Written Testimony of Joe Mariani

Technology and innovation research lead, Deloitte Center for Government Insights

Deloitte Services LP

United States House of Representatives Select Committee on the Modernization of Congress

July 28, 2022

Deloitte Center *for*
Government Insights

Chair Kilmer, Vice-Chair Timmons, and members of the Committee, thank you for the opportunity to testify today. I am Joe Mariani and I lead research into emerging technology for Deloitte's Center for Government Insights. I bring a broad range of experiences to this role — ranging from high school teacher to Marine Corps intelligence officer to government consultant to commercial technology researcher. As a result, this task of mining the breadth of industry and academia for big ideas that can improve the mechanics of government is exactly what gets me out of the bed in the morning.

Breadth is also important to Deloitte, with more than 120,000 employees in the United States working in nearly every industry – including more than 25,000 dedicated employees who help federal, state, and local government agencies advance their strategic initiatives and critical missions.

One product of that breadth is our ability to take solutions from other industries or academia and begin to understand in detail how they may work in government. For the past five years, we have studied the potential impact of Artificial Intelligence (AI) on government. We have looked at everything from how much time AI could save workers in each Federal agency to the rate of AI adoption in Federal, state, and local governments.[1,2]

From that research, we are confident AI has the potential to transform the legislative branch.

AI and human judgement as a team

Understanding how AI could benefit the legislative process starts with understanding AI's strengths and weaknesses.

AI is a powerful tool. It can handle massive volumes of data, run hundreds of thousands of repetitions, and do so with extreme speed and precision. You just have to take one look at my hand-written notes to know that precision, volume, and repetition are not traits at which most humans excel.

But AI does not, and cannot, replace human judgement. Although AI is a powerful tool, it still is only a tool that does what it is told. AI struggles to deal with tasks that can change dramatically or circumstances outside its programming. And, importantly for the purposes of legislating, AI cannot make value judgements.[3] AI can calculate the fastest, cheapest, or largest solution to a problem. But it cannot tell you if that solution is good or bad, right or wrong, desirable or undesirable.

Since good legislation needs both precision and judgements about what outcomes are desirable, AI can be an important addition to the process. The key is to think about humans and AI tools as a team, with each playing to their strengths.[4]

Applying AI to the legislative process

Like forming any team, the first step is defining the task at hand. After all, a football team is going to have a very different composition and practices than a chess team. Creating human-machine teams is no different. What you want to assess will change the type of AI tools and the governance structures needed – as well as the potential pitfalls that the human-machine team will face.

While AI can help many different areas of the legislative process — from AI assistants answering members' questions about legislation to Natural Language Processing (NLP) analyzing the US Code for

contradictions (which we know the Committee has already looked at) — we have identified two key applications of AI to help Congress:

1. AI as a microscope: Assess the impact of existing legislation

Legislation can impact every aspect of American life. However, the size and complexity of issues tackled by Congress can be difficult for even a group of experts to understand the impacts of a particular program or piece of legislation.

The same broad scope and volume of data that make assessing legislation a difficult problem for humans makes it an ideal challenge for AI. Machine Learning (ML) models can find patterns in inputs and outputs without having to specify ahead of time how those inputs and outputs are likely to be linked. Think of these type of ML models like a microscope. Just like a microscope can examine a leaf to find structures and patterns invisible to the human eye, so too can ML models find patterns in the outcomes of programs that may be invisible to humans.

There are already examples of ML models examining public policy in exactly this way. For several years, many different researchers have been using ML to understand the risk factors for infant mortality in childbirth. With all the data available in Electronic Health Records, many of these models can predict the likelihood of complications with 95% or greater accuracy.[5] Researchers from RAND then took those models the next step and, using data from Alleghany County, Pennsylvania, they used ML to evaluate what interventions had the biggest impact on reducing infant mortality.[6] The strength of ML models is that they are not only able to say what outcomes each intervention is likely to produce, but also among which groups. For example, researchers found that those mothers who used social services were less likely to use prenatal medical services and vice versa. These findings can then guide policy recommendations. Because both social and medical services are important to improving child health, policies aimed at building awareness of other services can have significant benefits.

While anything that improves the lives of infants is clearly a good policy outcome, other issues are less clear-cut. ML models can uncover hidden outcomes of policies or programs, but only humans can decide if those outcomes would qualify as successes or failures. In the spirit of human-machine teaming, once the ML model has uncovered the hidden outcomes of a program or piece of legislation, members or staff can then look at those outcomes and determine 1) if they are positive or negative and 2) if the overall benefits are worth the cost and difficulty of the program.

2. AI as a simulator: Test the potential impacts of future legislation

The ability of ML models to predict outcomes of policies begs the question: *What if we did something differently? How would things change?* In essence, this question seeks to create an AI "simulator" for a problem. Think of Apollo 13. When the crew suffer an explosion, they have to figure out new interventions, new ways of doing things. To see what would work and what would not, they used the ground-based simulator to try procedure after procedure until they found one that worked. Imagine having an AI simulation run through hundreds of thousands of possible interventions in minutes, instead of locking astronaut Ken Mattingly in a dark box for days.

In place of ML models trained on historical data and projecting trends into the future, these simulations are designed to capture the dynamics of complex systems like the economy or the health care system.

Simulations are based on models for how a portion of a complex system operates. For example, one form of simulation well-suited to legislative tasks are agent-based models that replicate how individual actors would respond to and interact in different situations. These models are good at capturing the "emergent properties" of complex systems where individual decision add up in unusual ways. Think about flocking birds: each bird just tries to stay close to the bird next to them, but all together, they make intricate patterns in the sky as they avoid obstacles and predators. The big human systems that Congress is often interested in like healthcare, the economy, or national defense exhibit similar traits.[7]

Researchers in Europe have created an agent-based model designed to help policymakers understand the likely impacts of different interventions on the Irish economy. The Innovation Policy Simulation for the Smart Economy (IPSE) uses data from patents, knowledge flows, and other economic data to model how individual companies and investors are likely to react to different policies.[8] For example, researchers can examine if different funding methods or tax incentives would help support the creation of new small businesses in a specific city or high-tech industry. Similar models could be a great benefit as Congress examines which policies could help spur domestic semiconductor manufacturing or other advanced technologies.

The benefits of policy simulation

In 1964 economist George Stigler said, "we do not know the relationship between the public policies we adopt and the effects these policies were designed to achieve."[9] ML models can help uncover the relationship between public policies and their effects. But that relationship is only helpful in making future policies if we assume the future is like the past. So, when we find ourselves in an era we know is different from the past – during a global pandemic, for example – or when we see a model based on historical assumptions drifting away from current data – such as with the widespread adoption of new technology, for example – then simulations become critical tools for understanding the likely outcomes of new public policies.

As the Apollo 13 simulator helped astronauts, AI simulations can help policymakers to:
- Uncover the drivers of a particular problem, whether that is the amperage limitations on the lunar module in Apollo 13 or the causes of regulatory noncompliance.[10]
- Understand which interventions could be effective, whether sequencing systems start-up for Apollo 13 or organizing national airspace to allow for more on-time flights.[11]
- Understand the trade-space of a given issue. Of all the effective interventions, how much is required, at what cost, and to achieve what outcome?[12]

Yet even these complex AI simulations cannot make value judgements. They cannot determine the best option. They can only assess the optimal choice for the given values and assumptions that humans specified at the start. However, by forcing the human side of the human-machine team to be specific about those values and assumptions, AI simulations may hold the potential to transform legislative processes.

For example, modeling the impact that of different policies on economic development can help validate the assumptions that undergird our positions. We may find that we assume in the model that government research and development (R&D) spending will crowd out private R&D in that industry. But this assumption can be tested, providing ground for more constructive debates. Similarly, simulations can help uncover human values that may not be well-articulated in a policy debate. For example, running a simulation to optimize economic growth may yield undesirable consequences leading

members to realize that, while economic growth is a goal, it may only be desirable when it improves living conditions for the public in a particular area. Testing assumptions and uncovering hidden values can help provide a more firm foundation for data-driven policy debates.

In fact, there is evidence that experimenting with models in itself may help drive consensus.[13] As members examine the values, potential interventions, and trade-space of a topic, they are likely to see more of the factors that they agree on, rather than the few on which they do not. This will certainly not bridge all ideological divides, but it can offer fertile ground for productive debate on evidence-based policies.

Potential challenges

While human-machine teaming has the potential to bring transformational benefits to the legislative process, it is not without risk. However, by focusing on the tasks we want those teams to perform, whether that be a microscope-like assessment of existing legislation or simulator-like testing of potential legislation, we can carefully control for those risks while still realizing the transformational benefits.

Data and model governance

AI's outputs are only as good as the reliability of the model and the accuracy of the data. If the data is not accurate and fit-for-purpose or the model is not robust and explainable, it can create significant issues for the privacy, security, or fairness of an AI tool. There are already several frameworks for understanding and managing the risks of using AI in government. The National Institute of Standards and Technology and US Government Accountability Office have issued several important reports on the topic, and organizations like the Department of Defense are already operationalizing much of that guidance.[14,15,16]

These guidelines are by no means one-size-fits-all. The unique tasks that any given model performs can lead to different challenges that require different controls. For example, the central role played by historical data in the "microscope-like" use of ML to assess existing programs means that those **ML models need clean, accurate data that is matched to their task**. Open public data can help ensure the availability of good data.[17] Similarly, tagging data sources with the use cases for which they are suitable can help avoid instances where data gathered in one context is used in another. For example, data that is representative for race and gender may not be representative for income level, so it should not be used in models where that is an important parameter.[18]

While ML models are based on historical data, AI simulations are primarily based on assumptions of how factors relate to each other – whether that is how individuals will react to a given choice or how smoking rates vary with rates of physical activity.[19] These assumptions can and should be based on real data, but they are still assumptions and there is never a guarantee that the future will look like the past. Therefore, when we see models based on historical assumptions drifting away from expectations, it can be a sign that those assumptions need to be adjusted to better match a changing world. For example, many mass transit models assumed only a few major modes of transit such as car, bus, train, or bike. However, the sudden, massive growth of e-scooter ridership in 2018 and 2019 would have altered these models, forcing them to reevaluate assumptions about how people get around.[20]

As a result, **AI simulations need** special attention paid to **transparency of assumptions**. Model governance procedures can help ensure the transparency of assumptions so human team members can understand the context by which the simulation reached its conclusions.

Security

The important role played by legislation means that security needs to be a prime consideration for any tool that can influence it. Whether it is used to assess or shape legislation, AI tools need protection beyond typical cybersecurity considerations. The potential for adversaries to manipulate the outcomes of these AI models to tip policymaking to their advantage calls for careful safeguards.[21]

The central importance of data to "microscope-like" ML models means that they can be particularly vulnerable to the poisoning of training data—that is, tampering with data used to train ML models with a goal of having those models produce undesirable results.[22] As a result, **ML models used to assess the impact of legislation need controls on the access to and quality of data**. On the other hand, **AI "simulator-like" models need safeguards placed on the variables, assumptions, and even outputs of the models** to avoid manipulation.

New processes, new skills, new training

The AI portion of the human-machine team is not the only aspect of the partnership that needs attention. Introducing new tools to the legislative process will require human team members to learn new skills, adapt to new processes, and work together in new ways.

For example, as "microscope-like" ML models uncover new outcomes of public policies, members of Congress will quickly find themselves consuming new types of data beyond bar charts and budget trends. New forms of information such as geospatial data, statistical relationships, and more are likely to become important to decision making. To ensure these new sources of information are easily consumable, members and staff may **need new data visualization tools**. Similarly, the ability to analyze, create, and present those visualizations will require greater data science skills for staff members.

The "simulator-like" AI models may bring even more radical changes. In place of an "analyze-then-present" form of giving information to decision makers, these types of models can allow for **real-time decision support** where members can sit side-by-side with staff to adjust models and examine conclusions as new data comes in. This shift has already taken place in industries such as auto racing, where Formula One race teams adjust strategy models in real-time based on thousands of data points as cars are racing around the track.[23] The shift to this mode of decision support can bring significant changes to how staffers spend their time. When Deloitte applied this concept to prototypical analysts in the intelligence community, for example, our model suggested that analysts could spend up to [39 percent more time advising decision-makers](#) with the adoption of AI at scale.[24]

The way forward

Implementing AI in the legislative process can seem like a seismic transformation. But the shift is possible with the right commitment and investment. Similar AI models are already at work in other industries and even other parts of government. However, the experiences of other industries highlight that while the change is eminently possible, it will take considerable leadership. Leadership not just to

put the technologies in place, but also to incorporate the training, education, and business practices needed to make them work for Congress and the American people.

Lessons learned from other industries can help Congress get started on its AI journey:

Don't try to model everything

The scale of the issues that Congress tackles is often tremendous and trying to model every aspect of each issue is impossible. The Formula One example shows that even relatively simple models can quickly get out of hand. For a single race, there are more race outcomes possible than there are electrons in the universe.[25] This is where the human part of the human-machine team can help. Rather than trying to model everything, using human value judgements prior to modeling can help identify the core aspects of the problem that need to be modeled. In short, it all starts with deciding what the problem is and understanding what is important. Then the tech can get to work.

Make a platform, not a solution

As the controller of the nation's finances, Congress also has a financial duty to the public. How can Congress get the most out of AI without having to build a new tool from the ground up for every new policy debate? The answer is to build an AI-enabled platform, rather than a single point solution. This is the approach Singapore took with its Virtual Singapore 3D model of the city-state. Virtual Singapore not only models the 3D layout of the city, but it also allows for hosting all other manner of data sources such as census and geospatial data.[26] That way when a new problem emerges, developers can simply create a new app within Virtual Singapore to run simulations about the new issue. Such an approach would allow Congress tap into AI in a way that is cost-effective, efficient, and able to evolve over time as technology changes.

Invest in the human dimension

Finally, the human element of the human-machine team is critical to long-term success of digital transformations. Leaders must pay attention to the new tasks that may take up more time, new skills that may require difficult retraining, and even new career paths that may change employees' life goals. Taking care of the people will help take care of the technology.

AI is a powerful tool for the assessments and simulations that Congress needs in its legislative processes. Pairing AI with the right people and the right processes can help provide common foundation for debate, encourage consensus, and deliver meaningful results for the American people.

Endnotes

[1] https://www2.deloitte.com/us/en/insights/focus/cognitive-technologies/artificial-intelligence-government-analysis.html

[2] https://www2.deloitte.com/us/en/insights/industry/public-sector/artificial-intelligence-impact-on-future-intelligence-analysis.html

[3] https://link.springer.com/article/10.1007/s41649-019-00096-0

[4] https://www2.deloitte.com/us/en/insights/focus/human-capital-trends/2020/human-ai-collaboration.html

[5] https://bmcpregnancychildbirth.biomedcentral.com/articles/10.1186/s12884-022-04594-2

[6] Peet, Evan D., Dana Schultz, and Susan L. Lovejoy, Using an Innovative Database and Machine Learning to Predict and Reduce Infant Mortality. Santa Monica, CA: RAND Corporation, 2021. https://www.rand.org/pubs/research_briefs/RBA858-1.html.

[7] https://www.annualreviews.org/doi/10.1146/annurev-publhealth-040617-014317

[8] https://www.nesta.org.uk/feature/smarter-policy-through-simulation/innovation-policy-simulation-for-the-smart-economy/

[9] https://www.chicagobooth.edu/review/danger-making-policy-based-assumption

[10] https://www.technologyreview.com/2020/05/05/1001142/ai-reinforcement-learning-simulate-economy-fairer-tax-policy-income-inequality-recession-pandemic/

[11] https://www2.deloitte.com/us/en/insights/focus/tech-trends/2020/digital-twin-applications-bridging-the-physical-and-digital.html

[12] Cerdá M, Tracy M, Keyes KM. 2018. Reducing urban violence: a contrast of public health and criminal justice approaches. Epidemiology 29(1):142–50

[13] https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2021EF002283

[14] https://www.nist.gov/system/files/documents/2022/03/17/AI-RMF-1stdraft.pdf

[15] https://www.gao.gov/products/gao-21-519sp

[16] https://media.defense.gov/2022/Jun/22/2003022604/-1/-1/0/Department-of-Defense-Responsible-Artificial-Intelligence-Strategy-and-Implementation-Pathway.PDF

[17] https://www2.deloitte.com/us/en/insights/industry/public-sector/open-data-ai-explainable-trustworthy.html

[18] https://www2.deloitte.com/us/en/insights/industry/public-sector/open-data-ai-explainable-trustworthy.html

[19] https://www.cdc.gov/pcd/issues/2021/20_0225.htm#:~:text=The%20Prevention%20Impacts%20Simulation%20Model%20(PRISM)%2C%20a%20system%20dynamics,of%20community%20initiatives%20to%20reduce

[20] https://www.sciencedirect.com/science/article/pii/S1361920921003941

[21] The Director of National Intelligence has included "minimizing the potential for adversarial influence" as one of the key design principles for AI in the Intelligence Community: https://www.dni.gov/files/ODNI/documents/Principles_of_AI_Ethics_for_the_Intelligence_Community.pdf For an introduction to adversarial AI, see the World Economic Forum's short primer: https://www.weforum.org/agenda/2018/11/what-is-adversarial-artificial-intelligence-is-and-why-does-it-matter/

[22] https://proceedings.neurips.cc/paper/2017/file/9d7311ba459f9e45ed746755a32dcd11-Paper.pdf

[23] Joe Mariani, "Racing the future of production: A conversation with Simon Roberts, operations director of McLaren's Formula One Team," *Deloitte Review* 22, January 22, 2018.

[24] https://www2.deloitte.com/us/en/insights/industry/public-sector/artificial-intelligence-impact-on-future-intelligence-analysis.html

[25] Formula 1, "What does an F1 strategist do?." YouTube video, 10: 29, posted January 9, 2020.

[26] https://www.nrf.gov.sg/programmes/virtual-singapore