Prototyping an IoT Transparency Toolkit to support Communication, Governance and Policy in the Smart City

Naomi Jacobs*[1], Fi Loveday[2], Milan Markovic[2], Caitlin D Cottrill[3], Raffaele Zullo[2], Peter Edwards[2]

[1]Imagination Lancaster, LICA, Lancaster University, Lancaster, UK;

[2]School of Natural and Computing Sciences, University of Aberdeen, Aberdeen, UK;

[3]School of Engineering, University of Aberdeen, Aberdeen, UK;

*Corresponding author. Naomi.jacobs@lancaster.ac.uk

Author ORCIDs:

Naomi Jacobs: 0000-0003-2079-1658

Caitlin D Cottrill: 0000-0002-1638-8113

Milan Markovic: 0000-0002-5477-287X

Peter Edwards: 0000-0002-4527-9186

Author biographies

Naomi Jacobs is Lecturer in Design Policy and Futures Thinking at Lancaster University, researching technology and society, and the nature of digital public spaces.

Fi Loveday is a freelance artist and designer. Her research explores the use of digital technology within the visual arts as a vehicle for developing connection and interaction between viewer and artist.

Milan Markovic is a Research Fellow in Computing Science at the University of Aberdeen. His research focuses on enhancing transparency and accountability in

complex socio-technical systems through intelligent processes based on provenance data models and semantic web technologies.

Caitlin D Cottrill is a Senior Lecturer in the School of Engineering at the University of Aberdeen. Her research explores the use of technology in urban spaces, with particular focus on transport and data privacy.

Raffaele Zullo is a Researcher in Computer Science and Engineering.
His research interests include network measurement and protocol analysis with a special focus on the impact of novel protocols and technologies on network security.

Peter Edwards is a Professor of Computing Science at the University of Aberdeen. His research interests encompass provenance, trust and information quality, with a particular emphasis on the role of semantic web technologies in helping realise transparent and accountable systems.

Prototyping an IoT Transparency Toolkit to support Communication, Governance and Policy in the Smart City

Abstract

This paper describes the design of a prototype transparency tool for communities, public and private organisations, and other stakeholders to use when developing, deploying and interrogating public space Internet of Things implementations as part of wider smart city initiatives. These deployments need to be accompanied by effective public communications and policy, and consideration of factors such as data privacy, so as not to negatively impact citizens. The intention of the tool is to guide transparency of these deployments in order to support trusted IoT ecosystems. The tool, which was developed as both a physical card deck and a digital implementation, was developed through a combination of participatory work and other research to understand the information requirements of citizens and other stakeholders. We suggest that further development of such tools can effectively support trustworthy IoT in smart cities.

Keywords: internet of things; transparency; governance; participatory design; smart cities

Introduction

Since the late 1990s, the concept of the 'smart city' has grown to increasing prominence worldwide (Suzuki and Finkelstein, 2019). Many smart city projects, both successful and unsuccessful, incorporate significant use of ubiquitous connected devices: the Internet of Things (IoT). In practice, the push to rapidly incorporate these technologies into urban spaces may result only in local pilot interventions not yet at full market-roll out stage (Kitchin et al, 2017). These can be implemented rapidly, with minimal time for consultation or associated policy design, in order to be 'ahead of the crowd' or even just to keep up with the perceived successes of comparable regions. Connected devices and systems are often introduced within cities in conjunction with non-governmental actors from grass-roots citizen led projects to corporate, commercial ventures. Such programmes can be wide ranging in scope, scale and intentions. Cardullo and Kitchin (2018) suggest that wider neoliberal structures influence smart city participation, limiting active citizen engagement, and leading localities to compete against each other in order to attract supra-national investments. If implications in areas such as privacy, transparency and trust are not fully considered, there can be serious consequences (Read et al, 2016).

Design research has an important contribution to make when proposing new distributed technological solutions in the urban space. Immediate and future ramifications must be considered, as well as trust factors, particularly when the collection of data about, and from, the public is involved. Design approaches are particularly suited to addressing such complexities in an inclusive and thoughtful manner, and are increasingly being used in policy contexts (Whicher, 2021). While the rhetoric of policy implementations

and decision making is often that of empowerment and serving the public good, the totality of social actors may not be fully considered in complex technology-centric projects which are led by techno-solutionism and presumptions of benefit (Cardullo and Kitchin, 2018).

The research presented in this paper was undertaken as part of an interdisciplinary project with the aim of enabling trusted IoT ecosystems. A key component of trust is transparency (Castelluccia et al, 2018). For many reasons, IoT deployments may lack the transparency needed for an interested party to obtain information about the system. This work explores the desires and needs of citizens with regard to public space IoT transparency, and describes the design of prototype transparency mechanisms that capture appropriate information. The objective is to support transparency of IoT deployments in public spaces throughout their development and deployment, by exposing information that may *not* initially be transparent, enabling clearer decision making about transparency and how it might be implemented to increase trust. The intention is that by using such tools, an individual framework of recommendations and guidelines for increased transparency can be produced for each unique deployment. With this in place, privacy and accountability concerns become visible at early stages of deployments where they can be mitigated. Consequently, a stronger trust relationship can be facilitated between those implementing such deployments, and those impacted by them. In the first section of this paper, we set out the context of this research and give an overview of related work. In the second, we describe how an index of transparency questions was sourced through a combination of literature reviews and participatory activities. Following this, we describe the design process undertaken to develop prototype tools which use these questions to provide guidance on transparency. We present initial results from a limited evaluation of the

tool, before concluding with a discussion including next steps for this work. Through this research, we ask whether collating key deployment information and exposing potential transparency limitations can enable greater transparency, greater trust, and active citizen engagement in public space IoT.

Context

IoT and the Smart City

Despite 'smartness' covering a variety of factors including not just information and communications technology, but also sustainability, resilience, innovation and business, (Cavada and Rogers, 2014), the role of new technologies in so-called smart city initiatives is well established in the public mind and media coverage. Close links are often made between smart cities and ICT infrastructure at various levels (Anthopoulos and Vakali, 2012). Since Weiser (1999) coined the term 'ubiquitous computing' in the late 1980s, there have been enormous advancements in the capabilities and use of connected, responsive devices which have come to be known as the IoT (Yoneki, 2005, Barbosa, 2015). These devices are vastly heterogenous; but have in common that they are familiar, everyday objects that have been equipped with additional networked capabilities (Lindley et al, 2019). Increasingly, these technologies are incorporated not just into private or semi-public spaces, but also into the public spaces of the data-driven urban environment (Foth et al, 2021). In many cases, techno-solutionist approaches presume the benefits of this increasing use of IoT for civic services, and the deployment of sensors for purposes such as smart lighting, parking or air quality measurement are considered to be operational matters falling outside direct political oversight and not requiring extensive citizen consultation (Cardullo and Kitchin, 2018). However, many scholars have argued that such technologies are not neutral, and there can be structural

imbalances in where the benefit is seen, and in agency, power, and who gets to decide what data is used, and how (Foth et al, 2021). Interrogative tools to support public scrutiny are therefore required to uncover the impacts of technology deployments, to support transparency and privacy by design, and in some cases, to establish whether the preferred course of action may be a decision not to use such solutions at all (Veale, 2020).

Privacy, trust and transparency

Concerns relating to IoT privacy, data management and mass surveillance are as old as the technology in question; Weiser's work in ubiquitous computing described how individual privacy might be compromised by existing in a shared spare where data was transferred indiscriminately (Weiser, 1999). Peppet (2014) suggests that not only are IoT devices particularly prone to privacy risks, but there are also challenges to meaningful consumer consent. For example, it may be necessary to make decisions about data sharing at a time when long-term harms that develop gradually over time are unknown. Choices made at this time may subsequently remain largely fixed (Lehtiniemi and Kortesniemi, 2017). With distributed systems in public spaces, the ability to give informed consent may be difficult or even impossible. Privacy policies, if they are available, may be complex and lengthy, hindering understanding (Luger et al, 2013). Even if members of the public are able to access and understand full details of public space deployment policies, and have the ability to interrogate data collected about themselves, there may be no practical way to opt out without restricting themselves from use of said spaces. Many thus advocate for a privacy-by-design approach to networked technologies (Langheinrich, 2001).

Read et al (2016) suggest that "building trust and confidence among all stakeholders, including the public, is considered to be essential to gain acceptance and

ensure continued development of a new technology." Transparency is a key component of trust (Castelluccia et al, 2018); in order to adequately assess whether a particular deployment is acceptable and meets criteria for a responsible use of technology, it is necessary to have access to the pertinent information; this requires transparency at all stages of the process, not just once the deployment is in place. Weber and Weber (2010) identify multiple categories of IoT transparency which should be considered. Cottrill et al. (2020) similarly highlight that IoT transparency functions at various levels; that of each individual device and its operation, the system as a whole, and the surrounding management and governance processes. Several projects and initiatives have discussed whether devices should display product labelling which would indicate their capabilities, functionality and privacy protection features (Emami-Naeini, 2020, Lu, 2019). System transparency should expose wider functionality such as data transfer and management which goes beyond individual devices. Transparency of management and governance processes additionally allows engagement with the underlying technological affordances at the previous two levels. The Chicago Array of Things is an example of a project which initially did not provide transparency, and had to revisit its privacy policies and also the hardware and software used in their devices (reducing the transmission of sensitive data by increasing on-device processing) following public pushback (Jacobs et al, 2020a). Such failures of transparency are often not intentional, but these potential issues may not be considered at early stages of the project, or are held simply as implicit knowledge and not codified or communicated.

In this work we designed transparency tools which expose elements of IoT deployments which may currently be opaque. By utilising design methods, we aimed to uncover knowledge which may be implicit, or identify gaps in knowledge which must be filled. This has two objectives, both of which aim to increase transparency and trust.

Firstly, by exposing these deployments to public scrutiny, it enables increased democratic participation (Weber and Weber, 2010) thereby allowing interested, impacted parties to interrogate deployments. By this mechanism, it is easier to understand how to enquire about what the systems consist of, what data are collected and how they might be used, and how such systems are governed. Secondly, by introducing those organisations that are planning or implementing deployments to questions which the public and other stakeholders might ask, we aim to expose knowledge gaps and areas where activities may be undertaken without full consideration of the impacts. By interrogating projects at an early stage, the intention is to prevent scenarios like the Chicago Array of Things example described above. The intention is that those considering or carrying out deployments will increase transparency in the deployments, in turn increasing public trust. We do not prescribe the mechanisms for transparency, or suggest that all information should be made transparent in every circumstance. Rather, we wish to provide tools which can expose what information exists about a particular system and *might* be made transparent to various stakeholders.

Developing Transparency Questions

The first stage of this process was to use a mixed methods approach to develop a comprehensive list of questions representing information that might be made available about a distributed urban IoT deployment. A system with the highest level of transparency would have answers to all applicable questions available. The research process to develop these questions took several forms, outlined below:

Questions arising from the literature

To develop an initial corpus of questions, a review of relevant literature was undertaken. This included examination of a variety of available documentation including industry guidelines on trust and IoT, published privacy regulations including the GDPR, data models for describing sensor networks such as the W3C Semantic Sensor Network Ontology, and an IoT Trust Framework published by the Online Trust Alliance. Information required for transparency included in these resources was restated in the form of 93 questions, which were manually clustered thematically by the project team into seven categories: *data collection; use of data; sharing of data; data storage; IoT system composition; sensing capabilities of the IoT system;* and *deployment of the IoT system.*

For example, questions in the category of data collection included:

- What data are collected?

- Are data being collected for purposes other than supporting the functionality/services being provided?

Nine questions were broader and did not relate specifically to any of these themes, therefore were categorised as *miscellaneous*, for example:

- Is there an independent dispute resolution body that a data subject can contact?

Participatory design fiction research

Questions were also gathered from prior work conducted as part of a wider programme of participatory research examining trust factors in public IoT deployment (described in Jacobs et al, 2020b), including a participatory workshop on public space IoT which discussed real-world IoT deployments, and used design fiction objects as prompts - diegetic fictional prototypes to allow the nine participants, who were local stakeholders

(including members of the public) to immerse themselves in a plausible fiction of a place-based deployment of IoT technology.

Data resulting from the workshop included transcriptions of user comments and responses to discussion questions recorded via worksheets. Key transparency requirements discussed by the participants were restated in the form of questions, the majority of which aligned with one of the seven core categories previously identified from the literature (see above). However, an additional category of questions was identified which centred on *communication*. An example of a question in this category is:

- Is a plain English explanation available?

Governance research

In parallel to the activities outlined above, we also examined current governance processes and policy surrounding the implementation of public space IoT deployments. This included a literature review of national and international smart city policy documents and standards, attendance at policy-related smart city events to gain an overview of the current landscape, case study examinations of particular deployments, and interviews conducted with representatives of IoT projects in two UK cities: Bristol and Aberdeen.

We also observed activities of local IoT projects initiated by different types of actors. These included:

1) the Air Aberdeen group[1], a citizen led grass roots community effort to build and distribute low-cost air quality monitoring devices to understand air pollution,

2) activities undertaken by Aberdeen City Council as part of their digital strategy, and

3) pilot projects developed at the University of Aberdeen through their IoT and Data working group, established to consider how to best utilise IoT technologies for campus management.

From our research with these individuals and organisations, a picture emerged of a general trend towards the encouragement of IoT deployments, often under the banner of smart city initiatives or inspired by specific challenges. We observed that these are often heavily siloed, with limited inter-locality communication even between projects with similar scope and remit, beyond celebratory sharing of successes. Much less frequently were negative outcomes or challenging processes shared, leading to a tendency for similar challenges to be encountered many times by different organisations. Examples of this included difficulties encountered when moving from pilot projects to larger scale initiatives which required much higher levels of infrastructure and management to support the technology.

As a result of this research, questions were formulated based primarily on interview transcripts, but also supplemented by observational work, which represented some of the key information that those managing such deployments might wish to have access to during the process; some of which came from post-hoc speculation by

---

[1] https://www.airaberdeen.org/

individuals and organisations about what they wish they had known earlier in the process.

Prototyping a Governance Tool

Having generated a comprehensive set of questions, the next step was to design a tool which would present appropriate questions that those wishing to interrogate their own proposed or extant IoT deployment could use to check whether all the appropriate information is available and supports transparency.

Question categorisation

A manual sorting process (Fig 1) was used to remove duplicate questions and further categorise the questions. In addition to the eight initial categories mentioned above, seven additional categories were identified: *Accountability, People, Purpose, Data Quality, Technical,* and *Logistics*. Individual questions within these categories were also grouped based on specific dependencies. For example, the relevance of the question 'How long will the deployment last?' is dependent on the answer to the question 'Does the deployment have a planned end point?' Some questions were reworded to increase clarity for a general audience.

[Insert Fig 1 here]

Fig 1. Manual sorting of questions into categories including linked question sets

Given the large number and broad range of possible questions, any tool must allow for filtering, in order to access only those relevant to specific circumstances. As an initial test of the utility of such a restricted question set, a manually filtered set of around 80 questions relevant to a particular deployment context (the Air Aberdeen project) was transcribed onto file cards which were bound into sets related by topics of interest (Fig 2). A preliminary discussion of these questions with members of this

group[2] took approximately two hours and received positive feedback, with the group suggesting that such an exercise would be particularly useful for developing a FAQ section of the project website to provide information for participants and those curious about the project.

[Insert Fig 2 here]

Fig 2. Manually filtered questions bound into topic sets

Initial Prototypes

Three initial 'paper' prototypes were designed with the aim of presenting these to a focus group whose feedback would allow the refinement into a working prototype for further testing and evaluation. At this stage, it was undecided whether a physical or digital tool would be a more effective mechanism for allowing interaction with the question set, so a variety of media were chosen for presentation. For the purposes of these prototypes, the question filtering process was designed around two layers of separation. The first of these layers identified which organisation type was responsible for the deployment – a public sector organisation, a private sector organisation, or a citizen group; the second identified the stage the project was at, with four stages at which users might apply the tool:

- we have a problem we want to solve;

- we have the opportunity to use technology;

- we have decided which devices we are using, or

- we have communicated to people about the project.

---

[2] This was undertaken as part of a 'Code the City' hack weekend event: https://codethecity.org/

Physical card game

The first format selected for the tool was a physical set of cards (Fig 3), as an extension of the earlier file card test version. Card-based tools for ideation and design activities are common, for example the IDEO method cards[3]. Friedman and Hendry (2012), in discussing the use of cards as a genre of design toolkit, note that "the physical format allows for persistence and recombination of the discrete ideas represented on individual cards." Several recent projects have used card-based formats for activities to explore and examine new technologies and their implications, for example, to explore legal aspects of data protection (Luger et al, 2015), algorithmic bias (Koene et al, 2019) or the so-called sharing economy (Fedosov et al, 2019). A number of recent projects have also used cards for ideation in the development of IoT devices and systems (Mora, 2017).

In our prototype, each question was written on an individual card. The cards had coloured 'tags' at the top representing which of the four project stages the question was relevant to, and the different organisations which might be initiating the project. The questions were grouped by topic, with the back of each card marked with ideograms indicating the topic.

[Insert Fig 3 here]

Fig 3. Physical card game paper prototype

Some questions in the set were those on which other questions depended, as described above. In this case, the initial questions were constructed as 'envelopes' within which the follow up questions were placed, to be answered only if the first question had a positive answer.

---

[3]http://www.ideo.com/work/item/method-cards/.

Adventure Game

The second prototype was inspired by 'adventure game' formats of text-based narrative fictions which lead players through a series of statements and decisions that build on previous choices. This was developed using a game prototyping tool called *Twine* (Engström et al, 2018), one of a number of recent platforms that allow non-technical people and those who are not professional coders to experiment with game making. It has been specifically designed to encourage the development of interactive experiences for purposes such as to "learn more about a certain topic" or "change perspectives in a story" (Friedhoff, 2013). When constructing a 'game', short segments of text are input and connected in a graphical interface that resembles a pin-board and notecards connected by string, providing a reflection for our ordering system of linked questions displayed similarly. The platform therefore allowed the question set to be replicated digitally and provided a quick tool to prototype a visual representation of a text-based, context sensitive question delivery system. The questions within each topic were displayed sequentially as a series of bullet points allowing the user to navigate through appropriate questions at their leisure. The particular questions revealed were selected through previous user choices. For example, selecting the stage 'we have decided which devices we are using' followed by the topic 'system deployment and logistics' reveals first questions in this topic: 'Who decides where devices are located? Is any bias involved in this?'

Bubble Map

The third prototype was a (non-functional) representation of a visualisation based digital interface, based on expanding questions in a 'bubble' format (Fig 4). The design was inspired by visualisation tools such as those described by McCandless (2012), which

use a radial tree structure (Draper et al, 2009) to display extending lists of content as needed. These visualisations combine the exploratory, fluid capabilities which characterise card-based tools with a digital format that allows more rapid information sifting and the ability to rapidly hide or reveal information as appropriate. Visualisations encourage exploration and sharing of data, and can be useful tools to initiate group discussion (Viegas and Wattenberg, 2006).

[Insert Fig 4 here]

Fig 5. Bubble prototype

The central circles represented a category topic, with the surrounding circles indicating the top-level questions relevant for each topic, which could be expanded by clicking to reveal second level questions. Indicator bars showed how far through the questions the user had progressed, and percentages of questions that had been answered versus questions which could not currently be addressed.

Prototype evaluation

These three prototypes were presented to a focus group consisting of five individuals including representatives from groups engaged with during our initial participatory and governance research (Air Aberdeen, the University of Aberdeen IoT development group), and local representatives who previously participated in workshops. After prototype demonstrations, a discussion was held which provided feedback on the various formats for presenting the questions.

There was no overall consensus on which prototype was preferred. The physical card game version was praised for being easy to manipulate and understand, particularly if it were to be used by those who did not have a high level of familiarity with technology. It was suggested that this format might be particularly useful for data subjects to identify questions that they should be asking about deployments, for

example members of the public querying devices deployed locally. The *Twine*-based adventure game prototype was praised for being simple to understand and presenting only relevant questions in a clear and unambiguous manner. Feedback on this version related to specific presentation, for example, that questions should be presented sequentially rather than under a category heading. The bubble map was less favoured, and comments suggested that the participants thought it was confusing and presented too much information at once in a way that was not intuitive to users. However, some aspects of this prototype were praised, such as graphics to indicate previous choices about stage of the process and identity of the user, and what proportion of questions had been answered so far.

A final point of feedback given more broadly across the prototypes was that it would be beneficial to log both answered questions, and those marked as requiring further consideration. Questions in the latter category might, for example, require additional research by those undertaking or wishing to understand a deployment, or be aspects which had not previously been considered and thus required further action before being fully addressed.

Governance Tool Design

Based on the feedback from the focus group, a full prototype was developed functioning in two different modes: A physical card deck was developed primarily for those who wished to learn about deployments, but who themselves may not possess expert knowledge of IoT technologies, using only the subset of questions relevant to data subjects. A digital version of the tool was also developed which gave access to both this set of questions and the wider corpus, for use by groups representing other stakeholder types.

The Physical Tool

[Insert Fig 5 here]

Fig 5. Physical tool card set

The 15 question categories were refined for clarity and conciseness by combining and renaming some categories, resulting in 10 categories. These were: *People, Purpose, Public Relations, Communication Strategy, System Development and Logistics, System and Devices, Collection and Use of Data, Data Storage and Sharing, Governance and Accountability,* and *Legal and Financial.*

For each topic category, a set of cards were designed as well as a topic category box to store them (Fig 5). The cards were laminated with dry-wipe plastic, with a space for writing comments. The use of dry-wipe meant that questions could not only be edited, but also erased completely for re-use of the tool.

The placement of these topic category boxes within the larger box provided a clear pathway for navigating through the questions in a suitable order. The questions were numbered sequentially from 1 to 153, reinforcing the suggested pathway through the questions.

It was also necessary to provide users with the ability to answer or omit questions depending on their relevance to the stage or stakeholder group. As with the earlier prototype, the cards included colour indicators at the top which displayed the deployment stage at which the questions were applicable. Within the card set box, there is a '*Discard'* box, in which cards that do not display the appropriate colour for the stage in progress, or are deemed irrelevant in that particular deployment, might be placed. The cards of dependent follow up questions were attached behind the question to which they related, and these could be disregarded if the lead question was answered

negatively. Attached *Prompt cards* led the user to decide whether or not to discard related questions in other category card sets that would be no longer relevant.

Within each category set, placed within the category box, are two divider cards named '*Resolved*' and '*Review*'. Each card in turn is placed behind the appropriate divider card to assist in working through the topic questions. Based on feedback from the focus group, the separation of questions which need further consideration (*Review)* and those which have answers available (*Resolved)* gives an estimation of how well the project is understood, allowing users to return to questions that may require further attention, and also provides a collection of answers for use in communication tools such as the construction of an FAQ, or to refer to internally when considering the project.

The box included an inlay tray to fit all contents comprising:

- 10 category card sets, each containing 2 divider cards '*Resolved*' and '*Review*'

- 1 *Discard* box

- 1 dry wipe pen

- 1 dry wipe cloth

- 1 set of instructions, also including an introduction to the IoT and explaining the project origins.

The Digital Tool

Providing a digital version of the tool allows more sophisticated sorting and presentation of the questions based on the requirements and specific situation of the user. The digital tool maintains a complete list of questions, and users are only presented with those which are relevant. The design of the tool was intended to be a combination of the simple text-based presentation of the *Twine* prototype, with an interface intended to resemble the physical aspects of the card deck and based on the same design principles. The colour palette used for the question topics is the same as

that used for the physical version, and in order to maintain similarity between the two versions, the question frame is styled as a card.

After an initial page containing instructions similar to those in the paper card deck, users are presented with three questions to provide context for the initial question filtering. These ask the user to indicate the roles of those present completing the exercise, the stage the project is at (see above), and which topic to start with. (see Table 1). An indication of who is present during the exercise was decided to be a more useful filtering exercise than organisation type, in part because differences between the question sets for different organisations were found to be minimal, and also because we found that the ability to answer certain questions was dependent on knowledge which may be linked to specific roles. For example, technical questions about the functioning of the sensors and devices require the presence of someone with knowledge of their detailed operation, which is not always the case when governance and management questions are being considered, though one can impact the other. By providing those using the tool with the opportunity to see that there are questions applicable to different roles, we hope to encourage more collaboration between those with different knowledge bases.

The main interface (Fig 6), displayed after the initial filtering questions, is composed of three main elements: the top bar displaying the current roles of those using the tool and the project stage, the middle section enabling navigation across questions and progress monitoring, and the bottom section for viewing and answering questions.

[Insert Fig 6 here]

Fig 6. Main digital tool interface

The user can view a question by clicking on the red, yellow or green tiles in the middle navigation bar corresponding to unanswered, marked for review, and resolved questions,

respectively. Only the first of the unanswered questions from each topic can be viewed, which prevents users from 'skipping forward', thus the preferred question order is preserved. Questions marked as light grey require additional roles or different project stages to be selected at the initial filtering stage or by changing these answers via controls on the top bar. Questions marked as dark grey require specific answers to prior questions before viewing.

When a question is displayed, the user must provide an answer to the question and mark it as *resolved* or *for review*. This echoes the system of dividers used to allocate questions in the physical card set. Depending on the type of question, the answer may be either a multiple-choice selection or free text. After a question is completed, clicking the next button displays the next question in the order. As described above, some questions can be linked to each other even if they belong to different topics. Any such linked questions are displayed as smaller cards next to the current question, allowing the user to navigate to them. Three summary progress bars on the right-hand side of the navigation pane display the overall percentages of questions resolved, for review or not yet examined.

The tool allows users to print questions and answers, allowing review and discussion of the answers without being bound to the Web application itself. The print function allows for separation of *resolved* and *for review* questions before producing a PDF file for printing.

Evaluation and Results

Initial responses to the physical cards from three local residents who attended a public engagement event and had not seen the tools before, were positive. Use trials with the wider community to evaluate whether the cards were beneficial in considering IoT deployments were planned to take place in early 2020, but were unable to be completed due to restrictions caused by the Covid-19 pandemic.

To evaluate the digital tool, separate evaluation workshops were held with the University of Aberdeen IoT and Data group and Air Aberdeen, who were presented with the opportunity to use the tool in the context of real IoT deployments that were planned or in progress. In each 1.5 hour session, three to five representatives of the groups were presented with the digital tool and asked to use it to evaluate their own deployment. Minimal instructions were provided, and the participants were expected to explore and navigate the system on their own.

Positive responses were gained from participants who took part in the workshops. The general consensus from the discussion was that this would be a useful tool, particularly when planning deployments and considering how to structure communication about the deployment, and also to identify missing information or issues requiring further consideration. All those who trialled the tool gave positive responses when asked about the look of the tool, and whether the instructions made sense. When asked to indicate ease of use on a 5 point Likert scale, all agreed that the tool was 'easy to use', though some also suggested that the tool was unnecessarily complex (Fig 7).

[Insert Fig 7 here]

Fig 7. Evaluation responses

Other feedback included a number of comments that some of the questions were repetitive or inconsistent, and a positive response to the printable PDF function which allowed users to keep a record of the answers they had given to the questions. The prototype version of the tool is hosted on Github[4]. Future work is in progress to refine the digital tool and user interface based on the evaluation feedback, and trial it with a local council in a different region.

---

[4] https://trustlens.github.io/TrustlensPolicyToolkit/

Discussion

We began this paper by asking whether exposing potential transparency limitations can enable greater transparency, greater trust, and active citizen engagement in public space IoT. In developing the transparency tools, we have found that in many projects there is a significant amount of information about the deployment that is not available explicitly, but may be held as implicit knowledge or shared between multiple individuals within an organisation and thus not available in totality to any one individual, or easily transferred as learning points to other projects. We hope that by using this or similar tools, the externalisation of this knowledge and the actualisation of gaps where things are not currently known will lead to a more considered approach to IoT deployments.

This approach challenges technological solutionism, particularly when it is applied at the very early stages of a project before the deployment has taken place or even had the details confirmed. We found that the information requirements of stakeholders (including the public) were diverse, reaching beyond considerations of data management and storage, but also including governance and legal considerations, financial motivations and outcomes, and how communications around a deployment and its data were managed. Not every stakeholder will need access to all of this information.

Such approaches to IoT transparency do incur risks. Using such tools should not take place in isolation, and we must be cautious of tokenism – that by carrying out such an exercise the problem is considered solved. In addition, there could be a temptation to describe a system in which all of the questions in this tool have answers as having 'full transparency'. We suggest that no such thing as full transparency exists, only higher and lower levels. Too much information is as bad as too little, and we would suggest that

this tool be used for defining an *appropriate* level of transparency for communication to various stakeholders. This must not take the form of excessive policies and forms which will be challenging and off-putting to read and understand (Luger et al, 2013).

Taking this work forward will require understanding more fully what transparency means to communities, and how tools such as this in association with technology may support trust in public space IoT. Collecting the information via this tool is only the first step; organisations must then decide what to do with it. One option could be that collecting the information in this way allows it to be organised in a machine-readable format supporting the development of additional tools allowing citizens to interrogate IoT systems based on individual needs and level of understanding. Technology supported approaches such as consent intermediaries (Lehtiniemi and Kortesniemi, 2017) could select appropriate information to provide based on level of transparency required. These tools could also support regulation and policymaking, wherein new proposals could be assessed by civil authorities and/or required to meet certain levels of transparency.

Conclusions

In this work, we have used a design approach to expose the complexities of IoT systems at all levels, in a way that we hope can be used to support both the developers of such systems, and stakeholders at all levels including policymakers. By using tools such as these, there is not only the potential for an increase in transparency with regards the public being able to have knowledge about public space deployments, but also internal transparency within organisations responsible for managing and deploying these systems. However, the evaluation process made it clear that what we have developed is still a prototype, and requires further development work to create a usable, shareable output that can be of use to companies, public bodies and other stakeholders.

For example, some of the questions appear similar, and participants challenged the apparent repetition, yet particular nuances change the nature of the interrogation and must be carefully considered in terms of whether to combine, keep, or edit them to make them more distinct and remove duplication. We intend to carry out further work to develop the prototype into a finished tool, and believe that such tools will be effective in supporting transparency, and thus trust in IoT deployments.

Acknowledgements

Declaration of Interest Statement

No potential conflict of interest was reported by the author(s).

References

Anthopoulos, L. G., & Vakali, A. (2012). Urban planning and smart cities: Interrelations and reciprocities. In The Future Internet Assembly (pp. 178-189). Springer, Berlin, Heidelberg

Barbosa, J. L. V. (2015). Ubiquitous computing: Applications and research opportunities. In 2015 IEEE international conference on computational intelligence and computing research (ICCIC) (pp. 1-8). IEEE

Castelluccia, C., Cunche, M., Le Métayer, D. and Morel, V., (2018). Enhancing transparency and consent in the IoT. In 2018 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW) (pp. 116-119). IEEE.

Cardullo, P., & Kitchin, R. (2018). Smart urbanism and smart citizenship: The neoliberal logic of 'citizen-focused' smart cities in Europe. Environment and Planning C: Politics and Space, 37(5), 813-830.

Cavada, M., Hunt, D. V., & Rogers, C. D. (2014). Smart cities: Contradicting definitions and unclear measures. In World Sustainability Forum (pp. 1-12). MDPI AG.

Cottrill, C. D., Jacobs, N., Markovic, M., & Edwards, P. (2020). Sensing the City: Designing for Privacy and Trust in the Internet of Things. Sustainable Cities and Society, 63, 102453.

Draper, G.M., Livnat, Y. and Riesenfeld, R.F., 2009. A survey of radial methods for information visualization. IEEE transactions on visualization and computer graphics, 15(5), pp.759-776.

Emami-Naeini, P., Agarwal, Y., Cranor, L. F., & Hibshi, H. (2020). Ask the Experts: What Should Be on an IoT Privacy and Security Label?. In 2020 IEEE Symposium on Security and Privacy (SP) (pp. 447-464). IEEE.

Fedosov, A., Kitazaki, M., Odom, W., & Langheinrich, M. (2019). Sharing economy design cards. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (pp. 1-14).

Foth, M., Anastasiu, I., Mann, M., & Mitchell, P. (2021). From Automation to Autonomy: Technological Sovereignty for Better Data Care in Smart Cities. Automating Cities, 319-343.

Friedhoff, J. (2013). Untangling Twine: A Platform Study. Presented at DiGRA conference 2013

Friedman, B., & Hendry, D. (2012). The envisioning cards: a toolkit for catalyzing humanistic and technical imaginations. In Proceedings of the SIGCHI conference on human factors in computing systems (pp. 1145-1148).

Jacobs, N. Edwards, P., Markovic, M., Cottrill, C. D., Salt, K. (2020a) Who Trusts in the Smart City? Transparency, Governance and the Internet of Things. Data & Policy

Jacobs, N., Markovic, M., Cottrill, C. D., Edwards, P., Corsar, D., & Salt, K. (2020b). Made-Up Rubbish: Design Fiction as a Tool for Participatory Internet of Things Research. The Design Journal, 23(3), 419-440.

Kitchin, R., Coletta, C., Evans, L., Heaphy, L., & MacDonncha, D. (2017). Smart cities, epistemic communities, advocacy coalitions and the 'last mile' problem. It-Information Technology, 59(6), 275-284.

Koene, A., Dowthwaite, L., Lane, G., Webb, H.M., Portillo, V. and Jirotka, M.D.A., (2018). UnBias: emancipating users against algorithmic biases for a trusted digital economy.

Langheinrich, M. (2001). Privacy by design—principles of privacy-aware ubiquitous systems. In International conference on Ubiquitous Computing (pp. 273-291). Springer, Berlin, Heidelberg.

Lehtiniemi, T., & Kortesniemi, Y. (2017). Can the obstacles to privacy self-management be overcome? Exploring the consent intermediary approach. Big Data & Society, 4(2), 2053951717721935.

Lindley, J., Coulton, P., & Alter, H. (2019). Networking with Ghosts in the Machine. Speaking to the Internet of Things. The Design Journal, 22(sup1), 1187-1199.

Lu, Jacqueline (2019) How can we bring transparency to urban tech? These icons are a first step. Sidewalk Labs. https://www.sidewalklabs.com/blog/how-can-we-bring-transparency-to-urban-tech-these-icons-are-a-first-step/

Luger, E., Moran, S., & Rodden, T. (2013). Consent for all: revealing the hidden complexity of terms and conditions. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 2687-2696 https://dl.acm.org/doi/10.1145/2470654.2481371

Luger, E., Urquhart, L., Rodden, T., & Golembewski, M. (2015). Playing the legal card: Using ideation cards to raise data protection issues within the design process. In Proceedings of the 33rd Annual ACM conference on human factors in computing systems (pp. 457-466).

McCandless, David. Information is beautiful. London: Collins, 2012.

Mora, S., Gianni, F., & Divitini, M. (2017, June). Tiles: a card-based ideation toolkit for the internet of things. In Proceedings of the 2017 conference on designing interactive systems (pp. 587-598).

Peppet, S. R. (2014). Regulating the internet of things: first steps toward managing discrimination, privacy, security and consent. Tex. L. Rev., 93, 85.

Read, S. A., Kass, G. S., Sutcliffe, H. R., & Hankin, S. M. (2016). Foresight study on the risk governance of new technologies: The case of nanotechnology. Risk Analysis, 36(5), 1006-1024.

Suzuki, L. and Finkelstein, A., 2019. Data as infrastructure for smart cities (Vol. 23). Computing and Networks.

Veale, M. (2020). A Critical Take on the Policy Recommendations of the EU High-Level Expert Group on Artificial Intelligence. European Journal of Risk Regulation.

Viegas, F. B., & Wattenberg, M. (2006). Communication-minded visualization: A call to action. IBM Systems Journal, 45(4), 801.

Weber, R. H. & Weber, R. (2010). Internet of Things. Berlin: Springer Berlin Heidelberg.

Weiser, M. (1999). The computer for the 21st century. ACM SIGMOBILE mobile computing and communications review, 3(3), 3-11.

Whicher, A. (2021). Evolution of policy labs and use of design for policy in UK government. Policy Design and Practice, 1-19.

Yoneki, E. (2005). Evolution of ubiquitous computing with sensor networks in urban environments. In Ubiquitous computing conference, metapolis and urban life workshop proceedings (pp. 56-59).