Understanding accountability in algorithmic supply chains

Jennifer Cobbe
University of Cambridge
United Kingdom

Michael Veale
University College London
United Kingdom

Jatinder Singh
University of Cambridge
United Kingdom

ABSTRACT

Academic and policy proposals on algorithmic accountability often seek to understand algorithmic systems in their socio-technical context, recognising that they are produced by 'many hands'. Increasingly, however, algorithmic systems are also produced, deployed, and used within a supply chain comprising multiple actors tied together by flows of data between them. In such cases, it is the working together of an *algorithmic supply chain* of different actors who contribute to the production, deployment, use, and functionality that drives systems and produces particular outcomes. We argue that algorithmic accountability discussions must consider supply chains and the difficult implications they raise for the governance and accountability of algorithmic systems. In doing so, we explore algorithmic supply chains, locating them in their broader technical and political economic context and identifying some key features that should be understood in future work on algorithmic governance and accountability (particularly regarding general purpose AI services). To highlight ways forward and areas warranting attention, we further discuss some implications raised by supply chains: challenges for allocating accountability stemming from *distributed responsibility* for systems between actors, limited visibility due to the *accountability horizon*, service models of use and liability, and cross-border supply chains and regulatory arbitrage.

CCS CONCEPTS

• **Social and professional topics → Computing / technology policy**; **Socio-technical systems**.

KEYWORDS

Algorithmic accountability, supply chains, AI as a Service, general purpose AI, political economy, accountability horizon

1 INTRODUCTION

Accountability for complex systems is a challenge in many domains. The 'many hands' problem holds that accountability is difficult where many people have together contributed to activity or outcome, as it may be impossible to allocate responsibility to any one of them [11, 77, 104]. Writing in 1996, Nissenbaum argued that computer systems raise a particular form of this problem – they are usually not the product of an individual programmer, but of groups or organisations, and may include components that were developed by others [77]. Addressing the 'many hands' problem is increasingly a concern in the algorithmic accountability literature (implicitly or explicitly), with proposals in recent years for accountability for algorithmic systems to operate at an organisational level [25, 40, 60, 62, 67, 75, 89, 98, 110].

Yet today's computer systems—including AI technologies—are increasingly modular, driven or underpinned by cloud-based technologies, and interconnected. The 'agile turn' of recent decades brought several key changes in software development, distribution, and infrastructure, directly influencing how businesses are organised and computing resources are distributed [46]. First *agile development* means that software (including 'AI' models) is now produced in short development cycles with continuous testing, revision, and iterative improvement after deployment [46]. Second, software is now generally modularised and distributed *as a service*, with a client-server model in which the server performs the computation [46, 74]. Third, the challenges of scaling services and increasingly portable client devices have driven advances in data centres with flexible resources and software becoming increasingly *cloud-based* [46, 74]. Consequently, software development now often involves, to various degrees, integrating pre-built modular components provided *as a service* and controlled by others into a complete product: not simply a system, but a *system-of-systems* [99].

As a result, digital technologies across society and the economy are increasingly organised around *data-driven supply chains* involving several interconnected actors and their systems. In these supply chains, data flows from actor to actor linking systems designed, developed, owned, and controlled by different people and organisations and often offered as services [79]: a sensor system (controlled by one actor) might connect to an analytics system (controlled by another) which itself outputs into a decision-making system (controlled by a third). This is often so even for seemingly simple applications; for example, a home thermostat can be driven by data from a national weather service, which is itself fed data from thermometers owned and operated by different actors. In such supply chains, the *working together* of services and systems controlled by different actors produces particular outcomes—hardware capabilities, software functionalities, the workings of commercial and industrial processes, 'AI' decisions and outputs, and more. Supply chains are *data-driven* in that the flow of data between actors links them together, allowing a system controlled by one actor to interact with those controlled by others and together produce some functionality [24, 99]. In the context of AI and algorithmic systems, **algorithmic supply chains** are those where several actors contribute towards the production, deployment, use, and functionality of AI technologies. In these supply chains, AI 'as a service' providers often play key roles [26].

By reconfiguring software development and distribution, the agile turn also had significant political economic ramifications and other implications [24, 46]. In bringing services together to produce functionality through supply chains, developers now delegate control over much of the underlying technologies to other actors, complicating the governance of such technologies and the products they are part of. It is no longer necessarily the case that computer systems are produced by a group of developers or an organisation, or by a vendor simply integrating various standalone components into one product (which itself raises the 'many hands' problem [77]).

To appear at 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)
https://doi.org/10.1145/3593013.3594073

Instead, as we argue, they often now involve a group of organisations arranged together in a data-driven supply chain, *each retaining control over component systems they provide as services to others.* Moreover, certain key actors in these supply chains—in particular, major cloud providers who often control underlying technologies—provide many kinds of services to millions of customers, taking up important positions across supply chains in many sectors [24, 26]. The agile turn, therefore, has reorganised many areas of social and economic life – now reconstituted around service-based supply chains with a few systemically important actors providing the core infrastructure that underpins contemporary society.

The agile turn's shift to data-driven supply chains brings significant implications for governance and accountability frameworks and mechanisms relevant to algorithmic systems (including privacy and data protection [46]). Allocating accountability across supply chain actors for producing, deploying, and using algorithmic systems is relevant to general academic, policy, and regulatory discussions and proposals around algorithmic accountability, as well as to more specific legislative efforts around regulation of AI. Here we argue that *governance of and accountability for algorithmic systems* as deployed and used in the real world needs to *operate across the supply chains* that already and will increasingly underpin, drive, and produce the outputs and effects of many AI technologies.

Much of the policy and academic literature, however, is grounded in an organisation-focused understanding of digital technologies. Even more recent work on algorithmic accountability which seeks to address the 'many hands' problem through a relatively broad view of accounting for algorithmic systems is typically focused on making specific stages of system lifecycle more transparent and understandable [7, 30, 40, 75, 85] or framed around the perspective of a single organisation [25, 89, 109, 110]. This attention now paid to the accountability of organisations for their algorithmic systems was long overdue, but the focus on organisational accountability has largely obscured the dynamics of contemporary algorithmic supply chains. We therefore still lack ways to conceive of these chains, to bring them within legal, regulatory, and governance mechanisms, and to appropriately distribute responsibility and accountability.

This paper contributes to understanding these challenges. First (§2), we discuss recent trends in algorithmic accountability and identify limitations regarding supply chains. Next (§3), we describe AI services and algorithmic supply chains, locate them within their broader context, and identify four key features: interdependence (§3.1), process (§3.2), integration (§3.3), and consolidation (§3.4). Then (§4) we discuss important implications for accountability these features raise, around the distributed nature of responsibility in supply chains (§4.1); the limited understanding individual actors may have of the broader chain due to the 'accountability horizon' (§4.2); the move to a services model and attempts by providers to maximise control and minimise liability (§4.3); and the cross-border nature of some supply chains (§4.4).

In all, we argue, algorithmic accountability work must urgently address the technological, legal, and political economic dynamics of algorithmic supply chains. We do not offer concrete suggestions to implement at a technical or organisational level to improve accountability in these supply chains, but instead hope to produce a shift in focus for algorithmic accountability as a field and indicate new research directions

2 ACCOUNTABILITY IN ALGORITHMIC SYSTEMS

Significant academic and policy work has sought various forms of accountability of and for algorithmic systems [110]. Accountability may be seen either as a *mechanism* or a *virtue* [13]. As a mechanism (an understanding particularly held in Europe and in non-US Anglophone countries), it is an institutional arrangement whereby an *actor* provides an account to a *forum*, who deliberates on that account and may impose consequences [12]. A developer might provide information to a regulator about their system, for example, with the regulator then issuing a penalty or requiring design changes. Some algorithmic accountability literature explicitly views accountability as a mechanism for holding actors to account for systems they are responsible for [1, 25, 61, 109]. By contrast, accountability as a virtue (typically a US interpretation) is as a normative concept, a set of standards for evaluating behaviour—often tied to being transparent, responsible, and responsive—with 'being accountable' seen as a positive quality of particular actors [13]. Some (predominantly technical) work has pursued it as a virtue and sought to improve the accountability of certain technologies by imbuing them with such positive qualities. Yet applying accountability to algorithmic systems in this way—rather than to the organisations responsible for them—often equates accountability with technical functionality (for example, building 'Accountable AI') rather with human virtues which are not reducible to technically tractable concepts [61].

We treat accountability as a mechanism, whereby actors are held accountable for technologies they are responsible for. However, accountability for digital technologies is typically challenging. The 'many hands' problem—that it is difficult to hold any one person responsible for an outcome where multiple people helped produce it—has long been recognised: computer systems are rarely produced by an individual who can be held responsible, but by teams and organisations with many people contributing [77]. Moreover, modular software development—where software developed by one organisation uses a library developed by another, for example—further complicates things [77]. Much software is too complex, relying on too many components, for any one person to fully understand or account for all of its workings.

In the context of algorithmic accountability, specifically, a key conceptual shift has been in understanding these systems not as 'algorithms' but as socio-technical *algorithmic systems*: "intricate, dynamic arrangements of people and code" [96]. This recognises that 'algorithms' are produced and work within broader human contexts and in practice cannot be understood separately from them. Simultaneously, explanations of (ML) model workings are increasingly recognised as insufficient to account for algorithmic systems [25, 35, 109]. Much research has therefore gradually moved away from seeking transparency or explanations of models (though this remains an important area of work) to understanding algorithmic systems more broadly as socio-technical phenomena. Much of this reflects—implicitly or explicitly—an understanding that algorithmic systems are often the result of 'many hands': produced by and deployed and used within teams and organisations. To understand and account for an algorithmic system, one needs to understand and account for the collective efforts of the organisational processes involved in producing, deploying, and using it.

The term ‘algorithmic system’ is now widely used in algorithmic accountability discussions, with academic and policy literature commonly suggesting ways to improve accountability for their organisational aspects. Some proposals seek lower-level mechanisms to document the choices and decisions made by people in developing, deploying, or using a system, such as proposals for datasheets [40] or data cards [86] to describe datasets, or model cards [30, 75] and factsheets [7] to describe model specification and capabilities. Such proposals often recognise accountability as a positive quality (i.e. a virtue) and seek ways to improve transparency of algorithmic production and deployment processes. Other proposals are higher-level, seeking to integrate lower-level mechanisms and provide ways of understanding and interrogating holistically the *process* of producing, deploying, and using an algorithmic system, such as proposals for auditability [110], reviewability [25], contestability [56], traceability [62], and others [109]. Broadly speaking, these have mainly reflected accountability as a mechanism, and sought ways to support institutional mechanisms and accountability relationships between actors and forums. Though coming from different perspectives, these various lower-level and higher-level mechanisms are all essentially grounded in understandings that algorithmic accountability—either as a virtue or a mechanism—must reflect the ‘many hands’ nature of AI technologies.

More recently, a ‘second wave’ of algorithmic accountability research has sought to address more structural concerns around the development, deployment, and effects of algorithmic systems [81]. This work moves from creating better methods and processes to scrutinise systems *in situ* to considering *whether* such systems should be built at all, how, for what purposes, and who gets to govern them. This wave echoes longer-standing critical work in fields such as surveillance studies, which has considered the structural impacts of technologies of sorting and profiling on societies, and in which arguments exist against using these technologies altogether [39, 48, 66]. While literature in these fields considers issues such as the cumulative effects of systems on individuals and communities [39], they typically consider systems themselves through an organisation-centric lens – equating the functionality of interest (credit scoring, criminal profiling, airport screening, targeting advertising), with either the actor authorising the action (bank, police department, interior ministry, online platform), or at times, a particular technology provider or contractor.

Yet following the agile turn, an organisation-centric view is a less meaningful frame for analysis. Consequential algorithmic systems are commonly produced, deployed, used, and have effects through and within supply chains. It is therefore no longer the case that software is generally developed by particular teams or organisations (who may have integrated components developed by others into their finished product). Instead, as we argue, functionality results from the *working together of multiple actors* across various stages of production, deployment, and use of AI technologies (connected by data flows across organisational, legal, technical, visibility boundaries). This does not mean that a particular single organisation will never be appropriate to hold to account, but that identifying the actors and processes that led to the functionality of any particular algorithmic system becomes significantly less straightforward.

We next (§3) explore key features of algorithmic supply chains, followed by their challenges and implications for the mechanism of algorithmic accountability (§4).

3 AI SERVICES AND ALGORITHMIC SUPPLY CHAINS

Significant barriers to entry limit the number of organisations that can produce bespoke state-of-the-art AI technologies in-house, either for their own use or to bring to market [26]. Developing, maintaining, and renewing advanced AI technologies typically requires large and relevant quantities of data, potentially from multiple sources and labelled or moderated, relating to many use-cases, contexts, and subjects. Cutting-edge model development requires scarce expertise in model training, testing and deployment, all with significant storage, compute, and networking needs.

Companies with these capabilities now offer commercial access to cloud-based AI technologies ‘as a service’ (AIaaS) [26, 64]. Major companies including Amazon [3], Microsoft [72], Google (Alphabet) [43], and IBM [51] offer networked access to various state-of-the-art AI capabilities, including both model-building services and pre-built (and ‘general purpose’) models in areas such as language, speech, vision, and analytics (see [64]), or generative models for producing text, images, audio, or video. Some companies offer specific services to customers, such as facial recognition [22], hiring [49, 87], or medical diagnostics [52, 65]. And some operate as platforms for all the above, looking to connect developers, clients and infrastructure providers, among others, in a multi-sided market – Amazon and Microsoft, for example, offer access to models from other providers alongside their own [73, 100], whereas other platforms are primarily an intermediary (such as HuggingFace [38]). AI services can be integrated into apps and Web services, analytics systems, business and industrial processes, workflows, and with IoT devices with real-world physical effects (collectively: ‘*applications*’). Low marginal cost and effort means this will likely become the primary way that organisations integrate AI capabilities [26].

AI services take various forms [64]. Here we focus on services offering access to pre-built ‘general purpose’ models and to customised models tailored using tools offered by providers. In these, providers take major roles in the technology’s production and distribution, developing and hosting systems on their (owned or managed) infrastructure. Services are typically accessed through application programming interfaces (‘APIs’) controlled by providers, which allow the underlying system’s capabilities to be integrated into applications by customers (Fig. 1). This client-server model thus allows providers’ algorithmic systems to run on their infrastructure, under their control, even while they are deployed by customers in applications across many contexts and use-cases. There are typically few (if any) checks on customers’ identities or intentions, services use standard-form contracts (at least for smaller clients), and customers are billed on the API calls made [26].

An application’s supply chain may involve several AI and non-AI services , potentially from multiple providers. Indeed, an AI service may be only one part of a broader, more complex chain for a given application, which may integrate multiple AI and other services. Actors in these chains are broadly ‘upstream’ or ‘downstream’ from the perspective of others – though this distinction can blur where

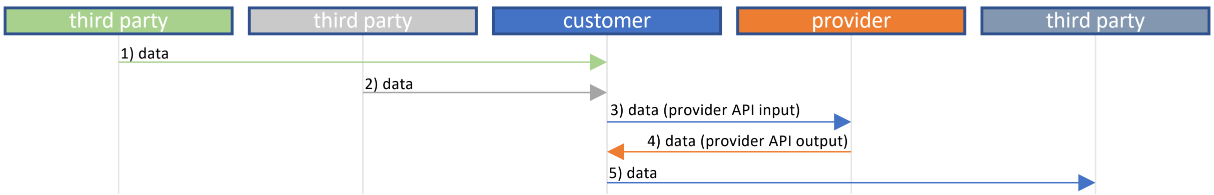

Figure 1: Sequence diagram of a simplified data-driven supply chain with an AI service. The customer sends input data to provider's API, who performs some kind of computation on that data, before returning the results to the customer. Some of the broader supply-chain is illustrated, where the customer has previously received data from other third parties, and later, sends some data to another.

actors take multiple positions in a chain. AI services themselves have supply chains, such as for dataset production activities like data gathering and labelling [69] (see §4.4). Firms using AI services may themselves provide services to their own customers [74], such as proctoring software sold to universities repackaging Amazon's facial recognition service [4], or copywriting software repackaging OpenAI's text generation model (GPT-3) [103].

We now identify several key features of algorithmic supply chains:

- production, deployment, and use are split between several interdependent actors;
- supply chain actors and data flows perpetually change;
- major providers' operations are increasingly integrated across markets and between production and distribution; and
- supply chains are increasingly consolidating around systemically important providers.

We draw out their implications for accountability in §4.

3.1 Supply chains split production, deployment, and use between interdependent actors

In algorithmic supply chains, different aspects of production, deployment, and use of AI technologies are split between multiple actors tied together by data flows. As a result, the activities of the various actors in supply chains each depend on the actions performed by others. This may involve various interacting AI and non-AI technologies, such as cloud services, servers, networks, data centres, data sources, and content delivery networks, controlled by different actors. The *working together* of the various actors who control these technologies—each doing something that enables, supports, or facilitates the actions of others—produces a particular outcome (see Fig. 2). Each actor in a supply chain may not be aware of the others, nor have consciously decided to work together towards that outcome – indeed, they may have limited understanding of actors even one or two steps removed (see §4.2). However, each depends on something done by others, and their role in a supply chain is contingent on the activities of actors both up– and downstream of them.

Actors in algorithmic supply chains are thus *interdependent*, each doing something to fulfil the needs of others (such as processing a particular data input and returning an output, or providing infrastructure to support application deployment). The interdependence of the various actors responsible for developing, deploying, and operating algorithmic systems in supply chains means they are not individual, independent actors *as such*. Instead, these actors, their relations, and their role in the workings and effects of AI technologies can only be understood *in the context of that supply chain*. Studying an actor and their systems in isolation from supply chain contexts is akin to studying an algorithmic model in isolation from its broader organisational context (the limitations of which are increasingly recognised (see §2)). The dynamics of interdependence in algorithmic supply chains—how they are structured, the relative importance of actors, and how problems spread—are therefore key considerations for algorithmic accountability, as we now explore.

3.1.1 Supply chain interdependencies are structured by technological, legal, and political economic factors. Interdependence between actors gives algorithmic supply chains their structure and functionality. Certain actors—typically (AI and non-AI) cloud service providers—have leveraged AI technologies they own, production processes they control, and cheaply-accessed networking technologies to pursue particular interdependencies with others and strategically position themselves in markets and supply chains of many kinds. Technologies afford certain capabilities to those who use or control them [31, 42, 78]. They can therefore also afford the ability to do things that fulfil the needs of others. Because, as we discuss in §3.1, people doing things for each other produces interdependence between them [36], different technologies can afford different kinds of interdependencies. Networking and data processing technologies, for example, allow the stages of production and deployment of AI technologies to be distributed geographically. They can therefore be done by different people, each of whom does something for the others, producing interdependence between them.

However, technologies and their affordances cannot *determine* interdependencies or the structure of supply chains. Affordances are not objective properties of technologies, but depend on context and perspective [31, 33, 42, 78, 108]. How providers can strategically position themselves is thus shaped both by their technologies' affordances *and* by social, legal, and political economic factors which also influence how actors relate to each other, what they do for each other, and the interdependencies that arise. Accordingly, to position themselves in supply chains and markets, providers have also leveraged political economic factors such as economies of scale and favourable legal frameworks such as intellectual property, intermediary liability, and data protection [26, 28]. Political economic and legal factors—not just technological—are thus important in producing and structuring algorithmic supply chains.

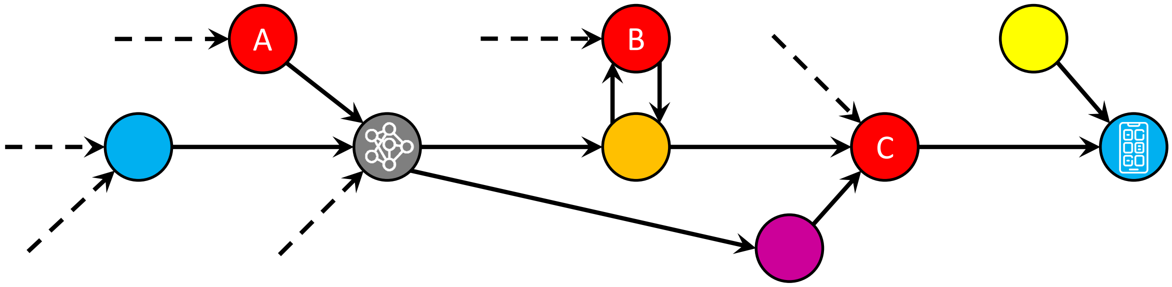

Figure 2: A representative AI supply chain. The application developer (blue) initiates a series of data flows by sending input data to an AI service provider (grey). One AI service provider (red) appears at multiple key points in the supply chain – providing infrastructure (A) for an AI service offered by (grey); providing an AI service (B) to another cloud service provider (orange); and providing technical infrastructure (C) for application deployment.

3.1.2 Some actors are core players in supply chains. Supply chain actors are generally not equal in their interdependence with each other, and some may do things that others particularly depend on. Their services, for example, may be relied upon by multiple others, as is often the provider's aim in offering general purpose services. Providers may depend on each customer only a little, while customers may depend on the provider for business-critical application functionality. Supply chain interdependencies are thus often *asymmetric*, with certain actors—typically including at least those responsible for production of AI technologies—performing core functions for others, while others are more peripheral. Various contextual signs might indicate that an actor is core in a particular chain. For example, they may be the application developer who calls the supply chain into existence. They may perform some function (such as providing an AI service) which is crucial to application functionality. They may provide a key step between actors (such as offering access to another provider's technology through their API) upon which subsequent steps of the chain depend. Or they may appear at multiple different points in the chain providing cloud-based technical architecture on which functions performed by other actors rely. Some actors may also be more interchangeable and replaceable than others – the barrier to entry for applying a specific off-the-shelf API is typically substantially lower than to generate the underlying technology to begin with.

These asymmetries of interdependence produce asymmetries in power [36]: where one actor depends more on things done by another than that other depends on them, the balance of power between them favours the second actor [36]. An application developer, for example, who uses a major provider's AI service, will depend more on that provider than the provider—who has many customers—will depend on that one developer. Power balances in algorithmic supply chains thus arise relationally yet asymmetrically and change over time as the relations and interdependencies between actors evolve [36]. These power balances are not determined by actors' relations to the *technologies* involved, but through their relation to and interdependence with *each other*. As we note (§3.1.1), this is subject to many potential influences, of which factors like control of production processes and APIs are just some. But, by leveraging their technologies alongside legal, social, and political economic factors, providers can hold significantly asymmetrical positions as core actors in many supply chains, with power balances between them and others heavily in their favour.

3.1.3 Interdependence helps problems propagate. Supply chain interdependencies mean that problems with one actor's technologies can propagate through other actors' systems. Where an AI service is biased in some way (such as facial recognition performing poorly on particular demographics [15]), that bias will be inherited by downstream applications relying on that service [64]. Such a cascade's effects may be complex and unpredictable given the largely undocumented and potentially dynamic set of actors and interdependencies found in many chains. It cannot be assumed that statistical guarantees will hold when systems are composed together, and it is not straightforward to evaluate a whole system when each individual component may have been evaluated under different threat models (or other critiera) [58, 64]. Unless identified by the provider, actors 'downstream' from them may be unaware of a problem until they notice some unexpected behaviour. Even then, because they have delegated key aspects of production (and possibly deployment) of the technologies their application relies on to other actors (such as AI service providers), customers may struggle to understand where in a supply chain the problem has arisen, why, and what they can do to mitigate it.

3.2 Supply chains are transient and dynamic with unstable interdependencies

The rise of agile development combined with the shift to a services-based distribution model for software has produced algorithmic supply chains which operate as dynamic *processes* of data flow between a changing number and arrangement of actors. Just as critical engagement with algorithmic systems must recognise that they change over time [96], so do algorithmic supply chains. Indeed, because supply chain actors are tied together by data flow, a chain may differ each time it is instantiated. At various points there may be multiple possible directions for data to flow between actors depending on the outcomes of analyses performed on it. A face detected in a video stream using one service, for example, might trigger a flow to a separate facial recognition service to identify the person (which perhaps relies on its own supply chain with associated data flows) and back again. Depending on their reported identity, this might trigger a flow to a third system to record and alert of the presence of that particular individual. Supply chains can thus be dynamically instantiated, and their structure may vary depending on the input data and the outputs of component systems. A supply chain's structure—the actors involved, what they are doing for each other, the interdependencies and power balances between them—may therefore only be fully apparent once the functionality or outcome has been been produced.

However, technical, legal, or economic relations between actors do often persist across multiple instances of a particular supply chain. An application developed to use a particular provider's service will typically use that service repeatedly, even if the path of data flow between actors differs between instances. As such, while supply chains may change overall, bilateral relations between particular actors may remain relatively consistent. However, the nature of their relationship—the services provided and used, for example—may still change over time. An application developer may introduce new features, for instance, which use additional services offered by the particular provider. They may deprecate features such that particular services are no longer needed. They might employ additional support services for rapid growth in application resource requirements (for example, where an application 'goes viral'). The provider may change their terms of service (altering the legal relationship between them) or withdraw particular services (resulting in changes in the developer's application). These are just some of the ways that relationships between actors may change.

3.3 Some actors' operations are integrated across markets and between production and distribution

Some providers of AI and other cloud services commonly found in algorithmic supply chains have reached high levels of integration; both horizontally (across markets and sectors), and vertically (across production and distribution processes). This has implications for their positioning and role in algorithmic supply chains.

3.3.1 Horizontal integration. Horizontally integrated companies operate across markets and sectors. The most prominent cloud providers (Amazon, Microsoft, Google, Alibaba, IBM) offer various services across many related and adjacent markets and may appear repeatedly in a given supply chain. Some of these services are AI-related; others are infrastructure for applications (storage, database, content delivery, identity and credential management, and so on); still others are user-facing, from business and consumer web-based services (such as maps or photo backup) to software packages for customers and their users (such as Microsoft 365). This allows a single provider to offer a range of services supporting customer applications, while also adapting the underlying technologies for their own user-facing services. It is also common for providers to purchase potential competitors and new market entrants, either to obtain intellectual property which may be useful in new markets, to expand their services across markets, or to stifle emerging competition in existing markets. Providers can also simplify the process for existing customers to bring AI services within their applications by providing tools to facilitate integrating them with their other services. Providers may financially incentivise customers to use several of their services instead of those of competitors.

3.3.2 Vertical integration. Vertically integrated companies control multiple stages of production and distribution. Several major AI providers—primarily Amazon, Microsoft, and Google—own key infrastructure for the production and distribution of their services across markets: data centres and servers; high performance computing systems; content delivery networks; APIs and customer-facing interfaces; and network infrastructure. Such vertical integration offers providers bespoke technical infrastructure for their specific needs which they can use for many services across markets to exploit economies of scale. High resolution media (requiring significant resources), for example, thus encourages vertical integration, as does state-of-the-art AI production (requiring more data, bigger and more complex models, intensive compute, and sophisticated training and testing processes). By controlling production and distribution, providers can also integrate their AI technologies into other services offered to customers and end-users. They can therefore test and further refine those AI technologies using customers' input data, applications, and real-world use cases [26]. This allows providers to reduce the resources needed to improve models, while offsetting some research and development costs by bringing it into a process paid for by customers [26]. They can thus lower the net cost of developing more accurate and more generalisable systems [26].

However, vertical integration has limits. AI providers might not operate their own in-house data cleaning and labelling processes, for instance (a key part of training, testing, and updating models). As we argue in §4.4, the business benefits to providers of bringing these processes 'in house' are potentially outweighed by the commercial advantages of extending supply chains across borders to exploit differences in laws. As a result, certain aspects of AI production are often outsourced to low-paid and insecure workers in the Global South [47, 84] (with data cleaning and labelling itself operating as a service offered by companies like Sama AI [95], or through Amazon's Mechanical Turk [2]). Moreover, some major providers—including Microsoft [73] and Amazon [100]—now offer access through their services to generative (foundation) models produced and controlled by others, marking a shift towards *less* integration in some emerging product sectors.

3.3.3 Providers all the way down. Though some prominent AI providers are both horizontally and vertically integrated, most others are not. Instead, they tend to specialise in a few closely-related services, such as algorithmic recruitment, processing legal documents, certain medical processes [64], and even 'algorithmic governance' and 'ethical AI' (see [34]), and do not operate across traditional cloud service markets. These specialist providers typically do not own their own infrastructure, but may 'rent' technological capacity from a larger provider (OpenAI, for example, exclusively uses Microsoft's Azure cloud services [73]). This reflects the fact that those developing advanced AI technologies and operating them at scale—whether providers or others—will in many cases require technical resources beyond the means of all but the biggest providers. As a result, whether through their own AI services or through those of others who depend on their cloud infrastructure, major providers like AWS, Microsoft Azure, and Google Cloud will be crucial players in future AI development and distribution.

Some AI-specific providers' services can be accessed only through a larger provider's interface and brought by customers into applications through that specific provider's cloud, rather than through a competitor (OpenAI's commercial services can be accessed *only* through Azure [73]). The larger provider's cloud offering thus operates as a platform through which they facilitate and can gatekeep market access to the smaller provider's service. In some cases, one cloud provider's interface may be used to access a specialist AI provider's model [100], where that specialist provider itself uses a *different* cloud provider for their supporting infrastructure for development [6]. That is to say, several larger cloud providers may be involved at different stages of production and deployment of specialist AI providers' services (and indeed, those of others).

3.4 Supply chains are increasingly consolidating around systemically important providers

The dynamics of interdependence and integration mean that algorithmic supply chains are increasingly consolidating around (primarily) Amazon, Microsoft, and Google [24, 26, 93]. Several factors tend towards consolidation, including competitive advantages offered by integration. These companies span markets, offering 'all-in-one' packages to application developers with easy access to state-of-the-art technologies, which readily scale and enable 'global' reach. In AI production, they leverage bespoke and advanced computing resources and expertise, significant quantities of data representing real-world deployments and use-cases, and economies of scale across AI and non-AI customer bases. They can therefore offer their services at lower cost, broader scale, greater technical sophistication, and with potentially easier access for customers than many competitors. Moreover, their substantial financial resources help consolidate their position through purchases of and investments in potential competitors (such as Google's purchase of DeepMind [41], or Microsoft's investment in OpenAI [73]).

As a result, the major AI providers are *systemically important* for the political economy, governance, and accountability of AI. Even where an application does not use a major provider's AI services (where the developer uses their own AI technology, for example, or obtains it from a smaller provider), major providers' non-AI services may form significant parts of the supply chains for either that application or the AI service it uses (or both). These providers are thus core actors in many supply chains. The leveraging of certain technological and other factors by these providers to strategically position themselves across markets can therefore be understood as a process of enclosure of AI-technological infrastructure and, by extension, of businesses, institutions, organisations, and sectors relying on supply chains involving their services. They are therefore positioned in commercially beneficial interdependencies both with other actors in particular supply chains, but also in a broader sense – a few dominant providers underpin important social and economic processes while themselves depending to various degrees on many actors in social, legal, technological, and political economic processes which help produce and maintain their position.

4 (IMPLICATIONS FOR) ACCOUNTABILITY IN ALGORITHMIC SUPPLY CHAINS

Algorithmic supply chains bring difficult implications for governance and accountability. Much algorithmic accountability research often reflects an organisation-focused, 'many hands' understanding of accountability (§2). Yet the production, deployment, and use of AI technologies in supply chains is split between multiple interdependent actors who together produce its workings and effects and whose part in producing functionality cannot be understood separately from the chain (§3.1). Organisation-focused framings cannot properly capture this distribution of responsibilities between actors across the stages of the AI lifecycle, which also challenges assignments of accountability in relevant legal frameworks (§4.1). Moreover, problems with systems can propagate widely downstream through supply chains (§3.1.3), yet particular actors are often unaware of the broader chain, and the limits of visibility across supply chains make interventions like risk assessments difficult (§4.2). It is therefore crucial for governance and accountability mechanisms to understand the actors in supply chains, what they do for each other, which of them take core roles, and the interconnections and interdependencies between them. At the same time, however, the dynamic, transient nature of supply chains (§3.2)—which can potentially be instantiated each time and unfold differently as data is processed—is also challenging.

Moreover, algorithmic supply chains are structured through interactions between technological, legal, social, and political economic factors (§3.1.1). It is therefore not enough to attend only to ways of making the technology more transparent or understandable (though this can help understand specific points in particular systems' lifecycle). Instead, algorithmic accountability work must consider broader factors: how providers leverage technology and law to structure interdependencies, integrate their operations (§3.3), consolidate their position (§3.4), increase their control and power while minimising legal accountability (§4.3), and extend their supply chains across borders to minimise cost and legal risk and maximise commercial benefit (§4.4). The dynamics of supply chains, the legal and political economic factors influencing their structure, and the relations and interdependencies between actors that result are all significant considerations from a view of accountability *as a mechanism*—one, in particular, for investigating, understanding, assessing, challenging, and contesting power. They are also important in considering who should be accountable, to whom, for what, and through which mechanisms and institutional arrangements

4.1 Responsibility for algorithmic systems is distributed between several actors

Governance and accountability mechanisms around algorithmic systems should address the *distributed responsibility* in algorithmic supply chains. Different actors control aspects of commissioning, designing, developing, deploying, using, or monitoring a particular AI technology. Responsibility for the workings and outcomes of supply chains is thus distributed among several actors who may not be straightforward to identify nor consistent across instances. Even when some actors are influential, there is therefore typically no one actor in overall control of a supply chain. Existing accountability literature, however, typically assumes that (while models or input data might change) the actors and components are not an important source of instability. Yet directing governance and accountability mechanisms at, or allocating accountability to, the wrong actors in supply chains risks undermining the stated goals of these mechanisms.

4.1.1 Legal accountability and distributed responsibility. Some legal frameworks have sought to address distributed responsibility in data-driven supply chains more generally. The Court of Justice of the European Union (CJEU) has attempted to contend with this in data protection law, for example. A key question in data protection law is who is a *data controller* – in *de facto* control of, and therefore primarily responsible in law for, personal data processing [37]. The CJEU has repeatedly held that multiple parties can be controllers for some or all aspects of a chain of processing [17–19, 26, 68]. Where several actors have common interests in the same processing, they may be *joint controllers* [37]; where their interests in the processing diverge, they may be *separate controllers*. In doing so, the Court made several useful observations: actors can be controllers if they have influence despite not having actual access to the personal data [17–19], controllers are not typically responsible for parts of the chain that precede or are subsequent to those they actually control [19]; controllers are responsible for data protection law within the framework of their 'responsibilities, powers and capabilities' [16, 20]; and using another actor's platform does not exempt a controller from their obligations [17].

Recognising the plurality of actors in chains of processing is welcome, but even data protection law's more nuanced assignment of roles and responsibilities may not readily map to algorithmic supply chains [26, 46, 68]. Under current understandings, AI service customers are likely data controllers (the dominant party, ultimately responsible for compliance and accountability), while providers may be *data processors* (the subordinate party, acting only under the instruction of a controller, with limited obligations) [26, 68, 74]. Yet this assignment of legal roles and responsibilities does not describe the real interdependencies and power relations between AI service providers (who are in control of their technologies, often core actors in supply chains, potentially systemically important more generally, typically presenting customers with 'take-it-or-leave-it' contracts, and to a large extent determining AI-driven functionality in customers' applications through their production processes) and their customers (potentially small companies without AI expertise, typically with no access to the provider's systems, control over them, or knowledge of how they work) [26]. Even where providers *are* likely controllers for aspects of the service—such as where they use customer data for service improvement—they typically attempt to minimise responsibility by claiming in their service agreements to be processors [26]. Yet the CJEU has consistently held that the factual situation outweighs contractual or other arrangements, and regulators have contradicted claimed assignments of legal roles in other kinds of data-driven supply chains [106]. Challenging providers' claims, however, would involve litigation or regulatory investigation. Moreover, given the need for joint controllers to agree the division of controllers' duties and responsibilities between themselves [37], it is not clear how joint controllership can work where actors don't necessarily know of each other or have any direct relationship.

The EU's proposed AI Act suffers from related tensions. It recognises that the 'user' of an AI system (in this context, generally the customer of an AI service) may differ from its 'provider', and envisages circumstances where a user of an AI service does so for a purpose not intended by the provider, and thus in law becomes responsible for the underlying system [29]. However, this does not reflect supply chain interdependencies and dynamics, where production, deployment, and use are distributed between actors. Instead, in this circumstance, the Act would potentially make actors several steps downstream from production responsible for ensuring that the AI technology complies with production-related legal requirements around training and testing, accountability, and risk management. While the user would inevitably be unable to comply (due to the actual distribution of practical responsibilities in algorithmic supply chains), the actor who developed and controls that technology and is thus factually responsible for production would face no obligations. This may incentivise actors who can never provide assurance of compliance to pretend they can – easily done due to the Act's self-regulatory framework and limited planned regulatory capacity [107]. Regulatory systems which hold supply chain actors downstream of production to account for design and development may do little to regulate those who are factually responsible for production and who benefit financially from potentially unlawful API queries, effectively shielding them from liability.

4.1.2 Allocating accountability. Governance and accountability mechanisms should therefore be grounded more clearly in and emphasise an understanding of the distribution of responsibility in algorithmic supply chains. Not *every* actor in a supply chain will be responsible for the outcome of the algorithmic system – some will provide only supporting services which do not meaningfully affect outcomes. Neither will actors who *are* in some way responsible be *equally* responsible – some play a bigger role than others in determining outcomes. Nor will they be responsible for the *whole* supply chain – different actors control different aspects of it. Accountability should thus be allocated to actors across supply chains based on a proper understanding of their technological and political economic dynamics. This requires processes and criteria for identifying the distribution of responsibility across supply chains and allocating accountability to those actors, for which activities, accounting to whom, and with what possible consequences.

It will therefore be important to understand the distribution of responsibility in algorithmic supply chains in terms of who is doing what for whom, who is performing what key functions for others, who is core to certain supply chains, and who is systemically important. Particular attention is needed due to systemically important actors – at the time of writing, primarily Amazon, Microsoft, Google, and perhaps a few others. Though technological and political economic dynamics tend towards consolidation around these companies, and though non-AI services often provide supporting infrastructure, it will still be important to have ways to determine which aspects of supply chains are key to their outcomes and effects, as opposed to those which could be interchanged without affecting those things. The latter, while potentially significant in their own right, are perhaps less of an urgent subject of governance and accountability mechanisms than the former.

4.2 The *accountability horizon* limits visibility across supply chains

A significant challenge for governance and accountability mechanisms in algorithmic supply chains is the **accountability horizon** – the point beyond which an actor cannot 'see', which depends on the actor and the supply chain. Supply chain actors will generally be able to know whom they interact with directly (a first 'step' in the chain), and perhaps whom those first step actors interact with in turn (a second 'step'), but may not understand (or be able to find out about) the data flows and interconnections beyond [24, 99]. Moreover, distributing responsibility for production and deployment between actors (§4.1) means each has incomplete information even if they *do* know who is up- and downstream of them. AI service providers, in control of production, may therefore lack knowledge of downstream contexts and use-cases of application deployments [64]. Those responsible for deployment and use typically lack access to models and often to information about their specification, training, testing, validation, and so on (and thus may have limited understanding of their capabilities and limitations).

The accountability horizon is thus a problem for producers of algorithmic systems in the earliest stages of developing their technologies (*problem framing*, §4.2.1) and for legal and other governance frameworks based around *risk management* (§4.2.2).

4.2.1 The accountability horizon makes problem framing difficult. Many algorithmic issues stem from choices around problem definition and framing that inform system design. Complex concepts may be formalised poorly, tasks may be incompletely captured, and different contexts may be insufficiently considered [97]. The 'many hands' problem made critical questions of identifying who framed the problem and when it was framed difficult to answer [82]. Supply chain dynamics *giving rise to the accountability horizon* complicate this further. Those responsible for production have limited capacity to understand the contexts of deployment and use by others, while the actors closest to the problem—those deploying or using the system—are generally unable to influence its design. Moreover, due to the split between production and deployment, application developers necessarily engage in their own problem framing – determining whether they need an AI service to address a particular problem and, if so, which is most suitable. Yet they may lack capacity to determine which service (if any) is most appropriate to their needs (particularly if organisations swap organisational and IT know-how for license managers [10]). This is further complicated by the fact that not all services are fungible, or adaptable to a range of different framings. Services may only accept certain kinds of input data, produce certain kinds of output data, or be amenable to certain kinds of alteration and customisation. They may be developed with particular underlying assumptions which can (or should) preclude their deployment or use in other contexts [64]. Supply chain integration may further reduce flexibility in problem framing, as technical hurdles to limit interoperability and cost implications make components less readily swappable (particularly where services are strategically bundled by providers).

More cynically, actors may encourage problem framing which increases demand for their own products and services. For example, organisations selling technologies for input data, such as cameras and other environmental sensors, may also sell workplace monitoring tools which take advantage of the data produced by these sensors. Application developers with low problem framing capacity might adopt these tools without properly identifying whether they need workplace monitoring at all. The dependency of application developers on supply chains may therefore risk the autonomy of those organisations [10]. Indeed, using AI services leaves healthcare, education and other established sectors vulnerable to unbundling and rebundling of their fundamental operations, leaving each stage amenable to value extraction through servitisation [10].

4.2.2 The accountability horizon makes risk management difficult. Many academic, policy, and legislative initiatives in recent years have proposed impact assessments, risk assessments, and risk management mechanisms to mitigate harms caused by or with AI technologies (for example, [1, 8, 25, 29, 32, 35, 40, 44, 50, 53, 54, 56, 57, 59, 60, 63, 70, 75, 88, 89, 92, 94, 110]). Data controllers' management of risks to data subjects' fundamental rights is also a core feature of the EU's data protection regime [37], applicable where AI services process personal data [26]. Typically, certain actors—who may differ between legal frameworks—have some obligation to identify and mitigate risks to individuals or their rights and interests arising from technologies they develop or control.

However, the accountability horizon makes governance and accountability mechanisms based around risk management difficult if not impossible to effectively implement. These mechanisms require knowledge of both the AI technology's specification and development (i.e. production) *and* the purpose and context of its application (deployment) [26, 64]. Yet without knowing in advance about their customers' application contexts and use-cases (which will be many, varied, and changing), providers cannot properly account for the range of potential risks that might arise [26, 54]. Similarly, without knowledge of or influence over production, customers cannot reliably assess how systems are developed, nor ensure that systems are appropriate to the risks arising in their context. Even where they have some knowledge, models are regularly updated, and customers may lack visibility or capacity to reassess. In many cases, therefore, no actor will have sufficient knowledge of or control over both production and deployment to be able to reliably assess or mitigate the impacts and risks. Risk management approaches to governance and accountability of AI technologies are therefore arguably not appropriate in this context (despite their importance in emerging laws applying to algorithmic systems, such as the EU's AI Act [29]).

4.2.3 Expanding the accountability horizon. The accountability horizon thus poses major problems for accountability. Interventions are needed to help expand the accountability horizon and better place actors to (i) know more about their own supply chains, and (ii) support others in knowing more about theirs. Yet organisation-focused tools to provide information on points in the AI lifecycle (such as [40, 75]) are of limited help where information about interconnections between actors is needed. Alternatively, tracking data flow between actors could help understand interconnections beyond the first few steps [99], as could legal and institutional mechanisms requiring information about arrangements. 'Know your customer' requirements around customer on-boarding for AI services (common in financial services) could help providers understand customers' purposes and intentions [26] (though these may only give some visibility over one or two steps in the chain). Moreover, recent CJEU data protection jurisprudence regarding transparency rights confirms that data subjects have the right to know the identity of any recipients of their personal data [21], which may help understand data flows. However, where the data controller does not know the recipients' identity—which may be common due to the accountability horizon—data subjects can instead be told about the *categories* of recipients of the data [21] (significantly less useful information).

A particular difficulty, however, is that accountability is contextual [5, 12, 25, 80, 109, 110]. The information needed to account for an algorithmic system depends on the actors responsible for its development, deployment, and use, on the forum owed the account, and on the broader context [25]. As such, the mechanisms needed to record, process, and provide information about algorithmic systems—such as [7, 23, 30, 40, 62, 75]—are also contextual. Yet the accountability horizon makes understanding context difficult for those who account for their part in supply chains. They may not know whom they need to account to, so may not know what information to retain, about what aspects of their processes, and in what form. They may also not know which actors upstream from them they can obtain accounts from. In general, the difficulties raised by the accountability horizon are not easily overcome.

4.3 Servitised distribution models give providers control beyond deployment

Nissenbaum observed that, in the mid-1990s, software vendors often demanded property protection for their products while denying, as far as possible, accountability for them [77]. Software would come with agreements which precluded *ownership* by the user (instead licensed to use a copy) and emphasised the producer's rights, while disclaiming their legal accountability for the software or anything it might do – even where harms resulted directly from defects in it [77] (though such attempts to shift liability may be invalid in some jurisdictions [71, 90]). This can be understood as developers attempting to maintain control over their software to the extent possible given the distribution model at the time (typically physical media), generating artificial scarcity for an information product to maximise revenue, with minimal risk and responsibility. This produces, as Nissenbaum puts it, a 'vacuum' of accountability [77].

Afforded in part by advances in data processing and networking technologies (§3.1.1), the agile turn changed software's distribution model away from (licensed) physical media to the service-based, API-centric models we describe above [46]. When combined with asymmetrical interdependence in algorithmic supply chains (§3.1.2), this service-based distribution model offers providers new avenues to extend control past the point of deployment. In particular, because providers depend less on individual customers than those customers may depend on them, providers can use standard form contractual service agreements and APIs as tools for structuring their relations with customers and others in supply chains in advantageous ways. Providers impose standard-form service agreements for most customers on a take-it-or-leave-it basis [74], with terms which favour the commercial interests of providers. Where vendors once sought expansive intellectual property protections, for instance, providers today seek to use their service agreements to maximise control over deployment of their technologies by reserving rights to dictate terms of use and change, withdraw, or cancel products and services at will. Moreover, providers disclaim legal accountability for things that happen through use of their services [55, 71, 105], and attempt to position themselves as data processors (§4.1) even where they are using customer data for their own purposes and will thus likely be the controller for that processing [26]. Providers can also use APIs as 'projections' [14] of the asymmetric balances of power with customers to destabilise attempts to hold providers to account: using APIs as tools to shutter businesses, destabilise research relations, and evade scrutiny [9], while contractually giving themselves those rights and exercising the power through changes in information policy as a form of control [91]. Providers can thus use a combined techno-legal strategy to position themselves advantageously in markets and shape supply chains to maximise revenue and reduce risk [27, 28].

4.4 Cross-border supply chains permit regulatory arbitrage

As we describe, data processing and networking technologies afford a geographical distribution of AI production and deployment (§3.1.1). The same technologies allow various production-related activities to themselves be distributed geographically, incentivised by jurisdictional differences in cost and regulation. This allows regulatory arbitrage, where companies in one jurisdiction exploit legal and political economic conditions in other jurisdictions to maximise commercial benefit while minimising legal accountability. This often involves contracting third-parties (such as Sama AI [95] or Supahands [102]) to undertake some aspects of production, such as data cleaning or labelling. For example, differences in privacy and data protection laws and labour protections can lower the cost and legal risk of dataset production activities like data collection, cleaning, and labelling [45, 69, 83]. Environmental factors like cheap water and energy and lax planning and waste laws can influence the location of compute and storage [76].

While some laws—such as the EU's data protection law [37] and AI Act [29] and California's Consumer Privacy Act [101]—have sought extra-territorial effect to address regulatory arbitrage, the cross-border nature of supply chains and difficulties of enforcement remains a significant accountability challenge.

5 CONCLUSIONS AND FURTHER RESEARCH

The 'many hands' problem has motivated efforts to provide information about the production, deployment, and use of algorithmic systems by teams and organisations (§2). The emergence of AI 'as a service' (or 'general purpose AI') and developments associated with cloud computing and the services model of software distribution (§3) challenge organisation-focused understandings of algorithmic accountability (§4) in ways that have not been widely addressed by the research community.

AI technologies now often involve *algorithmic supply chains*, with their production, deployment, and use split between multiple actors who *together* produce the technology's outcomes and functionality (§3.1). Major providers—now highly integrated both horizontally and vertically (§3.3)—are systemically important players (§3.1.2), and supply chains are increasingly consolidating around them (§3.4). Issues with particular systems can propagate through supply chains (§3.1.3), while they often change between instances, making it difficult to understand how they operate or who is involved (§3.2). Together, these dynamics of interdependence, perpetual change, integration, and consolidation produce supply chains in which responsibility for algorithmic systems is distributed between interdependent actors (§4.1) and visibility across the actors involved is low (§4.2). This challenges existing legal accountability frameworks while limiting the effectiveness of mechanisms like risk assessments. Moreover, splitting production and deployment makes it difficult to appropriately develop or choose AI services (§4.2.1). At the same time, the services distribution model allows providers to use terms of service and APIs to minimise legal accountability and maximise control over technologies beyond deployment (§4.3), while simultaneously extending their own production processes across borders to exploit differences in regulatory regimes (§4.4).

In all, the characteristics of algorithmic supply chains we have identified and the implications they raise challenge existing approaches to algorithmic accountability. Future algorithmic accountability research must therefore contend with supply chain dynamics: how they are structured, how they develop over time, how AI's functionality and effects are produced through them, and—importantly—how distributed responsibility challenges governance mechanisms and the accountability horizon limits visibility. This requires a broad view of supply chains, seeking to understand who is involved, what they are doing, and how to allocate accountability between them. Importantly, supply chains are structured by legal and political economic factors, which must be properly understood, as well as technological ones. If governance and accountability mechanisms are to hold those responsible for developing, deploying, and using AI technologies to account for their workings and effects, the dynamics of supply chains must be urgently addressed.

ACKNOWLEDGMENTS

The Cambridge-affiliated authors are members of the Compliant & Accountable Systems Group, which acknowledges the financial support of UK Research & Innovation (grants EP/P024394/1, EP/R033501/1, ES/T006315/1), The Alan Turing Institute, and Microsoft, through the Microsoft Cloud Computing Research Centre.

REFERENCES

[1] Ada Lovelace Institute. 2020. Examining the Black Box: Tools for Assessing Algorithmic Systems.
[2] Amazon. [n.d.]. Mechanical Turk. https://www.mturk.com.
[3] Amazon. 2021. Artificial Intelligence Services. https://aws.amazon.com/machine-learning/ai-services. (Accessed on 04/11/2021).
[4] Amazon. 2021. Case Study: WeShine. https://perma.cc/2V2T-NNKT.
[5] Mike Ananny and Kate Crawford. 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media and Society* 20, 3 (2018), 973–989. https://doi.org/10.1177/1461444816676645
[6] Anthropic. 2023. Anthropic Partners with Google Cloud. https://www.anthropic.com/index/anthropic-partners-with-google-cloud-2.
[7] Matthew Arnold, Rachel KE Bellamy, Michael Hind, Stephanie Houde, Sameep Mehta, Aleksandra Mojsilović, Ravi Nair, K Natesan Ramamurthy, Alexandra Olteanu, David Piorkowski, et al. 2019. FactSheets: Increasing trust in AI services through supplier's declarations of conformity. *IBM Journal of Research and Development* 63, 4/5 (2019), 6–1.
[8] Alejandro Barredo Arrieta et al. 2020. Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI. *Information Fusion* 58 (2020).
[9] Jef Ausloos and Michael Veale. 2020. Researching with Data Rights. *Technology and Regulation* 2020, 1 (2020), 136–157.
[10] Agathe Balayn and Seda Gürses. 2021. Beyond Debiasing: Regulating AI and Its Inequalities. https://perma.cc/4UAV-3UFB.
[11] Mark Bovens. 1998. *The quest for responsibility. Accountability and citizenship in complex organisations.* Cambridge University Press.
[12] Mark Bovens. 2006. Analysing and Assessing Public Accountability. A Conceptual Framework. *EUROGOV, European Governance Papers* No. C-06-01 (2006).
[13] Mark Bovens. 2010. Two Concepts of Accountability: Accountability as a Virtue and as a Mechanism. *West European Politics* 33, 5 (2010), 946–967. https://doi.org/10.1080/01402382.2010.486119
[14] Tania Bucher. 2013. Objects of Intense Feeling: The Case of the Twitter API. *Computational Culture: A Journal of Software Studies* 3 (2013).
[15] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability, and Transparency*. PMLR, 77–91.
[16] CJEU. 2014. Case C-131/12, Google Spain. ECLI:EU:C:2014:317.
[17] CJEU. 2018. Case C-210/16, Unabhängiges Landeszentrum für Datenschutz Schleswig-Holstein v Wirtschaftsakademie Schleswig-Holstein GmbH. ECLI:EU:C:2018:388.
[18] CJEU. 2018. Case C-25/17, Jehovan todistajat. ECLI:EU:C:2018:551.
[19] CJEU. 2019. Case C-49/17, Fashion ID GmbH & Co.KG v Verbraucherzentrale NRW eV. ECLI:EU:C:2019:629.
[20] CJEU. 2019. Case C-136/17, GC and Others. ECLI:EU:C:2019:773.
[21] CJEU. 2023. Case C-154/21, Österreichische Post. ECLI:EU:C:2023:3.
[22] Clearview AI. [n.d.]. https://www.clearview.ai. https://www.clearview.ai.
[23] Richard Cloete, Chris Norval, and Jatinder Singh. 2022. Auditable Augmented/Mixed/Virtual Reality: The Practicalities of Mobile System Transparency. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 4, Article 149 (dec 2022), 24 pages. https://doi.org/10.1145/3495001
[24] Jennifer Cobbe et al. 2019. What lies beneath: Transparency in online service supply chains. *Journal of Cyber Policy* 5, 1 (2019).
[25] Jennifer Cobbe, Michelle Seng Ah Lee, and Jatinder Singh. 2021. Reviewable Automated Decision-Making: A Framework for Accountable Algorithmic Systems. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 598–609.
[26] Jennifer Cobbe and Jatinder Singh. 2021. Artificial Intelligence as a Service: Legal Responsibilities, Liabilities, and Policy Challenges. *Forthcoming in Computer Law & Security Review* (2021).
[27] Julie E Cohen. 2012. "Piracy," "Security," and Architectures of Control. In *Configuring the Networked Self: Law, Code, and the Play of Everyday Practice*. Yale University Press, New Haven, CT.
[28] Julie E Cohen. 2019. *Between Truth and Power: The Legal Constructions of Informational Capitalism*. Oxford University Press.
[29] EU Commission et al. 2021. Proposal for a regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts. *COM (2021)* 206 (2021).
[30] Anamaria Crisan, Margaret Drouhard, Jesse Vig, and Nazneen Rajani. 2022. Interactive Model Cards: A Human-Centered Approach to Model Documentation. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) *(FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 427–439. https://doi.org/10.1145/3531146.3533108
[31] Jenny Davis. 2020. *How Artifacts Afford: The Power and Politics of Everyday Things*. MIT Press.
[32] Demos et al. 2020. Algorithm Inspection and Regulatory Access.
[33] Laurence Diver. 2018. Law as a User: Design, Affordance, and the Technological Mediation of Norms. *SCRIPTed: A Journal of Law, Technology & Society* (2018). Issue 1.
[34] EAIDB:The Ethical AI Database . 2023. Market Map. https://www.eaidb.org/map.html.
[35] Lilian Edwards and Michael Veale. 2017. Slave to the algorithm: Why a right to an explanation is probably not the remedy you are looking for. *Duke L. & Tech. Rev.* 16 (2017), 18.
[36] Norbert Elias. 1984. *What Is Sociology?* (revised ed.). Columbia University Press.
[37] European Union. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *OJEU* L119 (4 May 2016), 1–88.
[38] Hugging Face. 2022. The AI community building the future. https://huggingface.co.
[39] Oscar H. Gandy. 2010. Engaging Rational Discrimination: Exploring Reasons for Placing Regulatory Constraints on Decision Support Systems. *Ethics and Information Technology* 12, 1 (March 2010), 29–42. https://doi.org/10/bzwqrx
[40] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2018. Datasheets for datasets. *arXiv preprint arXiv:1803.09010* (2018).
[41] Samuel Gibbs. 2014. Google buys UK artificial intelligence startup Deepmind for £400m. https://www.theguardian.com/technology/2014/jan/27/google-acquires-uk-artificial-intelligence-startup-deepmind.
[42] James J Gibson. 2014. *The Ecological Approach to Visual Perception* (Classic ed.). Taylor Francis.
[43] Google. 2021. Cloud AI Building Blocks. https://cloud.google.com/products/ai/building-blocks. (Accessed on 04/11/2021).
[44] Government of Canada. 2020. Algorithmic Impact Assessment.
[45] Mary L Gray and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt, Boston.
[46] Seda Gürses and Joris van Hoboken. 2017. Privacy After the Agile Turn. In *Cambridge Handbook of Consumer Privacy*, Jules Polonetsky, Omer Tene, and Evan Selinger (Eds.).

[47] Kori Hale. 2019. Google & Microsoft Banking On Africa's AI Labeling Workforce. https://www.forbes.com/sites/korihale/2019/05/28/google-microsoft-banking-on-africas-ai-labeling-workforce.
[48] Bernard E Harcourt. 2007. *Against Prediction.* University of Chicago Press, Chicago.
[49] HireVue. 2022. End-to-End Hiring Experience Platform: Video Interviewing, Conversational AI & More | HireVue. https://www.hirevue.com/.
[50] Ben Hutchinson, Andrew Smart, Alex Hanna, Emily Denton, Christina Greer, Oddur Kjartansson, Parker Barnes, and Margaret Mitchell. 2021. Towards accountability for machine learning datasets: Practices from software engineering and infrastructure. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency.* 560–575.
[51] IBM. 2021. IBM Watson products and solutions. https://www.ibm.com/uk-en/watson/products-services.
[52] Infermedica. 2022. Call Center Triage. https://infermedica.com/product/call-center-triage.
[53] Information Commissioner's Office and The Alan Turing Institute. 2020. Explaining decisions made with AI.
[54] Seyyed Ahmad Javadi, Richard Cloete, Jennifer Cobbe, Michelle Seng Ah Lee, and Jatinder Singh. 2020. *Monitoring Misuse for Accountable 'Artificial Intelligence as a Service'.* Association for Computing Machinery, New York, NY, USA, 300–306. https://doi.org/10.1145/3375627.3375873
[55] Dimitra Kamarinou, Christopher Millard, and Kuan Hon, W. 2015. Privacy in the Clouds: An Empirical Study of the Terms of Service and Privacy Policies of 20 Cloud Service Providers. *Queen Mary School of Law Legal Studies Research Paper No. 209/2015* (2015). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2646447
[56] Margot Kaminski and Jennifer M Urban. 2021. The Right to Contest AI. *Columbia Law Review* 121, 7 (2021).
[57] Ansgar Koene et al. 2019. A governance framework for algorithmic accountability and transparency. *European Parliamentary Research Service, Panel for the Future of Science and Technology* PE 624.262 (April 2019).
[58] Blagovesta Kostova, Seda Gürses, and Carmela Troncoso. 2020. Privacy Engineering Meets Software Engineering. On the Challenges of Engineering Privacy ByDesign. https://doi.org/10.48550/arXiv.2007.08613 arXiv:2007.08613 [cs]
[59] PM Krafft, Meg Young, Michael Katell, Jennifer E Lee, Shankar Narayan, Micah Epstein, Dharma Dailey, Bernease Herman, Aaron Tam, Vivian Guetler, et al. 2021. An Action-Oriented AI Policy Toolkit for Technology Audits by Community Advocates and Activists. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency.* 772–781.
[60] Joshua A. Kroll, Joanna Huey, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson, and Harlan Yu. 2017. Accountable algorithms. *University of Pennsylvania Law Review* 165, 3 (Feb. 2017), 633–705.
[61] Joshua A. Kroll. 2020. Accountability in Computer Systems. In *The Oxford Handbook of Ethics of AI*, Markus D. Dubber, Frank Pasquale, and Sunit Das (Eds.). Oxford University Press, 0. https://doi.org/10.1093/oxfordhb/9780190067397.013.10
[62] Joshua A. Kroll. 2021. Outlining Traceability: A Principle for Operationalizing Accountability in Computing Systems. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (Virtual Event, Canada) *(FAccT '21).* Association for Computing Machinery, New York, NY, USA, 758–771. https://doi.org/10.1145/3442188.3445937
[63] Michelle Seng Ah Lee and Jatinder Singh. 2021. Risk identification questionnaire for unintended bias in machine learning development lifecycle. *Available at SSRN* (2021).
[64] Kornel Lewicki, Michelle Seng Ah Lee, Jennifer Cobbe, and Jatinder Singh. 2023. Out of Context: Algorithmic Fairness in "Artificial Intelligence as a Service". *arXiv:2302.01448* (2023). https://arxiv.org/abs/2302.01448
[65] Lunit. 2022. AI will be the new standard of care. By Lunit. https://www.lunit.io/en.
[66] David Lyon (Ed.). 2003. *Surveillance as Social Sorting: Privacy, Risk, and Digital Discrimination.* Routledge, London ; New York.
[67] Michael A Madaio, Luke Stark, Jennifer Wortman Vaughan, and Hanna Wallach. 2020. Co-designing checklists to understand organizational challenges and opportunities around fairness in ai. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems.* 1–14.
[68] René Mahieu, Joris van Hoboken, and Hadi Asghari. 2019. Responsibility for Data Protection in a Networked World: On the Queston of the Controller, "Effective and Complete Protection" and its Application to Data Access Rights in Europe. *Journal of Intellectual Property, Information Technology and E-Commerce Law* 10 (2019). Issue 1.
[69] Kira J. M. Matus and Michael Veale. 2022. Certification Systems for Machine Learning: Lessons from Sustainability. *Regulation & Governance* 16, 1 (2022), 177–196. https://doi.org/10.1111/rego.12417
[70] Danaë Metaxa, Joon Sung Park, Ronald E Robertson, Karrie Karahalios, Christo Wilson, Jeff Hancock, Christian Sandvig, et al. 2021. Auditing algorithms: Understanding algorithmic systems from the outside in. *Foundations and Trends® in Human–Computer Interaction* 14, 4 (2021), 272–344.
[71] Johan David Michels, Christopher Millard, and Felicity Turton. 2021. Standard Contracts for Cloud Services. In *Cloud Computing Law* (second ed.), Christopher Millard (Ed.). Oxford University Press, Oxford, 49–99. https://doi.org/10.1093/oso/9780198716662.003.0003
[72] Microsoft. 2021. Cognitive Services. https://azure.microsoft.com/en-gb/services/cognitive-services. (Accessed on 04/11/2021).
[73] Microsoft. 2023. Microsoft and OpenAI extend partnership. https://blogs.microsoft.com/blog/2023/01/23/microsoftandopenaiextendpartnership.
[74] Christopher Millard. 2021. *Cloud Computing Law* (2 ed.). Oxford University Press.
[75] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Conference on Fairness, Accountability, and Transparency.* 220–229.
[76] Vincent Mosco. 2014. Dark Clouds. In *To the Cloud: Big Data in a Turbulent World*. Paradigm, Boulder, CO.
[77] Helen Nissenbaum. 1996. Accountability in a Computerized Society. *Science and Engineering Ethics* (1996).
[78] Donald Norman. 1988. *The Design of Everyday Things.* Basic Books.
[79] Chris Norval, Jennifer Cobbe, and Jatinder Singh. 2020. Towards an accountable Internet of Things: A call for 'reviewability'. In *Privacy by Design for the Internet of Things: Building Accountability and Security.* The Institution of Engineering and Technology.
[80] Chris Norval, Kristin Cornelius, Jennifer Cobbe, and Jatinder Singh. 2022. Disclosure by Design: Designing Information Disclosures to Support Meaningful Transparency and Accountability. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) *(FAccT '22).* Association for Computing Machinery, New York, NY, USA, 679–690. https://doi.org/10.1145/3531146.3533133
[81] Frank Pasquale. 2019. The Second Wave of Algorithmic Accountability. *Law and Political Economy Project* (2019). https://lpeproject.org/blog/the-second-wave-of-algorithmic-accountability/
[82] Samir Passi and Solon Barocas. 2019. Problem Formulation and Fairness. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19).* ACM, New York, NY, USA, 39–48. https://doi.org/10/gftmpv
[83] Billy Perrigo. [n.d.]. OpenAI Used Kenyan Workers on Less Than $2 Per Hour to Make ChatGPT Less Toxic. *TIME* ([n. d.]). https://time.com/6247678/openai-chatgpt-kenya-workers/
[84] Billy Perrigo. 2023. Exclusive: OpenAI Used Kenyan Workers on Less Than $2 Per Hour to Make ChatGPT Less Toxic. https://time.com/6247678/openai-chatgpt-kenya-workers.
[85] Mahima Pushkarna, Andrew Zaldivar, and Oddur Kjartansson. 2022. Data Cards: Purposeful and Transparent Dataset Documentation for Responsible AI. In *Conference on Fairness, Accountability, and Transparency.*
[86] Mahima Pushkarna, Andrew Zaldivar, and Oddur Kjartansson. 2022. Data Cards: Purposeful and Transparent Dataset Documentation for Responsible AI *(FAccT '22).* Association for Computing Machinery, New York, NY, USA, 1776–1826. https://doi.org/10.1145/3531146.3533231
[87] Pymetrics. 2021. Talent Matching Platform. https://www.pymetrics.ai/.
[88] Emilee Rader, Kelley Cotter, and Janghee Cho. 2018. Explanations as mechanisms for supporting algorithmic transparency. In *Proceedings of the 2018 CHI conference on human factors in computing systems.* 1–13.
[89] Inioluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency.* 33–44.
[90] Chris Reed and Laura Edgar. 2021. Consumer Protection in the Cloud. In *Cloud Computing Law* (second ed.), Christopher Millard (Ed.). Oxford University Press, Oxford, 218–254. https://doi.org/10.1093/oso/9780198716662.003.0007
[91] Joel R. Reidenberg. 1998. Lex Informatica: The Formulation of Information Policy Rules through Technology. *Texas Law Review* 76, 3 (1998), 553–594.
[92] Dillon Reisman et al. 2018. Algorithmic Impact Assessments: A practical framework for public agency accountability (AI Now).
[93] Felix Richter. 2022. Amazon, Microsoft & Google Dominate Cloud Market. https://www.statista.com/chart/18819/worldwide-market-share-of-leading-cloud-infrastructure-service-providers.
[94] Pedro Saleiro, Benedict Kuester, Loren Hinkson, Jesse London, Abby Stevens, Ari Anisfeld, Kit T Rodolfa, and Rayid Ghani. 2018. Aequitas: A bias and fairness audit toolkit. *arXiv preprint arXiv:1811.05577* (2018).
[95] Sama AI. [n.d.]. https://www.sama.com. https://www.sama.com.
[96] Nick Seaver. 2013. Knowing Algorithms. *paper presented at Media in Transition 8* (2013).
[97] Andrew D. Selbst, Danah Boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and Abstraction in Sociotechnical Systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19).* ACM, New York, NY, USA, 59–68. https://doi.org/10/gftmp3

[98] Jatinder Singh et al. 2016. Responsibility & Machine Learning: Part of a Process.
[99] Jatinder Singh, Jennifer Cobbe, and Chris Norval. 2019. Decision Provenance: Harnessing Data Flow for Accountable Systems. *IEEE Access* 7 (2019), 6562–6574.
[100] Swami Sivasubramanian. 2023. Announcing New Tools for Building with Generative AI on AWS. https://aws.amazon.com/blogs/machine-learning/announcing-new-tools-for-building-with-generative-ai-on-aws/.
[101] State of Califorina. 2018. The California Consumer Privacy Act (CCPA).
[102] Supahands. [n.d.]. https://www.supaagents.com. https://www.supaagents.com.
[103] The Economist. 2022. Artificial Intelligence Is Permeating Business at Last. https://www.economist.com/business/2022/12/06/artificial-intelligence-is-permeating-business-at-last
[104] Dennis F. Thompson. 1980. Moral Responsibility of Public Officials: The Problem of Many Hands. *American Political Science Review* 74, 4 (1980), 905–916. https://doi.org/10.2307/1954312
[105] Felicity Turton, Dimitra Kamarinou, Johan David Michels, and Christopher Millard. 2021. Privacy in the Clouds, Revisited: An Analysis of the Privacy Policies of 40 Cloud Computing Services. *Queen Mary Law Research Paper No. 354/2021* (2021). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3823424
[106] Michael Veale, Midas Nouwens, and Cristiana Santos. 2022. Impossible Asks: Can the Transparency and Consent Framework Ever Authorise Real-Time Bidding After the Belgian DPA Decision? *Technology and Regulation* 2022 (2022), 12–22. https://doi.org/10.26116/techreg.2022.002
[107] Michael Veale and Frederik Zuiderveen Borgesius. 2021. Demystifying the Draft EU Artificial Intelligence Act. *Computer Law Review International* 22, 4 (2021), 97–112. https://doi.org/10/gns2s9
[108] Peter-Paul Verbeek. 2005. *What Things Do: Philosophical Reflections on Technology, Agency, and Design.* Penn State Press.
[109] Maranke Wieringa. 2020. What to account for when accounting for algorithms: a systematic literature review on algorithmic accountability. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 1–18.
[110] Rebecca Williams et al. 2022. From transparency to accountability of intelligent systems: Moving beyond aspirations. *Data & Policy* 4 (2022).