A Room With an Overview: Towards Meaningful Transparency for the Consumer Internet of Things

Chris Norval and Jatinder Singh

Abstract-As our physical environments become ever-more connected, instrumented and automated, it can be increasingly difficult for users to understand what is happening within them and why. This warrants attention; with the pervasive and physical nature of the IoT comes risks of data misuse, privacy, surveillance, and even physical harm. Such concerns come amid increasing calls for more transparency surrounding technologies (in general), as a means for supporting scrutiny and accountability. This paper explores the practical dimensions to transparency mechanisms within the consumer IoT. That is, we consider how smart homes might be made more meaningfully transparent, so as to support users in gaining greater understanding, oversight, and control. Through a series of three user-centric studies, we (i) survey prospective smart home users to gain a general understanding of what meaningful transparency within smart homes might entail; (ii) identify categories of user-derived requirements and design elements (design features for supporting smart home transparency) that have been created through two co-design workshops; and (iii) validate these through an evaluation with an altogether new set of participants. In all, these categories of requirements and interface design elements provide foundation for understanding how meaningful transparency might be achieved within smart homes, and introduces several wider considerations for doing so.

Index Terms—Internet of Things (IoT), transparency, accountability, user experience, design, technology impacts, smart homes

I. INTRODUCTION

Our environments are becoming increasingly connected, instrumented, and automated, amid an ever-growing myriad of network-enabled consumer Internet of Things (IoT) devices claiming to offer convenience, comfort, safety, and control [1], [2], [3], [4]. From lighting, heating, and appliances to home security and beyond, surveys indicate that the average US household has eleven connected devices, with 28% of households having at least one home automation device [5]. In this way, the IoT is already impacting the lives of many.

Most consumer IoT devices typically have domestic and lifestyle aims, enabling interactions between the user and their surroundings. For example, a smart thermostat might work to automatically manage the temperature of a home, turning the heating on when it is cold, or opening the windows if warm. However, the inner workings (i.e. the operation) of these systems can quickly become complex and opaque; a given deployment might entail numerous interactions (and data flows) between a range of components, where, for example, even fairly simple automated events (such as smart windows opening in response to weather reports) can be driven by a supply chain of data sources and dependencies [6]. In short, what data is collected, how it is used, and where it goes are often little-known by users [7], [8], and are typically obscured within and as part of the broader IoT ecosystem [9]. This hinders users in understanding and overseeing what is going on within their IoT deployment (which they may do so for reasons such as verifying that the deployment is operating as expected, in response to particular incidents or events, or for general curiosity). As a result, it may not be clear why the system operates in the way that it does, potentially leading to unexpected or unintended behaviours.

This opacity is problematic, not least given that the home a space in which may consumer IoT devices target—represents a private space (where such devices often feature sensors that can readily capture personal, sensitive and intimate information [10]). Furthermore, given the physical nature of such devices, one can readily envisage scenarios in which serious consequences might arise; smart ovens might automatically turn on during the night [11], or smart locks could prove obstructive if functioning unexpectedly, such as during a house fire [12]. As the IoT continues to pervade various aspects of our everyday lives, it is crucial that users are best placed to oversee, understand, and act upon any such issues, if and when they arise.

A. The role of transparency in the consumer IoT

It follows that **transparency** in the context of the consumer IoT is important. Greater visibility over the consumer IoT can help users in a variety of ways, including satisfying curiosity and concerns about how things are behaving, ensuring that their data is appropriately handled, monitoring system behaviour, helping to ensure correct operation and functionality. revealing what lead to a particular fault or incident, and so on [6]. That is, if users are able to better understand what is happening (or has happened) within their connected environment, then they are better positioned to take action and respond, as and when required. Such actions might involve, for example, re-configuring parts of the system to prevent undesirable data sharing, perhaps ceasing to use a particularly problematic IoT device altogether, seeking to contact a company to exercise data rights, begin litigation, and so on (see Fig. 1). Indeed, this need for oversight is increasingly recognised, as we see growing demands for transparency over technology, and emerging regulations requiring such (§VII-C). Yet, the practical dimensions of transparency, as a means for supporting accountability, is an under-considered area.

Importantly, simply providing information about a system is not necessarily useful ([13], [14], [15], [16], [17], [18]); pro-

The authors are part of the Compliant & Accountable Systems Group at the University of Cambridge, William Gates Building, Cambridge, CB3 0FD, United Kingdom e-mail: (cjn41@cam.ac.uk; jatinder.singh@cl.cam.ac.uk).

Meaningful Transparency				
Supports Helping to facilitate				
 Accountability Comprehension Oversight Review Scrutiny Contestation Autonomy 	 Reconfiguration of the system/devices Restricting unwanted data sharing The identification of misbehaving devices Exercising data subject rights Legal compliance The challenge of organisations Litigation and redress 			

Fig. 1. Meaningful transparency supports stakeholders in understanding, taking action, contestation and enacting change.

viding *too much* information, can even act to further obscure and obfuscate (the so-called 'transparency paradox' [16]). Rather, transparency must work to support users in understanding what's happening, so as to enable appropriate actions in response [14]. In other words, through ensuring that users are able to understand and oversee the happenings of their smart environment (what data is being collected, how it is being processed, transmitted locally and/or shared remotely, why particular actions and outcomes are occurring, etc.), they can then take appropriate and informed decisions on what to do in reponse.

It follows that there is a need for **meaningful transparency** [17]: the provision of information in a manner that caters to the needs and expectations of the user to facilitate their effective oversight, scrutiny, and review over these technologies [19]. In practice, such information should be *contextually appropriate* [19] to the user and their situation, and support wider aims of contestation [20], accountability [6], autonomy [21], and legal compliance [22] (see Fig. 1).

In an IoT context, while there has been considerable focus on specific issues (such as those around security, privacy, complexity and supporting users with specific deployments), *tackling opacity* so to support scrutiny and broader accountability aims has thus far been under-considered. Tackling such opacity in the consumer IoT is important, particularly as the prevalence of these devices increases [5], and as the consumer, social, and regulatory demands for greater transparency and accountability regarding technology continues to grow [23].

B. Exploring meaningful transparency in the consumer IoT

This paper explores the needs, expectations, and desires of participants as regards transparency mechanisms within the consumer IoT. Given that we consider transparency for general consumers, and that many consumer IoT products aim at domestic use, our focus is on the *smart home*. We consider the design of transparency mechanisms that are meaningful for their intended recipients (the users), exploring aspects such as what information (from, and about, their smart homes) they would want to know, how they believe this should be represented and communicated, and what such transparency mechanisms might look like in practice. In doing so, our goal is to provide practical ways forward by articulating types of *user-derived requirements* (specifications that formalise the needs of users [24]) and design approaches for those looking to build effective transparency mechanisms into the IoT.

To do this, we undertake three user studies, working with participants to explore the kinds of transparency mechanisms they would want and expect from smart home systems. This is to reveal what they consider a meaningfully transparent smart home might entail, and provide tangible ways forward for researchers and designers alike. Study 1 takes a broader approach to develop an understanding of what users want to know (across a selection of scenarios where something within the smart home warrants attention). We do this with a survey of 126 respondents, probing into what types of information they would care about, how they would expect that information to be communicated to them, and what follow-on actions they believe that this information would enable them to take. To further explore this in practice, we conduct two co-design workshops (with 5–6 participants each) in Study 2. During these sessions, the two groups of participants work together to create a list of user requirements and design prototype interfaces for what they would want and expect from a transparent smart home system. The result of these workshops is i) a general set of design elements (i.e., key design features for enabling smart home transparency), and ii) a selection of categories representing the types of user requirements that our participants felt important for bringing about transparency. Study 3 then looks to generally test, validate and scrutinise these design elements and user requirement categories with a (new) set of 56 participants, indicating the generalisability of these findings.

C. Contributions

Exploring how transparency mechanisms can better support transparency, through enabling scrutiny and oversight, is an emerging area of general importance [9], [25]. We consider this within the context of the consumer IoT. Towards this, we offer practical ways forward for supporting meaningful transparency in smart homes, by providing:

i) categories of user-derived requirements to assist and guide developers in implementing transparency mechanisms within their offerings;

ii) a collection of co-designed transparency 'design elements' – features of design for supporting smart home transparency; and

iii) insights on some of the broader considerations and challenges facing the design of transparency mechanisms for bringing about a more accountable consumer IoT.

These contributions aim to better support users in gaining more effective oversight over their connected environments, and developers and researchers in facilitating such. Moreover, we argue the importance of transparency mechanisms that cater to the needs and expectations of their users, and discuss various practical ways forward (for the IoT, and, indeed, beyond) toward enabling more meaningfully transparent—and thereby supporting more accountable—technologies.

II. BACKGROUND

The *Internet of Things* (IoT) typically refers to "the extension of the Internet [...] into the physical realm, by means of the widespread deployment of spatially distributed devices

with embedded identification, sensing and/or actuation capabilities" [26]. Importantly, however, there is much more to the IoT than just devices; a given IoT deployment will often comprise a socio-technical ecosystem of devices, software, systems, and organisations, in which the flow of data between, through, and across systems and organisations work to deliver overarching functionality [9], [27]. This may involve, for example, remote weather sensors (by means of web APIs), local sensors (proximity, temperature, etc.), and cloud-based web services, all working together to determine whether to actuate (turn on) a smart home's heating system.

IoT deployments raise the prospect of a range of potential issues occurring [27]. These might relate to system interoperability [12], a lack of manageability [28], [29], emergent behaviours [30], a lack of knowledge as to what data is being collected and how it is being used [7], [8], outright system failure [31], [32], [33], and so on. However, the IoT entails opacity [34]; it can be challenging for even the most technical users to effectively monitor, oversee, and/or diagnose how their smart environments are operating, be it to understand a particular issue or occurrence, or to satisfy curiosity. This can raise real challenges for the users of smart homes, with considerations relating to system reliability being said to play a crucial factor in smart home acceptance [35], [36].

As such, there appear clear benefits for mechanisms that can support users with greater oversight and understanding over these complex and opaque systems; naturally, transparency can assist in this regard [15], [17]. Transparency mechanisms can help illuminate the inner workings of a system - they can support users in verifying that their system is operating in line with expectations, identify particular areas of interest, and conduct targeted investigations into specific components or behaviours that arise. Moreover, effective transparency mechanisms also support individuals in making informed choices [10], allowing users to take further action, if and when required [9]. This might include the provision of information that results in users changing device settings to restrict undesired behaviours, removing problematic devices from use in the home, or challenging particular parties or contesting their actions [20], [37] (such as through complaints, legal means, etc.). That is, effective transparency mechanisms will often be a precursor to pursuing accountability aims [6], [19], playing a key role in supporting users in identifying, mitigating, and/or rectifying a range of different issues and concerns that may arise with regard to the consumer IoT.

A. Specific user concerns

In recent years, there have been numerous research studies exploring consumer attitudes to the IoT (see [38], [39] for two systematic literature reviews). Such research can provide valuable insights into the challenges and concerns facing prospective IoT users, across a range of different contexts and scenarios, and highlight areas where meaningful transparency may be able to assist. As such, we now briefly discuss three particularly prominent concerns from across the literature.

1) Understanding smart homes: Given that smart homes may be driven by a range of different systems (perhaps

including various sensors, actuators, and online services, incorporating a range of organisational ecosystems) [27], there is real potential for IoT deployments to quickly become complex. Even simple scenarios can lead to a disconnect between users' understanding of what is happening within their smart homes and what is actually occurring; in a user study with 20 participants, Yarosh and Zave [29] outlined a hypothetical smart lock (with four features of automated behaviour), and asked participants whether the door would be locked or unlocked within 20 given scenarios. They found that participants' mental models did not appear to align with the operation of the system, despite the participants claiming to understand how these system interactions should work [29]. This can be problematic, not least given the propensity for tech related issues within the home to be consequential (sensitive data may be captured and leaked, smart doors or windows may not secure when a person leaves the house, etc.).

Users misunderstanding what is occurring within their smart home has led to work which attempts to explore and mitigate this issue from a conceptual perspective. For example, Despouys et al. [40] have proposed a model for 'sensemaking' within the context of smart homes, in order to identify and manage potential scenarios where the expectations of users and autonomous systems are misaligned [40]. Similarly, Chuang et al. [41] analysed 'concept videos' of IoT products to develop a design vocabulary for human-IoT systems [41], as a means for supporting developers in better communicating and explaining the intended purpose of their IoT systems, such that their users might have a stronger understanding of how they operate. Some have explored the potential of methods for 'explainable AI' within smart home contexts, given that consumer IoT systems will often employ some form of machine learning [42], [43]. Such works that support the comprehension of smart home systems offer one way forward for how the developers and designers of IoT systems might better manage and communicate the behaviours of their systems to their users, helping to ensure that such systems act in line with expectations.

Relatedly, other bodies of work have explored practical techniques or prototypes aimed at supporting users in understanding how their system is, or has been, operating. Some of these have been more conceptual, such as work by Desjardins et al. [44], [45], who have used literary authors to translate IoT device data logs into fictional novels ('data epics') for their users [44], [45]). Others have explored various techniques and modalities for presenting information about how smart devices are operating to their users, e.g. through providing descriptions of apps and devices [10], creating data visualisations for sensor feeds [46], using 'nutrition labels' [47], [48], designing interfaces for rule editors [49], [50], [51], using voice assistants to query smart home logs and provenance data [34], and recommendations to help users with issues of consent [52]. The proposal of general mechanisms to help better inform users (as to the status and operation of their IoT devices) acts to highlight this topic as an ongoing area of concern.

It is clear that any potential disconnect between how users understand and expect their smart environments to operate, and how those environments may actually function can have negative implications. However, while some of the above work offers empirical insights into how we might better support users in understanding how their smart environments operate, much of the research into this topic has thus far involved participants only at the evaluation stage of developing a particular proposed solution or approach. Less considered are the needs and expectations from the perspective of users themselves – *what they want to know, how they believe that this should be communicated to them*, etc. Engaging users, however, as part of a design-process is important for helping ensure that its outcomes are helpful and effective [53], [54]. Therefore there are clear opportunities for research that is more formative in nature, focusing on how best we can support potential users—taking into account their own perspectives—in understanding how their IoT systems operate.

2) Privacy: Another specific issue for the IoT relates to user privacy; privacy is undoubtedly an important consideration within a consumer IoT context, where concerns include who data is being shared with [55], [56], [57], [58], [59], [60], [61], [4], how that data is being used by organisations [62], [1], [55], [56], [57], the intrusiveness of device sensors (e.g. cameras, microphones) [1], [55], [56], [57], [63], and broader concerns over a loss of control of personal information [1], [55], [56], [57], [58], [4]. Indeed, such concerns are for good reason, with recent studies showing that a considerable number of consumer-IoT devices are 'leaky', in the sense that they involve substantial, potentially non-essential, communication with third-parties [64], [65], [66]. Further exacerbating these issues is that consumer IoT devices will often be used within the home, where devices can potentially capture personal and intimate information [10], and 'privacy norms' may easily be violated [67]. As a result, such privacy-related concerns have been identified as key factors affecting users' purchase behaviour [55], [59], and whether they trust [58], [68], accept [69], [36], and/or adopt [68] smart devices within their homes.

While privacy remains an ongoing challenge in the IoT, there have been a plethora of research efforts that show promise in tackling particular privacy issues. For example, there have been a number of recent developments with regard to privacy enhancing technologies for the IoT (see [70], [71]), as well as researchers proposing methods for the developers of domestic IoT devices to better meet the privacy expectations of users, for example, by way of privacy norms [72], privacy-oriented design implications [73], [74], and recommendations [75], [76], [57], [77], [4]. Researchers have also developed and evaluated storyboards [78], prototypes [79], and tools [80] for assisting privacy-oriented concerns, privacy controls which provide varying levels of detail about why certain data is being requested [81], and nutrition labels [55] to inform consumers across a range of IoT-related privacy issues.

Importantly, while many of these concerns have (understandably) resulted in privacy-oriented outputs (be they interventions, prototypes, designs, resources, etc.), there is also scope for broader transparency mechanisms to assist here. That is, transparency mechanisms can support greater levels of oversight and understanding over smart home operations, and this information can, in turn, assist in privacy contexts. For example, such information might reveal the data being collected by an IoT device, where that information flows, how it is being used, and so on, which can indicate potential privacy concerns. Furthermore, such mechanisms also enable various actions to be taken in response, perhaps prompting users to change privacy settings, restrict where data is flowing to (e.g. see [65]), engage with particular parties to challenge or exercise data protection subject rights (as has been discussed at a high-level [52], [64], [80], [82], [83]), and so on. Once again, research which focuses on the transparency needs and expectations of users will have much to offer, and complement privacy-oriented research into the IoT.

3) Security: The IoT also gives rise to a range of security issues [84], [85], [86], [87], [27], [82], [88]. IoT deployments may comprise a large number of different devices (perhaps manufactured by a range of organisations, with some having more stringent security practices than others), and thus potentially introducing numerous possible points of failure [9]. Moreover, the IoT will often entail physical elements (e.g. actuations), where the consequences of security incidents may be severe by resulting in physical world harms [11], [12], [27]. And again, IoT devices can reside within intimate or sensitive locations within the home (e.g. bedrooms, bathrooms) [10], making the prospect of security issues particularly concerning. In all, security issues within the consumer IoT is an ongoing concern.

Unsurprisingly, IoT security is a notable concern for consumers; work exploring consumer attitudes to IoT adoption have identified a number of particular reasons for this, including concerns over data breaches [1], [56], [57], the perceived challenges of achieving a secure smart home [28], [89], [57], [90], access and power imbalances in multi-user smart homes [91], [77], the intrusiveness of another party being able to access sensor data (e.g. cameras, videos) [1], [92], [55], [56], the risk of having physical devices manipulated by a malicious actor (e.g. switching off the freezer, manipulating smoke detectors [1], [92]), and the perceived risk of physical safety (e.g. unlocking doors or windows) [92], [56], [57], among others.

Again, mechanisms for improving IoT transparency may have a role to play in relation to many of the above security concerns, for example, through providing information and oversight over the operation of devices (i.e. what drove particular outcomes and where data is flowing to [6]), to help confirm that devices are interacting and behaving in line with expectations [12], [93], and in identifying anomalous behaviours and possible security threats [94], [95]. Such information (be they access logs, provenance information [6], [96], reports, or other forms of ongoing system monitoring [97], [93]) can facilitate general oversight, targeted investigations, and subsequent follow-up actions to mitigate these concerns perhaps prompting the user to seek further information from device manufacturers, or even removing components from the IoT deployment altogether [9]. We are already seeing some proposed solutions for communicating such security-related information to users (again, such as nutrition labels [55], [98]).

B. Towards meaningful transparency: Addressing a gap

Importantly, in each of these three concerns, effective transparency mechanisms have much to offer, through better supporting users in overseeing and monitoring how their smart homes are operating, helping to identify and diagnose potential issues as they arise, and enabling informed responses as a result. Specifically, this can help mitigate issues of understandability (§II-A1), through providing relevant information about the system's operation in a way that is meaningful to the user, thereby supporting them in comprehending why the system is behaving in the way that it is. Such information can also help address *privacy* concerns (§II-A2), through providing clarity over when and where data is being transferred, which can facilitate follow-on actions to prevent, mitigate or seek recourse for issues relating to privacy. Finally, these oversight mechanisms can also help issues relating to security (§II-A3), through supporting users in verifying their correct operation (thereby assuaging such concerns), and in identifying, investigating, and seeking support and restitution in the event of security breaches. In all, meaningful transparency mechanisms allow users to monitor for, uncover, and act upon many of the types of concerns raised in the literature (as just discussed).

However, the question of how transparency mechanisms can best support users, such that they are meaningful and effective, is currently under-considered. There is therefore a clear role for research towards more effective transparency mechanisms by working with potential users throughout the design process to explore how transparency mechanisms can better reflect the needs and expectations of those that would stand to benefit from them. This means that there are research opportunities towards advancing our understanding of what meaningful transparency mechanisms in the consumer IoT might actually entail (as is our focus)-focusing on what users want to know, how they expect to interact with such, the benefits they perceive such mechanisms would provide, and so on-and to explore how such mechanisms might support wider concerns of understandability, privacy, and security that arise in the literature more broadly.

To reiterate, the topic of smart home transparency is important – through providing information over the operation of smart home systems, transparency mechanisms help support various follow-on actions in response to particular issues. Having greater levels of oversight might prompt users to change system settings to prevent egregious behaviours, verify correct system operation, remove devices from the IoT deployment entirely, or even identify and challenge particular parties (such as through legal means), etc. That is to say, having transparency mechanisms that are more effective, by being user-derived and user-centric, can work to assist users with a range of different challenges, issues, and concerns that they may face regarding the IoT.

III. METHODOLOGICAL OVERVIEW

Our approach entails working with prospective IoT users to explore, design, and evaluate mechanisms for bringing about greater levels of transparency within the consumer IoT. Specifically, we explore the general needs and expectations that users have with regard to transparency mechanisms, collect and categorise transparency-related user requirements, and develop and identify elements of design that the participants felt would enable meaningful oversight within a selection of scenarios. To do this, we undertake three user studies:

Study 1 (§IV) involves a survey to explore what transparency information they would consider useful within a smart home context, how they feel that information should be communicated to them, and how they consider one might use that information (i.e. what they foresee such information enabling).

Study 2 (\$V) comprises two independent workshops of 5–6 participants, representing the 'users' of smart homes. During each workshop, the participants work together to identify user requirements and design system prototypes that represent what they believe to be effective transparency mechanisms for the consumer IoT. This is to gain detailed insight into the types of transparency mechanisms that participants thought would be useful to realise, and how they foresaw using such interventions in response to particular interests, concerns, or events.

Study 3 (\$VI) takes the insights from the prior two studies, and validates their relevance and applicability through an evaluation survey with an altogether new set of 56 participants. Here, we look to explore the efficacy of our previous findings with a new sample, and find that our results from Study 2 appear to generalise (to this new cohort).

Our methodology was structured such that each study builds upon the findings and insights of those prior, while uncovering related insights and discussion points throughout the process. Note that all of our studies were approved by our departmental ethical review board, and all participants were compensated for their time in an amount reflecting the UK's 'living wage'.

1) Participant recruitment: For our surveys (Studies 1 & 3), we used Mechanical Turk (MTurk) for participant recruitment. MTurk is a widely popular recruitment method within academic research [99], and allows us to purposefully 'cast the net wide', to document a diverse set of different perspectives and ideas that we could find, enabling sampling at scale. It is for this reason we placed few restrictions on demographics so as to enable a range of participants, though we did enact some restrictions. For example, we restricted to those likely to have some command of English, and have a good MTurk task success rate (in line with guidance from other studies [100]). We also, purposefully, did not restrict the survey to IoT users (so as not to exclude perspectives of those who have not yet adopted such devices), nor did we limit participants based on levels of technical expertise (so to provide insights from various backgrounds).

For the co-design workshop (Study 2), we selected two groups of participants; the first group comprised respondents of Survey 1, those with a general background bringing with them a broad set of perspectives. However, given that Study 2 involves creating user requirements and involves visual design, we recognised that there were also advantages in recruiting participants with some level of experience in technology design, so as to provide a complementing and contrasting perspective to that of the other group. For this we recruited a second group comprising a 'convenience sample' [101] of undergraduate computer science students. This allowed us to leverage their experience of user experience (UX) principles, and offered a point of comparison (to explore the similarities and differences) between the students and the more 'general' users. Full details of participant recruitment are described with each study, and we further discuss the implications of our participant samples in §VII-B.

2) Grounding the studies: Across all studies, we make use of a selection of scenarios to ground and motivate our work. Each of these are carefully designed to reflect actual concerns that people have, with many motivated by examples discussed in the literature and observed in the real-world, including issues associated with data leakage [102], [65], [103], [104], physical harms [12], system malfunction [31], [32], [33], targeted advertising [103], system operation [28], [11], [34], [30], [12], [29], etc. Further information about each study are elaborated in their respective sections, and all data has been made available via GitHub [105].

IV. STUDY 1: SURVEYING USER INTERESTS AND EXPECTATIONS FOR TRANSPARENT SMART HOMES

Our first study entailed an online survey designed to uncover a broad understanding of what meaningful transparency within a consumer IoT context might entail. We focus on four main questions; given a scenario where something within a smart home warrants attention or goes wrong: i) what types of information do respondents feel is important to know?; ii) how do they expect this information to be communicated (in terms of system interaction)?; iii) how do they expect this information to be presented (in terms of design)?; and iv) what types of follow-on actions do they think this information would enable them to take? By exploring these questions, we obtain a better understanding (particularly within a smart home context) of the types of transparency mechanisms that users might come to expect, how these might work to meaningfully communicate the relevant information to the user, and how our participants would seek to use such mechanisms to support wider accountability aims.

A. Method

We recruited 126 participants from Mechanical Turk to take part in this questionnaire (survey). As Table I shows, these respondents reported a range of technical expertise: 23% claimed to have 'no' or 'some' knowledge; 32% had 'average' knowledge, and 45% had 'advanced' or 'expert' knowledge, thus reflecting various aims, understandings, and expectations. While we did not restrict the survey to only those actively using consumer IoT devices (§III-1), we nevertheless observed that 94% reported having smart devices within their homes, demonstrating that the vast majority of our respondents were actively being impacted by consumer IoT devices, and thus had some familiarity as to what the IoT represents. In this way, our sample is reflective of the aims earlier mentioned (§III-1) comprising a range of interests, expectations and endgoals, while still having a baseline understanding of consumer IoT products.

TABLE I	
DEMOGRAPHIC INFORMATION OF THE PARTICIPANTS FROM STUDY	1.

	% of respondents
Gender	
Female	31%
Male	69%
Other	0%
Age	
18–29	35%
30–39	37%
40-49	20%
50-59	8%
60+	0%
Technical Expertise	
No knowledge	5%
Some knowledge	18%
Average level of knowledge	32%
Advanced knowledge	34%
Expert knowledge	11%
Knowledge of Smart Devices	
No knowledge	1%
Some knowledge	25%
Average level of knowledge	53%
Advanced knowledge	39%
Expert knowledge	8%
Have Smart Devices in the Home	
None	6%
One or more	94%

After signing up and providing consent to participate, the main body of the questionnaire posed a selection of hypothetical scenarios to the participants, who were then asked open-ended questions relating to the above four questions. These scenarios, outlined in Table II, were based on real-world concerns and incidents; one related to suspected data leakage from voice assistants [102], [103]; one concerned smart windows opening when they shouldn't have [106]; one related to targeted advertising within a smart fridge [107], [104]; and one involved a smart lock not behaving as expected [12]. In this way, our scenarios were grounded in actual issues that users of such devices might face.

Participants were initially presented a brief description of one of the four scenarios (randomly selected), where some particular issue or concern warranted attention (see Table II). These, again derived from real-world scenarios (§III-2), were written in such a way as to have multiple possible reasons for the issue occurring, and the exact nature of what we had determined was happening within the smart home was withheld to participants at this stage of the survey. Participants were first asked what they thought was the most important information to know in that scenario. They were then taken to a new page, containing some further underlying information about the exact nature of that scenario, and were then asked: how should this information be communicated to the user; how might this information be structured or presented; and what follow-up actions would having access to such information facilitate. This process was then repeated with a second scenario (randomised ordering), giving us a range of responses across all of the scenarios while limiting the time commitment required of the participants.¹

After completing the two scenarios, participants were then

¹Note that as part of a pilot run, a few of the participants did complete this process for all four scenarios, before the number of scenarios presented to each participant was reduced to two to better reflect the anticipated time commitment that the study required.

Scenario	Brief Description	Underlying Information
Voice Assistant	You have seen news reports about certain smart voice assistants constantly recording audio and sending it to the manufacturer. Given this, you wish to check your own smart assistant to verify what, and when, information has been sent outside of your home.	After further investigation, you learn that i) information is only being sent to the voice assistant's manufacturer when it specifically hears its trigger word (its name being called); ii) an audio recording is then sent to the man- ufacturer for processing, and an audio response is sent back; iii) the voice assistant is seen to be communicating with a number of different advertising companies.
Smart Windows	You wake up in the middle of a particularly cold night, noticing that your smart windows have opened automatically. You know that the windows are automatically set to open when it is above a certain temperature inside. However, the indoor temperature feels far too cold for this occur, and the windows should therefore not have opened.	After further investigation, you know that your smart home contains three indoor temperature sensors, which are accessed by the smart home system. These readings are used to determine whether or not the windows should be opened (as well as for other purposes). One of these devices has been reporting temperatures far higher than the other two, suggesting it may have malfunctioned. These unusual readings began just after midnight.
Smart Fridge	Your Smart Fridge allows you to keep an inventory of what is stored inside, and can build a shopping list for use on-the-go. You receive an email from the fridge's manufacturer, updating the terms of service to allow this data to be used for advertising purposes with other companies. As a result, you are concerned about the privacy implications of your shopping habits being used for advertising.	After further investigation, you find that your fridge is categorising the types of items that you buy in order to predict characteristics about you (e.g. 'vegetarian', 'health conscious'), sending these the manufacturer. Fur- ther, you learn that the fridge is also sending this infor- mation to supermarkets and other advertisers, allowing them to send you adverts which they think will be of interest.
Smart Locks	You have a smart lock which should automatically lock your front door every evening after sunset. One evening, you notice that something is preventing the door from locking, despite it being dark outside.	Your smart home works to unlock the front door when- ever you arrive home. This works by a sensor that detects when your mobile phone is near the smart lock. Investigating the issue, you discover that your mobile phone is still being detected by this sensor, despite it being several rooms away from the front door. This appears to be preventing the front door from locking.

 TABLE II

 The four scenarios that participants were asked about.

asked some broader (general and demographic) questions about themselves and their attitudes to smart homes. These included questions relating to themselves and their technical expertise (Table I), and their broader concerns and interests relating to smart homes more generally (Fig. 2). Finally, we included an optional opt-in field where participants could enter their email address if they were happy to be contacted for subsequent stages of the research. All open-ended questions were analysed using thematic analysis [108]; answers (either in part, or in entirety) were categorised under multiple themes, and this was conducted iteratively until themes no longer changed as a result of new data.

B. Findings

Our initial survey uncovered a range of insights about meaningful transparency within a consumer IoT context.

1) Most respondents had concerns about the nature of smart devices, and wanted to know more about their operation: Asked directly, we found that over 80% of respondents were at least somewhat or extremely concerned about the nature and operation of smart devices in general (see Fig. 2). Of those who selected "I'm extremely concerned", further probing on their open-text elaborations for this answer revealed that this was largely either due to surveillance and privacy concerns, or device security and personal safety concerns (each of these categories making up ~1/3 of "extremely concerned" responses). This corresponds with the topics generally focused on in the literature (§II-A), while clearly indicating the importance of transparency mechanisms to help support such issues. In contrast, of the 19% that were "not at all" concerned, this was

typically due to them not believing that "there is reason to be concerned", with $\sim 1/3$ of these responses elaborating the sentiment that "there isn't any useful information someone is going to get out of my [data]". The remaining group of "somewhat concerned" respondents seemed open to the idea of using smart devices in the home, but still expressed caution over how they operated. Again, when asked directly, over 80% of respondents indicated that they were either somewhat or extremely interested in finding out more generally about how smart devices were operating and communicating (see Fig. 2). Such findings demonstrate the appetite for greater transparency, and for having the ability to oversee, inspect, and understand what is happening within smart homes.

2) Many respondents wanted to be able to oversee technical specifics: We again used thematic analysis [108] to explore what participants thought was important to know within each particular scenario. Open-text responses were categorised, and themes raised by at least two participants are presented in Table III, alongside their prevalence (the proportion of responses that were categorised under each theme). Note that given participants were only shown two of the four scenarios (randomised), each scenario has a different number of responses. The full set of responses and their codes are included within the supplementary materials.

One of the main themes to emerge, particularly for the window and smart lock scenarios, was the need to understand and verify what was going on. Indeed, nearly half of the responses for these two scenarios wanted information which allowed the participant to understand and verify what drove this actuation (49% and 45% respectively), such as accessing



Fig. 2. The majority of Study 1 respondents were concerned about the nature and operation of smart devices, and interested in finding out more about them.

TABLE III A THEMATIC ANALYSIS FROM STUDY 1. THEMES WERE CODED FROM PARTICIPANTS' OPEN-TEXT RESPONSES, INDICATING WHAT INFORMATION THEY FELT WAS IMPORTANT TO KNOW FOR A GIVEN SCENARIO.

% of responses categorised under each theme	
Voice assistant $(n = 67)$	
What information was recorded by the voice assistant	38.8%
Why the information was recorded by the voice assistant	16.4%
When the information was recorded by the voice assistant	10.4%
Where the information was stored	6.0%
What information was transferred over the network	52.2%
When information was transferred over the network	19.4%
Why information was transferred over the network	6.0%
How long information will be stored by other parties	6.0%
Where information will be stored by other parties	10.4%
Who has access to the information once sent to other parties	20.9%
How information is used by other parties	22.4%
How to prevent this from happening in the future	25.4%
Smart Windows (n = 74)	
Why the window opened	48.6%
When the window	12.2%
What data source triggered the actuation	48.6%
What data (readings) drove the actuation	47.3%
How to prevent this from happening in the future	14.9%
Smart Fridge (n = 69)	
What information was recorded by the smart fridge	7.2%
What information was transferred over the network	53.6%
When information was transferred over the network	8.7%
Where information will be stored by other parties	5.8%
Who has access to the information once sent to other parties	34.8%
How information is used by other parties	44.9%
How to prevent this from happening in the future	33.3%
Smart Locks (n = 62)	
What happened to the smart lock	3.2%
Why the smart lock didn't secure	35.5%
What data source prevented the door from locking	45.2%
What data (readings) prevented the door from locking	25.8%
How to prevent this from happening in the future	43.5%

"a history log" of the smart home, "to see why this triggered" and "whose fault it was". This is particularly interesting, given the literature on understanding smart homes (§II-A), where there was a potential disconnect between users' understanding of what was happening within smart homes and what was actually going on [29].

Looking to the voice assistant and smart fridge scenarios, there was also an appetite for overseeing technical specifics – though this time predominantly regarding data flows. For example, the most frequent theme emerging from these responses indicated the importance of information regarding *what* was being transferred outside of the smart home (52% and 54% of responses for the voice assistant and smart fridge scenarios respectively), *who* could access that information (21% and 35%), and *how* that information was being used (25% and 33%). However, these two scenarios also tended to raise the topic of *prevention* for this unexpected smart home behaviour; over a quarter of responses for the voice assistant and a third of

responses for the smart fridge scenarios specifically indicated that it was important to be able to "opt out", "leave this service", and "stop this level of invasive behaviour".

3) Interaction techniques and information visualisations offer one way forward: A common suggestion raised by participants with regard to how they expected this information to be conveyed to them was the desire for automatic notifications when a discrepancy or some anomaly was detected (such as the example of the faulty sensor providing readings far higher than would be normal). For example, one commented "[I'd want] to be automatically notified if a discrepancy between them is recorded at any point". This was typically described by respondents as an alert on the system, a mobile notification, an email, a phone call, etc. There were also descriptions of means to allow the user to perform more targeted investigations (i.e. in response to a particular incident or concern), alongside those driven by interest and more exploratory in nature. In terms of how that information should be communicated, responses were often contextual. For example, in the voice assistant scenario, some referred to asking the voice assistant, as a means to interrogate what is going on – a method previously suggested in the literature [34]. In contrast, for the window scenario, many of the responses described graphs of temperature readings (which have also been previously explored [46]). Other suggested responses included lists and tables (of raw numbers), summarised explanations, schematic diagrams, and other forms of data visualisation.

4) Enabling a means to take control: On the subject of what types of follow-on actions such information would enable, a few common themes that emerged related to the ability to disengage with the device or devices in question, should they be behaving in undesirable ways. For example, for the scenarios concerning user privacy (the voice assistant and smart fridge), respondents indicated that they would reassess their use of the device in question some sample quote include: "I would restrict my use of this voice assistant", "I'd have to seriously consider if owning a smart fridge was right for me", "this is a smart device that I just wouldn't have". Again, this demonstrates the key role that transparency information can play, for example as regards technology acceptance and adoption, by showing how it enables and supports effective oversight and actions in response to what occurs.

C. Summary

Throughout this survey, we probed respondents for their thoughts, interests, needs, and expectations as regards transparency mechanisms for the consumer IoT. From the results, we can see that participants did express concerns about how smart devices were operating, and there appeared real appetite for transparency measures that granted access to such information – both when things began to go wrong, and for wider aims of curiosity, validation, etc. The results indicate that there is much scope for research that explores how such transparency mechanisms might come to be expected by potential smart home users in practice. In all, the respondents appeared to recognise the benefits of having the ability to oversee their smart homes, and the many opportunities of doing so.

V. STUDY 2: CO-DESIGNING TRANSPARENCY REQUIREMENTS AND DESIGN ELEMENTS

In the previous study, participants had clear concerns over the nature and operation of smart devices, and overwhelmingly expressed an interest in having effective transparency mechanisms for the consumer IoT. Building on this, we next further probed on how IoT users believe that such transparency mechanisms should work in practice. To do this, we take a user-centric approach [53], undertaking two co-design workshops with participants to derive *their key considerations and design ideas* for transparent smart home systems. By doing so, we uncover the types of transparency mechanisms that our participants want and expect from the consumer IoT.

A. Method

Our two workshops entailed the same process, each lasting two hours and involving separate groups of participants. These participants were tasked with working together to complete two activities. The first activity involved the participants deriving user requirements for making smart homes more meaningfully transparent, allowing us to explore the types (or categories) of user requirements that they thought were important. The second activity involved the participants designing visual prototypes for investigating and understanding the operation of smart homes, which allowed us to derive a series of design elements (key aspects and features of design for enabling smart home transparency). These two activities were designed to provide tangible insights and ways forward-from the perspective of potential users-for how we might bring about consumer IoT systems with meaningful transparency in mind.

Our two groups were selected to bring a range of complementary skillsets and different perspectives, by having some with grounded experience in design vs. more 'general' users. For the latter, we began emailing all participants from the prior survey (§IV) that indicated that they would be interested in taking part in follow-up research. This email outlined the nature of the follow-up research, and asked those interested to complete a Doodle poll indicating dates that they were available. This resulted in six participants that were all available for a particular time slot (though one of these participants ultimately did not attend). This group therefore represented a fairly general group of prospective users (though with all having at least one IoT device within their homes).

Our second group comprised individuals with some knowledge of system requirements and interface prototyping, as a means of ensuring that their responses were grounded within some degree of systems design. To achieve this, we sent an email through our institution's computer science department (i.e. a 'convenience sample' [109], [101], [110], [111]), and recruited six undergraduate students to take part in this workshop. In addition to their knowledge of systems design, this cohort also allowed us to compare and contrast their outputs to that of the first (more 'general' or 'standard' user) group, to explore where similarities and differences might exist, and what insights might be learned as a result.

B. Activity 1: Co-designing transparency requirements

Participants were given an overview of the research, including what requirements are, what constitutes a 'good' requirement, and how they can be created. They were then asked to consider how smart homes might better cater to their interests and concerns, and to think about how these might be specified as user requirements. The activity itself was done on MURAL [112] (a collaborative whiteboard web app), and involved participants creating requirements (sticky notes with text), as well as moving and editing those created by themselves or others. Participants were also given access to some example quotes from the prior study as a means to stimulate discussion, though they were encouraged to include and contribute any of their own requirements should they have some in mind. Throughout this ~ 20 minute exercise, the participants were tasked with determining user requirements and prototype functionality as they saw fit. During this time, the researcher acted in a supportive role, chairing the codesign sessions careful to avoid biasing the outcomes with the researchers' preconceptions [113]. Once complete, participants then prioritised their requirements into three categories; 'Must', 'Should', and 'Could' (see Fig. 3), in line with the MoSCoW method of requirements prioritisation [114].

1) Activity 1 Findings – Categories of transparency requirements: This first activity led to 55 requirements being created across the two groups, with Group 1 creating 28, and Group 2 creating 27. To gain a better understanding of the types of requirements that were produced, each requirement was coded using thematic analysis [108] (in line with the process outlined in §IV-A) after the workshops had taken place. This was done so that we could explore not only the exact requirements that were identified, but the broader patterns and concerns that the participants focused upon. Through this process, a total of nine categories were identified; five comprised transparency-related concerns (Fig. 4), the remaining four concerned broader controls and mechanisms that the smart home should support (Fig. 5). These categories offer various insights into the types of transparency concerns that users may have, and the types of user requirements that may address these.

Given the focus of this work (on understanding how participants themselves felt transparency mechanisms should be achieved), our particular interest is in the five requirements presented in Fig. 4 (in contrast to those requirements concerning broader controls and mechanisms that the system should enable; Fig. 5). These five transparency requirements focus on *overseeing*, *exploring*, and *investigating* aspects of the smart

To Appear: C. Norval and J. Singh, "A Room With an Overview: Towards Meaningful Transparency for the Consumer Internet of Things," in IEEE Internet of Things Journal. DOI: 10.1109/JIOT.2023.3318369.

Users must be able to see the status of devices	The application should be able to alert the user if something suspcious occurs with the devices	Put data collection information on the product's listing before purchase	The system should be able to convey information visually and through sound (for accessibility)	Have of er colo	visual signs rror, such as our changes	The user should be able to add other users to view the status	Minimal technical language in high- level log explanations	System should record energy usage of each device in the home	The system could be accessible through a mobile and web app	Have moderated online forums where people can ask questions	There could be physical hub with a simple UI that shows information	The system could have a visual representation of the home
Have useful/safe default settings, such as your doors being locked	the system must allow users to set up alerts related to key safety devices behaviour	The system should have permissions for privacy control	Accurate logs of actual errors, not just erroneous actions	The sy visual what h	ystem should support lisations over nas happened	Users should be able to export information to other services/devices	The system should record the cause of actions	More information and notifications about system updates	Be able to get an overview of the device's features easily. e.g. by saying "help" or pressing a help button	Have digital setup guides, as well as paper ones, as these can easily be updated		
The system must be able to be manually overriden (with appropriate permissions, like keys)	The system should notify users when specified actions or errors happen	Record when updates are installed and by whom		The sy prov sup iss	ystem should vide further oport when sues occur	The user should be able to access more granular, hierarchical information about what happened						
Must				Should			 Could					

Fig. 3. A screenshot from Group 2's MURAL board. Participants could co-create requirements (sticky notes) and prioritise them as either 'Must', 'Should', or 'Could' by moving them into boxes.



Fig. 4. Derived categories of requirements (with examples) that focused on transparency.



Fig. 5. Other categories of derived requirements (with examples) regarding broader concerns.

home, as well as being *notified* when unusual behaviours or activities are identified, and conveying the relevant information in a *usable* way. As shown in Fig. 4, we provide a representative selection of four requirements of each category; see the supplementary materials [105] for the full set.

In all, this analysis gives us a broad set of categories reflecting user-derived requirements for bringing about greater levels of transparency regarding the operation of smart homes. Looking at which requirements were more closely associated with the 'Must', 'Should', or 'Could' of the MoSCoW priority system, we (anecdotally) observed some patterns whereby those categorised as 'must' tended to focus on system functionality, such as specific features that the system should facilitate, while those prioritised as 'should' often appeared more presentation-oriented, such as the use of terminology, data visualisations, and support tools. This may perhaps be because there are many different ways in which information can be communicated to users (c.f. functionality), however, future research may be able to probe further into this.

C. Activity 2: Co-designing transparency prototypes

The second activity involved tasking participants with creating a prototype for a 'transparency interface': a tablet/wall panel-based system, which interacted with devices within a smart home to provide greater transparency over how the devices were operating, what they were doing, etc. This acts as a means for analysing and exploring the *various types of transparency-related elements of design* that the participants created. That is, much like our process of identifying the

To Appear: C. Norval and J. Singh, "A Room With an Overview: Towards Meaningful Transparency for the Consumer Internet of Things," in IEEE Internet of Things Journal. DOI: 10.1109/JIOT.2023.3318369.



Fig. 6. The 'Misbehaving smart bulb' scenario. Top: Group 1's interface; Bottom: Group 2's interface.



Fig. 7. The 'Investigating adverts' scenario. Top: Group 1's interface; Bottom: Group 2's interface.

categories of requirements, we use these prototype designs to understand more about key aspects of system design for enabling smart home transparency that our participants' interfaces contained, and how the participants perceived the design of such features in practice, so to derive the design elements. Again, this activity took place via MURAL, using simple shapes, text, and icons to put together a simple set of storyboards showing what features such a system could have and how it would be used to investigate and explore the operation of smart homes.

The participants began by creating the prototype's home screen; that which would be seen when initially interacting with the tablet or wall panel. Participants were at liberty to add components (icons, text, shapes, etc.) to the prototype as they saw fit, and the process was chaired (i.e. moderated) by the researcher. Once the participants had created the home screen, we then performed two sub-activities in turn, each relating to a hypothetical scenario where something in the smart home warranted attention or investigation. The participants were first presented one of these scenarios, tasked with creating the subsequent screens that would 'storyboard' how their prototype could be used to investigate the scenario. They then repeated the process on the second scenario. In all, these activities helped to further derive a set of transparency design elements from the participants' prototypes, while providing a context to show how their systems might work to enable meaningful transparency from their perspectives.

The two scenarios were designed to explore and reflect realworld concerns (elaborated below), which work to ground and contextualise the activity for the participants. Both scenarios were formulated such that there could have been a number of potential underlying reasons for the concern, and a number of ways in which someone could use transparency mechanisms to investigate. As the participants designed new screens that would allow them to dive deeper into their prototypes, the researcher provided more information as to the actual nature of what was happening within the prototype. For example, when participants were designing the Troubleshooting (Group 1) and Log (Group 2) pages, the researcher informed them as to what these logs would report (remote requests were originating from overseas), as a way to simulate the discovery and diagnosis process for a previously unknown issue (a staged approach, similar to that used in Study 1; §IV-A). In this way, we gathered information about what design elements the participants expected such a system to have, what these features and pages within the system might look like, and how they would expect

Elements of design	Description
Status indication	Both groups created a 'Status' feature on the homepage, which presented information about the system's operational
	capabilities. This acted as a means to notify the user whenever potential issues were identified by the system, and directed
	them to either more information or tools to assist in understanding and rectifying the problem. In Group 1's case, the
	indicator entailed a direct route for navigating directly to details of the issue, whereas it appeared more of a prompt for
	Group 2 to dig into the interface and identify the issue at hand.
Devices list	Both groups also had a 'Devices' list, which presented a list of devices currently deployed within the smart home. This
	was used by Group 2 to navigate to the smart bulb during the first scenario. While not fully elaborated by Group I (given
	that they used a different means to navigate in both scenarios), they nevertheless included it as an option on the home
	screen. Means of navigating to the correct devices will likely be important in a transparent smart home context (particularly
	where several devices may be incorporated), and this approach is one potential way of enabling navigation in smart home
	contexts.
Rooms list	Similarly, both groups had a way of navigating through the physical spaces (i.e. rooms) of the house in which their smart
	devices were configured and deployed. That is, both prototypes allowed the user to group and explore devices in a manner
	relating to the environment's physical layout. This was part of the Devices list for Group 2 but was its own menu option
Drive ev. settings	In Choup 1.
Privacy settings	Again, boin groups nad a Settings option on the nomepage, and boin settings pages prominently realized a Privacy sub-
	option, while our scenarios and not lead to entire group intruer elaborating on what such a privacy settings page lingu
	contain, our groups used the privacy options at the top of the first in their settings page – periaps indicating the importance that the aparticipants plaqued on privacy options at the top for its monoament and control
Viewable system logs	that the participants placed on privacy and the means for his management and control.
viewable system logs	Bour groups had options for viewing system logs, which played a prominent role in investigating the first section of the secti
	more user friendly ways of understanding a system's operation (cf system logs) Howaver this nevertheless emphasizes the
	contextual nature of transparency mechanisms in that both 'lay' explanations as well as more technical system logs
	each be useful across different contexts and scenarios. Furthermore, the exact presentation of the system logs different
	Group 1's was slightly more tabular in nature, whereas Group 2 (the students) onted for more of a user-facing explanation
	Yet much of the information presented in the interface was the same for both groups.
Connectivity records	Similarly the groups both included features relating to evice connections connectivity and any internal/external interactions
Connectivity records	with other devices or services. While this was stored within the 'Settings' menu for Group 2 Group 1 had this menu option
	as a prominent option on their front page. Given the nature of our scenarios, this design element did not end up getting
	elaborated through either group's storyboards, but their inclusion in both interfaces indicates the potential importance of
	this feature.
Advertising inferences list	Both groups also outlined a page that would show advertising information (e.g. the inferences and profiles generated about
	the user, and how they were being used). In both cases, this page was accessible through the system's Settings page, and
	outlined a means to oversee and control aspects of the advertising profile that had been created within the smart home
	systems.

TABLE IV THE DESIGN ELEMENTS THAT EMERGED FROM THE PARTICIPANTS' PROTOTYPES.

to be supported in eventually finding the information that they were looking for. In all, this design exercise produces a rich set of insights into how our participants thought transparency mechanisms might better work to inform.

Scenario 1 – Misbehaving smart bulb: The first scenario was security-related, featuring a smart bulb that was being acted upon by a malicious actor (hijacking a smart bulb appears a common exemplar [115], [116], [117]). The participants were told that the bulb starts to behave strangely, changing colours seemingly unprompted and occasionally flashing rapidly, and were then tasked with co-designing the steps that they would take on their interface, to investigate this issue further.

Scenario 2 – Investigating adverts The second scenario was privacy-related, and involved suspected data leaks being used for targeted advertising (a real concern that many people have [102], and that has been observed happening in practice [118], [65]). This involved participants being told that their recently purchased smart doorbell (with a camera) had been placed on their porch overlooking their front garden, and at around the same time, they started receiving personalised adverts on their smart TV for gardening. Given that the camera would have seen them gardening, the implication is whether their new device might be invading their privacy, by sharing or leaking information about their gardening habits with the TV's advertising provider. As the participants created

the prototype interface screens, these would reveal (steered by the researcher) that the doorbell was not involved in the gardening adverts; rather, the gardening inference was made as a result of i) their smart TV's viewing history of gardening shows, and ii) online purchases for gardening equipment. As such, the scenario was designed to represent a case where initial suspicions led to further investigation, revealing that the system was operating in a different way from what was originally expected. Also in contrasting with the first, the second scenario did not relate to a specific technical issue or incident; rather it involved detailing the steps that the participants would expect to take more broadly, to diagnose and understand what was occurring within their smart home through a targeted investigation.

1) Activity 2 Findings – Design elements for transparent smart homes: This activity resulted in two prototype systems (one produced per group), each with two 'storyboards' [119] (for each scenario). These storyboards illustrated how the user would engage with the transparency interface to investigate and understand what was happening within their smart home. See Fig. 6 for the storyboards for the first scenario, and Fig. 7 for the second scenario).

Note that the two groups (and, by extension, the two workshops) were run independently, in that neither group saw the prototypes, requirements, or any other outputs from the other. Interestingly, however, the two sets of storyboards appeared to show a number of similarities and overlaps between the two groups in term of their transparency-oriented design feature and functionality, suggesting that these overlaps may be more generally applicable. As such, we opted to thematically categorise aspects of the prototypes' interfaces (much like we categorised responses in our earlier survey, and categorised the types of requirements in Activity 1), in order to explore this concept further. We call these categories of interface components *design elements*.

This involved thematically grouping aspects of the design, features, and functionality as they concerned the transparency mechanisms prototyped by our participants. For example, both groups' prototypes contained a 'Devices' page or tab, which presented the user with a list of devices that were deployed within the smart home. This resulted in a 'Devices List' design element, which we then created a description for, based on how the groups had implemented it.

As Table IV presents, there were a number of common transparency-related design elements that both groups had independently devised. As such, this list represents a set of corroborated, user-derived design considerations relevant for those designing and/or developing smart home systems. Note we discuss the implications of some of the more unique design decisions (i.e. those incorporated by only one group) in §VII-B.

D. Summary

Through these two co-design workshops, participants' perspectives, needs, and expectations were explored with regard to how they believed that transparency mechanisms for smart home technologies should be brought about in practice. These workshops resulted in range of user-defined insights toward the types of features and functionality that might help bring about greater levels of transparency within the consumer IoT. We discovered nine groups of requirements-five related directly to transparency-related concerns (Fig. 4), and four concerned broader controls and mechanisms that the smart home should support (Fig. 5). We also derived seven design elements (Table IV) representing the key aspects of system design (that both groups had outlined in their prototypes) for enabling smart home transparency. These outcomes provide us with tangible ways forward for the design and implementation of meaningful transparency mechanisms within the consumer IoT – as envisioned by our participants.

VI. STUDY 3: VALIDATING THE EFFICACY OF OUR DESIGN ELEMENTS AND TRANSPARENCY REQUIREMENTS

Through the previous studies, we have derived i) a selection of categories, representing the types of user requirements that our participants raised, and ii) a collection of design elements, offering practical ways forward for developers and designers wishing to bring about improved transparency within their IoT products. Here, we describe a further study to validate whether i) the design elements appear usable and effective in supporting meaningful transparency, and ii) the perceived relevance and general coverage of the types of transparency requirements that we have identified.

To explore whether the applicability of our categories of user requirements and design elements appeared to carry forward, we recruited 56 new participants from MTurk to take part in this questionnaire (survey). While we used MTurk filtering criteria to ensure that none had taken part in our earlier studies (§IV-A), our new cohort had similar characteristics to that of Study 1; participants had similar demographics, expressed a range of technical expertise, and overwhelmingly had smart devices in their homes (as shown in Table V). Furthermore, by chance, when asked the extent to which they were concerned about the nature and operation of smart devices, and whether they were interested in finding out more about how they were operating and communicating (as we did for the participants in Study 1), our new cohort responded with similar characteristics (shown in Fig. 8) to those of Study 1. That is, the samples of both surveys appeared similar (in terms of demographics, technical expertise, in owning IoT devices, and general concerns and interests relating to smart home transparency).

 TABLE V

 Key demographic information of the participants from the third study.

	% of respondents
Gender	
Female	36%
Male	64%
Other	0%
Age	
18–29	38%
30-39	43%
40-49	14%
50-59	4%
60+	2%
Technical Expertise	
No knowledge	4%
Some knowledge	18%
Average level of knowledge	38%
Advanced knowledge	36%
Expert knowledge	5%
Knowledge of Smart Devices	
No knowledge	2%
Some knowledge	11%
Average level of knowledge	48%
Advanced knowledge	32%
Expert knowledge	5%
Have Smart Devices in the Home	
None	2%
One or more	98%

A. Activity 1: Exploring the efficacy of the design elements

We begin by considering the design elements derived from the workshops of the prior study. Our focus is exploring whether and how these design elements assisted meaningful transparency through communicating relevant information to these new participants. Given the similarities of the two groups' interfaces from the design study (§V-C), we opted to combine the key design elements (identified from §V-C1) into one interface. In other words, we merged the key features of each interface, while maintaining a consistent and coherent design throughout (so as not to distract participants with aesthetic considerations). The result was a new set of storyboards (i.e. a new interface; see Fig. 9) with which we could explore the efficacy of the common design elements that emerged.



Percentage of responses

Fig. 8. Again, the majority of respondents from Study 3 were also concerned about the nature and operation of smart devices, and interested in finding out more about them.



Fig. 9. The 'merged' interface, combining the design elements from both groups' prototypes. Top: Scenario 1; Bottom: Scenario 2.

After agreeing to take part, participants were first presented with a broad description of one of the scenarios (randomised order) from §V-C, and each storyboard image corresponding to that scenario was then shown to the participant, alongside textual descriptions of each 'step' in the process. Again, this reflected the investigative process of using such an app, taking an exploratory approach to illustrate how one might use the system to investigate the scenario at hand.

After the first scenario's storyboard was presented to the participant, they were then asked (via open-text boxes) what had happened; those that started with Scenario 1 were asked why the bulb was acting erratically, whereas those that started with Scenario 2 were tasked with reporting whether the doorbell was involved in the gardening adverts. We then asked a follow-up question relating to how confident they were in their previous answer (3-point Likert), alongside a further open text box where they could elaborate on their confidence. By asking these questions, we probe the interface's ability to both convey the appropriate information, and how confident the users felt that they understood what happened both key aspects with regard to meaningful transparency. This process was then repeated with the other scenario, exposing the participant to how the aforementioned design elements could assist in investigating both particular issues or concerns, while mitigating for order effect biases [120].

After completing both scenarios, participants were then

asked to complete a System Usability Scale survey (SUS) [121] about the prototype. The SUS comprises a set of ten Likert questions on aspects such as ease of use, consistency, and complexity, and can be used to generate a 0-100 score representing "a general quality of the appropriateness to a purpose of any particular artefact" [121]. In this way, we use the SUS to gauge the extent to which the participants found the prototype to be appropriate, or usable, as a means for engaging with the transparency mechanisms and understanding what happened within the system.

Exploring i) what participants thought happened in the system, ii) how confident they were, and iii) whether the system received a 'good' SUS score (indicating general aspects of system usability), provides three metrics about the extent to which the prototype supported participants in meaningfully interrogating the information at hand.

1) Activity 1 Findings – Exploring the potential of design elements: To explore whether the design elements helped in enabling meaningful transparency, we look to see what our participants thought what was happening within the two scenarios. For the first scenario, 66% of participants correctly listed a security breach as the potential reason for the bulb changing colour, and 91% of participants correctly identified the source of the adverts in the second scenario (Table VI). It is worth noting that, similar to the responses of §IV-B, these answers were evaluated based on open-ended text, as

opposed to selecting from multiple options, etc. That is, \sim two-thirds of responses specifically deduced that the requests coming from overseas was the result of a security breach. This is interesting, showing that while a majority of participants correctly identified that a security breach was to blame, a non-negligible proportion did not (with the remaining \sim third typically suggesting technical faults or other issues more generic). This, again, shows the importance of assisting users in interpreting the information, and that different users may require varying levels of support in doing so.

TABLE VI

RESULTS FROM STUDY 3, ACTIVITY 1. PARTICIPANTS INDICATED (VIA OPEN TEXT ANSWERS) WHAT THEY THOUGHT WAS THE UNDERLYING REASON FOR EACH SCENARIO, AND THEIR CONFIDENCE IN THE ANSWER.

Scenario 1: Misbehaving smar	t bulb
Correct $(n = 37)$	66.1%
Extremely confident	30.4%
Somewhat confident	58.9%
Not at all confident	10.7%
Scenario 2: Investigating adver	rts
Correct $(n = 51)$	91.1%
Extremely confident	41.1%
Somewhat confident	55.4%
Not at all confident	3.6%

Also relevant is the degree of confidence that users felt about their answers, as shown in Table VI. Interestingly, the two sets of proportions appear somewhat similar, despite proportion of correct responses being different; while ensuring user confidence is high is undoubtedly a key aim, it is worth recognising that those who are extremely confident in their interpretation of smart home data won't always be correct in their answers. In other words, while having transparency mechanisms that inspire confidence in users' understanding over how their smart homes operate is important, so too is working with various types of users—across different use cases and contexts—to ensure that these work to speak to users and their levels of expertise, and indeed, work to inform and empower, rather than mislead or oppress.

We then look to the SUS scores as another indicator as to the design elements' potential for facilitating meaningful transparency. Low SUS scores might indicate that the participants did not find the resulting functionality particularly useful or usable. In contrast, higher SUS scores indicate a generally more 'appropriate' [121] system for the task at hand (interrogating the transparency information to determine what had happened). Prior work has determined that a "poor" SUS score is around 35.7, "okay" is 50.9, and "good" is 71.4 [122]. Our merged interface received an mean SUS score of 72.1 (σ = 16), faring well according to these SUS benchmarks. Again, the SUS explores aspects such as consistency, complexity, ease of use, and perceived confidence in using the system, and offers several relevant dimensions toward enabling meaningful transparency (§II-B).

B. Activity 2: Evaluating the coverage and relevance of transparency requirements

We next consider the types (or 'categories') of requirements that were identified in §V-B. To recap, the co-designed requirements that were created in the second study were categorised into nine types; five related directly to transparencyrelated concerns (allowing the user to 'oversee', 'explore', and 'investigate' the transparency information, for the system to 'notify' the user when necessary, and several features relating to 'usability'; Fig. 4), whereas the remaining four concerned broader controls and mechanisms that the smart home should support (allowing the users to retain 'control', 'safety', and 'security' with regards to their smart home, with additional suggestions for user 'engagement'; Fig. 5). Given that our focus is on how we might bring about meaningful transparency mechanisms in practice, we focus on the former—those transparency-related—and explore the extent to which these appear representative and applicable to our new participants.

This activity began by presenting participants with a brief summary of each of the five transparency-related categories of requirements (see Fig. 4). This contained the five categories (oversee, notify, explore, investigate, and usability), alongside a short description and examples for each. Participants were first asked to prioritise these categories in terms of importance through allocating 100 points across each of these five categories (with more points representing higher importance). We also asked (via open-ended text) whether there were other categories or specific requirements that the participant thought were absent, but should be included. These questions allowed us to gain a sense of priority over these categories, and broadly to what extent they covered the transparency concerns or interests that the participants could identify.

Participants were then presented, in turn, three of the five categories (randomly selected, and presented in randomised order) alongside questions relating to the category in question. These questions asked for participants' general thoughts regarding that category; other examples of requirements that they thought would fit into this category; and how the participant thought those requirements could best be illustrated or implemented in practice. Again, these questions help to build up our understanding of the five categories that emerged from the co-design workshop and their general applicability across a wider sample.

1) Activity 2 Findings – Exploring the requirements' coverage: We start by looking at the prioritisation of requirements. Recall that participants are tasked with allocating 100 points across the five categories (which would result in 20 points for each category, if all were seen as equally as important). Each of these five categories appeared to receive similar scores (as seen in Table VII), and these proportions did not appear statistically different to each other. As such, it appears as though our five categories were all seen as broadly equal in importance by our participants, with no category appearing significantly more or less important (though further research with larger sample sizes could work to investigate whether and how these differences may manifest). In other words, each of our five categories appear to show value toward understanding the types of requirements that IoT developers may wish to aim for.

When asked if there were any other categories of transparency requirements that the participants thought were miss-

TABLE VII Users were asked to allocate 100 points across each of the five categories of requirements. By default, these were distributed evenly, with each category receiving 20 points.

Category	\bar{x}	σ
Oversee	20.3	9.4
Explore	15.1	7.0
Investigate	18.4	8.2
Notify	24.1	9.6
Usability	22.2	13.0

ing or should be included, 23.7% (n = 18) of responses included a suggestion (all of which are included within the supplementary materials). However, when using thematic analysis [108] to categorise these responses, all of the suggestions appeared to fall into our existing categories (i.e. those shown in Fig. 4 and Fig. 5). For example, one participant suggested "a category dedicated to protecting my smart home from hackers" (i.e. 'Security'); whereas another asked for "info on data is a big one and controlling who it's shared with" ('Control'). These findings are again interesting, given that they appear to corroborate the types of requirements and concerns being raised by our participants of Study 2, and suggest the coverage of our requirement-types is representative.

C. Summary

In all, these findings appear promising; they demonstrate: i) that the design elements (realised through the prototypes) appear to have potential in supporting users in meaningfully interrogating transparency information – demonstrated through a large proportion of participants correctly, and confidently, determining what was going on within the smart home; ii) that the merged interface was considered to have a 'good' [122] amount of usability, receiving an average SUS score of 72.1 (out of 100); and iii) that the categories of requirements appear to represent the transparency-related interests and concerns of an altogether new set of participants, finding that no new categories emerged from the exercise, and demonstrating the broad coverage, relevance, and applicability of these requirement types.

These findings offer a foundation for understanding how meaningful transparency might be achieved within smart homes (and, indeed, beyond). Next, we elaborate some of the implications of our work, and of transparency in the consumer IoT more broadly.

VII. DISCUSSION

Earlier, we described how there is a clear need for work which attempts to understand how we might bring about improved transparency mechanisms within a consumer IoT context, and the importance of working with prospective users throughout the process (§II-B). Towards this, we have undertaken a set of user studies that i) demonstrate the appetite for greater levels of transparency surrounding consumer IoT deployments (§IV); ii) identify paths forward toward the practical development of more meaningful transparency mechanisms, through understanding more about (a) the types of requirements and (b) design elements that our participants felt that smart homes should provide (\$V); and iii) validate the coverage of these types of requirements, and the efficacy of these design elements, with an altogether new set of participants (\$VI). Our findings provide practical insights for IoT developers and researchers alike, toward enabling more meaningful transparency mechanisms within the consumer IoT. In realising more effective transparency mechanisms, we can help to support scrutiny, and thereby accountability – a concern which will only grow in importance given increasing consumer demand and emerging regulatory requirements for such. As such, we now discuss some of the broader aspects of our work.

A. The importance of meaningful transparency

Across our three studies, we have demonstrated the clear desire that users have for greater levels of transparency within the consumer IoT. This was found rather explicitly in both the first and third study, where the vast majority of participants were somewhat or extremely concerned about the nature and operation of smart devices and environments, and were interested in finding out more about how their IoT systems were operating. Such findings demonstrate *the importance that many of our participants placed upon transparency mechanisms*, and their role in bringing about an IoT that more closely aligns with their needs and expectations.

However, as discussed, it is well-established that simply 'dumping' information on users will not necessarily be effective [13], [14], [15], [16], [17]. Indeed, this was recognised by some participants, for example, with one warning of the risks of *"information overload"*. Furthermore, also crucial is that transparency mechanisms do not work to mislead, distract, or to otherwise provide users with artificially inflated levels of confidence [25] (as alluded to in §VI-A1). Therefore, careful consideration into the design and evaluation of transparency mechanisms will likely be of the utmost importance.

It is for this reason that we have emphasised the importance of transparency mechanisms that are *meaningful* for users; such that they directly cater to the needs, requirements, experiences, and levels of expertise of a broad range of people [25]. Towards this, our research has focused on elucidating the transparency mechanisms that *the participants themselves* felt were important, and how such mechanisms could better work to support their aims and interests within a smart home context. In doing so, we present the types of user requirements and design elements that our participants derived, which they thought would better allow them to understand their smart environments and to support them in taking action in response when necessary.

Furthermore, our findings appear to indicate a broader consensus – not only over the importance of meaningful transparency, but also with what designing for meaningful transparency might mean within a smart home context. This consensus was demonstrated through our validation study (Study 3), and was also apparent from our co-design workshops, where we observed considerable overlaps between the outcomes of our two (independent) groups of participants. In this way, our findings may represent a promising starting point for illustrating the importance of meaningful transparency within smart homes, and how users might be better supported in understanding why, and how, these systems operate in the way that they do.

B. Next steps and future research opportunities

Our findings represent a starting point for understanding how transparency mechanisms might better meet the needs and expectations of users. While our categories of requirements and design elements offer tangible ways forward in this regard, our results are intended as indicative; our goal is not to argue that our findings comprise a fully representative set of requirements, design elements, or considerations that might exist in the wider population. Rather, our findings reflect how our participants believed that transparency mechanisms for smart homes might be achieved in practice. In this way, our research provides a foundation for developers and researchers alike to consider, use and build upon. Towards this, we now identify a few areas where future research might be able to assist.

1) Research methods and contexts: Our aim was to document and explore the types of requirements and interface components that our participants thought would assist transparency. Future work could explore the deployment of these outcomes 'in the wild', with actual consumer IoT devices and users. This might entail building upon the requirements and design elements presented here, perhaps exploring how best our outcomes can be translated across the various contexts and scenarios that might arise in any given IoT deployment.

Importantly, however, is that issues of transparency are contextual, and what is needed will often depend on circumstances. Recall that our studies were scenario-led (§IV-A, §V-C); while these scenarios reflected grounded concerns and interests that real people have with the consumer IoT (§III-2), future research may consider exploring different scenarios and contexts. Similarly, while we mainly focused on the context of a 'control panel' for consumer smart homes (§V-C), and thus, outlined elements suitable for such a modality, there are many ways that transparency-related information can be presented and communicated. In all, while these decisions provided some necessary scoping and grounding for our research studies, there are many additional research opportunities, such as those focusing on different scenarios and use cases, different types of transparency mechanisms, system modalities, interaction techniques, and other types of connected environments.

2) Participant samples: There is also scope for considering our research within different cohorts and samples of participants. Recall that we used Mechanical Turk for recruitment to our surveys; while we placed few restrictions on who could take part (§III-1), our sample did appear somewhat limited in representation (e.g. fairly 'techy', with the vast majority having at least one smart device in their home). Similarly, for our co-design workshops, we used a group of computer science undergraduates as a point of comparison. While the use of such a 'convenience sample' allowed us to recruit participants with some knowledge of system requirements and interface prototyping (as a means of ensuring that their responses were grounded within some degree of systems design), important is that we do not intend it to be reflective of the population at large.

While we did verify our findings with an altogether new set of participants (\$VI), follow-up research could explore the extent to which they generalise to new audiences, and how different samples may express different (or similar) characteristics. For example, note that for Study 3, we only considered the design elements that showed consensus across both groups (\$V-C) – omitting, for example, elements such as a 'Settings' menu and a 'Routines' tab for creating simple automated script (see Table IV). This indicates that there is potential for future work that explores other potential design elements, across a range of cohorts.

3) Features and functionality: Our work was user-focused, where participatory methods were used to have the participants themselves develop, determine, and create various transparency-related requirements and interventions. Naturally, there are opportunities for future work to probe on specific features and functionality related to transparency and other related issues. For example, it may be useful to explore particular methods for privacy preservation, including techniques to perturb or obfuscate data, and methods to restrict unintended data access [123]. Similarly, exploring specific means for meaningfully describing the purposes for which smart home data is collected, used and transferred is an area for further exploration. Indeed, this could entail probing or adapting various approaches from literature, such as providing descriptions of apps and devices [10], creating data visualisations for sensor feeds [46], or 'nutrition labels' [47], [48]. While out of scope for this particular paper, exploring how participants might look to influence the design of such approaches might represent promising areas for future research.

More broadly, given the importance of transparency mechanisms (and the oversight, scrutiny, and accountability that they can enable), there is a clear need for research which furthers our understanding of the risks and implications associated with transparency mechanisms. This includes work on ensuring that transparency mechanisms are effective in communicating specific risks and concerns, and importantly that they operate to empower and inform, not mislead or distract. Further, work relating to concerns over malicious design practices (i.e. 'dark patterns') [124], which we see being discussed across a range of technical contexts, also warrants consideration here [125].

4) Applicability beyond the consumer IoT: Lastly, though we have focused on the consumer IoT, our research has broader relevance; transparency and accountability—as they relate to technologies more broadly, beyond that of the IoT are topics of growing importance [23]. This is because transparency will often be a precursor to broader accountability aims, and users seeking a greater understanding of technologies will often be doing so in response to particular issues or concerns. That is to say, there is also much scope for similar research that focuses on transparency mechanisms for different types of technologies (beyond the consumer IoT). Indeed, there are various other technologies facing calls for increased accountability—including AI and algorithmic systems [19], [6], augmented, mixed and virtual reality [126], [127], [128], and cloud services [19], [129], [130], to name a few—and research which works with users to derive meaning-ful transparency mechanisms may have much to offer. Such research might provide new contributions (some of which will transcend across the specific technologies in question), and can therefore help to further our understanding of meaningful transparency mechanisms in general, and how they can work to facilitate greater levels of oversight and understanding.

In all, our findings are but one piece of a much larger puzzle; through working with our participants to design novel transparency mechanisms for smart homes, our work lays the foundations for, and aims to bring more attention to, this nascent research area. There is real potential for work—across various different user groups, scenarios, contexts, and, indeed, technologies—to build upon that which we have found. Furthermore, understanding the similarities and differences across these different samples, contexts, and technologies might help us to build a more comprehensive understanding of how we might realise more meaningful transparency mechanisms. As such, the above represent but a few of the areas where new research has much to offer, amid the growing demands for more transparency surrounding the technologies that are becoming ever more commonplace within our lives.

C. Drivers for change: Encouraging better practice

Questions around the motivations of tech organisations, and drivers for change, are worth considering. While one might suggest that it is not in the tech organisations' interests to facilitate greater scrutiny into their actions, doing so can offer some advantages that organisations may wish to consider going forward. This is because there are growing pressures for increased transparency regarding technologies in general, and organisations themselves can also stand to benefit by being proactive in delivering transparency mechanisms that are more in line with consumer demands. This may be, for example, for reasons of reputation (showing that they take their responsibilities as tech developers/manufacturers/services seriously), or perhaps for reasons of competitive advantage (as we have already seen raised within the context of data protection [131]).

Nevertheless, these pressures for increased transparency within the tech sector go beyond consumer demand alone; transparency demands are also increasingly arising from law (e.g. the EU's GDPR [132]), standards bodies [133], and civil society [134], [135]. Indeed, when asking workshop participants whether they thought the co-designed system would be useful, it was questioned whether governments could do more to "require that [organisations] are open" – an argument that appears to be growing in prominence [75], [136], [57], [82].

In fact, issues relating to transparency and accountability are already prompting regulatory attention, and we do see transparency requirements playing a key role in a number of emerging regulatory regimes, relevant to the IoT and beyond. These include various regulations that are emerging from the European Union, including the GDPR, the AI Act (AI is commonly used in IoT contexts to process data from sensors, enable automation through trigger actuators, and so on [43]), the Internet of Things Policy [137] and proposed Cyber Resilience Act [138]. Other examples include the UK's proposed Consumer IoT Regulation [139] and Code of Practice for Security of IoT [140]. More broadly, governments and legislators appear to be playing an increasingly prominent role in helping to encourage and define better practices (e.g. for transparency and security), and this trend of regulatory intervention looks set to continue.

Again, however, there appears a growing recognition that regulatory obligations for transparency transcend the provision of information or data alone [25]. Within the context of the GDPR, for example, representatives from EU data protection authorities have specifically recognised that "the quality, accessibility and comprehensibility of the information is as important as the actual content of the transparency information" [141]. In other words, it is not enough that such information is 'dumped' on users, and important is that they can effectively understand, engage with, and act upon the information provided through transparency mechanisms. It is therefore apparent that work such as ours-focusing on transparency mechanisms that are *meaningful and usable* for their intended users-could have a significant role to play in helping to shape better practices with regard to how such transparency mechanisms come to be expected.

VIII. CONCLUDING REMARKS

As the IoT continues to proliferate, it is important to ensure that the technology operates in a manner that meets the needs and expectations of its users. Here, transparency plays an important role – by providing information for users about the operation of these systems, transparency mechanisms work to support consumer oversight, understanding, accountability, autonomy, and control. However, simply providing information can often be of limited benefit; rather, there is a need to ensure that transparency mechanisms are more effective in catering to the needs and expectations of their users.

Through a range of user studies, we provide several tangible ways forward on this under-considered topic, bringing users into the design process, and documenting how they perceive that transparency mechanisms can best support their aims. We uncover the types of information that prospective IoT users want to know, how they expect that to be communicated, and what follow-on actions such information might enable. We also present sets of participant-derived requirements and design elements for bringing about a more transparent consumer IoT, finding that these appeared to support meaningful transparency aims with an altogether new group of participants. That is, by giving some insight into how transparency mechanisms can better serve the needs of users within the consumer IoT, our broader aim is to help bring about more transparent and accountable technologies, and call for more attention to be brought into this important area of research.

ACKNOWLEDGMENT

We acknowledge the financial support of UK Research & Innovation (grants EP/P024394/1, EP/R033501/1), The Alan Turing Institute, and Microsoft, through the Microsoft Cloud Computing Research Centre.

REFERENCES

- [1] G. Chalhoub, M. J. Kraemer, N. Nthala, and I. Flechais, "It did not give me an option to decline: A longitudinal analysis of the user experience of security and privacy in smart home products," in Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, ser. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021
- [2] S. Davidoff, M. K. Lee, C. Yiu, J. Zimmerman, and A. K. Dey, "Principles of smart home control," in Proceedings of the 8th International Conference on Ubiquitous Computing, ser. UbiComp'06. Berlin. Heidelberg: Springer-Verlag, 2006, pp. 19-34.
- [3] T. Jakobi, C. Ogonowski, N. Castelli, G. Stevens, and V. Wulf, "The catch(es) with smart home: Experiences of a living lab field study," in Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems. New York, NY, USA: Association for Computing Machinery, May 2017, pp. 1620-1633.
- [4] S. Zheng, N. Apthorpe, M. Chetty, and N. Feamster, "User perceptions of smart home IoT privacy," Proceedings of the ACM on Human-Computer Interaction, vol. 2, no. CSCW, pp. 200:1-200:20, Nov. 2018.
- [5] K. Westcott, J. Loucks, D. Littmann, P. Wilson, S. Srivastava, and D. Ciampa, "Build it and they will embrace it: Consumers are preparing for 5G connectivity in the home and on the https://www2.deloitte.com/content/dam/insights/us/articles/6457_ go." Mobile-trends-survey/DI_Build-it-and-they-will-embrace-it.pdf, June 2021, accessed: 2021-11-15.
- [6] J. Singh, J. Cobbe, and C. Norval, "Decision provenance: Harnessing data flow for accountable systems," IEEE Access, vol. 7, pp. 6562-6574, 2019.
- [7] T. Jakobi, G. Stevens, N. Castelli, C. Ogonowski, F. Schaub, N. Vindice, D. Randall, P. Tolmie, and V. Wulf, "Evolving needs in IoT control and accountability: A longitudinal study on smart home intelligibility," Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, vol. 2, no. 4, pp. 171:1-171:28, Dec. 2018.
- [8] M. Tabassum, T. Kosinski, and H. R. Lipford, "I don't own the data: End user perceptions of smart home device data practices and risks," in Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019). Santa Clara, CA: USENIX Association, Aug. 2019, pp. 435-450.
- [9] C. Norval, J. Cobbe, and J. Singh, "Towards an accountable Internet of Things: A call for 'reviewability'," in Privacy by Design for the Internet of Things: Building accountability and security. The Institution of Engineering Technology, 2021.
- [10] A. Crabtree, T. Lodge, J. Colley, C. Greenhalgh, K. Glover, H. Haddadi, Y. Amar, R. Mortier, Q. Li, J. Moore et al., "Building accountability into the Internet of Things: the IoT Databox model," Journal of Reliable Intelligent Environments, vol. 4, no. 1, pp. 39-55, 2018.
- [11] A. Carman, "Smart ovens have been turning on overnight and preheating to 400 degrees," https://www.theverge.com/2019/8/14/20802774/ june-smart-oven-remote-preheat-update-user-error, August 2019, accessed: 2021-11-15.
- [12] R. Trimananda, S. A. H. Aqajari, J. Chuang, B. Demsky, G. H. Xu, and S. Lu, "Understanding and automatically detecting conflicting interactions between smart home IoT applications," in Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ser. ESEC/FSE 2020. New York, NY, USA: Association for Computing Machinery, 2020, pp. 1215-1227.
- [13] A. Acquisti, I. Adjerid, and L. Brandimarte, "Gone in 15 seconds: The limits of privacy transparency and control," IEEE Security & Privacy, vol. 11, no. 4, pp. 72-74, 2013.
- [14] M. Bovens, "Analysing and assessing accountability: A conceptual framework1," European Law Journal, vol. 13, no. 4, pp. 447-468, 2007.
- [15] J. A. Obar, "Sunlight alone is not a disinfectant: Consent and the futility of opening big data black boxes (without assistance)," Big Data & Society, vol. 7, no. 1, pp. 1-5, 2020.
- [16] C. Stohl, M. Stohl, and P. M. Leonardi, "Managing opacity: Information visibility and the paradox of transparency in the digital age," IJoC, vol. 10, no. 2016, pp. 123-137, 2016.
- [17] N. Suzor, S. West, A. Quodling, and J. York, "What do we mean when we talk about transparency? Toward meaningful transparency in commercial content moderation," International Journal of Communication, vol. 13, no. 0, 2019.
- [18] D. Kamarinou, C. Millard, and J. Singh, "Machine learning with personal data," in Data protection and privacy: The age of intelligent machines. Hart Publishing, 2017.

- [19] J. Cobbe, M. S. A. Lee, and J. Singh, "Reviewable automated decisionmaking: A framework for accountable algorithmic systems," in Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, ser. FAccT '21. New York, NY, USA: Association for Computing Machinery, 2021, pp. 598-609. M. Kaminski and J. M. Urban, "The right to contest AI," Columbia
- [20] Law Review, vol. 121, no. 7, 2021.
- [21] J. Pridmore and A. Mols, "Personal choices and situated data: Privacy negotiations and the acceptance of household intelligent personal assistants," Big Data & Society, vol. 7, no. 1, 2020.
- [22] J. Ausloos and M. Veale, "Researching with data rights," Technology and Regulation, vol. 2020, pp. 136-157, Jan. 2021.
- F. Pasquale, The Black Box Society: The secret algorithms that control [23] money and information. Harvard University Press, 2015
- [24] K. Holtzblatt and H. R. Beyer, "Requirements gathering: The human factor," Commun. ACM, vol. 38, no. 5, pp. 31-32, May 1995.
- [25] C. Norval, K. Cornelius, J. Cobbe, and J. Singh, "Disclosure by design: Designing information disclosures to support meaningful transparency and accountability," in 2022 ACM Conference on Fairness, Accountability, and Transparency, ser. FAccT '22. New York, NY, USA: Association for Computing Machinery, 2022, pp. 679-690.
- [26] D. Miorandi, S. Sicari, F. De Pellegrini, and I. Chlamtac, "Internet of Things: Vision, applications and research challenges," Ad Hoc Networks, vol. 10, no. 7, pp. 1497-1516, 2012.
- [27] J. Singh, T. Pasquier, J. Bacon, H. Ko, and D. Eyers, "Twenty security considerations for Cloud-supported Internet of Things," IEEE Internet of Things Journal, vol. 3, no. 3, pp. 269-284, 2016.
- [28] A. B. Brush, B. Lee, R. Mahajan, S. Agarwal, S. Saroiu, and C. Dixon, "Home automation in the wild: Challenges and opportunities," in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. New York, NY, USA: Association for Computing Machinery, May 2011, pp. 2115–2124.
- [29] S. Yarosh and P. Zave, "Locked or not? Mental models of IoT feature interaction," in Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, ser. CHI '17. New York, NY, USA: Association for Computing Machinery, May 2017, pp. 2993-2997.
- [30] D. Roca, D. Nemirovsky, M. Nemirovsky, R. Milito, and M. Valero, "Emergent behaviors in the Internet of Things: The ultimate ultra-largescale system," IEEE Micro, vol. 36, no. 6, pp. 36-44, 2016.
- [31] W. He, J. Martinez, R. Padhi, L. Zhang, and B. Ur, "When smart devices are stupid: Negative experiences using home smart devices," in 2019 IEEE Security and Privacy Workshops (SPW), 2019, pp. 150-155
- [32] M. Palekar, E. Fernandes, and F. Roesner, "Analysis of the susceptibility of smart home programming interfaces to end user error," in 2019 IEEE Security and Privacy Workshops (SPW), 2019, pp. 138-143.
- [33] L. Xing, "Cascading failures in Internet of Things: Review and perspectives on reliability and resilience," IEEE Internet of Things Journal, vol. 8, no. 1, pp. 44-64, 2021.
- [34] C. Norval and J. Singh, "Explaining automated environments: Interrogating scripts, logs, and provenance using voice-assistants," in Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers, ser. UbiComp/ISWC '19 Adjunct, 2019, pp. 332-335.
- [35] E. Park, Y. Cho, J. Han, and S. J. Kwon, "Comprehensive approaches to user acceptance of Internet of Things in a smart home environment,"
- *IEEE Internet of Things Journal*, vol. 4, no. 6, pp. 2342–2350, 2017. E.-M. Schomakers, H. Biermann, and M. Ziefle, "Users' preferences [36] for smart home automation - Investigating aspects of privacy and trust," Telematics and Informatics, vol. 64, 2021.
- [37] M. Yurrita, T. Draws, A. Balayn, D. Murray-Rust, N. Tintarev, and A. Bozzon, "Disentangling fairness perceptions in algorithmic decisionmaking: The effects of explanations, human oversight, and contestability," in Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, ser. CHI '23. New York, NY, USA: Association for Computing Machinery, 2023.
- [38] D. Marikyan, S. Papagiannidis, and E. Alamanos, "A systematic review of the smart home literature: A user perspective," Technological Forecasting and Social Change, vol. 138, pp. 139-154, 2019.
- D. Mocrii, Y. Chen, and P. Musilek, "IoT-based smart homes: A [39] review of system architecture, software, communications, privacy and security," Internet of Things, vol. 1-2, pp. 81-98, 2018.
- [40] R. Despouys, R. Sharrock, and I. Demeure, "Sensemaking in the autonomic smart-home," in Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication. Seattle Washington: ACM, Sep. 2014, pp. 887-894.

- [41] Y. Chuang, L.-L. Chen, and Y. Liu, "Design vocabulary for human–IoT systems communication," in *Proceedings of the 2018 CHI Conference* on Human Factors in Computing Systems. New York, NY, USA: Association for Computing Machinery, Apr. 2018, pp. 1–11.
- [42] J. Dai, C. Zhang, D. Aliakseyeu, S. Peeters, and W. A. Ijsselsteijn, "The effect of explanation design on user perception of smart home lighting systems: A mixed-method investigation," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, ser. CHI '23. New York, NY, USA: Association for Computing Machinery, 2023.
- [43] İ. Kök, F. Y. Okay, Ö. Muyanlı, and S. Özdemir, "Explainable artificial intelligence (XAI) for Internet of Things: A survey," *IEEE Internet of Things Journal*, vol. 10, no. 16, pp. 14764–14779, 2023.
- [44] A. Desjardins, H. R. Biggs, C. Key, and J. E. Viny, "IoT data in the home: Observing entanglements and drawing new encounters," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, Apr. 2020, pp. 1–13.
- [45] A. Desjardins and H. R. Biggs, "Data epics: Embarking on literary journeys of home Internet of Things data," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, ser. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021.
- [46] N. Castelli, C. Ogonowski, T. Jakobi, M. Stein, G. Stevens, and V. Wulf, "What happened in my home? An end-user development approach for smart home data visualization," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, May 2017, pp. 853–866.
- [47] A. Railean and D. Reinhardt, "Let there be LITE: Design and evaluation of a label for IoT transparency enhancement," in *Proceedings* of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct, ser. MobileHCI '18. New York, NY, USA: Association for Computing Machinery, Sep. 2018, pp. 103–110.
- [48] A. Railean and D. Reinhardt, "OnLITE: On-line label for IoT transparency enhancement," in *Secure IT Systems*, ser. Lecture Notes in Computer Science. Springer International Publishing, 2021, pp. 229– 245.
- [49] M. Manca, F. Paternò, C. Santoro, and L. Corcella, "Supporting enduser debugging of trigger-action rules for IoT applications," *International Journal of Human-Computer Studies*, vol. 123, pp. 56–69, 2019.
- [50] V. Zhao, L. Zhang, B. Wang, S. Lu, and B. Ur, "Visualizing differences to improve end-user understanding of trigger-action programs," in *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, ser. CHI EA '20. New York, NY, USA: Association for Computing Machinery, 2020, pp. 1—10.
- [51] V. Zhao, L. Zhang, B. Wang, M. L. Littman, S. Lu, and B. Ur, Understanding Trigger-Action Programs Through Novel Visualizations of Program Differences. New York, NY, USA: Association for Computing Machinery, 2021.
- [52] C. Castelluccia, M. Cunche, D. Le Metayer, and V. Morel, "Enhancing transparency and consent in the IoT," in 2018 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW), Apr. 2018, pp. 116–119.
- [53] C. Abras, D. Maloney-Krichmar, J. Preece et al., "User-centered design," Bainbridge, W. Encyclopedia of Human-Computer Interaction. Thousand Oaks: Sage Publications, vol. 37, no. 4, pp. 445–456, 2004.
- [54] Y. Yao, L. Huang, Y. He, Z. Ma, X. Xu, and H. Mi, "Reviewing and reflecting on smart home research from the human-centered perspective," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, ser. CHI '23. New York, NY, USA: Association for Computing Machinery, 2023.
- [55] P. Emami-Naeini, H. Dixon, Y. Agarwal, and L. F. Cranor, "Exploring how privacy and security factor into IoT device purchase behavior," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, May 2019, pp. 1–12.
- [56] J. M. Haney, S. M. Furman, and Y. Acar, "Smart home security and privacy mitigations: Consumer perceptions, practices, and challenges," in *HCI for Cybersecurity, Privacy and Trust*, ser. Lecture Notes in Computer Science, A. Moallem, Ed. Cham: Springer International Publishing, 2020, pp. 393–411.
- [57] J. Haney, Y. Acar, and S. Furman, "It's the company, the government, you and I: User perceptions of responsibility for smart home privacy and security," in *30th USENIX Security Symposium (USENIX Security* 21). USENIX Association, Aug. 2021, pp. 411–428.

- [58] O. Kulyk, K. Milanovic, and J. Pitt, "Does my smart device provider care about my privacy? Investigating trust factors and user attitudes in IoT systems," in *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*, ser. NordiCHI '20. New York, NY, USA: Association for Computing Machinery, Oct. 2020, pp. 1–12.
- [59] K. Marky, S. Prange, F. Krell, M. Mühlhäuser, and F. Alt, "You just can't know about everything: Privacy perceptions of smart home visitors," in *19th International Conference on Mobile and Ubiquitous Multimedia*, ser. MUM 2020. New York, NY, USA: Association for Computing Machinery, 2020, pp. 83—95.
- [60] M. Williams, J. R. C. Nurse, and S. Creese, "Privacy is the boring bit: User perceptions and behaviour in the Internet-of-Things," in 2017 15th Annual Conference on Privacy, Security and Trust (PST), 2017, pp. 181–18109.
- [61] M. Windl, V. Winterhalter, A. Schmidt, and S. Mayer, "Understanding and mitigating technology-facilitated privacy violations in the physical world," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, ser. CHI '23. New York, NY, USA: Association for Computing Machinery, 2023.
- [62] N. Abdi, K. M. Ramokapane, and J. M. Such, "More than smart speakers: Security and privacy perceptions of smart home personal assistants," in *Fifteenth Symposium on Usable Privacy and Security* (SOUPS 2019). Santa Clara, CA: USENIX Association, Aug. 2019, pp. 451–466.
- [63] E. Zeng, S. Mare, and F. Roesner, "End user security and privacy concerns with smart homes," in *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. Santa Clara, CA: USENIX Association, Jul. 2017, pp. 65–80.
- [64] A. I. Hudig, C. Norval, and J. Singh, "Transparency in the consumer Internet of Things: Data flows and data rights," http://iot-transparency. org/, 2023, accessed: 2023-08-09.
- [65] A. M. Mandalari, D. J. Dubois, R. Kolcun, M. T. Paracha, H. Haddadi, and D. Choffnes, "Blocking without breaking: Identification and mitigation of non-essential IoT traffic," in *Privacy Enhancing Technologies Symposium (PETS)*, 2021.
- [66] J. Ren, D. J. Dubois, D. Choffnes, A. M. Mandalari, R. Kolcun, and H. Haddadi, "Information exposure from consumer IoT devices: A multidimensional, network-informed measurement approach," in *Proceedings of the Internet Measurement Conference*, ser. IMC '19. New York, NY, USA: Association for Computing Machinery, 2019, pp. 267—279.
- [67] N. Apthorpe, Y. Shvartzshnaider, A. Mathur, D. Reisman, and N. Feamster, "Discovering smart home Internet of Things privacy norms using contextual integrity," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 2, no. 2, jul 2018.
- [68] Y. Liao, J. Vitak, P. Kumar, M. Zimmer, and K. Kritikos, "Understanding the role of privacy and trust in intelligent personal assistant adoption," in *Information in Contemporary Society*, ser. Lecture Notes in Computer Science, N. G. Taylor, C. Christian-Lamb, M. H. Martin, and B. Nardi, Eds. Cham: Springer International Publishing, 2019, pp. 102–113.
- [69] E. Lafontaine, A. Sabir, and A. Das, "Understanding people's attitude and concerns towards adopting IoT devices," in *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, no. 307. New York, NY, USA: Association for Computing Machinery, 2021, pp. 1–10.
- [70] S.-C. Cha, T.-Y. Hsu, Y. Xiang, and K.-H. Yeh, "Privacy enhancing technologies in the Internet of Things: Perspectives and challenges," *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 2159–2187, 2019.
- [71] C. Li and B. Palanisamy, "Privacy in Internet of Things: From principles to technologies," *IEEE Internet of Things Journal*, vol. 6, no. 1, pp. 488–505, 2019.
- [72] N. Abdi, X. Zhan, K. M. Ramokapane, and J. Such, "Privacy norms for smart home personal assistants," in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, no. 558. New York, NY, USA: Association for Computing Machinery, May 2021, pp. 1–14.
- [73] Y. Yao, J. R. Basdeo, S. Kaushik, and Y. Wang, "Defending my castle: A co-design study of privacy mechanisms for smart homes," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, ser. CHI '19. New York, NY, USA: Association for Computing Machinery, May 2019, pp. 1–12.
- [74] Y. Yao, "Designing for better privacy awareness in smart homes," in Conference Companion Publication of the 2019 on Computer Supported Cooperative Work and Social Computing, ser. CSCW '19.

New York, NY, USA: Association for Computing Machinery, 2019, pp. 98—101.

- [75] G. Chalhoub, I. Flechais, N. Nthala, R. Abu-Salma, and E. Tom, "Factoring user experience into the security and privacy design of smart home devices: A case study," in *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, ser. CHI EA '20. New York, NY, USA: Association for Computing Machinery, 2020, pp. 1—-9.
- [76] C. Chhetri and V. G. Motti, "Eliciting privacy concerns for smart home devices from a user centered perspective," in *Information in Contemporary Society*, ser. Lecture Notes in Computer Science, N. G. Taylor, C. Christian-Lamb, M. H. Martin, and B. Nardi, Eds. Cham: Springer International Publishing, 2019, pp. 91–101.
- [77] E. Zeng and F. Roesner, "Understanding and improving security and privacy in multi-user smart homes: A design exploration and in-home user study," in 28th USENIX Security Symposium (USENIX Security 19). Santa Clara, CA: USENIX Association, Aug. 2019, pp. 159– 176.
- [78] H. Jin, B. Guo, R. Roychoudhury, Y. Yao, S. Kumar, Y. Agarwal, and J. I. Hong, "Exploring the needs of users for supporting privacyprotective behaviors in smart homes," in *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, ser. CHI '22. New York, NY, USA: Association for Computing Machinery, Apr. 2022, pp. 1–19.
- [79] C. Chhetri and V. Motti, "Designing and evaluating a prototype for data-related privacy controls in a smart home," in *Human Aspects of Information Security and Assurance*, ser. IFIP Advances in Information and Communication Technology, N. Clarke and S. Furnell, Eds. Cham: Springer International Publishing, 2022, pp. 240–250.
- [80] W. Seymour, M. J. Kraemer, R. Binns, and M. Van Kleek, "Informing the design of privacy-empowering tools for the connected home," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, Apr. 2020, pp. 1–14.
- [81] K. Marky, V. Zimmermann, A. Stöver, P. Hoffmann, K. Kunze, and M. Mühlhäuser, "All in one! User perceptions on centralized IoT privacy settings," in *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, ser. CHI EA '20. New York, NY, USA: Association for Computing Machinery, 2020, pp. 1—-8.
- [82] L. Urquhart and J. Chen, "On the principle of accountability: Challenges for smart homes & cybersecurity," *CoRR*, vol. abs/2006.11043, 2020.
- [83] S. Wachter, "The GDPR and the Internet of Things: A three-step transparency model," *Law, Innovation and Technology*, vol. 10, no. 2, pp. 266–294, Jul. 2018.
- [84] N. Komninos, E. Philippou, and A. Pitsillides, "Survey in smart grid and smart home security: Issues, challenges and countermeasures," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 4, pp. 1933– 1954, 2014.
- [85] F. Meneghello, M. Calore, D. Zucchetto, M. Polese, and A. Zanella, "IoT: Internet of threats? A survey of practical security vulnerabilities in real IoT devices," *IEEE Internet of Things Journal*, vol. 6, no. 5, pp. 8182–8201, 2019.
- [86] B. K. Mohanta, D. Jena, S. Ramasubbareddy, M. Daneshmand, and A. H. Gandomi, "Addressing security and privacy issues of IoT using blockchain technology," *IEEE Internet of Things Journal*, vol. 8, no. 2, pp. 881–888, 2021.
- [87] A. Sabir, E. Lafontaine, and A. Das, "Hey Alexa, who am I talking to?: Analyzing users' perception and awareness regarding third-party Alexa skills," in *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, ser. CHI '22. New York, NY, USA: Association for Computing Machinery, 2022.
- [88] K. Zhao and L. Ge, "A survey on the Internet of Things security," in 2013 Ninth International Conference on Computational Intelligence and Security, 2013, pp. 663–667.
- [89] G. Chalhoub, "The UX of Things: Exploring UX principles to inform security and privacy design in the smart home," in *Extended Abstracts* of the 2020 CHI Conference on Human Factors in Computing Systems, ser. CHI EA '20. New York, NY, USA: Association for Computing Machinery, Apr. 2020, pp. 1–6.
- [90] Y. H. Hwang, "IoT security & privacy: Threats and challenges," in Proceedings of the 1st ACM Workshop on IoT Privacy, Trust, and Security, ser. IoTPTS '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 1.
- [91] C. Geeng and F. Roesner, "Who's in control? Interactions in multi-user smart homes," in *Proceedings of the 2019 CHI Conference on Human*

Factors in Computing Systems. New York, NY, USA: Association for Computing Machinery, May 2019, pp. 1–13.

- [92] R. Duezguen, P. Mayer, B. Berens, C. Beckmann, L. Aldag, M. Mossano, M. Volkamer, and T. Strufe, "How to increase smart home security and privacy risk perception," in 2021 IEEE 20th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom), Oct. 2021, pp. 997–1004, iSSN: 2324-9013.
- [93] T. Pasquier, J. Singh, J. Powles, D. Eyers, M. Seltzer, and J. Bacon, "Data provenance to audit compliance with privacy policy in the Internet of Things," *Personal Ubiquitous Comput.*, vol. 22, no. 2, pp. 333—344, apr 2018.
- [94] X. Han, T. Pasquier, A. Bates, J. Mickens, and M. Seltzer, "UNICORN: Runtime provenance-based detector for advanced persistent threats," in *Network and Distributed System Security Symposium (NDSS'20)*. Internet Society, 2020.
- [95] J. Wang, S. Hao, R. Wen, B. Zhang, L. Zhang, H. Hu, and R. Lu, "IoT-Praetor: Undesired behaviors detection for IoT devices," *IEEE Internet* of Things Journal, vol. 8, no. 2, pp. 927–940, 2021.
- [96] T. F. J.-M. Pasquier, J. Singh, D. Eyers, and J. Bacon, "Camflow: Managed data-sharing for Cloud services," *IEEE Transactions on Cloud Computing*, vol. 5, no. 3, pp. 472–484, 2017.
- [97] J. Singh, C. Millard, C. Reed, J. Cobbe, and J. Crowcroft, "Accountability in the IoT: Systems, law, and ways forward," *Computer*, vol. 51, no. 7, pp. 54–65, 2018.
- [98] P. Emami-Naeini, Y. Agarwal, L. Faith Cranor, and H. Hibshi, "Ask the experts: What should be on an IoT privacy and security label?" in 2020 IEEE Symposium on Security and Privacy (SP), May 2020, pp. 447–464, iSSN: 2375-1207.
- [99] Pew Research Center, "Research in the crowdsourcing age, a case study," https://www.pewresearch.org/internet/2016/07/11/ research-in-the-crowdsourcing-age-a-case-study/, July 2016, accessed: 2022-11-14.
- [100] F. M. Shipman and C. C. Marshall, "Ownership, privacy, and control in the wake of cambridge analytica: The relationship between attitudes and awareness," in *CHI* '20, 2020, pp. 1—12.
- [101] J. Jager, D. L. Putnick, and M. H. Bornstein, "More than just convenient: The scientific merits of homogeneous convenience samples," *Monographs of the Society for Research in Child Development*, vol. 82, no. 2, pp. 13–30, 2017.
- [102] G. Chalhoub and I. Flechais, "Alexa, are you spying on me?: Exploring the effect of user experience on the security and privacy of smart speaker users," in *HCI for Cybersecurity, Privacy and Trust*, A. Moallem, Ed. Cham: Springer International Publishing, 2020, pp. 305–325.
- [103] D. Winder, "How to stop your smart home spying on you," https://www.theguardian.com/technology/2020/mar/08/ how-to-stop-your-smart-home-spying-on-you-lightbulbs-doorbell-ring-google-assista March 2020, accessed: 2022-08-14.
- [104] R. Yus and P. Pappachan, "Smart devices spy on you 2 computer scientists explain how the Internet of Things can violate your privacy," https://theconversation.com/ smart-devices-spy-on-you-2-computer-scientists-explain-how-the-internet-of-things-March 2022, accessed: 2022-08-14.
- [105] C. Norval and J. Singh, "Supplementary data A room with an overview: Towards meaningful transparency for the consumer internet of things," https://github.com/cnorval/meaningful_IoT, 2023, accessed: 2023-09-18.
- [106] H. S. Alavi, E. F. Churchill, M. Wiberg, D. Lalanne, P. Dalsgaard, A. Fatah gen Schieck, and Y. Rogers, "Introduction to Human-Building Interaction (HBI): Interfacing HCI with architecture and urban design," *ACM Trans. Comput.-Hum. Interact.*, vol. 26, no. 2, mar 2019.
- [107] D. Bastos, F. Giubilo, M. Shackleton, and F. El-Moussa, "GDPR privacy implications for the Internet of Things," in *4th Annual IoT Security Foundation Conference*, vol. 4, 2018, pp. 1–8.
- [108] V. Braun and V. Clarke, "Using thematic analysis in psychology," *QRP*, vol. 3, no. 2, pp. 77–101, 2006.
- [109] M. H. Bornstein, J. Jager, and D. L. Putnick, "Sampling in developmental science: Situations, shortcomings, solutions, and standards," *Developmental Review*, vol. 33, no. 4, pp. 357–370, 2013.
- [110] P. Sedgwick, "Convenience sampling," BMJ, vol. 347, 2013.
- [111] S. J. Stratton, "Population research: Convenience sampling strategies," *Prehospital and Disaster Medicine*, vol. 36, no. 4, pp. 373–374, 2021.
- [112] MURAL, "MURAL," https://www.mural.co/, January 2022, accessed: 2022-01-21.
- [113] J. Smith and H. Noble, "Bias in research," *Evidence-Based Nursing*, vol. 17, no. 4, pp. 100–101, 2014.

- [114] M. A. A. Elsood, H. A. Hefny, and E. S. Nasr, "A goal-based technique for requirements prioritization," in 2014 9th International Conference on Informatics and Systems, 2014, pp. SW–18–SW–24.
- [115] A. Cuthbertson, "Hackers can hijack your house through your light bulb, researchers discover," https://www.independent.co.uk/ tech/philips-hue-smart-light-bulb-hack-cyber-security-a9317456.html, February 2020, accessed: 2022-08-13.
- [116] E. Ronen and A. Shamir, "Extended functionality attacks on IoT devices: The case of smart lights," in 2016 IEEE European Symposium on Security and Privacy (EuroS&P), 2016, pp. 3–12.
- [117] E. Ronen, A. Shamir, A.-O. Weingarten, and C. O'Flynn, "IoT goes nuclear: Creating a ZigBee chain reaction," in 2017 IEEE Symposium on Security and Privacy (SP), 2017, pp. 195–212.
- [118] E. McGowan, "Here's what your Ring doorbell knows about you," https://blog.avast.com/what-amazon-ring-knows-about-you-avast, May 2021, accessed: 2022-08-14.
- [119] K. N. Truong, G. R. Hayes, and G. D. Abowd, "Storyboarding: An empirical determination of best practices and effective guidelines," in *Proceedings of the 6th Conference on Designing Interactive Systems*, ser. DIS '06. New York, NY, USA: Association for Computing Machinery, 2006, pp. 12—21.
- [120] W. D. Perreault, "Controlling order-effect bias," *The Public Opinion Quarterly*, vol. 39, no. 4, pp. 544–551, 1975.
- [121] J. Brooke, "SUS: A quick and dirty usability scale," in Usability evaluation in industry. Taylor and Francis, 1996.
- [122] A. Bangor, P. Kortum, and J. Miller, "Determining what individual SUS scores mean: Adding an adjective rating scale," *JUS*, vol. 4, no. 3, pp. 114–123, 2009.
- [123] F. Loukil, C. Ghedira-Guegan, A. N. Benharkat, K. Boukadi, and Z. Maamar, "Privacy-aware in the IoT applications: A systematic literature review," in *On the Move to Meaningful Internet Systems*. *OTM 2017 Conferences*, H. Panetto, C. Debruyne, W. Gaaloul, M. Papazoglou, A. Paschke, C. A. Ardagna, and R. Meersman, Eds. Cham: Springer International Publishing, 2017, pp. 552–569.
- [124] C. M. Gray, Y. Kou, B. Battles, J. Hoggatt, and A. L. Toombs, "The dark (patterns) side of UX design," in *CHI '18*, 2018.
- [125] M. Kowalczyk, J. T. Gunawan, D. Choffnes, D. J. Dubois, W. Hartzog, and C. Wilson, "Understanding dark patterns in home IoT devices," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, ser. CHI '23. New York, NY, USA: Association for Computing Machinery, 2023.
- [126] R. Cloete, C. Norval, and J. Singh, "A call for auditable virtual, augmented and mixed reality," in *Proceedings of the 26th ACM Symposium on Virtual Reality Software and Technology*, ser. VRST '20. New York, NY, USA: Association for Computing Machinery, 2020.
- [127] R. Cloete, C. Norval, and J. Singh, "Auditable augmented/mixed/virtual reality: The practicalities of mobile system transparency," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 5, no. 4, dec 2022.
- [128] C. Norval, R. Cloete, and J. Singh, "Navigating the audit landscape: A framework for developing transparent and auditable XR," in *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, ser. FAccT '23. New York, NY, USA: Association for Computing Machinery, 2023, pp. 1418—1431.
- [129] S. A. Javadi, C. Norval, R. Cloete, and J. Singh, "Monitoring AI services for misuse," in *Proceedings of the 2021 AAAI/ACM Conference* on AI, Ethics, and Society, ser. AIES '21. New York, NY, USA: Association for Computing Machinery, 2021, pp. 597–607.
- [130] C. J. Millard, *Cloud computing law*, 2nd ed. Oxford University Press Oxford, 2021.
- [131] C. Norval, H. Janssen, J. Cobbe, and J. Singh, "Data protection and tech startups: The need for attention, support, and scrutiny," *Policy & Internet*, vol. 13, no. 2, pp. 278–299, 2021.
- [132] European Union, "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation)," *Official Journal of the European Union*, vol. L119, pp. 1–88, May 2016.
- [133] A. F. T. Winfield, S. Booth, L. A. Dennis, T. Egawa, H. Hastie, N. Jacobs, R. I. Muttram, J. I. Olszewska, F. Rajabiyazdi, A. Theodorou, M. A. Underwood, R. H. Wortham, and E. Watson, "IEEE P7001: A proposed standard on transparency," *Frontiers in Robotics and AI*, vol. 8, 2021.
- [134] Electronic Frontier Foundation, "Electronic Frontier Foundation," https: //www.eff.org/, January 2022, accessed: 2022-01-20.

- [135] Open Rights Group, "Open Rights Group," https://www. openrightsgroup.org/, January 2022, accessed: 2022-01-20.
- [136] J. Chen and L. Urquhart, "They're all about pushing the products and shiny things rather than fundamental security: Mapping sociotechnical challenges in securing the smart home," *arXiv preprint arXiv:2105.11751*, 2021.
- [137] European Commission, "Europe's Internet of Things Policy," https:// digital-strategy.ec.europa.eu/en/policies/internet-things-policy, October 2022, accessed: 2022-12-19.
- [138] European Commission, "Cyber Resilience Act," https: //digital-strategy.ec.europa.eu/en/library/cyber-resilience-act, September 2022, accessed: 2022-12-19.
- [139] GOV.UK, "New cyber security laws to protect smart devices amid pandemic sales surge," https://www.gov.uk/government/news/ new-cyber-security-laws-to-protect-smart-devices-amid-pandemic-sales-surge, April 2021, accessed: 2021-11-15.
- [140] GOV.UK, "Code of practice for consumer IoT security," https://www.gov.uk/government/publications/ code-of-practice-for-consumer-iot-security, October 2018, accessed: 2022-12-19.
- [141] Article 29 Working Party, "Guidelines on transparency under Regulation 2016/679," no. WP260, April 2018.