

Date of Hearing: June 27, 2023

ASSEMBLY COMMITTEE ON PRIVACY AND CONSUMER PROTECTION

Jesse Gabriel, Chair

SCR 17 (Dodd) – As Introduced January 30, 2023

SENATE VOTE: 33-0

SUBJECT: Artificial intelligence

SYNOPSIS

The provisions of this resolution are based upon the principles outlined in the Blueprint for an AI Bill of Rights, a white paper issued last October by the White House Office of Science and Technology Policy. The Blueprint identifies five principles that should “guide the design, use, and deployment of automated systems to protect the American public in the age of artificial intelligence.” These principles are Safe and Effective Systems; Algorithmic Discrimination Protections; Data Privacy; Notice and Explanation; and Human Alternatives, Consideration, and Fallback.

Despite the use of the term “AI” in its title, the Blueprint focuses on one application of artificial intelligence, specifically, “automated systems”: any system, software, or process that (i) uses computation as whole or part of a system to determine outcomes, make or aid decisions, inform policy implementation, collect data or observations, or otherwise interact with individuals and/or communities, and (ii) have the potential to meaningfully impact the American public’s rights, opportunities, or access to critical resources or services.

The Blueprint is a set of principles meant to inform both public- and private-sector decisionmaking when addressing automated systems; it neither provides specific policies to implement, nor does it prescribe statutory text. Each of the five principles is summarized and discussed in the analysis, together with real-world examples illustrating why policy guided by these principles is so important.

Despite the fact that the Blueprint was issued by the federal government, adoption of this resolution will in no way increase the possibility that California’s regulation of automated systems will be preempted by federal law.

This resolution was generated by ChatGPT in January 2023. The analysis concludes with a look at how large language models, which have advanced significantly in the intervening five months, would generate a resolution based on the Blueprint if asked to do so today.

SUMMARY: Commits the Legislature to examining and implementing the principles outlined in the *Blueprint for an AI Bill of Rights*, published in October 2022 by the White House Office of Science and Technology Policy. Specifically, **this bill:**

- 1) Finds that the use of technology, data, and automated systems poses significant challenges to democracy and the rights of the public, as evidenced by incidents of unsafe, ineffective, or biased systems in health care, discriminatory algorithms in hiring and credit decisions, and unchecked data collection that threatens privacy and opportunities.

- 2) Finds that automated systems also have the potential to bring about extraordinary benefits, including increasing efficiency in agriculture and revolutionizing industries through data analysis.
- 3) Finds that President Joseph R. Biden has affirmed civil rights and democratic values as a cornerstone of his administration and has ordered the federal government to work toward rooting out inequity and advancing civil rights, equal opportunity, and racial justice.
- 4) Finds that the White House Office of Science and Technology Policy has developed the *Blueprint for an AI Bill of Rights*, a set of five principles to guide the design, use, and deployment of automated systems in a manner that protects the rights of the public while leveraging the benefits of AI.
- 5) Resolves that the California Legislature affirms its commitment to President Biden’s vision for safe AI and the principles outlined in the *Blueprint for an AI Bill of Rights*, including: Safe and Effective Systems, Algorithmic Discrimination Protections, Data Privacy, Notice and Explanation, and Human Alternatives, Consideration, and Fallback.
- 6) Resolves, further, that the California Legislature commits to examining and implementing these principles in its legislation and policies related to the use and deployment of automated systems in the State of California.
- 7) Resolves that the Secretary of the Senate transmit copies of this resolution to the author for appropriate distribution

COMMENTS:

1) **Background.** The provisions of this resolution are based upon principles outlined in the *Blueprint for an AI Bill of Rights*, a white paper issued last October by the White House Office of Science and Technology Policy. The full 73-page *Blueprint* is available to read at <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>. A summary is available at <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>.

The *Blueprint* identifies five principles that, as it explains, should “guide the design, use, and deployment of automated systems to protect the American public in the age of artificial intelligence.” (*Blueprint*, p. 3.) These principles are Safe and Effective Systems; Algorithmic Discrimination Protections; Data Privacy; Notice and Explanation; and Human Alternatives, Consideration, and Fallback. Both the term “automated systems” and these five principles are discussed below.

At the outset, it is important to realize the *Blueprint* is a set of principles meant to inform both public- and private-sector decisionmaking when addressing automated systems. The *Blueprint* neither provides specific policies to implement, nor does it prescribe statutory text. Rather:

The *Blueprint for an AI Bill of Rights* is [meant] to help guide the design, use, and deployment of automated systems to protect the rights of the American public in the age of artificial intelligence. Developed through extensive consultation with the American public, these principles are a blueprint for building and deploying automated systems that are aligned with democratic values and protect civil rights, civil liberties, and privacy. The *Blueprint* [includes] the five principles, notes on Applying the [*Blueprint*], and a Technical Companion

that gives concrete steps that can be taken by many kinds of organizations—from governments at all levels to companies of all sizes—to uphold these values. [...]his framework provides a national values statement and toolkit that is sector-agnostic to inform building these protections into policy, practice, or the technological design process. Where existing law or policy—such as sector-specific privacy laws and oversight requirements—do not already provide guidance, the *Blueprint*...should be used to inform policy decisions. (*Blueprint*, p. 4.)

This resolution would affirm the Legislature’s commitment to President Biden’s vision for safe AI and the principles outlined in the *Blueprint*. It would also commit the Legislature to examining and implementing these principles in its legislation and policies related to the use and deployment of automated systems in the State of California.

2) **Author’s statement.** According to the author:

Artificial intelligence has the potential to transform every aspect of our society and improve our lives in countless ways, from enabling personalized healthcare to enhancing the safety and efficiency of our transportation systems. But as with any new technology, it also presents new challenges and risks, particularly with respect to data privacy, cybersecurity, and bias. This resolution acknowledges these risks and affirms our commitment to addressing them in a thoughtful and deliberate manner. By doing so, we can help ensure that the benefits of AI are realized by all, and that we build a future together that is not only smart, but also safer and more just.

3) **What is an automated system, and how does it differ from artificial intelligence more generally?** Wikipedia defines “artificial intelligence” as “intelligence—perceiving, synthesizing, and inferring information—demonstrated by machines, as opposed to intelligence displayed by humans or by other animals.” (Wikipedia, *Artificial intelligence*, available at https://en.wikipedia.org/wiki/Artificial_intelligence.) “Artificial intelligence” is an umbrella term that encompasses many different technologies, such as computer vision, expert systems, genetic algorithms, knowledge representation, machine learning, neural networks, and robotics. (The interested reader can look up each of these topics on Wikipedia.) Many of these technologies date as far back as the 1950s: GPT-4, the subject of much news coverage over the past three months, is at its core, a neural network; the first neural network, the Perceptron, was constructed in 1958.

This resolution and the *Blueprint* are focused on a specific application of artificial intelligence, known as automated systems. The *Blueprint* defines “automated system” to mean:

[A]ny system, software, or process that uses computation as whole or part of a system to determine outcomes, make or aid decisions, inform policy implementation, collect data or observations, or otherwise interact with individuals and/or communities. Automated systems include, but are not limited to, systems derived from machine learning, statistics, or other data processing or artificial intelligence techniques.... (*Blueprint*, p. 10 [emphasis added].)

However, the *Blueprint* does not apply to all automated systems; only those that “**have the potential to meaningfully impact the American public’s rights, opportunities, or access to critical resources or services.**” (*Blueprint*, p. 8 [emphasis added].)

So, for example, while Netflix's recommendation algorithm is an automated system (since it collects data about what viewers watch, and aids them in deciding what to watch next), it would not qualify as the sort of automated system covered by this resolution (since it does not meaningfully impact viewers' rights, opportunities, or access to critical resources or services).

A chatbot like ChatGPT could be deemed an automated system (since it collects data and interacts with individuals) if it were assisting a doctor in making diagnoses or recommending treatment (since it then has the potential to meaningfully impact individuals' opportunities to heal or their access to medical treatment). But if it were being used to help write a press release or a homework essay, then it would not be considered an automated system under the *Blueprint's* framework.

4) What are the five principles regarding automated systems that the resolution endorses?

This resolution endorses the adoption of the five principles set forth in the *Blueprint*.

- Principle #1: Safe and Effective Systems

The *Blueprint's* summary of this principle emphasizes two components. The first component is public input:

Automated systems should be developed with consultation from diverse communities, stakeholders, and domain experts to identify concerns, risks, and potential impacts of the system. [...] Automated systems should not be designed with an intent or reasonably foreseeable possibility of endangering your safety or the safety of your community. (*Blueprint*, p. 15.)

The second component involves ongoing testing for safety, accompanied by risk mitigation in response to that testing:

[Automated s]ystems should undergo pre-deployment testing, risk identification and mitigation, and ongoing monitoring that demonstrate they are safe and effective based on their intended use, mitigation of unsafe outcomes including those beyond the intended use, and adherence to domain-specific standards. Outcomes of these protective measures should include the possibility of not deploying the system or removing a system from use. [...] Independent evaluation and reporting that confirms that the system is safe and effective, including reporting of steps taken to mitigate potential harms, should be performed and the results made public whenever possible. (*Ibid.*)

The *Blueprint* provides a number of real-world examples of automated systems that have been deployed and found to be unsafe and/or ineffective.

An unsafe system: "A proprietary model was developed to predict the likelihood of sepsis in hospitalized patients and was implemented at hundreds of hospitals around the country. An independent study showed that the model predictions underperformed relative to the designer's claims while also causing 'alert fatigue' by falsely alerting likelihood of sepsis." (*Blueprint*, p. 16.)

An ineffective system: "A company installed AI-powered cameras in its delivery vans in order to evaluate the road safety habits of its drivers, but the system incorrectly penalized drivers when

other cars cut them off or when other events beyond their control took place on the road. As a result, drivers were incorrectly ineligible to receive a bonus.” (*Blueprint*, p. 17.)

- Principle #2: Algorithmic Discrimination Protections

The *Blueprint* describes algorithmic discrimination as occurring when “automated systems contribute to unjustified different treatment or impacts disfavoring people based on their race, color, ethnicity, sex (including pregnancy, childbirth, and related medical conditions, gender identity, intersex status, and sexual orientation), religion, age, national origin, disability, veteran status, genetic information, or any other classification protected by law.” (*Blueprint*, p. 23.)

The *Blueprint* stresses the need to act proactively, rather than reactively, to prevent algorithmic discrimination:

Those responsible for the development, use, or oversight of automated systems should conduct proactive equity assessments in the design phase ... or during its acquisition to review potential input data, associated historical context, accessibility for people with disabilities, and societal goals to identify potential discrimination and effects on equity resulting from the introduction of the technology. The assessed groups should be as inclusive as possible of the underserved communities mentioned in the equity definition: Black, Latino, and Indigenous and Native American persons, Asian Americans and Pacific Islanders and other persons of color; members of religious minorities; women, girls, and non-binary people; lesbian, gay, bisexual, transgender, queer, and intersex (LGBTQI+) persons; older adults; persons with disabilities; persons who live in rural areas; and persons otherwise adversely affected by persistent poverty or inequality. Assessment could include both qualitative and quantitative evaluations of the system. (*Blueprint*, p. 26.)

As with the first principle (Safe and Effective Systems), the *Blueprint* stresses the importance of ongoing testing for algorithmic discrimination, coupled with mitigation measures:

Automated systems should be tested using a broad set of measures to assess whether the system components, both in pre-deployment testing and in-context deployment, produce disparities. [...] For every instance where the deployed automated system leads to different treatment or impacts disfavoring the identified groups, the entity governing, implementing, or using the system should document the disparity and a justification for any continued use of the system. [...] Disparities that have the potential to lead to algorithmic discrimination, cause meaningful harm, or violate equity goals should be mitigated. When designing and evaluating an automated system, steps should be taken to evaluate multiple models and select the one that has the least adverse impact, modify data input choices, or otherwise identify a system with fewer disparities. If adequate mitigation of the disparity is not possible, then the use of the automated system should be reconsidered. [...] Automated systems should be regularly monitored to assess algorithmic discrimination that might arise from unforeseen interactions of the system with inequities not accounted for during the pre-deployment testing, changes to the system after deployment, or changes to the context of use or associated data. [...] Riskier and higher-impact systems should be monitored and assessed more frequently. (*Blueprint*, p. 27)

The *Blueprint* provides a number of real-world examples of harmful discrimination that has resulted from the use of automated systems. For example: “An algorithm designed to identify patients with high needs for healthcare systematically assigned lower scores (indicating that they

were not as high need) to Black patients than to those of white patients, even when those patients had similar numbers of chronic conditions and other markers of health.” (*Blueprint*, p. 25.)

- Principle #3: Data Privacy

The *Blueprint* identifies data privacy as a “foundational and cross-cutting principle required for achieving all others in this framework.” (*Blueprint*, p. 31.) As will be seen, California is ahead of most states in seeking to ensure residents’ data privacy, though this principle could stand to be more comprehensively incorporated throughout state policy.

Key concepts underlying this principle are as follows:

- “You should be protected from violations of privacy through design choices that ensure such protections are included by default, including ... that only data strictly necessary for the specific context is collected.” (*Blueprint*, p. 30.)

This concept is known as “data minimization,” and is incorporated in laws such as the California Consumer Privacy Act (CCPA). (*See* Civ. Code § 1798.100(c).)

- “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.” (*Blueprint*, p. 30.)

These concepts are also embodied in the CCPA. (*See* Civ. Code §§ 1798.105, 1798.115, 1798.120, and 1798.121.)

- “Systems should not employ user experience and design decisions that obfuscate user choice or burden users with defaults that are privacy invasive.” (*Blueprint*, p. 30.)

A growing number of California laws, such as the Genetic Information Privacy Act, incorporate this concept through restrictions on the use of “dark patterns” to obtain user consent. (*See* Civ. Code § 56.18.)

- “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.” (*Blueprint*, p. 30.)

AB 1215 (Ting, Chap. 579, Stats. 2019) put in place a three-year moratorium on the use of biometric surveillance systems in connection with police-worn body cameras. If enacted, AB 1034 (Wilson, 2023) would reinstate that lapsed moratorium for another three years.

Despite California’s leadership in safeguarding residents’ data privacy, the laws cited above do not specifically apply to the use of automated systems. This is likely to be an area of future legislative action. However, at the agency level, the Civil Rights Council (which issues regulations under the Fair Employment and Housing Act) is engaged in rulemaking regarding the use of automated decision systems in employment. Similarly, the California Privacy Protection Agency is currently engaged in rulemaking regarding “access and opt-out rights with respect to businesses’ use of automated decisionmaking technology,” as directed by Proposition 24 (2020).

Ultimately, as noted in the *Blueprint*, “While a patchwork of laws exists to guide the collection and use of personal data in specific contexts, including health, employment, education, and credit, it can be unclear how these laws apply in other contexts and in an increasingly automated society. Additional protections would assure the American public that the automated systems they use are not...surveilling them without [requiring] context-specific consent or legal authority.” (*Blueprint*, p. 31.)

- Principle #4: Notice and Explanation

The *Blueprint* summarizes this principle as meaning that “you should know that an automated system is being used and understand how and why it contributes to outcomes that impact you.” (*Blueprint*, p. 40.)

The first component of this principle is notice:

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. (*Ibid.*)

The second component is explanation:

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. Automated systems should provide explanations that are technically valid, meaningful and useful to you and to any operators or others who need to understand the system, and calibrated to the level of risk based on the context. (*Ibid.*)

Ultimately, “[T]he American public deserve to know consistently and across sectors if an automated system is being used in a way that impacts their rights, opportunities, or access. This knowledge should provide confidence in how the public is being treated, and trust in the validity and reasonable use of automated systems.” (*Blueprint*, p. 41.)

The *Blueprint* provides a number of real-world examples of automated systems that fail to provide adequate notice and explanation, including the following: “A formal child welfare investigation is opened against a parent based on an algorithm and without the parent ever being notified that data was being collected and used as part of an algorithmic child maltreatment risk assessment. The lack of notice or an explanation makes it harder for those performing child maltreatment assessments to validate the risk assessment and denies parents knowledge that could help them contest a decision.” (*Ibid.*)

- Principle #5: Human Alternatives, Consideration, and Fallback

The *Blueprint* summarizes this principle as meaning that “[y]ou should be able to opt out, where appropriate, and have access to a person who can quickly consider and remedy problems you encounter.” (*Blueprint*, p. 46.)

In many ways, this may be the most controversial of the five principles. A major reason to deploy an automated system is efficiency: the system likely costs less than a human would to do the same task; it never tires; and it follows its programming exactly. Requiring a human option is decidedly inefficient, because humans are (relatively) expensive, need rest, and behave differently at different times and in different situations. It is therefore important to understand the motivation behind this principle:

There are many reasons people may prefer not to use an automated system: the system can be flawed and can lead to unintended outcomes; it may reinforce bias or be inaccessible; it may simply be inconvenient or unavailable.... Yet members of the public are often presented with no alternative, or are forced to endure a cumbersome process to reach a human decision-maker once they decide they no longer want to deal exclusively with the automated system or be impacted by its results. As a result of this lack of human reconsideration, many receive delayed access, or lose access, to rights, opportunities, benefits, and critical services. (*Blueprint*, p. 47.)

The *Blueprint* provides a number of real-world examples of the harms that ensue when automated systems fail to provide human alternatives, e.g., “A patient was wrongly denied access to pain medication when the hospital’s software confused her medication history with that of her dog’s. Even after she tracked down an explanation for the problem, doctors were afraid to override the system, and she was forced to go without pain relief due to the system’s error.” (*Blueprint*, p. 48.)

5) Adopting this resolution will not implicitly support federal preemption of California law. California has done nation-leading work on issues related to the regulation of the harms of technology, as can be seen in the discussion of the Data Privacy principle above. While the *Blueprint* is a product of the federal government, adoption of this resolution will not, either implicitly or explicitly, endorse federal preemption of the state’s efforts to address automated systems.

The *Blueprint* specifically includes the following legal disclaimer:

The *Blueprint for an AI Bill of Rights* is non-binding and does not constitute U.S. government policy. It does not supersede, modify, or direct an interpretation of any existing statute, regulation, policy, or international instrument. It does not constitute binding guidance for the public or Federal agencies and therefore does not require compliance with the principles described herein. It also is not determinative of what the U.S. government’s position will be in any international negotiation. Adoption of these principles may not meet the requirements of existing statutes, regulations, policies, or international instruments, or the requirements of the Federal agencies that enforce them. These principles are not intended to, and do not, prohibit or limit any lawful activity of a government agency, including law enforcement, national security, or intelligence activities. (*Blueprint*, p. 2.)

While this paragraph largely refers to the federal government, the fact that the *Blueprint* “does not constitute U.S. government policy [and] does not require compliance with the principles described” means that these principles cannot have any preemptive effect on state policy.

Adoption of this resolution, therefore, should in no way legally limit California’s ability to address automated systems if there is a conflict between state and federal law on this topic in the future.

6) **Exploring the world of large language model-based chatbots.** According to the author's office, the text of this resolution was generated in January 2023 by providing ChatGPT with the following prompt: "Draft a California Senate Concurrent Resolution reaffirming California's commitment to President Biden's vision and the five principles that the White House Office of Science and Technology Policy has identified that should guide the design, use, and deployment of automated systems to protect the American public in the age of artificial intelligence, as outlined in the following White House Report:" and then providing ChatGPT with the text of the White House's web page summarizing the *Blueprint*, which is available at <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>.

In the intervening five months, there has been significant development in the world of large language model (LLM)-based chatbots. ChatGPT now runs on top of the next-generation GPT-4 LLM (whereas previously it ran on the older GPT-3.5 model), and Google has released a competing LLM-based chatbot, named Bard. As it may be of historical interest, here is the text of the resolution generated by the current version of ChatGPT:

Resolved by the Senate of the State of California (the Assembly concurring), that

SECTION 1. The Legislature hereby affirms its commitment to the vision for a safe artificial intelligence (AI) environment put forth by President Biden, as articulated in the White House Report titled "Blueprint for an AI Bill of Rights".

SECTION 2. The Legislature emphasizes the importance of safeguarding the rights of Californians in the age of AI, recognizing the necessity of preventing harmful uses of automated systems, and promoting their benefits, in line with the principles laid out in the aforementioned report.

SECTION 3. The Legislature pledges to guide its policy decisions and actions by the five key principles identified in the "Blueprint for an AI Bill of Rights":

1. ****Safe and Effective Systems****: Californians should be protected from unsafe or ineffective systems and the systems should be developed and deployed following rigorous pre-deployment testing, risk identification, and ongoing monitoring.
2. ****Algorithmic Discrimination Protections****: Californians must not face discrimination from automated systems. Measures should be proactively taken to protect individuals from algorithmic discrimination, ensuring fair and equitable design and usage.
3. ****Data Privacy****: Californians deserve robust protections against abusive data practices. Respect for personal data decisions, minimization of data collection, and strict adherence to privacy-by-design practices should be the norm.
4. ****Notice and Explanation****: Californians have the right to know when an automated system is being used, and they deserve clear and accessible explanations about the impacts of such systems on their lives.
5. ****Human Alternatives, Consideration, and Fallback****: Californians should have the option to opt out from automated systems where appropriate and have access to human consideration and remedy when automated systems fail or produce an error.

SECTION 4. The Legislature commits to incorporate these principles into policy and practice, thereby encouraging private sector companies, local governments, and other organizations in California to do the same.

SECTION 5. The Legislature calls for the Blueprint for an AI Bill of Rights to be used as a guiding framework for future decisions related to AI policy in California, particularly where existing laws or policies do not provide explicit guidance.

SECTION 6. The Legislature reaffirms its commitment to ensuring that the advancement of AI technology in California aligns with democratic values and protects civil rights, civil liberties, and privacy, in line with President Biden's vision.

SECTION 7. The Secretary of the Senate shall transmit copies of this resolution to the author for appropriate distribution.

And here is the text of the resolution generated by the current version of Google Bard:

Resolving to affirm the Legislature's commitment to President Biden's vision for a safe AI and the five principles that the White House Office of Science and Technology Policy has identified that should guide the design, use, and deployment of automated systems to protect the American public in the age of artificial intelligence, as outlined in the White House Report titled "Blueprint for an AI Bill of Rights" available at <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>.

Whereas, artificial intelligence (AI) is a rapidly developing technology with the potential to revolutionize many aspects of our lives; and

Whereas, the potential benefits of AI are great, but there are also risks associated with this technology, such as bias, discrimination, and job displacement; and

Whereas, President Biden has called for a "safe AI" that is developed and used in a responsible and ethical manner; and

Whereas, the White House Office of Science and Technology Policy has identified five principles that should guide the design, use, and deployment of automated systems to protect the American public in the age of AI:

- Human-centered design: AI systems should be designed with the needs and interests of people at the forefront.
- Transparency: People should be able to understand how AI systems work and make informed decisions about how to interact with them.
- Accountability: There should be mechanisms in place to hold those responsible for developing and using AI systems accountable for their actions.
- Fairness: AI systems should not discriminate against individuals or groups.
- Safety: AI systems should be designed to be safe and secure.

Now, therefore, be it resolved by the Senate of the State of California,

- That the Legislature affirms its commitment to President Biden’s vision for a safe AI.
- That the Legislature supports the five principles identified by the White House Office of Science and Technology Policy for the design, use, and deployment of automated systems to protect the American public in the age of AI.
- That the Legislature urges the state’s policymakers, businesses, and educational institutions to work together to ensure that AI is developed and used in a responsible and ethical manner.

7) **Related legislation.** AB 302 (Ward) would require the California Department of Technology (CDT) to conduct an inventory of all high-risk automated decision systems being used in state agencies. Status: Senate Governmental Organization Committee.

AB 331 (Bauer-Kahan) would draw on the principles in the *Blueprint* to establish a statutory framework for furthering the safe and informed development and use of automated decision tools by both public and private entities in California. Status: Held, Assembly Appropriations Committee.

SB 313 (Dodd) would establish an Office of Artificial Intelligence within CDT, with “the powers and authorities necessary to guide the design, use, or deployment of automated systems by a state agency to ensure that all AI systems are designed and deployed in a manner that is consistent with state and federal laws and regulations regarding privacy and civil liberties and that minimizes bias and promotes equitable outcomes for all Californians.” Status: Held, Senate Appropriations Committee.

SB 721 (Becker) would create the California Interagency AI Working Group, which would deliver a report to the Legislature regarding artificial intelligence, include a recommendation for a statutory definition of “artificial intelligence” for use in legislation. Status: Two-year bill, Assembly Privacy and Consumer Protection Committee.

AB 13 (Chau, 2021) would have established the Automated Decision Systems (ADS) Accountability Act of 2021. The Act would have required state agencies seeking to procure ADS for high-risk applications to consider, among other things, steps taken by a prospective contractor to identify and mitigate potential disparate impacts that could result from use of that ADS; required a prospective contractor for an ADS for a high-risk application to submit an ADS impact assessment containing specified information about the ADS; and required the contracting agency to submit to CDT a high-risk ADS accountability report containing specified information regarding their proposed use of the ADS. Status: Gut-and-amend for another purpose.

ARGUMENTS IN SUPPORT: The Santa Monica Democratic Club commends this resolution as follows:

We are concerned about the burgeoning deployment of fast-developing artificial intelligence (AI) technologies and believe that the need to assess the various risks they pose to the general public and to devise appropriate guardrails is urgent. While specific safeguards governing the use of particular applications of AI—such as automated decision tools that can have consequential effects on individuals and communities and synthetic media that have enormous potential for destabilizing and weakening the democratic process—are needed, SCR 17 is a good first step in acknowledging the risks of these fast-developing technologies.

REGISTERED SUPPORT / OPPOSITION:

Support

Santa Monica Democratic Club

Opposition

None on file.

Analysis Prepared by: Jith Meganathan / P. & C.P. / (916) 319-2200