamps, vacuum cleaners, plug sockets, washing machines, baby monitors, and many more (Kastrenakes 2017a).

The term 'Information Appliance' was coined by the famed usability engineer Don Norman (Norman 1998). This term describes physical devices which, rather than performing *many* functions in a potentially confusing manner, aim to do fewer things more aptly and appropriately. Three axioms described his vision for Information Appliances: simplicity, versatility, and pleasureability. The simplicity axiom suggests that the complexity of the appliance is that of the task, not the tool. Perhaps inspired by Arthur C. Clarke's '3rd law', this bootstraps the notion that the technology itself should be invisible. The versatility axiom says that appliance should allow and encourage novel and creative interactions; the technology should be improvisational. The pleasureability axiom posits that Information Appliances be fun, enjoyable and rewarding to use. Although the term 'appliance' carries connotations of hardware the most recognizable manifestation of Norman's concept is in the ubiquity of the specialized 'apps' which populate content stores for smartphone platforms.

In contrast to the previous generation of software, smartphone apps tend to focus on specific tasks, exemplify the axioms listed above, and supersede the generalized software paradigm of desktop computing whereby a software package has a vast array of features. Apps such as, for example, Instagram, inherit the complexity of the *task* (where the task is taking, sharing and commenting on photos) but are not complex as *tools*. Continuing with the Instagram example, although versatile (for example filters can be used in all kinds of improvisational and creative ways), this versatility does not impair a pleasurable user experience by making interactions overly complicated. The adoption of these software-based Information Appliances (smartphone apps) was contingent and reliant on the availability of a suitable means of access: the smartphone itself. Put differently, in order for IoT Information Appliances to make sense, there needs to be a universal way to access them. Smart speakers incorporating VUIs are increasingly playing this enabling role for the domestic IoT.

Making Sense of the Internet of Things by Looking at the Stars

Although this paper's primary concern is the GDPR and VUIs, it is the IoT's pervasiveness that brings GDPR and its relationship to voice interfaces into relief. The individual 'things' that make up the IoT are a network of heterogeneous interconnected objects that are readable, recognizable, locatable, addressable and/or controllable via the Internet (Coulton 2015). However, it is IoT devices' ability to interoperate with each other and with cloud-based services by generating, sharing and processing data, which underpins the depth of IoT's true value. Inspired by Walter Benjamin's writings we re-appropriate his conception of constellations. For Benjamin, 'ideas are to objects as constellations are to stars' (Benjamin 1999). Taking inspiration from this poetic sentiment, we posit that it is useful to consider any given IoT use case as a constellation. In the same way the individual stars that make up stellar constellations are simultaneously entities in their own right *as well as* part o*f* the constellation entity, IoT 'things' are simultaneously things in their own right, as well as part of an IoT constellation. While constituents of a constellation do exist individually, their meaning and significance is augmented by virtue of being part of a bigger system.

To concretise this metaphor and give it an IoT context, consider a smart heating system enabled by IoT service providers (e.g. Hive or Nest). Multiple separate entities must come together to make this system work: one or more pieces of smart home hardware, software to control that hardware, the actual central heating/cooling system that is being controlled, and oftentimes a smart speaker VUI such as Amazon Echo which would enable voice control. In our metaphor, each of these components is considered a 'star', while the system as a whole is the constellation. Appreciating that these individual entities work together enabled by their intermediary networks, is a key factor defining how individual 'things' become the *Internet of Things*. We hope is an intuitive and productive position to consider. However, the cosmological metaphor is applicable in additional ways too.

As we view them from the surface of the planet Earth, the constellations of stars have a different appearance depending on where the observer stands. The same constellation will appear the 'right' way up in the southern hemisphere, the 'wrong' way up in the northern half of our planet, and when viewed from the equator the same collection of stars appears on its side. It is also worth noting that different *cultures* observing the same constellations of stars interpret them very differently. For example, the collection of stars known widely as the Big Dipper in the United States is referred to as The Plough in the United Kingdom, and has a wide variety of other names too: The Saucepan, Bear, Stretcher, Wise Men and Drinking Gourd. Hence, in part due to constellations appearing differently on account of geometry and variance in observation-position, and in part due to other factors such as cultural subjectivity, the same constellation of entities is *interpreted* quite differently. Depending on who you are and your particular perspective, even though any given constellation is made up from the same individual things, how those things coalesce and develop meaning, as a collective constellation, is very fluid. Finally, we should note that, just as with the North Star (Polaris) when viewed from Earth, constituent parts of constellations are not *necessarily* visible from all points of observation. Because of the North Star's position, south of the equator it ceases to be visible. This is true in some IoT situations too, an entity which, from some perspectives is integral to the constellation, may be totally obscured from another perspective.

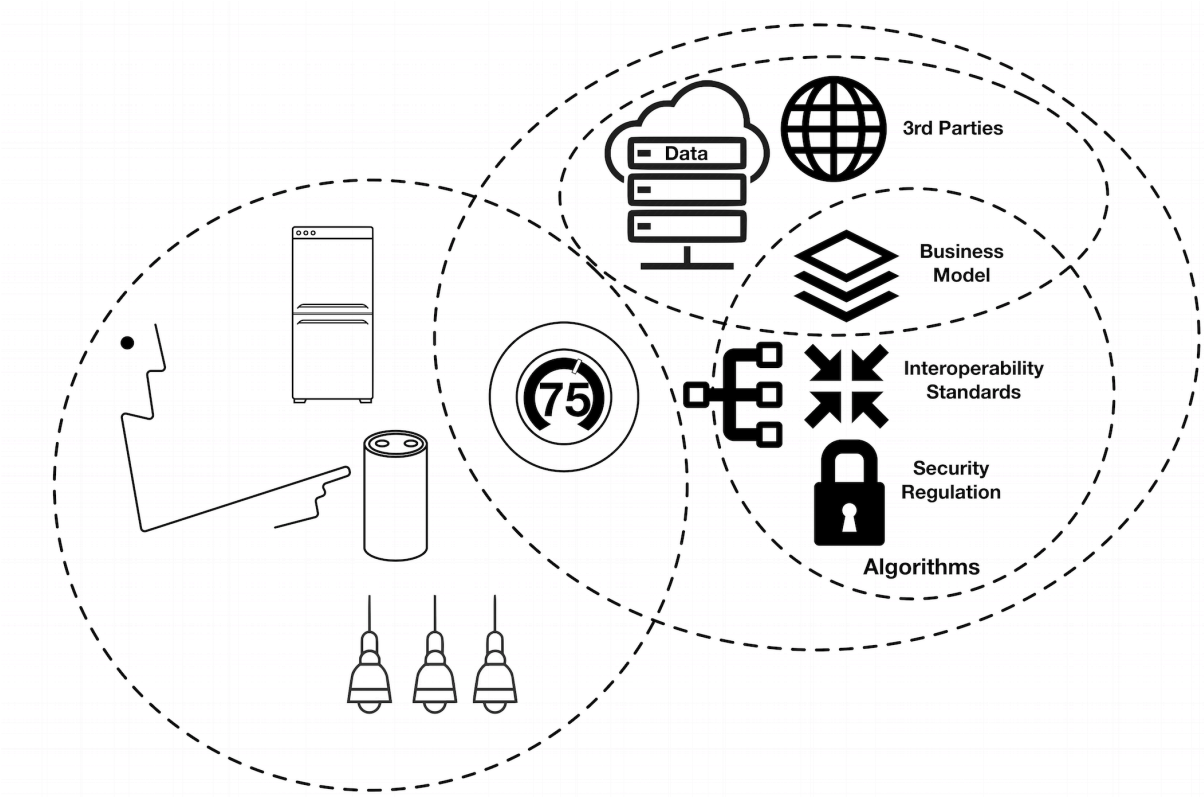

Figure 2. Visualisation of various overlapping IoT constellations.

We suggest that this metaphor can be used to describe most, if not all, IoT situations. Let's return to our exemplar situation of a domestic smart heating system to demonstrate the perspective element of the constellation metaphor. From the point of view of the homeowner who has installed the smart heating system, the elements mentioned previously (smart home hardware, control software, heating/cooling system, smart speaker) are the relevant constituent parts of the constellation. Together these individual points make up a constellation, which to the homeowner likely carries several possible meanings. For example, hypothetically, the constellation may represent a modern, technologically enabled house that can be remotely controlled and automatically learns about its users' habits. Such a house may be more convenient and more energy efficient than it would otherwise be. However, as this is a constellation, of course there are multiple possible perspectives to consider. For example, an Airbnb guest in an abode with a smart heating system would arguably see the system somewhat differently given that their host can control the temperature of their lodgings and monitor their usage entirely remotely. Last, although this part of the constellation is largely out of view, the perspective of IoT service providers is also of relevance. Continuing with the thermostat example, 'learning' thermostats collect data from the system's sensors, runs analytics on that data, then use algorithms to optimize heating and cooling in the user's home for maximum comfort. Although entirely out of sight, and with inner workings the details of which are rarely known to the user, this element of the constellation has a large and demonstrable influence on how the constellation is perceived to any other observer.

From Human Centred Design via Simplicity and Arriving at Obfuscation

The constellation metaphor for the IoT we feel is helpful in that it reflects the true heterogeneity of the things, entities, services, people and perspectives which make up the IoT networks themselves are our perceptions of them. Despite potentially bottomless depths of interwoven relationships between them and their data, it is *also* the case that the majority of IoT devices are designed and engineered to have a high degree of simplicity for the user. The methods by which designers focus on how users *actually* use devices, as opposed to how designers *say* they should use them; with the aim of arriving at devices which are easy and simple to use, may loosely be described as a Human Centred Design (HCD) approach (Giacomin 2014; Krippendorff 2004).

Not only is HCD is the modus operandi for the majority of technology designers including those working in IoT contexts, one might easily infer from its name that the approach would naturally build upon and produce Human-Centric Data. It is true that HCD has been highly influential in the design of technology for over 30 years, and its pervasiveness and longevity may be seen as a reflection of how HCD's methods have been successfully leveraged to help design a myriad of devices and services that are efficient, effortless, and edifying when in use and profitable for the corporations which design them (Giacomin 2014). As per the earlier discussion about Information Appliances (which, are themselves part of the HCD discourse) pivotal to being human centred is a technology's ability to fade into the background, and to ensure that any complexity remaining is essential to the task being achieved and *not* a by-product of the tool being used to achieve it. This makes sense if we assume that the human user of the device or service is *only* interested in the activity, that their motivation is *only* around completing the task with as little effort as possible, and that they are uncaring as to *how* it gets done (or, perhaps, what unseen side-effects there are of it being doing). Hence, simplifying that task by removing complexity which is not directly relevant to the completion of the task places the human's needs central to the factors motivating design choices.

To give an example, when a smartphone is used to place a call, the telephone software checks to see if the speaker is in use (by a music playing app, for instance). If it were the case that the speaker was busy then the telephone software automatically suspends whatever other application is using the speaker so the phone call can be placed. Thus, the complexity of the *task* is reduced, and the only remaining complexity is that of the *tool*, which, in this example, is a big green button in the dialler software saying, "Call Dad". The human, or user, simply has to worry about pressing the call button, if the user's only concern is placing the phone call, then, their needs have been placed at the centre of the design.

There are a litany of examples of HCD's simplicity axiom being applied and resulting excellent technologies which put the human at the centre effectively without any particular cause for concern. However, the IoT introduces, and makes commonplace, challenges to the *apparent* need for simplicity that is inherent to many human centred approaches. As Norman pointed out within a paper critiquing misinterpretations of HCD, the focus of the simplicity axiom should be on making the complex more understandable rather than masking it entirely (Norman 2005; Norman n.d.). We contend that in respect of the things which make up the IoT, simplicity-driven HCD has demonstrably resulted in end-user

interfaces and devices that, although making things simple for the user, obfuscate an underlying complexity. In many cases the details which this simplification process masks are actually crucial to understanding the *true* nature of the device, and hence when HCD is applied in this way the devices and services it is applied *to* become conflicted. They may have interfaces that are studied, validated, and refined (i.e. 'human centred') yet, as we will discuss below, those very same devices – by virtue of the IoT's particular properties – may also turn out to be, invasive, unreliable, or untrustworthy (i.e. *not* 'human centred').

Examples of this internal conflict manifests in many different forms. For instance, devices which perform one function for their user, but are dependent on their network connectivity often fulfil some other function for their manufacturer. The most prevalent example of this is smart televisions that monitor their users' habits and transfer data about those users to their manufacturers. Moving from privacy and trust to security, some shortcomings often found in IoT devices may be explained by the designers and engineers responsible for the technologies wanting to shield their users from the underlying intricacies of how devices work in order to make them easy to connect to a network, set up, and ultimately, use. Compromising security over usability has resulted in myriads of insecure IoT devices which, among other possible issues, are easily enslaved into botnets such as Mirai without their users becoming aware or having a chance to prevent enslavement. Building devices on top of cloud services allows engineers and designers to 'black box' (Latour 1999) huge amounts of complexity and processing power. For example, the voice recognition used by Google's Home and Amazon's Echo products relies on sophisticated machine learning algorithms and vast arrays of hardware. For the user however, the small flickering lights on top of their devices gives no indication that, after using a wake word, they are momentarily connected to computer systems (as well as sometimes humans!) many thousands of miles away. Beyond delivering a rewarding and usable interface, system architectures such as this invite manufacturers to leverage the data that is gathered in ways that are superfluous to the core functionality of the device. Also, this cloud-dependent architecture results in products that, if the cloud service is unavailable, becomes almost entirely useless.

Made famous in part because of Germany's Federal Network Agency classifying it as a surveillance device, the IoT doll My Friend Cayla exemplifies these three issues combined together in a single device. Both the doll's hardware, and its supporting smartphone app, have significant security issues. Most notably the doll, which relies on a Bluetooth connection, has no means of setting a password. Although this simplifies the process of pairing the doll to its smartphone app, it also makes it technically very easy to 'hijack' either the voice recordings that the doll makes and the words that the doll speaks. The doll is also highly dependent on a cloud service, if the cloud service is not available or ceases to be compatible with the latest version of smartphone operating systems (as is currently the case with the latest version of Android) then most of the doll's functionality ceases. Finally, when using the cloud service recordings are sent to the doll's manufacturer, who then shares the data with a third party, who have reported they may, in turn, use the data for various other purposes (what these other purposes are is not specified).

Although a full exploration of HCD is beyond the scope of this particular paper we do make several assertions about the methods and ideology which together make up HCD. HCD is demonstrably successful at helping develop rewarding and effective interfaces for computer

systems. Part of the process that helps arrive at this outcome is a pursuit of simplicity in interactions with those systems. The methods which develop this simplicity were, by and large, developed before the advent of the web, and certainly before the widespread adoption of the IoT. Because of the unavoidable complexity of the IoT's networks, these same methods, which almost always have worthy intentions, act to obfuscate sometimes critical elements of the constellation. Because HCD's methods arrive at this outcome without directly impeding the visible elements of the interactive system, those same methods can be described as simultaneously helping develop a 'Human Centred' system that *also* reduce some humans' ability to have agency within that particular constellation, and therefore, are counterproductive in terms of HCD ideals.

Object Oriented Ontology and Design Fiction

Our reference to Object Orientated Ontology (OOO) is made in response to the various constructs we have previously introduced. To recapitulate the position we've rhetorically built so far, the adoption of the IoT and VUIs, anticipated growth of their influence, and the changing regulatory environment is the *context* this paper is concerned with. Our constellation metaphor for the IoT is useful in two ways. First it tries to communicate the multiplicity of the IoT, conveying how 'the' IoT is in fact not a singular 'thing' at all, but a gamut of things positioned amongst many intertwined spectra of interests, incentives, and influences. Second, the constellation metaphor provides a conceptual bridge between the IoT and the subsequent section of the paper which critiques HCD by challenging the notion of what a centre is through noting constellations - things made up from multiple singles - do not have a universal centre. The critique argues that HCD infused designs often have a well-meaning attachment to the worthy notion of there being a 'human centre', the worthiness *does* work effectively in terms of making interfaces accessible. However, HCD's methods are inadvertently duplicitous in IoT contexts, a situation underpinned by the tendency of HCD methods, when applied to IoT constellations, to lead to obfuscation. HCD, in the context of IoT constellations, can *apparently* serve a user's needs, while properties of the same design, in one or more unseen other domains, create devices or services that are likely to be acting on behalf of some other external entity's interest. In the following we discuss OOO. A thorough explication of OOO in this paper is beyond its scope, yet our intention is to include enough of a description such that we articulate why and how this ontology is useful as an adhesive for the constructs discussed thus far and a platform for responding to them.

OOO is a contemporary strand of philosophy, which, as with other speculative realist perspectives, rejects 'correlationism'. Correlationism takes the view that things are only real insofar as they are sensible to a human subject in terms of humans' correlation between what it is to think and what it is to be. Thus, by *rejecting* correlationism agency is theoretically assignable to non-human actants. Graham Harman, one of the proponents of OOO, extends the Heideggerian position that tools and other objects make sense in terms of their utility or purpose to humans. Harman suggests that 'things' need not be defined by *human* interactions alone, but their definition is in fact much more elusive, and should be construed on each object's own terms or in terms of interactions betwixt objects. However, for the majority of objects these interactions are not intimate encounters: Harman points out that a rock's existence is defined in the most part by its *own* reality and not so much by its interaction with other rocks, other *non*-rocks, or even by the shoes of humans walking on top of it. For Harman the interactions which define any given object's ontological truth are

ultimately particular, they make sense only for specific objects at a specific moment in time; "objects only unlock each other's realities to a certain extent" (Harman 2002, p.2). The resultant ontology is one that appreciates that objects are not defined by human perceptions and that each object's reality is mostly imperceptible to other objects. The conclusion of this thesis is that no single object, class of object, or collection of objects should have inherently more or less agency than another. Hence, OOO is a flat ontology, devoid of any inherent or implicit hierarchy. Given our argument that constellations are essentially 'centreless', this ontologically level (or 'centreless') landscape is a useful theoretical plane to consider and to bolster our metaphor.

We should note that OOO is as widely criticised as it is lauded. Charlesworth points out that by aligning with this philosophy any given human may say they are no more or less significant than any other object, and hence may absolve themselves from responsibility to act (Charlesworth 2015). The rhetorical conclusion of this is to ask 'if the OOO is a zero-sum game, why is it useful'? The ongoing debates around the virtues of speculative realism, along with sceptical counterpoints such as that above, have some considerable substance. However, it is not the purpose of this paper to provide a justification for OOO. Rather our task is to consider how the OOO thesis might be used as a way of re-evaluating, and extending, approaches for designing technology in such a way that it can become *more than* human centred. Cast in the shadow of HCD's tendency to inadvertently 'simply obfuscate' the full appearance of IoT constellations, we believe that OOO is a powerful ontological 'jumping off point' from which to consider how we might take account for the constructs discussed thus far, operationalize our critique, and propose realistic and optimistic strategies for designing better regulatory regimes, legal frameworks, and technologies.

In order to achieve this these bold aims, and in order to provide a segue from hifalutin rhetoric into the tangibility of practice, we draw inspiration from the work of video game designer Ian Bogost. In his book Alien Phenomenology, Bogost argues a practical engagement with OOO, which is otherwise a rather otherworldly and mainly cognitive endeavour, can be made possible by using video game design to build artificial worlds (2012). The approach is founded on the notion that despite philosophical discussions of metaphysics being undoubtedly interesting, we should be distrustful of practitioners (of metaphysics) who cannot develop an empirical base for their conclusions (put differently, esoteric philosophy that is *purely* rhetorical is something Bogost treats with an element of incredulity).

"If a physician is someone who *practices* medicine, perhaps a metaphysician ought be someone who *practices* ontology. Just as one would likely not trust a doctor who had only read and written journal articles about medicine to explain the particular curiosities of one's body, so one ought not trust a metaphysician who had only read and written books about the nature of the universe." (Bogost 2012, p.96)

Being a practicing metaphysician is not easy. Bogost wishes to 'play God', and, by becoming demiurgic allow the intangibility of metaphysics to be made tangible. If one can have a direct experimentation with ontology, a *material* engagement with philosophy, then, Bogost argues, we might better understand the true nature of the ontological rhetoric in the first place. He proposes to achieve this by crafting artificial worlds using video games, worlds

which have their own attributes, properties, quirks and idiosyncrasies – limited only by the *God's* (i.e. designer's) imagination – a interested scholar can practice ontology, and, in Bogost's eyes make ontologically-inspired conclusions that are more trustworthy and inspiring than otherwise.

Our intention in this paper is to slightly abstract Bogost's argument. The abstraction goes thus; the attribute that gifts game designers this totemic ability is their power to construct worlds from scratch. These worlds, because their rulesets only exist in the virtual electronic domain, are unconstrained by any of our usual conceptions of reality; in a video game mathematics, the law, material properties, or anything else you can to think of can be redefined as the designer sees fit. But these 'world building' abilities are not constrained to game designers exclusively. Writers, filmmakers, physicists, and philosophers are but some other professions that of occasion are required to craft imaginary worlds in pursuit of their aims. Thus we look to the maturing speculative design and research technique Design Fiction as an appropriate method for exploring the issues discussed so far in this paper. Design Fiction is a collection of methods and approaches that pivots around world building (Coulton et al. 2017), and, thus, is a ripe way of enacting Bogost's thinking and, for our purposes, to *practice* the metaphysics of OOO (Joseph Lindley, Coulton & Cooper 2017).

Speculative design, a relative of critical design, is a future focused design endeavour that focuses design practice on question asking, rather than answering. That is, these approaches are not attempting to create a products for sale, or that necessarily solve a problem, but rather to elicit a deeper understanding of a particular issue or selection of issues (Dunne & Raby 2013; Dunne 2006; Auger 2013). Within this speculative design family of movements, and an ever-growing body of Design Fiction practice, this relatively young field is very much 'pre-paradigmatic'. Hence there are concurrent yet incongruent perspectives on what Design Fiction is, what it aims to achieve, and how it does that. Among the developing field, we align with a particular thesis known as 'Design Fiction as World Building' (Coulton et al. 2017). This approach uses a variety of examples of Design Fiction practice to demonstrate that the means by which Design Fiction derives value is by constructing one-or-more artefacts that, when viewed together, describe the coordinates, or 'entry points', into a fictional world (e.g. Joseph Lindley, Coulton & Cooper 2017; Lindley & Coulton 2015; J. Lindley et al. 2017). Each entry point (which, we note is an artefact of one sort or another) tends to depict one part of that world. Usually these depictions work at a particular scale either representing a large area of the world, but without a huge amount of detail ('zoomed out'), or a very detailed depiction, but of a smaller area ('zoomed in').

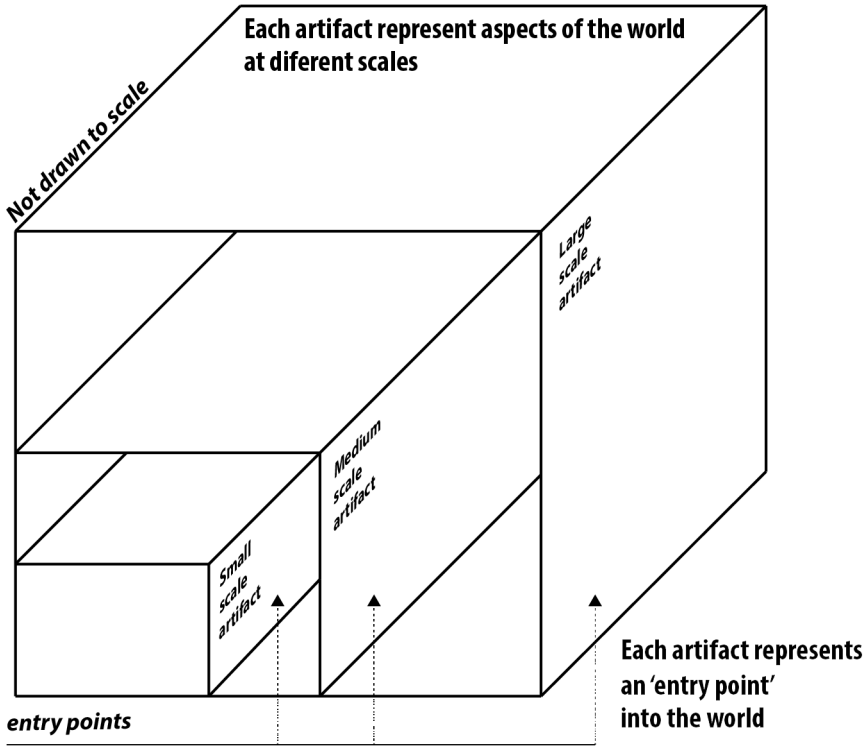

Figure 2. Diagrammatic representation of Design Fiction as World Building (cf. Coulton et al. 2017). The artefacts which make up a Design Fiction (represented here by the nested squares) provide 'entry points' to the fictional world. Each one of these entry points gives clues about that world at a different 'scale'.

Design fiction is a research method (Lindley 2015a). There is not, however, one single process by which Design Fiction can be used to do research or one single *type* of research that can be done with Design Fiction. For example, in one configuration the actual *process* of creating a Design Fiction can produce the kind of contingent knowledge associated with 'Research through Design' (cf. Gaver 2012; Lindley & Coulton 2016; Lindley 2015b). Design fiction can also be used as part of the contextual search or background research task associated with supporting a different creative endeavour or design process (Lindley 2015a). Finally, Design Fiction may also be used as a component in a methodological assemblage (cf. Law 2004) such as a stimulus or intervention in focus groups (Duggan & Lindley 2015) or as part of a 'speculative enactment' process (cf. Elsden et al. 2017). Although there is scope for the particular Design Fictions we present here to be expanded and appropriated for means beyond those contained in this paper, what we provide is a reflexive account of the creative and design process, hence this is 'Research through Design Fiction' (Lindley 2015b). We have derived insights in several ways: first, the meaning embodied in the designed artefacts themselves; second, by discussing the practicalities of the plausible future world ('Design Fiction world') in which these designs would make sense; and finally by describing reflexively (cf. Schön 1992) the logics and rationality behind the design decisions we took throughout the process.

Speculating about Privacy: Crafting the Design Fiction

The discussions in this paper so far are representative of the design space and theoretical framework that the speculation works within. Within the context of a fast-moving post-industrial world - a world that is pervaded by VUIs controlling IoT devices (themselves part

of constellations) and that suffer from HCD's simplicity/obfuscation dichotomy – our task was to infuse the spirit of the incoming GDPR legislation, combined with the notions of IoT as constellations/OOO, into a plausible Design Fiction. As with many design processes resisting the urge to try and answer the brief's questions too early (Cross 2011) was a significant task. Instead we committed to an evolving and iterative process of experimentation and exploration. For us, a significant part of this involved understanding how users enter into agreements with the vendors of the technologies historically and exploring how GDPR may impact upon this. As we discuss below, it seems quite obvious that the historic way of doing things is not at all acceptable in terms of the *spirit* behind the GDPR. In parallel it was necessary to explore the technical and practical aspects of designing and creating voice powered interfaces.

The GDPR addresses a wide variety of issues, much of which pertains to individuals' rights, and how organisations collecting and processing their data should respect them (also significant are the penalties for noncompliance mandated by GDPR, although we do not discuss this here). This includes any individuals right to be aware of what data is held about them, the right to access that data, the right to rectify incorrect data, erase any data and/or restrict processing of data. Additional rights include data portability (i.e. individual should be able to take data held in one place and reuse it elsewhere) and to object (i.e. to refuse consent for profiling and decision making based on one's own data). Superficial reading of many existing terms of use and privacy agreements quickly shows *why* GDPR is necessary; the ambiguity and vacuity of the 'legalese' style language employed appears designed to indemnify corporations legally at the expense of any given individual's ability to exert influence over their data (perhaps unsurprising as the regulatory environment stems from a time when data was not so ubiquitous or powerful). Beyond the impenetrable language present in user agreements, there are other common issues around consent. For example asking users to tick a box to indicate they have read and agree to the legal agreement does not, in reality, have *any* correlation to whether they have read the agreement, and even less so to whether they have *understood* the agreement. In practical terms ticking such boxes tends to infer "I wish to use the device" and has very little to do with the content of the agreement.

We use the information collected about and from you for a variety of business purposes, including for example, to:

- respond to your questions and requests;
- provide you with access to certain functions and features of the Services;
- verify your identity and seek your consent;
- communicate with you about your account and activities using the Services;
- communicate changes to any of our policies or Services;
- improve our Services;
- to provide you with the most user-friendly navigation experience
- for internal business purposes (including calculating statistics); process applications and transactions;
- to meet our legal and regulatory obligations and protect our legitimate interests;
- to carry out research and analysis, training and quality assurance;
- if you agree, to contact you about other services and products that we think may be of interest to you; and
- for any other purposes which we clearly explain to you at the time you provide your personal information or to which you otherwise consent.

Figure 3. Extract from the Privacy Policy relating to My Friend Cayla[1] demonstrating the linguistically ambiguous style of language we found to be common among existing privacy policies.

Another facet of these issues, in contemporary IoT devices, manifests around presentation; the means by which the legal agreement is showed to a user is oftentimes not fit for

[1] https://www.myfriendcayla.com/privacy-policy

purpose if the purpose was for the agreement to be read and understood. A practical example of this is the June (internet connected) Oven. Within seconds of being turned on the oven's control panel presents a new user with a legal agreement on a relatively tiny oven control panel; this is not the ideal medium to read such an agreement!

Figure 4. The consent interface on a June Oven, customers *must* agree to this before they are able to use the oven at all.

This interaction (which is representative of domestic IoT devices) also begs the question 'why must I give *all* the consent necessary to use *all* features of this device, in order to use *any* of its features?' Put differently, why is consent a one-off thing? In healthcare for example consent is given many times throughout a treatment, the level to which consent is attained relates direct to the potential significance of the activity. Consenting for major surgery is a lengthy and time-consuming process, but given the high risks of surgery, this makes complete sense: patients must understand the risk and doctors must be able to prove that consent has been meaningfully attained. However, when a healthcare professional needs to carry out a more basic activity, providing oral painkilling medication, taking some blood for, or even rearranging a patient's pillows, verbal consent is usually sufficient (but must still be obtained at each juncture in treatment). This proportionality is entirely absent in 'one size fits all and one time only' user agreements. A translation of the healthcare model into the realm of the IoT would likely resonate with several GDPR

requirements such as the right to be informed, the right of rectification, and the right to restrict processing. The right to be informed, in particular, stipulates that data controllers must specify in a concise, transparent, accessible, and intelligible manner what personal data is being processed. We hypothesised that conversation via VUIs may be a viable and proportionate way of aligning with GDPR requirements, particularly given that VUI-based conversational agents will likely become the conduit for IoT Information Appliances.

During the design process we ran two workshops with computer science and design students in order to help us develop a wide spectrum of understanding around (1) how users may interact via VUIs and (2) how the requirements of GDPR may be met. Using paper prototyping tools (see figures 5 and 6) we explored the detail of implementing voice interactions for consent. During these workshops perhaps the most worrying response was the observation from computer science students that there were a series of technical workarounds which were relatively easy to implement. For example, rather than designing an interaction which would facilitate a privacy agreement or consent process which is compliant with the spirit of GDPR, our workshop participants realised it may be easier to ensure that encryption or suitable levels of pseudonymity exempt data from being classified as personal. We include this observation as a point of interest, and as a discussion point vis-à-vis the adoption and domestication (cf. Silverstone 2006; Joseph Lindley, Coulton & Sturdee 2017) of GDPR. However, how *real* organisations and their information governance professionals respond to GDPR will be an evolutionary process that develops as consensus emerges around how to interpret the wording of the regulations. Whilst it seems plausible that systems developers may find and utilise 'loopholes' in the GDPR which allow collection and manipulation of data that is against the *spirit* of the regulations but within the law, beyond acknowledging the possibility we did not make this a central feature of our Design Fiction speculations.

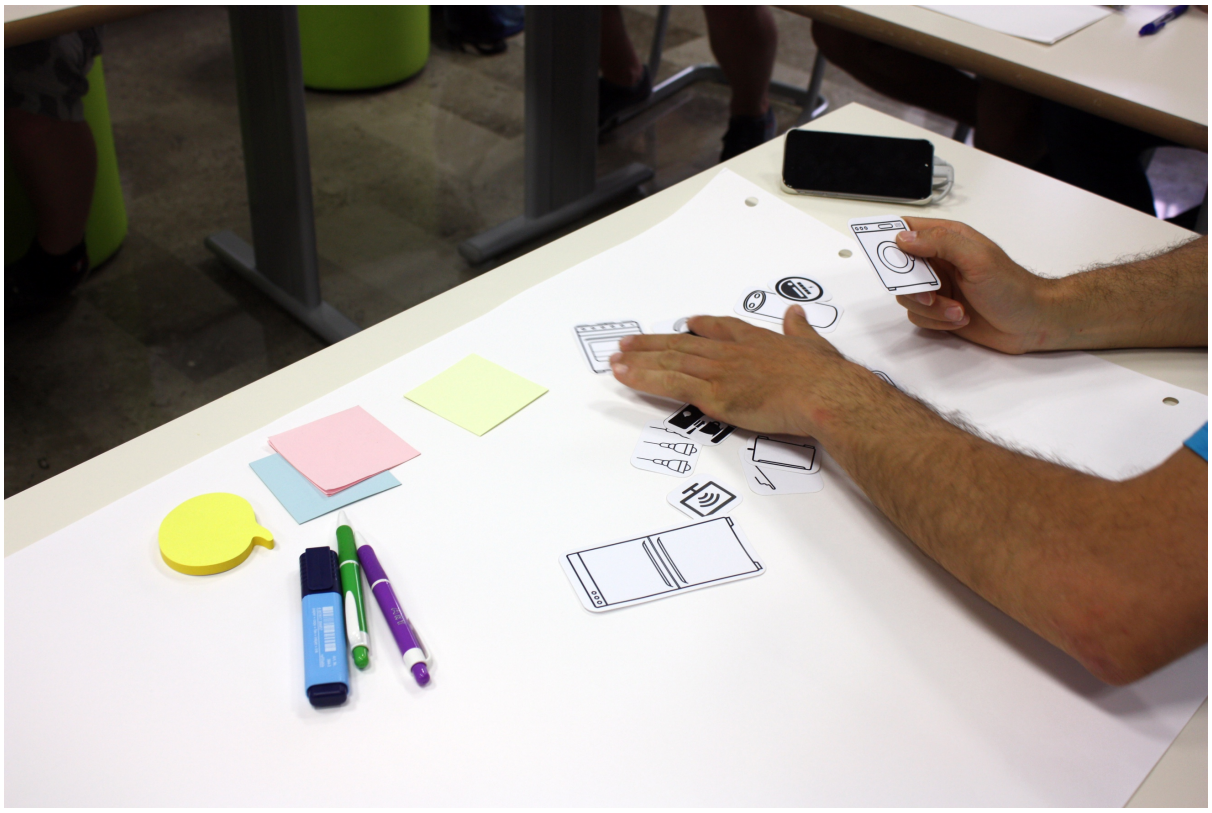

Figure 5. Workshop materials for prototyping GDPR compliant voice interactions with house hold appliances.

Along with direct experimentation with a variety of voice agents (Google Home, Amazon Echo, Apple Siri, Microsoft Cortana) the workshops helped us to understand the complexity of crafting voice interactions, and the limitations of the medium. In our prototyping sessions it became apparent that long passages of text and lists (such as those commonly found in privacy agreements), because they are harder to comprehend and recall than their written equivalents, are not viably conveyed through voice. Whilst short voice interactions are rewarding and efficient, non-binary decisions and complex information is extremely difficult to convey via a VUI alone.

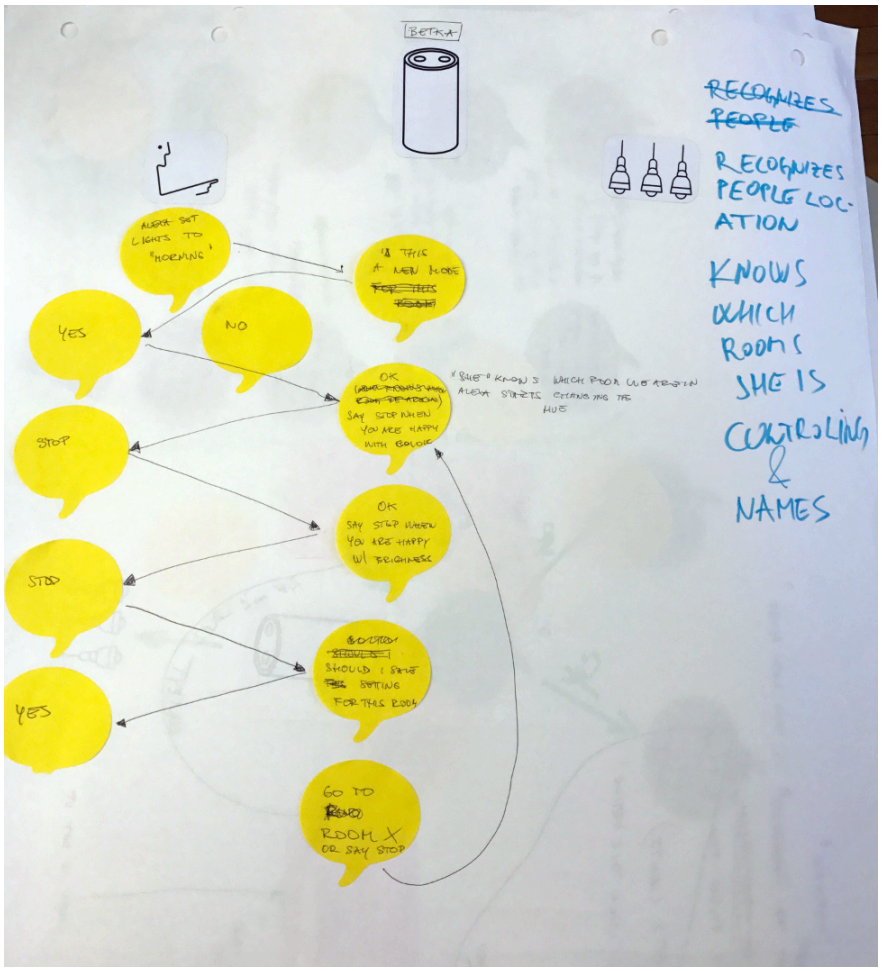

Figure 6. Prototype conversation flow for a voice interaction based around a binary decision tree.

In fact, although VUIs are becoming widely adopted extremely quickly, a range of research discusses the problematic aspects of designing interfaces for them. Reeves uses conversation analysis in order to critique the design of VUIs identifying various lines of further enquiry, but also highlighting that, fundamentally, voice interactions with VUIs are not conversational in the same way a human to human voice interaction would be (Reeves 2017). This finding is supported by analyses of chat bots; these computer interfaces (or 'Botplications') use the same machine learning algorithms as VUIs to allow them to appear conversational but rely on text-based input. Chat bot interfaces tend towards specialisation,

but are usually are somewhat less intuitive to use than VUIs (Klopfenstein et al. 2017, p.564). User identification and security is another problematic area for VUIs. Notably Google's Home product has implemented voice-based user authentication, although Amazon's market leading Echo product has not (Barrett 2017). It is not currently clear how being able to recognise a particular user's voice could protect against voice impersonators or recordings of authorised users.

A previous Design Fiction project, which provided the foundations for the ideas presented in this paper, built a world around a fictional internet connected kettle named Polly (Joseph Lindley, Coulton & Cooper 2017; Lindley & Coulton 2017). Also adopting the world building approach to Design Fiction the project comprised a wide range of artefacts including press release, marketing and promotion materials, details of a successful crowdfunding campaign, and details of functionality. Much of these resources were created in order to boostrap the world building machinery, but the fictional functionalities depicted included several OOO-inspired interface designs. One of these was the kettles ability to, using its primary interface (which is a VUI) converse with its owner in order to arrive at a bespoke privacy agreement, customised for each user. In practice what this means is that when the kettle is being configured it will describe to the user what data it gathers, why, and what features (of the kettle) that data supports. For the kettle, because it is a specialised Information Appliance with very specific functionality, these options are relatively few and hence the interaction is viable via voice. Once the user has agreed (or refused) to the various points of consent a machine-readable version of this information is shared with the users router. Within the fictional world of Polly the kettle this machine readable privacy policy is part of a protocol called Minimum Necessary Datagram Protocol (MNDP), and MNDP is a requirement for certification from a fictional IoT regulator (named OfIoT).

Figure 7. Example conversation to negotiate a customised privacy policy (cf. Joseph Lindley, Coulton & Cooper 2017).

Although Polly's voice interface-enabled privacy policy system appears viable, and aims to work within the spirit of GDPR, on reflection informed by our workshops and other emerging research (e.g. Reeves 2017) we noted that the speech above (figure 7) is unhelpfully long. Such long passages of speech remove the illusion of conversation and reduce the otherwise intuitive quality of a VUI. Another critical reflection on Polly is around our Information Appliance hypothesis; it seems extremely unlikely that *every* IoT device would have its *own* VUI. Rather, a single point of access will likely become established. From this starting position, and with the constructs discussed previously, we began using Design Fiction to prototype an Information Appliance with a voice-based consent process.
Our aim then, was to use Design Fiction to prototype a GDPR compliant Information Appliance's consent procedure. We intended the primary means of interaction to be through a standalone VUI (which we anticipate will become the enabler of IoT Information Appliances). In contrast to established paradigms for providing consent we wished to create an atemporal mechanism which can accommodate different levels of consent at different times. Based upon *avoiding* the issue of a VUI delivering an elongated soliloquy our aim was to design a conversational interaction. Finally we conscious not to 'cheat' the system by creating something that would be technically GDPR-compliant but would shirk the *spirit* of the legislation.

We elected to base this Design Fiction around the premise of an IoT door lock. Several such devices exist on the market today which offer a range of functionality including geofencing (automatic locking/unlocking dependent on your location), guest access (giving somebody access via their smartphone), and voice activation (Delaney 2017). Our initial designs revolved around a conversational approach to interacting with the lock via a VUI, however we quickly realised that purely verbal interactions could not meaningfully convey the complexity of the necessary privacy-based concepts. This problem, in hindsight, most likely stems from the cognitive bandwidth required. Human senses are not all equal. Whilst our senses of taste, hearing and smell, convey relatively 'low bandwidth' information, our sense of touch and sight can convey much more higher bandwidth information (Coulton 2017). Hence, conveying the amount of information traditionally contained in a privacy policy (frequently several thousand words worth of information) via a VUI would make the voice-based interaction with an Information Appliance both useless *and* impractical.

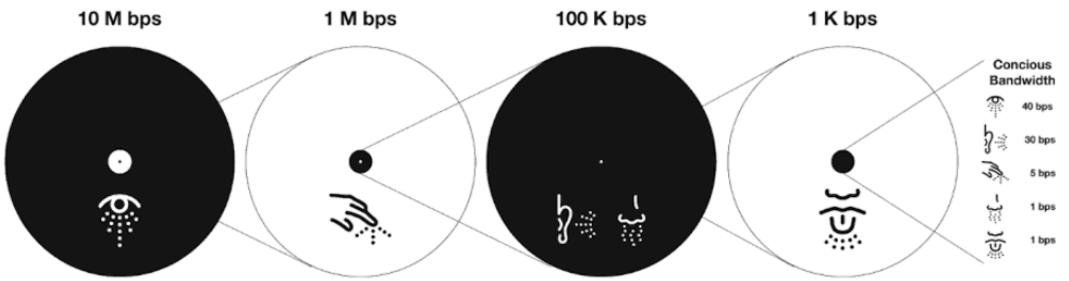

Figure 8. Bandwidth of the senses.

Having realised that voice was *not* a viable way to communicate the detail of information pertaining to GDPR compliance, privacy, and data governance relating to the device – partly because our ears cannot deliver enough bandwidth - we began to consider how to visually convey the relevant information. This too, posed a challenge, however. The most immediate issue related to the potential complexity of what we wanted to convey. Although our example device, a lock, is relatively simple, with only a small selection of possible features and associated data collection/processing, there is a geometric increase in complexity depending on the number of actants in the constellation.

If we consider our smart lock and possible integration with two other services: Amazon Echo and IFTTT (a company that allows over 500 electronic services to be connected together using simple 'If *this* then *that*' logic). The lock company's system will likely collect data locally, which sits somewhere on the user's home network between the lock, the router, and perhaps a smartphone. In order to support services such as remote unlocking, guest access and geofencing data is likely shared with the lock company's internet servers too. It would not be uncommon for some element of usage data to be shared *outside* of the lock company (perhaps pseudonymised, but perhaps not) with at least one extra company for some kind of analytics or external service provision. That extra company will arguably include a clause in their own privacy statement which says they may share data with *another* company. Now, if we consider that Amazon Echo and IFTTT *must* move data pertaining to usage between the user's local network and their own cloud-based services, then even in this *tiny* ecosystem of a single device with basic functionality, by the time we consider what *could* have happened to users' data there are hundreds of *possible* flows for a user's data. Though these are not all probable, but they *are* possible. This complexity curtailed our early attempts to design a visual 'map' of where data may be stored and/or processed; there were simply too many potential connections to be able to make any sense of the information. It also helped us to realise that representing the *uncertainty* of what *might* happen to data was a key requirement of this design.

In order to help cut through the confusion of this entanglement we introduced new constraints to our creative process. Rather than trying to represent all aspects of privacy relevant to GDPR – a very broad category – we needed to make the information we wished to present simpler. Hence, we elected to make our visual 'map' of data flows only represent data which could *identify* the user. Although vastly reducing our design space, identifiability itself is still not a straightforward category, in particular because *partial* identifiability can quickly become full identifiability when two disparate datasets are combined. We also quickly noted a qualitative difference between data held locally (on the users' network, on devices under their control), data held by a known provider (for example, the lock company, or Amazon) and data that is 'elsewhere' (for example an analytics or sales company), in terms of potential users' trust.

Designing the Interaction

Our revised speculation revolved around using a VUI to trigger device *detection*, before deferring to a screen-enabled device to help configure privacy settings with a visual aid at the relevant stage of the interaction. We elected to showcase this, as part of the Design Fiction, in the form of a companion smartphone app. The first challenge with this app was to

devise a method to articulate identifiability, at local, known provider, and other levels, which would also convey the necessary element of uncertainty/probability. The series of circles below shows our early prototype for how this may be achieved.

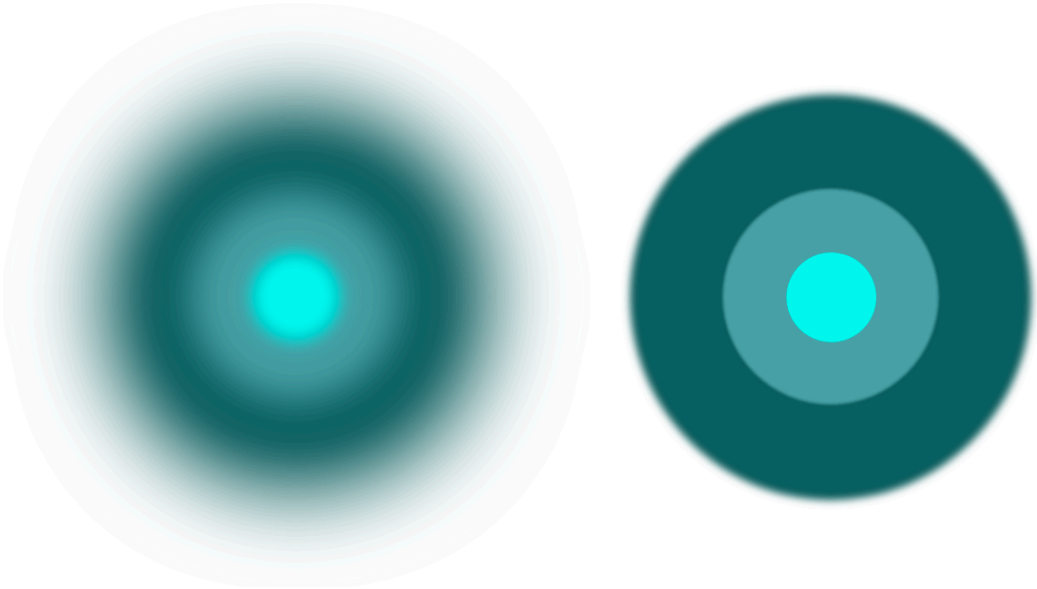

Figure 9. Early prototype of privacy app for visually representing probability of identifiability for data held by various stakeholders.

The inner most circle on each one of these diagrams represents data held on the user's *own* network and hardware. The middle circle represents data held by providers *known* to the user (i.e. the manufacturer of devices they have attached to their network). Meanwhile the outer circle represents data held by *anyone else* – in practice these would be companies with a right to access the user's data according to a privacy agreement (data may or may not be pseudonymised) but the identity of the company would not *necessary* be known to the user. At each of these three levels the edge of the representative circle can be soft or hard, clear or blurry, focused or fuzzy. This variance represents the *probability of identifiability.* Hence, if the data held on my local network can *definitely* identify me, then the inner most circle would have a clearly defined, hard edge (e.g. the right most diagram above). If, on the other hand, data held by an unknown third party is very *unlikely* to identify me, then the outer-most circle would be extremely fuzzy and blurred (e.g. the left most diagram above).

The next challenge was to imagine how these diagrams may relate to service provision and functionality for our IoT lock. To investigate this we considered the necessary data flows to support four possible features for a smart lock:

- Using a smartphone or other NFC-enabled device as a key
- Using geofencing to automatically lock the door when the user leaves the house
- Using Amazon Echo to voice control door-locking
- Using IFTTT to integrate with any other IFTTT compatible service (e.g. Tweet "I'm home" every time the door is unlocked after a period of being locked)

Each of these levels of service corresponds to a different degree of required data-sharing and potential identifiability. We hypothesised that:

- To unlock the door using NFC there is no *requirement* for any information that would *definitely* identify the user to be stored at any point of the system. It could plausibly be configured without even user registration.
- To use geofencing it would be necessary to employ the lock company's servers as an intermediary between a user's smartphone and the lock itself. In this case it is probably that the lock company would require registration, and that this information would be held on their servers.
- In order to use Amazon Echo we *know* that users must be registered with Amazon, and that Amazon holds a significant amount of data about those users. If voice-activating the lock then Amazon would also – probably – have access to the meta data relating to locking/unlocking.
- IFTTT is a bridging service, connecting the services of over 500 companies and organisations. The majority of these services hold identifiability information, although whether they need to (or do) share it with one another is not clear in most circumstances.

With these possible scenarios considered and their associated data flows researched we considered how they would manifest in the form of a smartphone app. The prototype below depicts a simple slider interface which allows the user to select settings on a spectrum of most private to most functional. As the function is increased the four icons along the top which represent the lock's features become coloured/grey according to whether they are enabled or not. As the features are turned on and off the three spheres change blurriness to represent how likely it is the user is identifiable at each of the three levels for the given level of functionality. For any given IoT device this screen would require bespoke research and configuration to fully understand what features are reliant on the sharing of what data.

Figure 10. Prototype app for selecting a balance between privacy and functionality including icons depicting the specific functionality available and associated circles representing probable identifiability.

To further prototype the necessary interactions and to further construct the Design Fiction world that this prototype exists within we developed a short film that depicts the process in context along with voice interaction. In this film version we refined the design and began to develop potential branding for the system. We envisaged this service to be a modular add on that device developers may buy in to (as opposed to a proprietary system that each IoT development company would have to devise; although either model is plausible). Hence the app would be a standard interface users would be accustomed to using (akin to, perhaps, the 'Login with Facebook' buttons that are frequently seen in smartphone apps today). Within this facet of the world we named the developers of the privacy app 'Orbit' in part in reference to the similarity between the app's three concentric circles and the 'orbits' in atomic nuclei. The addition of an eye and ghost motif is intended to semiotically invoke notions of privacy. Our video prototype follows thus:

Figure 11. “Alexa, discover new devices”. Alexa scans for new smart home devices.

Figure 12. The lock, shown here, flashes to show it has been detected by the Amazon Echo.

Figure 13. (Alexa says) “Discovery complete. Smart lock found. Check your Orbit privacy app”

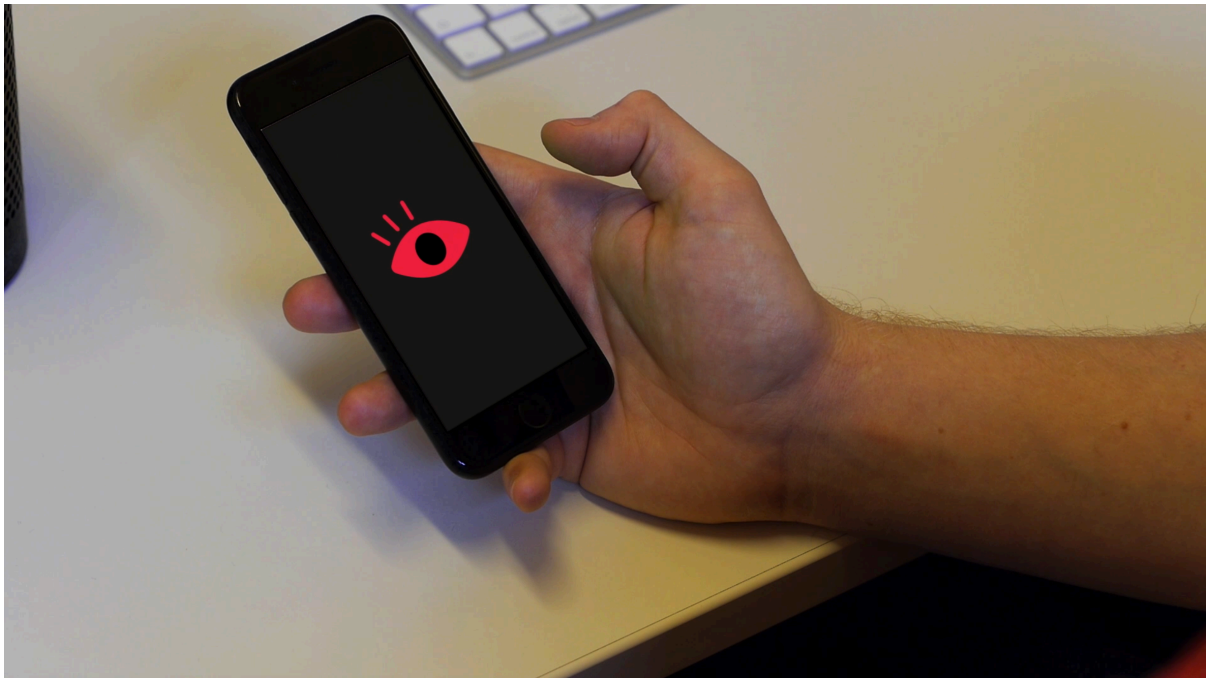

Figure 14. The app automatically loads once the user’s phone is unlocked.

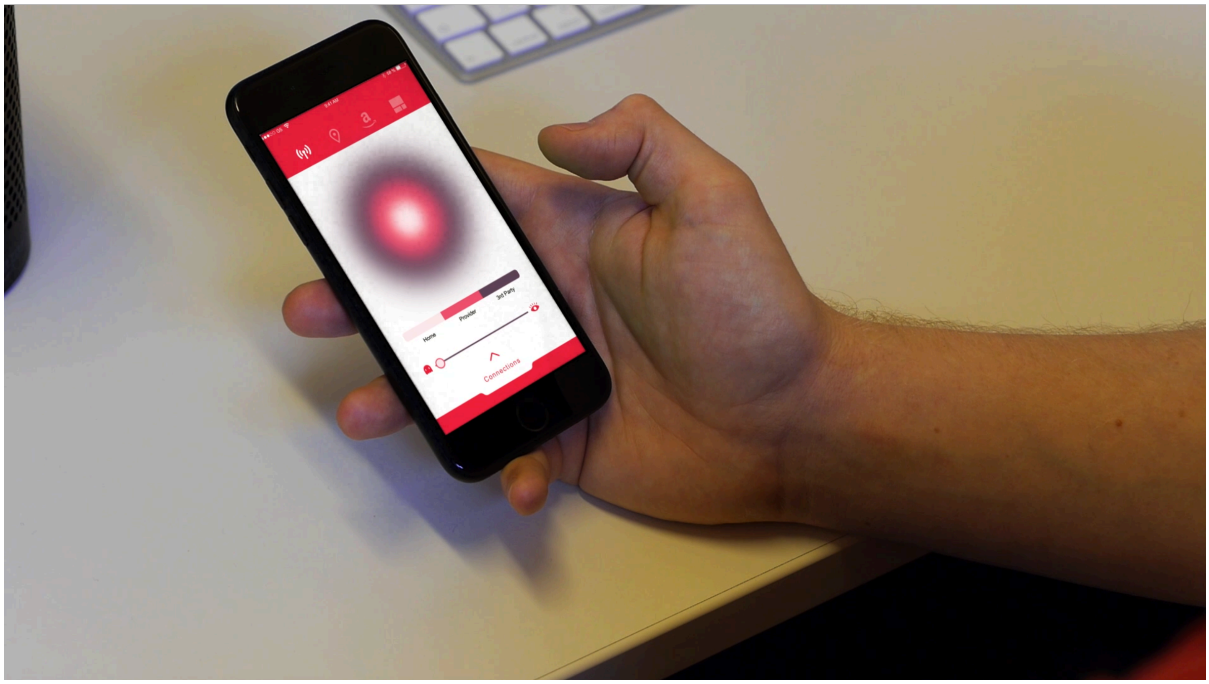

Figure 15. Using the slider the user can select suitable privacy settings before returning to further voice-based interaction via the Amazon Echo.

Reflecting on the Speculation and Extending the World

Whilst part and parcel of the reason for creating a Design Fiction is to prototype designs and interactions as demonstrated above, of equal importance is to consider the world within which such designs and prototypes would make sense. The prototypes which help to invoke a Design Fiction's artificial world are as important as the world itself (Coulton et al. 2017). Although this Design Fiction begins this process, it is preliminary work and crafting a larger more coherent Design Fiction would result in significantly more concrete findings.

Nonetheless, a key learning from producing the prototype thus far is the fiendish difficulty of meaningfully programming the 'orbits' (circles) according to data flows associated with IoT devices. Whether some of the specificity inferred in the wording of GDPR will reduce the complexity inherent in understanding where data flows, or not, will be a big factor dictating how the regulations will ultimately impact upon interface design.

We hypothesised expanding the Design Fiction world with the addition of a fictional job advertisement for a data governance professional whose role it would be to do the research necessary to populate the Orbit app's 'orbits' for specific applications. Through the creation of such a Design Fiction artefact it would be possible to prototype a specific interpretation of the GDPR and begin to infer tangible insights about the fictional world that Orbit makes sense within.

We also contemplated what would motivate device developers to use such a system (which presumably would have a cost associated with it). It seems plausible to imagine that the underlying motivation for developers to provide costly configuration interface such as that depicted in the Orbit app would come via some form of regulation. Just as food companies re-specify their recipes depending on how regulators insist they demonstrate nutrition

information, perhaps IoT developers may be regulated in a similar way, perhaps resulting in IoT “nutrition” labelling.

Figure 16. Example IoT ‘nutrition’ labelling.

When considering the concept of IoT labelling a pertinent question, which derives from our own decision to simply represent identifiability in the Orbit app, is; what properties would such labels represent, and how would the figures be derived? This is a particularly intractable problem because, unlike with nutrition, there are rarely discrete values for the relevant factors. Also with pertinence is the question, could the extent, size, and heterogeneity of a constellation be somehow represented in discrete and quantifiable terms?

Through the design process we considered ways to augment the Orbit app. Three specific concepts were discussed. First, we considered the idea of adding dot-like-graphics which literally ‘orbit’ the circles. Each dot would represent a specific piece of data and by tapping it a user could view information about that data point (e.g. what it contains, what it is for, who has it, and services dependent on it). Second we discussed the possibility of leveraging colour and shape in the diagrams. Perhaps colour, shape, or some other semiotic separation, could allow multiple privacy-relevant concepts to be displayed on a single diagram (e.g. layering pseudonymity atop identifiability). Last we considered further utilising the size of the circles (or other shapes) displayed on the app as a means to represent the aggregate volume of data at any given level.

In summary, our Design Fiction built around the Orbit concept reflects the difficulty of building a meaningful, GDPR-compliant, consent system into the VUI-powered smart home of the future. Significant compromises were necessary in order to arrive at a usable interface which met our criteria; specifically the decision to *only* represent identifiability and also to include a significant element of probability (as opposed to certainty) within the visual feedback. Another significant issue, a deeper understanding of which is likely dependent on *actual* interpretations of GDPR over the coming years, is to do with the research necessary in order to arrive at reliable values for each possible Orbit configuration. In other words, we do not know whether it will be possible to quantify factors such as ‘identifiability’, at varying levels, with any degree of reliability. Even so, with these limitations accepted, we do feel that this design *is* a viable way of empowering users to proactively and dynamically play a role in deciding how they wish to participate in the flows of data amongst constituent parts

of any given IoT constellation. At the very least it may provide an indication of how to make concrete the otherwise esoteric IoT as constellation/OOO discussion. Further development of this particular Design Fiction world, we suggest, will also be a useful tool to pre-emptively conceive and prototype ways to implement and interpret GDPR.

Concluding Remarks

In this paper we presented elements of theory, method, and practice, relevant to the Internet of Things (IoT), Voice User Interfaces (VUIs), Human Centred Design, and the incoming General Data Protection Regulations (GDPR). A term that is open to interpretation and that covers an incredibly diverse array of concepts, we interrogated the concept of IoT, and built an accessible and meaningful metaphorical frame to reflect its heterogeneity; IoT as constellations. We bootstrap our metaphorical framing with contemporary philosophical thoughts around non-anthropocentrism and Object Oriented Ontology (OOO) and put this in to practice using the burgeoning speculative design technique of Design Fiction. Our contributions are threefold. First, we hope to promote and advocate for our constellation metaphor as a robust framework to help analyse and conceive of specific IoT situations in terms of their GDPR-relevant aspects. Second, we propose and demonstrate Design Fiction as a viable means to prototype how regulatory frameworks such as GDPR may manifest themselves at the interface-side of technology. Third, we present our Design Fiction around the Orbit privacy app, not as a posited 'solution' to the issues discussed, but as a means to pose more meaningful and powerful questions, *as well as* an indicative design proposal suggesting how to begin approaching GDPR-compliant, voice-based, user interface design.

Acknowledgements

This work was carried out at Imagination Lancaster as part of the PETRAS Cyber Security for the IoT Research Hub. It was funded by EPSRC grant with reference EP/N02334X/1. Many thanks to Dan Burnett for assistance developing Orbit prototypes, Leon Cruickshank for detailed discussed about the constellation metaphor, and all our other colleagues at Imagination Lancaster and PETRAS.

References

Anon, 2017. Parks Associates: Voice Control is a Key Differentiator for Smart Home. Available at: http://www.parksassociates.com/events/connections-summit/media/cs-2017-pr3 [Accessed February 27, 2017].

Auger, J., 2013. Speculative design: crafting the speculation. *Digital Creativity*, 24(1), pp.11–35.

Baldini, G. et al., 2016. Ethical Design in the Internet of Things. *Science and Engineering Ethics*, pp.1–21.

Barrett, B., 2017. Google Home Now Supports Multiple Users For Smarter Personalization | WIRED. *Wired*. Available at: https://www.wired.com/2017/04/google-home-multiple-accounts/ [Accessed August 22, 2017].

Benjamin, W., 1999. *The Arcades Project (trans. Eiland, H and McLaughlin, K)*, Cambridge, MA: Belknap.

Bogost, I., 2012. *Alien phenomenology, or, what it's like to be a thing*, U of Minnesota Press.

Carroll, R., 2015. Goodbye privacy, hello Alexa: here's to Amazon echo, the home robot who hears it all. *The Guardian*. Available at:

https://www.theguardian.com/technology/2015/nov/21/amazon-echo-alexa-home-robot-privacy-cloud [Accessed August 3, 2017].
Cecchinato, M. & Harrison, D., 2017. Degrees of Agency in Owners & Users of Home IoT devices. In *Making Home: Asserting Agency in the Age of IoT CHI 2017 Workshop*.
Charlesworth, J.., 2015. The End of Human Experience. *Art Review*. Available at: https://artreview.com/opinion/summer_2015_opinion_jj_charlesworth/ [Accessed January 30, 2017].
Coulton, P. et al., 2017. Design Fiction as World Building. In *Proceedings of the 3nd Biennial Research Through Design Conference*. Edinburgh, UK.
Coulton, P., 2015. *Playful and gameful design for the Internet of Things. In More Playful User Interfaces*, Springer Singapore.
Coulton, P., 2017. Sensing Atoms and Bits. In I. Heywood, ed. *Sensory Arts and Design*. Bloomsbury, pp. 189–203.
Cross, N., 2011. *Design Thinking: Understanding How Designers Think And Work*, Bloomsbury.
Delaney, J., 2017. The Best Smart Locks of 2017. *pcmag.com*. Available at: http://uk.pcmag.com/surveillance-cameras/77460/guide/the-best-smart-locks-of-2017.
Duggan, J. & Lindley, J., 2015. *Sans Duty - Making Tax Visisble*,
Dunne, A., 2006. *Hertzian Tales: Electronic Products, Aesthetic Experience, and Critical Design*, The MIT Press.
Dunne, A. & Raby, F., 2013. *Speculative Everything*, London: The MIT Press.
Elsden, C. et al., 2017. On Speculative Enactments. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp.5386–5399.
Gaver, W., 2012. What should we expect from research through design? *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12*, p.937.
Giacomin, J., 2014. What is human centred design? *Design Journal*, 17(4), pp.606–623.
Harman, G., 2002. *Tool-being: Heidegger and the metaphysics of objects*, Open Court Publishing.
Kastrenakes, J., 2017a. Amazon's Alexa is everywhere at CES 2017 - The Verge. *The Verge*. Available at: http://www.theverge.com/ces/2017/1/4/14169550/amazon-alexa-so-many-things-at-ces-2017 [Accessed February 27, 2017].
Kastrenakes, J., 2017b. Burger King's new ad forces Google Home to advertise the Whopper - The Verge. *The Verge*. Available at: https://www.theverge.com/2017/4/12/15259400/burger-king-google-home-ad-wikipedia [Accessed August 3, 2017].
Klopfenstein, L.C., Malatini, S. & Bogliolo, A., 2017. The Rise of Bots : A Survey of Conversational Interfaces , Patterns , and Paradigms. , pp.555–565.
Krippendorff, K., 2004. Intrinsic Motivation and Human-Centered Design. *Theoretical Issues in Ergonomic Science*, 5(1), pp.43–72.
Lanier, J., 2013. *Who owns the future*, Simon and Schuster.
Latour, B., 1999. *Pandora's Hope: Essays on the Reality of Science Studies*, Harvard University Press.
Law, J., 2004. *After Method: Mess in Social Science Research*, London & New York: Routledge.
Lindley, J., 2015a. A pragmatics framework for design fiction. In *Proceedings of the European Academy of Design Conference*.

Lindley, J., 2015b. Researching Design Fiction With Design Fiction. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition - C&C '15*. New York, New York, USA: ACM Press, pp. 325–326.

Lindley, J. & Coulton, P., 2015. Game of Drones. In *Proceedings of the second ACM SIGCHI annual symposium on Computer-human interaction in play*.

Lindley, J. & Coulton, P., 2017. On the Internet No Everybody Knows You ' re a Whatchamacallit (or a Thing). *Making Home: Asserting Agency in the Age of IoT Workshop*.

Lindley, J. & Coulton, P., 2016. Pushing the Limits of Design Fiction. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*. New York, New York, USA: ACM Press, pp. 4032–4043.

Lindley, J., Coulton, P. & Cooper, R., 2017. Not on demand: Internet of things enabled energy temporality. In *DIS 2017 Companion - Proceedings of the 2017 ACM Conference on Designing Interactive Systems*.

Lindley, J., Coulton, P. & Cooper, R., 2017. Why the Internet of Things Needs Object Orientated Ontology. In *Proceedings of EAD12 - Design for Next*.

Lindley, J., Coulton, P. & Sturdee, M., 2017. Implications for Adoption. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*.

Norman, D., HCD harmful? A Clarification - jnd.org. Available at: http://www.jnd.org/dn.mss/hcd_harmful_a_clari.html.

Norman, D., 2005. Human-centered design considered harmful. *interactions*, 12(4), p.14.

Norman, D.A., 1998. *The invisible computer: why good products can fail, the personal computer is so complex, and information appliances are the solution*, The MIT Press.

Oltermann, P., 2017. German parents told to destroy doll that can spy on children. *The Guardian*. Available at: https://www.theguardian.com/world/2017/feb/17/german-parents-told-to-destroy-my-friend-cayla-doll-spy-on-children [Accessed August 3, 2017].

Reeves, S., 2017. Some conversational challenges of talking with machines. In *Talking with Conversational Agents in Collaborative Action Workshop*.

Sauer, G., 2017. A Murder Case Tests Alexa's Devotion to Your Privacy | WIRED. *Wired*. Available at: https://www.wired.com/2017/02/murder-case-tests-alexas-devotion-privacy/ [Accessed August 3, 2017].

Schön, D.A., 1992. Designing as reflective converstion with the materials of a deisgn situation. *Knowledge-Based Systems*, 5(1), pp.3–14.

Silverstone, R., 2006. Domesticating domestication. Reflecting on the life of a concept. In T. Berker et al., eds. *Domestication Of Media And Technology*. Open University Press, pp. 229–247.

Toffler, A., 1990. *Future Shock*, Random House.