

*Competition and Markets Authority*  
**Algorithms, competition and consumer harm: call for information**

The Compliant and Accountable Systems research group is based in the Dept. Computer Science & Technology, University of Cambridge. We are an interdisciplinary team exploring the interplays of law and emerging technology, including questions of regulation of digital technologies, algorithmic systems, and online platforms.

We welcome the CMA's consultation on the issues of harms arising from algorithmic systems. These are important areas with wide-ranging implications for society, including for consumers and the broader population. We are therefore keen to share our research in this area and to assist with the CMA's work.

1. **Meaningful review:** Many of the points raised in the CMA's paper require ways of finding out whether, how, and why various problems with algorithmic systems occur.

We noted a considerable focus on the technology itself – footnotes 123 (data, design, and outcome implementation) and 177 (provenance) as examples. Technical approaches for 'algorithmic audits' are potentially useful, and will have a role to play; indeed, we have proposed mechanisms by which provenance techniques can give insight into the behaviour of algorithmic systems.<sup>1</sup>

However, we have argued that a broader, systematic and more holistic approach to transparency is what is needed for the meaningful review and assessment of algorithmic systems – these systems being inherently *socio-technical*. There is an opportunity for regulators to push for broader forms of transparency that support meaningful accountability as it relates to algorithmic systems, including the human and organisational processes relating to them.

To that end, we have developed a framework, which we term *reviewability*, to facilitate the (legal, regulatory or other) review of algorithmic systems – from the conception of the system, through its commissioning, design, development, deployment, and ongoing use, as well as any subsequent assessments of its functioning. This reviewability framework helps those responsible for algorithmic systems to design, develop, and use their systems in such a way that assists with regulatory oversight and compliance. The framework was published at this year's Fairness, Accountability, and Transparency (FAccT) conference, and though it is discussed in the context of automated decision-making, the framework is relevant to algorithmic systems in general.<sup>2</sup>

The CMA's paper discusses a range of algorithmic harms. The reviewability framework we propose applies across these – offering a systematic, practical foundation for meaningful accountability by providing a view of the existence and causes of issues arising from algorithmic systems. The framework has the potential to significantly benefit regulators, such as the CMA, who are seeking to identify and intervene on algorithmically-driven harms.

---

<sup>1</sup> Decision Provenance: Harnessing Data Flow for Accountable Systems:

<https://doi.org/10.1109/ACCESS.2018.2887201>

<sup>2</sup> Reviewable Automated Decision-Making: A Framework for Accountable Algorithmic Systems:

<https://doi.org/10.1145/3442188.3445921>

2. **Fairness and discrimination:** The CMA’s paper rightly raises issues of fairness, given that algorithmically-driven discrimination has the potential to cause significant problems in many areas.

Information on algorithmic processes – including the commissioning, design, development, and use of algorithmic systems and any assessment of their potentially discriminatory effects – will often be of significant benefit in determining whether, how, and why discrimination is taking place or is likely to arise. Again, our work on reviewability, as mentioned above, is very relevant here.

The CMA’s paper makes good reference to on-going research in the area of algorithmic fairness. However, one area less discussed, and that we seek to highlight is the rapidly growing area of *‘fairness tooling’*: methods and tools that purport to support those building algorithms to identify and mitigate bias and fairness issues. In recent work exploring the landscape of open-source fairness toolkits, we found that there are significant gaps between the requirements, needs and expectations of practitioners and the functionality provided by the toolkits, with a number of practitioners raising concerns that these toolkits may give a false sense of security, and possibly engender a false confidence in flawed algorithms, by making what is a complex and contextual problem appear trivial or easily fixed.<sup>3</sup> We seek to draw the CMA’s attention to this area, given that it is likely that organisations will rely on such tooling to assert that bias and fairness concerns were fully and properly considered.

3. **Recommender systems:** Recommender systems involve building (often detailed) profiles of people’s interests, preferences, and behaviours, and can affect what people might see in content feeds; media; what choices they have and might make in e-commerce contexts; the employment opportunities, utility pricing plans, and public services deemed open to them, and so on. As such, whoever is responsible for designing and deploying a recommender system has significant influence over consumers and others who interact with such systems. This raises significant policy questions around the power and gatekeeping ability of the companies responsible for platforms or services that build and use recommender systems, particularly for platforms that are dominant or otherwise widely used in certain markets.

As the CMA’s paper makes clear, there is growing evidence of a wide range of potential problems arising from the use of recommender systems. We have published work on regulatory issues around recommender systems, including high-level principles to inform the development of regulatory interventions.<sup>4</sup> Though our discussion focused on social platforms, it is equally applicable in other contexts, including e-commerce, and provides a general basis for regulation.

4. **Dark patterns:** We welcome the CMA’s discussion of user manipulation through design. These design practices, including dark patterns, can have a direct and potentially significant impact on consumers and their interests. Yet, this is an area comparatively under-considered. One theme of our research concerns the methods and incentives for encouraging better practices. We have explored, for instance, how the actions of a German regulator were, in effect, a push back against problematic user-interface design, and considered the general relationship between regulation and transparent design practices.<sup>5</sup>
5. **Technical platforms and infrastructure:** Though the CMA’s paper covers some harms to consumers arising from platforms, we note that relevant issues relating to broader technical ecosystems were little discussed.

Specifically, the technical infrastructure (‘cloud’) that underpins consumer-facing services is heavily consolidated around relatively few companies – primarily Amazon (with Amazon Web Services), Microsoft (Microsoft Azure), and Google (Google Cloud), and a range of specific infrastructural service providers such as Cloudflare (offering, among other things, content delivery and DDoS protection

---

<sup>3</sup> The Landscape and Gaps in Open Source Fairness Toolkits: <http://doi.org/10.2139/ssrn.3695002>

<sup>4</sup> Regulating Recommending: Motivations, Considerations, and Principles: <https://doi.org/10.2139/ssrn.3371830>

<sup>5</sup> Regulating transparency? Facebook, Twitter and the German Network Enforcement Act: <https://dl.acm.org/doi/abs/10.1145/3351095.3372856>

services), those providing security and identity services (e.g. Qualys, Okta), and so on. We have undertaken work identifying consolidation at this infrastructural level, indicating some issues this raises.<sup>6</sup>

Importantly, what is offered to consumers depends on the nature of the infrastructural services on which consumer-facing services rely. This means that though there may appear to be market diversity at the consumer-facing level, there is considerably less diversity below the surface – potentially limiting consumer choice.

This is because organisations that leverage such infrastructures to provide their consumer-facing offerings will need to do so by selecting from one of a decreasing number of platforms. The nature of the platform on which they rely will, in turn, determine the nature of what they can offer consumers. Further, there are also gatekeeping concerns – platform providers wield much power given organisations necessarily rely on their services in order to function. This can, again, lead to platforms (directly or indirectly) shaping the nature of consumer offerings, which is of particular concern given that many platforms operate across a range of different markets. Further, from a consumer perspective, it may be difficult, if not impossible, to choose, for example, services that are not underpinned by Amazon without losing access to a typical range of digital products and services.

Particularly relevant to issues of algorithmic harms is the significant consolidation regarding a subset of these infrastructure services termed '*Artificial Intelligence as a Service*' ('AlaaS'). This warrants consideration, as unlike other cloud services, AlaaS directly enables and facilitates functionality in consumer-facing platforms, services, and applications, by allowing them to add advanced language (e.g. translation), speech (e.g. speech transcription or voice recognition), vision (e.g. image or facial recognition), or analytics capabilities. Similar to how other cloud infrastructure has become the default for organisations wishing to create digital or online services, it is likely that AlaaS will become the primary way that many companies and organisations leverage AI in future.

AlaaS therefore raises a number of concerns with respect to algorithmic harms. Problems with these AI systems, such as around bias and discrimination, will likely propagate across the provider's entire customer base (i.e. consumer-facing organisations), with wide-ranging effects on consumers. Moreover, the potential for direct consumer harms arising from the misuse or abuse of these services – typically readily available to deploy at scale cheaply, with only a few clicks, and with few checks as to who is obtaining the service and for what purpose – is significant.<sup>7</sup> These concerns are in addition to the general concerns of infrastructure consolidation earlier described. We have a paper exploring the legal and policy dimensions to AI services, which we can make available on request.

The above represents some of our team's initial reflections on the CMA's paper. Note that our research covers a broad range of topics very relevant to the regulation and governance of technology, much of which relevant to the CMA's focus; indeed, we became aware of the CMA's call for engagement at the FAccT session on inspecting social media algorithms. As such, we'd be very happy to discuss these, and other topics, in either a formal or informal capacity.

Best regards,

Dr Jennifer Cobbe and Dr Jat Singh  
Compliant & Accountable Systems Group ([www.compacctsys.net](http://www.compacctsys.net))  
University of Cambridge  
[jennifer.cobbe@cst.cam.ac.uk](mailto:jennifer.cobbe@cst.cam.ac.uk) / [jatinder.singh@cst.cam.ac.uk](mailto:jatinder.singh@cst.cam.ac.uk)

---

<sup>6</sup> What lies beneath: transparency in online service supply chains: <https://doi.org/10.1080/23738871.2020.1745860>

<sup>7</sup> Monitoring Misuse for Accountable 'Artificial Intelligence as a Service': <https://doi.org/10.1145/3375627.3375873>