September 24, 2024

Chair Nily Rozic and Chair Steven Otis
Legislative Office Building
Albany, NY 12210

Re: Ensuring consumer protection & safety relating to the use of artificial intelligence

Dear Honorable Asssemblymembers,

Consumer Reports[1] appreciates the opportunity to submit testimony about consumer protection issues related to artificial intelligence. Consumer Reports (CR) is an independent, nonprofit and nonpartisan organization that works with consumers to create a fair and just marketplace.

States are leading the way on artificial intelligence policy, and Consumer Reports has been pleased to see bipartisan consensus on the need to safeguard consumers' rights.

Artificial intelligence tools present challenges for consumers related to transparency, fairness, privacy, validity, testability, and more. We outline issues and some potential remedies below, and would be happy to provide committee members additional resources.

Transparency. Generative AI products – including chatbots, AI-altered images, and persuasive AI voice clones – have the potential to deceive and confuse consumers. For example, Consumer Reports has heard from members who believe scammers used AI to impersonate the voices of family members in desperate need of money. Customer service chatbots that aren't clearly labeled as bots can cause frustration. AI-altered images, like deepfake celebrity product endorsements, can mislead consumers, and cause reputational harm to the person depicted. CR believes that it should be clear to consumers when they are interacting with AI, and that AI-generated content should be labeled when it might cause confusion or deception. Companies and platforms should take steps to reduce the likelihood that their products are used to impersonate, deceive, or otherwise break the law.

Transparency is also an issue for other types of AI, sometimes known as analytical (or discriminative) AI. Analytical AI refers to a broad category of products that includes everything from fraud detection tools, to software that uses AI to analyze job applications and make hiring recommendations, to AI and algorithmic tools that insurance companies use in underwriting and pricing decisions. Analytical AI tools may draw on vast quantities of consumers' personal data—such as location data, social media profiles, and online shopping habits—to make predictions. These tools can be faulty, relying on inaccurate data,

[1] Founded in 1936, Consumer Reports (CR) is an independent, nonprofit and nonpartisan organization that works with consumers to create a fair and just marketplace. Known for its rigorous testing and ratings of products, CR advocates for laws and company practices that put consumers first. CR is dedicated to amplifying the voices of consumers to promote safety, digital rights, financial fairness, and sustainability. The organization surveys millions of Americans every year, reports extensively on the challenges and opportunities for today's consumers, and provides ad-free content and tools to 6 million members across the U.S.

entrenching bias, or making predictions not rigorously supported by the data. For example, AI background check companies have been sued for making mistakes that cost people job opportunities and income.

When companies use AI and algorithmic tools to help make high-stakes decisions about consumers' lives—such as decisions about access to employment, housing, insurance, lending and financial products, and more—we believe consumers have a right to know how the system is evaluating them. Once the decision is made, consumers should be presented with an explanation of the key factors, and have an opportunity to correct any incorrect personal information.

This comports with results of a recent nationally representative survey of 2,022 US adults CR conducted on the use of AI and algorithms in high-stakes decisions about consumers. When asked about a scenario in which an AI program or algorithm had been used to determine whether or not they would be interviewed for a job, most Americans (83%) said they would want to know specifically what information about them the program used to make that decision. Even more (91%) said that they'd want the opportunity to correct any incorrect data if they found out that that program made the decision by relying, in part, on incorrect information.

Recommendations related to transparency:

- **Require clear disclosure when an algorithmic tool is being used to help make a consequential decision about a consumer,** like whether they qualify for a loan, are selected for a rental apartment, get a promotion, or see their insurance rates or utility bills go up.
- **Require companies to explain why a consumer received an adverse decision** when an algorithmic tool was used to help make a consequential decision about a consumer. Explanations must be clear enough that, at a minimum, a consumer could tell if the decision was based on inaccurate information. Explanations should include actionable steps consumers can take to improve their outcome. **If a tool is so complex that the company using it cannot provide specific, accurate, clear, and actionable explanations for the outputs it generates, it should not be used in consequential decisions.**
    - Already, lenders are required to provide specific and accurate reasons when they deny someone credit or take other adverse actions under the federal Equal Credit Opportunity Act. In 2022, the Consumer Financial Protection Bureau clarified that creditors using complex algorithms still have to comply with this law, and explained that if an algorithm is so complex that creditors cannot provide specific and accurate reasons why a person was denied a loan, they cannot use that algorithmic tool. However, notices provided under the Equal Credit Opportunity Act are still often insufficient in practice.
- **Require disclosure of AI-generated content when it might cause confusion or deception**, including text, video, audio, and images. For example, chatbots that a reasonable person might misunderstand as a human-to-human interaction should be labeled.
- **Prohibit some harmful categories of deepfakes**, including non-consensual pornographic content, materials that use someone else's likeness for commercial gain without permission, or fraudulent material.

Fairness: Problems with bias in AI systems are well documented. For example, when Amazon developed a machine-learning based recruiting engine, engineers discovered that the system was downgrading applicants with the word 'women's' in their application, such as 'women's chess captain.' The engine had been trained on patterns in the past decade of applications to Amazon, most of which came from men. A different resume screening program identified two factors as the best predictors of future job performance: having played high school lacrosse and being named Jared.

State and federal laws already prohibit discrimination in many contexts, and there's no exemption for AI and algorithmic decision tools. However, it is worth identifying any potential gaps, and clarifying how these laws apply to companies developing and deploying AI and algorithmic tools, as well as expectations for compliance. Companies should also assess their products for discrimination throughout the development process.

Transparency for consumers will help address biased systems too. Consumers should be provided with a specific and accurate explanation when they are denied an important life opportunity. This is already the norm under federal consumer credit law: when creditors take adverse actions against you, they are legally required to provide you an explanation under the federal Equal Credit Opportunity Act.

Some artificial intelligence tools are so complex that the companies using them do not know why they produce the specific recommendations they do. Those tools should not be used to make high-stakes decisions about consumers. This principle is recognized under federal guidance related to the Equal Credit Opportunity Act as well.

Recommendations related to fairness:

- **Prohibit algorithmic discrimination.** Existing civil rights laws prohibit many forms of discrimination, but it is worth identifying any potential gaps and clarifying how these laws apply to companies developing and deploying AI and algorithmic tools, as well as how they are expected to comply and how these laws will be enforced when it comes to certain uses of AI.
- **Establish an affirmative duty for companies to search for and implement less discriminatory algorithms used for consequential decisions.** Often, there are alternative algorithms that are similarly effective and less discriminatory. Companies should be required to consider fairness as a performance metric and search for less discriminatory alternatives throughout the model development process.
- **Require companies to have appropriate data governance frameworks in place to address bias.** This includes ensuring appropriateness and suitability of data and addressing any issues with accuracy, completeness, and representativeness of data or historical bias in data.
- **Clarify acceptable, appropriate use of data on sensitive characteristics, such as race, for purposes of assessing bias.**
- **Require companies that make tools used in consequential decisions** to have internal processes to assess and remediate the risk of harm at all stages of product development.
- **Require companies that make tools used in consequential decisions to undergo independent, third-party testing** for bias, accuracy and more pre-deployment, and regularly after deployment.

Clear and consistent standards should be developed for testing and third-party auditors should be regulated by relevant regulators or bodies.

- **Limit self-preferencing:** Big Tech shouldn't use their AI models to preference their own products or services when doing so would materially harm competition.

Privacy: Artificial intelligence has the potential to supercharge targeted advertising, crafting unique messages based on what companies know about a specific consumer. These practices could cross the line into being manipulative. Consumer Reports has been involved in New York's effort to pass privacy legislation, and hopes to see the Empire State sign a strong privacy bill into law. Our core principle when it comes to consumer data is "data minimization"; that is permitting companies only to collect and process the data necessary to fulfill the services consumers request. This same principle applies to AI products.

Recommendations related to privacy:

- **Require data minimization across AI tools**. Companies should be required to limit data collection, use, retention, and sharing to what is reasonably necessary to provide the service or conduct the activity that a consumer has requested, with limited additional permitted uses.
    - Consumer Reports has done extensive advocacy and research on pro-consumer privacy laws. To read more about data minimization, check out our model state privacy bill and our white paper about how data minimization could work in practice.
- **Prohibit retaliation against consumers who exercise privacy or procedural rights,** such as the right to a clear, actionable, specific, and accurate explanation of an algorithmic decision.
- **Prohibit the sale and sharing of personal data collected by generative AI tools to third parties.**
- **Ban remote biometric identification in publicly accessible spaces**, including retail, with limited exceptions. Consumers cannot consent to be tracked if their only alternative is to not enter public/semi-public spaces.
- **Require companies use reasonable cyber and data security measures when developing their products** and continue to provide security support, like updates and patches, over time.
    - Again, this is arguably already covered by Section 5 of the FTC Act, which the agency has used to bring cases against companies that fail to adequately protect consumers' personal information.
    - However, the law is less clear on the question of companies' liability for failing to maintain product security over time.

Snake oil and substantiation: Some AI products do not live up to their marketing hype. When MIT Technology Review tested Curious Thing, a product that was marketed at the time as an AI job interview assessment tool, their tester received a 6 out of 9 for English competency despite answering all of the questions in German. Companies such as HireVue have developed AI assessments used by major employers to analyze candidates' choice of words, facial expressions, speaking voice, perceived "enthusiasm," and more to generate an "employability" score. AI researchers have described using facial expressions and speaking voice to predict job performance as "pseudoscience" and argue that applicants may be penalized for traits that may not ultimately relate to job performance, such as unusual mannerisms, or having a stutter.

When companies over promise or deceive consumers about what their AI products can do, they are likely violating state consumer protection laws that prohibit deceptive practices. Consumer Reports believes that companies should have to substantiate the claims they make about their AI products.

Recommendations related to snake oil and substantiation:

- **Require companies to meaningfully substantiate the claims they make when describing or marketing their AI products.**
    - State consumer protection laws already prohibit deceptive practices. Section 5 of the Federal Trade Commission Act also prohibits unfair or deceptive acts or practices.

Testability: Artificial intelligence poses some unique challenges for third-party testing. Companies can control researchers' access to software and technology in a way they cannot control access to physical products available in stores, like washing machines and blenders. For example, companies can cut off researcher's access to data, as Facebook did when it shut down the accounts of the researcher's behind New York University's Ad Observatory, a project that studied the ads people see on Facebook. Tech companies can also frustrate external testing by adding clauses to their terms of service that prohibit users from sharing information about their accounts with researchers. Other legal impediments exist: it's unclear whether companies would be successful in claims that public interest auditors have violated state and federal anti-hacking laws, such as the Computer Fraud and Abuse Act.

Policymakers should consider targeted reforms of these state and federal anti-hacking laws to ensure that good faith public interest research that does not meaningfully tax a company's resources or compromise other interests (such as privacy) is allowed. Legislators should also consider enacting legislation that explicitly prohibits contractual language unfairly limiting researchers' ability to audit algorithms for bias. Lastly, we recommend enacting protections for whistleblowers. When individuals bring up issues internally to upper management those issues don't get addressed in a reasonable time period, those individuals should be protected from retaliation if they come forward.

Recommendations related to testability:

- **Require companies that develop algorithmic tools that help make consequential decisions to provide access to vetted, public interest researchers.** To get a thorough understanding of how these tools work and where they fall short, independent, high quality research must be conducted.
- **Whistleblower protections and incentives:** Whistleblowers can expose problems to the public that companies have no real incentive to disclose or address. But right now, there are few protections for people who want to disclose issues surrounding AI.
- For more, read Consumer Reports' research on legal barriers to public interest algorithmic auditing

Other issues: There are other issues and remedies that policymakers ought to consider. One such issue is platform accountability. It is trivially easy to run a Google search and find software to clone other people's voices without their consent, unlocking new avenues for scammers. A person using AI to commit fraud should bear responsibility for violating the law, but the companies that provide these tools and the

platforms that make them easily accessible should perhaps also bear some liability and should be required to exercise reasonable care.

Another issue is intellectual property. Generative artificial intelligence models are trained by digesting vast swaths of the internet, including teenager's video blogs, published books, chart-topping songs, and original journalism. This practice may prove to be in violation of intellectual property law: several lawsuits, including cases brought by authors such as John Grisham and publications such as the New York Times, are currently pending. It may also violate consumers' reasonable expectations about what happens to the photos, words, and other information they share online.

We appreciate the opportunity to share our perspective and the work New York policymakers have already undertaken to create a fairer digital marketplace.

Best regards,

Grace Gedye
Policy Analyst
Consumer Reports