SENATE JUDICIARY COMMITTEE
Senator Thomas Umberg, Chair
2023-2024 Regular Session

AB 302 (Ward)
Version: May 18, 2023
Hearing Date: July 6, 2023
Fiscal: Yes
Urgency: No
CK

SUBJECT

Department of Technology: high-risk automated decision systems: inventory

DIGEST

This bill requires the California Department of Technology (CDT), on or before September 1, 2024, to conduct a comprehensive inventory of all high-risk automated decision systems (ADS) that have been proposed for use, development, or procurement by, or are being used, developed, or procured by, any state agency.

EXECUTIVE SUMMARY

ADS are algorithm-driven applications that can assist or supplant human decisionmaking processes in areas such as credit decisions, employment screening, insurance eligibility, and the delivery of government services. ADS process enormous datasets and make decisions with speed and reliability that vastly exceed human capabilities. However, poorly designed or poorly understood systems can create unfair, biased, and inaccurate results. When deployed by government agencies, flawed ADS may disproportionately harm low-income families and communities of color and undermine trust in the public sector. Moreover, norms of participatory governance and due process may be jeopardized when ADS affect agency policymaking, adjudications, or enforcement.

The bill seeks to promote the oversight of automated decisionmaking in California by requiring CDT to conduct a comprehensive inventory of all high-risk ADS, which is ADS that is used to assist or replace human discretionary decisions that have a legal or similarly significant effect, including decisions that materially impact access to housing and employment. CDT is then required to submit a report of the inventory to the Legislature. The bill is sponsored by the Greenlining Institute and supported by civil rights and consumer protection advocates, including the TechEquity Collaborative and Equality California. There is no known opposition. It passed the Senate Governmental Organization Committee by a vote of 15 to 0.

PROPOSED CHANGES TO THE LAW

Existing law:

1) Establishes, within the Government Operations Agency, the CDT, and generally tasks the department with the approval and oversight of information technology (IT) projects, and with improving the governance and implementation of IT by standardizing reporting relationships, roles, and responsibilities for setting IT priorities. (Gov. Code § 11545, et seq.)

2) Expresses the intent of the Legislature that policies and procedures developed by CDT and Department of General Services (DGS) pertaining to the acquisition of IT goods and services provide for all of the following: the expeditious and value-effective acquisition of IT goods and services to satisfy state requirements; the acquisition of IT goods and services within a competitive framework; the delegation of authority by DGS to each state agency that has demonstrated to DGS's satisfaction the ability to conduct value-effective IT goods and services acquisitions; and the review and resolution of protests submitted by any bidders with respect to any IT goods and services acquisitions. (Pub. Con. Code § 12101.)

3) Provides that no person in the State of California shall, on the basis of sex, race, color, religion, ancestry, national origin, ethnic group identification, age, mental disability, physical disability, medical condition, genetic information, marital status, or sexual orientation, be unlawfully denied full and equal access to the benefits of, or be unlawfully subjected to discrimination under, any program or activity that is conducted, operated, or administered by the state or by any state agency, is funded directly by the state, or receives any financial assistance from the state. (Gov. Code § 11135 et. seq.)

4) Provides, pursuant to the Unruh Civil Rights Act, that all persons within the jurisdiction of this state are free and equal, and no matter what their sex, race, color, religion, ancestry, national origin, disability, medical condition, genetic information, marital status, sexual orientation, citizenship, primary language, or immigration status are entitled to the full and equal accommodations, advantages, facilities, privileges, or services in all business establishments of every kind whatsoever. (Civ. Code § 51.)

This bill:

1) Requires CDT, on or before September 1, 2024, to conduct, in coordination with other interagency bodies as it deems appropriate, a comprehensive inventory of all high-risk ADS that have been proposed for use, development, or procurement by, or are being used, developed, or procured by, any state agency.

2) Requires the comprehensive inventory to include a description of all of the following:
    a) any decision the ADS can make or support, the intended benefits of that use, and the alternatives to that use;
    b) the results of any research assessing the efficacy and relative benefits of the uses and alternatives of the ADS described above;
    c) the categories of data and personal information the ADS uses to make its decisions;
    d) the measures in place, if any, to mitigate the risks, including cybersecurity risk and the risk of inaccurate, unfairly discriminatory, or biased decisions, of the ADS, including performance metrics, cybersecurity controls, privacy controls, risk assessments or audits for potential risks, and measures or processes in place to contest an automated decision.

3) Requires CDT, on or before January 1, 2025, and annually thereafter, to submit a report, as specified, of the comprehensive inventory to the Assembly Committee on Privacy and Consumer Protection and the Senate Committee on Governmental Organization. This requirement expires on January 1, 2029.

4) Defines the following terms:
    a) "Automated decision system" means a computational process derived from machine learning, statistical modeling, data analytics, or artificial intelligence that issues simplified output, including a score, classification, or recommendation, that is used to assist or replace human discretionary decisionmaking and materially impacts natural persons. ADS does not include a spam email filter, firewall, antivirus software, identity and access management tools, calculator, database, dataset, or other compilation of data.
    b) "High-risk automated decision system" means an ADS that is used to assist or replace human discretionary decisions that have a legal or similarly significant effect, including decisions that materially impact access to, or approval for, housing or accommodations, education, employment, credit, health care, and criminal justice.
    c) "State agency" includes every state office, department, division, bureau, and the California State University. It does not include the University of California, the Legislature, the judicial branch, or any board, except as provided.

COMMENTS

1. Background

   a. *Algorithmic bias*

Owing to recent advances in processing power and the rise of big data, artificial intelligence's capacity and the scope of its applications have expanded rapidly, impacting how we communicate, interact, entertain ourselves, travel, transact, and consume media. In ways we may not fully comprehend, artificial intelligence empowers and encumbers us. It has been used to accelerate productivity, achieve efficiencies, liberate us from drudgery, help us understand and enjoy the world, connect with each other, and live longer, fuller lives. It has also been used to constrain personal autonomy, compromise privacy and security, foment social upheaval, exacerbate inequality, spread misinformation, and subvert democracy. For good or ill, its transformative potential seems boundless.

The rapid proliferation of algorithm-driven applications reflects advances in a subset of artificial intelligence known as "machine learning," a technique that "aims to help computers discover fuzzy rules by themselves, without having to be explicitly instructed every step of the way by human programmers."[1] Machine learning "enables computer systems to learn and make predictions based on historical data. The machine learning process is powered by a machine learning algorithm, a function that is able to improve its performance over time by training itself using methods of data analysis and analytical modelling."[2] The most prominent type of machine learning is "deep learning," which "uses artificial neural networks – simplified computer simulations of how biological neurons behave—to extract rules and patterns from sets of data."[3]

Algorithms process enormous datasets and make decisions with speed and reliability that vastly exceed human capabilities. "They determine everything from what ads we see on the Internet, to whether we are flagged for increased security screening at the airport, to our medical diagnoses and credit scores. They lie behind two of the most powerful products of the digital information age: Google Search and Facebook's Newsfeed."[4] The most sophisticated algorithms need no supervision and use deep neural networks to "discover hidden patterns in data, typically those unrecognizable to, or difficult to discern by, humans."[5] In addition to organizing vast troves of data,

---

[1] *How machine learning works* (May 14, 2015) The Economist, https://www.economist.com/the-economist-explains/2015/05/13/how-machine-learning-works. All internet citations are current as of June 21, 2023.
[2] *Intro to AI for Policymakers: Understanding the Shift* (March 2018) Brookfield Institute, http://brookfieldinstitute.ca/research-analysis/intro-to-ai-for-policymakers/.
[3] *How machine learning works, supra,* fn. 1.
[4] Jacob Weisberg, *The Digital Poorhouse* (June 7, 2018) The New York Review of Books, https://www.nybooks.com/articles/2018/06/07/algorithms-digital-poorhouse/.
[5] *AI for Policymakers*, *supra*, fn. 2 at p. 5.

algorithms offer the possibility of eliminating human biases in areas such as hiring decisions, credit scores, and criminal sentencing.

However, an algorithm is only as good as the information it is analyzing. Flawed inputs will produce flawed outputs. And an algorithm may key in on factors other than those intended by its designer. In one example, a software student was dismayed to learn his program that could reliably distinguish dogs from wolves had, in actuality, learned to recognize snow in the background of the pictures rather than the canine's features.[6] When we do not fully understand how an algorithm works, we are unable to determine which aspects of data it is focusing on.

And in many cases algorithms may inadvertently pick up human biases. In *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*, Cathy O'Neill shows that ultimately, a person shapes an algorithm, mediating the datasets gathered and deciding how to weigh them. These decisions are colored by our inherent biases and cultural predilections. Jacob Weisberg writes that "[c]orrelations reflected in historical data become invisibly entrenched in policy without programmers having ill intentions. Quantified information naturally points backward."[7] Rather than eliminating bias, some algorithms reinforce it, cloaking discrimination with mathematical neutrality.

ProPublica recently explored this phenomenon in the field of criminal justice.[8] Some jurisdictions factor algorithm-driven risk assessments into criminal bail, sentencing, and parole decisions. In 2014, U.S. Attorney General Eric Holder warned that the risk scores might be injecting bias into the courts. "Although these measures were crafted with the best of intentions, I am concerned that they inadvertently undermine our efforts to ensure individualized and equal justice," he stated, adding, "they may exacerbate unwarranted and unjust disparities that are already far too common in our criminal justice system and in our society." After obtaining the risk scores assigned to more than 7,000 people arrested in Broward County, Florida, ProPublica found that the scores "proved remarkably unreliable in forecasting violent crime." ProPublica's study validated Holder's fears: "Black defendants were still 77 percent more likely to be pegged as at higher risk of committing a future violent crime and 45 percent more likely to be predicted to commit a future crime of any kind."

A lack of transparency reduces accountability, again underscoring the importance of oversight. "[A]lgorithms simply grind out their results, and it is up to humans to review and address how that data is presented to users, to ensure the proper context and

---

[6] *Husky or Wolf? Using a Black Box Learning Model to Avoid Adoption Errors* (August 24, 2017) UCI Applied Innovation, http://innovation.uci.edu/2017/08/husky-or-wolf-using-a-black-box-learning-model-to-avoid-adoption-errors/.
[7] *The Digital Poorhouse*, *supra*, fn. 4.
[8] Angwin, et al., *Machine Bias* (2016) ProPublica, https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

application of that data."[9] New York University School of Law Professor Sarah Valentine puts a finer point on it: "Helpful as algorithms may be, they inevitably target marginalized populations and exacerbate the social stratification and vast inequality that already exists in our society."[10]

Illustrating this dynamic in the field of education, Meredith Broussard, a data journalism professor at New York University and author of "Artificial Unintelligence: How Computers Misunderstand the World," wrote an op-ed piece in the *New York Times* detailing how "the International Baccalaureate—a global program that awards prestigious diplomas to high school students—canceled its usual in-person exams because of the [COVID-19] pandemic" and instead "used an algorithm to 'predict' student grades based on an array of student information, including teacher-estimated grades and past performance by students in each school."[11] Tens of thousands of students, surprised to find out they failed, protested the results. "High-achieving, low-income students were hit particularly hard: many took the exams expecting to earn college credit with their scores and save thousands of dollars on tuition."[12]

Coining the term "technochauvinism"—the idea that technological solutions are almost always superior to ordinary human decisionmaking—Broussard writes:

> Computers are excellent at doing math, but education is not math — it's a social system. And algorithmic systems repeatedly fail at making social decisions. Algorithms can't monitor or detect hate speech, they can't replace social workers in public assistance programs, they can't predict crime, they can't determine which job applicants are more suited than others, they can't do effective facial recognition, and they can't grade essays or replace teachers.
>
> In the case of the International Baccalaureate program, grades could have been assigned based on the sample materials that students had already submitted by the time schools shut down. Instead, the organization decided to use an algorithm, which probably seemed like it would be cheaper and easier.
>
> The process worked like this: Data scientists took student information and fed it into a computer. The computer then constructed a model that outputted individual student grades, which International Baccalaureate

---

[9] Keith Kirkpatrick, *Battling Algorithmic Bias* (2016) Communications of the ACM Vol. 59, No. 10, pp. 16-17, https://cacm.acm.org/magazines/2016/10/207759-battlingalgorithmic-bias/abstract.
[10] *Artificial Intelligence and Predictive Algorithms: Why Big Data Can Lead to Big Problems* (2019) 46 Fordham Urb. L.J. 364, 365.
[11] *When Algorithms Give Real Students Imaginary Grades* (Sept. 8, 2020) New York Times https://www.nytimes.com/2020/09/08/opinion/international-baccalaureate-algorithm-grades.html.
[12] *Id.*

> claimed the students would have gotten if they had taken the standardized tests that didn't happen. It's a legitimate data science method, similar to the methods that predict which Netflix series you'll want to watch next or which deodorant you're likely to order from Amazon.
>
> The problem is, data science stinks at making predictions that are ethical or fair. In education, racial and class bias is baked into the system — and an algorithm will only amplify those biases.
>
> Crude generalizations work for Netflix predictions because the stakes are low. If the Netflix algorithm suggests a show and I don't like it, I ignore it and move on with my day. In education, the stakes are much higher.[13]

b. Examples of harmful uses of ADS in state government

Nationally, there have been several examples of state governments' use of ADS that have disproportionately harmed disadvantaged communities:

- Between 2013 and 2015, a Michigan ADS operating with minimal employee oversight wrongly accused 40,000 people of employment insurance fraud, many of whom were forced to pay heavy fines. Upon appeal, less than eight percent of those fraud charges were found to be legitimate.[14] The ADS cost the state $47 million and millions more as a result of lawsuits.
- In 2016, the state of Arkansas implemented an algorithm to assign access to Medicaid benefits. However, an estimated 19 percent of Medicaid beneficiaries had their benefits inappropriately cut, losing access to home care, nursing visits and medical treatments. In a lawsuit filed by Arkansas Legal Aid, the courts ultimately found those who were denied benefits could not effectively challenge the system, since there was no way of knowing what information factored into the algorithm's opaque decision-making process leading to that result. That case ultimately revealed the algorithm featured several design flaws, miscoding and incorrect calculations.[15]

---

[13] *Id.*

[14] Alejandro de la Garza, *States' Automated Systems Are Trapping Citizens in Bureaucratic Nightmares With Their Lives on the Line* (May 20, 2020) *Time Magazine* https://time.com/5840609/algorithm-unemployment/.

[15] Colin Lecher, *What happens when an algorithm cuts your healthcare* (Mar. 21, 2018) *The Verge* https://www.theverge.com/2018/3/21/17144260/healthcare-medicaid-algorithm-arkansas-cerebral-palsy.

- A market analysis algorithm used in Detroit to direct public housing subsidies, tax breaks and housing development redirected critical funding away from Detroit's poorest and predominantly black neighborhoods.[16]
- Since 2016, social workers in a Pennsylvania county have relied on an algorithm to help them determine which child welfare calls warrant further investigation. Now, the Justice Department is reportedly scrutinizing the controversial family-screening tool over concerns that using the algorithm may be violating the Americans with Disabilities Act by allegedly discriminating against families with disabilities, including families with mental health issues.[17]

In California, the CalWIN system, which provides a means for applying for CalFresh benefits, Medi-Cal and CalWORKS, included incorrectly translated policy in its code, which caused overpayments, underpayments, and improper terminations of public benefits, including the denial of Medicaid to foster children in contravention of federal law.[18] More recently, the ACLU found that the California Department of Public Health's COVID-19 vaccine distribution algorithm could "leave more than 2 million vulnerable Californians—many of them from Black and Latinx communities—without additional supply, despite the state's core goal of equity in vaccine distribution."[19]

An article by two professors from the University of California at Berkeley argues that government officials are increasingly purchasing ADS with insufficient knowledge of their design and operation, and how this aligns with public values.[20] The article states:

> At every level of government, officials contract for technical systems that employ machine learning—systems that perform tasks without using explicit instructions, relying on patterns and inference instead. These systems frequently displace discretion previously exercised by policymakers or individual front-end government employees with an opaque logic that bears no resemblance to the reasoning processes of agency personnel. However, because agencies acquire these systems

---

[16] Le, Vinhcent, *Algorithmic Bias Explained: How Automated Decision-Making Becomes Automated Discrimination* (Feb. 23, 2021) The Greenlining Institute, http://greenlining.org/publications/reports/2021/algorithmic-bias-explained/

[17] Ashley Belanger, AI tool used to spot child abuse allegedly targets parents with disabilities (Jan. 31, 2023) Ars Technica, https://arstechnica.com/tech-policy/2023/01/doj-probes-ai-tool-thats-allegedly-biased-against-families-with-disabilities/.

[18] Daniel Keats Citron, *Technological Due Process* (2007) 85 Wash. U. L. Rev. 1249, 1249 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1012360.

[19] *ACLU research suggests that California's vaccine distribution plan may leave more than 2 million vulnerable residents without additional supply* (May 6, 2021) ACLU Northern California, https://www.aclunc.org/news/aclu-research-suggests-california-s-vaccine-distribution-plan-may-leave-more-2-million.

[20] Diedre Mulligan and Kenneth Bamberger, *Procurement As Policy: Administrative Process for Machine Learning* (2019) 34 Berkeley Technology Law Journal 781 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3464203#page=4.

> through government procurement processes, they and the public have little input into—or even knowledge about—their design or how well that design aligns with public goals and values.
> […] When the adoption of these systems is governed by procurement, the policies they embed receive little or no agency or outside expertise beyond that provided by the vendor. Design decisions are left to private third-party developers. There is no public participation, no reasoned deliberation, and no factual record, which abdicates Government responsibility for policymaking.[21]

If the state agency is unable to understand or explain how an ADS that supplants a human decisionmaking process works, this raises due process concerns, as denials of rights or deprivations of property would be difficult to meaningfully challenge. (*See Mathews v. Eldridge* (1976) 424 U.S. 319, 333.) "Automated systems jeopardize due process norms. Their lack of meaningful notice, and a hearing officer's tendency to presume a computer system's infallibility, devalue hearings."[22]

c. Frameworks for addressing ADS

In 2018, New York City enacted the nation's first algorithmic accountability law, which regulates New York City agencies' use of algorithms by creating a task force to oversee the government's use of algorithms, examine how error and bias enter into their design, and recommend measures that ensure accuracy and fairness.

In 2019, the Canadian government adopted a Directive on Automated Decision-Making and an accompanying algorithmic impact assessment tool to guide the use of automated decision making at the federal level.[23] The Directive defines ADS as "any technology that either assists or replaces the judgement of human decision-makers." The Directive establishes impact assessment levels for ADS, based on the anticipated impact on the rights of individuals or communities, the health or well-being of individuals or communities, the economic interests of individuals, entities, or communities, and the ongoing sustainability of an ecosystem.

Article 22 of the European Union's General Data Protection Regulation provides that a "data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her." Exceptions are allowed when necessary for entering into or performing a contract or when the person has granted explicit consent, provided that "the data controller … implement[s] suitable measures to safeguard the data subject's rights and freedoms and legitimate interests, at least the

---

[21] *Id.* at 781.
[22] *Technological Due Process*, *supra*, fn. 17.
[23] *Directive on Automated Decision-Making*, Government of Canada website, https://www.tbs-sct.gc.ca/pol/doc-eng.aspx?id=32592.

right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision."

More recently the Biden Administration has published its Blueprint for an AI Bill of Rights, which is a set of five principles and associated practices to help guide the design, use, and deployment of AI to protect the rights of the American public:

- *Safe and Effective Systems*: You should be protected from unsafe or ineffective systems. Automated systems should be developed with consultation from diverse communities, stakeholders, and domain experts to identify concerns, risks, and potential impacts of the system.

- *Algorithmic Discrimination Protections*: Designers, developers, and deployers of automated systems should take proactive and continuous measures to protect individuals and communities from algorithmic discrimination and to use and design systems in an equitable way. This protection should include proactive equity assessments as part of the system design, use of representative data and protection against proxies for demographic features, ensuring accessibility for people with disabilities in design and development, pre-deployment and ongoing disparity testing and mitigation, and clear organizational oversight.

- *Data Privacy*: You should be protected from abusive data practices via built-in protections and you should have agency over how data about you is used. You should be protected from violations of privacy through design choices that ensure such protections are included by default, including ensuring that data collection conforms to reasonable expectations and that only data strictly necessary for the specific context is collected. Designers, developers, and deployers of automated systems should seek your permission and respect your decisions regarding collection, use, access, transfer, and deletion of your data in appropriate ways and to the greatest extent possible; where not possible, alternative privacy by design safeguards should be used. Systems should not employ user experience and design decisions that obfuscate user choice or burden users with defaults that are privacy invasive. Consent should only be used to justify collection of data in cases where it can be appropriately and meaningfully given. Any consent requests should be brief, be understandable in plain language, and give you agency over data collection and the specific context of use; current hard-to-understand notice-and-choice practices for broad uses of data should be changed. Enhanced protections and restrictions for data and inferences related to sensitive domains, including health, work, education, criminal justice, and finance, and for data pertaining to youth should put you first. In sensitive domains, your data and related inferences should only be used for necessary functions, and you should be protected by ethical review and use prohibitions. You and your communities should be free from unchecked surveillance; surveillance technologies should be subject to heightened oversight

that includes at least pre-deployment assessment of their potential harms and scope limits to protect privacy and civil liberties. Continuous surveillance and monitoring should not be used in education, work, housing, or in other contexts where the use of such surveillance technologies is likely to limit rights, opportunities, or access. Whenever possible, you should have access to reporting that confirms your data decisions have been respected and provides an assessment of the potential impact of surveillance technologies on your rights, opportunities, or access.

- *Notice and Explanation*: You should know that an automated system is being used and understand how and why it contributes to outcomes that impact you. Designers, developers, and deployers of automated systems should provide generally accessible plain language documentation including clear descriptions of the overall system functioning and the role automation plays, notice that such systems are in use, the individual or organization responsible for the system, and explanations of outcomes that are clear, timely, and accessible. Such notice should be kept up-to-date and people impacted by the system should be notified of significant use case or key functionality changes. You should know how and why an outcome impacting you was determined by an automated system, including when the automated system is not the sole input determining the outcome.

- *Human Alternatives, Consideration, and Fallback*: 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 focus on ensuring broad accessibility and protecting the public from especially harmful impacts.[24]

2. Taking inventory of high-risk ADS used by state agencies

The bill defines "automated decision system" as a computational process derived from machine learning, statistical modeling, data analytics, or artificial intelligence that issues simplified output, including a score, classification, or recommendation, that is used to assist or replace human discretionary decisionmaking and materially impacts natural persons. Under the bill, an ADS does not include a spam email filter, firewall, antivirus software, identity and access management tools, calculator, database, dataset, or other compilation of data.

The bill defines "high-risk" ADS as an ADS that is used to assist or replace human discretionary decisions that have a legal or similarly significant effect. This includes decisions that materially impact access to, or approval for, housing or accommodations,

[24] *Blueprint For An AI Bill Of Rights* (Oct. 2022) Office of Science and Technology Policy, https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf.

education, employment, credit, health care, and criminal justice. Examples of high-risk ADS would likely include those that determine a person's eligibility for public benefits and services, are used for state employment screening, or used by law enforcement authorities for profiling individuals or predicting the likelihood of recidivism.

The bill requires, on or before September 1, 2024, CDT to conduct a comprehensive inventory of all high-risk ADS that have been proposed for, or are being used, developed, or procured by, state agencies. CDT is to coordinate with other interagency bodies, as it deems appropriate.

The inventory must include a description of the decisions that the ADS can make or support and both the intended benefits and alternatives to that use. Detail must be provided about the categories of data and personal information used by the ADS. In order to meaningfully assess the ADS listed in the inventory, the bill requires a description of any research into its efficacy and benefits and of any measures in place to mitigate the risks. These risks include those discussed at length above – the risk of inaccurate, unfairly discriminatory, or biased decisions – as well as cybersecurity risks. Examples of this would be any performance metrics that gauge the accuracy of the system; cybersecurity and privacy controls; and risk assessments or audits for potential risks. The inventory must also highlight any measures or processes that are in place to contest an automated decision to better assess the due process being afforded affected individuals.

The CDT must submit a report of the comprehensive inventory to the Assembly Committee on Privacy and Consumer Protection and the Senate Committee on Governmental Organization by January 1, 2025, and annually thereafter until January 1, 2029.

3. Stakeholder positions

According to the author:

> California is leading the way in adopting Automated Decision Systems (ADS) across state agencies to modernize and deliver services more efficiently. State agencies are using ADS in various ways, including to detect fraud in unemployment and tax filings, speed up document processing at the Department of Motor Vehicles, and help make better decisions in welfare services and healthcare reimbursements and California's climate investments.
>
> When used properly, these systems can benefit Californians. However, if these systems are not designed and implemented correctly, they can create unfairly biased or inaccurate results that harm Californians and reduce trust in these systems. These results can disproportionately harm low-

> income families and communities of color given the number of government services and programs that impact them.
>
> AB 302 would ensure that Californians have transparency into the government's use of high-risk ADS and provide state agencies with the information to analyze their use of high-risk ADS. By requiring the CDT to establish guidelines identifying ADS that have a high-risk of adverse impacts and conduct an inventory of those high-risk ADS, this bill will help state agencies identify and minimize the risk of adverse and discriminatory impacts that result from their design and implementation of ADS.

The Greenlining Institute, the sponsor of the bill, writes:

> California is taking the lead in modernizing government systems, utilizing ADS to identify fraud, streamline document processing, and to make decisions for public benefits programs. However, poorly designed automated systems can result in biased, unfair or inaccurate results that disproportionately impact low-income families, and communities of color. For example, the IRS' tax auditing algorithm is three times more likely to flag Black Americans for audits compared to other taxpayers, despite having no evidence that these taxpayers are at a higher risk of tax evasion. In Michigan and Arkansas, poorly designed algorithms wrongly denied tens of thousands of low-income families access to unemployment benefits and Medicaid. AB 302 grants California the opportunity to ensure the adoption of government ADS is done in a way that is transparent and promotes proper accountability and oversight over the use of these systems.

A coalition of advocacy groups, including Oakland Privacy and Privacy Rights Clearinghouse, writes in support:

> As California works to improve and modernize its government functions through the use of big data and automated systems, it must also lead in transparency around where these systems are used, what decisions they can support and how each agency is managing the risks associated with the use of these systems.
>
> AB 302 will authorize the California Department of Technology to conduct a comprehensive inventory of all high-risk Automated Decision Systems used by state agencies. The inventory will also identify the decisions these systems make or support, the potential benefits of these systems and the measures taken to reduce the risk of inaccurate, unfairly discriminatory, or biased automated decisions. We therefore support AB 302 to provide

much needed information to Californians and state agencies around the use of ADS.

SUPPORT

Greenlining Institute (sponsor)
California Civil Liberties Advocacy
Electronic Frontier Foundation
Electronic Privacy Information Center
Equality California
Media Alliance
Oakland Privacy
Privacy Rights Clearinghouse
Secure Justice
TechEquity Collaborative

OPPOSITION

None known

RELATED LEGISLATION

Pending Legislation:

SCR 17 (Dodd, 2023) affirms the California Legislature's commitment to President Biden's vision for a safe AI and the principles outlined in the "Blueprint for an AI Bill of Rights" and expresses the Legislature's commitment to examining and implementing those principles in its legislation and policies related to the use and deployment of automated systems. SCR 17 is currently in the Assembly Privacy and Consumer Protection Committee.

SB 313 (Dodd, 2023) establishes the Office of Artificial Intelligence. It requires state agencies to disclose when they are using generative artificial intelligence to communicate with a person and to provide them an option to speak with a natural person at the agency. SB 313 was held on suspense in the Senate Appropriations Committee.

SB 398 (Wahab, 2023) establishes the Artificial Intelligence for California Research Act, which requires CDT to develop and implement a comprehensive research plan to study the feasibility of using advanced technology to improve state and local government services. SB 398 is currently in the Senate Governmental Organization Committee.

SB 721 (Becker, 2023) creates, until January 1, 2030, the California Interagency AI Working Group made up of members with varied expertise to deliver a report to the

Legislature regarding AI. The group is directed to specified tasks, including recommending a definition of AI and determining the relevant agencies to develop and oversee AI policy and implementation of that policy. SB 721 is currently in the Assembly Privacy and Consumer Protection Committee.

AB 331 (Bauer-Kahan, 2023) prohibits "algorithmic discrimination," that is, use of an automated decision tool to contribute to unjustified differential treatment or outcomes that may have a significant effect on a person's life. It requires any deployer of an automated decision tool to perform an impact assessment for those tools and to notify any natural person that is the subject of the consequential decision that an automated decision tool is being used to make, or be a controlling factor in making, the consequential decision. AB 331 was held on suspense in the Assembly Appropriations Committee.

Prior Legislation:

AB 13 (Chau, 2021) would have established the Automated Decision Systems Accountability Act, which, in the context of the State's procurement policies, promotes oversight over ADS that pose a high risk of adverse impacts on individual rights. The bill was eventually gutted and amended to address a different topic.

AB 858 (Jones-Sawyer, 2021) would have provided that the use of technology—defined to include algorithms derived from the use of health care-related data—shall not limit a worker who is providing direct patient care from exercising independent clinical judgment in the assessment, evaluation, planning, and implementation of care, nor from acting as a patient advocate. AB 858 was vetoed by Governor Newsom at the request of the author and sponsor.

AB 976 (Chau, 2020) would have established the AI in State Government Services Commission to gather input on how AI and data science could be used to improve state services. The bill was held on the Senate Appropriations Committee suspense file.

AB 2269 (Chau, 2020) the Automated Decision Systems Accountability Act of 2020, among other things, would have required a business in California that provides a program or device that uses an ADS to take affirmative steps to ensure that there are processes in place to continually test for biases, as specified; and, would have established an ADS Advisory Task Force, as specified. The bill died in the Assembly Privacy and Consumer Protection Committee.

ACR 125 (Jones-Sawyer, 2020) would have urged policymakers in both federal and state government to explore ways to promote the development and use of new technologies to reduce bias and discrimination in hiring and employment. The measure died in this Committee due to the limits placed on the Legislature because of the COVID-19 pandemic.

SB 348 (Chang, 2019) would have required the Director of CDT to develop a strategic plan to aid departments and agencies with incorporating AI into state IT operations, as specified. The bill died in the Assembly Appropriations Committee.

SB 444 (Umberg, 2019) would have requested the Regents of the University of California (UC) to enact a resolution authorizing the law schools at UC Berkeley and UC Irvine to participate in a pilot project to develop AI or machine-learning solutions to address access to justice issues faced by self-representing litigants in their respective courts. The bill died in the Assembly Higher Education Committee.

AB 1576 (Calderon, 2019) would have required the Secretary of Government Operations to appoint participants to an AI working group to evaluate the uses, risks, benefits, and legal implications associated with the development and deployment of AI by California-based businesses. The bill was held on the Senate Appropriations Committee suspense file.

SJR 6 (Chang, Res. Ch. 112, Stats. 2019) urged the President and the Congress of the United States to develop a comprehensive AI Advisory Committee and to adopt a comprehensive AI policy.

AB 594 (Salas, 2019) would have authorized the Director CDT to designate a position within the department to evaluate the uses of AI in state government and to advise the Director of Technology on incorporating AI into state IT strategic plans, policies, standards and enterprise architecture, and would have required CDT to adopt guidelines by January 1, 2021, to govern the use and implementation of AI technologies in state government functions, as specified. The bill was vetoed by Governor Newsom.

ACR 215 (Kiley, Resolution Ch. 206, Stats. 2018) expressed the Legislature's support for a set of principles for the governance of AI known as the 23 Asilomar AI Principles.

PRIOR VOTES:

Senate Governmental Organization Committee (Ayes 15, Noes 0)
Assembly Floor (Ayes 79, Noes 0)
Assembly Appropriations Committee (Ayes 11, Noes 0)
Assembly Privacy and Consumer Protection Committee (Ayes 11, Noes 0)
****