
A Design Philosophy for Agents in the Smart Home

William Seymour

University of Oxford
Oxford, OX1 3QD, UK
william.seymour@cs.ox.ac.uk

Abstract

The home is often the most private space in people's lives, and not one in which they expect to be surveilled. However, today's market for smart home devices has quickly evolved to include products that monitor, automate, and present themselves as human. After documenting some of the more unusual emergent problems with contemporary devices, this body of work seeks to develop a design philosophy for intelligent agents in the smart home that can act as an alternative to the ways that these devices are currently built. This is then applied to the design of privacy empowering technologies, representing the first steps from the devices of the present towards a more respectful future.

CCS Concepts

•Human-centered computing → Human computer interaction (HCI); Empirical studies in HCI;

Introduction

My home is my castle. The home is an extremely important place in people's lives, and one of the few places over which an individual truly has control. The home also represents a place where one can withdraw without fear of being watched or judged.

At the same time, we have seen the encroachment of connected intelligent agents into the home. These devices

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
CHI '20 Extended Abstracts, April 25–30, 2020, Honolulu, HI, USA.
Copyright is held by the author/owner(s).
ACM ISBN 978-1-4503-6819-3/20/04.
<http://dx.doi.org/10.1145/3334480.3375032>

Author Bio

William Seymour is a fourth year DPhil student at the University of Oxford's Department of Computer Science and Centre for Doctoral Training in Cyber Security.

As part of the Human Centred Computing group he uses speculative design and technology probes to explore alternate design ideologies for intelligent agents in the smart home.

have powerful sensing and computational capabilities, and as a result are able to offer increased safety, automation, and convenience. Devices such as voice assistants are taking on roles previously only described in science fiction, as teachers helping with homework or as carers for the memory impaired. But there is increasing concern that these devices do not primarily act in the interests of their owners, but rather those of developers and manufacturers whose business models revolve around the shaping of behaviour and collection of personal data.

However this is not the only way that such agents could be designed. To this end, my research explores other ways of designing intelligent agents for the connected smart home. Beginning by mapping out the ethical concerns generated by contemporary smart devices, my work develops an alternate design philosophy of what these devices should do, who they should serve, and how they might better meet users' needs. This is achieved in three distinct but complementary ways:

1. *Theoretical*—using philosophical accounts of respect to develop a design philosophy, suggesting ways that intelligent agents might better integrate with people's lives and living spaces
2. *Speculative*—using speculative design and design fiction to envision scenarios where devices and agents adhere to this design philosophy, showing how the theoretical components inform their design, highlighting how they differ from contemporary products, and critically reflecting on how such products are built
3. *Empirical*—taking the first steps towards the futures depicted in the speculative elements, my design philosophy will be applied to privacy-empowering tech-

nologies in the home, creating prototypes that explore the design space for near-future devices

Background & Motivation*Smartness as a Marketing, Historical, and Ideological Concept*

As with any loosely defined technological concept (such as cloud computing), perceptions of smartness and what it means for devices to be smart are influenced and guided by a number of interwoven (and often competing) narratives. Marketing narratives of smart devices often present polar views of users and the threats or hazards that they mitigate. Tropes around smartphone usage, for example, are often clustered around the benefits of integration and the perils of dis-integration, pitching the ideal of the connected and productive phone user against fears of becoming addicted or out of touch [7]. Manufacturers also frequently fall into the trap of 'solutionism' [15], situating devices as solutions to problems without acknowledging the risks and drawbacks that come with them, a process designed to "bring [detectors] into the fold while keeping [vendors'] central mission of capital accumulation and technocratic governance intact" [10].

To an extent, these marketing narratives represent the hopes and dreams of contemporary society, often echoing the depictions of intelligent agents and artificial intelligence in science fiction. Portrayals of future cultures where artificial intelligence enables prosperity and flourishing (e.g. the works of Asimov and others from the 'golden age' of science fiction), can be contrasted with more dystopian visions (e.g. works by Gibson, who pioneered the cyberpunk genre of science fiction) where technology is associated with "human pollution, global destruction, individual control, mental regression, and dehumanisation" [9]. This extends to depictions of individual agents or intelligences, such as J.A.R.V.I.S (*Iron Man*, 2008) and HAL 9000 (*2001: A Space*

Odyssey, Arthur C. Clark) respectively. Unfortunately, these (literal) narratives, as well as those told by designers and marketers, often come at the expense of more nuanced understandings of how people might or do use technology, pitching practices and values against each other (like the smartphone tropes of integration and dis-integration) and polarising perceptions of those technologies [7].

Ethical Concerns in the Smart Home

Many prior works have investigated the ethical concerns arising from the use of smart home devices, from smart toys [13], smart TVs [12], and smart meter agents [5, 22], to robotic home assistants [27]. Common themes include how automation intended as a convenience can still diminish one's sense of autonomy, and the adverse effects on the relationships and social orders in the home of technology that does not recognise them. Other work has also begun to unpick the subtle long term effects of living with continuously sensing technology designed either explicitly or implicitly as surveillance devices [4, 17]. Research focused more specifically on voice assistants has examined the Media Equation [21] as applied to voice assistants [14, 20], as well as how families adapt existing structures of authority to voice assistants [19].

A substantial body of work addresses perceptions of and threats to privacy in the smart home. Built on the foundation of Altman's theory of privacy as "selective control of access to the self" [1], Palen and Dourish propose that privacy violations take place along boundaries of private/public disclosure, identity, and temporality [18]. Nissenbaum's theory of contextual integrity similarly demarcate acceptable data flows along the lines of contextual information sharing norms in relation to the subject, sender, recipient, type, and transmission of data [16]. But in order to evaluate data flows, users need to be aware of them. Other research has

sought to explore potential threats that users are *unaware* of (e.g. [25]) or lack the required knowledge to fully understand [26, 29]. Sensemaking and similar approaches have shown promise in closing this gap, helping to give users a greater situational awareness of what their devices are recording and sharing about them [28, 30].

Research Approaches & Methods

Technology Probes

Technology probes are fully- or semi-functional prototypes designed be deployed and used over an extended period of time. The data gathered from the probe and it's users is then used to meet a range of research goals: the social science goal of understanding how new technology is used in context, the engineering goal of testing implementations of that technology, and the design goal of generating creative new ways to use technology to meet people's needs and desires [8].

Speculative Design

Falling under the umbrella of critical design methods, speculative design prompts reflection on contemporary culture and design processes by taking familiar objects and exaggerating select features until they become strange or uncomfortable. By asking questions about the norms that these design artefacts violate, one is forced to ask what kind of society would create such a product, and the extent to which that society differs from or mirrors our own. In this way, speculative design objects act as catalysts for debate about what people really want, rather than suggesting themselves what is desirable [6].

Design Fictions

Like speculative design, design fictions also offer strange or uncomfortable visions of the future, but do so without creating physical artefacts. This allows for the creation of




Figure 1: The smart functionality cards as arranged by an interview participant.




Figure 2: The range of possible responses to under-specified user requests.

spaces, organisations, and entire worlds that support and contextualise design artefacts, occupying the space between “the arrogance of science fact, and the seriously playful imaginary of science fiction” [2].

Results to Date

(1) Understanding Current Smart Home Devices

In order to develop a design philosophy for agents in the smart home, it was first necessary to understand the concerns that users have about devices, and how this is related to what makes them ‘smart’.

From surveys of smart device owners ($n=120$), eight different characteristics were identified that were seen as core to devices being ‘smart’. These then formed the basis for semi-structured interviews ($n=15$), where participants were introduced to a set of cards with the eight functionalities. After being asked to arrange cards next to devices they thought had that capability (see Figure 1), participants were then asked if they would have any concerns about using the device, and whether their concerns would be alleviated by adding/removing functionality cards. In order to go beyond immediate personal reactions to the devices, we then asked them to respond to a set of vignettes depicting smart home contexts (e.g. rental accommodation with devices provided by the landlord).

Thematic analysis of the interviews revealed a variety of interactions between device functionalities and ethical concerns around the use of technology to observe others (particularly children and older adults), transparency around how devices were exercising their capabilities, and the role that devices played as social actors in the home.

This work is currently under review.

(2) Using Respect to Inform Device Design

Having seen how devices, apps, and services balance the needs and interests of users against those of other stakeholders, how might this calculus be changed in order to facilitate the building of trust and comfort? In this work, philosophical accounts of respect are used to precisely characterise and make salient the ways users’ needs can be treated by increasingly adaptive and intelligent systems.

Respect is a natural and integral part of human relationships, and an essential skill in social cognition. Respect is also intrinsically related to perceptions of trustworthiness and loyalty. As devices and systems become increasingly advanced, our interactions with them start to resemble human social interactions in several ways. Natural modalities such as speech introduce the possibility for ill-posed or under-specified requests (see figure 2), introducing a gulf of interpretation where devices are tasked with inferring what users really need. In this context, we unpack respect into a rich and descriptive typology that describes how future systems might better serve users’ needs.

This work is currently under review.

(3) Privacy Therapy: What if Your Firewall Could Talk?

In order to provide an example of how the thesis’ alternate design philosophy might be realised, this project presents a speculative prototype of a ‘respectful Alexa’, a smart assistant that is loyal to its users and prioritises their needs over those of its manufacturer and developers. The prototype is engineered using a real Echo unit running a custom Alexa skill. Conversing with the prototype highlights the different ways that such a device might accommodate its user’s needs, and is designed to be interoperable with the Aretha technology probe (research item 4 below).

Accompanying the speculative design artefact is an exploration of the technical measures that this class of ‘respectful’ devices could implement. Examples include anonymisation by running extraneous requests (e.g. fetching weather reports from five extra cities), or means by which devices could gain the trust of users (e.g. by stopping heartbeats and analytics data for a short period of time).

See [23, 24] for more detail.

(4) Informing the Design of Privacy-Empowering Tools for the Smart Home

Prior work shows that despite privacy concerns over connected devices in the home, users lack the insight, knowledge, and options essential for taking effective action. To start to understand the potential for new kinds of tools to address these gaps, we developed Aretha, a privacy assistant technology probe that combines a network disaggregator, personal tutor, and firewall to empower users on their home networks.

To better understand how this combination of resources could allow users to gain awareness of data disclosures by their devices, form educated privacy preferences, and control unwanted data flows, Aretha was deployed in three households over six weeks. During this time participants were instructed to keep track of what was happening on the visualisations (see figure 3) and follow the content in the curriculum. Towards the end of the study the firewall control allowed them to block network traffic between specific devices and companies of their choosing.

The probe encouraged participant households to talk together about security and privacy, and lines of questioning evolved from *where* devices were sending data to *why* those destinations required it (and what were they doing with it). While the firewall ultimately did not meet user’s

needs, it did prompt them to find their own controls and suggested strategies to improve similar tools in the future.

This work is currently under review.

Next Steps

(5) The Power of Giving Devices a Voice

Speech activates the same locations of the brain regardless of whether its origin is human or machine. As voice assistants become more sophisticated and more devices support voice control, it is important to understand the effects this might have on those using these technologies. This paper will explore the presence of cognitive shifts that might result from long term use of devices that present themselves as human-like.

A survey of voice assistant users will be used to explore the interplay between three axes along which users perceive their voice assistants as: mechanical devices (inc. reliability, confidence, and consistency); social actors (how interactions relate to social rules and norms governing interpersonal relationships; and anthropomorphism (how much users ascribe human characteristics and intents to voice assistants).

From this I hope to discover how these axes interrelate, as well as how they correlate with ownership factors (e.g. length of ownership). It is expected that this will highlight potential concerns around deploying human-like devices, as well as suggesting strategies that can be taken to mitigate the potentially harmful effects of the Media Equation in voice assistants.

(6) Voice Assistants with Multiple Personalities

When you buy a voice assistant, you receive a physical object. The warranty covering the assistant applies to this physical object, but what is really being purchased is ac-



Figure 3: Time series visualisation in the Aretha probe. Shows traffic volume by company and by devices over a 24 hour period.

cess to the agent that ‘inhabits’ the plastic and silicon. The intelligent agents that drive current commercial offerings are designed to speak and act like humans, often engendering feelings of social presence in users [3]. Inspired by Japanese Shinto spiritual practices, this design fiction imagines voice assistants as hosts for different spirits concerned with human beings—kami. These assistants (mostly) want us to be happy; if they are treated properly they can bring us numerous benefits such as automating tasks for us, accessing information, or providing help and advice (known as the harmonious soul), but if they are disregarded then they can sow discord and cause trouble (known as the wild soul).

Through this lens I hope to contrast these fluid and sometimes capricious responses with contemporary notions of voice assistants as polite and subservient, as well as highlighting the gulf of interpretation in a (pseudo) real-world context. The design fiction asks us to imagine what kind of a society might adopt these kami assistants, and how they might adapt to them as they change between harmonious and wild.

(7) Informing Data Protection by Design and Default in Future Smart Homes

From its beginnings as a means of self-regulation, the privacy by design paradigm has recently been enshrined in the GDPR as ‘data protection by design and default’ (Art. 25). Properly enacting the spirit of these regulations in smart home devices requires an understanding of several areas that manufacturers have traditionally paid little attention to, such as the social dynamics surrounding shared or communal devices, and the use of technology by and for vulnerable groups.

To address this, I am part of a project funded by the UK Information Commissioner’s Office that will design, prototype,

and evaluate a number of novel devices that address the above issues. With the aim of informing the design of the smart home of the future, we will also explore the ways that hardware and software design can be used to raise awareness of the legal rights of data subjects, and how these might be customised to fit users’ needs.

See [11] for more information.

Expected Contributions

- Exploration of ethical concerns caused by current devices, highlighting the key tensions at play in the context of the connected home. Comprises research items (1) and (5).
- An alternate design philosophy for future intelligent agents in the smart home, illustrated by examples of speculative and fictional device designs that adhere to this philosophy. Comprises research items (2), (3), and (6).
- Design and implementation of near-future devices that bridge the gap between contemporary products and the design artefacts previously described. This will be done through the creation of prototypes that are then evaluated by users and industry practitioners. Comprises research items (4) and (7).

Acknowledgements

This work is supported by EPSRC through the PETRAS subproject *Respectful Things in Private Spaces* (grant number N02334X/1), and the ICO grants program (2018-2019).

REFERENCES

- [1] Irwin Altman. 1975. *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding*. (1975).

- [2] Julian Bleecker. 2009. Design Fiction: A short essay on design, science, fact and fiction. *Near Future Laboratory* 29 (2009).
- [3] Eugene Cho. 2019. Hey Google, Can I Ask You Something in Private?. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 258.
- [4] Eun Kyoung Choe, Sunny Consolvo, Jaeyeon Jung, Beverly Harrison, Shwetak N Patel, and Julie A Kientz. 2012. Investigating receptiveness to sensing and inference in the home using sensor proxies. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 61–70.
- [5] Enrico Costanza, Joel E Fischer, James A Colley, Tom Rodden, Sarvapali D Ramchurn, and Nicholas R Jennings. 2014. Doing the laundry with agents: a field trial of a future smart energy system in the home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 813–822.
- [6] Anthony Dunne and Fiona Raby. 2013. *Speculative everything: design, fiction, and social dreaming*. MIT press.
- [7] Ellie Harmon and Melissa Mazmanian. 2013. Stories of the Smartphone in everyday discourse: conflict, tension & instability. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1051–1060.
- [8] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, and others. 2003. Technology probes: inspiring design for and with families. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 17–24.
- [9] Dominic Idier. 2000. Science fiction and technology scenarios: comparing Asimov's robots and Gibson's cyberspace. *Technology in Society* 22, 2 (2000), 255–272.
- [10] Rob Kitchin. 2015. Making sense of smart cities: addressing present shortcomings. *Cambridge Journal of Regions, Economy and Society* 8, 1 (2015), 131–136.
- [11] Martin J Kraemer, William Seymour, Reuben Binns, Max Van Kleek, and Ivan Flechais. 2019. Informing The Future of Data Protection in Smart Homes. *arXiv preprint arXiv:1910.01973* (2019).
- [12] Nathan Malkin, Julia Bernd, Maritza Johnson, and Serge Egelman. 2018. “What Can’t Data Be Used For?” Privacy Expectations about Smart TVs in the US. In *European Workshop on Usable Security (Euro USEC)*.
- [13] Emily McReynolds, Sarah Hubbard, Timothy Lau, Aditya Saraf, Maya Cakmak, and Franziska Roesner. 2017. Toys that listen: A study of parents, children, and internet-connected toys. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 5197–5207.
- [14] Sarah Mennicken, Oliver Zihler, Frida Joldaschewa, Veronika Molnar, David Aggeler, and Elaine May Huang. 2016. It’s like living with a friendly stranger: perceptions of personality traits in a smart home. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 120–131.
- [15] Evgeny Morozov. 2013. *To save everything, click here: The folly of technological solutionism*. Public Affairs.

- [16] Helen Nissenbaum. 2009. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.
- [17] Antti Oulasvirta, Aurora Pihlajamaa, Jukka Perkiö, Debarshi Ray, Taneli Vähäkangas, Tero Hasu, Niklas Vainio, and Petri Myllymäki. 2012. Long-term effects of ubiquitous surveillance in the home. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. ACM, 41–50.
- [18] Leysia Palen and Paul Dourish. 2003. Unpacking privacy for a networked world. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 129–136.
- [19] Martin Porcheron, Joel E Fischer, Stuart Reeves, and Sarah Sharples. 2018. Voice Interfaces in Everyday Life. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 640.
- [20] Amanda Purington, Jessie G Taft, Shruti Sannon, Natalya N Bazarova, and Samuel Hardman Taylor. 2017. Alexa is my new BFF: social roles, user satisfaction, and personification of the amazon echo. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 2853–2859.
- [21] Byron Reeves and Clifford Ivar Nass. 1996. *The media equation: How people treat computers, television, and new media like real people and places*. Cambridge university press.
- [22] Tom A Rodden, Joel E Fischer, Nadia Pantidi, Khaled Bachour, and Stuart Moran. 2013. At home with agents: exploring attitudes towards future smart energy infrastructures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1173–1182.
- [23] William Seymour. 2018. How Loyal is Your Alexa?: Imagining a Respectful Smart Assistant. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, SRC20.
- [24] William Seymour. 2019. Privacy Therapy with Aretha: What If Your Firewall Could Talk?. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, SRC12.
- [25] Irina Shklovski, Scott D Mainwaring, Halla Hrund Skúladóttir, and Höskuldur Borgthorsson. 2014. Leakiness and creepiness in app space: Perceptions of privacy and mobile app use. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 2347–2356.
- [26] Madiha Tabassum, Tomasz Kosinski, and Heather Richter Lipford. 2019. "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)*.
- [27] Lachlan Urquhart, Dominic Reedman-Flint, and Natalie Leesakul. 2018. Responsible Domestic Robotics: Exploring Ethical Implications of Robots in the Home. (2018).
- [28] Max Van Kleek, Reuben Binns, Jun Zhao, Adam Slack, Sauyon Lee, Dean Ottewell, and Nigel Shadbolt. 2018. X-ray refine: Supporting the exploration and refinement of information exposure resulting from smartphone apps. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 393.
- [29] Max Van Kleek, Ilaria Liccardi, Reuben Binns, Jun Zhao, Daniel J Weitzner, and Nigel Shadbolt. 2017. Better the devil you know: Exposing the data sharing

practices of smartphone apps. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 5208–5220.

- [30] Max Van Kleek, William Seymour, Michael Veale, Reuben Binns, and Nigel Shadbolt. 2018. The need for sensemaking in networked privacy and algorithmic responsibility. In *Sensemaking in a Senseless World: Workshop at ACM CHI'18, 22 April 2018, Montréal, Canada*. Association for Computing Machinery (ACM).