



Federal Trade Commission Request for Comments re: Notice of Hearings on Competition and Consumer Protection in the 21st Century

**Comments of AI Now Institute
August 20, 2018**

The AI Now Institute respectfully submits the following comments in response to the Federal Trade Commission’s (FTC) announced hearings and comment period on Competition and Consumer Protection in the 21st Century. Artificial intelligence (AI) is beginning to power a variety of new and emerging commercial and consumer technologies, but we are only just beginning to understand both its positive and negative impacts, especially on vulnerable populations. This stems from a variety of reasons, including the rapid pace and diversity of deployment, the complexity and scale of the technologies, and the lack of established methodologies for understanding how and when they help or harm consumers. Because of these uncertainties and the magnitude of the risks they pose, AI Now urges the Commission to keep serious watch over these technologies and to develop sound policy, guidance, and when necessary enforcement actions.

The AI Now Institute is an interdisciplinary research institute that focuses on the social implications of artificial intelligence. Housed at New York University, our current research domains include examining AI systems in the context of, rights and liberties, labor and automation, bias and inclusion, and safety and infrastructure.¹ AI Now works with researchers, legal advocates, policymakers and other stakeholders on investigating emerging research trends and concerns, and identifying policy interventions for these concerns.

In April, AI Now published an Algorithmic Impact Assessment (AIA) report,² which offers a practical transparency and accountability framework for assessing the use and impact of government algorithmic systems, including AI based systems. AIAs draw directly from impact assessment frameworks in environmental protection,³ human rights,⁴ privacy,⁵ and data protection⁶ policy domains by combining public agency review and public input.

¹ AI Now Institute, <https://ainowinstitute.org>.

² AI Now Institute, Algorithmic Impact Assessment: A Practical Framework for Public Agency Accountability (April 2018), <https://ainowinstitute.org/aiareport2018.pdf>.

³ Leonard Ortolano and Anne Shepard, “Environmental impact assessment: challenges and opportunities,” *Impact assessment* 13, no. 1 (1995): 3-30.

<https://www.tandfonline.com/doi/abs/10.1080/07349165.1995.9726076>.

An AIA gives both the agency and the public the opportunity to evaluate the potential impacts of the adoption of an algorithmic system before the agency has committed to its use. This allows the agency and the public to identify concerns that may need to be negotiated or otherwise addressed before a contract is signed. AIAs also require ongoing monitoring and review, recognizing that the dynamic contexts within which such systems are applied. In implementing AIAs, agencies should consider incorporating AIAs into the processes they already use to procure algorithmic systems or any existing pre-acquisition assessment processes the public authority already undertakes.⁷

An AIA covers any algorithmic system before it is deployed, no matter how it was acquired or if it was developed internally. A pre-procurement AIA gives agencies the opportunity to engage the public and proactively identify concerns, establish expectations, and draw on expertise and understanding from relevant stakeholders. It also requires vendors to provide more information, documentation and assurances, for the agency to perform the impact assessment.

The key elements of an AIA are:

1. Agencies should conduct a self-assessment of existing and proposed algorithmic systems, evaluating potential impacts on fairness, justice, bias, or other concerns across affected communities;
2. Agencies should develop meaningful external researcher review processes to discover, measure, or track impacts over time;
3. Agencies should provide notice to the public disclosing their definition of “algorithmic system,” existing and proposed systems, and any related self-assessments and researcher review processes before the system has been acquired;
4. Agencies should solicit public comments to clarify concerns and answer outstanding questions; and
5. Governments should provide enhanced due process mechanisms for affected individuals or communities to challenge inadequate assessments or unfair, biased, or otherwise harmful system uses that agencies have failed to mitigate or correct.

⁴ United Nations, “Guiding Principles on Business and Human Rights: Implementing the United Nations ‘Protect, Respect and Remedy’ Framework,” 20-24 (2011), http://www.ohchr.org/Documents/Publications/GuidingPrinciplesBusinessHR_EN.pdf.

⁵ Kenneth A. Bamberger and Deirdre Mulligan, “Privacy Decision Making in Administrative authorities,” *Chicago L. Rev.* 75(1):75 (2008), <https://www.truststc.org/pubs/258.html>.

⁶ “Data Protection Impact Assessments,” Information Commissioner’s Office, accessed March 16, 2018, <https://ico.org.uk/for-organisations/guide-to-the-general-data-protection-regulation-gdpr/accountability-and-governance/data-protection-impact-assessments/>.

⁷ Catherine Crump, “Surveillance Policy Making by Procurement,” *Wash. L. Rev.* 91 (2016): 1595.

The AIA framework goes beyond just the components described above. Those parts of an AIA — the definition of “algorithmic system,” public authority self-assessment, public participation, and external meaningful researcher access — must be structured into a process that ensures pre-acquisition review of algorithmic systems and the opportunity for public input to be solicited and addressed. In practice, the AIA process will not necessarily look identical between different local or national contexts because of existing government procurement requirements and relevant local interests or laws. Local, state, and international governments are contemplating implementation of the AIA framework, so it is imperative that the FTC understand how to best use their authority to protect consumers, if this framework is implemented.

AIAs provide a key mechanism for developers, vendors, and providers to assess new algorithmic technologies for consumer risks

While the AIA framework was initially developed for government agencies to assess their own algorithmic systems, it could easily be applied to private developers, vendors, and providers of algorithmic systems (including AI systems). When engaging in the AIA process, these companies are given the opportunity to share with their customers and users key information about how these systems work, including the benefits and risks they may pose. AIAs include specific technical information and documentation, testing and audit results, risk assessments, and reasonable assurances as to what the system is capable of delivering and optimized to achieve, as well as ongoing monitoring to ensure the system continues to produce intended results. For example, if a company were to develop an algorithmic system to assess applicants for credit lending, employment, admission to an educational institution, or other key social necessities, the AIA process would require the company to consider whether its system would, at a minimum, have any disparate impacts on different groups such as groups based on race, gender, age, sexuality, or ability. If such risks were present, the AIA process would also require that the company discuss any potential strategies to mitigate these risks. For government agencies, the AIA process also includes a public notice and comment period, plus a challenge period for agencies that produce inadequate AIAs. While these specific provisions would not likely apply to private companies, any system they provide through a government procurement process would then become part an AIA for the procuring entity.

AIAs also provide essential information for FTC investigations into potential deception, unfair business practices, or other violations of consumer rights.

Once produced, AIAs become part of our understanding of how the developers, vendors, and providers of these systems are representing the systems to consumers and the public writ large. Similar to statements made concerning the privacy and security of technological systems, we believe the FTC will have authority to seek enforcement actions under relevant laws for deceptive, misleading, or otherwise unfair statements within AIAs, such as under its Section 5

authority or through specific statutory regimes such as the Equal Credit Opportunity Act or the Fair Credit Reporting Act (FCRA).⁸

For example, many algorithmic systems could easily fit the FCRA definition of a consumer report, with vendors of these systems fitting the corresponding definition of a consumer reporting agency (CRA).⁹ The Fair Credit Reporting Act defines consumer reports as a communication from a credit reporting agency (1) bearing on a consumer's personal characteristics or mode of living, which can include almost any information about consumers,¹⁰ (2) that is “used or expected to be used...for purpose of serving as a factor in establishing the consumer’s eligibility.”¹¹ Algorithmic systems are often designed and trained on purchased datasets about consumers or government data about constituencies relevant to the implementation of the system. Thus, vendors of algorithmic systems, may well qualify as CRAs if the systems contain consumer information that will be used to assist or make determinations about eligibility for government benefits or services. When implemented, these systems often use detailed information about individuals to assist various service providers (including public agencies) in making actual determinations regarding who receives benefits (e.g. housing, healthcare or government subsidies), the amount or type of benefit an individual should receive, or resource allocation decisions (e.g. where fire departments or libraries are located).

Under the AIA framework, developers and vendors would provide public agencies (or, if expanded, private companies) with the information necessary for the provider to evaluate potential impacts on fairness, justice, bias, or other concerns across affected communities. This should include, but is not limited to, information about the data sets used to design and train the system, documentation of design decisions that could affect the fairness and accuracy of the system, and evidence of bias, fairness, or other validation studies of the system. In accordance with FCRA, entities that were deemed to fall within the definition of CRAs would need to ensure the accuracy of information used in algorithmic systems, as well as assist service providers in developing procedures for consumers to access their own information, along with the ability to correct any errors.

It is worth noting, however, that when some consumers have attempted¹² to access information about or question the accuracy of decisions regarding eligibility for government resources based on the use of algorithmic systems, some vendors have hindered access to relevant information

⁸ 15 U.S.C. §§ 1681–1681x (2014).

⁹ Kate Crawford & Jason Schultz, *Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms*, 55 B.C.L. Rev. 93 (2014), <http://lawdigitalcommons.bc.edu/bclr/vol55/iss1/4>

¹⁰ *Trans Union Corp. v. FTC*, 81 F.3d 228, 231 (D.C. Cir. 1996).

¹¹ 15 U.S.C. § 1681a(d)(1).

¹² Virginia Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*, (New York: St. Martin's Press, 2018)

alleging overbroad intellectual property claims.¹³ Thus, there are opportunities for the FTC to examine this issue and determine how best to maintain the FCRA's effectiveness in this arena.

Another area for potential FTC attention is the question of when algorithmic system/AI developers, vendors, and providers are required to deliver adverse action notices to consumers. We have already seen this issue arise in early instances when consumers have attempted to challenge denial of certain government benefits.¹⁴ Given that such systems and their opacity are likely to increase quickly and broadly over time, it is imperative that the FTC establish strong guidance, and if necessary, enforcement of the adverse action notice requirement for algorithmic/AI systems.

In addition to the above recommendations, we encourage the FTC to also consider effective education opportunities including studies of government use of AI systems and policy statements or advisory opinions of the practices and obligations of government vendors.

Respectfully submitted,

Rashida Richardson
Jason Schultz
AI Now Institute
60 Fifth Avenue
New York, NY 10011
rashida@ainowinstitute.org

¹³ See Robert Brauneis and Ellen P. Goodman, "Algorithmic transparency for the smart city," *20 Yale J. L. & Tech.* 103 (2018); *K.W. v. Armstrong*, 298 F.R.D. 479 (D. Idaho Mar. 25, 2014).

¹⁴ See Press Release, American Federation of Teachers, Federal Suit Settlement: End of Value-Added Measures for Teacher Termination in Houston (October 10, 2017), <https://www.aft.org/press-release/federal-suit-settlement-end-value-added-measures-teacher-termination-houston>; Colin Lecher, "What Happens When An Algorithm Cuts Your Health Care", *The Verge*, Mar. 21, 2018, <https://www.theverge.com/2018/3/21/17144260/healthcare-medicaid-algorithm-arkansas-cerebral-palsy>.