



Artificial Intelligence: Policy and Practice

November 2024

JCOTS Membership

Delegate C.E. Cliff Hayes, Jr., Chair
Delegate Bonita G. Anthony
Delegate Mike A. Cherry
Delegate Michelle Lopes Maldonado*
Delegate David A. Reid
Delegate Anne Ferrell H. Tata*
Delegate Michael J. Webert

Senator Adam P. Ebbin, Vice-Chair*
Senator Lashrecse D. Aird*
Senator Christie New Craig
Senator Barbara A. Favola
Senator Ghazala F. Hashmi

*Indicates membership in the Artificial Intelligence Subcommittee

Executive Director

Jodi Kuhn

Report Author

Gates Palissery, PhD Candidate, Translational Biology, Medicine, and Health Program,
Virginia Tech
COVES Fellow for Senator Aird (Summer 2024) and JCOTS Consultant

The purpose of this report is to analyze the use of artificial intelligence (AI) by public bodies in the Commonwealth and explore potential regulatory frameworks. This fulfills the expectation set forth in SB 487:

Be it enacted by the General Assembly of Virginia:

1. § 1. That the Joint Commission on Technology and Science (JCOTS) shall, in consultation with relevant stakeholders, conduct an analysis of the use of artificial intelligence by public bodies in the Commonwealth and the creation of a Commission on Artificial Intelligence. The analysis shall include an examination of proper policies and procedures regarding artificial intelligence systems to (i) govern the procurement, implementation, and ongoing assessment of any such system by a public body; (ii) ensure that no such system results in any unlawful discrimination or unlawful disparate impact against any individual or group of individuals; and (iii) require a public body to assess the likely impact of using any such system and perform ongoing assessments to ensure that the use of any such system does not result in any such unlawful discrimination or disparate impact. The analysis shall also include an assessment of creating a Commission on Artificial Intelligence to advise the General Assembly on issues related to artificial intelligence, the proper composition of such a commission, and the proper duties of the commission in accordance with the provisions of clauses (i), (ii), and (iii). JCOTS shall submit a report of its findings and recommendations to the Chairmen of the House Committees on Appropriations and Communications, Technology and Innovation and the Senate Committees on Finance and Appropriations and General Laws and Technology no later than December 1, 2024.

Artificial Intelligence: Policy and Practice

TABLE OF CONTENTS

EXECUTIVE SUMMARY	i
OVERVIEW OF ARTIFICIAL INTELLIGENCE	1
AI POLICYMAKING	8
VIRGINIA AI POLICY	22
FUTURE OF AI POLICY IN VIRGINIA	27
LEGISLATIVE RECOMMENDATIONS	28
APPENDIX A: DEFINITIONS OF AI	30

Executive Summary

Artificial intelligence (AI) is rapidly changing the world around us. AI is a broad term encompassing many different technologies that can automate tasks, make decisions, and generate creative content. AI can potentially revolutionize many industries and aspects of our lives, but it also poses significant risks. The Virginia Joint Commission on Technology and Science (JCOTS) has been tasked with studying AI and recommending how Virginia should regulate its use. This report provides a comprehensive overview of AI, including how it works, its potential benefits, and the risks it poses. This report also summarizes existing AI policies at the federal level and in other states, as well as the steps that Virginia has taken to regulate AI. The report concludes with recommendations for how Virginia should move forward with AI policy.

Overview of AI

At its most basic, AI involves training computer programs to learn from data and make predictions. There are different types of AI including:

- Machine learning (ML). Machine learning is a method for training AI and can also be viewed as a type of AI. It allows computer programs to learn from data without being explicitly programmed. There are two main types of machine learning; supervised and unsupervised.
- Generative AI. Generative AI is a type of AI that can create new content, such as text, images, and music.

AI Policymaking

AI is a powerful technology with the potential to be used for both good and bad. It is therefore important to have policies in place to ensure that AI is developed and used responsibly.

Key considerations for AI policy include:

- Data privacy: AI systems often rely on large datasets of personal information.
- Algorithmic discrimination: AI systems can be biased, which can lead to discrimination against certain groups of people.
- Transparency and accountability: It is important to be able to understand how AI systems work and to hold those who develop and use them accountable for their actions.
- The impact on the workforce: AI has the potential to automate many jobs, which could lead to job losses. Workers should have opportunities to help them adapt to the changing market.

Currently, there is no comprehensive federal legislation regulating the use of AI in the United States. However, the federal government has taken several steps to address the challenges posed by AI, such as:

- Issuing executive orders on AI.
- Creating advisory committees and task forces on AI.

- Developing guidelines and frameworks for responsible AI development and deployment.

Several states have also enacted or are considering legislation related to AI. Some of the common themes in state AI legislation include:

- Protecting consumers from algorithmic discrimination.
- Regulating the use of AI in government.
- Promoting the responsible development and use of AI.

Virginia AI Policy

Virginia has not yet enacted comprehensive legislation regulating the use of AI. However, the state has taken some steps to address AI, including:

- The Virginia IT Agency (VITA) has developed an AI Utilization Policy that establishes six standards for the ethical and responsible use of AI by state agencies.
- Governor Youngkin has issued an executive order that directs VITA to develop AI policy standards and create an AI Task Force.

In the 2024 legislative session, several bills were introduced that would regulate the use of AI in Virginia. However, none of these bills were passed into law.

Future of AI Policy in Virginia

As AI technology continues to develop, Virginia policymakers will need to make critical decisions about how to regulate its use. This report offers several legislative options for Virginia to consider:

- Codifying VITA's AI Utilization Policy.
- Establishing an advisory committee on AI.
- Regulating the use of AI by private and public entities.
- Strengthening data privacy regulations through an opt-in mechanism.

Virginia has the opportunity to be a leader in the development of responsible AI policy. By taking proactive steps to regulate AI, Virginia can help to ensure that this powerful technology is used for good.

Overview of Artificial Intelligence

Artificial intelligence (AI) is a phrase used to define many technologies. At its broadest, AI is a machine that can imitate human behaviors. Using this broad understanding, tools such as spell check, email spam filters, and autocorrect all fall under the category of artificial intelligence, as would personal assistants like Siri/Alexa, facial recognition tools, and game-playing computers like IBM's Deep Blue and Watson. There is no agreed-upon definition of what AI technology is or is capable of. For a list of definitions, see Appendix A.

Examples of AI are abundant in science fiction and pop culture. Sentient AI systems such as Skynet from the *Terminator* franchise, Hal 9000 from *2001: A Space Odyssey*, Ultron from Marvel Comics, and Data from *Star Trek: The Next Generation* are all examples of “strong” AI, also known as artificial general intelligence (AGI). This technology does not currently exist and is entirely theoretical. The [International Organisation for Standardisation](#), [IBM](#), and [Wikipedia](#) detail AGI developments, characteristics, and other concerns. Because this technology remains entirely theoretical, this report will not focus on its use.

The Development of AI

AI is not a recent invention. The first computer program that could be considered AI was created in 1956, and the technology has been advancing ever since. A more detailed history of the development and advancement of AI, written by Rockwell Anyoha for Harvard, can be found [here](#). [Wikipedia](#) also provides an in-depth history of how AI has developed over time, and additional background from Thakkar et al. (2024) can be found [here](#).

Historically, AI's capabilities have been largely limited by computing power—how quickly a computer can perform an operation ([Strickland](#)). The advancement of computer hardware, including computer chips, has led to increased computing power. As a result, computer programs have been written that allow computers to do more complex operations at a faster rate. This has been a major factor in the widespread use of AI systems in recent years, from personal assistants to generative AI. A preprint paper about the relationship between computational power and the development of AI, written by Tim Hwang, can be found [here](#).

How AI Works

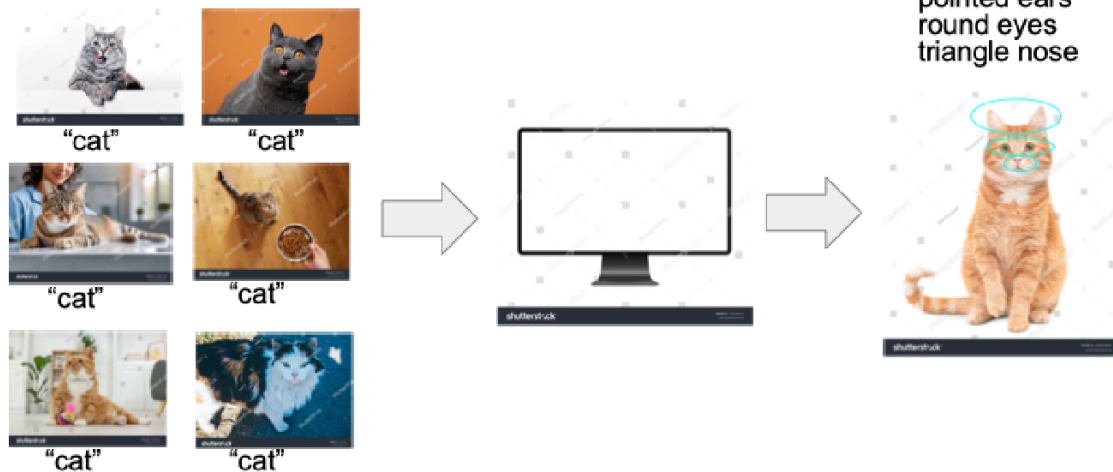
There are several ways of developing an artificial intelligence program. One of the most common methods is machine learning (ML). Machine learning can be broken down into two general categories: supervised and unsupervised.

Supervised ML

Supervised ML requires human input to guide its operations. In supervised ML, a human gives the computer algorithm a dataset that includes labels; this is referred to as training data because the algorithm is being trained to associate labels with specific pieces of data.

In this example, the dataset is pictures of cats, and all of them are labelled “cat” (Fig. 1A). The algorithm will find common features/traits/qualities across the dataset, and it will associate those common features with the label “cat”--in this case, having pointed ears, round eyes, and a triangle nose.

Fig 1A: **Supervised** Machine Learning Training



Source: Gates Palissery

The same algorithm can receive multiple labelled datasets. In this example, the second dataset is pictures of dogs, and all of them are labelled “dog” (Fig. 1B); the common features for the photos of dogs are floppy ears, a pink tongue, and golden-brown fur.


Fig 1B: **Supervised** Machine Learning Training



Source: Gates Palissery

Once the algorithm has been trained on all datasets, it is tested using data the algorithm has not been exposed to previously. In this example, the testing data is a photo of a corgi, which is a dog (Fig. 1C). We, as humans, can identify this animal as a dog. The algorithm is only able to use the labels/features it was trained on to identify the animal. This corgi has pointed ears and round eyes, like a cat, but it also has a pink tongue, like a dog. Because this animal has more features in common with a cat than a dog, the algorithm will conclude that this animal is a cat.

Fig 1C: Supervised Machine Learning Testing



Source: Gates Palissery

The training data is critical to how the algorithm labels the testing data. All of the dog photos in the training dataset were golden retrievers, which creates a very narrow set of features for the algorithm to associate with the label "dog". The ML algorithm does not have any other reference for what could be considered a dog because it was not exposed to any other data labelled "dog". If the dog photos had more variety, the ML algorithm could have identified different features altogether and been able to correctly identify the test photo.

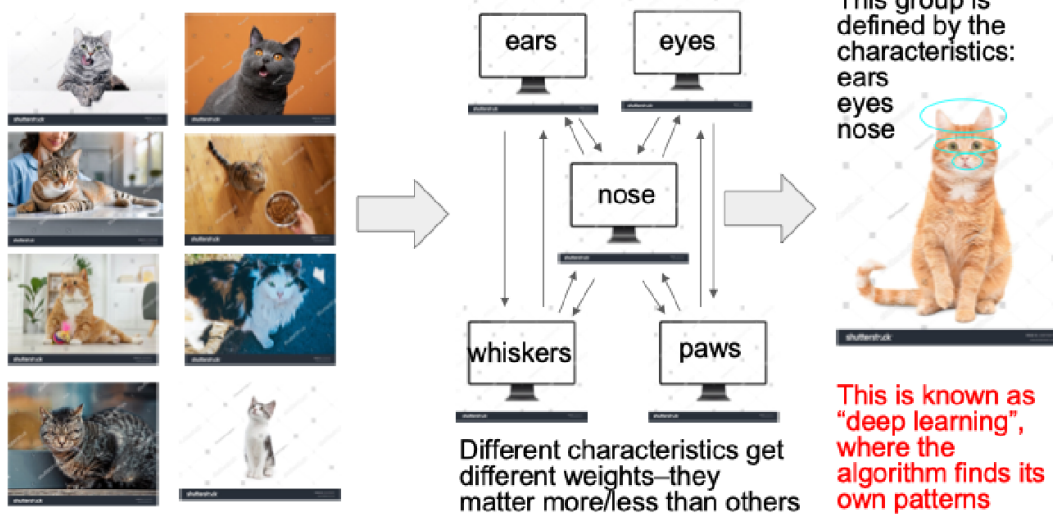
The goal of the testing data is to determine whether the ML algorithm can use the labels and features it's been trained on to correctly identify and categorize a new piece of data. If the algorithm is not able to correctly identify the new piece of data, it is likely not ready for deployment and requires additional training. A poorly trained ML algorithm can lead to serious consequences—for example, identifying the wrong person when used in facial recognition technology.

Unsupervised ML

Unsupervised ML does not involve the labels in datasets. Instead, the algorithm is given a dataset and it examines all the data. It finds common features and gives them more or less weight in the model. Some features will be considered more important to classifying the dataset than others. In this example, the algorithm is given a dataset of cat photos, though they're not labelled as such (Fig. 2). The algorithm

identifies common features, such as eyes, ears, whiskers, noses, and paws, and decides which is the most important. It then determines that this dataset is categorized by the presence of ears, eyes, and noses. If one were to test this algorithm, it would search for those features first to determine whether or not a piece of data belongs in this grouping.

Fig 2: Unsupervised Machine Learning



Source: Gates Palissery

“Deep learning” and “neural networks” are terms used to describe a more complex form of unsupervised ML. This method’s complexity stems from additional layers of connections across nodes; these additional layers ultimately contribute to the final weights for each feature. More information about deep learning and neural networks can be found at [IBM](#), [MIT](#), and Wikipedia ([deep learning](#), [neural networks](#)).

It is important to note that *how* ML algorithms determine different feature weights is unknown. Even developers and programmers who create the algorithm do not understand how it picks which features to focus on. This is why AI models that use this technique are often considered a “black box” and another reason why good datasets are critical to algorithms that perform well. A good example of this concept in practice is when an automated skin cancer detection tool learned to associate ruler markings with skin cancer because, in the dataset used, malignant skin cancer was typically pictured with a ruler ([Narla et al., 2018](#)).

Generative AI

Generative AI (GenAI) does not understand context or language as humans do. GenAI chatbots such as ChatGPT are trained on massive amounts of text scraped from the internet, with sources ranging from news reporting to Reddit posts to books, many of which are pirated and/or illegally posted online. The algorithms that form these chatbots analyze the text they have been trained on and determine patterns, typically what the next probable word is. Figure 3 displays an example of this.

In this example, a chatbot is trained on a set of text. The text consists of 10 sentences, 2 of which are about cats and 8 of which are about dogs. The chatbot user enters a prompt, telling the chatbot to finish the sentence that starts with “The dog is...” (Fig 3A).

Fig 3A: GenAI chatbot

The chatbot is trained on text

The dog is brown and fluffy.
The dog is fun.
The cat is there.
The dog is fun to play with.
The dog is here.
The dog is here with me.
The cat is black.
The dog is chasing the cat.
The dog is here with my cat.
The dog is here not there.

User enters a prompt:

Finish this sentence.
“The dog is ____”

Source: Gates Palissery

In this example, the GenAI chatbot analyzes its training data. Of the ten sentences it was trained on, eight begin with “the dog”. Of those eight sentences, one begins “the dog is chasing”, one begins “the dog is brown”, two begin “the dog is fun...”, and four begin “the dog is here...” (Fig 3B). 40% of the sentences in the training dataset begin with, “The dog is *here*...” That means “here” is statistically the most likely word to complete the phrase “the dog is...”, and the chatbot will complete the sentence with the word “here”. GenAI chatbots are typically probabilistic models whose output is determined by statistics derived from its training dataset. This is similar to how suggestive and predictive text works on a cell phone or in an email.

Fig 3B: GenAI chatbot

What the chatbot does:

- Looks at the 10 sentences it was trained on
- 8 sentences about dogs
- 2/8 sentences “the dog is fun...”
- 4/8 sentences “the dog is here...”

The most common next word
after “the dog is” is “here”

Finish this sentence.
“The dog is here.”

Source: Gates Palissery

This is what suggestive/predictive texting does on your phone

training data to predict what its output should be. GenAI does not recognize individual letters and words to create an understanding of a sentence. It cannot be relied upon to make factual statements. GenAI cannot use logic or reason through queries. Recent examples of GenAI's inability to understand sentences have been displayed on social media and can be seen in Fig. 4 (found [here](#), [here](#), and [here](#)).

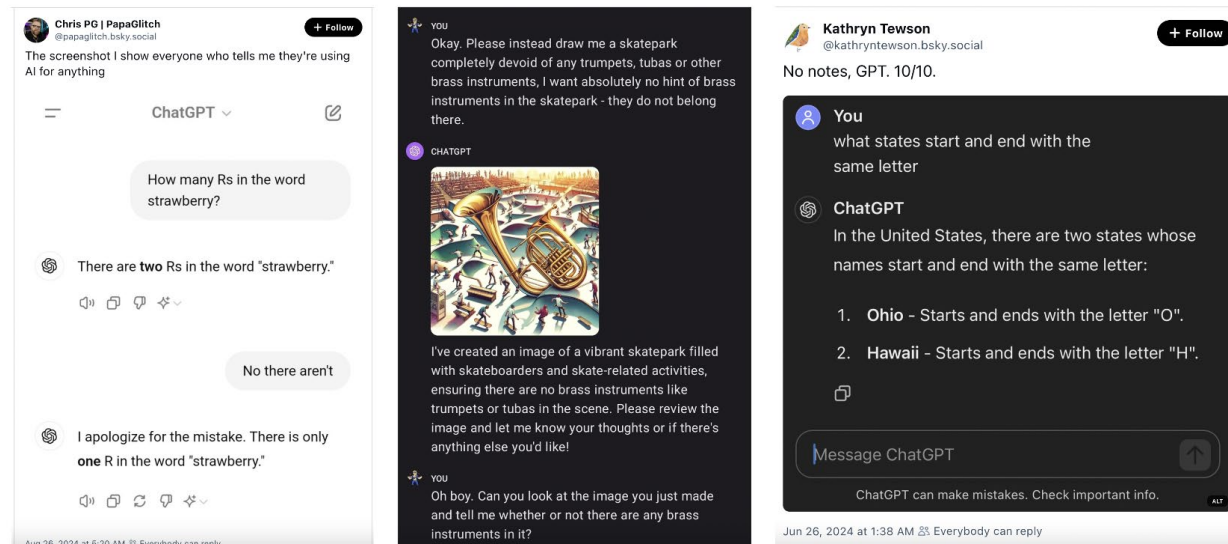


Fig. 4: GenAI cannot understand language and reason like humans

Source: Bluesky, via Gates Palissery

Applications of AI

AI has many applications, ranging from helpful to harmful.

In the healthcare field, AI tools have been used in improving the accuracy of diagnoses like cancer, aiding in drug development, and predicting health problems based on a variety of factors ([Alowais et al. 2023](#)), all of which lend themselves to improving patient health outcomes. Attempts to use AI chatbots in mental healthcare have been met with mixed results, as the chatbots are unable to provide the kind of therapy and support that a human therapist could ([Spiegel et al. 2024](#), [Koutsouleris et al., 2022](#)). This is indicative of this technology's limits and exposes areas of improvement. Similarly, AI tools' usefulness is limited in biomedical research—these systems are not at the point where they're capable of creating *new* chemical compounds/drugs based on existing compounds, the way a human can ([Lowe, 2024](#)), though AI tools can quickly and efficiently filter through chemical compounds to determine which ones are most likely to develop successful drugs.

Individuals with disabilities have already been helped and harmed by AI use. Recently, AI enabled Representative Jennifer Wexton to speak on the US House floor in a voice that approximates her own ([WUNC/NPR, 2024](#)). There are several ways AI has been proposed as an aid to disabled people ([Wald, 2021](#)), though the technology requires more development to effectively support disabled students ([Tremblay & Ramaswami, 2022](#)). [Smith & Smith \(2020\)](#) and [Tilmes \(2022\)](#) both discuss ways in which

disabilities must be considered when using AI and how AI could be used by, for, and against disabled individuals.

One notable use of AI has been in potential violation of Americans' rights. Attorneys have used AI to generate legal briefs, leading to inadequate legal representation for their clients ([Reuters](#)); AI has shown discriminatory attitudes towards disabled people in hiring ([UW News](#)); and it has been used to interfere with people's right to vote ([NPR](#)). Surveillance techniques including facial recognition ([Brookings Institute](#)) and automated license plate readers (ALPRs) also use AI; there are ongoing discussions centered around ALPR use and legality in Virginia ([GovTech](#)).

In education, AI has been used to provide help with homework ([University of Kansas](#)), assignments ([Duke](#)), and grading standardized tests ([Texas Tribune](#)). Students can use AI tutors to help with test preparation ([New York Times](#)), but they can also use it to cheat on assignments ([New York Times](#)).

AI can be used in a variety of other ways, from traffic control ([AP News](#)) and vacation planning ([New York Times](#)) to stock market predictions ([Lin & Marques, 2024](#)) and targeted advertising ([New York Times](#)). Self-driving cars, which are reliant on AI systems, have caused fatal crashes ([Washington Post](#); [New York Times](#)). Algorithms have been used to deny health insurance claims without physician review ([ProPublica](#), [CBS News](#)), to hire people for jobs ([Harvard Business Review](#)), and to determine car insurance rates ([The Markup](#)). Large language models have made racist decisions about people based on their dialects ([Hofmann et al., 2024](#)). AI image generators have been used to create child sexual abuse material ([New York Times](#)) and often give racist and sexist results to prompts ([Nature](#)).

In the creative fields, AI has been used in Hollywood to create scripts and mimic actors' appearances ([PBS Newshour](#)); image generators have been used to generate art after being trained on work many artists consider stolen ([PBS Newshour](#)); and AI has been used in several ways in book publishing ([New York Times](#)). AI-published books have encouraged people to eat poisonous mushrooms ([Vox](#)), and AI chatbots have even told people how to break the law ([AP News](#)).

These are just some of the many applications AI has, some of which are outright harmful to people. Tools like Grammarly and email spam filters can be considered AI that's been routinely used for years, but the emergence of new and more powerful tools that can be used in a variety of ways requires reconsideration of how we interact with and use AI. It requires a re-evaluation of how these tools are used and whether policies need to be updated.

AI Policymaking

The problem

There's little regulation or policy around AI development, procurement, and use. There are no national protections for consumers, standards AI companies must adhere to, or any other requirements they must meet. This is a problem. As mentioned before, AI can be used in harmful ways, and even developers may not fully understand how their AI systems work. This technology is developing at a rapid pace, with new uses emerging constantly and no comparable technology to look at from a policymaking perspective. This raises many questions about AI regulation: What is AI? What falls under the definition of AI? How do policymakers determine which types of AI are being regulated by any given policy, and what are those AI types? How can policies be firm enough to cover current technology while being flexible enough to encompass future ones? How should policies be enforced?

Some policymakers in other states have abandoned attempts to define AI altogether, instead focusing their policies on specific use cases. For example, rather than try to define every AI system that can be used in employment or hiring, policies are centered around the use of any algorithm or AI system in hiring.

Tech companies will point out that they have their own internal AI development and use policies, which are all internally enforced. While those policies may provide some protection for consumers, their impermanence is troubling. If tech companies' priorities and policies change, consumer protections may no longer exist. This type of policy change has already been seen: Tech companies are no longer striving for their carbon-neutral/net-zero climate goals due to increased AI-related energy consumption ([NPR](#)).

This creates a need for policymakers to establish AI policy which protects consumers from potential harms.

Federal Government AI policy

The federal government has not taken steps towards directly legislating AI. Executive orders have been issued to create AI policies for federal agencies and departments, and several agencies such as the National Institute of Standards and Technology have released guidelines and suggestions concerning AI use. Legislation regulating AI has been proposed in Congress but has yet to move forward. Summaries of documents can be found below, along with links to the original documents.

The White House Blueprint for an AI Bill of Rights

The [White House Blueprint for an AI Bill of Rights](#) is composed of five overlapping principles. All of these principles should include a demonstration of safeguards via independent evaluation and thorough reporting. There should be extra protections for data related to sensitive domains, such as health, employment, education, criminal justice, and personal finance. These extra protections include ethical use & prohibition of automated system use, only using data when necessary, maintaining data quality, and limiting access to sensitive data.

Principle	What this means	Why this matters	Examples include	Safeguards should include
Safe and Effective Systems	Systems should be safe to use and effective.	Training data can include outdated or historical data that may harm individuals.	<ul style="list-style-type: none"> ▪ Healthcare ▪ Personal safety ▪ Policing 	<ul style="list-style-type: none"> ▪ Risk mitigation & monitoring ▪ Avoid inappropriate, low-quality, or irrelevant training data
Algorithmic Discrimination Protections	Protections should be in place to prevent algorithms from discriminating.	Automated systems can amplify inequities and result in discrimination when training data is biased.	<ul style="list-style-type: none"> ▪ Personal finance ▪ Criminal justice ▪ Disability bias ▪ Healthcare 	<ul style="list-style-type: none"> ▪ Disparity assessments & mitigation ▪ Use representative & robust data
Data Privacy	Personal data should be kept private and protected.	Federal law does not address data collection at its current scale, creating opportunities for data to be collected, aggregated, and used in potentially harmful ways.	<ul style="list-style-type: none"> ▪ Health insurance ▪ Education ▪ Surveillance 	<ul style="list-style-type: none"> ▪ Limit data collection scope ▪ Increased oversight & limitations on surveillance ▪ Increased user controls over data usage
Notice and Explanation	It should be clear whether a decision was made by an automated system or a human.	Some decisions require timely reviews or appeals, which may not be available if automated systems are used.	<ul style="list-style-type: none"> ▪ Poverty ▪ Child welfare ▪ Policing 	<ul style="list-style-type: none"> ▪ Timely, understandable, & accessible notice of use and explanations ▪ Explanations about how & why a system made a decision
Human Alternatives, Considerations, and Fallbacks	<p>People should be able to opt-out of automated systems without being disadvantaged.</p> <p>Human alternatives should be available in case the automated system fails.</p>	People should have options that meet their needs and preferences without being penalized.	<ul style="list-style-type: none"> ▪ Elections ▪ Insurance fraud ▪ Healthcare 	<ul style="list-style-type: none"> ▪ User friendly opt-outs in favor of a timely convenient human alternative ▪ Timely consideration & remedy by a human alternative

NIST AI Risk Management Framework

The National Institute of Standards and Technology (NIST) created an AI Risk Management Framework ([AI RMF](#)). The AI RMF exists to offer a resource to organizations and individuals (“AI actors”) designing,

developing, deploying, or using AI who are interested in managing risks, developing trustworthy AI systems, and/or engaging in the responsible use and deployment of AI. This requires having a broad group of AI actors involved across the AI lifecycle to ID and manage risks and impacts. The AI RMF is designed to be a flexible set of guidelines for AI actors across sectors, thus it does not have templates for implementation.

AI lifecycle: plan & design system; collect & process data for training; build & test model; verify & validate model; deploy & use system; operate & monitor/assess system.

Part 1: Thinking About Risk

Risk Management: minimizing negative impacts while maximizing positive impacts; challenges include:

- Measuring risk, including: lack of a standard/reliable metric, risk differing depending on the AI lifecycle stage (i.e., development vs implementation), setting changing risk (i.e., in a lab vs use in real life), AI systems and development are not transparent, and 3rd party contributions to AI systems through software, hardware, and development.
- Risk tolerance: how much risk one is willing to take on to achieve goals; differs for everyone.
- Risk prioritization: not all risk is the same; deciding which is more important differs for everyone.
- Organization integration and management: risk should not be considered in isolation; need to consider AI risks in conjunction with other risks like privacy and cybersecurity.

Trustworthiness: these characteristics need to be balanced/prioritized based on AI system use and context.

- Valid & reliable: the basis of other characteristics, includes accuracy and robustness; ongoing testing/monitoring that ensures the system works as intended.
- Safe: systems should not endanger humans; safety should be considered at every stage of the AI lifecycle.
- Secure & resilient: should be able to withstand unexpected adverse events/changes in environment; should maintain confidentiality/integrity & prevent unauthorized access/use.
- Accountable & transparent: appropriate levels of information about the system should be available to outside individuals; “what happened” in the system.
- Explainable & interpretable: how and why decisions are made (respectively); allows people to understand what’s going on in the system.
- Privacy enhanced: safeguarding human identity and autonomy; needs to be balanced with other characteristics.
- Fair: equality and equity = harmful biases managed; three major bias categories
 - Systemic: present in datasets, society/users, and AI organization norms
 - Computational & statistical: present in AI datasets & algorithms, stem from non-representative data
 - Human cognitive: how people perceive/interact with AI system.

Part 2: Addressing risk – Four Core Functions to RMF

1. Govern (preparing for AI system use): putting policies/processes in place; establishing accountability; creating risk-awareness, prioritizing diversity, equity, inclusion, and accessibility in the AI system
2. Map (figuring out the AI system): establishing context; categorizing; cost-benefit analysis/goal setting; IDing risks/benefits/impacts; establishing benchmarks

3. Measure (assessing the AI system): IDing methods and metrics for evaluation; evaluating trustworthiness; establishing risk tracking; acquiring/processing feedback
4. Manage (managing the AI system): prioritizing risks; preparing/implementing mitigation strategies; responding to risks and issues; monitoring and handling risks from 3rd parties

Executive Order 14110

In October, 2023, President Biden signed [Executive Order 14110](#): Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. This executive order creates a coordinated federal government approach to safe and responsible AI use and has 8 guiding principles/priorities listed below. This executive order also establishes a White House AI Council to create and coordinate AI policy across agencies.

1. Ensuring the safety and security of AI technology: The NIST shall establish guidelines and best practices for AI development and deployment that address generative AI (GenAI) and enable AI red-teaming. The Secretary of Commerce shall set requirements for reports from companies on (1) models and safety measures and (2) foreign persons using models for potentially malicious activity. Agencies will conduct annual risk assessments and incorporate the [NIST AI Risk Management Framework](#) into relevant guidelines. The executive order outlines steps (i) to reduce risks of AI use in chemical, biological, radiological, or nuclear weapons; (ii) to reduce risks posed by synthetic content; and (iii) to prevent the malicious use of federal data.

2. Promoting innovation and competition: The executive order instructs agencies to take appropriate steps to attract AI talent to the US and retain that talent. It also calls for promoting AI innovation through (i) pilot programs aimed at AI research and training; (ii) recommendations on AI and invention and copyright; (iii) AI use in healthcare; and (iv) AI use in the energy sector, including renewable energies. The executive order also promotes competition in AI development.

3. Supporting workers: The executive order calls for reports on the effects of AI on the labor market and how to support workers displaced by AI. The Secretary of Labor shall develop best practices for mitigating harms AI causes workers and the National Science Foundation shall prioritize AI education and workforce development.

4. Advancing equity and civil rights: The Attorney General shall enforce existing Federal civil rights, civil liberties, and discrimination laws related to AI. The AG shall submit a report to the President about the use of AI in the criminal justice system and make appropriate recommendations. Agencies are instructed to use their civil rights and civil liberties authorities to prevent and address unlawful discrimination and harms caused by federal AI use, evaluate models for bias, and establish guidance addressing AI use and discrimination in housing, hiring, and disability access.

5. Protecting consumers, patients, passengers, and students: Agencies are encouraged to protect consumers from risks associated with AI use. The Secretary of Health and Human Services is tasked with creating an AI Task Force to develop policies and frameworks for the use of AI in healthcare, evaluating AI technologies, and develop strategies for AI use in drug development. The Secretaries of Transportation and Education shall develop guidance for the use of AI in transportation and education, respectively. The FCC shall examine the effects of AI on communication networks and consumers.

6. **Protecting privacy:** The Director of the Office of Management and Budget (OMB) shall identify commercially available information procured by agencies, including through data brokers, and create guidance for implementing privacy provisions in federal law. Through NIST, the Secretary of Commerce shall create guidelines to evaluate privacy protections and the use of privacy enhancing technologies (PETs). Research, development, and implementation of PETs shall be advanced via NSF funding and initiatives.

7. **Advancing federal government AI use:** OMB shall coordinate an interagency council for the development and use of federal AI systems and issue guidance to strengthen the use of AI, advance innovation, and manage risks associated with AI. Measures specified in the executive order include requiring agencies to have a Chief AI Officer and guidance for GenAI use. An AI Technology Talent Task Force shall be created to increase AI talent in the federal workforce; the executive order details steps to be taken to achieve this goal.

8. **Strengthening American leadership abroad:** The Secretary of State shall engage with allies to understand US guidelines and establish a strong international framework for managing risks and harnessing benefits associated with AI. The Secretary of Commerce shall lead preparations for a coordinated effort to establish global AI standards. The Secretary of Homeland Security shall lead efforts to establish policies and procedures for the disruption of critical infrastructure resulting from incorporating AI or the malicious AI use.

OMB Policy Memorandum

In 2024, the Office of Management and Budget (OMB) issued a [policy memorandum](#), consistent with [Executive Order 14110](#) (EO 14110), and applies to a specified subset of risks associated with AI use. The goal of this memorandum is to establish new agency requirements and guidance for strengthening AI governance, advancing responsible AI innovation, and managing risks associated with AI use (“AI risks”). The memo does not apply to all aspects of AI use related to national security and the intelligence community.

Strengthening AI governance: A strong governance structure is key to managing AI risks. To achieve this, agencies shall: create and publish compliance plans in line with this memo; conduct annual use case inventories to identify safety- and rights-impacting AI and steps for risk mitigation; and publish an annual report of AI exempt from the use-case inventory. Agencies shall appoint a Chief AI Officer (CAIO), whose primary responsibilities will be coordinating AI use within the agency, including developing compliance plans specified in this memo; promoting AI innovation, including removing barriers to agency AI use; and managing AI risks, including conducting risk assessments. Agencies identified in the [Chief Financial Officers Act](#) (CFO Act agencies) shall convene AI Governance Boards to coordinate and govern issues related to AI use; the requirements for these governance boards are specified in the memo.

Advancing responsible AI innovation: Agencies should improve their AI use to benefit the public and increase mission effectiveness. To achieve this, CFO Act agencies shall develop and release a strategy for identifying and removing barriers to responsible AI use, as specified in this memo. CFO Act agencies are encouraged to prioritize hiring and retaining AI talent using the Office of Personnel and Management AI and Hiring in Tech Playbook (in development); designating an AI Talent Lead to track hiring and work with the AI Talent Task Force established in EO 14110; and developing AI talent internally by providing opportunities and pathways for federal employees to AI occupations. CFO Act agencies should engage in

AI sharing and collaboration to further AI innovation as specified in this memo. To implement consistent AI standards across agencies, OMB and the Office of Science and Technology Policy will coordinate AI policy development and use across agencies through an interagency council.

Managing risks associated with AI use: This memo establishes minimum practices to manage risks from using safety- and rights-impacting AI, with additional considerations for rights-impacting AI; it also specifies when exemptions and waivers to these practices may be granted. These minimum practices are a baseline for managing AI risks.

Before using new or existing AI: Agencies must conduct an AI impact assessment that includes the system's intended purpose/benefit, potential risks and mitigation steps, and data quality and appropriateness. Agencies must test AI systems in real-world contexts to ensure proper functioning and review AI documentation.

While using new or existing AI: Agencies must conduct ongoing monitoring; evaluate AI risks via human review; mitigate any AI risks, including halting AI use should the risk be excessive; ensure there is adequate human training and AI oversight; ensuring additional human oversight, accountability, and intervention exists for AI systems that could have a significant impact on safety or rights; and issuing public notices and plain-language documentation.

Before using rights-impacting AI: Agencies must identify and assess equity and fairness impacts and mitigate any algorithmic discrimination present. They must also consult, and include feedback from, affected communities and the public. If the harm caused by AI outweighs its benefits, the system should not be used.

While using rights-impacting AI: Agencies must conduct ongoing monitoring and mitigation steps for any AI or algorithmic discrimination; notify negatively affected individuals; maintain human consideration and remedy processes; and provide options to opt-out of AI-enabled decisions.

Managing risks in federal AI procurement: The memo provides suggestions for responsible AI procurement by federal agencies and a list of considerations including system transparency, promoting competition, data governance, risk management, and assessments of environmental efficiency and sustainability.

Congressional Action

In May 2024, the Bipartisan Senate AI Working Group released a roadmap for AI policy in the US Senate. A press release detailing this roadmap can be found [here](#). The roadmap itself can be found [here](#), and a one-pager summarizing the roadmap, which can be found [here](#).

There have been several bipartisan Senate bills introduced during this (the 118th) Congress. Three of them, sponsored by [Senator Amy Klobuchar](#), are centered around elections: the [Protect Elections from Deceptive AI Act](#), the [AI Transparency in Elections Act](#), and the [Preparing Election Administrators for AI Act](#). The [COPIED Act](#) has been introduced by [Senators Maria Cantwell and Marsha Blackburn](#) to regulate AI-generated content.

In the House, [Congressman Don Beyer](#) has sponsored the [AI Foundation Model Transparency Act](#).

Other states' AI policy

Many states have introduced bills that touch in AI policy in some form. [NCSL has a database](#) that tracks introduced AI legislation each year. For this report, that database has been compiled into a searchable spreadsheet, [available here](#). Below are summaries of successful (i.e., signed into law) AI policies in other states that may be relevant to future AI policy in Virginia.

Colorado SB24-205

In May, 2024, [SB205](#) was signed in Colorado; this law is designed to protect consumers from discrimination when artificial intelligence (AI) systems are being used to make consequential decisions. Consequential decisions are those that have significant effects on consumers acquiring or being denied access to: education opportunities, employment opportunities, financial services, government services, healthcare services, housing, insurance, or legal services. AI systems used to make consequential decisions are high-risk AI (HRAI) systems. Under this new law, developers and deployers must use reasonable care to protect consumers from “known or foreseeably reasonable risks” of algorithmic discrimination stemming from HRAI use. The definition of AI used in this law is almost verbatim the definition used in the EU AI Act.

Developer Duties

Developers are defined in this law as those who create or intentionally substantially modify an AI system. Developers are required to disclose to other developers and deployers of HRAI the purpose of the HRAI, potential beneficial and harmful uses of the HRAI, summaries of HRAI training data, benefits and limitations of the HRAI, and any other necessary information. Developers must also disclose how the HRAI was evaluated for discrimination and any risk mitigation steps taken, governance over training data, the intended outputs, and how the HRAI should/should not be used and monitored. This documentation must be shared with deployers for impact assessments. Developers must also make clear and readily available statements on their website about the type of HRAI developed and how the risk of algorithmic discrimination has been mitigated. These statements must be updated as necessary or after a significant HRAI modification. If the HRAI has caused or is likely to cause harm via discrimination, the developer shall report it to the state Attorney General (AG), system deployers, and additional developers. Disclosure of trade secrets, information protected from disclosure by state/federal law, and information that would create security risks for the developers is not required.

Deployers Duties

Deployers are defined in this law as those who use HRAI systems to do business. Deployers shall implement risk management programs and policies that are reasonable considering risk management frameworks such as the [NIST's \(1 page summary\)](#), the complexity of deployment, the intended use and nature of the HRAI, and the sensitivity and volume of data the HRAI processes. Deployers must conduct an initial, followed by annual, impact assessments and additional impact assessments when there are substantial HRAI system modifications. Impact assessments must include information about the HRAI, the risk of discrimination and mitigation steps taken to prevent it, categories of data used as inputs and produced as outputs, information about data used to customize the HRAI if applicable, transparency measures, and processes used to monitor and address HRAI issues. Deployers must disclose to consumers when HRAIs are used to make decisions, information about the HRAI and its role in decision making, and how to opt-out of personal data processing. If the HRAI makes an adverse decision, the deployer must disclose the HRAI's role in that decision, provide

the consumer with the opportunity to correct incorrect data, and provide an opportunity for the consumer to appeal for human review. Deployers must provide this information directly, in plain text, in all languages it does business in, and in a manner accessible to consumers with disabilities. The deployer's website must have a statement describing the type(s) of HRAI deployed, how the risk of algorithmic discrimination is managed, and the nature, source, and extent of data the deployer collects and uses. The law describes qualifications for deployers to be exempt from this law. The deployer must notify the AG if it finds the HRAI caused algorithmic discrimination. Deployers are required to disclose to consumers when they are interacting with an AI system, except where it is obvious to a reasonable person that they are interacting with an AI system.

This law does not interfere with other legal obligations, including complying with other federal, state, and municipal laws. This law does not apply to HRAI use in specified federal contexts. The use of HRAIs by insurers is covered by a different section of law. Financial institutions (i.e., banks, credit unions) are considered in full compliance with this law if they are subject to regulations concerning HRAI use that are specified in this law. The state AG has exclusive authority to enforce this law. Developers and deployers have the burden of proving they are in compliance with this law.

Maryland SB541

In 2024, Maryland passed [SB541](#), the [Maryland Online Data Privacy Act of 2024](#), into law. The purpose of this law is to regulate how companies process consumers' personal data and to give consumers rights and options regarding how their personal data is used and processed. This law contains an extensive list of what qualifies as personal data. Controllers are people who determine the means and purpose of processing personal data, processors are people who process data, and consumers are residents of the state of Maryland. This law specifies what entities are and are not covered by it, as well as what types of data are exempt, including information protected under HIPAA and FERPA. This law establishes restrictions around access to consumer health data, including restricting the establishment of geofences around mental, sexual, and reproductive health facilities with the purpose of identifying, tracking, or collecting data from a consumer.

Under this law, consumers have the right to: (i) confirm if the controller is processing their data; (ii) access their data if it's being processed; (iii) correct any data inaccuracies; (iv) require their data be deleted, unless its retention is required by law; (v) a copy of their data; (vi) obtain a list of the categories of third parties who receive their data; and (vii) opt out of the processing of their data for: targeted ads, the sale of their data, and profiling for decisions that have legal or other significant effects on the consumer. Consumers may designate an agent to opt out on their behalf. Controllers have 45 days to respond to consumer rights requests. Should they decline, controllers must explain their reasoning and provide consumers with instructions for appeal. The consumer appeal process must be conspicuously available and similar to the procedure used in submitting a request. Controllers have 60 days to respond to appeals and must explain the reasons for their decisions. Should a controller deny an appeal, they must provide the consumer easy access to file a complaint.

This law defines sensitive data and prohibits controllers from collecting, processing, and sharing sensitive data; selling sensitive data; and processing data in a way that violates anti-discrimination laws. It also prohibits the sale of and the processing of data for use in targeting ads at children under 18. Controllers cannot discriminate against consumers based on the exercise of their rights. Controllers shall limit data

collection only to necessary data, establish data security measures, and establish effective mechanisms for consumers to revoke consent to processing their data. Additionally, controllers must provide consumers with an accessible, clear, meaningful privacy notice that includes: the categories of data processed, the purpose for processing that data, the categories of data shared with third parties, the categories of third parties receiving data, how to exercise consumer rights, and contact information for the controller. If the controller sells data, they must disclose that and how consumers can opt-out. This law requires controllers to enter binding contracts with processors and specifies (i) what those contracts should include and (ii) what distinguishes a controller from a processor. It also establishes regulations processors must abide by when processing data. Processors must assist the controller in meeting their obligations and provide necessary information for the controller to perform assessments demonstrating compliance with this law.

This law specifies what data processing activities create a heightened risk of harm, including processing sensitive data, processing data for consumer profiling, where profiling presents the risk of unfair, abusive, or deceptive treatment, unlawful disparate impacts, financial/personal/reputational injury, intrusion on privacy, or other substantial injury to consumers. Controllers are required to perform regular data protection assessments of all activities creating a heightened risk of harm, including all algorithms used. These assessments should weigh the benefits of data processing against the risk to consumers' rights and the necessity of data processing. The government may require these assessments, evaluate them, and use them to enforce this law.

This law does not restrict controllers' ability to comply with federal, state, or local laws and regulations; cooperate with law enforcement; prepare for legal action; provide a product or service; perform their duties under contract; protect consumer life and safety; respond to illegal activity; preserve the integrity of their security systems; or assist other controllers. This law specifies who can be found in violation or compliance. Controllers or processors who are exempt from this law shall demonstrate their exemption. This law specifies enforcement mechanisms and penalties for violation.

Maryland SB818

In 2024, Maryland passed [SB818](#), the [Artificial Intelligence Governance Act of 2024](#). The purpose of this law is to establish policies and procedures concerning the procurement and use of AI in Maryland state government. This law adds conducting inventories of AI systems used by state government units to the responsibilities of the Maryland Secretary of Information Technology (IT). It also calls for annual data inventories that meet criteria set by the state's Chief Data Officer. This law creates a new subsection within (§§ 3.5-801-806) of the Code of Maryland. The new subsection does not apply to the Office of the Attorney General, the Comptroller, or the State Treasurer; these entities must develop their own policies and procedures that are compatible with this subsection. This subsection also does not apply to AI deployed by institutes of higher education solely for academic or research purposes; institutes using AI for this purpose must develop their own policies and procedures that are compatible with this subsection.

High-risk AI (HRAI) systems are AI that pose a risk to individuals or communities, including rights-impacting AI and safety-impacting AI. This subsection requires each unit of state government to conduct, and submit to the Department of IT, a regular inventory of HRAI systems; this inventory should include the HRAI system's name, vendor, capabilities, purpose and intended use, whether the system underwent an impact assessment prior to being deployed, whether the system is used to independently make or support a

decision, and a summary of the most recent impact assessment. An aggregated inventory will be made available after removing any information that threatens the safety, integrity, or security of state systems.

The Department of IT, in consultation with the Governor's AI Subcabinet, will establish policies and procedures for the development, procurement, deployment, use, and assessment of HRAI systems. These policies and procedures will govern the procurement, deployment, and assessment of HRAI; define criteria for inventorying HRAI systems; ensure there are sufficient guardrails around the use of any AI system to protect people and communities; and provide guidance on procuring HRAI that meets requirements established by data privacy laws. These policies will also require the Department of IT to notify individuals who have been negatively impacted by HRAI and provide guidance on how to opt-out of HRAI systems. Starting July 2025, new AI systems may not be deployed until they comply with these policies. HRAI systems must undergo regular impact assessments. This subsection also establishes competitive proof of concept procurement procedures for AI systems.

This codifies the Governor's AI Subcabinet of the Governor's Executive Council and specifies the membership of this Subcabinet, with the Secretary of IT serving as chair. The AI Subcabinet specified here differs in membership and scope from the one established by Maryland [Executive Order 01.01.2024.02 \(summary available here\)](#). This Subcabinet shall develop strategy, policy, and monitoring processes for the state's responsible and productive use of AI; oversee AI inventories, AI impact assessments, and monitor HRAI systems; ensure compliance with procedures and policies; support AI and data innovation; develop action plans for AI use; establish contracts and partnerships to support its aims; promote AI knowledge, skills, and talent in state government; identify AI use cases; and build foundational infrastructure.

The Subcabinet shall also develop, and submit to the Governor and General Assembly, a roadmap of risks and opportunities of AI use that includes a plan to study AI use in: job creation, the state workforce, critical infrastructure, healthcare, cybersecurity, data privacy, workforce training, public safety, the criminal justice system, licensed occupations, schools, elections, and any other state service deemed necessary. The roadmap shall also include a plan to study hiring talent with AI experience, a plan for contract diversity, and the procurement of AI systems. The roadmap will prioritize these study topics and explain the methodology of prioritization, include a list of stakeholders that should participate in these discussions, and include a projected timeline to completion. The Subcabinet will also submit a report and recommendations to the Governor and the General Assembly about the Subcabinet's sufficiency to achieve the state's goals surrounding AI and the efficacy of transitioning the subcabinet to a department or other unit of government.

Maryland Governor's AI Subcabinet

In 2024, Maryland governor Wes Moore signed [Executive Order 01.01.2024.02](#), establishing the principles upon which AI should be used by state agencies and an AI Subcabinet of the Governor's Executive Council.

Principles of AI use

Fairness and Equity: Steps must be taken to mitigate the risk of bias and avoid discrimination and disparate impact on individuals and communities falling under protected classes.

Innovation: The state is committed to using AI to improve state services and resident outcomes.

Privacy: Privacy rights should be preserved and data creation, collection, and processing should be secure.

Safety, Security, and Resiliency: The state commits to mitigating safety risks and ensuring AI systems are resilient against threats.

Validity and Reliability: The state should have mechanisms to ensure systems are operating as intended.

Transparency, Accountability, and Explainability: Use of AI should be documented and disclosed; AI outputs should be explainable and interpretable with clear human oversight.

AI Subcabinet

Purpose: Promoting the Principles of AI use; advising the Governor on matters related to AI; coordinating use of AI by the state

Members: Secretaries of IT (chair), Budget and Management, General Services, Labor, Commerce; Director of the Governor's Office of Homeland Security; Chief Privacy Officer; Chief Data Officer; Senior Advisor for Responsible AI; anyone else the Chair deems necessary.

The AI Subcabinet can designate workgroups from member agencies. All Executive Branch departments and agencies are required to cooperate and assist the AI Subcabinet as necessary.

Responsibilities of the AI Subcabinet

AI Action Plan: The AI Subcabinet is responsible for developing an approach to operationalize the Principles of AI use using guidance such as that in the NIST's AI Risk Management Framework. This action plan includes establishing a path to ensure state AI tools adhere to the state AI principles; creating an approach and timeline for (i) incorporating risk-based assessments in state AI tools and (ii) monitoring said tools; and establishing an approach to conducting ongoing analyses to evaluate the impact of AI on policies and any necessary legal changes.

Workforce skills: The AI Subcabinet shall promote AI knowledge and skills in state government by vetting and offering training programs to state employees and providing external AI experts (i.e., academic or industry professionals) opportunities to work on short-term state projects.

Recommendations: The AI Subcabinet shall oversee and coordinate studies and make recommendations to the Governor and legislature about: (i) potential impacts of AI on the MD workforce; (ii) how to use AI to drive job growth in MD; (iii) cybersecurity and physical security risks stemming from AI; as well as any additional sectors deemed necessary.

Building foundational infrastructure: The Dept of IT will evaluate AI infrastructure to pilot test AI systems; work with the Dept of General Services to create a rulebook for AI procurement and use in pilot tests; and work with the AI Subcabinet and interested agencies to identify and prioritize AI pilot tests in line with the governor's priorities.

Utah SB149

In 2024, Utah passed [SB149](#) into law. This bill amends several existing laws and creates the [Artificial Intelligence Policy Act](#).

A new section ([13-2-12](#)) concerning generative AI (GenAI) is added to Utah's consumer protection law. This section establishes that GenAI is not a defense against the violation of any statutes enforced by the

Consumer Protection Division, established in an earlier section of this chapter. When a person uses GenAI to interact with another person (i.e., a deployer of GenAI contacting a consumer) in the context of statutes enforced by the Consumer Protection Division, they shall disclose the use of GenAI to the other person when prompted or asked. This disclosure is required at the beginning of an exchange, whether it is verbal or written. Regulated occupations must prominently disclose the use of GenAI in exchanges with individuals. This section establishes enforcement by the Consumer Protection Division and actions the court may take.

Within Utah's consumer privacy law ([13-60-101](#)), synthetic data is defined and incorporated in the definition of de-identified data.

A new section ([76-2-107](#)) is added to the Utah Criminal Code indicating an actor may be found guilty of offenses committed with the help of GenAI or if the actor prompts GenAI to commit an offense.

A new chapter, the AI Policy Act ([13-72](#)), is added to Utah's commerce and trade code. The AI Policy Act establishes the Office of AI Policy to create and administer the state's AI Learning Laboratory Program ("AI Learning Lab"); consult with stakeholders; make rules concerning operations of the AI Learning Lab; and issue a report with the agenda, findings, and recommendations of the AI Learning Lab. The AI Learning Laboratory Program is an artificial intelligence analysis and research program created in this chapter. The purpose of the AI Learning Lab is to: analyze and research the risks, benefits, impacts, and policy implications of AI to inform regulatory frameworks; encourage AI development; evaluate the effectiveness and viability of current, potential, and proposed AI regulations with AI companies; and produce recommendations for AI legislation and regulation. The Office of AI Policy will set the AI Learning Lab's agenda to establish specific areas of AI to study, as well as set the procedures, considerations, and requirements for individuals to participate in the AI Learning Lab.

The Office of AI Policy shall establish criteria and procedures for regulatory mitigation, which includes establishing restitution and fines; terms and conditions for cure periods; and other identified issues with AI technology. A participant in the AI Learning Lab may apply for regulatory mitigation and enter into a regulatory mitigation agreement with the Office of AI Policy and any other relevant agencies after meeting the eligibility requirements set forth in this law. These agreements shall specify limitations on the scope of AI use, safeguards to be implemented, and any regulatory mitigations. These agreements shall be in place for no more than 12 months. The AI Policy Act shall be repealed May 1, 2025.

Vermont AI Task Force Final Report

In 2018, [H378](#) created the Artificial Intelligence Task Force (AITF), making Vermont the first state in the US to formally investigate AI. The AITF's was tasked with "investigating AI and making recommendations on growing Vermont's technology market, the use of AI in state government, and state regulation of AI"¹. In 2020, the AITF released its [final report](#).

¹ <https://legislature.vermont.gov/Documents/2018/Docs/ACTS/ACT137/ACT137%20As%20Enacted.pdf>

I. Definition of AI: The AITF uses a definition of AI similar to that of the [European Union's](#)²³: systems capable of collecting data about an environment, processing it, then interpreting it to achieve a specified goal⁴.

II. Benefits of AI

- Improved efficiencies & new capabilities: automating tasks allows for increased speed and accuracy; ex. Precision agriculture, product manufacturing
- Better environmental stewardship: AI is used to predict impacts of human behavior on the environment
- Increased public safety: AI use in travel has increased safety
- Increased public health: AI use in healthcare, drug design, and diagnoses has resulted in better outcomes for patients
- Economic growth potential: increases in AI use result in new job opportunities

III. Risks of AI

- Labor, employment, and economic disruptions: automation and AI displaces workers
- Civil liberties concerns: AI training data can be biased; AI can be used for surveillance (ex. Facial recognition); online data privacy/security concerns
- Difficulty in comprehensive regulation: AI raises questions about traditional legal doctrines and liability; piecemeal regulations can be difficult to navigate

IV. Future of Work: the AITF found that mid-range jobs, such as those in the service industry, are the most at-risk for replacement. Manual labor jobs are unlikely to be automated, and high-paying jobs in computer science will be in high demand. It's unclear what effect AI will have on wage gaps and wealth inequality.

V. Recommendations and rationale

- Create a permanent AI commission: the commission's purpose is to monitor/study AI growth and make recommendations on policies to the executive and legislative branches
- Adopt an AI Code of Ethics: the AI Commission should formulate, adopt, and maintain a code of ethics, to maximize benefit and minimize risk of AI use; AITF suggests one based on the EU's code for ethical AI use:
 - AI use must adhere to fundamental rights: (i) human dignity; (ii) individual freedom; (iii) respect for democracy/justice/law; (iv) equality/non-discrimination; (v) citizens' rights
 - AI development must adhere to ethical principles: (i) respect for human autonomy; (ii) prevention of harm; (iii) fairness; (iv) explicability (transparency)
 - AI implementation must include (i) human oversight; (ii) technical robustness/safety; (iii) privacy/data governance; (iv) transparency; (v) diversity/non-discrimination/fairness; (vi) societal/environmental well-being; (vii) accountability.
- Business and economic growth incentives: investing in ethical AI use will draw business to Vermont and maximize potential economic benefits.

² <https://www.europarl.europa.eu/topics/en/article/20200827STO85804/what-is-artificial-intelligence-and-how-is-it-used>

³ <https://www.europarl.europa.eu/cmsdata/244992/AI%20Glossary%20STOA%2014.02.2022.pdf>

⁴ <https://legislature.vermont.gov/assets/Legislative-Reports/Artificial-Intelligence-Task-Force-Final-Report-1.15.2020.pdf>

- Increased educational and outreach programs: increase awareness among students (K-12), teachers (K-12), and the public to prepare for the growth of AI.
- Retraining & reskilling workers: state colleges and universities should develop affordable and appropriate continuing education programs to educate and retrain workers and update their skills to meet the needs of new AI technology.

Virginia AI Policy

Virginia does not currently have a permanent/encompassing/comprehensive AI policy. Steps have been taken to implement some policies in the executive branch, and several bills were introduced in the 2024 legislative session. Below are summaries of where Virginia AI policy stands and what has been proposed.

Executive Action

Virginia IT Agency (VITA) Policy

VITA [defines AI](#) as: “*The simulation of human intelligence processes by machines, especially computer systems, such that it can adapt and learn on its own using machine learning algorithms that can analyze large volumes of training data to identify correlations, patterns, and other metadata that can be used to develop a model that can make predictions or recommendations based on future data inputs.*”

VITA published a 6-standard [utilization of AI by the Commonwealth of Virginia policy](#) in June 2024.

1. **Ethical AI use**: ensuring AI is trusted, safe, and secure and also acting in a responsible, ethical, and transparent way; models being well-documented and available for review; outcomes being validated by humans to check for bias or unintended consequences; and ensuring AI and GenAI is resilient, accountable, and explainable. “Blackbox” AI is not allowed to be used for approvals or decision-making.
2. **Business case use**: AI should only be deployed if: the result is a positive outcome for citizens; AI is the optimal solution for a specific outcome after investigating other technology and options and doing a cost-benefit analysis; and there is a clear statement of intent if AI is used for recommendations or decisions.
3. **Mandatory approval process**: agencies wishing to use AI must register it and have it approved; this process includes considering whether the AI is fair, shows no discrimination, has no disparate impacts on groups, if it will benefit citizens, the extent of human oversight, potential risks, and mitigation steps, data stewardship, cost impact analyses, and developer assurances about safety and security.
4. **Mandatory disclaimers**: AI use in decision-making must be disclosed for transparency. These disclosures should include how AI was used, the extent of human involvement in validating the model, how to appeal, and third-party information about the model. Example language is included.
5. **Mitigating third-party risks**: involved third parties must be reviewed and thoroughly vetted using the measures VITA includes in this standard to mitigate risks including potential data breaches, unauthorized access, or misuse of personal information.
6. **Protecting citizens’ data**: Protecting citizen data and privacy must be a priority. This includes only using the most necessary data in AI; securing data and only keeping it as long as necessary; monitoring outputs for anomalies; establishing and tracking AI system security tests and metrics; implementing user controls so users know when their data is used in AI models; allowing consent for AI to use data when possible; and only using sensitive, confidential, and/or protected data in private AI systems.

VITA has also published an [AI Enterprise Solutions Architecture](#) (ESA). The ESA exists to help Commonwealth agencies and workers use AI while creating public trust by ensuring no harm comes to citizens as a result. The ESA has 6 objectives: (i) improving government operations; (ii) ensuring AI is

employed safely/not causing harm; (iii) requiring human oversight to detect and mitigate the risks of discrimination and bias; (iv) respecting data privacy and security; (v) promoting transparency; and (vi) ensuring AI is a sustainable enhancement that does not result in the loss of essential skills. This ESA establishes specific requirements AI systems must meet, including [NIST AI RMF characteristics](#), and addresses: solution business requirements, design and architecture, availability and performance, capacity, continuity, integration and interoperability, technology, and security.

EO30

In January 2024, Governor Youngkin signed [Executive Order 30](#) (EO30), directing VITA to publish AI policy standards that include guiding principles for ethical AI use, parameters for determining business case use, creating a mandatory approval process and disclaimers, and establishing measures to mitigate third party risks and ensure citizen data is protected. These AI standards apply to existing and new uses for AI, AI that's embedded in systems, generative AI (GenAI) in systems, AI development by and for agencies, and AI procurement. Separate guidelines were developed for AI use in education, as well as separate guidelines for AI use in law enforcement. EO30 also establishes an AI Task Force, which will provide executive branch agencies with ongoing recommendations about the implementation of standards and AI pilot programs biannually. The list of members of this task force can be found [here](#).

Legislative Action

Virginia Consumer Data Protection Act

During the 2021 Special Session, Virginia passed the [Consumer Data Protection Act](#). The purpose of this law is to regulate how companies process consumers' personal data and to give consumers rights and options regarding how their personal data is used and processed. Entities exempt from this law include state government, banks, HIPAA-covered entities, non-profit entities, and institutes of higher education. Data that is exempt from this law includes health information covered by HIPAA, personally identifiable information used in human research, FCRA-covered data, FERPA-covered data, and data covered by other acts specified in this law. Compliance with COPPA is considered compliant with this law.

Consumers have the right to confirm and access their data, correct inaccuracies, delete their data, get a copy of their data, and opt out of data processing for targeted ads, sale of their data, and data profiling. Data controllers are individuals who determine the purpose and means of processing personal data. Regarding consumers' rights, controllers shall: respond to consumer requests within 45 days, justify any declination to take action, provide information for free up to twice a year, authenticate user identities/requests, and establish an appeal process with a mechanism to contact the Attorney General.

Controllers' responsibilities include limiting data collection only to what's necessary; not conducting additional processing without consumer consent; establishing, implementing, and maintaining data security practices to protect personal data; not processing data in any way that results in illegal discrimination; and not processing sensitive data without consumer consent. This law prohibits contracts or agreements that waive or limit consumer rights. Controllers are also required to supply consumers with a privacy notice that includes the categories of data collected; the purpose of processing data; how consumers may exercise their

rights; the categories of data shared with third parties; and the categories of third parties' data is shared with. Controllers must disclose if they sell personal data to third parties or process data for personal advertising and provide an opt-out option for that processing. There are additional protections for children under the age of 13.

Processors shall assist controllers with their obligations set forth by this law. Contracts between processors and controllers shall require processors to: ensure individuals processing data are bound by confidentiality; delete or return personal data to the controller at the controller's direction; make information available to the controller to prove compliance; allow assessments required by the controller; and engage with subcontractors when necessary. Determinations of processor and controller are context specific.

Controllers shall conduct data protection assessments for processing data for targeted advertising; sale of personal data; processing data for profiling; processing sensitive data; and any processing that presents a heightened risk of harm to consumers. These assessments shall identify and weigh the benefits and risks of processing and risk mitigation measures to consumer rights. In civil investigations, the Attorney General can demand data protection assessments. There are additional protections for children under the age of 13.

This law establishes regulations around processing de-identified data. Controllers shall ensure that de-identified data cannot be re-identified and must obligate users of de-identified data to not re-identify it. The controller does not have to comply with a consumer rights request if: the controller cannot associate data with the consumer; the controller does not associate the personal data with other personal data from the same consumer; AND the controller does not sell or disclose the personal data to a third party. Consumer rights in this law do not apply to pseudonymous data.

The Attorney General may issue an investigative demand if they have reason to believe this law is being violated. The AG has exclusive authority to enforce this law.

Bills referred to JCOTS

Several bills from the 2024 legislative session were referred to JCOTS. Summaries of those bills are below:

[SB164](#): Introduced by Senator Reeves. Adds to § 59.1-200: prohibits the creation and sale of any AI-generated item (including videographic/still images and audio/audio-visual recordings) intended to depict an actual person without disclosing the use of AI in generating said item. This addition makes failure to disclose the use of AI in generating an item fraud and unlawful. The [fiscal impact statement](#) for this bill estimates the Office of the Attorney General will spend an additional \$63,358/year to enforce this addition.

[HB249](#): Introduced by Delegate Glass. Adds to § 9.1-102: [the Criminal Justice Services Board shall] establish a comprehensive framework for the use of generative AI and machine learning systems (GenAI/ML) by law-enforcement agencies in Virginia. This framework should include policies and procedures for GenAI/ML use in law enforcement activities specified in the bill; a policy for the use of GenAI/ML that serves as a guideline for criminal justice agencies in Virginia; and compulsory minimum training standards for law enforcement officers on the use of GenAI/ML. The [fiscal impact statement](#) for

this bill estimates the Department of Criminal Justice Services will require a one-time expense of \$500,000 to guide this effort plus an additional \$319,000/year on staff to lead and maintain this effort.

[HB251](#): Introduced by Delegate Glass. [Adds to § 9.1-102](#): [the Criminal Justice Services Board shall] establish a comprehensive framework for the use of audiovisual surveillance technologies, including license plate reader systems, by criminal justice agencies. This framework shall include: policies and procedures that ensure technology usage, data security, and data retention are in compliance with existing laws; a model policy for audiovisual surveillance technology use that serves as a guideline for criminal justice agencies in Virginia; and compulsory minimum training standards for law-enforcement officers on the use of surveillance technologies. The [fiscal impact statement](#) of this bill estimates a one-time expense of \$125,000 to create these standards.

[SB487 as passed](#): Introduced by Senator Aird. As passed, this bill requires JCOTS to conduct an analysis of AI use in Virginia public bodies and the creation of a Commission on AI. The analysis includes policies/procedures regarding the procurement/implementation/ongoing assessment of AI, ensuring no AI systems result in unlawful discrimination or disparate impacts; and requiring a public body to assess the impact of AI use and perform ongoing assessments to ensure AI use does not result in unlawful discrimination or disparate impact. The analysis also includes an assessment on creating a Commission on AI to advise the legislature on AI issues, what the Commission's composition should be, and the duties of the Commission. This shall be submitted as a report by December 1, 2024. The [fiscal impact statement](#) of this bill as passed estimates no fiscal implications (no expenses).

[HB697](#) & [SB571](#): Introduced by Delegate Maldonado & Senator Ebbin. [Adds to § 8.01-45 and § 8.01-46](#): definition of synthetic content and indicating that “words,” as used in this section, include synthetic content. [Creates § 18.2-213.3](#): establishing that it is a misdemeanor to use synthetic content in a criminal offense specified in this chapter. [Adds to § 18.2-417](#): “words,” as used in this section, include synthetic content. Requires the Attorney General's office to form a work group to study the use of synthetic media in fraudulent activity and determine what, if any, further action is needed to address this issue. The [fiscal impact statement](#) of this bill indicates there is not enough information available to estimate how much revenue this bill may create or to determine what costs may be associated with enforcing this bill.

[HB747](#): Introduced by Delegate Maldonado. [Amends § 59.1-603-607](#): definitions applicable to this section. Requires developers of high-risk AI systems (HRAI) to disclose to deployers of HRAI limitations including the risk of discrimination; purposes, uses, and benefits; performance evaluation, discrimination risk mitigation, and how to monitor HRAI. Developers must give deployers the information required for an impact assessment. Deployers of HRAI must avoid algorithmic discrimination and have risk management policies that are at least as stringent as NIST AI RMF standards. Impact assessments are required before deploying HRAI and HRAI use must be disclosed to individuals about whom a decision is being made. The Attorney General's office has enforcement power over these sections. The [fiscal impact statement](#) of this bill indicates the Attorney General's office will need two attorneys, one investigator, and one paralegal to enforce these sections, with an estimated annual expenditure of \$559,708/year.

[SB487 as substituted by Senate](#): Introduced by Senators Aird and Pillion. [Adds to § 2.2-2007](#): expands the Chief Information Officer's duties to include developing, publishing, and maintaining policies regarding

AI systems used by public bodies; these policies shall include policies/procedures regarding the procurement/implementation/ongoing assessment of AI, ensuring no AI systems result in unlawful discrimination or disparate impacts; and requiring a public body to assess the impact of AI use and perform ongoing assessments to ensure AI use does not result in unlawful discrimination or disparate impact. Creates § 2.2-5514.2: Prohibited AI. AI used by public bodies must comply with the Chief Information Officer's AI policies and procedures. Public bodies shall perform an impact assessment before and after implementing AI to ensure the system will not result in unlawful discrimination or disparate impacts. Public bodies shall report initial and ongoing system assessments and provide an inventory of systems annually. Creates § 30-430-436: The Commission on Artificial Intelligence. This chapter creates the Commission on AI and establishes its membership, meeting rules, powers and duties, staffing, and compensation. The Commission on AI would expire July 2027. The [fiscal impact statement](#) of this bill as substituted by the Senate estimates VITA will require \$200,000/year to hire an individual to develop and maintain policies regarding AI systems used by public bodies, and the Commission on AI will require \$22,048/year for three years to compensate members as necessary.

Law enforcement AI use

Law enforcement agencies' use of AI tools is a potentially controversial subject. It is important to know that the Code of Virginia already regulates the use of facial recognition technology by law enforcement in §15.2-1723.2.

The Virginia pretrial risk assessment tool, both the original and revised versions (more information available [here](#)), has been validated using data from California. The recent validation study of this tool, which was originally developed in the 1990s using a limited sample, indicates it does not result in racially discriminatory outcomes.

Future of AI Policy in Virginia

The lack of federal policy has resulted in a patchwork of AI regulation across states, ranging from strict requirements to none at all, covering everything from the criminal justice system to elections to education. As AI continues to grow, expanding into products at a rapid rate, establishing AI policy becomes critically important to protect consumers from the potential harm that can result from AI use.

There are many things to consider when drafting AI policy including data privacy, racial justice, civil rights, criminal activity, unlawful discrimination, misinformation/disinformation, deepfakes and other AI-generated content such as child sexual abuse material, intellectual property rights, copyright infringement, the effects on learning and education, physical safety, and the environment/climate change. AI consumes massive amounts of resources and energy and has been shown to be as harmful as it is helpful, creating concerns about how this technology can/should be used and what accountability looks like.

One advantage that policymakers in Virginia have comes from the resources and experiences of other policymakers. Several states have policies that can be used as templates that Virginia policymakers build upon, and policymakers in other states have offered to share their experience and expertise in crafting legislation.

Legislative Recommendations

Recommendation 1: Codify VITA's AI Utilization Policy

The General Assembly may wish to consider codifying VITA's current AI Utilization Policy, requiring all public bodies that wish to use AI technology to adhere to the policy's standards. As AI technology changes and develops, VITA's AI Utilization Policy may change. By codifying the current policy and its standards, the General Assembly will create a consistent minimum standard of regulation that public bodies wishing to use AI must meet, independent of future changes. The AI Utilization Policy, if codified, may serve as a scaffolding upon which to build out more detailed regulations as AI technology and applications change.

One element to review when codifying VITA's AI Utilization Policy is its mandatory approval process: this process requires considerations closely resembling the NIST AI RMF. Because of this, the General Assembly may wish to modify the approval process considerations to match the NIST AI RMF, which is a nationally accepted framework, when codifying the AI Utilization Policy.

Recommendation 2: Establish an Advisory Committee on AI

The General Assembly may wish to consider establishing an Advisory Committee on Artificial Intelligence to monitor developments in AI technology and make legislative recommendations to the General Assembly accordingly. Should an Advisory Committee be established, membership should include members of the General Assembly, civilian experts in AI technology and applications, and the Chief Information Officer of the Commonwealth or his designee. This Advisory Committee should be encouraged to share resources and collaborate with the Governor's AI Task Force and VITA where appropriate.

Recommendation 3: Regulate AI use by private and public entities

The General Assembly may wish to consider legislation regulating the use of AI technology by private and public entities. In the 2024 session, HB747 and SB487 both included language to accomplish this goal. Should the General Assembly craft legislation regulating AI use, a risk-based approach may be of interest. By using a risk-based approach, the General Assembly can ensure technologies that pose minimal risk to individuals and their rights (e.g., spell check, spam filters, tools like Grammarly) are not unnecessarily regulated while technologies that may pose a risk (e.g., algorithms used to determine whether individuals receive public benefits) require a level of human oversight. The General Assembly may consider risk-based frameworks such as the NIST AI RMF as guidelines for establishing requirements for developers and deployers of AI technology in the private sector. The General Assembly may also consider that software companies are not always forthright with their incorporation of AI technology in their Terms & Conditions, which may create concerns for consumers and any public bodies using that technology.

Recommendation 4: Strengthen data privacy regulations through an opt-in mechanism

The General Assembly may wish to consider modifying Virginia's current data privacy laws to give consumers the ability to opt-in, rather than opt-out. Currently, consumers have the ability to opt-out of data collection practices and request their data be deleted by companies. However, consumers may not know of these protections, understand their purpose, or have the ability to implement opt-out requests. Creating an opt-in system would maximize consumer protections by prohibiting companies from using consumers' data without their explicit informed consent. Such an approach has not been taken by states with data privacy laws, which would make Virginia a leader in this approach.

Appendix A: Definitions of AI

This is a non-exhaustive list of definitions for artificial intelligence, as used by a variety of bodies across the country and the world.

International Definitions

Organisation for Economic Co-operation and Development [2024 updated definition](#): An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.

European Union [2018 definition](#): “Artificial intelligence (AI) refers to systems that display intelligent behavior by analyzing their environment and taking actions – with some degree of autonomy – to achieve specific goals. AI-based systems can be purely software-based, acting in the virtual world (e.g. voice assistants, image analysis software, search engines, speech and face recognition systems) or AI can be embedded in hardware devices (e.g. advanced robots, autonomous cars, drones or Internet of Things applications).”

European Union [2020 study on what AI is definition](#): AI refers to systems that display intelligent behavior by analyzing their environment and taking action – with some degree of autonomy – to achieve specific goals

European Union [2022 glossary](#), based on language from a [2021 proposal to establish rules around using AI](#): ‘Artificial intelligence system’ (AI system) means software that is developed with one or more of the techniques and approaches listed in Annex I and can, for a given set of human-defined objectives, generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with

ANNEX I

ARTIFICIAL INTELLIGENCE TECHNIQUES AND APPROACHES

referred to in Article 3, point 1

- (a) Machine learning approaches, including supervised, unsupervised and reinforcement learning, using a wide variety of methods including deep learning;
- (b) Logic- and knowledge-based approaches, including knowledge representation, inductive (logic) programming, knowledge bases, inference and deductive engines, (symbolic) reasoning and expert systems;
- (c) Statistical approaches, Bayesian estimation, search and optimization methods.

European Union [2024 AI Act](#): AI system’ means a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments

Definitions used by US Federal bodies

National Institute of Standards and Technology AI Risk Management Framework ([NIST AI RMF](#)) definition:

[the RMF refers to an] AI system as an engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy (Adapted from: OECD Recommendation on AI:2019; ISO/IEC 22989:2022).

Office of Management and Budget [2020](#) and [2024](#) definitions, as [defined in federal legislation](#):

Section 238(g) of the John S. McCain National Defense Authorization Act for Fiscal Year 2019, Pub. L. No. 115-232, 132 Stat. 1636, 1695 (Aug. 13, 2018) (codified at 10 U.S.C. § 2358, note), defined AI to include the following:

- (1) Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets.
- (2) An artificial system developed in computer software, physical hardware, or another context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action.
- (3) An artificial system designed to think or act like a human, including cognitive architectures and neural networks.
- (4) A set of techniques, including machine learning, that is designed to approximate a cognitive task.
- (5) An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision-making, and acting.

[Executive Order 14110 \(15 USC 9401\)](#): The term "artificial intelligence" means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to-

- (A) perceive real and virtual environments;
- (B) abstract such perceptions into models through analysis in an automated manner; and
- (C) use model inference to formulate options for information or action.

Definitions used by US states

Connecticut, as defined in [SB1103](#): "Artificial intelligence" means (A) an artificial system that (i) performs tasks under varying and unpredictable circumstances without significant human oversight or can learn from experience and improve such performance when exposed to data sets, (ii) is developed in any context, including, but not limited to, software or physical hardware, and solves tasks requiring human-like perception, cognition, planning, learning, communication or physical action, or (iii) is designed to (I) think or act like a human, including, but not limited to, a cognitive architecture or neural network, or (II) act rationally, including, but not limited to, an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communication, decision-making or action, or (B) a set of techniques, including, but not limited to, machine learning, that is designed to approximate a cognitive task

Vermont, as defined in [H410](#): "Algorithm" means a computerized procedure consisting of a set of steps used to accomplish a determined task

“Automated decision system” means any algorithm, including one incorporating machine learning or other artificial intelligence techniques, that uses data-based analytics to make or support government decisions, judgments, or conclusions.

“Automated final decision system” means an automated decision system that makes final decisions, judgments, or conclusions without human intervention.

“Automated support decision system” means an automated decision system that provides information to inform the final decision, judgment, or conclusion of a human decision maker

California, as defined in [AB302](#): “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 decision-making and materially impacts natural persons. “Automated decision system” does not include a spam email filter, firewall, antivirus software, identity and access management tools, calculator, database, dataset, or other compilation of data.

“High-risk automated decision system” means an automated decision system 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.

Texas, as defined in [HB2060](#): “Artificial intelligence systems” means systems capable of:

(A) perceiving an environment through data acquisition and processing and interpreting the derived information to take an action or actions or to imitate intelligent behavior given a specific goal; and
(B) learning and adapting behavior by analyzing how the environment is affected by prior actions.

Indiana, as defined in [S150](#): “artificial intelligence” has the meaning set forth in [IC 4-13.1-5-1](#): “artificial intelligence” means computing technology that is capable of simulating human learning, reasoning, and deduction through processes such as:

- (1) acquiring and analyzing information for the purpose of improving operational accuracy through improved contextual knowledge;
- (2) identifying patterns in data; and
- (3) improving operational outcomes by analyzing the results of a previous operation and using the analysis to modify the operation to achieve an improved result.

Maryland, as defined in [SB818](#): “artificial intelligence” means a machine-based system that: (1) can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments;

(2) uses machine and human-based inputs to perceive real and virtual environments and abstracts those perceptions into models through analysis in an automated manner; and

(3) uses model inference to formulate options for information or action.

(1) “high-risk artificial intelligence” means artificial intelligence that is a risk to individuals or communities, as defined under regulations adopted by the department in consultation with the governor’s artificial intelligence subcabinet.

“High-risk artificial intelligence” includes rights-impacting artificial intelligence and safety-impacting artificial intelligence.

“Rights-impacting artificial intelligence” means artificial intelligence whose output serves as a basis for decision or action that is significantly likely to affect civil rights, civil liberties, equal opportunities, access to critical resources, or privacy.

“Safety-impacting artificial intelligence” means artificial intelligence that has the potential to meaningfully significantly impact the safety of human life, well-being, or critical infrastructure.

Utah, as defined in [SB149](#): "Generative artificial intelligence" means an artificial system that:

- (i) is trained on data;
- (ii) interacts with a person using text, audio, or visual communication; and
- (iii) generates non-scripted outputs similar to outputs created by a human, with limited or no human oversight.

Colorado, as defined [SB24-205](#): "artificial intelligence system" means any machine-based system that, for any explicit or implicit objective, infers from the inputs the system receives how to generate outputs, including content, decisions, predictions, or recommendations, that can influence physical or virtual environments.

Definitions used in Virginia

[VITA](#) definition of AI: "The simulation of human intelligence processes by machines, especially computer systems, such that it can adapt and learn on its own using machine learning algorithms that can analyze large volumes of training data to identify correlations, patterns, and other metadata that can be used to develop a model that can make predictions or recommendations based on future data inputs."