

Testimony of the Center for Democracy & Technology on Senate Bill 2

Testimony for the Connecticut Senate General Law Committee

February 29, 2024

The Center for Democracy & Technology (CDT) thanks the General Law Committee for considering our testimony on this legislation. CDT is a nonprofit, nonpartisan organization fighting to advance civil rights and civil liberties in the digital age.

Companies and government agencies are increasingly using artificial intelligence (AI) tools and other automated decision systems (ADSs) to make decisions that dramatically impact the lives of consumers and workers. CDT has been closely watching as policymakers across the country have grappled with how to manage the potential benefits and risks that AI and ADSs pose.

We greatly appreciate the effort that Sen. Maroney and the Committee members have clearly put into crafting this sweeping legislation, which would cover algorithmic decision-making in settings as diverse as employment, housing, health care, and criminal justice. A great deal of effort and thought clearly went into assembling this ambitious bill. The suggestions below reflect the importance of the issues and how essential it is to get this legislation exactly right--or as close to exactly right as possible. We hope the Committee will take the below suggestions in that constructive spirit.

Take the time to consult carefully with key stakeholders before moving this bill

Given the bill's broad scope, we urge the Committee to listen closely to advocates for consumers and workers with expertise in the various subject-matter areas that the bill covers. It is true that artificial intelligence is already being widely used. It is equally true, however, that many existing laws already protect workers' and consumers' rights and that moving to enact legislation without taking the time necessary for adequate consultation with key stakeholders risks creating unintended conflicts with those laws. It would be far better to wait a few months--or even until the next legislative session--to craft a bill that builds on existing civil rights, consumer, and labor protections rather than confusing or undermining them. As some of our suggestions below indicate, the bill's current text would, unfortunately, do precisely that.

Definitions and Scope

Eliminate the requirement that a covered system be a "controlling factor" in a consequential decision.

The current language of the bill would only apply to systems that are "specifically developed and marketed, or specifically modified, to make, or be a controlling factor" (automated decision tool) or that "when deployed, make[s], or [are] a controlling factor in making," a consequential decision. This is inconsistent with existing laws. Civil rights protections (including in

employment, education, and housing) generally apply not only to decisions based *solely* on unlawful discrimination, but to any decision where impermissible factors influenced the outcome. This approach is essential because otherwise, a plaintiff's case becomes virtually impossible to prove. It would be inappropriate to subject AI systems to a lower standard.

Moreover, the "controlling factor" requirement would have the practical effect of giving companies the ability to opt out of complying with the bill. That's because, in employment, criminal justice, housing, and many other contexts, vendors almost invariably say that their systems are designed merely as tools to *assist* humans, and deployers always say that humans have final say in decisions--even if, in reality, the tools' "recommendations" are decisive and human "reviewers" defer to AI outputs. It would be trivially easy for developers to avoid compliance by including a disclaimer in their marketing materials stating that a tool "is not designed to be a controlling factor in any decision," and for deployers to avoid compliance by having a human rubber-stamp algorithmic recommendations.

Indeed, the "controlling factor" requirement creates a catch-22 that effectively makes it impossible for anyone to challenge a deployer or developer's assertion that an AI system is exempt from the law. Once a company chooses to assert that a tool does not meet the "controlling factor" requirement, it need not even disclose the existence or use of the AI tool. In that case, consumers, workers, and regulators may not even be aware of the tool, and thus will not be able to challenge the employer's assertion that the tool is not being used in a manner that has a controlling impact on a decision. In effect, that means deployers and developers would have the unilateral ability to opt out of complying with the law.

Relatedly, [a recent study by researchers at Cornell University, Consumer Reports, and Data & Society](#) showed that few companies are conducting audits or making disclosures required by New York City's AI hiring law, which similarly is limited to tools that have a dominant effect on the decision-making process.

Recommended Revisions

- Revise the definition of "automated decision tool" and "high-risk artificial intelligence system"¹ so that they cover any system that "makes, or assists in making," a consequential decision.

Delete the definition of "high-risk artificial intelligence system" and simply use the definition of "automated decision tool"

The definitions of "automated decision tool" and "high-risk artificial intelligence system" are already nearly redundant; the deletion of the "controlling factor" requirement from both definitions, as recommended above, would make them wholly redundant. Consequently, the definition of "high-risk artificial intelligence system" should be deleted, and the more straightforward term "automated decision tool" should be used instead.

¹ But we recommend deleting the latter definition, as discussed immediately below.

Recommended Revisions

- Delete the definition of “high-risk artificial intelligence system”
- Replace “high-risk artificial intelligence system” with “automated decision tool” wherever the former term appears

Rely on existing definitions of unlawful discrimination instead of creating a new and untested definition of “algorithmic discrimination”

Antidiscrimination laws have a long history, with extensive interpretative case law at both the state and federal levels in employment, housing, education, and the other consequential decision settings that SB 2 covers. It is unnecessary and potentially misleading to create a new, untested definition of “algorithmic discrimination,” particularly since the proposed definition uses the undefined term “unjustified differential treatment or impact.” Courts will be left to grapple with what makes differential treatment or impact “[un]justified,” since that term is not defined either in this statute or in existing antidiscrimination laws.

Introducing this new definition, rather than relying on definitions and interpretations developed through existing laws and regulations, creates the harmful impression that discrimination by a machine somehow deserves different, special treatment as compared to discrimination through other means. This will create confusion and could easily lead to different and potentially contradictory standards for algorithmic and non-algorithmic discrimination. That would benefit neither businesses nor consumers and workers.

Recommended Revisions

- Delete the definition of “algorithmic discrimination”
- Add definition of “unlawful discrimination” that simply reads: “Unlawful discrimination’ means any discriminatory practice or discriminatory employment practice, as defined in Sec. 46a-51, or any other practice or act that otherwise violates a law against discrimination in this state.”
- Replace “algorithmic discrimination” with “unlawful discrimination” wherever the former term appears

Delete the redundant and potentially confusing definition of “artificial intelligence” and use the same definition of “artificial intelligence system” throughout the bill

As it stands, the bill includes definitions of both “artificial intelligence” and “artificial intelligence system.” The former term is only used in isolation once in sections 1-7: in the definition of “automated decision tool.” There is no reason to use a special definition just for that one term. Moreover, the definition of “artificial intelligence system” is based on the OECD definition, which has gained currency as a preferred definition in legislation, regulations, and policy documents across the country and, indeed, the world. The bill should be amended to eliminate the redundant and less widely used definition.

Relatedly, the same definition of “artificial intelligence system” should be used throughout the bill. It is not clear why the term should mean one thing in the context of automated employment decisions (Sec. 1 to Sec. 3) but something else in the context of deceptive media (Sec. 11) or in government agency settings (Sec. 12 and Sec. 13).

Recommended Revisions

- Delete the definition of “artificial intelligence.”
- Amend the definition of “automated decision tool” so that it refers to “any system or service that (A) is an artificial intelligence system” rather than “any system or service that (A) uses artificial intelligence.”

Notice and Transparency

Expand the notice requirements so that consumers and workers actually receive meaningful notice of how a tool works

Unlike with traditional decision-making processes, people frequently do not know when AI is evaluating them, much less how they will be evaluated. Without strong notice provisions, citizens will be unable to exercise their rights under existing laws. SB 2’s requirement that deployers provide impacted individuals with a pre-evaluation notice that includes “a plain language description of” the system is unhelpfully vague, and the bill includes no post-assessment notice or explanation requirements whatsoever. The existing notice provisions would not provide consumers or workers with the information needed to determine whether a tool has violated or may violate their legal rights.

Consumers and workers need true transparency. This means, at a minimum, disclosure of each ADS a company uses on them, what types of decisions an ADS makes, what personal data and attributes an ADS uses, how it uses that information to make decisions, and explanations of adverse decisions. The bill should be amended to ensure individuals receive such notice.

Recommended Revisions

- Expand pre-evaluation notice requirements in section 3(e)(1)(B) to give individuals the right to know what information the tool uses and how the tool uses that information to make a decision:

[Deployer must provide the consumer with] a description, in plain language, of such high-risk artificial intelligence system, which description shall, at a minimum, include a description of

- (I) the personal characteristics or attributes that the system will measure or assess, the method by which the system measures or assesses those attributes or characteristics, how those attributes or characteristics are relevant to the consequential decisions for which the system should be used,*
- (II) the system’s outputs,*

- (III) The logic used in the system, including the key parameters that affect the output of the system;*
- (IV) the type(s) and source(s) of data collected from natural persons and processed by the system when it is used to make, or assists in making, a consequential decision,*
- (V) the results of the most recent impact assessment, or an active link to a webpage where a candidate can review those results,*
- (VI) any human components of such system, and*
- (VII) how any automated components of such system are used to inform such consequential decision.*

Right to specific, accurate, and actionable explanation and barring tools whose output are not easily explainable

An algorithmic tool that is used to make consequential decisions about consumers should be able to produce specific and accurate explanations of why it generates the outputs it generates, such that individuals understand the output and whether that output is based on inaccurate information or inferences about them. When an ADS is used to make, or assist in making, a consequential decision, consumers should thus receive a simple, actionable explanation of why the decision was made. When a deployer cannot provide such an explanation, it should not use the ADS.

Recommended Revisions

- Add the following text to section 3(e):
 - (2) A deployer shall, within 14 days after an automated decision tool is used to make, or assists in making, a consequential decision, provide any natural person that was the subject of the consequential decision:
 - (A) A simple and actionable explanation that identifies the principal factors, characteristics, logic and other information related to the individual that led to the consequential decision;*
 - (B) The role that the automated decision tool played in the decision-making process; and*
 - (C) A meaningful opportunity to submit corrections or otherwise provide supplementary information relevant to the consequential decision.**
 - (3) No deployer shall use an automated decision tool to make, or assist in making, a consequential decision if it cannot provide an accurate notice that satisfies the requirements of paragraphs (1) and (2) of subdivision (c) of this section.*

Require notices/explanations to be prepared and presented in a manner that ensures effective disclosure

The present language of the bill allows deployers to choose the manner in which they provide notices, subject only to the requirement that they choose a method of disclosure that is “clear

and readily accessible.” This language is far too vague and permissive; it does not ensure that consumers or workers will receive actual notice of the information that the bill ostensibly requires deployers to provide them, much less ensure that it is presented in a manner that consumers and workers will easily understand. Instead, The bill should require deployers to take specific steps to ensure that Connecticut residents subjected to consequential decisions by AI systems receive actual notice in languages and formats that clearly and effectively convey the required information.

Recommended Revisions

- Replace language in section 3(e) allowing deployer notices to be presented “in any manner that is clear and readily accessible” with a requirement that notices be:
 - *Transmitted directly to the subject of the consequential decision when possible, or else made available in a manner reasonably calculated to ensure that the subjects of consequential decisions receive actual notice;*
 - *Provided in English, in any non-English language spoken by at least one percent (1%) of the population of this state as of the most recent United States Census, and in any other language that the deployer regularly uses to communicate with the subjects of consequential decisions;*
 - *Written in clear and plain language;*
 - *Made available in formats that are accessible to people who are blind or have other disabilities; and*
 - *Otherwise presented in a manner that ensures the communication clearly and effectively conveys the required information to subjects of the relevant consequential decisions.*

Impact Assessments

Require that deployer impact assessments be conducted by an independent third party

The bill seeks to address potential harms of AI systems via impact assessments, but the current permissive language will allow those impact assessments to be ineffective. There have already been multiple instances where vendors published misleading impact assessments, in which the company either conducted the impact assessment themselves and seemed to cherry-pick the data points to present or retained a third party that was only granted partial access to relevant data. Such an impact assessment is not reliable. Third-party auditors who have full access to AI systems and are free of conflicts of interest are more likely to analyze and publish truthful assessments.

Recommended Revisions

- Revise section 3(c) so that deployer impact assessments must be conducted by an independent entity, using the definition of independent auditor from [the Lawyers' Committee for Civil Rights Under Law's Model AI Bill](#):

(c)(2)(D) The impact assessment required by this subdivision must be completed by an independent person or entity who exercises objective and impartial judgment on all issues within the scope of the impact assessment. Each deployer or developer of the automated decision tool shall provide the independent auditor with all information and data regarding the design, functionality, testing, and performance of the automated decision tool. A person is not independent for purposes of this subparagraph if they:

(1) Are or were involved in using, developing, offering, licensing, or deploying the automated decision tool;

(2) At any point during the impact assessment, has an employment relationship with a developer or deployer that uses, offers, or licenses the automated decision tool; or

(3) At any point during the impact assessment, has a direct financial interest or a material indirect financial interest in a developer or deployer that uses, offers, or licenses the automated decision tool.

- Make related edits to prefatory language in section 3(c)(1) so that it reads:

(A) No deployer may deploy a high-risk artificial intelligence system on or after July 1, 2025 unless the automated decision tool has been the subject of an impact assessment for the high-risk artificial intelligence system; and

(B) Beginning on July 1, 2025, impact assessments must be completed for a deployed high-risk artificial intelligence system not later than ninety days after any intentional and substantial modification to such high-risk artificial intelligence system is made available.

Require deployers to mitigate identified discrimination risks before deploying a decision tool

Currently, the bill would require developers to inform deployers of “known or reasonably foreseeable risks” of algorithmic discrimination and the “measures the developer has taken to mitigate” those risks. Similarly, deployers would have to conduct impact assessments to identify potential discrimination risks and analyze “the steps that have been taken to eliminate such risks.” Nothing in the bill, however, would actually require developers or deployers to resolve discrimination risks in their AI systems.

Recommended Revisions

-
- Add the following as a new subparagraph to section 3(c)(2): “If the analysis required by subparagraph (A)(ii) of this subdivision reveals a risk of algorithmic discrimination, or the analysis required by subparagraph (A)(viii) of this subdivision reveals a potential limitation on accessibility, the deployer shall not deploy the automated decision tool until the risk is mitigated.”

Require impact assessments to be specific to the system and the setting in which the decision is made

The current text would allow companies to avoid conducting an impact assessment specific to the system and decision setting in question if they (1) conduct an impact assessment addressing a “comparable” set of systems or (2) have previously conducted an impact assessment on the system that is “reasonably similar in scope” to that required by SB 2. The quoted terms are undefined and give developers and deployers far too much latitude to decide whether they must conduct an impact assessment specific to the system and the setting in which a consequential decision is made.

Moreover, the undefined “comparable” and “reasonably similar” standards risk creating inconsistencies with existing antidiscrimination laws. For example, in employment, a worker can make a *prima facie* case of disparate impact discrimination by demonstrating that an employer’s selection practice has an adverse impact in the context of a specific job at a specific employer. Under federal law, an employer can rebut that *prima facie* case only by showing that a selection procedure is “job related *for the position in question* and consistent with business necessity.”² Federal regulations make clear that employers cannot demonstrate job-relatedness or otherwise defeat a plaintiff’s *prima facie* case merely by showing that the selection procedure has been validated for a “comparable” job; instead, the employer must demonstrate that the position for which a previous validation study was conducted was “substantially the same” with respect to all major work behaviors for the position in question.³ Setting a new, seemingly looser standard for impact assessments in the context of algorithmic tools would be incongruous and inappropriate.

Recommended Revisions

- Section 3: Delete subdivisions (3) and (4) in subsection (c)
- Section 4: Delete subdivisions (3) and (4) in subsection (c)

Right to Opt Out of Automated Decisions

Give the subjects of automated decisions the right to opt out and have the decision made through other means

When consequential decisions are left to automated systems, consumers and workers should have the right to access the information upon which the decision was based, to obtain an explanation as to the reasons for the decision itself, and to opt out of automated decision-making and request human review.⁴ The preceding recommendations cover the former

² 42 U.S.C. § 2000e-2(k)(1)(A)(i).

³ 29 C.F.R. § 1607.7(B).

⁴ See White House Office of Science & Technology Policy, Blueprint for an AI Bill of Rights (2022), <https://www.whitehouse.gov/ostp/ai-bill-of-rights/> (“You should be able to opt out, where appropriate, and have access to a person who can quickly consider and remedy problems you encounter. You should be able to opt out from automated systems in favor of a human alternative, where appropriate. Appropriateness should be determined based on reasonable expectations in a given context and with a

two elements; the additional provision of opt-out rights would further help reduce the risk of discrimination and other harms, such as by giving disabled consumers and workers the right to opt out of decision-making processes for which they cannot obtain adequate accommodation, or where they otherwise believe the automated system will not make a fair and accurate decision due to their disability. When combined with proper notice and explanation, opt-rights allow affected individuals an opportunity to raise concerns, request accommodation, and make an informed decision about whether, when, and how to proceed with the automated decision-making process.

Recommended Revisions

- Add the following as a new subsection in section 3:

(1) If a consequential decision is made in whole or in part based on the output of an automated decision tool, a deployer shall comply with a natural person's request to not be subject to the automated decision tool and to be subject to an alternative decision process or accommodation.

(2) After a request pursuant to subdivision (1) of this subsection, a deployer may reasonably request, collect, and process information from a natural person for the purposes of identifying the person and the associated consequential decision. If the person does not provide that information, the deployer shall not be obligated to provide an alternative decision process or accommodation.

(3) A deployer shall not retaliate against any natural person who makes a request pursuant to subdivision (1) of this subsection.

(4) For purposes of employment, this subsection shall not apply in any instance where any other law or regulation of this state provides stronger protections.

Enforcement, Defenses, and Burdens of Proof

Eliminate provisions creating unique burdens of proof or rebuttable presumptions in favor of developers or deployers

The proposed language creates a rebuttable presumption that “reasonable care” was taken to avoid “algorithmic discrimination” if an employer complies with the bill’s impact assessment and notice provisions. This unacceptably undermines Connecticut’s existing laws against discrimination (and is especially damaging given the weaknesses in the impact assessment and notice provisions as described above). At best, SB 2’s new approach to burden-shifting would create confusion if an AI system were the subject of a claim under Connecticut’s laws against discrimination. At worst, it would signal that algorithmic discrimination is somehow more acceptable, and thus can meet lower standards, than discrimination through other means. The existing burdens of proof and presumptions that apply in the relevant decisional settings should continue to apply in the context of automated decisions.

focus on ensuring broad accessibility and protecting the public from especially harmful impacts. In some cases, a human or other alternative may be required by law.”).

Recommended Revisions

- Delete the “rebuttable presumption” provisions in sections 2(a), 3(a), and 4(a)(2) and replace it with language making it clear that compliance with SB 2 does *not* provide a defense to a claim under any other law, *i.e.*: “A [developer’s/depoyer’s] compliance with this section shall not constitute a defense in a civil or administrative action regarding claims that the developer violated any other section in this act or any other law.”

Make violations of section 3 a discriminatory practice subject to enforcement under Connecticut’s existing human rights laws

Giving exclusive enforcement authority to the Attorney General--a generalist office without specific expertise in the discrimination laws that undergird Section 3--may result in an overburdened AG and an underenforced law. The Commission on Human Rights and Opportunities is the state agency with expertise in Connecticut’s antidiscrimination laws. Depriving that agency of authority to investigate and enforce this law will create confusion and inefficiency if an AI system ends up being subject to a discrimination claim. Additionally, the Human Rights Law provides a private right of action, which is by far the most effective way to ensure companies are held accountable for violations of this law.

Recommended Revisions

- Delete references in section 7 to the Attorney General having “exclusive” or “sole” enforcement authority
- Add the following as a new subdivision in Section 7: “A violation of the requirements established in section 3 of this act shall constitute a discriminatory practice for purposes of section 46a-51 of the general statutes.”

Eliminate the sweeping affirmative defense provided in section 7(f)

Similarly, Section 7(f) creates an affirmative defense that would allow developers and deployers to escape liability if they adopt completely untested risk-management frameworks, provide a *pro forma* mechanism for feedback from users (with no obligation to make it easy for users to submit such feedback--much less to review or act on such feedback), or engage in certain acts designed to mitigate harmful deployments of automated systems. Passing a law with this affirmative defense would dramatically undermine enforcement. It would also create potential for confusion and inconsistencies, forcing courts to grapple with the issue of whether companies complied with unfamiliar risk management frameworks that were not designed to form the basis for legal defenses in cases involving civil rights or consumer and worker protections.

Recommended Revisions

- Delete section 7(f)

Eliminate the right to cure for Attorney General enforcement

The right to cure dramatically weakens enforcement, especially when enforcement is already narrowly cabined in the bill. The right to cure is also highly unusual under the law—wrongdoers rarely get a free bite at the apple, to simply stop lawbreaking once they are caught and thus avoid any consequences. And it is unfair to the victim who is denied a remedy simply to give the wrongdoer a chance to cure. This concept is foreign to the antidiscrimination laws this bill is designed to build upon, and it should not be introduced through this legislation.

Recommended Revisions

- Delete the first two sentences of section 7(b)

Exemptions from Disclosure and FOIA Obligations

Eliminate provisions (1) allowing developers and deployers to avoid making crucial disclosures by claiming confidential information protections and (2) creating special protections for deployers and developers under public record laws

It is understandable to balance consumers' and workers' need for disclosure against companies' need to protect their intellectual property rights. However, the current bill goes too far in creating exceptions to its developer disclosure requirements. It creates carve-outs not only for trade secrets, but also for undefined "confidential or proprietary information." Similarly, the bill contains special exemptions that would give developers and deployers of automated tools special protection from having materials lawfully obtained by the Attorney General exempted from disclosure under Connecticut's Freedom of Information Act (CT FOIA).

These protections are unwarranted and would severely undermine the effectiveness of both this bill's disclosure requirements and of Connecticut's public records laws. The broad carveout for "confidential and proprietary information" is so vague and broad that the exception will likely swallow the rule, as companies may attempt to designate virtually all information concerning their decision-making processes as "confidential." Additionally, CT FOIA already contains provisions protecting trade secrets and certain other sensitive information from disclosure. See Ct. Gen. Stat. Sec. 1-210(5). Developers and deployers of AI systems can rely on those protections and do not need or deserve new, special protections that other businesses do not enjoy.

Recommended Revisions

- Delete "or other confidential or proprietary information" from section 2(f)
- Delete FOIA exemptions in sections 2(g), 3(i), and 4(d)

Artificial Intelligence Advisory Council

Task the proposed Council with reviewing the effectiveness of SB 2’s transparency and impact assessment requirements

The notice and impact assessment requirements of SB2 are an important part of the bill’s regulatory objectives, and for the reasons described above, it is essential that they function well. The bill should thus explicitly task the Artificial Intelligence Advisory Council with reviewing the effectiveness of those components of this legislation as part of its work, with a goal of revisiting the provisions to add further implementing guidance if they are not achieving their intended goals.

Recommended Revisions

- Add the following subdivision to Sec. 8(b): “Review and make recommendations related to improving the utility and effectiveness of notice, disclosure, and impact assessments in addressing known and reasonably foreseeable harms arising from the use of artificial intelligence systems and automated decision tools;”
- Add the following to the end of current Sec. 8(b)(3): “make recommendations concerning the adoption of other legislation concerning artificial intelligence and automated decision tools”

Prohibition against “deceptive media” within 90 days of election

Limit application of Section 11 to depictions of candidates for office

Section 11 would prohibit using AI to depict people doing things they did not do within 90 days of an election unless the use of AI is labeled prominently. While this is very much a worthy objective, the scope of the bill is potentially overbroad--applying not only to depictions of candidates for office, but to depictions of all “human beings” regardless of whether or not that person is a candidate for any office. Further, its execution risks creating different standards for similar types of misleading materials. For instance, this section seemingly would prohibit the use of AI generated images of people to dramatize examples of policy concerns without labeling the images, but would not apply in the same way to using actors to dramatize the same example. It is not clear why these situations should be treated differently.

It would make more sense to apply the restriction only to depictions of candidates for office. That would bring this section closer to “electioneering” restrictions and to the policy goal of preventing the creation of false AI generated images that mislead voters about candidates.

Recommended Revisions

- Replace “human being” with “candidate” in the definition of “deceptive media”

State Agency Studies of AI Use

Tighten review requirements in sections 12 and 13 for state agency uses of AI and strengthen associated documentation and reporting requirements

Sections 12 and 13 provide an excellent foundation for studying the ways in which state agencies are using artificial intelligence systems. We suggest a number of revisions, detailed below, that would strengthen these review requirements and ensure that the public can more readily access information regarding how various state agencies are using artificial intelligence systems.

Recommended Revisions

- Add “and shall document the privacy, civil rights, and equity implications of such uses” after “to improve efficiencies” in Sec. 12(b)
- Add new subdivision (3) after subdivision (2) in section 12(b): “Any such pilot project shall measure how generative artificial intelligence . . . (3) *is likely to contribute to or reduce discrimination and disparities in access to government services and benefits*”
- Add the following to the final sentence of Sec. 12(d): “Such report shall include a summary of all pilot projects approved by the commissioner under this section, *the impact assessments associated with such programs*, and any recommendations for legislation...”
- In Sec. 13(b)(1), add the following as additional subparagraphs:
 - The date upon which each use case was added and updated
 - In what context(s) the system has been deployed
- In subparagraph (d)(2) of Sec. 13, add “documenting” and “reporting” so that the beginning of the subparagraph reads: “(2) methods for identifying, *documenting, reporting*, and mitigating potential output inaccuracies, fabricated text...”

Expand scope of review to cover uses of all artificial intelligence systems, not just those that use generative artificial intelligence

There are several references in Sections 12 and 13 to the need to study potential uses or impacts of “general artificial intelligence” systems, but not to other types of AI. It is not clear why the studies required in these sections should be limited only to generative systems. Many types of AI, both generative and non-generative, could improve the efficiency of state operations or be used in ways that harm Connecticut residents. Sections 12 and 13 should be amended to consider potential uses and impacts of all types of artificial intelligence.

- Replace “general artificial intelligence tools” with “artificial intelligence systems” in Sec. 13(d)(1)
- Replace “general artificial intelligence” with “artificial intelligence systems” in Secs. 12(b), 12(c) (both occurrences), and 13(d)(2)⁵

⁵ See also our recommendation in the “Scope and Definitions” section of this testimony to use the same definition of “artificial intelligence system” throughout the bill.

Conclusion

It is not merely possible, but imperative, to do AI regulation well. To that end, we urge the Committee members to consider the recommendations in this document, as well as those submitted by other labor and consumer advocates, as it considers this bill. Please do not hesitate to contact Matthew Scherer, Senior Policy Counsel on CDT's Privacy & Data Project, at mscherer@cdt.org if you have any questions regarding this testimony or if we can provide additional resources or information that will assist in your consideration of this legislation. Thank you for your attention.