nature machine intelligence

Perspective

https://doi.org/10.1038/s42256-025-01017-7

Materiality and risk in the age of pervasive AI sensors

Received: 26 April 2024

Accepted: 28 January 2025

Published online: 20 March 2025

Check for updates

Mona Sloane[1] ✉, Emanuel Moss[2], Susan Kennedy[3], Matthew Stewart[4], Pete Warden[5], Brian Plancher[6] & Vijay Janapa Reddi[4]

Artificial intelligence (AI) systems connected to sensor-laden devices are becoming pervasive, which has notable implications for a range of AI risks, including to privacy, the environment, autonomy and more. There is therefore a growing need for increased accountability around the responsible development and deployment of these technologies. Here we highlight the dimensions of risk associated with AI systems that arise from the material affordances of sensors and their underlying calculative models. We propose a sensor-sensitive framework for diagnosing these risks, complementing existing approaches such as the US National Institute of Standards and Technology AI Risk Management Framework and the European Union AI Act, and discuss its implementation. We conclude by advocating for increased attention to the materiality of algorithmic systems, and of on-device AI sensors in particular, and highlight the need for development of a sensor design paradigm that empowers users and communities and leads to a future of increased fairness, accountability and transparency.

Over the past decade, several overlapping, multidisciplinary communities of research and development have emerged to analyse and address the implications of AI systems operating across society, particularly through ethics, engineering, governance, critical academic and advocacy perspectives (see, in particular, the *ACM Conference on Fairness, Accountability, and Transparency*, the *AAAI/ACM Conference on AI Ethics and Society*, and the *ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*). This work has focused crucial attention on how datasets, algorithms and machine learning systems deployed at scale produce impacts for human rights, equity, already-disadvantaged populations and communities, and society at large[1–8]. Furthermore, this work has expanded the scope of investigations of these impacts beyond a relatively narrow focus on, for example, bias in algorithms[9,10] or datasets[11,12] to the situated interactions of complex computational systems in society[1,13]. Only recently, this scope has also broadened to include the material dimensions of AI systems, particularly around the environmental impacts of computation in terms of carbon-intensive energy consumption and the geographic footprint of server infrastructure[14–16], as well as e-waste[17,18]. However, relatively scant attention has been paid to the materiality of AI in terms of data collection devices and particularly sensors that interact with the physical environment (that is, the sensing of physical phenomena to produce data) and enact algorithmic inferences (for example, AI-enabled Internet of Things (IoT) devices).

There are material dimensions to problems of AI fairness, accountability and transparency that can be addressed by understanding how ubiquitous, algorithmically enabled sensors produce risks around bias, equity, privacy, accountability, transparency and consent. Materiality is a complex term, constituted by both the physical existence of an object and the social treatment thereof. There are many things that concretely materialize in our social world without taking on physical properties: the perception (or atmosphere) of a space, the potential of an idea or our social status. Other things assert themselves clearly through the physical presence: bridges, buildings, heirloom objects or people. Because materiality is integrated into social practices and social institutions, it profoundly shapes sociotechnical systems, like AI[19–22]. This becomes obvious in the context of the pervasive data collection that is facilitated by the material properties of the components

[1]University of Virginia, Charlottesville, VA, USA. [2]Intel Labs, Hillsboro, OR, USA. [3]Santa Clara University, Santa Clara, CA, USA. [4]Harvard University, Boston, MA, USA. [5]Stanford University, Stanford, CA, USA. [6]Barnard College, Columbia University, New York, NY, USA. ✉e-mail: mona.sloane@virginia.edu

of the many technological devices we use in our daily life: cameras, microphones, batteries and so on.

Today's sensor technology has shrunk the material components needed to turn physical phenomena into data onto a very small footprint. Increasingly small devices are equipped with microphones, cameras and machine learning abilities, bringing AI to the materiality of the sensor (rather than communicating data with the cloud). At the smallest scale, and at the lowest power of operation, this is referred to as 'tinyML'[23,24]. The materiality of these new types of sensor is deeply implicated in the pervasive data collection that underpins AI systems. In this Perspective, we follow Lievrouw in framing materiality as "the physical character and existence of objects and artifacts that makes them useful and usable for certain purposes under particular conditions" (page 25 in ref. 25). We build on this definition to draw attention to the material design of sensors and its impact on how physical phenomena are transformed into data. By doing so, we propose that there are certain risks associated with the designed materiality of sensors that common ethical approaches (for example, technomoral virtues[26]) or recent AI risk frameworks (for example, the Risk Management Framework (RMF) developed by the US National Institute of Standards and Technology (NIST)[27] or the European Union (EU) AI Act[28]) do not sufficiently attend to. Similarly, past reviews on the impact of AI on sensing technology mostly focus on the technical challenges of model compression at the edge and large-scale IoT architectures. And, while some may mention privacy and security in passing, they pay little attention to the embodied material challenges of such technological shifts[29–38].

We stipulate that this blind spot in AI ethics and governance primarily stems from ignoring the material affordances of technical objects, and sensors specifically. Foundational texts in science and technologies studies define affordances as the properties of objects that "are compatible with and relevant for people's interactions"[39] (see also ref. 40). Here, we build on more recent work that positions the affordances of technologies as "mediating between a technology's features and its outcomes"[41] to focus our attention on the properties of sensors that are relevant for data collection (for example, sensitivity to physical phenomena, onboard processing, storage and transmission), noting that affordances can both enable and constrain interactions[42]. While originally framed as the way in which design features enable and constrain user engagement and social action[41], we propose to understand materiality as the material affordances of sensors that deeply affect data ontology (that is, considerations of what phenomena can and ought to be data-fied), data collection and data processing—all of which affect how the benefits and risks of AI systems unfold further downstream. We note that the benefits of such material affordances are often clearly identified by those who design and market sensors, but risks are less frequently identified by designers or articulated for customers or the broader public. Drawing on a growing body of work on AI risks, we define AI risks as concrete harms that can be experienced by individuals or communities or that can be afflicted on the environment through the deployment and use of AI systems[43–46].

For illustration, we can consider the way physical properties change the performance of sensors for varying skin tones[47–49]. For example, the physical properties of the charge-coupled devices inside digital cameras, a paradigmatic sensor and the algorithms that process their outputs, contribute to how skin tones are rendered in digital files[47], which in turn contribute to computer vision applications such as face detection and facial recognition[50].

Adopting a material lens and an affordances approach also makes it possible to apprehend the risks that adhere not just to a single sensor, device or AI system but instead emerge from their widespread adoption. We stipulate that adoption is driven by how material properties and functionalities combine with calculative behaviours and calculative models, that is, the collectively shared assumptions and practices about both usefulness and economic value that render sensors commercially viable (these include, but are not limited to, cost of energy for production and use, sensor size, supply chain considerations of material components, as well as profit projections and pricing schemes)[51,52]. Calculative models directly affect sensor price, availability and ubiquity, making the concept particularly useful for understanding drivers behind the proliferation of sensors. As we will show, low production cost and particular calculative models of cameras, microphones and other sensors make them increasingly common in a wide range of contexts.

The calculative models that drive the increased affordability and accessibility of sensors have led to positive developments. The integration of sensors in home appliances and other consumer goods can improve their utility and performance. For example, there is a small camera in the Keurig K-Supreme Plus Smart coffee maker that detects the type of pod being inserted to adjust the water temperature for optimal brewing. Sensors have also enabled beneficial applications such as predictive maintenance for public infrastructure[53] and industry[54], providing support for safe driving practices[55], and optimizing energy usage in offices and homes[56]. Most notably, sensors coupled with edge computing have opened up novel applications that can make progress towards the 17 Sustainable Development Goals outlined in the United Nations' 2030 Agenda for Sustainable Development[57,58]. For example, sensors are being used to optimize agriculture[59,60] and aid wildlife conservation efforts[61–63].

However, sensors also give rise to significant concerns as their proliferation into the public, professional and intimate dimensions of our daily routines enables unprecedented data mining and commercialization of once private or anonymous moments and behaviours[64,65]. The material affordances of sensors embedded in mobile devices (such as cameras, microphones, gyroscopes, GPS antennae and inertial measurement units) and home electronics are often opaque and unaccountable in ways that make it difficult for anyone to understand when and what information they might be collecting and analysing. In addition, it is well known that sensor-driven surveillance technologies are more likely to be deployed in already over-surveilled communities and professions[66–68], exacerbating disparate AI impacts and inequities[69,70].

In this Perspective, we examine the materiality of AI risk production by paying close attention to how sensors are incorporated into AI systems and vice versa. We suggest that the material affordances and calculative models of sensors contribute to the growing ubiquity of sensors and the general risks of AI systems. To do so, we first demonstrate how the material affordances and calculative models of sensors co-evolved over the past half-century by laying out the 'evolutionary history' of sensors that culminates in today's environment of pervasive AI-enabled sensing. Leading on from this analysis, we build on the two most prominent and widely adopted AI risk management frameworks in the USA and in Europe—the RMF developed by the US NIST[27] and the EU AI Act[71]—to propose a sensor-sensitive AI risk identification framework. In a third step, we raise a call for the development of a new sensor design paradigm that addresses the risks posed by the convergence of ubiquitous sensing and AI technologies, particularly in the context of tinyML. Overall, the key contribution of our work is to expand attention to the material affordances and calculative models of sensor-based AI systems when engaging in AI risk diagnostics.

Sensors everywhere

Sensors are devices that convert "a physical phenomenon into an electrical signal" (page 1 in ref. 72) that can then be used to quantify environmental aspects such as light, heat and pressure. They are designed to monitor phenomena within and beyond our human perception, such as imagery, movement, sound or chemical composition. They transform physical phenomena into numerical representations, adding to the 'avalanche' of numbers[73] that make it possible for the world—people, commodities, communities and nature—to be represented computationally, for behaviours to be analysed. Crucially, they are the foundation for many of the datasets that undergird powerful AI technologies. Digital cameras are the sensors that produced the datasets that enabled image recognition[74]. Speech recognition would be impossible without the datasets generated using digital microphones[75]. And many more AI applications, from autonomous vehicles to personal athletic trackers, rely on sensor data as well. However, technical innovations do not gain widespread adoption merely because of their technical superiority or consumer demand. Rather, they shape and are shaped by a wide range of (often competing) social and commercial interests[76]. To develop a sensor-sensitive approach to AI risk diagnostics, it is key to unpack how the material affordances of sensor technology co-evolved over time in tandem with the calculative models that motivated their spread and uptake.

The 'evolution' of sensing

Sensors have distinct material characteristics that generally remain unchanged: sensitivity to the intensity of an input, dynamic ranges within which phenomena can be accurately transformed into meaningful signals, and a propensity to produce noise alongside meaningful signals[72]. But sensors' material affordances—the social actions they enable rather than their abstract characteristics[77]—have changed dramatically over the past few decades, with significant implications for how they are used, and for the overall AI risk landscape. As technology has advanced, sensors have become smaller, more efficient, more sensitive and more connected, allowing them to be embedded in various environments and objects, thus providing real-time data and insights that were previously unattainable. This 'evolution' has not only extended our ability to understand and interact with the world around us but also paved the way for innovations that continue to shape our future. Sensors have evolved from simple analogue devices into 'intelligent' systems capable of analysing data and making decisions at the edge. The advent of wireless sensor networks in the late 1990s marked an important milestone, allowing remote collection and analysis of sensor data[78]. However, sensors were still reliant on external processing power. The subsequent rise of smartphones and IoT devices brought sensors out of isolation, interconnecting them through the internet[79]. For the first time, distributed networks of sensors could coordinate to achieve larger goals. However, even as billions of sensors flooded the environment with data, their capabilities remained confined to data collection rather than interpretation.

This changed with the emergence of edge AI sensors and tinyML sensors, which move machine learning out of data centres and closer to the sensor devices themselves. Edge AI sensors communicate with nearby processors—either separate components within the same device that contains the sensor, or a mobile phone or computer wirelessly connected to the sensor device—where machine learning models run locally[29,30]. TinyML sensors incorporate machine learning models directly on the sensors based on new advances in small but powerful models that can run directly on such devices—so-called tinyML models[23,24]. In particular, these models can run directly on the microcontrollers commonly found on physical sensor hardware. This allows for real-time data processing and analysis at the sensor level, without needing to transmit data to a separate computing device. Furthermore, tinyML sensors usually limit the information shared from the sensor device to only application-specific instructions, for example, whether or not a person is standing in front of the device, rather than an entire video data stream that could then be put to many purposes off-device[23,80]. Rather than indiscriminately streaming all collected data to cloud servers, edge AI and tinyML sensors interpret their environment and surface insights, and make decisions not in distant server centres but at the edge, that is, on the device or a nearby connected device.

Below, we show how successive stages of sensor development, with their evolving material affordances and calculative models (which combine considerations around sensor size, energy use, supply of material components and so on into profit projections, pricing schemes and sales strategies), shape the risk profiles of sensors as components of AI systems, providing common examples of each in Fig. 1, which we then discuss in the following section ('A sensor-sensitive AI risk diagnostics framework'). We note that each of the stages of sensor development we identify build on earlier material affordances of sensors while adding additional properties. Accordingly, when thinking about the risks sensors contribute to algorithmic systems, we see that many of those risks accumulate from one stage to the next.

- Traditional sensor[72,81–83]. A device that affords the ability to generate data from or about the physical world. The sensor itself may transmit analogue or digital signals, and signals produced by the sensor may need post-processing on- or off-device to be useful. Data may be stored on- or off-device, or may be entirely ephemeral. Traditional sensor devices do not require an internet connection to function and typically have no internet connection capability. Data must be purposefully transferred off-device and stored using specific protocols such as Serial Peripheral Interface (SPI) or Inter-Integrated Circuit (I^2C), both of which are common communication protocols used to transfer data between microcontrollers and peripherals. No statistical inferencing happens on-device.
- IoT sensor[84–88]. A device that affords the ability to generate data from or about the physical world, and to access that data in real time or near-real time using internet transfer protocols. IoT sensors also afford the ability to collect data not directly related to the physical phenomena they are designed to sense like a list of connected devices, strength of WiFi signals, battery status of connected devices and other metadata related to the device's performance. IoT sensors may be able to operate without an active internet connection, but most are designed with capabilities that require internet connectivity, particularly to generate or store data. No statistical inferencing happens on-device, and rarely even happens in the cloud. Data transmitted over the internet are (almost) always stored in a centralized location. This type of device acts primarily as an interface to cloud-based data storage systems.
- **AIoT sensor**[29,31,32,34,36,89]. A device that affords the ability to conduct cloud-based AI decision-making based on data from or about the physical world. This is done through real-time or near-real-time access to AIoT (AI of Things; that is, a combination of AI systems with Internet of Things infrastructure) sensor data using internet transfer protocols, and integrated with AI/machine learning inferencing techniques in the cloud (and importantly not on or near the device). As such, this type of device requires an active internet connection to operate at full functionality. Importantly, full datastreams generated by the sensor are (temporarily) stored in centralized locations (that is, cloud or server device). This type of device acts primarily as an interface to cloud-based decision-making systems powered by AI.
- **Edge AI sensor**[30,35–38,90–92]. A device that affords the ability to conduct AI decision-making at the location of the device, based on data from or about the physical world. This is done by processing sensor data using machine learning techniques through a combination of on-device and near-device edge processing (for example, the Apple Watch offloads some computations to a user's cell phone). These data may afford remote access and data reuse instances where it is transmitted off-device, and possibly over the internet, in real time or near-real time, in its raw or processed form. Transmitted data may or may not be stored in a centralized location. This type of device extends the cloud-based intelligence of AIoT sensors to the edge.
- **TinyML sensor**[80,93]. A device that affords the ability to accomplish a predetermined task with only the minimal amount of information needed to accomplish that task. For example, a tinyML camera designed for person detection may only afford the ability to read a single bit from the sensor device ('1' if there is a person within view of the camera and '0' if not). This is done by using machine learning inferencing techniques on data from or about the physical world

		Traditional sensors (1970s–1990s)	IoT sensor (1990s–2010s)	AIoT sensor (2010s)	Edge AI sensor (Late 2010s–present)	TinyML sensor (Present day)
	Examples	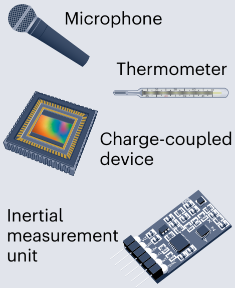 Microphone Thermometer Charge-coupled device Inertial measurement unit	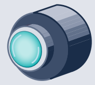 CCTV camera	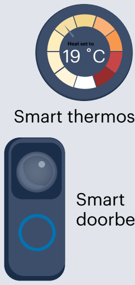 Smart thermostat Smart doorbell	Smart glasses Smart watch	Useful sensors (person detector) 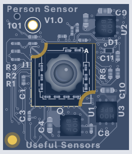
Material affordances	Data generated	• Digital or analogue signals about physical phenomena	• Digital signals about physical phenomena • Metadata about device and location • Additional data not related to physical phenomena being sensed	• Digital signals about physical phenomena • Metadata about device and location • Additional data not related to physical phenomena being sensed	• Digital signals about physical phenomena • Metadata about device and location • Additional data not related to physical phenomena being sensed • Data can be stored off-device but affords ability to limit or prohibit off-device data storage	• Digital signals related to predetermined tasks
	Data storage	• Data stored directly on-device or offloaded through direct access	• Data rarely stored on-device • Data often stored on cloud servers	• Data stored on cloud servers	• Data storage optional near- or on-device	• No data stored
	Connectivity	• No internet connectivity needed or even possible	• No internet connectivity needed but often internet-connected	• Internet connectivity needed for full feature functionality	• Full functionality possible without internet connectivity	• Transmits only minimal stream of data to accomplish predetermined task
	Statistical inferencing	• No statistical inferencing happens on-device	• No statistical inferencing happens on-device, inferencing may occur using cloud-stored data	• Statistical inferencing off-device	• Statistical inference near- or on-device	• Statistical inference occurs on-device only
Calculative models	Object based	✕				
	Data based	✕	✕	✕	✕	
	Subscription based	✕	✕	✕	✕	✕

Fig. 1 | Timeline of sensor evolution from passive analogue detectors to intelligent IoT and machine learning-enabled systems. Examples of devices, their material affordances and their calculative models (indicated by crosses) are also included.

entirely on-device. Crucially, such sensors do not afford the ability to read the raw image data from the sensor device. These devices are typically self-contained and store minimal or no sensed data.

It is important to note that these successive stages are 'evolutionary' only in a descriptive sense; sensors of each stage remain in production and use today and build on many of the characteristics of the sensors they were preceded by. But their material affordances are quite different from stage to stage. Where traditional sensors require direct access (physically, over wire, or through radio or infrared signals), IoT and AIoT sensors enable data access from anywhere in the world with an internet connection, data aggregation and real-time monitoring of many distant locations. Edge AI sensors can afford similar capabilities, but also afford autonomous operation of nearby devices connected to these sensors. TinyML sensors, in contrast, afford the ability to accomplish tasks without needing to support high-bandwidth data flows or process data off-device. Importantly for understanding the AI risks associated with different stages of sensors, each stage has different affordances for misuse (whether intentional or not); traditional sensors and TinyML sensors are more difficult for bad actors to gain unauthorized access to than IoT devices, and do not require the production of datastreams that might make users and passersby vulnerable to privacy breaches or other malicious behaviour.

Calculative models and sensor development

The wide adoption and evolutionary transformation of sensors can be explained in part by their sheer utility for gathering data about the physical world, and its ascribed usefulness and commercial viability. In addition, each stage of sensor development adds capabilities that lead to new products and services. But the proliferation of sensors cannot be explained entirely by the technical needs they satisfy. Instead, the integration of sensors into consumer devices, industrial machinery and civil infrastructure—like any other product—is steered more by the dynamics of calculation for circulation and trade, rather than by technical features. Economic markets are not independent agential entities that are external to social and material life, but rather are collectively organized tools that facilitate the calculation of the value of goods[51,94]. Calculation bridges quantitative and qualitative aspects in an effort to make goods tradable. The practices and cultures of calculation[95] differ by market, but always evolve around certain sets of calculative models, that is, repeatable ways of calculating an object or good. These calculative models combine cultural knowledges and assumptions about people and consumption (for example, about convenience in the home, stipulating that people will prefer to operate their light switch with voice commands over physically switching on lights) with determinations of the social and behavioural impact of the material affordances of an object or product, costs of manufacture, actual and projected volume of sales, and cumulative profits in the context of the many possible additional factors that affect price, costs and earnings[96,97].

There are three interrelated material aspects affecting the evolution of calculative models of sensor technology and their recently accelerated proliferation. The first is a reduction in cost for the production of sensors due to advancements in manufacturing technologies, the development of more affordable material components and the reduction in cost driven by production at scale[98]. The second, relatedly, is the miniaturization of sensors[99–101], often considered an innovation and adoption driver[102–105]. The third one is, again relatedly, a measurable reduction in energy consumption in sensor deployment. Especially for large networks of sensors, vastly increased energy efficiency of sensors, even for edge AI and tinyML sensors that run machine learning models at the edge, enhance efficiency by removing the need for cloud services, making sensors capable of real-time data analysis for a wide range of applications where an internet connectivity is impractical or latency restrictions are severe[58,106–111].

These three interrelated and material aspects have given rise to three distinct calculative models of sensor technology, with one emerging before the other. While in the 1970s, 1980s and 1990s, the calculable and tradable unit was the object of a traditional sensor itself (thermometer, microphone, charge-coupled device or inertial measurement unit), the calculative model changed with the arrival of IoT sensors and sensor networks, starting in the 1990s. Here, the calculable unit is not the sensor or sensor network itself, but the value of the data that it collects. Pricing emerges around the ability of collecting and interpreting that data, for example, in the context of predictive maintenance, energy management and consumer-behaviour analysis. A third calculative model later emerged, spanning across all sensor types, focused on subscription-based and service-oriented models as calculative units. Companies are increasingly offering sensors as part of a service package, where the initial cost is low, but users pay a recurring fee for data analysis, cloud storage and other associated services.

Today, even mundane appliances have increasing numbers of sensors built into them. There has been an exponential growth in the market size of deployed smart IoT devices, with a total 14.3 billion devices now in use globally in 2022 and a projected 29 billion devices by 2027[112]. This proliferation of AI-connected sensors brings new forms of risk: the likelihood, frequency and severity of a harm occurring to individuals, communities or the environment. As sensors variously proliferate, connect to AI services and embed machine learning capabilities on-device, they shape and reshape how and when algorithmic harms may occur. This has important implications for how we should diagnose AI risk, given the many ways different sensors interact with social worlds and the physical environment, how data are shared between and across devices containing different types of sensor, and how the material affordances of sensors enable the activities that lead to harm.

A sensor-sensitive AI risk diagnostics framework

Increased attention to the harms and risks of AI systems has led to the development of several risk management frameworks for AI (see, for example, refs. 113,114) that are variously oriented towards specific business purposes or regulatory conformity. In the USA, NIST has developed an AI RMF that is designed as a general purpose framework, applicable across many domains[27]. The EU has developed a risk framework for AI through the EU AI Act that identifies specific uses of AI as presenting different levels of risk, with some uses prohibited, others high risk enough to be ‘regulated’, and still others requiring varying safety features and transparency mechanisms[71]. Additional frameworks have been produced by governmental organizations, such as the Organisation for Economic Co-operation and Development[115], or consultancy groups, for example, refs. 116,117. The NIST AI RMF developed seven characteristics of what they refer to as ‘trustworthy AI systems’, and which are quickly becoming key elements of AI governance. Such systems are valid and reliable, safe, secure and resilient, accountable and transparent, explainable and interpretable, privacy enhanced, and fair. The NIST AI RMF points out the need for “balancing each of these characteristics based on the AI system’s context of use” (page 12 in ref. 27). The EU AI Act similarly focuses on contexts of use, by drawing attention to deployments that are explicitly prohibited, those that are regulated as ‘high-risk AI systems’, those that have ‘fundamental safety and transparency’ concerns, those that require transparency only, and those that are not in the scope of the act[71]. However, these risk management frameworks set out to provide comprehensive and often too broad guidance on AI risk identification, mitigation and management, and do not adequately attend to the role that pervasive sensing has in contributing to the risks of AI systems. Therefore, we propose a sensor-sensitive AI risk diagnostics framework that attends to how the material affordances of sensors, and the calculative models that drive their features and deployment patterns, produce specific AI risks within wider AI systems.

Our approach proceeds from an analysis of the materiality of objects—here, the material affordances of sensors—in contrast to frameworks like NIST’s, which proceeds from impacts to users and communities and without attention to the impacts and dynamics of materiality and proliferation. Emerging from our preceding analysis of material affordances and calculative models of sensors, our sensor-sensitive framework for AI risk diagnostics focuses on five key aspects that must be considered as part of assessing AI risk of sensors: (1) calibration, (2) documentation, (3) proprietary data profusion, (4) privacy and (5) waste. Each of these aspects are intended to be considered in addition to other characteristics of an AI system. Overall, our diagnostic ought to be read as complementary to government-mandated AI risk management approaches as we suggest it be used specifically for sensor-intensive AI systems and applications. Our argument is that sensors contribute to the risks of AI systems through their material affordances and calculative models. This contribution is worthy of analysis separate from an analysis focused exclusively on how the data produced by sensors contribute to the risks of AI systems.

Calibration

All sensors must be properly calibrated to return validated, reliable, accurate, robust, comparable data about the physical world. But the challenge of both achieving and confirming these characteristics is different for various types of sensor[118]. Traditional sensors are often calibrated before they leave the factory and are engineered either to provide users a means of recalibrating the sensor to a known value (for example, using standardized weights to adjust a scale or boiling distilled water for a digital thermometer) or to include an adjustment curve to adjust observed measures to known measures. Calibration also has a sociotechnical dimension, in that a sensor is calibrated for an intended purpose and can be miscalibrated for other purposes. A thermometer designed for personal health use may need to be calibrated to be most accurate around normal human body temperature, whereas a thermometer designed for industrial use might need to be calibrated differently to be most accurate around the melting point of various substances (for example, iron, ammonia or copper). There are numerous instances where sensors have been miscalibrated for their social uses; colour film was long miscalibrated to lighter skin tones despite the fact that it was being used for photographing people of a wider range of skin tones, and digital photography was developed to emulate colour-film calibration curves[48]. In addition, sensors can become miscalibrated over time, as communities and the environment change; so-called shot-spotter technology requires on-site recalibration to function properly, particularly when urban geography changes in ways that affect acoustics[119].

Networked sensors—IoT and AIoT sensors—do not always afford the direct physical access required to conduct or confirm recalibration, so more sophisticated computational techniques may be needed to calibrate already-deployed sensors over a wide area without direct access to the device itself or a reliable reference measurement to calibrate sensor readings[120]. Edge AI and tinyML sensors also raise validity concerns, as their designed-in autonomy often means calibration cannot be remotely maintained once deployed. The calculative models that underwrite the development and pervasive deployment of sensors also affects their calibration to their environment; designing calibration affordances into sensors adds to their cost, maintaining calibration of a fleet of distributed sensors requires substantial organizational infrastructures and costs, and data collected from networked sensors may continue to hold commercial value even if it falls out of calibration.

Sensors' material affordances shape the data they produce as a condition of its production. Unless they are properly designed for the phenomena they are deployed to sense, the data they produce will be inherently unreliable; oxygen-saturation sensors, for example, must be properly calibrated to an individual's skin tone to work properly[49]. Furthermore, because pervasive sensing places sensors into contexts for which they may or may not have been properly calibrated, they are implicated in the reliability and safety of AI systems, as well as their ability to function fairly and manage bias appropriately. This risk is more acute in some use cases than in others, and so risks associated with the mis/calibration of sensors should be evaluated both in relation to the NIST characteristics of 'trustworthy AI' and the various levels of risk associated with use cases in the EU AI Act.

Documentation

Traditional sensors have long been developed and deployed through a commitment to safety engineering[121], subject to safety testing, and—like other components—accompanied by standardized datasheets that layout their safe and unsafe use cases (for example, ref. 122). AIoT, edge AI and tinyML sensors are similarly accompanied by such datasheets; however, these types of sensor afford tight coupling to machine learning, either in the cloud, on-device or nearby. For these tightly coupled systems, datasheets do not provide adequate documentation for key trustworthiness characteristics such as interpretability and transparency. As these sensors become more pervasive, their under-documentation makes it increasingly difficult to manage their appropriate deployment, identify cases that require closer scrutiny and monitor their impacts for potential incidents of harm.

Existing proposals for documenting edge AI and tinyML sensors suggest including model cards for their machine learning models[93,123]. While this would help mitigate the risks associated with inadequate documentation, effectively managing these risks requires a comprehensive evaluation of whether AI systems—including the sensors that produce data for them—are adequately documented. As the use of sensors becomes more widespread, documentation is also crucial for understanding the material affordances of data production and the calculative models under which data may be collected, bundled and sold. While existing documentation might account for specific data production from a networked sensor, such sensors may afford the production of additional datastreams beyond that which is explicitly documented, producing unanticipated risks[124]. The production and leakage of undocumented data can cause considerable harms that may be magnified in specific use cases. Transparency is a key dimension of AI governance; however, inadequate documentation of the material affordances of pervasive sensor networks comprises a discrete risk to be identified and managed as well.

Proprietary data profusion

The calculative models that underlie a sensor-saturated world tend towards reducing the unit costs of sensors and ensures their omnipresence. The material devices are becoming cheaper to produce and any excess unit costs are often rationalized by the calculative models discussed above. This has the result that the amount of data collected, which can be used for providing services (for free or on a subscription plan) as well as sold on a secondary market through data brokers[125], is always maximized. Such a model leads observers to conclude that the scale of data collection and aggregation often exceeds individuals' expectations or ability to control[126,127]. Given the shift in the calculative models attached to sensors, particularly the perceived need to recoup per-unit costs, there is increasing pressures for private- and public-sector organizations to keep data proprietary. This constitutes a risk that the material affordances and calculative models of sensors contributes to, which in turn makes AI systems more risky.

Data profusion brought by sensors has the potential to amplify the divergence between data quantity and quality and its associated risks. Calculative models are oriented towards the deployment of low-cost sensors which, unlike industry grade sensors, are more susceptible to factors that will result in incomplete or inaccurate data[128]. Sensor data quality can also suffer owing to variations in hardware manufacturing, sensor drift or the state of the battery—as sensor data tend to becomes less reliable as the battery nears the end of its lifespan[129]. Critiques of 'big data' practices underscore the failure of datasets to offer an objective, accurate and comprehensive portrayal of real-world phenomena[130–132]. Given that policy-making decisions and resource allocation often rely on quantifiable information, the invisibility or inaccurate portrayal of specific individuals or phenomena within datasets contributes to their marginalization. An increase in dirty data generated in a sensor-saturated world risks perpetuating these issues. Moreover, data profusion propelled by sensors might exacerbate the risks associated with poor data quality as it will create an overwhelming demand for data cleaning. This is a time-consuming process that requires a domain expert to be done effectively and cannot be easily scaled to meet demand[133].

Privacy

The materiality of sensors and the scale of their deployment present a unique set of challenges with respect to privacy that are worthy of explicit attention. The physical characteristics of sensors in terms of their small size allows them to be inconspicuously integrated into one's surroundings in novel and unexpected ways. This presents a formidable obstacle to safeguarding privacy through the mechanisms of notice and consent. Individuals who are unaware of sensors may be involuntarily subjected to data collection and algorithmic processing. But even if individuals were notified of sensors, the abundance of these devices would render the practice of consent untenable. The time and attention required to provide informed consent in every case would exceed an individual's finite capacities[134,135].

The NIST AI RMF notes how AI systems can present new privacy risks such as enabling inferences that jeopardize de-identification efforts. This risk will become especially salient in a sensor-saturated world, as an increase in the volume and breadth of data is positively correlated with the ability to draw such inferences. Moreover, the process of sensor fusion, where data from multiple sensors are combined, enables inferences that would otherwise not be possible from a single data stream[136]. Pervasive sensing, therefore, provides material affordances for expanding the potential to draw inferences and cross-reference data not only contribute to re-identification risks of de-identified data[137] but also threaten an individual's ability to exercise control over their data in at least two respects. First, the potential to draw inferences means that individuals may not fully understand either the fact of data collection (that is, the fact that data were being collected by a sensor at all) or the implications of data collection, casting doubt on the informed nature of their consent. Second, an individual's decision to opt out of data collection may be rendered futile, as the decisions of others to permit data collection can enable inferences that inadvertently implicate those who seek to abstain, eroding the notion of individual agency in data sharing[138,139].

Waste

The environmental impact of AI is addressed in risk management frameworks such as the NIST AI RMF, which explicitly identifies risks of harm to the environment within its definition of 'safe' AI systems. This framework briefly references "conditions [under which] the environment is endangered" (page 14 in ref. 27), although it lacks operational guidance for assessing these risks. While broader discussions about the environmental costs of computing technologies are more developed elsewhere[140–143], these discussions predominantly focus on the operational activities associated with product use, such as the energy consumption required for training and running inference on a machine learning model. We contend that this framing does not fully capture the risks posed by the material production and disposal of sensors[17,144,145]. Below, we detail the unique environmental challenges associated with sensors to illustrate why waste is a distinct and critical aspect of risk for these devices.

First, the environmental risks of sensors differ significantly from other AI systems owing to their reduced carbon output associated with operational activities. The operation of battery-powered sensors, even at a massive scale, can be expected to consume significantly less energy than other computing technologies[58]. Instead, the environmental impact of these devices lies primarily in the manufacturing and disposal phases, where carbon emissions and other risks are tied to the supply chain and end-of-life processing[146]. The biggest contributing factor to their embodied footprint is the batteries that power them[58]. Coin cell batteries, for example, pose significant challenges due to their small size, short lifespan and toxic components (for example, lithium, mercury, cadmium). As these batteries might not be easily recyclable or biodegradable, they may accumulate in landfills or, worse yet, contribute to pollution and environmental hazards as a result of their improper disposal. The second largest factor of sensors' carbon footprint is the sensor components themselves, which often rely on rare earth elements and carry high extraction costs including habitat destruction, water pollution and other ecological impacts. Recent research has shown that significant improvements can be made with regards to reducing the waste impact of sensors by way of integrating sustainable materials into sensor design[147].

Although these environmental risks may resemble those associated with other consumer electronics, the distributed nature of sensors nevertheless raises distinct concerns. Sensor deployments in agriculture, wildlife or environmental monitoring can involve hundreds or even thousands of devices spread across large or remote areas. Unlike consumer electronics, which are typically consolidated in urban areas with access to recycling programmes and associated infrastructure, retrieving and responsibly disposing of sensors may be impractical. Thus, there is a risk that sensors will contribute to 'stranded e-waste' that accumulates in inaccessible or remote environments. Furthermore, sensor devices often operate invisibly or autonomously due to their miniaturization and edge computing design. This materiality of sensors contributes to the risk of 'invisible pollution', where discarded devices contribute to environmental harm in ways that may go unnoticed by regulators or consumers. Unlike centralized computing technologies, which are easier to oversee and optimize for sustainability, the decentralized and distributed nature of sensors presents novel challenges for managing the associated environmental risks.

In addition, sensor-based AI systems are often viewed as environmentally friendly compared with other AI systems owing to their low energy requirements during operation. When deployed towards sustainable aims, their overall carbon footprint may even be net negative[58]. However, the energy efficiency of sensors can, counterintuitively, contribute to their broader environmental risk. When sensors are seen as energy efficient, this framing can obscure the environmental impact associated with their end-of-life processing and disposal. Moreover, when applied to sensors, a theory in economics known as Jevons' paradox suggests that increased efficiency can lower the perceived environmental cost, encouraging widespread adoption and ultimately increasing resource consumption[148].

Finally, the low cost and scalability of sensors further compound these risks. While the affordability of sensors is often celebrated as an innovation that can enable their global adoption, it may inadvertently promote an overreliance on sensor-based solutions. Given the low cost and accessibility of sensors, they may be preferred over non-technological alternatives such as policy reforms or community-based interventions, even when the latter have lower environmental costs. The calculative models and materiality of sensors leads to increased affordability, energy efficiency and smaller-size devices, spurring their widespread application and adoption. While sensors hold promise for supporting sustainability efforts, the distributed nature of their deployment and their potential to have a paradoxical effect on resource consumption present distinct environmental risks that must not be overlooked.

Implementation

To implement the above risk diagnostics framework, efforts must be made to map, measure and mitigate the ways in which sensors contribute to the overall risk of an AI system. Mapping risks consists of identifying how the specific sensors employed as part of that system might become miscalibrated or be poorly calibrated, how they might need to be documented, what data they produce, how they contribute to or undermine privacy, and what forms of waste their production and use represents. Particular attention ought to be paid to the material affordances of these sensors, to identify what features or components of the sensors are associated with each risk category. Measuring risks consists of identifying the severity and scope of each risk category. A privacy risk might be slight, with respect to any one user, but have widespread implications for thousands or even millions of users. Conversely, it may be severe but occur rarely. Both dimensions must be evaluated to accurately assess the overall risks. Special attention should be given to the calculative models at play and how they shape the severity and scope of identified risks. Mitigating risks consists of identifying the roles that material affordances and calculative models play in each identified risk and exploring how these elements could be modified. For instance, if a particular privacy risk is given broader scope owing to a data retention policy that enables data brokers to buy and sell user data, implementing a data deletion policy could help reduce some of that risk. In addition, alternative designs may be developed to partially mitigate the risk of waste, and active monitoring of reference devices might be undertaken to mitigate the risk of miscalibration over time.

Discussion and future work

As sensors become interwoven into the fabric of everyday life, they may fade into the background of conscious experience. Nevertheless, these devices are capable of exerting powerful influences over individuals' choices and behaviour. AI sensor-based systems can subtly or overtly exert power, whether through persuasion or coercion[149]. Activity trackers, for instance, employ a subtle approach by using data-driven feedback to encourage users to exercise more. In contrast, a seat-belt sensor that locks the car ignition until the driver buckles up represents a more coercive form of influence. The evolving material affordances of sensors—particularly their shrinking size and enhanced processing capabilities—will allow for their integration into devices that are ever more intimately linked to individuals' daily routines[150]. This reality, combined with the calculative models driving the creation of a sensor-saturated world, suggests that individuals will probably be increasingly subjected to technological influences.

AI sensor-based systems can offer significant utility that make them attractive to consumers. For instance, despite users' awareness of the privacy and security risks associated with IoT devices, research reveals a prevalent 'I want it anyway' attitude among users[151]. However, as AI sensor-based systems become more pervasive, users may reassess their willingness to accept the associated risks. One study indicates that users are less inclined to embrace data collection if they perceive it as excessive within the broader context of their lives: “Smart home technology was often viewed as a further invasion of, or threat to, privacy in a society where already too much personal information is collected and stored” (page 369 in ref. 152). Notably, individuals’ attitudes are context sensitive; their perception of risk associated with a particular IoT application is intertwined with their views on the broader proliferation of AI sensor-based systems.

The proliferation of sensors gives rise to new dimensions of risk that are capable of being felt by the general public. Yet existing approaches to AI ethics and governance overlook these risks owing to their narrow focus on specific applications of AI. For example, the NIST RMF for Information Systems and Organizations—which pertains to information systems ranging from cloud-based systems to IoT devices—explicitly avoids considerations related to material affordances of technological objects such as sensors[153].

Similarly, the EU AI Act focuses on regulating AI systems deemed high risk owing to their use case, but overlooks broader risks associated with the proliferation of sensors that contribute to proprietary data profusion. The EU Data Act[71] represents a notable effort to address this gap, aiming to promote equity by allowing individuals and businesses to access data generated by their use of AI sensor-based systems. This recent development in the EU regulatory landscape demonstrates that piecemeal progress is possible. However, a more unified approach to risk management can be achieved by adopting a lens of analysis grounded in both the material affordances and calculative models of sensors. This perspective allows for a comprehensive understanding of the risks associated with AI sensor-based systems by considering the roles that their technical and physical properties have in shaping their impacts.

Beyond risk management, further work is necessary to support efforts towards the responsible design, development and deployment of sensors. This includes technical teams developing sensors to incorporate interdisciplinary collaborations with social scientists and behaviourists[154] into their work, and potentially leads to harnessing the potential benefits of sensors (in part through attention to their material affordances)[155], particularly for decentralizing power by widening access to tinyML. With sufficient sensor and AI literacy, individuals and communities may be able to build sensor-driven AI systems that genuinely benefit them, under terms they define themselves. These could include entirely private and closed health monitoring systems, weather and crop monitoring employed in farming communities, or sensor deployment for citizen research projects focused on environmental justice or other community concerns.

To animate and harness these benefits alongside strengthened sensor governance through the proposed sensor-sensitive AI diagnostics framework, it is essential to meaningfully engage stakeholder communities to contribute to the creation of inclusive guidelines and best practices. This ensures that the deployment of sensor technologies considers a broad spectrum of perspectives and needs, balancing technological advancement with societal well-being[156,157]. The focus on material affordances of sensors demonstrated above is just as well suited to the exploration of how sensors might benefit such communities, on their terms, as it is to the enumeration and analysis of risks.

A range of different stakeholders, all of whom are experts in various applied fields as well as dimensions of risk discussed above, should focus on collaborating to create a sensing paradigm that is aiming to alleviate the negative impacts stemming from the material affordances and calculative models prevalent in today’s sensor ecosystem, and developing community-driven approaches to sensor and data use. In addition, they should work on scoping transparency in ways that are relevant to the lived experience of interacting with sensors to promote the creation of transparent systems that make it easy for users to understand how their data are being used and for what purpose. Incorporating methodologies such as the machine learning technology readiness levels (MLTRL) framework can provide a structured approach to ensuring that these systems are robust, reliable and responsible from development through deployment[158].

In addition, the machine learning sensor paradigm[80] can provide a technical frame for risk-aware and community-driven sensor and data use. This approach suggests that sensors should process all data internally and transmit only abstracted, high-level data through a streamlined interface. This adheres to the principle of data minimization, ensuring that raw data remain exclusively accessible to the onboard sensor processor. This architecture not only enhances system self-containment, thereby improving auditability and accessibility, but also empowers users by maintaining control over their raw data, reducing the likelihood of unwarranted data exploitation by commercial and governmental entities. The proliferation of such a paradigm, or the development of alternatives, will be critical in ensuring that responsibility is a core design principle for future sensor systems. While substantial additional work needs to be done to adequately address the risks posed by AI sensor-based systems[159], this approach provides a strong starting point. Parallel efforts can also use the lens of calculative models to intervene in how sensor-focused AI is underwritten financially. Taxation and increased regulatory scrutiny of such devices can shift the development logics away from integrating sensors that facilitate unconstrained data collection in every device.

These efforts can be combined with a focus on interpreting the societal impact of these technologies, advocating for the rights of affected communities and helping to draft robust regulatory frameworks that govern the ethical use of sensor data. Sociotechnical approaches like this can also have a vital role in AI design[160,161], particularly in the public sector[162,163], as well as fostering global public awareness and education[164], ensuring that the implications of sensor technology are widely understood.

Conclusion

This Perspective highlights the dimensions of risk associated with AI systems that arise from the material affordances of sensors and their underlying calculative models. It proposes a sensor-sensitive framework for diagnosing these risks, complementing existing approaches such as the NIST AI RMF and the EU AI Act. A key advantage of this framework is its emphasis on a broader range of stakeholders involved in risk diagnostics and management. While guidelines like the NIST AI RMF focus primarily on actors throughout the AI lifecycle, and the EU AI Act primarily applies to providers and deployers of AI systems, the sensor-sensitive approach brings attention to often-overlooked stakeholders. For example, actors involved in the manufacturing, production and evaluation of sensors have a crucial role in mitigating risks related to calibration, documentation and waste management. In addition, the risk of proprietary data profusion illustrates how economic policy and regulation can complement technical solutions. Thus, the sensor-sensitive approach fills a gap in the existing AI ethics and governance discourse by considering how sensors contribute to the risks of AI systems and implicate key actors beyond those directly involved in the AI lifecycle.

We call for urgent attention to be directed towards developing responsible sensor architectures and regulatory frameworks concerning sensor development and associated data usage. While some recent work provides a commendable starting point, much remains to be done in this evolving field. Furthermore, engaging stakeholders and communities is essential to fully harness the benefits of sensor technologies within a new sensing paradigm. This is particularly important as the risks identified are not exhaustive and may evolve due to changes in the material affordances of sensors, shifts in calculative models or ongoing AI innovations leading to new use cases.

References

1. Metcalf, J., Moss, E., Watkins, E. A., Singh, R. & Elish, M. C. Algorithmic impact assessments and accountability: the co-construction of impacts. In *Proc. 2021 ACM Conference on Fairness, Accountability, and Transparency* 735–746 (ACM, 2021).
2. Shelby, R. et al. Sociotechnical harms of algorithmic systems: scoping a taxonomy for harm reduction. In *Proc. 2023 AAAI/ACM Conference on AI, Ethics, and Society* 723–741 (ACM, 2023).
3. Weidinger, L. et al. Ethical and social risks of harm from language models. Preprint at https://arxiv.org/abs/2112.04359 (2021).
4. Selbst, A. D. An institutional view of algorithmic impact. *Harv. J. Law Technol.* **35**, 117 (2021).
5. Birhane, A. Algorithmic injustice: a relational ethics approach. *Patterns* https://doi.org/10.1016/j.patter.2021.100205 (2021).
6. Barocas, S. & Selbst, A. D. Big data's disparate impact. *Calif. Law Rev.* **104**, 671–732 (2016).
7. Noble, S. U. *Algorithms of Oppression: Data Discrimination in the Age of Google* (New York Univ. Press, 2018).
8. Eubanks, V. *Automating Inequality: How High-tech Tools Profile, Police, and Punish the Poor* (St. Martin's Press, 2018).
9. Danks, D. & London, A. J. Algorithmic bias in autonomous systems. In *Proc. 26th International Joint Conference on Artificial Intelligence* Vol. 17, 4691–4697 (AAAI, 2017).
10. Mitchell, S., Potash, E., Barocas, S., D'Amour, A. & Lum, K. Algorithmic fairness: choices, assumptions, and definitions. *Annu. Rev. Stat. Appl.* **8**, 141–163 (2021).
11. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K. & Galstyan, A. A survey on bias and fairness in machine learning. *ACM Comput. Surv.* **54**, 1–35 (2021).
12. Hazirbas, C. et al. Towards measuring fairness in AI: the casual conversations dataset. *IEEE Trans. Biom. Behav. Identity Sci.* **4**, 324–332 (2021).
13. Ehsan, U., Singh, R., Metcalf, J. & Riedl, M. The algorithmic imprint. In *Proc. 2022 ACM Conference on Fairness, Accountability, and Transparency* 1305–1317 (ACM, 2022).
14. Dodge, J. et al. Measuring the carbon intensity of AI in cloud instances. In *Proc. 2022 ACM Conference on Fairness, Accountability, and Transparency* 1877–1894 (ACM, 2022).
15. Bender, E. M., Gebru, T., McMillan-Major, A &, Shmitchell, S. On the dangers of stochastic parrots: can language models be too big? In *Proc. 2021 ACM Conference on Fairness, Accountability, and Transparency* 610–623 (ACM, 2021).
16. Weidinger, L. et al. Taxonomy of risks posed by language models. In *Proc. 2022 ACM Conference on Fairness, Accountability, and Transparency* 214–229 (ACM, 2022).
17. Bridges, L. et al. Geographies of digital wasting: electronic waste from mine to discard and back again; https://www.geographiesofdigitalwasting.com/
18. Kidd, M. Energy and e-waste: the AI tsunamis. *DCD* https://www.datacenterdynamics.com/en/opinions/energy-and-e-waste-the-ai-tsunamis/ (2023).
19. Law, J. & Mol, A. Notes on materiality and sociality. *Sociol. Rev.* **43**, 274–294 (1995).
20. Pinch, T. Technology and instituions: living in a material world. *Theory Soc.* **37**, 461–483 (2008).
21. Lievrouw, L. A. & Livingstone, S. in *Handbook of New Media: Social Shaping and Social Consequences of ICTs* 1–14 (2006).
22. Miller, D. *Materiality* (Duke Univ. Press, (2020).
23. Warden, P. & Situnayake, D. *TinyML: Machine Learning with Tensorflow Lite on Arduino and Ultra-low-power Microcontrollers* (O'Reilly Media, 2019).
24. Janapa Reddi, V. et al. Widening access to applied machine learning with tinyML. *Harv. Data Sci. Rev.* https://doi.org/10.1162/99608f92.762d171a (2022).
25. Gillespie, T., Boczkowski, P. J. & Foot, K. A. in *Media Technologies: Essays on Communication, Materiality, and Society* 21–51 (MIT Press, 2013)
26. Vallor, S. *Technology and the Virtues: A Philosophical Guide to A Future Worth Wanting* (Oxford Univ. Press, 2016).
27. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (National Institute of Standards and Technology, 2023); http://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf
28. Veale, M. & Zuiderveen Borgesius, F. Demystifying the draft EU Artificial Intelligence Act—analysing the good, the bad, and the unclear elements of the proposed approach. *Comput. Law Rev. Int.* **22**, 97–112 (2021).
29. Mukhopadhyay, S. C. et al. Artificial intelligence-based sensors for next generation IoT applications: a review. *IEEE Sens. J.* **21**, 24920–24932 (2021).
30. Singh, R. & Gill, S. S. Edge AI: a survey. *Internet Things Cyber Phys. Syst.* **3**, 71–92 (2023).
31. Haick, H. & Tang, N. Artificial intelligence in medical sensors for clinical decisions. *ACS Nano* **15**, 3557–3567 (2021).
32. Zhang, L. & Zhang, L. Artificial intelligence for remote sensing data analysis: a review of challenges and opportunities. *IEEE Geosc. Remote Sens. Mag.* **10**, 270–294 (2022).
33. Masson, J. F. *Roadmap for the Use of Machine Learning and Artificial Intelligence in Sensing* (ACS, 2024).
34. Ullo, S. L. & Sinha, G. R. Advances in IoT and smart sensors for remote sensing and agriculture applications. *Remote Sens.* **13**, 2585 (2021).
35. Merenda, M., Porcaro, C. & Iero, D. Edge machine learning for AI-enabled IoT devices: a review. *Sensors* **20**, 2533 (2020).
36. Zhu, G. et al. Pushing AI to wireless network edge: an overview on integrated sensing, communication, and computation towards 6G. *Sci. China Inf. Sci.* **66**, 130301 (2023).
37. Abadade, Y. et al. A comprehensive survey on TinyML. *IEEE Access* **11**, 96892–96922 (2023).
38. Dutta, L. & Bharali, S. TinyML meets IoT: a comprehensive survey. *Internet Things* **16**, 100461 (2021).
39. Gaver, W. W. Technology affordances. In *Proc. SIGCHI Conference on Human Factors in Computing Systems* 79–84 (ACM, 1991).
40. Gibson, J. In *Perceiving, Acting and Knowing: Toward an Ecological Psychology* (eds Shaw, R. & Bransford, J.) 1st edn (1977).
41. Davis, J. L. *How Artifacts Afford: the Power and Politics of Everyday Things* (MIT Press, 2020).
42. Kennewell, S. Using affordances and constraints to evaluate the use of information and communications technology in teaching and learning. *J. Inf. Techol. Teacher Educ.* **10**, 101–116 (2001).
43. Acemoglu, D. *Harms of AI* (National Bureau of Economic Research, 2021).
44. Kusche, I. Possible harms of artificial intelligence and the EU AI Act: fundamental rights and risk. *J. Risk Res.* https://doi.org/10.1080/13669877.2024.2350720 (2024).
45. Watkins, E. A., Moss, E., Metcalf, J., Singh, R. & Elish, M. C. Governing algorithmic systems with impact assessments: six observations. In *Proc. 2021 AAAI/ACM Conference on AI, Ethics, and Society* 1010–1022 (ACM, 2021).
46. Bengio, Y. et al. Managing extreme AI risks amid rapid progress. *Science* **384**, 842–845 (2024).
47. Roth, L. Looking at Shirley, the ultimate norm: colour balance, image technologies, and cognitive equity. *Can. J. Commun.* **34**, 111–136 (2009).
48. Galdino, G. M., Vogel, J. E. & Vander Kolk, C. A. Standardizing digital photography: it's not all in the eye of the beholder. *Plas. Reconstr. Surg.* **108**, 1334–1344 (2001).

49. Guo, C. Y., Huang, W. Y., Chang, H. C. & Hsieh, T. L. Calibrating oxygen saturation measurements for different skin colors using the individual typology angle. *IEEE Sens. J.* **23**, 16993–17001 (2023).
50. Buolamwini, J. & Gebru, T. Gender shades: intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency* 77–91 (PMLR, 2018).
51. Muniesa, F., Millo, Y. & Callon, M. An introduction to market devices. *Sociol. Rev.* **55**, 1–12 (2007).
52. Callon, M. & Law, J. On qualculation, agency, and otherness. *Environ. Plan. D* **23**, 717–733 (2005).
53. IoT device detects wind turbine faults in the field. https://www.engineering.com/iot-device-detects-wind-turbine-faults-in-the-field/ (2020).
54. Restle, P. J. et al. The clock distribution of the power4 microprocessor. In *2002 IEEE International Solid-State Circuits Conference. Digest of Technical Papers* Vol. 1, 144–145 (IEEE, 2002).
55. Flores, T. et al. TinyML for safe driving: the use of embedded machine learning for detecting driver distraction. In *2023 IEEE International Workshop on Metrology for Automotive (MetroAutomotive)* 62–66 (IEEE, 2023).
56. Shah, A. S., Nasir, H., Fayaz, M., Lajis, A. & Shah, A. A review on energy consumption optimization techniques in IoT based smart building environments. *Information* **10**, 108 (2019).
57. *Transforming our World: The 2030 Agenda for Sustainable Development* (United Nations, 2015); https://sdgs.un.org/publications/transforming-our-world-2030-agenda-sustainable-development-17981
58. Prakash, S. et al. Is tinyML sustainable? Assessing the environmental impacts of machine learning on microcontrollers. *Commun. ACM* **66**, 68–77 (2023).
59. Mrisho, L. M. et al. Accuracy of a smartphone-based object detection model, PlantVillage Nuru, in identifying the foliar symptoms of the viral diseases of cassava-CMD and CBSD. *Front. Plant Sci.* **11**, 590889 (2020).
60. King, A. Technology: the future of agriculture. *Nature.* **544**, S21–S23 (2017).
61. Solana, A. Elephants vs trains: this is how AI helps ensure they don't collide. *ZDNET* https://www.zdnet.com/article/elephants-vs-trains-this-is-how-ai-helps-ensure-they-dont-collide/ (2020).
62. Temple-Raston, D. Using AI in Malawi to save elephants. *NPR* https://www.npr.org/2019/09/17/761682912/using-ai-in-malawi-to-save-elephants (2019).
63. Johnson, K. Google's AI powers real-time orca tracking in Vancouver Bay. *VentureBeat* https://venturebeat.com/ai/googles-ai-powers-real-time-orca-tracking-in-vancouver-bay/ (2020).
64. Elmqvist, N. Data analytics anywhere and everywhere. *Commun. ACM* **66**, 52–63 (2023).
65. Zuboff, S. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* 1st edn (PublicAffairs, 2018).
66. Fernback, J. Sousveillance: communities of resistance to the surveillance environment. *Telemat. Inform.* **30**, 11–21 (2013).
67. Monahan, T. Regulating belonging: surveillance, inequality, and the cultural production of abjection. *J. Cult. Econ.* **10**, 191–206 (2017).
68. Sevignani, S. Surveillance, classification, and social inequality in informational capitalism: the relevance of exploitation in the context of markets in information. *Hist. Soc. Res.* **42**, 77–102 (2017).
69. Gilman, M. & Green, R. The surveillance gap: the harms of extreme privacy and data marginalization. *NYU Rev. Law Soc. Change* **42**, 253 (2018).
70. Parsons, C. Beyond privacy: articulating the broader harms of pervasive mass surveillance. *Media Commun.* **3**, 1–11 (2015).
71. *AI Act* (European Commission, 2025); https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
72. Wilson, J. S. *Sensor Technology Handbook* (Elsevier, 2004).
73. Hacking, I. In *Biopower: Foucault and Beyond* (eds Cisney, V. W. & Morar, N.) 65–81 (Univ. Chicago Press, 2015).
74. Krizhevsky, A., Sutskever, I. & Hinton, G.E. Imagenet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **25**, (2012).
75. Ardila, R. et al. Common voice: a massively-multilingual speech corpus. Preprint at https://arxiv.org/abs/1912.06670 (2019).
76. Cowan, R. S. in *The Social Shaping of Technology: How the Refrigerator Got its Hum* (eds MacKenzie, D. A. & Wajcman, J.) 202–218 (Open Univ. Press, 1985).
77. Gibson, J. J. in *The People, Place, and Space Reader* (eds Gieseking, J. J. et al.) 56–60 (Routledge, 2014).
78. Buratti, C., Conti, A., Dardari, D. & Verdone, R. An overview on wireless sensor networks technology and evolution. *Sensors* **9**, 6869–6896 (2009).
79. Mainetti, L., Patrono, L. & Vilei, A. Evolution of wireless sensor networks towards the internet of things: a survey. In *SoftCOM 2011, 19th International Conference on Software, Telecommunications and Computer Networks* 1–6 (IEEE, 2011).
80. Warden, P., Stewart, M., Plancher, B., Katti, S. & Reddi, V. J. Machine learning sensors: a design paradigm for the future of intelligent sensors. *Commun. ACM* **66**, 25–28 (2023).
81. Zhang, Y., Gu, Y., Vlatkovic, V. & Wang, X. Progress of smart sensor and smart sensor networks. In *Fifth World Congress on Intelligent Control and Automation* Vol. 4, 3600–3606 (IEEE, 2004).
82. Vetelino, J. & Reghu, A. *Introduction to Sensors* (CRC Press, 2017).
83. Soloman, S. *Sensors Handbook* (McGraw-Hill, 2009).
84. Li, S., Xu, L. D. & Zhao, S. The Internet of Things: a survey. *Inf. Syst. Front.* **17**, 243–259 (2015).
85. Rose, K., Eldridge, S. & Chapin, L. *The Internet of Things: An Overview* (The Internet Society, 2015).
86. Sehrawat, D. & Gill, N. S. Smart sensors: analysis of different types of IoT sensors. In *2019 3rd International Conference on Trends in Electronics and Informatics* 523–528 (IEEE, 2019).
87. Krishnamurthi, R., Kumar, A., Gopinathan, D., Nayyar, A. & Qureshi, B. An overview of IoT sensor data processing, fusion, and analysis techniques. *Sensors* **20**, 6076 (2020).
88. Kocakulak, M. & Butun, I. An overview of wireless sensor networks towards Internet of Things. In *2017 IEEE 7th Annual Computing and Communication Workshop and Conference* 1–6 (IEEE, 2017).
89. Vincent, D. R. et al. Sensors driven AI-based agriculture recommendation model for assessing land suitability. *Sensors* **19**, 3667 (2019).
90. Fabre, W., Haroun, K., Lorrain, V., Lepecq, M. & Sicard, G. From near-sensor to in-sensor: a state-of-the-art review of embedded AI vision systems. *Sensors* **24**, 5446 (2024).
91. Wen, D. et al. Task-oriented sensing, computation, and communication integration for multi-device edge AI. *IEEE Trans. Wirel. Commun.* **23**, 2486–2502 (2023).
92. Sodhro, A. H., Pirbhulal, S. & De Albuquerque, V. H. C. Artificial intelligence-driven mechanism for edge computing-based industrial applications. *IEEE Trans. Industr. Inform.* **15**, 4235–4243 (2019).
93. Stewart, M. et al. Datasheets for machine learning sensors: towards transparency, auditability, and responsibility for intelligent sensing. Preprint at https://doi.org/10.48550/arXiv.2306.08848 (2024).
94. Callon, M. & Muniesa, F. Peripheral vision: economic markets as calculative collective devices. *Organ. Stud.* **26**, 1229–1250 (2005).

95. Hansen, K. B. Model talk: calculative cultures in quantitative finance. *Sci. Technol. Hum. Values* **46**, 600–627 (2021).
96. Besedovsky, N. Financialization as calculative practice: the rise of structured finance and the cultural and calculative transformation of credit rating agencies. *Socioecon. Rev.* **16**, 61–84 (2018).
97. MacKenzie, D. *An Engine, Not a Camera: How Financial Models Shape Markets* 1st edn (MIT Press, 2008).
98. *2019 Manufacturing Trends Report* (Microsoft Dynamics 365, 2019); https://info.microsoft.com/rs/157-GQE-382/images/EN-US-CNTNT-Report-2019-Manufacturing-Trends.pdf
99. Huck, C. W. In *Sense the Real Change: Proc. 20th International Conference on Near Infrared Spectroscopy* (eds Chu, X. et al.) 59–72 (Springer Nature, 2022).
100. Roy, R. & Miller, J. Miniaturization of image sensors: the role of innovations in complementary technologies in overcoming technological trade-offs associated with product innovation. *J. Eng. Technol. Manag.* **44**, 58–69 (2017).
101. Rodriguez-Saona, L., Aykas, D. P., Borba, K. R. & Urtubia, A. Miniaturization of optical sensors and their potential for high-throughput screening of foods. *Curr. Opin. Food Sci.* **31**, 136–150 (2020).
102. Tricoli, A., Nasiri, N. & De, S. Wearable and miniaturized sensor technologies for personalized and preventive medicine. *Adv. Funct. Mater.* **27**, 1605271 (2017).
103. Frazier, A. B., Warrington, R. O. & Friedrich, C. The miniaturization technologies: past, present, and future. *IEEE Trans. Industr. Electron.* **42**, 423–430 (1995).
104. Madou, M. J. *Fundamentals of Microfabrication: The Science of Miniaturization* (CRC Press, 2018).
105. Yang, Z., Albrow-Owen, T., Cai, W. & Hasan, T. Miniaturization of optical spectrometers. *Science* **371**, eabe0722 (2021).
106. Jiang, C. et al. Energy aware edge computing: a survey. *Comput. Commun.* **151**, 556–580 (2020).
107. Chen, Y. et al. Energy efficient dynamic offloading in mobile edge computing for Internet of Things. *IEEE Trans. Cloud Comput.* **9**, 1050–1060 (2019).
108. Rault, T., Bouabdallah, A. & Challal, Y. Energy efficiency in wireless sensor networks: a top-down survey. *Comput. Netw.* **67**, 104–122 (2014).
109. Sun, H. et al. MEMS based energy harvesting for the Internet of Things: a survey. *Microsyst. Technol.* **24**, 2853–2869 (2018).
110. Raha, A. & Raghunathan, V. Towards full-system energy–accuracy tradeoffs: a case study of an approximate smart camera system. In *Proc. 54th Annual Design Automation Conference 2017* 1–6 (ACM, 2017).
111. Schurgers, C. & Srivastava, M. B. Energy efficient routing in wireless sensor networks. In *2001 MILCOM Proceedings Communications for Network-centric Operations: Creating the Information Force* Vol. 1, 357–361 (IEEE, 2001).
112. *State of IoT—Spring 2023* (IoT Analytics, 2023); https://iot-analytics.com/product/state-of-iot-spring-2023/
113. Saif, I. & Ammanath, B. 'Trustworthy AI' is a framework to help manage unique risk. *MIT Technology Review* https://www.technologyreview.com/2020/03/25/950291/trustworthy-ai-is-a-framework-to-help-manage-unique-risk/ (2020).
114. Floridi, L. et al. capAI: a procedure for conducting conformity assessment of AI systems in line with the EU Artificial Intelligene Act. Preprint at *SSRN* https://doi.org/10.2139/ssrn.4064091 (2022).
115. *Advancing Accountability in AI: Governing and Managing Risks throughout the Lifecycle for Trustworthy AI* OECD Digital Economy Papers Vol. 349 (OECD, 2023); https://www.oecd-ilibrary.org/science-and-technology/advancing-accountability-in-ai_2448f04b-en
116. Baquero, J. A., Burkhardt, R., Govindarajan, A. & Wallace, T. *Derisking AI: Risk Management in AI Development* (McKinsey, 2020).
117. *AI and Risk Management* (Deloitte, 2018); https://www2.deloitte.com/content/dam/Deloitte/nl/Documents/innovatie/deloitte-nl-innovate-lu-ai-and-risk-management.pdf
118. Whitehouse, K. & Culler, D. Calibration as parameter estimation in sensor networks. In *Proc. 1st ACM International Workshop on Wireless Sensor Networks and Applications* 59–67 (ACM, 2002).
119. Hansen, J. H. & Boril, H. Gunshot detection systems: methods, challenges, and can they be trusted? In *Audio Engineering Society Convention* Convention Paper 10540 (AES, 2021).
120. Delaine, F., Lebental, B. & Rivano, H. In situ calibration algorithms for environmental sensor networks: a review. *IEEE Sens. J.* **19**, 5968–5978 (2019).
121. Leveson, N. G. *Engineering a Safer World: Systems Thinking Applied to Safety* (MIT Press, 2016).
122. Dewey, F. R. A complete guide to data sheets. *Sensors Magazine* (1998); https://www.allegromicro.com/-/media/files/technical-documents/ complete-guide-to-datasheets-pub26000.pdf
123. Mitchell, M. et al. Model cards for model reporting. In *Proc. Conference on Fairness, Accountability, and Transparency* 220–229 (ACM, 2019).
124. Mitev, R., Pazii, A., Miettinen, M., Enck, W. & Sadeghi, A. R. Leakypick: IoT audio spy detector. In *Proc. 36th Annual Computer Security Applications Conference* 694–705 (ACM, 2020).
125. Anthes, G. Data brokers are watching you. *Commun. ACM* **58**, 28–30 (2015).
126. Moyopo, S. Quantifying the data currency's impact on the profit made by data brokers in the Internet of Things based data marketplace. *Eur. J. Electr. Eng. Comput. Sci.* **7**, 7–16 (2023)
127. Crain, M. The limits of transparency: data brokers and commodification. *New Media Soc.* **20**, 88–104 (2018).
128. Teh, H. Y., Kempa-Liehr, A. W. & Wang, K. I. K. Sensor data quality: a systematic review. *J. Big Data* **7**, 11 (2020).
129. Ye, J., Stevenson, G. & Dobson, S. Detecting abnormal events on binary sensors in smart home environments. *Pervasive Mob. Comput.* **33**, 32–49 (2016).
130. D'ignazio, C. & Klein, L. F. *Data Feminism* (MIT Press, 2023).
131. O'Neil, C. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (Crown, 2017).
132. Crawford, K. The hidden biases in big data. *Harvard Business Review* (2013); https://hbr.org/2013/04/the-hidden-biases-in-big-data
133. Ridzuan, F. & Zainon, W. M. N. W. A review on data cleansing methods for big data. *Procedia Comput. Sci.* **161**, 731–738 (2019).
134. Scott, E. *The Trouble with Informed Consent in Smart Cities* (IAAP, 2019).
135. Froomkin, A. M. Big data: destroyer of informed consent. *Yale J. Law Technol.* **21**, 27 (2019).
136. Elmenreich, W. *An Introduction to Sensor Fusion* (Vienna Univ. Technology, 2002).
137. Sweeney, L. Simple demographics often identify people uniquely. *Health* **671**, 1–34 (2000).
138. Barocas, S. & Nissenbaum, H. Big data's end run around procedural privacy protections. *Commun. ACM* **57**, 31–33 (2014).
139. Ding, W., Jing, X., Yan, Z. & Yang, L. T. A survey on data fusion in Internet of Things: towards secure and privacy-preserving fusion. *Inf. Fusion* **51**, 129–144 (2019).
140. Dhar, P. The carbon impact of artificial intelligence. *Nat. Mach. Intell.* **2**, 423–425 (2020).
141. Wu, C. J. et al. Sustainable AI: environmental implications, challenges and opportunities. *Proc. Mach. Learn. Syst.* **4**, 795–813 (2022).

142. Van Wynsberghe, A. Sustainable AI: AI for sustainability and the sustainability of AI. *AI Ethics* **1**, 213–218 (2021).
143. Lannelongue, L., Grealey, J. & Inouye, M. Green algorithms: quantifying the carbon footprint of computation. *Adv. Sci.* **8**, 2100707 (2021).
144. Cooper, Z. G. T. Of dog kennels, magnets, and hard drives: dealing with big data peripheries. *Big Data Soc.* **8**, 20539517211015430 (2021).
145. Bridges, L. E. Material entanglements of community surveillance & infrastructural power. *AoIR Selected Papers of Internet Research, 2020* https://doi.org/10.5210/spir.v2020i0.11179 (2020).
146. Gupta, U. et al. Chasing carbon: the elusive environmental footprint of computing. In *2021 IEEE International Symposium on High-Performance Computer Architecture* 854–867 (IEEE, 2021).
147. Ozer, E. et al. Bendable non-silicon RISC-V microprocessor. *Nature* **634**, 341–346 (2024).
148. Sorrell, S. Jevons' paradox revisited: the evidence for backfire from improved energy efficiency. *Energy Policy* **37**, 1456–1469 (2009).
149. Verbeek, P. P. Ambient intelligence and persuasive technology: the blurring boundaries between human and technology. *Nanoethics.* **3**, 231–242 (2009).
150. Nafus, D. *Quantified: Biosensing Technologies in Everyday Life* (MIT Press, 2016).
151. Wang, X., McGill, T. J. & Klobas, J. E. I want it anyway: consumer perceptions of smart home devices. *J. Comput. Inf. Syst*. **60**, 1–11 (2018).
152. Balta-Ozkan, N., Davidson, R., Bicket, M. & Whitmarsh, L. Social barriers to the adoption of smart homes. *Energy Policy* **63**, 363–374 (2013).
153. Joint Task Force Transformation Initiative *Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy* NIST SP 800-37r2 (National Institute of Standards and Technology, 2018); https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-37r2.pdf
154. Rahwan, I. et al. Machine behaviour. *Nature* **568**, 477–486 (2019).
155. Soltoggio, A. et al. A collective AI via lifelong learning and sharing at the edge. *Nat. Mach. Intell.* **6**, 251–264 (2024).
156. *Algorithmic Impact Assessment: A Case Study in Healthcare* (Ada Lovelace Institute, 2022); https://www.adalovelaceinstitute.org/report/algorithmic-impactasssessment-case-study-healthcare
157. Metcalf, J. et al. A relationship and not a thing: a relational approach to algorithmic accountability and assessment documentation. Preprint at https://arxiv.org/abs/2203.01455 (2022).
158. Lavin, A. et al. Technology readiness levels for machine learning systems. *Nat. Commun.* **13**, 6039 (2022).
159. Huckelberry, J. et al. TinyML security: exploring vulnerabilities in resource-constrained machine learning systems. Preprint https://arxiv.org/abs/2411.07114 (2024).
160. Baxter, G. & Sommerville, I. Socio-technical systems: from design methods to systems engineering. *Interact. Comput.* **23**, 4–17 (2011).
161. Bauer, J. M. & Herder, P. M. in *Philosophy of Technology and Engineering Sciences* (ed. Meijers, A.) 601–630 (Elsevier, 2009).
162. Leslie, D. Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector. *Zenodo* https://doi.org/10.5281/zenodo.3240528 (2019).
163. Abbas, R., Pitt, J. & Michael, K. Socio-technical design for public interest technology. *IEEE Trans. Technol. Soc.* **2**, 55–61 (2021).
164. Plancher, B. et al. TinyML4D: scaling embedded machine learning education in the developing world. In *Proc. AAAI Symposium Series* Vol. 3, 508–515 (AAAI, 2024).

Author contributions

M. Stewart and V.J.R. organized the exploratory seminar that led to this paper. All authors developed the problem statement. M. Sloane developed the theoretical approach. M. Stewart, E.M., S.K., B.P., M.P.S. and V.J.R. designed and executed the analytical approach. S.K., M. Sloane and E.M. conducted the policy analysis. B.P., M. Stewart, V.J.R. and P.W. produced the technical framework and historical analysis. M. Sloane, E.M., S.K., M. Stewart, B.P. and V.J.R. wrote the paper with input from all authors. B.P. managed the layout, formatting and figure design. M. Sloane managed the authorship, submission and revision process.

Competing interests

P.W. is a founder and major shareholder of Useful Sensors Inc., which works on privacy-preserving sensor technology. The other authors declare no competing interests.

Additional information

Correspondence should be addressed to Mona Sloane.

Peer review information *Nature Machine Intelligence* thanks Carina Prunkl and Andrea Soltoggio for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.