Written Evidence Submitted by

Dr Jennifer Cobbe and Dr Jatinder Singh, Compliant & Accountable Systems research group, Department of Computer Science & Technology, University of Cambridge

(GAI0106)

Dr Jennifer Cobbe[1] and Dr Jatinder Singh[2]

The Compliant & Accountable Systems group[3] is an interdisciplinary research group of computer scientists, legal academics, and social scientists. We are particularly concerned with accountability for and governance of new and emerging digital technologies. We have published relevant research on accountability for and governance of algorithmic systems [1] [2] [3] [4] [5] [6] [7], AI services [1] [7] , AI supply chains [1] [8] [9], and other related topics [10] [11]. That work informs this submission.

Summary

- AI technologies increasingly involve complex data-driven supply chains. What is produced by 'AI' is often the product of an interconnected systems-of-systems with multiple actors performing different roles that together result in a particular output or functionality. These supply chains are complex and dynamic, but also increasingly consolidated around a few key actors who provide cloud-based AI systems and significant other technological infrastructure to others (including other providers of AI services) 'as a service'.

- The working of AI technologies through services and supply chains has potentially significant consequences. Failures or problems in one system may propagate through a supply chain and across an AI service provider's customer base, for instance, potentially with far-reaching effects. Models developed to be offered generically as a service may not work as expected when applied to specific contexts and applications, with potentially significant implications for affected people. And the assignment of roles and responsibilities in applicable legal and regulatory frameworks may not always map to the division of control in AI supply chains, potentially undermining the law's effectiveness.

- Accordingly, AI technologies should be regulated not just as single systems, but as supply chains. Governance and accountability mechanisms that account for supply chains are needed to help understand which actors are responsible for which aspects of AI technologies and how particular outputs and effects are produced. Such mechanisms should provide reviewability both of (a) the commissioning, development, deployment, and use of specific algorithmic systems, *and* of (b) the supply chains that connect them with other systems and which drive AI technologies.

---

[1] https://www.cst.cam.ac.uk/people/jc2106.

[2] https://www.cl.cam.ac.uk/~js573.

[3] https://www.compacctsys.net.

AI technologies often work through supply chains and ecosystems

1. **Computer systems – including AI technologies – are increasingly modular, cloud-driven, and interconnected in nature.** They often rely on a supply chain of software and hardware components made, owned, and controlled by others – an interconnected and interdependent *system-of-systems*. These supply chains are *data-driven* in that systems and actors are linked by flows of data between them: a sensor system (controlled by one actor) might connect to an analytics system (controlled by a different actor) which might itself output into a decision-making system (controlled by a third actor). As a result, much contemporary software development is a matter of integrating pre-built components provided by others into a complete product. In doing so, developers delegate control over much of the underlying technology to those others actors.

2. **In practice, 'AI' is often not a single, discrete thing: what is produced by AI technologies is produced through the workings of interconnected systems forming a complex supply chain.** A developer wishing to integrate AI technologies may rely on cloud-based services, servers, protocols, data centres, third-party data sources, and content delivery networks to do so. Some component systems in these supply chains might be 'AI' but others won't be, yet the various actors and systems *together* produce a particular output or functionality. Each actor and system within that system-of-systems may itself depend on a supply chain of other systems and processes.

3. **The complex data-driven supply chains that underpin AI technologies are not static: they change, develop, and evolve over time.** The supply chains for a particular AI technology can be brought into existence each time an AI technology is used, and the actors and systems involved may differ depending on the particular inputs to and outputs of component parts in that supply chain. Models themselves are typically revised and updated iteratively, based on training datasets and procedures that change over time. Other supporting technical infrastructure is similarly revised regularly and data flows between actors may change accordingly.

4. **It is often not feasible for most individuals, companies, organisations, and governments to themselves develop advanced AI technologies of sufficient quality for deployment.** The development, upkeep, and renewal of advanced AI technologies – such as speech transcription, image recognition, and analytics – requires large quantities of training data describing many contexts, use cases, and subjects. Also required are teams of engineers with advanced technical skills, extensive computing and other technological resources, and significant time and money.

5. **Instead, technology companies with AI expertise now offer commercial access to various AI technologies 'as a service'.** Major technology companies including Amazon[4], Microsoft Azure[5], Google[6], and IBM[7] offer access to a wide range of AI technologies, across language[8], speech[9],

[4] Amazon Web Services: https://aws.amazon.com/machine-learning/ai-services.

[5] Microsoft Azure: https://azure.microsoft.com/en-us/solutions/ai.

[6] Google Cloud: https://cloud.google.com/products/ai.

[7] IBM Watson: https://www.ibm.com/uk-en/consulting/artificial-intelligence.

vision[10], and analytics[11]. More specialist companies – such as Clearview AI[12] (facial recognition) – offer specific services to customers. Unlike many other cloud-based services which provide supporting infrastructure for applications, AI services drive core application functionality (such as content recommendation). 'AI as a Service' is rapidly growing in prominence, and will likely be how many organisations use AI in future.

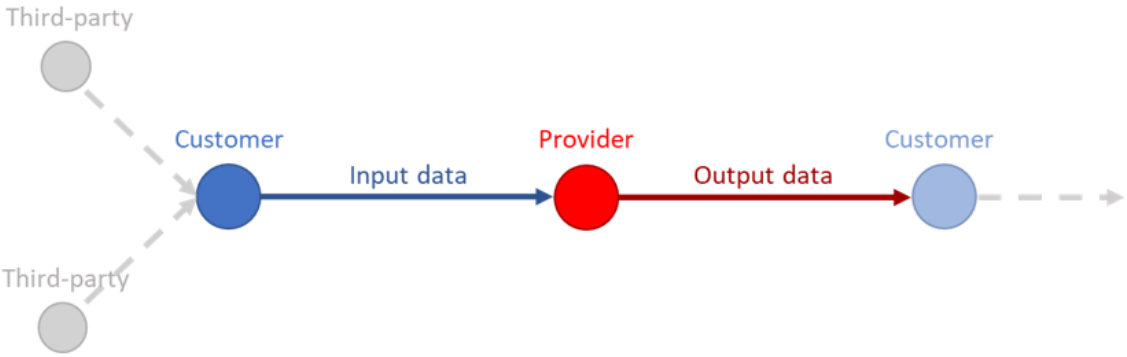

6. **AI services can be accessed through an interface[13] which gives access to the capabilities of pre-built models developed by the provider (Figure 1)**. Customers can thus bring AI functionality into websites, apps, workflows, and analytics and business processes of many kinds (which we collectively term 'applications') at low cost, with only a few clicks, on the basis of standard form contracts, and with few (if any) checks on the identity, purposes, intention, or expertise of the customer. Customers are primarily billed for actual usage. A single customer may use AI services (and other non-AI cloud services) of several providers in one application. Models can sometimes be customised by customers to suit their particular application. Some providers also offer services that allows customers to develop their own models using the providers' infrastructure, with similar access through an interface for deployment. For higher value customers, AI service providers may enter into bespoke arrangements.

***Figure 1**. A simplified AI service. Customers send input data to providers, who perform some kind of computation on that data, before returning the results of that computation to customers. Customers may receive input data from third-party sources, and may pass outputs to others.*

7. **The most prominent providers – Amazon, Microsoft, Google, and IBM – are now *horizontally integrated* in that they offer a wide range of cloud-based services of different kinds, across many related and adjacent markets.** This includes various AI services, but also a range of other (non-AI) cloud-based services. These other cloud-based services may be user-facing (including consumer web services and full online software suites for business) or infrastructural (computing resources that underpin applications). Such providers make it easy for customers using the provider's cloud services to also use the provider's AI services in their applications. Providers

[8] For example, text sentiment analysis; translation; and knowledge base creation.

[9] For example, speech transcription; speech synthesis; and voice recognition.

[10] For example (for both still images and video), image analysis and classification; object recognition; and facial detection, analysis, or recognition.

[11] For example, web usage; behavioural analysis; recommendations and personalisation; content moderation; and anomaly detection.

[12] https://www.clearview.ai.

[13] Known as an 'Application Programming Interface', or 'API'.

may also offer financial incentives for doing so. They often also purchase potential competitors and newer market entrants to expand services or stifle competition

8. **Amazon, Microsoft, Google, and IBM are also *vertically integrated* in that they own or control much of the infrastructure and supporting systems and processes on which their AI services rely.** Vertical integration allows them to exploit economies of scale with their own physical infrastructure and computing resources: data centres, servers, and analytics technologies; high performance computing systems; content delivery networks; and significant physical network architecture. Vertical integration combined with horizontal integration also allows companies to use their AI technologies and other infrastructural hardware and software systems in their other online services. This provides them with opportunities to test and further refine those AI technologies and related systems for offering to customers as a service.

9. **Smaller AI providers are typically not integrated horizontally or vertically**. Instead, they often specialise in one or two closely related services, and themselves rely on computing infrastructure belonging to one of the big three cloud service providers to underpin and deliver their services (Open AI uses Microsoft's Azure cloud services, for example[14]). In some cases, an AI technology offered as a service by a smaller provider can be accessed by customers only through a larger provider's interface (Open AI's services can be accessed only through Azure[15]).

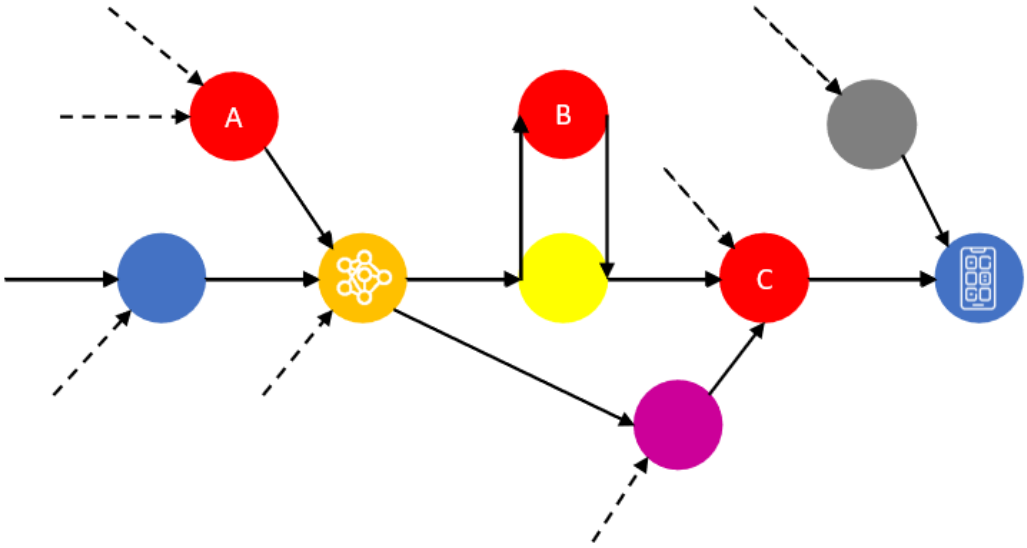

Figure 2. *A representative AI supply chain. The application developer (blue) initiates a series of dataflows by sending input data to an AI service provider (orange). One AI service provider (red) appears at multiple key points in the supply chain – providing infrastructure (A) for an AI service offered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

[14] https://openai.com/blog/infrastructure-for-deep-learning.

[15] https://openai.com/blog/infrastructure-for-deep-learning.

10. **As a result of the above, AI supply chains are increasingly consolidating around a few actors.** The major cloud service providers are core actors in many AI supply chains, strategically centring themselves alongside other more peripheral actors (Figure 2). Even where Amazon, Microsoft, or Google's AI services are not used in a particular AI-driven website or application – i.e. the customer develops their own AI technology or acquires it from a smaller provider – their other cloud-based services may be a significant part of the supply chain underpinning that AI technology. However, even with consolidation, many of these supply chains are dynamic, involving multiple actors with complex and changing interdependencies.

Implications and consequences of AI services and supply chains

11. **The interdependent nature of AI supply chains means that problems with one actor's system can potentially cascade through many other systems and actors, with complex and unpredictable effects.** As supply chains constantly change and develop, with a dynamic and largely undocumented set of actors and interdependencies, it may be difficult to foresee problems and their consequences or to identify their causes. Consequently, a problem with a provider's service can propagate across their customer base. Where a provider's AI system is biased in some way, for example, their customers' applications will inherit and propagate that bias. Customers are unlikely to know that a problem exists until their applications exhibit unexpected behaviour. Even then, customers are unlikely to be able to understand why the problem has arisen, where in the supply chain it has arisen, and what they can do to mitigate it.

12. **Certain problems stem from the generic nature of AI services when applied to specific applications and contexts.** AI service providers will build AI systems to be as generally-applicable as possible to maximise their customer base. Those systems will however be used by customers in specific contexts to perform specific tasks. Context matters for ensuring that AI systems are appropriate to the application, yet generic models cannot account for nor be tested across all possible applications and contexts where they might be used. As a result, systems that might seem okay when assessed generically can cause problems when deployed in specific customer applications.

13. **The introduction of AI technologies into previously unconnected physical spaces requires careful consideration of its potential effects and implications.** The ease with which advanced AI services can be integrated into applications potentially provides wide scope for AI-enabled surveillance and tracking in both public and private spaces, by both public and private actors. This has potentially significant implications for people's rights and freedoms and for the power dynamics in such spaces. Moreover, when used for surveillance purposes, intractable biases, errors, and other problems in AI technologies risk exacerbating divisions and hierarchies in these spaces along racial, ethnic, and gender lines.

14. **Assignments of roles and responsibilities under existing laws may not map easily to AI supply chains.** In some cases, responsibilities may aim at the wrong actors, undermining the effectiveness of legal and regulatory frameworks. In data protection, for example, AI services customers are generally data controllers (the dominant party, in control of processing) and providers are often processors (their subordinate). Yet this does not describe the true power

relationships between providers of AI services (in control of their technologies, often core actors in supply chains, and ultimately determining AI-driven functionality in customers' applications) and their customers (typically with no access to systems, control over them, or knowledge of how they work) [1]. The effect is that the law's obligations and responsibilities are often directed at the wrong party, undermining the ability of the law to achieve public policy goals.

Governance and accountability for AI supply chains

15. **To help address these issues, governance and accountability mechanisms are needed to ensure that AI technologies and processes are *reviewable* [3] and can support audit, review, and contestation**. These should include mechanisms for investigating the development, deployment, and use of AI technologies and for redress for harms and violations. Legal and regulatory obligations should require information about systems to be made available to certain actors – customers, end-users, people affected by automated decisions, regulators, and oversight bodies, as appropriate. Moreover such information should be *contextually appropriate* to properly account for the specifics of the situation and the particular needs and levels of expertise of the recipient. Legal powers, rights, and governance arrangements should bring consequences and changes in the technology, its design, its deployment, and its use where problems exist.

16. **Governance and accountability mechanisms for reviewable AI technologies should be grounded in an understanding of AI systems and supply chains as *socio-technical* in nature.** AI technologies and the systems that comprise AI supply chains are commissioned, designed, developed, deployed, and used by people, generally within particular organisational contexts and with particular goals and outcomes in mind. The choices and actions of those people within those contexts determine the working of individual AI system and should therefore be the primary focus of governance and accountability mechanisms. Interventions which seek transparency of models or other technical components alone are not sufficient to understand how AI technologies work or how outputs and effects are produced [3].

17. **Individual AI systems should be reviewable to allow understanding of key stages and steps across the whole algorithmic process [3].** This includes *commissioning and specification of systems as well as development, deployment, and use.* Across this process, record-keeping and logging requirements should mandate the retention and provision of information about technical aspects (inputs, outputs, software and hardware components, the underlying technical platforms and infrastructures, training data and models), as well as related organisational processes and workflows, the actions and decisions of individuals, and so on. Record-keeping requirements should also target systems' run-time operation – i.e. capturing information on how a system operates and behaves.

18. **Technical, legal, and institutional mechanisms are needed to understand the interconnections and interdependencies in supply chains.** Because the outputs and effects of 'AI' technologies are often in practice produced by multiple systems working together through a supply chain, reviewability of individual systems is not by itself sufficient. Supply chains also need to be reviewable. Reviewability mechanisms for supply chains should assist in understanding what

roles different kinds of actors are playing, which actors are doing what kinds of things for which others, and with what kinds of technologies. However, at present, there are few mechanisms for investigating supply chains beyond one 'link' or the first step [8]. Further research is needed on technical mechanisms to track data flow between systems [9], which would assist, as would legal mechanisms requiring information about arrangements and connections between actors.

19. **Ongoing impact assessments are needed with obligations to identify, mitigate, or eliminate risks to rights, freedoms, and interests of potentially affected people.** Systems should not be deployed or used where such risks cannot be satisfactorily mitigated or eliminated. However, the complex and dynamic supply chains that together produce the outputs and effects of AI technologies makes impact assessments difficult[16]. Although key aspects of supply chains are increasingly consolidated around a few actors (who control particular systems and technologies), there is usually no one actor in control of, or with a full view of, a supply chain as a whole. 'Upstream' actors (such as AI service providers) often lack knowledge of potential 'downstream' contexts and use cases for their many customers' applications and deployments. Requirements around customer onboarding – such as a form of 'Know Your Customer' – may help to some degree here. 'Downstream' actors (such as AI service customers) will often have no way of knowing about the 'upstream' supply chain or of the strengths, weaknesses, capabilities, and limitations of specific systems therein. Further work is needed to determine what kind of assessments should be required of which actors within AI supply chains, in what contexts, and to what end.

20. **Potential risks and implications of record-keeping should be understood and accounted for when designing accountability frameworks.** Record-keeping to support governance and accountability mechanisms may not necessarily give an accurate picture of how systems and supply chains work *in practice*. They may in some cases instead provide information about envisaged or idealised practices or processes; employees of particular actors may undermine or evade record-keeping; or recorded information may not be accurate or representative. Moreover, the keeping of records about decisions made by people in commissioning, developing, designing, or using AI systems brings risks of intensified employee monitoring and surveillance, as well as risks of revealing information that is commercially or personally sensitive. This information may relate either to AI serve providers, or – crucially – to their customers, to users of AI technologies, or to people affected by decisions made using AI technologies.

21. **Legal roles and responsibilities must be targeted at the correct actors within supply chains to achieve desired public policy goals.** Different actors may be in control of, or responsible in different applicable laws, for different stages in the process of commissioning, designing, developing, deploying, and using a particular AI technology. These actors may be either 'upstream' and 'downstream' in supply chains, depending on the particular arrangements for a given application. Future regulation will require processes and criteria for identifying control of particular aspects of algorithmic processes and supply chains and assigning responsibility and accountability to those actors. Further work is therefore needed to understand how to identify

[16] This is a key weakness of the risk management approach taken in the EU's proposed AI Act.

control and how to assign responsibility, for which activities, accounting to whom, with what possible consequences.

References

[1] J. Cobbe and J. Singh, "Artificial Intelligence as a Service: Legal Responsibilities, Liabilities, and Policy Challenges," *Computer Law & Security Review,* vol. 42, 2021 <https://doi.org/10.1016/j.clsr.2021.105573>.

[2] J. Cobbe, "Administrative Law and the Machines of Government: Judicial Review of Automated Public-Sector Decision-Making," *Legal Studies,* vol. 39, 2019 <https://doi.org/10.1017/lst.2019.9>.

[3] J. Cobbe, M. S. A. Lee and J. Singh, "Reviewable Automated Decision-Making: A Framework for Acccountable Algorithmic Systems," in *2021 ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT 2021)*, 2021 <https://doi.org/10.1145/3442188.3445921>.

[4] J. Cobbe and J. Singh, "Regulating Recommending: Motivations, Considerations, and Principles," *European Journal of Law and Technology,* vol. 10, no. 3, 2019 <https://ejlt.org/index.php/ejlt/article/view/686>.

[5] J. Cobbe, M. S. A. Lee, H. Janssen and J. Singh, "Centring the Rule of Law in the Digital State," *IEEE Computer,* no. 53, 2020 <https://doi.ieeecomputersociety.org/10.1109/MC.2020.3006623>.

[6] R. Williams, R. Cloete, J. Cobbe, C. Cottrill, P. Edwards, M. Markovic, I. Naja, F. Ryan, J. Singh and P. Wei, "From Transparency to Accountability of Intelligent Systems – moving beyond aspirations," *Data and Policy,* no. 4, 2022 <https://doi.org/10.1017/dap.2021.37>.

[7] S. A. Javadi, R. Cloete, J. Cobbe, M. S. A. Lee and J. Singh, "Monitoring Misuse for Accountable 'Artificial Intelligence as a Service," in *Proceedings of the 2020 AAAI/ACM Conference on AI, Ethics, and Society (AIES '20)*, 2020 <https://doi.org/10.1145/3375627.3375873>.

[8] J. Cobbe, C. Norval and J. Singh, "What Lies Beneath: Transparency in Online Services Supply Chains," *Journal of Cyber Policy,* vol. 5, no. 1, 2020 <https://doi.org/10.1080/23738871.2020.1745860>.

[9] J. Singh, J. Cobbe and C. Norval, "Decision Provenance: Harnessing data flow for accountable systems," *IEEE Access,* no. 7, 2019 <https://doi.org/10.1109/ACCESS.2018.2887201>.

[10] A. Ball-Burack, M. S. A. Lee, J. Cobbe and J. Singh, "Differential Tweetment: Mitigating Racial Dialect Bias in Harmful Tweet Detection," in *2021 ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT 2021)*, 2021 <https://doi.org/10.1145/3442188.3445875>.

[11] C. Norval, J. Cobbe, K. Cornelius and J. Singh, "Disclosure by Design: Document engineering for meaningful data disclosures," in *2022 ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT 2022)*, 2022 <https://doi.org/10.1145/3531146.3533133>.

(November 2022)