Accountability in the Internet of Things (IoT): Systems, law & ways forward

Jatinder Singh*♦, Christopher Millard+, Chris Reed+, Jennifer Cobbe*, Jon Crowcroft*

*Dept. of Computer Science & Technology (Computer Laboratory), University of Cambridge
+Centre for Commercial Law Studies, Queen Mary University of London

Abstract

Accountability is key to realising the full potential of the IoT. This is for reasons of adoption and public acceptability, and to ensure that the technologies deployed are, and remain, appropriate and fit for purpose. Though technology generally is subject to increasing legal and regulatory attention, the physical, pervasive and autonomous nature of the IoT raises specific accountability challenges; for instance, relating to safety and security, privacy and surveillance, and general questions of governance and responsibility. This article considers the emerging 'systems of systems' nature of the IoT, giving the broad legal context for these concerns, to indicate technical directions and opportunities for improving levels of accountability regarding technologies that will increasingly underpin and pervade society.

1. Introduction

The Internet of Things (IoT) is often grandly conceived as a highly connected environment in which the physical and digital worlds blend, transforming our homes, hospitals, streets, businesses, and cities.

We are at the early stages of any so-called "IoT revolution"—we are currently closer to an "Internet of Silos," as the technology tends to operate solely within a particular deployment (for example, an organization or house) or industry vertical. However, there are significant efforts underway toward realizing more holistic IoT projects, such as smart cities that seamlessly support a vast range of applications driven by the promise of new services, efficiencies, and trillions in value.[1]

Despite the technological advances making ubiquitous/pervasive computing a reality, many of the potential benefits will not be realised unless people are comfortable with and embrace the technologies. Accountability is crucial for trust, as it relates to the responsibilities, incentives, and means for recourse regarding those building, deploying, managing, and using IoT systems and services.

[♦] Contact: jatinder.singh@cl.cam.ac.uk

However, the nature of this broad IoT vision—a system of systems—poses particular challenges,[1,2] including those relating to **privacy and surveillance** given a highly instrumented, sensor-rich world; **safety and security** given the reliance on, and physical effects of, the technologies; and **governance and responsibility** given the scale of these complex, multi-vendor/stakeholder environments.

Such concerns are *sociotechnical* in nature, and both law and technology have clear roles to play. Law and governance frameworks provide a basis for establishing rights, liability, responsibilities, and mechanisms for compensation and holding entities to account. At the same time, technology can assist with accountability, providing means that facilitate the control and auditing of IoT technologies. Given the complexity and interdependence of the technical and legal issues relevant for IoT accountability, moving forward requires closer collaboration among the involved disciplines.

Toward this end, this article focuses on the systems-of-systems nature of the emerging IoT, providing legal context for key accountability challenges and technical opportunities for improving accountability in the IoT as it continues to develop. In doing so, we raise high-level issues for consideration and highlight ways for more collaborative IoT development and technical/legal research.

2. The nature of the IoT

At its core the IoT involves bringing the physical world online through sensors, which perceive aspects of the physical world, and actuators, which deliver physical effects.

The ISO/IEC define the IoT as “an infrastructure of interconnected objects, people, systems and information resources together with intelligent services to allow them to process information of the physical and the virtual world and react.”[3] In line with this definition, the focus of ubiquitous/pervasive computing is integrating sensors and actuators into large-scale (potentially global) systems environments—comprising networks, clouds, and so on—capable of reacting appropriately to context to deliver wide-ranging functionality and efficiencies. This entails physicality, scale, automation, and seamlessness. Data is the driver: sensors produce data representing aspects of physical environments; actuators respond to commands (also data) resulting in physical effects; and data transfer, storage, aggregation, and processing operations help realise desired functionality. It is the exchange of information—the flow of data—that determines what happens in the IoT.

As the Figure shows, in practice the IoT constitutes a system of systems, where functionality is realised via processing chains through which data is exchanged.[4] These chains consist of a series of components, each potentially a system itself, that interact to bring about functionality.[2] Components include physical devices (sensors, actuators), gateways (hubs, phones), clouds (public/private/hybrid/fog/edge), and software/services/agents for data access, processing, analytics, adaptation, and automation.

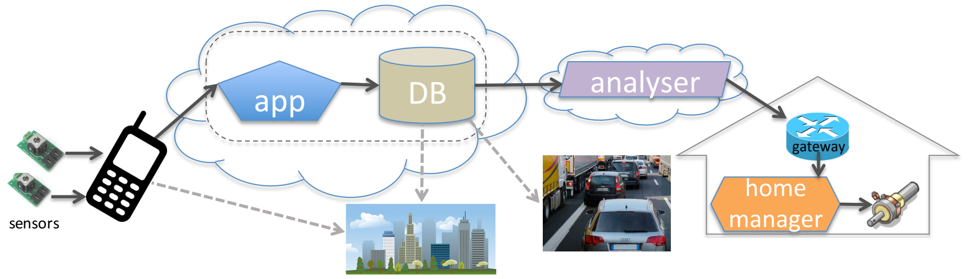

IoT as a system-of-systems: An example IoT processing chain whereby sensors (e.g. location and temperature) on a mobile device trigger a heating system in the home. The functionality results from an orchestration of a range of system components, *through a series of data exchanges between them*, where each component is potentially managed by a different organisation or user. Also illustrated (grey) is that data could be used/reused for different purposes – here flowing from a personal datastore to the city's transport system and other services.

2.1 System-of-systems aspects

An IoT systems-of-systems environment has several key characteristics relevant to accountability.

Diversity of governance. Components will be owned, managed, and operated by different people and organizations, perhaps in different geographies, with their own set of interests, incentives, responsibilities, and obligations.[5,6]

Dynamic interactions. Processing chains can form dynamically.[4] For instance, when you arrive in a new city, your wearables might automatically interact with services embedded in the surroundings. This will become common, given that customization to accommodate the preferences of individuals, organizations, or contextual/environmental concerns is central to many IoT applications. Further, the same components (and therefore their data) might be used and reused in various ways for different purposes and by different parties—for example, for example, an individual might leverage the thermometers from the heating systems of the buildings they visit to feed temperature readings to their health wearables. We can expect many instances in which individuals interact with the same environment, but each leveraging the same components for different purposes.

Such functionality enables new applications, systems, and services to emerge and be built organically and on demand; facilitates customization; and provides longevity, as components can be repurposed. However, it also brings complex accountability challenges, as those building and deploying the technologies might not envisage all their possible purposes, contexts, uses and users.

Data analytics. The ever-increasing deployment of sensors means the generation of volumes of rich and granular data, offering great potential for new insights and efficiencies from data analytics, including machine learning (ML). In a systems-of-systems context, data (including the results of analytics) from various sources will be combined and/or processed, possibly across domains of administrative control.[4,7] For instance, a retailer might seek inferences relating to product usage from data generated by a range of different sensors and systems in customers' homes.

Automation. The IoT and its supporting infrastructure increasingly reacts, adapts, and responds automatically, when and where necessary. Examples include alerting emergency services when someone falls, automatically rerouting traffic after an accident, or carrying over users' preferences regarding heat, lighting, and music as they move between their car and rooms in their home. Analytics and ML pave the way for uncovering detailed and complex representations of context, events, patterns, profiles, and preferences to drive automated responses. Although not all automation involves ML, it has a clear role to play—leveraging data to build models, and applying those models to data to take actions.

Automation entails interactions across components, to both provide data/contextual inputs and elicit responses. For instance, on detecting your entrance to an airport, various components might automatically reconfigure and adapt your devices (phone, wearables) to interact with local services to help you check in, guide you to the gate, and so on. Similarly, automated traffic management entails interactions among road infrastructure, individual vehicles, public transportation systems, and emergency services. Some actions and interactions might be predefined, others determined at runtime.

2.2 Accountability aspects

This complexity raises numerous accountability challenges, including:

Governance and responsibility. The question of who can and should be held accountable is particularly challenging for issues that arise through the composition of, and interactions between, components managed by different entities, rather than a particular entity failing to act appropriately, due to the opacity of some technology, and where components are used, reused or behave in ways that vary from the original intention or vary over time.

Privacy and surveillance. A highly instrumented world doubles as a surveillance apparatus for commercial and governmental interests. Considerations include empowering individuals regarding their personal data, transparency of data records, data transfer and usage, and entities' roles in the various data-processing chains.

Safety and security. The IoT's scale exacerbates security issues given the vast numbers of components, their possible interconnections (all potential points of failure), and the many actors/vendors involved. Moreover, actuation, or failure to actuate, could result in physical harm. Active failure prevention and risk mitigation are important, as is auditing to facilitate learning from failures.

3. Accountability: The legal dimension

The inherent complexity of the IoT as a system of systems does not easily fit traditional legal and regulatory constructs, raising many concerns beyond the subject of this article.[5,6,8–10] However, the law is also complex, multilayered, and distributed—a given law's meaning and effect is mediated by its interaction with other laws and the people impacted. This means that accountability is hard to assign until the risks of any new technology have been discovered through use. We therefore focus on the broader normative aims common to laws regulating technology, and for the IoT, two areas stand out: legal obligations and liabilities, and regulation of personal data.

3.1 Legal obligations and liability

For technical systems, the most important legal aspect of accountability is transparency about a system's workings. A technology producer who has caused loss or damage to a user is obligated to show that they acted reasonably or fairly, or faces liability.[8] This is often true even if there is a contract between the producer and the user. Although a producer is usually held liable if the technology does not perform properly irrespective of fault, most liability issues in the IoT arena will likely arise from the provisioning and processing of information and the associated (potentially physical) consequences, with the producer's primary obligations to use reasonable care and skill.

Law regulates the actions of legal entities, including individuals and companies, *not* technologies themselves. Therefore, a key focus of IoT accountability is an entity's decision whether to use a specific IoT technology, or how to assemble or integrate various components—such decisions potentially could be made at all stages of design, production, and supply. Moreover, the nature of transparency depends on the entity. For example, often IoT users are ignorant of the workings of a technology, making it a de facto black box. The users' only options are to trust the technology or not use it, as they generally cannot know in advance about potential failures. In contrast, laws generally consider what a technology producer ought to have known in advance—for example, through testing and evaluation processes—and what can subsequently be discovered in the event of failure. Users who compose their own systems from off-the-shelf components are in some sense producers (although still operating with apparent black boxes). The level of transparency required of those users clearly differs from what would and should be demanded from standard producers (entities in a processing chain). Legal requirements to incorporate transparency into IoT technologies must take differences of perspective into account.

The starting point in any legal discussion of IoT transparency must therefore be the recipient of information about the workings of the technology and the data flowing across it. In data protection law there are transparency obligations to data subjects and regulators, each with different information needs.

This is also true when assessing liability for an IoT technology failure. Often technology users' liability is based on negligence—failure to take reasonable steps to avoid foreseeable risks.[8] Users might be liable for deciding to use a technology in particular circumstances, but this will depend on what they ought to have known about the risks. If potential liability arises from how a component, or combination of components, was used, or from choices made by users when composing their own IoT system, again the question of foreseeability arises. We might thus expect technology users to demand enough information from technology producers to make this assessment.

However, other than specific data-protection requirements, generally technology producers are not currently legally obliged to explain how the technology works. Requiring those who produce IoT technologies to provide such technical transparency would be legally novel. Technology is gradually supplanting human decision making in many fields, such as vehicle automation. If IoT users' liability is based on their knowledge of the technological risks, many losses are likely to go uncompensated as the cause is no longer user negligence in any meaningful sense. Conversely, liability might more likely arise where various IoT components are combined or deployed in a manner that contravenes clear, and clearly communicated, design or safety specifications.

Foreseeability of risk could also be problematic when deciding if a technology producer should have liability. IoT devices are typically manufactured by different producers and often involve and are mediated by various software/services. Each component or device producer, and each service provider, including platforms and integration tools, will likely only have a partial picture of the potential risks.

Thus, transparency itself might not be the solution. If no one can foresee the likely risks of failure, basing liability on failure to take reasonable steps to guard against those risks seems inappropriate.

3.1.1 Potential liability approaches

One possibility is to focus regulatory effort on the design and code driving IoT activities to ensure that risks are properly assessed and mitigated. It is argued that in a digital environment code is effectively *the* law because it controls the activity in question, and therefore the law should focus on code.[11] This would require *ex ante* transparency, in which workings of a technology can be explained in advance of its use. But providing such transparency is difficult and expensive—perhaps infeasible—where systems are complex, distributed, and dynamic, as in the emerging IoT.

An alternative is to adopt *strict liability* in relation to particularly risky activities.[8] In such schemes, common in product liability, the law does not consider fault or intention but instead assigns liability based on whether harm is caused. This might, for example, be appropriate for self-driving cars, which are largely IoT systems interacting with other IoT systems. Drivers already must be insured against negligence, and transferring that obligation to autonomous vehicle operators and extending liability for all accidents caused by the car is legally straightforward. However, a liability scheme alone will not achieve social acceptance of self-driving cars; society might further demand ex ante transparency about how cars make driving decisions before allowing their widespread use.

Ex post transparency might suffice when the precise workings of a technology are not fully known in advance but can be discovered after the event by auditing logs, testing, and investigation. Means for discovering this information might need to be designed into a black-box technology, perhaps incentivised by presuming the producer was negligent but allowing this presumption to be rebutted by showing that the producer took reasonable care and skill or those using or combining various system components did so reasonably. This would allow existing legal frameworks to provide remedies and provide time for new IoT risks to become apparent and well understood. Indeed, in highly dynamic scenarios where composition and functionality are emergent, this might be the only practicable form of transparency. Lower-risk activities could continue to be dealt with under existing negligence and contract law, allowing the supply chain to allocate risks among its members in accordance with their commercial interests.

To recap, liability risks relate to the activity undertaken (for example, using a car on the road or putting a car on the market). A legal focus on transparency should therefore be on communicating known risks and incentivizing effective processes for identifying previously unknown risks. Beyond transparency, liability also incentivises the development of technologies to assist in managing risks.

3.2 Privacy and data protection

More than 120 countries have data-protection laws with broad application to businesses. (The US patchwork of privacy laws makes it a notable outlier).[12] Despite differences, most national and sub-national laws align with the OECD Privacy Guidelines (1980, revised 2013), and a growing number shadow the high-profile General Data Protection Regulation (GDPR) (EU) 2016/679 or its predecessor the Data Protection Directive (95/46/EC), in which regulatory obligations often apply on an “all or nothing” basis. The key test is whether information constitutes *personal data*, which is defined as “any information relating to an identified, or identifiable, natural person” (GDPR Art. 4). If information is personal data then *data controllers* and their agents, *data processors,* are subject to various compliance obligations, some of which are onerous or even impossible in specific situations. Even stricter rules apply to a subset of personal data described as *special category data* that relates to health, religion, racial origin, biometrics, and so on (GDPR Art. 9). Conversely, if information is not or ceases (for example, via anonymization) to be considered personal data, it may not or no longer be subject to data-protection laws. This binary approach can be problematic, especially in the IoT’s complex data-processing scenarios.

The basic principles underpinning data-protection laws—fairness, lawful processing, purpose limitation, data minimization, accuracy, storage limitation, integrity, and confidentiality—together with individual rights to data access, correction, erasure, and so on originate in various European laws from the 1970s and have not changed much since. Meanwhile, the technologies generating and processing information about people have evolved dramatically, with vast and rapidly expanding IoT ecosystems of sensors and components collecting or processing data about individuals.

Applying data-protection rules to IoT systems is fundamentally challenging. First, almost any type of data is or may become personal data, and the strict rules governing sensitive special category data are of surprisingly broad application. Moreover, the proposed EU ePrivacy Regulation (poised to repeal Directive 2002/58/EC) could extend regulatory requirements beyond personal data to cover information about devices and the content and metadata of transmissions between them. Second, it is often difficult, if not impossible, to map the detailed accountability and compliance obligations theoretically applicable to data controllers and processors onto the reality of the multiparty, multilayered, and dynamic supply chains of the emerging IoT. Third, restrictions on cross-border data transfers, perhaps realistic when data was stored on physical media, make little sense where IoT data could be collected, accessed, and processed from anywhere. Fourth, some rules might be incompatible with many IoT-based processes—for example, systems that use ML to facilitate significant “automated decisions” where it may be impossible to provide “meaningful information regarding the logic involved” (GDPR Arts.13–15).[13] Similarly, the idea that individuals can insist in certain circumstances that their personal data be erased—“the right to be forgotten”—might be infeasible in massively distributed environments like the IoT.

It follows that organizations will likely continue with a pragmatic approach to comply with data-protection laws. However, trends are toward increased regulatory powers; for example, for many breaches of GDPR regulators can impose penalties up to the greater of €20 million or 4 percent of global annual turnover. The prospect of severe penalties combined with high regulatory uncertainty could have a chilling effect on the IoT. A different outcome is, however, possible. Many data-protection laws contain mechanisms for risk assessment and mitigation, including privacy impact assessments and data protection by design. An industry is already emerging for technical compliance tools that facilitate data mapping, auditing, and other mechanisms to help data controllers fulfil their legal obligations. Perfect compliance in a complex IoT system will likely be impossible,[5,6] but the increasingly strong accountability focus of data-protection laws might encourage regulators to use their powers to promote a more privacy-aware IoT.

4. Enabling accountability: technical directions and opportunities

Technology can help enable accountability in the IoT, particularly as the IoT and its ecosystems are still developing. Although technology will not be a panacea, it can complement and enhance other accountability mechanisms such as laws, institutional governance frameworks and standards.[9,14]

Specifically, to better align the emerging IoT with legal and regulatory concerns, technical means can assist with

Control [to determine what happens]: Enabling active steps for meeting obligations or exercising rights; and/or

Audit [make visible what happens/happened]: Providing evidence or explanations to indicate the nature of systems and their runtime operations, which enables responsive actions and recourse.

Control and audit relate directly to the previously described accountability considerations: governance and responsibility, privacy and surveillance, and safety and security. Technical control mechanisms allow proactive steps to meet one's responsibilities, for example, to tackle safety and security issues and give agency over personal data, while increasing transparency about systems and data provides information and evidence for more informed decision making, uncovering (and dealing with) risks, managing system operation, aiding investigations, and helping to hold entities to account.

We now highlight some directions and opportunities where technology could assist in increasing levels of accountability. There are a range of possibilities, as technical means differ in terms of the functionality they offer, the accountability concerns they might assist, and the guarantees and levels of assurance they provide.

4.1 Federated architectures

The IoT naturally encourages more federated system architectures, where computation and storage are pushed "closer" toward the local environments they serve.[15] Smaller clouds are emerging (fog/edge computing),[14] an instance of which might, for example, mediate the devices within a home or room or a person's wearables with her surroundings as she navigates a city. Drivers for federation include the performance, latency, and transmission considerations of servicing the local interactions of IoT-enabled physical environments, avoiding the overhead of more traditional ("distant") cloud infrastructures.[15]

There is ongoing research into federated clouds and microservices that will support decentralization of IoT infrastructure.[16] Key challenges include mediating between devices and different clouds when considering issues such as how/when to offload data and processing given functional and performance considerations regarding resources (storage, processing power), reliability, resilience, and so on.

Federated architectures are promising for improving IoT accountability by potentially enabling greater control and agency. This is because the local infrastructure components can serve as points of transparency and intervention for those within and managing such environments—for example, a cloud dedicated to a person's home might allow him to dictate its operation: how data is shared and what data is sent beyond local boundaries, such as to a service provider.

In light of this, efforts are underway to build IoT-focused *personal data stores*[10] that act as brokers for an individual's data and enable "small data" analytics in which the computation is done locally and those performing analytics do not have access to the personal data itself. This, as with *secure multiparty computation,*[14] represent alternatives to more traditional approaches in which analytics and computation occurs over an accessible, aggregated dataset. These methods could facilitate compliance with requirements relating to processing justifications, data minimisation, confidentiality/integrity, and third-party data transfers.

4.2 Data flow management

The IoT's system-of-systems nature raises interesting accountability considerations with respect to data flows across application, network, and organizational boundaries.

Access controls tend to operate within a particular application scope, enforced at a particular point in a system to govern a particular exchange but typically not beyond that.[4] In other words, there is often little visibility of or means for control over data once it is released to (accessed by) another party.

How might accountability be delivered in the IoT's complex and composite component chains? Emerging *data provenance* methods that track the flow of data end-to-end, across technical and administrative boundaries, hold real promise.[7,18] Records of data flow between components can indicate who is involved, where they are located, and what occurred. This could facilitate IoT accountability by providing evidence—for example, how personal data is collected and subsequently processed, improper behavior such as unjustified personal data disclosure to an advertiser, and the application of security patches—and by aiding risk management, investigation and liability apportionment by revealing which interactions led to a data leakage, physical harm, or other failure.

There is also ongoing work on methods to manage data proactively as it moves across boundaries. Such approaches, such as sticky policies and information flow control, generally entail coupling a management policy with data and/or its processing agents to be enforced and respected end-to-end.[4] Such capabilities could help those with rights (for example, individuals), and responsibilities (for example, organisations) to manage and exercise these throughout IoT processing chains, even when the actual components involved are beyond their direct control. Related are advances in *homomorphic encryption*,[17] which allows computation over encrypted data (without revealing the plaintext). As such schemes become more practical, they will likely influence data flow and management policy.

We have previously discussed the potential of data flow tracking and management mechanisms to assist with compliance in the IoT.[4,7,18] Much more research is needed on managing data-provenance records (what is captured, volume, access), presenting it meaningfully, and specifying and managing the policy, contexts, and complexity for means of governing end-to-end data exchange.

4.3 Enabling Audit

Given that accountability requires transparency, there is a role for technical auditing mechanisms that capture and help explain the nature of systems.

Regarding ex ante (predeployment) auditing, there are testing methodologies and ongoing work on software verification methods that ensure software satisfies a specification. Important in an IoT context, and requiring further exploration, is aligning software aspects with the physical properties, limitations, and requirements of deployed sensors and actuators. Also relevant are sandboxing and simulation methods.

Given the IoT's complexity, ex post (operational) audit records are useful for accountability by indicating what occurred and those involved. Suitable records include data flow, as well as other aspects of system operation including the contexts in which processes were operating and the management policies applied.

Blockchain or distributed ledger technology (DLT) could play a role. Arguably, much interest in DLT beyond cryptocurrency and speculation relates to the desire for increased transparency through reliable, cross-entity auditing mechanisms.[19] DLTs can establish an "immutable" record that can operate across boundaries, without necessarily (depending on its flavor) requiring trusted third parties. *Smart contracts*—code tied to a DLT instance that executes on certain events (not necessarily contracts in a legal sense, nor necessarily "smart")—also appear relevant, as they could enable the logging of evidence and legally significant events and proactively taking appropriate actions. For example, a smart contract might automatically notify users/regulators of a data breach, as data-protection law can require.

As such, secure/append-only logs or audit mechanisms—DLTs being one of a myriad of possibilities—could record evidence of system operation, data flow, and component interactions, potentially end-to-end. Indeed, DeepMind is exploring ledger-based approaches to enable "verifiable data audit" as a response to increased transparency demands regarding their use of UK National Health Service data (https://deepmind.com/blog/trust-confidence-verifiable-data-audit).

Much work remains. As mentioned previously, knowing what to record is a challenge, as is regulating access to audit information—particularly where records are shared, as audit data itself might be sensitive and entail obligations. Meaningful presentation of audit information is crucial so as to be interpretable across boundaries and appropriate for the relevant parties—be they technical experts, regulators, or end users—and the circumstances. Also required are means for ensuring that components properly write to the relevant logs and ledgers, reconciling audit records across boundaries and various repositories, and accounting for ad hoc interactions and components with short lifespans.

4.4 Algorithmic black boxes (and algorithmic/automated decision-making)

Data analytics, particularly ML, has recently received much attention. ML entails producing models from data to derive new insights, make predictions, assist in decisions, and so forth. ML will be important for the IoT, as a sensor-rich environment produces rich, granular data for building models to realise efficiencies and to drive automation.

ML raises accountability challenges given its statistical nature and because it can be complex and opaque. Addressing ML's black-box aspects is a growing area of research, as the FAT-ML (fairness, accountability, and transparency in ML http://www.fatml.org) and other communities demonstrate. Work includes auditing methods for exposing the inner workings of ML models, processes, reasons for a decision, and so on; and, control methods to, say, prevent bias and undue discrimination. The legal considerations of such work include general transparency concerns as well as specific issues such as GDPR's so-called (and contested) "right to explanation" and anti-discrimination regulations, among others.[13]

Although much of the technical discussion currently focuses on ML specifics, the systems context is also relevant. This includes the sources and nature of the data used to build ML models, the data to which the ML models are applied, and the flow-on effects of the decisions.[7] This broader context is particularly important given the IoT's composite, systems-of-systems nature: data might come from and be processed by a variety of sources, and ML-driven outputs/decisions could go on to directly or indirectly ("butterfly effects") affect numerous other systems.

Considering the broader context could help enable IoT accountability, notwithstanding the deployment of algorithmic black boxes, as other parts of the system can work to audit, mediate, and constrain the black box's behavior. For instance, data flow methods show promise: *decision provenance*[7] proposes using provenance methods (capturing data flow/processing trails) to expose 'decision pipelines', bringing visibility over areas including how datasets were constructed, the nature of system composition, and the consequences of a particular ML-driven action/decision; while data flow controls could manage and constrain these aspects.[4]

Indeed, from a legal accountability perspective, the overall processes and their effects are often more important than complete transparency of each component's inner workings. Therefore, research should consider both the specific ML and broader systems aspects of transparency and control, especially as issues of "algorithmic accountability" continue to gain prominence.

4.5 Trusted computing

Highly relevant to IoT accountability are advances in hardware/crypto-based *trusted execution environments* (TEEs)[14] Also called *secure enclaves,* these provide for secure and isolated code execution and data processing (cryptographically sealed memory/storage), as well as remote attestation (configuration assertions).

The proliferation of secure enclaves could impact accountability in a systems-of-systems environment, providing the foundations for building technical assurance into many components and administrative domains. For instance, TEEs can assist secure multiparty computation by enabling verification that particular code was executed and that critical data was not leaked. This might facilitate, for example, privacy-preserving ML analytics that occurs through agreed-upon functions without revealing the data (plaintext) itself. TEEs could also assist with the dynamic formation of data-processing chains, wherein remote attestation establishes a trust relationship between components—for example, the identity and/or software configuration of the remote component—before exchanges occur. More generally, TEEs provide the foundations for building robust auditing and strong policy-enforcement (control) regimes. In addition, enclaves could enable service providers to be removed from the chain of trust; the hardware guarantees can avoid the need to trust the entire provider-managed technical stack.

This technology is already available; examples include ARM's TrustZone (www.arm.com/products/security-on-arm/trustzone) and Intel's Software Guard Extensions (SGX; https://software.intel.com/en-us/sgx). However, it is still maturing: attacks have been demonstrated,[20] and current resource requirements might impose design constraints in an IoT context. Further, there is a learning curve as developers must ensure they properly use the enclave as well as general concerns about over-reliance on hardware, particularly given the recent Meltdown/Spectre vulnerabilities (https://meltdownattack.com), and in placing trust in a particular hardware vendor. That said, TEE technology continues to develop, gaining in prominence and featuring in the roadmaps of major ICT providers, and therefore will become increasingly relevant in IoT accountability discussions.

5. A collaborative way forward

The physicality, scale, and systems-of-systems nature of the grand IoT vision raise numerous accountability challenges. As IoT ecosystems continue to develop, and pervasive/ubiquitous computing becomes a reality, accountability aspects must receive proper consideration to ensure the technologies built and deployed are acceptable, adopted, dependable, and fit for purpose.

Progress in accountability, as a sociotechnical issue, requires a multidisciplinary approach. We have argued that better aligning technology with legal realities, and *vice versa*, is necessary for addressing accountability concerns. By providing the legal context, and highlighting some relevant technical directions, we have indicated areas warranting further attention by the research community. Specifically, we believe technical means that assist in control and auditing and that synergise legal, regulatory, and other governance frameworks represent a collaborative way forward in tackling the concerns regarding the applications, systems, and services that will increasingly pervade our world.

Acknowledgments

We acknowledge financial support from Microsoft Corporation via the Microsoft Cloud Computing Research Centre and from the UK Engineering and Physical Sciences Research Council (EPSRC).

References

1. UK Government Office for Science, Internet of Things: Making the Most of the Second Digital Revolution, report GS/14/1230, 2014.
2. D. Miorandi et al., “Internet of Things: Vision, Applications and Research Challenges,” *Ad Hoc Networks*, vol. 10, no. 7, 2012, pp. 1497–1516.
3. ISO/IEC JTC 1, *Internet of Things (IoT): Preliminary Report 2014*, Int’l Org. for Standardization/Int’l Electrotechnical Commission, 2015.
4. J. Singh et al., “Big Ideas Paper: Policy-Driven Middleware for a Legally-Compliant Internet of Things,” *Proc. 17th Int’l Middleware Conf.* (Middleware 16), 2016.
5. C. Millard, W.K. Hon, and J. Singh, “Internet of Things Ecosystems: Unpacking Legal Relationships and Liabilities,” *Proc. 2017 IEEE Int’l Conf. Cloud Eng.* (IC2E 17), 2017.
6. W.K. Hon, C. Millard, and J. Singh, “Twenty Legal Considerations for Clouds of Things,” Queen Mary School of Law Legal Studies Research Paper no. 216, 2016; https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2716966.
7. J. Singh, J. Cobbe, and C. Norval, “Decision Provenance: Capturing Data Flow for Accountable Systems,” arXiv:1804.05741, 2018; https://arxiv.org/abs/1804.05741.
8. C. Reed, E. Kennedy, and S. Silva, “Responsibility, Autonomy and Accountability: Legal Liability for Machine Learning” Queen Mary School of Law Legal Studies Research Paper no. 243, 2016; https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2853462.
9. R.H. Weber, “Accountability in the Internet of Things,” *Computer Law & Security Rev.*, vol. 27, no. 2, 2011, pp. 133–138.

10. A. Crabtree et al., “Building Accountability into the Internet of Things: The IoT Databox Model,” *J. Reliable Intelligent Environments*, vol. 4, no. 1, 2018, pp. 39–55.
11. L. Lessig, Code and Other Laws of Cyberspace, Version 2.0, 2nd rev. ed., Basic Books, 2006.
12. G. Greenleaf, “’European’ data privacy standards implemented in laws outside Europe,” (2017) 149 Privacy Laws & Business International Report, 21 [2018] UNSWLRS 2.
13. D. Kamarinou, C. Millard, and J. Singh, “Machine Learning with Personal Data,” *Data Protection and Privacy: The Age of Intelligent Machines*, R. Leenes et al., eds., Hart Publishing, 2017.
14. J. Crowcroft and A. Gascón, "Analytics without Tears or Is There a Way for Data to Be Anonymised and Yet Still Useful?," in *IEEE Internet Computing*, vol. 22, no. 3, 2018, pp. 58-64.
15. R. Want, B.N. Schilit, and S, Jenson, “Enabling the Internet of Things,” *Computer*, vol. 48, no. 1, 2015, pp. 28–35.
16. Y. Elkhatib et al., “On Using Micro-Clouds to Deliver the Fog,” *IEEE Internet Computing*, vol. 21, no. 2, 2017, pp. 8–15.
17. Abbas Acar, Hidayet Aksu, A. Selcuk Uluagac, and Mauro Conti. 2018. “A Survey on Homomorphic Encryption Schemes: Theory and Implementation,” *ACM Comput. Surv.* 51, 4, Article 79 (July 2018).
18. T. Pasquier et al., “Data Provenance to Audit Compliance with Privacy Policy in the Internet of Things,” *Personal and Ubiquitous Computing*, vol. 22, no. 2, 2018, pp. 333–344.
19. G. Pogson, *Insight Report on Distributed Ledger Technologies*, Lloyds Register Foundation/Alan Turing Inst., 2017.
20. M. Schwarz et al., “Malware Guard Extension: Using SGX to Conceal Cache Attacks,” *Proc. 14th Int’l Conf. Detection of Intrusions and Malware, and Vulnerability Assessment* (DIMVA 17), 2017, pp. 3–24.