Workshop on Reviewable and Auditable Pervasive Systems (WRAPS)

Chris Norval
Computer Science & Technology
University of Cambridge
Cambridge, United Kingdom
chris.norval@cl.cam.ac.uk

Richard Cloete
Computer Science & Technology
University of Cambridge
Cambridge, United Kingdom
richard.cloete@cl.cam.ac.uk

Milan Markovic
Computing Science
University of Aberdeen
Aberdeen, United Kingdom
milan.markovic@abdn.ac.uk

Iman Naja
Computing Science
University of Aberdeen
Aberdeen, United Kingdom
iman.naja@abdn.ac.uk

Kristin B Cornelius
Information Studies
University of California, Los Angeles
California, United States
krisbcorn@g.ucla.edu

ABSTRACT

Pervasive systems are increasingly being deployed in new and innovative ways – be it in our homes, vehicles, or public spaces. Such systems have the potential to bring a wide range of benefits, blending advanced functionality with the physical environment. However, these systems also have the potential to drive real-world consequences through decisions, interactions, or actuations, and there is a real risk that their use can lead to harms (physical injuries, financial loss, or even death). These concerns appear ever-more prevalent, as a growing sense of distrust has led to calls for *more transparency and accountability* surrounding the emerging technologies that increasingly pervade our world.

A range of things can—and often *do*—go wrong, be they technical failure, user error, or otherwise. As such, means to effectively *review, understand, and act upon* the inner workings of pervasive systems is becoming increasingly important. Means for reviewing and auditing how these systems are built/developed and used are crucial to the ability to determine the cause of failures, prevent re-occurrences, and/or to identify parties at fault. Yet, despite the wider landscape of societal and legal pressures for record keeping and increased accountability, implementing such transparency measures faces a range of challenges.

This workshop will bring together a range of perspectives into how we can better audit and understand the complex, sociotechnical systems that increasingly affect us (whether directly or indirectly). From tools for data capture and retrieval, technical/ethical/legal challenges, and early ideas on concepts of relevance – we solicit submissions that help further our understanding of how pervasive systems can be built to be reviewable and auditable, helping them to be more transparent, trustworthy, and accountable.

CCS CONCEPTS

• **Social and professional topics** → **Computing / technology policy.**

KEYWORDS

auditability; interpretability; accountability; transparency; compliance

ACM Reference Format:

Chris Norval, Richard Cloete, Milan Markovic, Iman Naja, and Kristin B Cornelius. 2021. Workshop on Reviewable and Auditable Pervasive Systems (WRAPS). In *Adjunct Proceedings of the 2021 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2021 ACM International Symposium on Wearable Computers (UbiComp-ISWC '21 Adjunct)*, September 21–26, 2021, Virtual, USA. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3460418.3479265>

1 CONTEXT

Pervasive and emerging technologies (e.g. AI, the Internet of Things, virtual and augmented reality, robotics) are becoming increasingly deployed into a range of scenarios – within our homes, vehicles, and workplaces [10]. In many cases, they promise to make our lives easier, safer, and better.

Through interactions with the real world (decisions, actuations, or data transfer), such technologies have the potential to drive significant and far-reaching consequences for those involved [4, 9]. Examples may include algorithmic models for lending decisions [7], using AI to predict students' exam scores [11], and AR-assisted surgery [1] – to name a few. Naturally, these consequences aren't always positive, and *things can—and will—go wrong*. It will therefore become ever-more important to understand what happened, why, who was responsible, (ex post) as well as preventing such occurrences in the first place (ex ante).

Such concerns are challenging for a number of reasons. For one, these technologies are increasingly being deployed as a 'system of systems', which may amalgamate a range of cloud services, external APIs and data sources [12]. This results in new points of failure, where issues such as service downtime, unreliable inputs or faulty sensors can propagate throughout. Second, the data flowing into, within, and out of these systems tend to be opaque, making it

This work is licensed under a Creative Commons Attribution International 4.0 License.

UbiComp-ISWC '21 Adjunct, September 21–26, 2021, Virtual, USA
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8461-2/21/09.
<https://doi.org/10.1145/3460418.3479265>

increasingly difficult for stakeholders (be they developers, regulatory authorities, or users themselves) to understand why certain outcomes or actions occur.

The *reviewability and auditability* of a system concerns the ability to monitor, oversee, investigate, evaluate and interrogate what is happening, or has happened, with a system [9]. In essence, it entails capturing information of the system and its operational context sufficient for supporting such actions. This might include system debug logs, or equally it could involve capturing data as it enters and/or transfers out of the boundary of the system (recording camera feed, sensor readings, preprocessed data streams, etc.).

Mechanisms for such, be they technical or administrative, can help to capture relevant information so as to facilitate explorations, investigations or otherwise interrogation of the technologies and stakeholders involved. From technical approaches (such as provenance [2]) to means of recording the socio-technical decisions and other aspects surrounding design and deployment (datasheets [6], model cards [8], etc.), there are a range of ways that more transparency could be delivered, for the betterment of its users. By refocusing data capture regimes for transparency purposes, we can look to bring about more accountability surrounding how such systems operate, and therefore increase trust – a crucial factor in their uptake and adoption [14].

These challenges also come amid wider societal and legal calls for more transparency and accountability surrounding the technologies that pervade our everyday environments. We are already seeing the regulatory landscape shift towards record keeping obligations and providing strengthened rights for those whose personal data are being processed in order to gain oversight over how their data is being used. The EU's General Data Protection Regulation (GDPR) [5] and the California Consumer Protection Act (CCPA) [3] represent two notable examples within a data protection context alone, and similar trends toward strengthened data rights look set to continue.

It stands that reviewability and auditability is important – however, as a research topic, the wider implications of using technical data capture methods to provide oversight over pervasive systems appears nascent. Aspects such as what should be captured, how, when, where, and for whom; many larger questions are emerging, and the stakes are getting ever-higher as the technologies in which we increasingly rely become more advanced.

1.1 Understanding the Challenges

Naturally, any means of capturing system data will introduce trade-offs, and auditability is no exception. For one, there will likely be performance and storage overheads of recording vast quantities of data, and deciding what and when data should be captured may be a complex and highly contextual undertaking. Such concerns are not limited to technical (system) logging alone – increasingly, work is recognising the sociotechnical nature of pervasive systems [6, 8, 9], and the wider processes involved in bringing ubiquitous technologies to fruition.

Also relevant are the privacy implications of capturing and storing what could amount to personal and sensitive information. Legal questions might also emerge over the capturing of evidentiary

data, data subject rights, and contractual obligations when concerning the monitoring of interactions with third parties (e.g. cloud providers).

Questions also surround how such information, which itself is likely complex or technical in nature, can be presented in a meaningful way. The usability of this information can directly influence its efficacy in explaining why complex decisions were made, or how certain data processing occurred. Such challenges are contextual; what's 'usable' for a technical developer may be entirely different to that of an end-user. Also crucial is that vast quantities of data are not used to mask and obfuscate questionable practices (the so-called transparency paradox [13]).

Such challenges call for different thinking into how we can better understand (interpret) the emerging technologies that pervade our world. UbiComp's diverse range of topics, experienced academic and industry partners, and its prestige will undoubtedly offer a rare opportunity to draw more attention to this important topic. Furthermore, we hope to encourage more members of the UbiComp community to further consider the implications, transparency, and accountability of that which they produce.

2 NATURE OF THE WORKSHOP

We will seek submissions which explore how technology can better meet the needs and expectations of its operators, creators, and society as a whole.

We intend to run a half-day workshop that will bring together new ideas into this nascent topic, collating some of the potential solutions and outstanding challenges to implementing more meaningful transparency throughout pervasive systems. We look to bring together experts from a range of disciplines, including those technical, legal, and design-oriented. We are anticipating approximately 5–6 accepted papers, and intend to advertise the workshop widely. Our wider aim is to foster a community of researchers looking to advance how various parties can understand and oversee technologies as they continue to emerge, and going forward, we hope to run this workshop year on year.

2.1 Paper Presentations

We will solicit paper submissions (max 6 pages) which further our understanding of reviewable and auditable pervasive systems. Each paper will be reviewed by two members of our program committee, which already contains several notable researchers across a broad range of topics.

Authors of accepted papers will present their work. Talks will be 15 minutes each (followed by 5 minutes of questions), though this may be adapted depending on the final number of papers accepted.

2.2 Panel Discussion

The workshop will culminate in an interactive panel discussion with a selection of experts (selected authors of accepted papers, and notable guests) to stimulate discussion and debate. Attendees will be able to pose suggestions, ideas, or thoughts on how to bring about greater awareness over auditability challenges within the UbiComp community. The co-chairs additionally have a selection of cross-panel questions to chair and steer the discussion.

2.3 Topics

We will encourage submissions from a range of topics that would help broaden the debate, including (but not limited to):

- Tools, techniques, and frameworks to assist in providing greater transparency and oversight over the workings of pervasive systems
- Methods for explaining and understanding systems/models
- Methods for fostering trust and transparency in pervasive systems
- The usability of audit data
- Performance implications of capturing audit data
- Privacy, security and data protection implications of auditability mechanisms
- Vocabularies and frameworks for modelling relevant information to support auditability and explainability
- Data aggregation and consolidation
- Legal considerations relating to record keeping and auditing mechanisms
- Access controls and data sharing regimes
- Audit log verification methods

2.4 Proposed Schedule

The structure of the workshop will consist of:

[0900–0915] **Welcome and Introduction**
 [0915–1000] **Keynote**
 [1000–1010] **Coffee break**
 [1010–1200] **Paper presentations**
 [1200–1225] **Panel Discussion**
 [1225–1230] **Wrap up and closing**

2.5 Organisers

Chris Norval is a Research Fellow within the Compliant and Accountable Systems group within the Department of Computer Science and Technology at the University of Cambridge. His research involves exploring how technology can better align with regulatory and societal calls for transparency and accountability. As part of this, much of his work has advocated for techniques, and identified the challenges, toward making complex systems more understandable and usable for a broad range of stakeholders. Chris was Student Volunteer Co-Chair at the 2019 UbiComp/ISWC conference.

Richard Cloete is a Research Fellow within the Compliant and Accountable Systems group within the Department of Computer Science and Technology at the University of Cambridge. His research lies at the intersection of technology and law, with a particular focus on methods and processes that work to make emerging technologies safer, more transparent, and accountable.

Milan Markovic is a Research Fellow in Computing Science at the University of Aberdeen, UK. 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. His recent experience includes work on transparency models and IoT solutions for food supply chains, smart cities and autonomous systems.

Iman Naja is a Research Fellow in Computing Science at the University of Aberdeen, UK. Her current research focuses on using provenance and Semantic Web technologies to explore how to realise accountable and transparent intelligent systems.

Kristin B Cornelius received her PhD from the University of Los Angeles, California and is a visiting researcher affiliated with the Compliant and Accountable Systems group at the University of Cambridge. She investigates the transition of paper documents to digital documents, including the consequences of users' interpretations of electronic Terms of Service agreements that have become commonplace across the internet.

REFERENCES

- [1] Hasaneen Fathy Al Janabi, Abdullatif Aydin, Sharanya Palaneer, Nicola Macchione, Ahmed Al-Jabir, Muhammad Shamim Khan, Prokar Dasgupta, and Kamran Ahmed. 2019. Effectiveness of the HoloLens Mixed-Reality Headset in Minimally Invasive Surgery: A Simulation-Based Feasibility Study. *Surgical Endoscopy* (June 2019). <https://doi.org/10.1007/s00464-019-06862-3>
- [2] Khalid Belhajjame, Helena Deus, Daniel Garijo, Graham Klyne, Paolo Missier, Stian Soiland-Reyes, and Stephen Zednik. 2013. PROV Model Primer. <https://www.w3.org/TR/2013/NOTE-prov-primer-20130430/>. Accessed: 2021-04-01.
- [3] California State Legislature. 2018. California Consumer Privacy Act of 2018. *Cal. Civ. Code* §1798.100 (24 September 2018).
- [4] Richard Cloete, Chris Norval, and Jatinder Singh. 2020. A Call for Auditable Virtual, Augmented and Mixed Reality. In *26th ACM Symposium on Virtual Reality Software and Technology* (Virtual Event, Canada) (VRST '20). Association for Computing Machinery, New York, NY, USA, Article 16, 6 pages. <https://doi.org/10.1145/3385956.3418960>
- [5] European Union. 2016. 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* L119 (4 May 2016), 1–88.
- [6] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna M. Wallach, Hal Daumeé III, and Kate Crawford. 2018. Datasheets for Datasets. In *Proceedings of the 5th Workshop on Fairness, Accountability, and Transparency in Machine Learning*.
- [7] Michelle Lee, Luciano Floridi, and Jatinder Singh. 2020. Spelling Errors and Non-Standard Language in Peer-to-Peer Loan Applications and the Borrower's Probability of Default. *SSRN* (2020).
- [8] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (Atlanta, GA, USA) (FAT* '19). Association for Computing Machinery, New York, NY, USA, 220–229. <https://doi.org/10.1145/3287560.3287596>
- [9] Chris Norval, Jennifer Cobbe, and Jatinder Singh. 2021. Towards an accountable Internet of Things: A Call for 'Reviewability'. *Privacy by Design for the Internet of Things. The Institution of Engineering and Technology*. (2021).
- [10] Chris Norval and Jatinder Singh. 2019. 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* (London, United Kingdom) (UbiComp/ISWC '19 Adjunct). Association for Computing Machinery, New York, NY, USA, 332–335. <https://doi.org/10.1145/3341162.3343802>
- [11] Kelsey Piper. 2020. The UK used a formula to predict students' scores for canceled exams. Guess who did well. <https://www.vox.com/future-perfect/2020/8/22/21374872/uk-united-kingdom-formula-predict-student-test-scores-exams>. Accessed: 2021-04-01.
- [12] J. Singh, J. Cobbe, and C. Norval. 2019. Decision Provenance: Harnessing Data Flow for Accountable Systems. *IEEE Access* 7 (2019), 6562–6574. <https://doi.org/10.1109/ACCESS.2018.2887201>
- [13] Cynthia Stohl, Michael Stohl, and Paul M Leonardi. 2016. Digital age managing opacity: Information visibility and the paradox of transparency in the digital age. *International Journal of Communication* 10, 2016 (2016), 123–137.
- [14] The European Commission Independent High-Level Expert Group on Artificial Intelligence. 2019. Ethics Guidelines for Trustworthy AI. https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=60419.