

An Ethics Framework for Research into Heterogeneous Systems

*J Happa**, *JRC Nurse**, *M Goldsmith**, *S Creese**, *R Williams †*

**Department of Computer Science, University of Oxford, †Pembroke College, University of Oxford*

Keywords: Research, Ethics, IoT, Heterogeneous Systems.

Abstract

Heterogeneous systems often found in the Internet of Things (IoT) have a wide range of challenges in ethics and law. Any device with an IP address can potentially collect, process, store and share data and make automated decisions in unpredictable ways. When conducting research and development in IoT, it is necessary to have a comprehensive socio-technical understanding of decision-making and data-handling, as well as procedures in place to pre-empt and address unforeseen consequences. In this paper we propose a comprehensive conceptual-modelling approach to help researchers systematically identify, consider and respond to challenges in ethics and law when conducting research and development of heterogeneous systems. Our framework is a six-layered model that addresses these concerns with regards to proximity to the data and actions in question. Using our framework, researchers should be able to deliver use-case scenarios that should be peer-reviewed by a large number of experts in dissimilar domains with the aim of identifying issues to why the research and development proposed is not done responsibly, so researchers can address these concerns. Finally, we explore a IoT use-case scenario, and we propose future directions for this work.

1 Introduction

Real-world heterogeneous systems, such as embedded systems or general computing systems often found in the Internet of Things (IoT) has a wide range of challenges in ethics and law. These concerns are related to the novel capabilities in data-handling and decision-making often found in research and development of new ideas. IoT brings the promise that smart devices will be interconnected, giving rise to the potential of enhanced services and ease of use where previously this was impossible or at least very difficult to achieve. Connecting a plethora of devices together will not only increase the attack surface of the system, but without ethical and law considerations, sensitive and personal data may leak and breach confidentiality. There is also the reality that devices may make unpredictable decisions due to lack of policy or programming foresight. While work is being conducted towards standards and certification in the IoT space [1], we argue it is also important to develop a methodology to help researchers identify which ethics and law concerns may emerge during the research and development phases of new ideas in the IoT space.

Devices with an IP address can often, in principle, collect, process, store and share data in unpredictable ways or make unpredictable automated decisions. It is therefore necessary for researchers to have a comprehensive socio-technical understanding of how research and development can lead to challenges in ethics and law. This is particularly the case for the development of any novel tools or techniques to be deployed in IoT infrastructures, but also in the research itself (data collection, processing, analysis and publication). We make the assumption it is impossible to identify all ethical and law issues that can emerge from an IoT project, but by aggressively peer-reviewing the conceivable data-handling and decision-making, we may tackle ethical and law concerns proactively or reactively in more well-informed ways.

The information and communications-technology space is evolving. Ethical guidelines and legislation have to be continually updated as a result. These update according to real-world practices, examples and consensus to determine acceptable behaviour in society. This means it is necessary to be able to predict how research and development will play out in laboratory conditions as well as in the real world, which is not always a straightforward task. Being able to identify whether there is the potential for, for instance, social, financial, political or privacy consequences of investigating or deploying new methods or tools would be hugely valuable to any research project, particularly prior to any ethics review or before deploying or publishing any research output.

Today, it may not be sufficient for researchers alone to review ethical concerns in research projects involving interconnected machines such as IoT infrastructures. We postulate that in order to have a sound understanding of challenges in ethics and law, and how to address them, it is necessary to include as many stakeholders as possible to peer-review IoT projects. These would consider data-handling and automated decision-making concerns, specifically looking into:

- **Variety** of groups – include experts of widely different backgrounds to examine challenges in ethics and law (ethicists, policy makers, data protection authority, businesses etc.), giving rise to a comprehensive view of potential concerns.
- **Proximity** to data – include experts from different layers of proximity to the project, those who have a direct vested interest in the research output as well as those who are close to the project itself (researchers, controllers and processors) and related communities (e.g. researchers who are not involved in the research project in question, but may provide meaningful insight).
- **Approach** type – both proactive and reactive efforts will be necessary to combat potential concerns.

1.1 Paper Contributions

In this paper we propose a novel ethics framework for research, development and deployment of heterogeneous systems that may not yet be fully defined and understood. The purpose of our framework is to deliver a conceptual-modelling approach to help researchers systematically consider challenges in ethics and law when conducting research into (not well-understood) heterogeneous systems. Our framework is a six-layered model that addresses ethical concerns. Using our framework, researchers can deliver use-case scenarios to be peer-reviewed by a large number of different experts to conduct research more responsibly.

The remainder of the paper covers the following. Section 2 outlines the background and related work in the topic. Section 3 presents the framework itself, first with an overview, then detailing each component of the framework itself. Section 4 presents examples use cases. Section 5 discusses the benefits, disadvantages and future work with regards to the framework. Finally, Section 6 concludes this paper.

2 Background

There are risks involved with each device joining the IoT, but also when novel ideas are researched, developed and tested. The European Union (EU) for instance has taken a proactive approach to address ethical and privacy concerns in the business and research sectors which also affect IoT with critical ideas such as the “right to be forgotten”¹. The EU has set up risk-assessment procedures with a wide variety of stakeholders (including the Ethics Subgroup IoT) [2].

Furthermore, any research funded by the EU has to follow its rules and guidelines, such as the Charter of Fundamental Rights of the European Union², the European Convention on Human Rights³, the European Code of Conduct for Research Integrity⁴. If a research project involves personal data from Data Subjects (DS), it is also required to be General Data Protection Regulation (GDPR)⁵ compliant as of May 2018, with an exception being the Directive on security of Network and Information Systems (NIS Directive)⁶.

Any method or tool that handles personal data (as defined by the EU) to be used in the EU is required to deal with data-handling according to the GDPR. Additional procedures may involve ethical reviews, especially if the research at any point includes handling of personal data or automated decision-making that are real-world actions. Depending on the scope of the ethical concerns, the local Data Protection Authority (DPA) may also need to be informed.

¹ See article 17 in the General Data Protection Regulation (GDPR).

² http://ec.europa.eu/justice/fundamental-rights/charter/index_en.htm

³ http://www.echr.coe.int/Documents/Convention_ENG.pdf

⁴ https://ec.europa.eu/research/participants/data/ref/h2020/other/hi/h2020-ethics_code-of-conduct_en.pdf

⁵ <https://www.eugdpr.org/>

⁶ <https://ec.europa.eu/digital-single-market/en/network-and-information-security-nis-directive>

Van den Hoven [2] and Wachtel [3] compiled a summary of outputs of the IoT Expert group, including an Ethics IoT Subgroup⁷, outlining cybersecurity strategies, ethics, privacy, architectures, standards, identification and governance concerns, and present their findings and recommendations.

Key ethical issues discussed arising from its development and deployment include:

1. **ubiquity and pervasiveness** – there is no clear path to partially opt out of a fully-fledged IoT system.
2. **miniaturization and invisibility** – devices take many forms, which calls for measures to keep technology visible and amenable to inspection, audit, quality control and accountability.
3. **ambiguity and ontology** – it will be necessary to deal with unclear criteria of identity and IoT system boundaries as the distinctions between objects, artefacts and human beings blur together.
4. **identification** – what should the rules be for assigning, administering and managing IoT identities? This may become important as more and more seemingly insignificant objects are assigned identities.
5. **connectivity, distributed systems and big data** – there is a high degree of connectivity, often distributed, generating large quantities of data (and subsequent data transfers) between objects and persons in networks.
6. **mediation and autonomous agency** – IoT provides ways of extending and augmenting human agency, even to the point that it may exhibit artificial and spontaneous and emerging agency. How do we ensure appropriate data-handling and decision-making, especially keeping artificial intelligence both ethical and abiding by the law?
7. **embedded intelligence and extended mind** – smart objects may embed intelligence/knowledge function as tools become external extension to the human body. What does this mean for everyday use, and at what stage is having embedded extensions an unethical advantage or disadvantage in society?
8. **seamless transfers invoking unpredictability and uncertainty** – information flow may become not visible to the end-user, which may raise concerns about how people understand how their data is handled.

A key theme in the aforementioned issues relates to how people are simply not accustomed to new features in new and emerging technologies. Societal norms and attitudes to these capabilities have not been set. Real-world historical examples are needed to set guidelines. Currently, there is little empirical evidence to build a foundation on. Furthermore, the evolution of digital systems is making it increasingly difficult for all stakeholders to discuss cyber-related topics succinctly and proficiently. Politicians, ethicists, lawyers and business owners may lack the necessary background to describe or understand technical concerns, while cybersecurity analysts may not be able to identify societal concerns from their perspective alone [4].

⁷ Three groups were involved in their consultation: The European Group on Ethics, The ETICA project, and The Expert Group on Responsible Research and Innovation.

With domain-specific vocabulary having subtly different interpretation across subjects, it is difficult to get people engaged in describing cyber-concerns in a way that is unambiguous and easily understood by all experts.

Traditionally, devices that join a network are mainly subject to risks related to security vulnerabilities, as Confidentiality, Integrity and Availability (CIA) of systems can be compromised. While risk assessments are even more of a challenge in the IoT [5], ensuring people's safety is a critical issue given some application scenarios of IoT devices. In healthcare for instance, hacked connected devices could have the potential to be harmful to our health (or at worst be fatal), or programming bugs could lead to such devices misbehaving and compromising their user's safety. More recently, the EU is looking to prepare a "right to repair" legislation to combat short product lifespans (planned obsolescence). Whether the EU will add cybersecurity patching as part of this legislation remains to be seen.

Leverett et al. [1] present an in-depth discussion on the topic of standardisation and certification in IoT, pointing out concerns about safety, security, liability, transparency and privacy principles. They identify missing institutional resources and suggest a strategy for filling the gap. Specifically, they believe cybersecurity regulators will have to: ascertain, agree, and harmonise protection goals, set standards (whether these be policies or protocols), certify standards achievement, enforcing compliance, reduce vulnerabilities, reduce compromises, and reduce system externalities. Such changes will affect engineers in testing facilities, but also regulator committees, with a focus on sustainability of software and the necessary means to support it. Mobile phones today, for instance, are considered to have a shelf-life of 2-3 years. We expect companies to be able to provide security support throughout that timeframe. However, some IoT devices (e.g. cars with smart apps or self-driving vehicles) can have a significantly longer lifespan.

Leverett et al. propose the notion of cyber-covigilance with the creation of a European Safety and Security Engineering Agency to provide a shared resource for policymakers and regulators. Specifically, they suggest its mission should be to support the European Commission's (EC) policy work where technical security or cryptography issues are relevant; support sectoral regulators in the EU institutions and at the member state level; develop cross-sectoral policy and standards, for example arising from system integration; act as a clearing house for data from post-market surveillance and academic studies; work to promote best practice and harmonisation; and act as a counterweight to the national-security orientation of member-state security authorities.

Weber [6] highlight new security and privacy challenges that IoT bring, and argues that new regulatory approaches are necessary to ensure privacy and security become mandatory in IoT. In particular, he mentions that attacks have to be intercepted, data authenticated, access controlled and the

guaranteed privacy of customers. The nature of the IoT asks for a legal framework that adequately takes into account the globality, verticality, ubiquity and technicity of the IoT.

Van Kranenburg and Bassi [7] claim that broader challenges posed by the IoT cannot be managed with the current policy tools and research programs. These challenges are:

- **global cooperation and standards** giving examples of how different nations have different priorities for the IoT's future;
- **new business models** and new currencies;
- ethics, control society, surveillance, consent and data driven life, outlining how Privacy Enhancing Technologies (PETs) are a partial solution and points out how the EU's legislations are likely to become a subset of ethics thinking through "**ethics by design**" in real-world systems; and
- technological challenges driven by the **need to save energy**.

Baldini et al. [8] present an approach for users' interaction with the IoT, implemented through a policy-based framework. Specifically, users are provided with wider controls over personal data or the IoT services by selecting specific sets of policies, which can be tailored according to users' capabilities and contexts in which they operate. They also highlight several challenges and processes for ethical design in the IoT space. The challenges include:

- **economic incentives for data protection** of the DS are limited to the businesses creating the IoT applications and devices;
- how the DS has often **incomplete information about the consequences of disclosing data** either voluntarily (e.g., providing data) or involuntarily (e.g., collection of position information);
- the complete set of **information necessary to make a rational choice** with reference to data-handling could be so large that the DS may not be able to access the IoT service in an effective way;
- **psychological biases** affect the perception of immediate benefits and can fail to recognise impact the long-term negative impact (e.g., risk to users' privacy);
- **tensions** between businesses needs to collect and process data and rights to privacy;
- **cost** of implementing privacy enhancing or data protection solutions;
- **accountability** of the IoT applications regarding users' privacy;
- separating online from offline information and their **linkage can generate privacy breaches**;
- depending on level of technical proficiency, the DS can have different levels of **perceptions of risks**;
- ability and agility to **conform to regulatory frameworks**; and
- **context changes the uses of IoT** services and devices.

Moreover, the processes mentioned by Baldini et al. include:

- **understanding the need for and value of trust** in society at the level of public and private stakeholders;

- **translating** these needs and values into ethical design;
- **demonstrating** that these needs and values are taken into account; and
- **establishing a clear framework for transparency and accountability**. While existing frameworks have previously not been developed with IoT in mind, it will have to adopt to these emerging technologies.

2.1 Engineering Frameworks

Existing engineering frameworks have a focus on preserving privacy through engineering reference models and principles, and do not strictly speaking address challenges in ethics and law. These principles behind privacy preserving mechanisms are a first step towards aid in ethics and law challenges in engineering.

The Ensuring Consent and Revocation (EnCoRe) [9] project proposed several approaches to formalise rulesets to ensure consent and revocation of personal data by service providers.

The National Institute of Standards and Technology (NIST) [10] discusses the concepts of privacy engineering and risk management for federal systems. Their work aims to establish the basis for a common vocabulary to facilitate better understanding and communication of privacy risk within federal systems, and the implementation of privacy principles. It introduces two key components to support the application of privacy engineering and risk management: privacy engineering objectives and a privacy risk model.

The MITRE Corporation’s Privacy Engineering Framework [11] outlines how privacy engineering activities map to stages of the classic systems engineering life cycle. A mapping exists for every systems engineering cycle, including agile development, since every life cycle includes core activities in some form.

The Privacy Management Reference Model (PMRM) [12] provides a methodology for understanding and analysing privacy policies and their privacy management requirements in defined use cases, as well as selecting the technical services to support privacy controls. It is relevant for use cases in which personal information flows across regulatory, policy, jurisdictional, and system boundaries.

Fisk et al. [13] define three engineering privacy principles that guide sharing security information across organisations: Least Disclosure, Qualitative Evaluation, and Forward Progress. They break down these principles to concept, implementation, consequences for ignoring and design approaches to achieving implementation. They then discussed how these principles then apply to reduce risks of the data exposure and help manage trust requirements for data sharing.

3 Framework

3.1 Overview

Translating real-world concepts such as the law and ethical principles into computational rulesets is not a straightforward task, however, mainly because the law and ethics may have properties that are subjective in the real world. Computational systems are unlikely to be able to mirror human-level decision-making. As such it may not be possible to develop systems that deterministically compute decisions that are compatible with human reasoning in every circumstance. The use of the term ‘reasonable’ for instance is a term that may apply differently to different situations. Being able to develop a system capable of applying this term correctly in every circumstance in a legal setting would be a non-trivial (highly impossible) task.

Having said this, it is worth pointing out that in law, reasonableness is often relative to what other entities in the same situation would have done. In tort law for instance, it is a defence to show that other doctors would have done the same thing, and in public law we ask whether something was so unreasonable that *no reasonable decision-maker* could have arrived at that standard. So industry standards and norms are going to be relevant here and if one company can do something the question will then arise why other companies did not adopt the same safeguards.

Our framework aims to deliver a comprehensive conceptual reasoning and reference model to help researchers and developers systematically identify, consider and respond to challenges in ethics and law when conducting research and development of heterogeneous systems. The approach assumes it is necessary to consider challenges in ethics and law as seen from as many domain-expert perspectives as possible with an aggressive peer-review approach, aiming to find problems, and assuming that no research and development project will be perfectly capable of resolving ethics and law challenges on their own.

If no real data is available (because the system or algorithm is currently being implemented), it is all the more important that the use-case scenarios contain assumptions that carry an appropriate degree of verisimilitude. Different experts and community guidelines should enforce rigour in the project.

Specifically, during research and development, it is necessary for the various layers to consider how realistically use-case scenarios predict data handling, as well as provide decision-making use-cases. The objective of the framework is to identify concerns through use-case scenarios, but more importantly, also to identify how to address these concerns using a pessimistic and antagonistic approach.

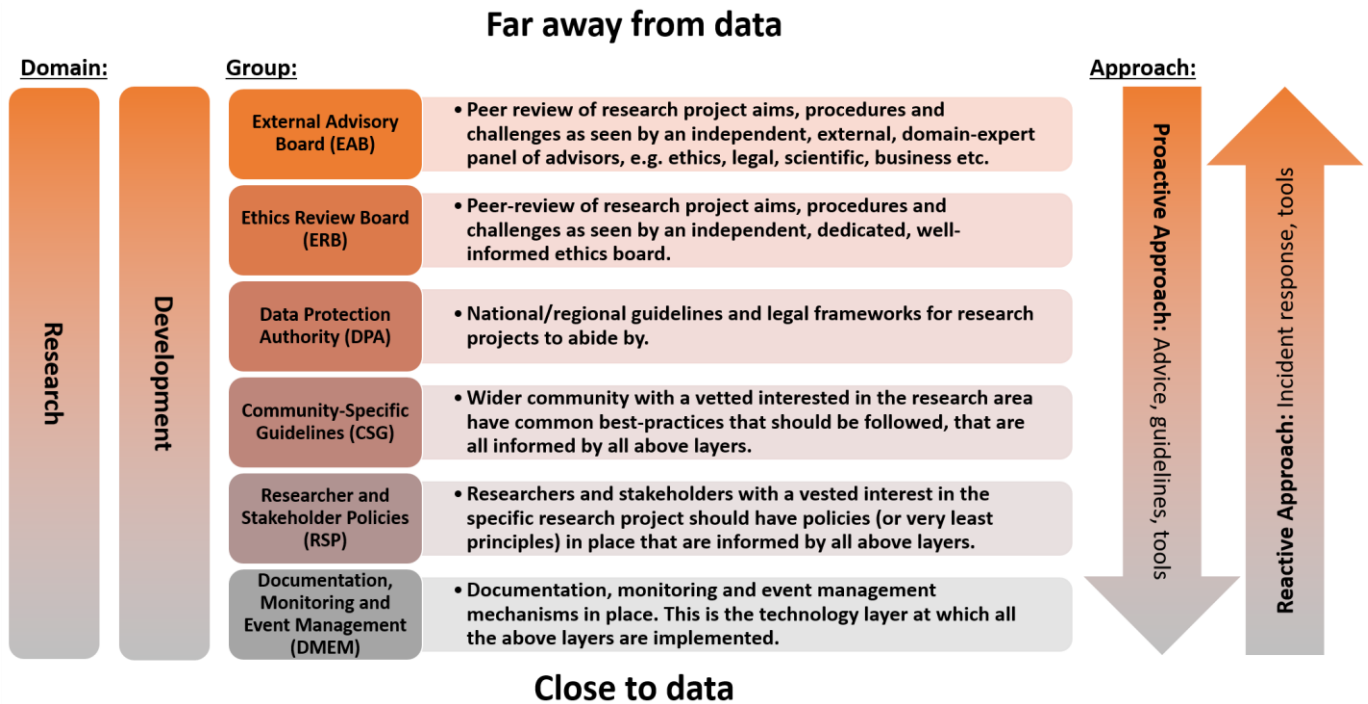


Figure 1: a summary of the ethics framework for conducting research in heterogeneous systems.

Through the peer-review at each of the different layers, in both research and development (which also includes deployment, i.e. not simply deploying IoT in laboratory conditions, but outside laboratory settings), the various experts should be able to identify concerns and propose solutions as seen from their perspective.

It is assumed that no expert in isolation is capable of identifying all challenges in ethics and law in the research and development of new technology or algorithm. Rather, a collection of experts from different backgrounds is likely to be able to provide more well-informed insight than the researchers, as well as any research ethics committee alone.

The framework should provoke researchers into asking meaningful questions such as: if we collect, process, store or share data (e.g., pertaining to a sensitive IoT device) in unpredictable ways – what consequences are likely to follow, and can we use our added insight to improve our methods?

The core of the framework proposes which **Domains** to consider, **Group** who should peer-review the data and data-handling in question, and the **Approach** (type of effort) involved. A six-layer representation of groups is assumed to cover anyone who may have any (even remotely or indirect) vested interest in the outcome of the project. The model assumes both top-down and bottom-up reasoning approach to identify and address possible concerns. We do so by considering the proximity to the data and decision-making in question, see Figure 1. We begin first by describing the main domains.

3.2 Domain

There are two key domains which we consider in our model, research and development. In both domains, it will be necessary to propose a **data-management plan for how to collect, process, store and share data** that is generated and handled throughout the project as well as **how to handle automated decision-making capabilities**. Specifically, we can consider the data that is generated by researchers from potential end-users as well as machines:

- **End-user data.** Researchers may conduct surveys, questionnaires, interviews, focus groups, usability studies and collect Key Performance Indicators (KPIs) about user activities, actions, etc. **During research:** we have to consider what insight we might be learnt from end-user data as well as how this might impact the handling of that data. **During development:** we have to consider how the research feeds into the implementation, and whether this has any cause for concern in terms of how research output is applied (e.g. into tool, method and algorithm implementations).
- **Machine data.** Researchers use machine data (e.g. logs or statistics about an IoT device or a connected machine's behaviour) to answer research questions, validate any hypothesis or use the machine data to refine the existing system's capabilities and functionalities (development). Case studies should help facilitate the modelling of expected and possibly unpredicted behaviour in the ecosystem. It should be possible to then test those expected behaviours with how closely they match reality.

We continue by describing which groups of people should peer-review the project data-management plan.

3.3 Group

The six layers of groups to peer-review the project approach can be summarised as (starting with farthest away from the data and decision-making in question).

3.3.1 External Advisory Board (EAB)

The board is a group of people who come from different backgrounds outside of the project, with no direct interest in the project, who are not involved in the types of communities in which the research and development takes place. They should, however, still have some broad interest in the outputs of the project, and are therefore willing to peer-review the approach of the project. Peer-review should be of research project aims, procedures and challenges as seen by an independent, external, domain-expert panel of advisors, e.g. ethics, law, scientific, business etc.

3.3.2 Ethics Review Board (ERB)

The ERB is a board or ethics committee, for instance those found at universities that aim to peer-review of research project aims, procedures, methodologies and challenges as seen by an independent, dedicated, well-informed ethics board. The purpose of this committee is to be able to put this project in the context of other projects and be able to provide some insight as to challenges that may arise by having insight into challenges that have emerged in other projects historically.

3.3.3 Data Protection Authority (DPA)

National/regional guidelines and legal frameworks for research projects to abide by help inform the project about which aspects are clear violations, but this group need to be contacted in case of confidentiality breaches or when seeking data protection advice.

3.3.4 Community-Specific Guidelines (CSG)

The wider community with a vetted interested in the research area have common best-practices that should be followed, that are all informed by all above layers. The communities themselves can be consulted for clarification and peer-review of the project's own aims and goals. In the context of public CSIRT research, an example of community-specific guidelines would be Trusted Introducer Service⁸ (TIS) in which forms the trusted backbone of infrastructure services and serves as clearinghouse for all security and incident response teams. TIS lists well-known teams and accredits and certify teams according to their demonstrated and checked level of maturity.

3.3.5 Researcher and Stakeholder Policies (RSP)

Researchers and stakeholders with a directly vested interest in the specific research project should have policies/principles in place that are informed by all above layers and affect those below. Ideally, these policies should take form of a formal Service Level Agreements (SLA) or some agreement akin to an End-User License Agreement (EULA). Such agreements

should also include use-cases of what to do in the event of plausible worst-case scenarios in terms of data-handling breaches and automated decision-making. These policies should be peer-reviewed by all parties directly affected.

3.3.6 Documentation, Monitoring and Event Management (DMEM)

Manual and automated documentation, monitoring and event-management mechanisms should be in place where possible. Specifically, research should contemplate means and mechanisms to collect information about how data is processed and decisions are made, e.g. through permission violation checkers and taking time to identify appropriate logging of machine and end-user activities for safety and security reasons. We envisage this being the implementation of all the above layers. As such, any high-level decisions made about data-handling and decision-making in IoT behaviour will need to trickle down to the implementation layer.

3.4 Approach

The arrows in Figure 1 indicate directionality of proactive and reactive approaches, and how these approaches should affect any future project decision-making. These concerns are considered at each layer, but feed to all subsequent layers above and below the group they were considered in. Those that feed insight downward are likely to be **proactive** approaches, meanwhile those that feed insight upwards to the upper layers are likely to be **reactive** approaches.

Proactive approaches are responses to seeking advice on ethics and law in advance of any incident having happened that relates to concerns in question in the project: for instance project policy changes or having to change data-handling practices to make the project abide by the law and remain ethically compliant. Proactive approaches aim to help better inform the project participants about likely emerging issues. The plausibility of any emerging issue should also be under scrutiny, which is why a peer-reviewed approach can help exclude inappropriate or otherwise implausible use cases.

Reactive approaches are response to an incident having already happened. Examples include data confidentiality breaches, attacks that have safety consequences for people, devices or infrastructure, etc. There exist, for instance, technological solutions such as intrusion-detection techniques that can be deployed and respond to attacks and limit harm; manual incident handling by cybersecurity researchers or analysts may also be necessary to contact affected parties and mitigate any situation. The consequence of any past incident should help inform proactive approaches with the aim to prevent such an incident from happening again, or at the very least severely limit future similar incidents. The lesson learnt may also affect how other groups above the technological layer (bottom layer) consider similar concern in the future.

⁸ <https://www.trusted-introducer.org/>

4 The IoT Use Case

In an IoT ecosystem, data-handling can become a major issue. Existing literature in Section 2 has hinted that miniaturisation, ubiquity, connectivity and agency can become major concerns for end-users. We argue that these aspects may disempower end-users' ability to understanding of how data is handled, especially if this information flow is communicated poorly. Attack surfaces may increase as a consequence on negligent consideration for interaction of different types of IoT devices. In many cases, we may have to rely manufacturer's ability to communicate data handling as well put our trust in their secure programming and data-handling philosophy (often all or nothing opt in/out) approach is safe, secure and preserves an end-user's privacy. The legality and ethical behaviour of such a project may come into question, esp. if researchers put these capabilities in the hands of third parties.

A simple example may for instance be a piece of wearable computing hardware that tracks GPS coordinates, obtain local weather and traffic data, and that is able to connect to a smart phone and a smart home in the interest of being able to send information back and forth to the home and phone such as: house temperature and a list of food and groceries that are in the fridge (to help users keep track of which items need restocking), while connecting to the end-users social media accounts. All of the information can be stored in the cloud as backup, but each IoT device manufacturer may wish to provide their own cloud service, while social media platforms may wish to publish this data automatically. There are several concerns with this scenario, esp. if researchers wish to trial new ideas in an already existing IoT ecosystem. From a technical standpoint: interaction between multiple IoT devices may yield unforeseen consequences, such as unintended data leaks. Connectivity may also significantly increase the attack surface of the end-user's IoT ecosystem, esp. if new and untested ideas are introduced into the environment.

Manufacturers are unlikely to be able to function and system test the interaction between multiple IoT devices and how they interact with each other. If an attacker is able to gain privileged access on any of the devices, they may be able to gain privileged access to connected devices. In several cases, banking apps, health data or other personal data may be stored on the phone, which may be within reach to an attacker. Also important to consider is, what information is share between apps, and how to assure that the information is stored securely in the way that end-users easily grasp.

This is particularly important as manufacturers, social media accounts and attackers can correlate different data sources and obtain meta-data insight about users and sell this information on to third parties without users knowing. From a research and development perspective, this is unlikely to happen in laboratory conditions, however, with mass deployment of IoT ecosystems, scenarios like this may be worth exploring as while they are in many cases improbable, they are not infeasible.

5 Discussion

As previously mentioned, translating real-world concepts such as law constraints and ethical principles into computational rulesets is non-trivial task. Unlike existing literature in this space that aim at highlighting where issues in IoT may rise more broadly speaking, we have specifically focused on delivering a systematic methodology to help researchers better ensure that their work is sound, and at the very least be able to demonstrate reasonable efforts have been made to ensure that the project does its best to protect users, their data in a safe and secure environment during research and development.

We believe the key benefits of this approach are:

- The framework is **an aggressive approach to identify and address challenges in ethics and law** that may rise from any IoT project. Specifically, the framework should be used in the effort to identify corner cases. Unlike existing methods, our approach aims to help researchers find failures in their existing approaches to ethics and the law in their projects as early on so these can be addressed appropriately. We assume that no project scope, data and decision-making handling will ever be perfect (in the eyes of ethics and the law), and that at some point the project is bound to not have thought through some aspect related to the law or ethics. Our aim is to empower researchers by preparing for this eventuality with a framework that anticipate failures and aims to learn from them in order to make more well-informed decisions about possible solutions through peer-review.
- **Peer-review is likely to identify concerns** that researchers and their ethics committees alone are unlikely to be able to find and address on their own.

We believe that a key limitation of our approach is that it is very conceptual as it stands, and that each of the layers will have to be tailored for each individual project: this open up the possibility of each research project using the framework differently, making it difficult to ensure that projects are applying the framework appropriately. Our approach aims to be generic enough to match any research and development project, at the cost of specificity.

The approach is currently being used in the PROTECTIVE⁹ project, a research and development project in cyber-threat intelligence sharing among public Computer Security Incident Response Teams (CSIRTs), with the project asking questions related to the challenges in ethics and law of cyber-threat intelligence sharing in a National Education Research Network (NREN) space. In the project we are identifying instances in which cyber threat intelligence may also include personal data, and what efforts are necessary to identify and anonymise personal data about to be shared between NRENs as well as identifying sharing of data that may otherwise breach the GDPR or other NDAs. We are in the process of identifying what automated actions the tool can conduct to

⁹ <https://protective-h2020.eu/>

provide assurance that the cyber threat intelligence about to be shared will be compliant with the GDPR and NDAs. In the future, we envisage similar data sharing capabilities possible in IoT ecosystems.

Future work should look to explore other use cases of our framework and investigate how it can be expanded and improved. We assume that templates for lessons learnt could be developed to decrease the time necessary to identify ethical issues. The more real-world use-cases and projects that we apply this approach in may help us towards validation that this approach is the most appropriate methodology to employ in identifying and addressing challenges in ethics and law in IoT and other heterogeneous system projects. In the IoT space, we envisage cloud and fog computing and home IoT infrastructures can aid in the cyber-threat intelligence-generation and processing stages. The case studies presented are with extension of the PROTECTIVE project in mind, on a larger scale. We envisage the miniaturisation and invisibility of computer devices, and believe our approach will aid in the development of the inspection, audit, quality control and accountability of devices and their developers, esp. in a world where new data sources and new devices (with previously unpredictable new capabilities in agency and autonomy in decision-making and data-handling).

6 Conclusion

Understanding concerns in ethics and law in unexplored research and development territory can be challenging in itself, especially when the research attempts to investigate novel ideas. In this paper we have proposed a novel ethics framework for research, development and deployment of heterogeneous systems that may not yet be fully understood. Our framework delivers a conceptual-modelling approach to help researchers and developers systematically consider challenges in ethics and law when conducting research, development and deployment in IoT systems.

Our approach aims to help researchers identify and best address most challenges in ethics and law. The framework is a six-layered model that addresses ethical concerns with regards to proximity to the data in question. We propose that researchers can use our framework to deliver use-case scenarios to be peer-reviewed by a large number of different experts. We facilitate the understanding of who to approach for peer-review as well as approaches in dealing with challenges in ethics and law. We also explore IoT use-case scenarios, and propose future directions for this work.

Acknowledgements

This research was conducted as a part of the PROTECTIVE project. This project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No. 700071. This output reflects the views only of the author(s), and the European Union cannot be held responsible for any use which may be made of the information contained therein.

References

- [1] É. Leverett, R. Clayton and R. Anderson, "Standardisation and Certification of the 'Internet of Things'," Technical Report, 2017.
- [2] J. v. d. Hoven, "Fact sheet- Ethics Subgroup IoT - Version 4.0," Technical Report, 2012.
- [3] T. Wachtel, "10th Meeting of the Internet of Things Expert Group," Technical Report, 2012.
- [4] J. Happa and G. Fairclough, "A Model to Facilitate Discussions About Cyber Attacks," in *Ethics and Policies for Cyber Operations*, Springer International Publishing, 2016, pp. 169-185.
- [5] J. R. C. Nurse, S. Creese and D. De Roure, "Security risk assessment in Internet of Things systems," *IEEE IT Professional*, vol. 19, no. 5, pp. 20-26, 2017.
- [6] R. H. Weber, "Internet of Things - New security and privacy challenges," *Computer Law & Security Review*, vol. 26, no. 1, pp. 23-30, 2010.
- [7] R. van Kranenburg and A. Bassi, "IoT Challenges," *Communications in Mobile Computing*, vol. 1, no. 9, 2012.
- [8] G. Baldini, M. Botterman, R. Neisse and M. Tallacchini, "Ethical design in the internet of things.," *Science and engineering ethics*, pp. 1-21, 2016.
- [9] I. Agrafiotis, S. Creese, M. Goldsmith and N. Papanikolaou, "Applying formal methods to detect and resolve ambiguities in privacy requirements.," in *IFIP PrimeLife International Summer School on Privacy and Identity Management for Life*, 2010.
- [10] S. Brooks, M. Garcia, N. Lefkowitz, S. Lightman and E. Nadeau, "An Introduction to Privacy Engineering and Risk Management in Federal Systems," National Institute of Standards and Technology Internal Report, 2017.
- [11] S. Shapiro, N. Washington, K. Miller, J. Snyder and J. McEwen, "Privacy Engineering Framework," MITRE Corporation Technical Report, 2014.
- [12] A. Cavoukian, D. Jutla, F. Carter, J. Sabo, F. Dawson, J. Fox and S. Fieten, "Privacy Management Reference Model," Organization for the Advancement of Structured Information Standards (OASIS).
- [13] G. Fisk, C. Ardi, N. Pickett, J. Heidemann, M. Fisk and C. Papadopoulos, "Privacy principles for sharing cyber security data," in *IEEE Security and Privacy Workshop*, 2015.