

Dispelling Myths About Artificial Intelligence for Government Service Delivery

Artificial intelligence (AI) has been promised to revolutionize every aspect of society, [including government services](#). Amidst the excitement, it may be difficult to separate fact from fiction: What is AI exactly? How does it differ from other technologies, and how can it really improve government services? These questions are especially important as government agencies [allocate resources to execute contracts with private vendors](#) based on the promises of AI. This brief aims to dispel common misconceptions about AI and government service delivery by establishing a common understanding that:

1. There is not a shared definition of what AI is;
2. Some kinds of AI have made more technological progress than others;
3. AI is not inherently objective or fair; and
4. AI requires significant resources in order to function well.



There is Not a Shared Definition of What AI is

These days, it is [commonplace to call nearly any application](#) of data or technology “artificial intelligence,” no matter how sophisticated it might be. That means different actors might talk past each other by using the term AI instead of a more specific, concrete one.

Some technologies might be referred to as AI even though they do not involve any AI at all:

- *Data sharing*: Organizations, including public agencies, collect and exchange information all the time. Just transmitting, warehousing, and aggregating data might be a [smart thing for your agency to do](#) to accomplish its goals, but that does not mean it is artificial intelligence.
- *Computer programs*: Computer programs are great at doing mechanical, repetitive tasks, like automatically getting new copies of data from a database or performing a large number of mathematical operations. While the rote functions that computers perform assist greatly with deploying AI, the use of computers alone does not mean that AI is at play.
- *Descriptive statistics*: Descriptive statistics are numerical figures like mean, median, and standard deviation. Computing descriptive statistics is not really AI either, even though it might be really important to crunch those numbers as part of AI.

Examples are worth a thousand words, so consider the following hypothetical: the U.S. Department of Transportation (DOT) wants to accurately estimate the number of automobiles in the U.S. Here is how you could solve this problem using the techniques surveyed above:¹ DOT

- 1 As an aside, though the hypothetical problem here seems innocuous, there are still privacy concerns. In particular, when state Department of Motor Vehicles (DMV) share data with DOT, what information about car owners are they also sharing with it? How long does DOT retain that data for? How do they plan on protecting that data? If DOT plans on using satellite imagery, do they also plan on re-identifying who might own a particular car by looking at Census Bureau tract information?

uses *computer programs* to pull and aggregate vehicular registration data from state DMVs (*data sharing*) and count the number of vehicles, as well as compute averages by region (*descriptive statistics*).

Moreover, even when considering technologies that could count as AI, a different term might be more specific and informative. The word “AI” might mean one of the following approaches:

- *Data science*: Data science aims to identify and explain patterns in data. Others might call data analytical techniques “business analytics” or simply “data analysis.” Sometimes, statistical analysis also involves predicting unknown quantities, but predicting accurately takes a back seat to identifying trends. Data science requires a person to interpret the results and create a narrative between different pieces of information.
- *Machine learning (ML)*: Machine learning is about using old data to teach computers how to analyze new information. For instance, one common goal of ML is to predict what outcome is most likely, given a set of information. While some might use ML and AI interchangeably, ML is often considered by those in the computer science field to be a subset of AI.²
- *Deep learning*: Deep learning is a subfield of machine learning that has led to breakthroughs in image recognition, speech recognition, text generation, and image generation. Whenever computers tackle one of those problems, deep learning is very likely at work. The phrase “deep learning” is a term of art for how these kinds of models process inputs, but there is not anything inherently superior or more profound about how deep learning models “learn.” Depending on the task at hand, deep learning might perform better or worse.

We can also illustrate how DOT might use these techniques to count cars:

- *Data science and machine learning*: Some DMVs have missing data for some municipalities. DOT pulls as much data as it can from DMVs but also gathers demographic data for each municipality from the census. Data scientists can teach computers about the statistical relationship between demographic variables and the number of cars in a municipality. The end result is a model that predicts the number of cars for municipalities missing that information, so long as demographic variables are available.
- *Deep learning*: DOT works with partners to count the number of cars in satellite imagery. Because this is an image recognition task (recognizing the roofs of cars in overhead pictures), DOT has to use deep learning.

2 So what is the difference between data science versus machine learning? Here is the rundown: data science tries to explain the data, can be used on smaller datasets, and can not be performed automatically by computers. In contrast, machine learning just tries to make predictions about new or unobserved inputs, often requires large amounts of data, and can eventually be run automatically by computers (after a significant amount of manual building and expiration).



Some Kinds of AI Technologies Are Advancing More Quickly Than Others

All AI has benefitted from the information technology explosion, which resulted in the widespread adoption of computers and the internet. Improvements in information technology and data collection have enabled a wide variety of applications, like applying for government benefits online or analyzing gigabytes of data using a personal computer.

But the [most significant advancements in AI have come from deep learning](#). Deep learning has enabled technologists to make significant strides in, for example, image recognition and [playing board games like Go](#). Those two tasks are hallmarks of progress because they have been hard for computer scientists [to crack for decades](#). However, progress on these two tasks should not lead technologists and policymakers to believe that current AI techniques can solve any problem thrown at them. Predicting human behavior in particular remains a [hard problem that deep learning does not offer much help with](#). That is an important lesson for government agencies, since AI might not be much better than humans at tasks like determining [which defendants are more likely to commit crimes while released on bail](#).³



AI is Not Necessarily Objective or Fair

Another important lesson for the government is that AI is not necessarily [objective](#) or [fair compared to alternatives](#). One reason for this is that many uses of AI involve data, but data is inherently biased. This is especially true for government agencies that want to change historical trends in data like [student achievement gaps](#) or [unhoused rates](#).

Another reason is that choices about how to build AI can also introduce biases. The choice of data is one example, but another common pitfall is the extent to which measures of success account for all individuals, not just the majority demographic group. If an agency evaluates an AI program based on how well it works on average for all individuals, that favors individuals who comprise the majority demographic group and might mask harms to certain communities.

Even the example of the Department of Transportation hypothetical might have biases. The deep learning approach uses satellite imagery to count cars, which does not account for cars hidden in parking garages. Hypothetically, those hidden cars might be owned more often by some demographics than others.



AI Requires Resources

Proponents of AI claim that it will [automate human tasks and require less resources](#), but deploying AI still requires significant resources that government agencies should be aware of. Agencies [may lack technical staff with expertise in AI](#), so they might have to [offload key decisions and judgments to partner vendors](#). As mentioned previously, making AI work well often requires data, which agencies might face [legal, ethical, or logistic obstacles in analyzing and sharing](#). Finally, deploying

³ There might still be other important applications of deep learning for government services though, like [offering chatbot services to guide applicants](#).

AI is not a one-off engagement and typically requires [continuous oversight](#). All of these additional considerations add substantially to the time and money required to deploy AI successfully.

In the context of the Department of Transportation hypothetical problem, additional resources are absolutely necessary: DOT needs to coordinate with state DMVs to get data, work with lawyers to ensure that all relevant laws and policies have been respected, and employ technologists who can process the data.



Spotting AI in Your Government Agency

In light of these ideas about AI and public service delivery, what should government agencies who are interested in AI do? The first step is to understand whether the suggested approach involves AI or not. Two key questions to ask are:

1. Does the application automatically make assessments, predictions, or judgments?
2. Does the application need to be “trained” or “tuned” on other data before it can be used?

If the answer to both of these questions are yes, then AI is likely to be involved. The agency might want to then do research on how appropriate the use of the technology is to the task at hand.

There are many questions that public agencies should answer before moving forward with using AI for service delivery. While not exhaustive, the following questions that agencies should consider include:

- How well can key information be captured in data?
- What potential biases might the AI tool create or reinforce?
- How much improvement can AI offer versus status quo or non-AI approaches?
- How will the AI tool be embedded in the agency decision making process?
- What role will human oversight continue to play?
- How will the agency ensure decisions are transparent and explainable to their clients, customers, and members of the public?

To address potential bias issues, the agency might want to ask for a bias audit or conduct their own. The agency should consider the overall cost of the AI application, including indirect ones like additional staff, staff training, ongoing evaluation tools, and governance mechanisms.

Readers might also be interested in the recommendations from other CDT reports, such as using [AI for identity verification](#) and [making government data publicly available](#). The recommendations in those resources can also broadly be applied to government usage of AI.

This is one in a series of information sheets designed by CDT's Equity in Civic Technology team to give practitioners inside public agencies clear, actionable guidance on how to most responsibly use technology in support of the communities they serve. More info: <https://cdt.org/civic-tech-inventory>.