



Legislative Council Staff

Nonpartisan Services for Colorado's Legislature

Memorandum

September 25, 2025

TO: Interested Persons

FROM: [Matt Bishop](#), Principal Fiscal Analyst

SUBJECT: Overview of Artificial Intelligence

Summary

Artificial intelligence (AI) is an emerging technology that has the potential to reshape the economy. Although AI research has been conducted for decades, products are now proliferating due in part to advances in computing power. This memorandum explores basic topics in contemporary AI. It describes several limitations and misconceptions about AI and discusses AI regulation in Colorado.

AI Basics

AI programs, or models, differ greatly from traditional computer programming in how they solve problems. In traditional computing, a programmer writes a series of instructions based on what inputs are available and what outputs are desired. The program is a shortcut, using the computer to quickly perform the tasks defined by the programmer. A calculator or word processor are examples of traditional computer programs.

In AI models, the programmer provides sample inputs and a method for evaluating outputs to a machine learning program (such as a neural network), which performs its own analysis through iteration. Once the model is sufficiently trained, it can use its analytical training on new inputs to deliver new outputs. This allows AI models to address problems that the programmer cannot solve, and therefore can be useful for problems with which traditional computer programs may struggle. This includes large dataset analytics, pattern recognition, and language processing. Although it is a powerful tool, it has some limitations and drawbacks. A few of these are described below.



The Black Box Problem

The analysis that AI models perform are opaque, even to the programmer. What goes on between the inputs and outputs is a “black box” where the programmer cannot see what the model is doing and the model cannot explain afterwards how it achieved the outputs. It can be challenging to evaluate whether an AI model gets the right answer for the right reason. For example, researchers discovered that an AI model trained to detect cancerous tumors from photographs was actually [looking for rulers](#): the images of real tumors in the training images contained a ruler to indicate the tumor’s size.¹ This made the model successful when evaluated on the training data but not as useful for detecting new tumors. If a model gets the right answer for the wrong reason, it may deliver inaccurate results when used on new data.

Bias

One goal of using computer programs—both traditional software and AI models—is that they can theoretically provide objective perspectives and root out implicit bias, such as in hiring or health care. AI models require large amounts of training data. If programmers use real-world data, the model can learn and propagate existing biases, and the black box problem can make those biases difficult to detect. Therefore, AI developers require not only large quantities of data, but high-quality data. In some cases, developers clean existing data sets to address these concerns before training the AI models. In other cases, developers can generate artificial data sets that meet their quality criteria.

Hallucination

AI models create outputs based on what programmers and users say they want, but the systems cannot understand the content of what they provide and therefore may “hallucinate” content that appears valid but does not exist. For example, a [Colorado lawyer was disciplined](#) in 2023 for presenting a legal brief generated by AI because the cases it cited were not real.² The AI model knew what a legal brief sounds like, but it did not understand what legal citations are and invented them. Users of generative AI systems should be careful to review the output they receive.

¹ When AI Flags the Ruler, Not the Tumor, VentureBeat.

² [Disciplinary Judge Approves Lawyers Suspension for Using ChatGPT to Generate Fake Cases](#), Colorado Politics.



Uses of AI Systems

Existing AI models are better at solving some problems than others. Even as their use proliferates in our society and economy, they struggle to adapt to contexts without specific training.

Narrow Applications

Current AI models are best used to address narrowly defined tasks. The more specific the task, the easier the system is to train and the more reliable the outputs. AI systems perform poorly on problems for which they are not trained. For example, a model trained to predict the outcomes of baseball games may not do well at predicting the outcomes of elections, and vice versa.

Generative AI

Generative AI systems have significantly advanced in recent years. These are programs that can create text, images, videos, or other data from user prompts. Commonly known examples include chatbots, such as ChatGPT. While they are capable of more uses than previous systems, they still have narrow applications and are subject to the limitations described above, including bias and hallucination. These models continue to learn from prompts they are given, so it is important never to upload sensitive data to a generative AI system that you do not own.

Artificial General Intelligence

Also known as General AI, artificial general intelligence (AGI) possesses reasoning that meets or exceeds human capabilities. It is the type of AI depicted in science fiction stories such as *2001: A Space Odyssey* and *The Terminator*. Whereas existing AI models specialize in narrow tasks, AGI may be able to accomplish a wide variety of tasks under its own judgment. While technology companies are pursuing AGI, it may be years or decades before it is achieved, if ever. Concerns about a technology singularity or conflict between humans and AI are based on theoretical AGI, and do not pertain to existing AI systems.

AI for Consumer Products

Like computers generally, AI is very good at processing large amounts of complex data, especially for determining narrow outcomes. But as with other automated systems, it cannot replace human judgment and can struggle to incorporate unique circumstances or context.

Technology firms are competing with each other and with other countries to release products and gain market share. Computing power continues to increase, which means that a new



“generation” of AI software is always right around the corner and old products are quickly outclassed. Meanwhile, the question of how much testing is “enough” to ensure that systems are operating as intended and without adverse outcomes is subjective. These factors encourage firms to release products in a manner that shifts risks to consumers. A lack of industry standards around consumer protection and data privacy creates an opportunity for regulation. Yet firms are concerned that regulation that is applied unevenly in a global marketplace may impact competitiveness.

Environmental Impacts

AI outputs are computed “in the cloud,” which means they happen at large data centers owned by AI companies. These facilities, which are rapidly deploying as AI companies grow, have drawn criticism for their environmental impacts, especially their energy and water usage. Data centers use energy both for the servers themselves and to cool the facilities; these high-powered computers generate lots of excess heat and must be kept cool. Cooling also requires water, and the size of the data centers can strain local water resources. AI companies often invest in renewable energy or water-recycling technologies to reduce these impacts. However, critics warn that rollout of sustainability infrastructure lags behind the pace of data center deployment.

AI Regulation in Colorado

The Colorado General Assembly passed [Senate Bill 24-205](#), the first AI consumer protection bill in the country, and extended the effective date of that bill’s requirements to June 30, 2026, through [Senate Bill 25B-004](#). The bill establishes requirements of companies that develop or deploy AI systems, mechanisms for consumers to challenge certain outcomes from AI systems, and penalties for violations.

Responsibilities of AI Developers

Those who develop or deploy certain AI systems must use reasonable care to protect consumers from the risks of using these systems. This includes maintaining documentation on the intended purpose of the AI system, the type of data used to train it, and how it was evaluated to mitigate discrimination. If the developer or deployer discovers a risk of discrimination, either through their own testing or from a credible report, they must report the risk to the Attorney General within 90 days.



Consumer Appeals

The bill affects AI systems used to make “consequential decisions,” which affect the provision or cost of:

- education or employment opportunities;
- financial or legal services;
- health care services;
- housing;
- essential government services; or
- insurance.

When a consequential decision is adverse to a consumer, they must receive information on why the decision was made, how AI contributed to the decision, the type and source of data used by the AI system, and an opportunity to appeal the decision. Any appeal must involve human review of the decision if feasible and in the best interest of the consumer.

Enforcement

Each violation of the bills requirements is a deceptive trade practice under Colorado’s consumer protection laws. The Attorney General may take legal action against an AI developer or deployer based on complaints they receive. Each violation carries a possible penalty of up to \$20,000. The bill does not create a private right of action.