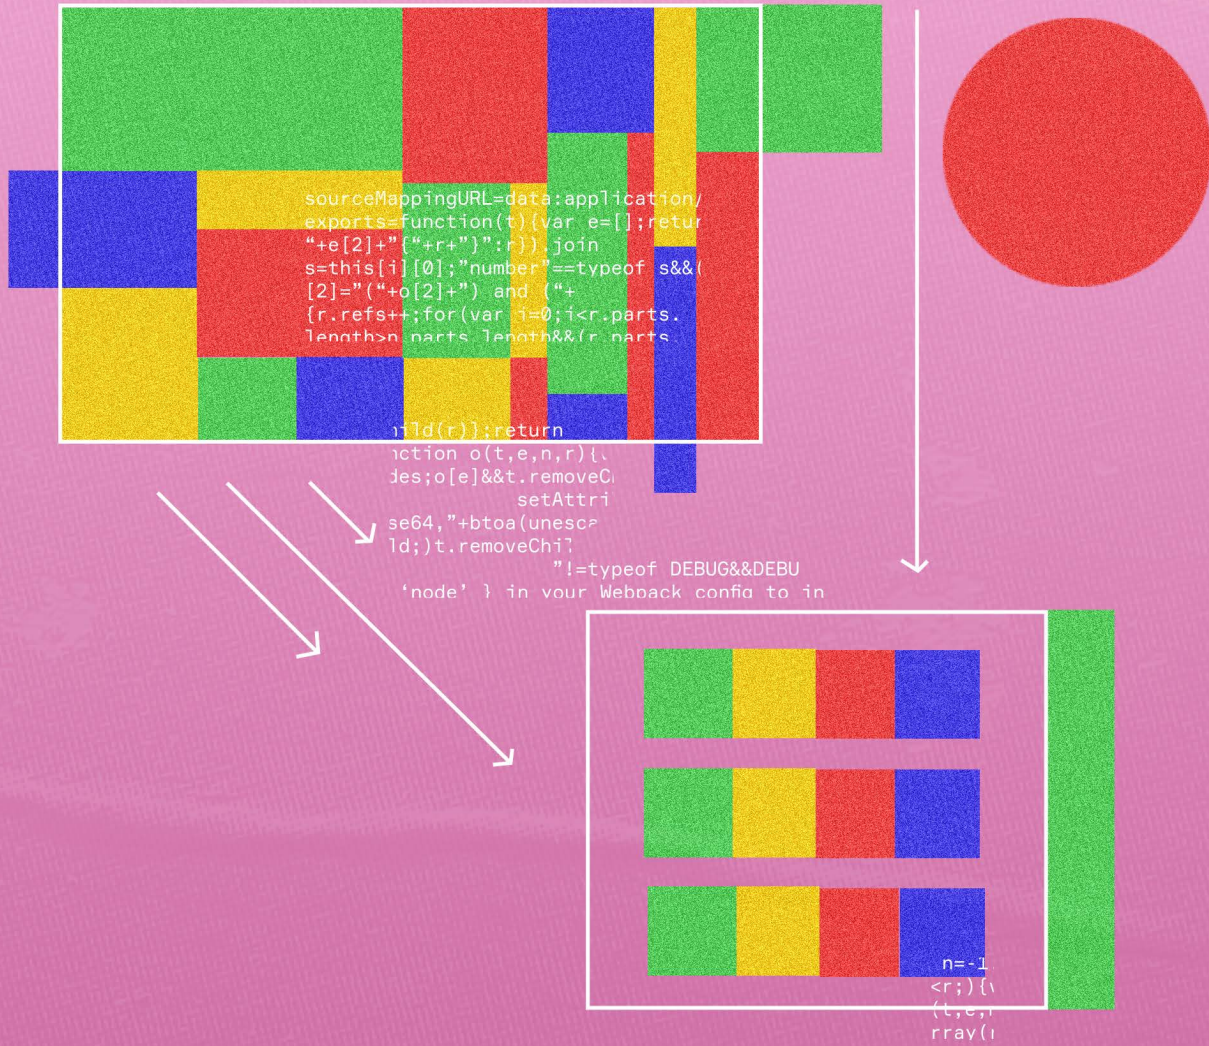


AI can turbocharge government regulatory operations

Armed with powerful AI tools, government regulators can change how they work and interact with businesses and citizens alike.

A report from the Deloitte Center for Government Insights



*02 . . . AI holds enormous promise to help transform
regulatory operations*

*03 . . . Three considerations for realizing AI's potential in
regulatory operations*

04 . . . So how can AI transform regulatory operations?

08 . . . Getting started

09 . . . Toward a future of better, faster regulatory processes

10 . . . Endnotes



AI holds enormous promise to help transform regulatory operations

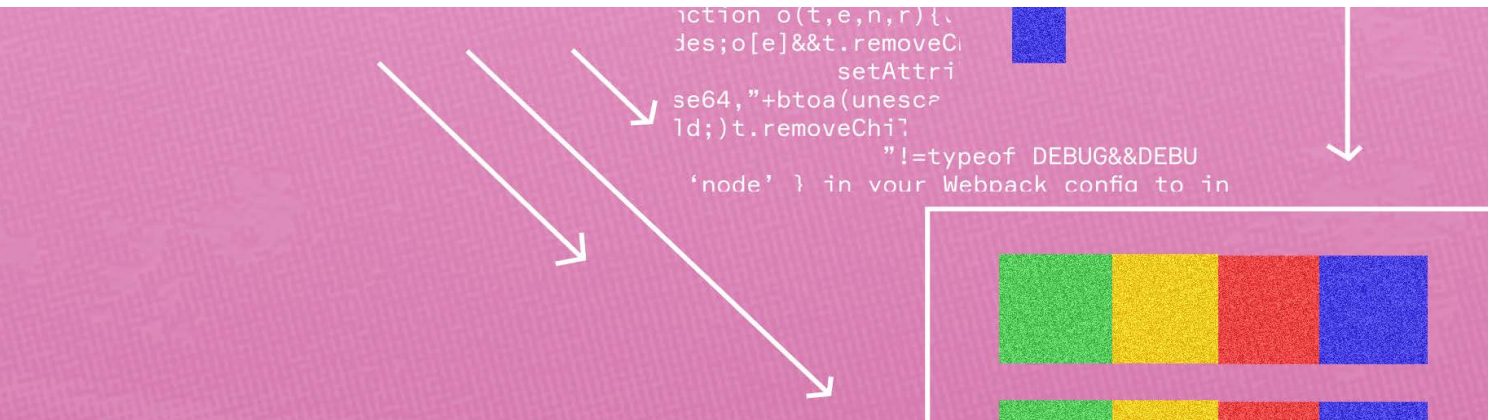
For decades, the economy has digitized. Airlines today feature more lines of code in their software than they have physical parts.¹ Films can be made on digital tape with completely virtual surroundings.² And financial transactions can be made so quickly that some stock exchanges have to slow them down with miles of fiber optic cable.³

But even as these and other industries experience the incredible speeds of digitization, they still need to operate safely. Airlines need to be safe, film copyrights need to be protected, and financial transactions need to be kept secure. Yet, the regulators charged with overseeing these industries often find themselves with outdated tools that can limit effectiveness. For example, we used artificial intelligence (AI) to analyze more than 217,000 sections of the 2017 US Code of Federal Regulations and found nearly 18,000 sections containing duplicate or similar passages.⁴

Modern regulators need modern tools. If AI can help identify those duplicate rules, it can create a huge opportunity to eliminate redundant or conflicting regulations. And that could just be the beginning. The recent emergence of large language models that can process vast amount of text and other forms of generative AI bring powerful new tools to bear that can produce coherent, contextually relevant, and arguably even artistic outputs.⁵

Armed with these powerful AI tools, regulators can change how they work and interact with businesses and citizens alike.⁶ For example, generative AI can improve interactions with citizens and businesses, analyze and summarize large volumes of stakeholder input, automate administrative tasks like report generation, code software solutions, and even suggest tailored solutions.

While AI isn't the answer to every problem, when applied to the right tasks and paired with appropriate human judgment, AI, can be a big part of the solution.



Three considerations for realizing AI's potential in regulatory operations



To harness generative AI's potential, regulatory organizations should consider the following steps: 1) understand the unique capabilities of different AI tools; 2) employ multiple tools tailored to unique tasks, and; 3) adapt business processes to integrate AI and human judgment. Embracing these principles can help regulatory agencies keep pace with the industries they oversee by speeding up regulatory processes, automating internal processes to reduce burden on staff, and improving compliance rates through data-driven decision-making.

Understand the unique capabilities of different tools

Different AI tools work differently. Generative AI, for example, can do creative tasks that other forms of machine learning cannot, but at the cost of some degree of accuracy—for example, the hallucinations you may have heard about where generative AI responds with plausible, but incorrect facts. On the other hand, more traditional forms of machine learning can discern hierarchical relationships from data, but can't easily show why it reached a certain conclusion.

This interplay of strengths and weaknesses means that some AI tools may be better suited than others for different tasks. Tasks requiring repetitively creating content, like transcribing hearing notes or generating reports or analyzing user sentiment on a call, could be given to generative AI. Tasks involving finding patterns in large volumes of data with a high degree of precision, like identifying duplicative legislation or detecting fraud, could be assigned to traditional machine

learning technologies. Finally, the strengths of human judgment in dealing with high variability or social components means that AI should augment human expertise. Applying AI tools can free up humans to accomplish more challenging tasks more efficiently.

This is similar to a system already in place in Denmark. The Danish Business Authority uses AI to do the computationally heavy task of analyzing more than 230,000 financial statements a year for possible fraud and inconsistencies.⁷ But once an issue has been found, determining if fraud has occurred requires awareness of content that human judgment should provide.

Look for multiple tools tailored to their tasks

Because different AI models do different tasks, adopting the most sophisticated and most expensive tool may not yield the best results. Rather, the best results are likely to come from multiple, smaller tools, each playing to its particular strengths. Take AI augmented tax operations. An AI digital assistant can prepare tax returns in roughly five minutes.⁸ A natural language processing (NLP) enabled chatbot can ask tax-related questions to a taxpayer and autofill a tax return with the help of RPA. Finally, a taxpayer can review the tax return, validate the data, and then file the return. Further, simple tax returns can be done by the chatbot, leaving more complicated returns for humans.

A single generative AI tool would need extremely strict parameters for this task to avoid getting too creative with tax advice. The combination of different AI tools, each playing to their strengths, can generate results with a smaller footprint than a single massive tool.

Adapt business processes to integrate AI and human judgment

It can be counterproductive to layer new technologies onto regulatory processes built for the paper age. Agencies should rethink their regulatory processes with an understanding of what AI can deliver. A Deloitte survey found that when AI projects are accompanied by significant workflow changes, they are 36% more likely to succeed.⁹

An effective process can emerge by first determining the regulatory outcome and then identifying the most suitable technologies to accomplish the associated tasks. The Securities and Exchange Commission (SEC) has invested in a tool that incorporates NLP and machine learning to analyze regulatory filings of investment advisers. Instead of just aiming to make existing processes faster, the SEC

is trying to use this tool to do something new: to find patterns that could flag potential violations. Backtesting analyses found algorithms are five times better than random inspections of filings at identifying language in a filing that qualifies for referral to enforcement.¹⁰

Such a monitoring tool could support risk-based targeting of compliance activities by identifying where the likelihood or impact of noncompliance is the highest. In other words, this tool could allow regulators to work differently. Rather than waiting for violations and issuing penalties or remediations, regulators could proactively inspect filings.



So how can AI transform regulatory operations?

Even from these three considerations, it is clear that AI can have a tremendous impact on regulatory operations. AI can help improve the efficiency of existing tasks, such as the Danish Business Authority's use of AI to identify fraud. But it can also allow regulators to achieve desired outcomes in fundamentally new ways, such as the SEC's use of NLP and AI to detect infractions by investment advisers. AI can also help improve both efficiency and effectiveness across the entire regulatory life cycle—including public consultation, developing regulation, and enforcing regulations (figure 1).

Using AI for counter manipulation of public comments

Public comment periods are a major element of most regulatory decision-making. The importance and openness of public comment also, however, makes it vulnerable to manipulation.

When the Federal Communications Commission (FCC) proposed new rules for net neutrality in 2017, their online comment system received over 22 million comments. Researchers used tools like NLP and clustering

Figure 1

AI can improve both efficiency and effectiveness across the regulatory life cycle

The examples below show how AI could help transform regulatory life cycle

● AI improving effectiveness

● AI improving efficiency

<ul style="list-style-type: none"> ● Predictive public sentiment ● Stakeholder engagement chatbots 	<ul style="list-style-type: none"> ● Initial drafts of regulations ● Coding adaptive regulatory tools 	<ul style="list-style-type: none"> ● Auto-generating lists of inspection targets ● Generating response for approvals and rejections 	<ul style="list-style-type: none"> ● Assessing policy effectiveness and simulation ● Identifying and rectifying conflicting regulations
<ul style="list-style-type: none"> ● Screening AI-generated public comments ● Analyzing and summarizing input 	<ul style="list-style-type: none"> ● Identifying overlapping and duplicative regulation ● Simulating policy and regulatory scenarios 	<ul style="list-style-type: none"> ● Predictive and risk-based enforcement ● Detecting fraud 	<ul style="list-style-type: none"> ● Regulatory performance dashboards ● Tracking compliance rates
Consultation and engagement	Developing new regulations or legislation	Enforcing regulations	Reviewing regulations

Regulatory life cycle

Source: Deloitte analysis.

algorithms to determine that more than 1 million of those comments—maybe far more—were written by bots.¹¹ The FCC hired a contractor to search comments for repeated phrases, word use, and sentence structures. They found up to 1.3 million “unique” messages in which words had been swapped for synonyms within identical sentences and paragraphs.¹² Unlike a cut-and-paste email campaign, where real people copy a message composed by an advocacy group, these messages had each been

slightly changed, likely by a bot, to bypass filters and appear “unique.” One analysis found that only 800,000 comments were organically written rather than being spam messages or part of a larger email campaign.¹³

To help combat bot-generated pressure campaigns, leaders can sort through comments with a variety of tools. Reading all 22 million comments was infeasible for humans, but AI can cluster comments into groups,

essentially narrowing those millions of comments down to 25 or so bulk messages. The top 25 comments represented 98% of the comments in favor of repealing net neutrality.¹⁴ “Sentiment analysis” can sort comments into pros and cons, providing a rough poll.

AI can also allow regulators to reimagine the process of open comment periods. Spam detection can help filter out bot attacks and AI-generated comments. NLP and sentiment analysis could summarize legitimate comments. Generative AI can cluster their arguments, and perhaps even make inferences about the constituencies at play. Finally, generative AI can summarize findings of multi-stakeholder discussions to help allow policymakers to focus on the most important issues.

Using appropriate safeguards, AI can assess not just the content of comments, but also the metadata. Factors like operating system, IP address, time of submission, and browser all help identify spam. A data analytics firm found a batch of suspected FCC bot comments based on invisible line breaks called “\n strings.”¹⁵

AI can include such factors in analysis as long as the system takes care to protect privacy. Deloitte’s

trustworthy AI framework (see sidebar, “Six dimensions of trustworthy AI”), which demands transparency and limits on how data is used, can help AI assess relevant data without compromising privacy.¹⁶

Maintaining institutional memory

Regulatory agencies often rely on human memory. Regulators test regulatory strategies, close loopholes, and write regulations. They become experts in technologies, facilities, and industries. They learn how to provide a steady and predictable service while navigating government procedures. When a long-serving regulatory expert leaves an agency, it can mean a vast loss of knowledge. Turnover can also make it difficult for incoming regulatory and policy staff to hit the ground running, which could delay rulemaking or enforcement of regulations.

AI could help mitigate these problems. First, as employees leave, exit interviews can capture their tacit knowledge. Generative AI can then use those interviews to create onboarding and training documents tailored to a new hire’s specific job. Generative AI could be trained on existing legislation and regulations, and historical

SIX DIMENSIONS OF TRUSTWORTHY AI

Deloitte’s trustworthy AI framework lays out six key dimensions that can help build trust in AI. The framework is designed to help regulatory agencies identify and mitigate potential risks related to AI.

1. **Fair and impartial:** AI should be designed and trained to follow a fair and consistent process that makes fair decisions. It must include internal and external checks to reduce discriminatory bias, including in regulatory decisions like approving or rejecting permits or licenses.
2. **Transparent and explainable:** Users should understand how technology is being leveraged, particularly in making decisions. Regulatory agencies should

emphasize creation of algorithms that are transparent and can be explained to business entities who are being impacted by those algorithms.

3. **Responsible and accountable:** AI systems should have policies outlining who is responsible and accountable for system output or decision-making. This will likely become increasingly more important as generative AI is used in regulatory applications such as drafting legislation and summarizing stakeholder input on draft policies.
4. **Safe and secure:** For AI to be trustworthy, it should be protected from cybersecurity risks that could

manipulate the system and result in digital harm.

5. **Respectful of privacy:** Privacy is critical for AI since the sophisticated insights generated by AI systems often stem from data that is more detailed and personal. Trustworthy AI should comply with data regulations and only use data for the stated and agreed-upon purposes.
 6. **Robust and reliable:** AI should be at least as robust and reliable as the traditional systems, processes, or people it is augmenting or replacing. It should generate consistent and reliable outputs, especially as it is scaled.
-

enforcement decisions so existing staff and new hires can query chatbots to better understand complex policy documents and speed up regulatory research processes.

Improving effectiveness of inspections

AI offers myriad options for improving inspections throughout the process: prioritizing where to inspect, reducing paperwork during inspections, producing first drafts of reports, and uncovering insights into historical patterns.

The tools for targeting inspections are well-established. Some health departments use AI to determine which restaurants are most likely violating code.¹⁷ First, a program assesses risk factors using everything from 311 complaints about garbage to Twitter posts about food poisoning.¹⁸ These risk assessments help target inspections toward more likely offenders, resulting in infractions that are more than a third of a control group with random inspections.¹⁹ Researchers estimated the software could save 9,000 food poisoning incidents and upwards of 500 hospital visits a year.²⁰

Likewise, building inspectors in New York City have used an AI tool to determine which factors correlate to structure fires and assign every building a risk score. The tool helps focus inspections, much like the tool for health inspectors targeting potentially unsafe restaurants. It automatically lists buildings in order of priority, including incorporating lower-risk buildings with mandatory annual inspections, like schools. While the department started out considering about six risk factors, the New York City Fire Department's latest iteration tracks over 7,500 datapoints from 17 city agency data streams.²¹ The AI learns from each new fire, identifying new factors to consider, and reassesses risk in real time. This feeds a real-time dashboard of fire risk, which fire station leaders can use to plan. Similar algorithms could predict high-risk situations in nearly any regulatory domain, from workplace hazards to environmental spills.


Generative AI could come into place once inspections begin. A new permit for a manufacturing facility might take days for an air quality inspector to understand. Does the provision about gas boilers larger than X cubic feet running at night apply at this facility? Generative AI could summarize the permit into key points to inspect

and highlight pertinent sections. The inspector would still have to read it all, but at least they'd start with a summary. Here, again, is where institutional knowledge comes in. A chatbot with memory of past infractions from a specific facility can suggest where to look, and information gleaned from past inspectors can remind the new guy to check behind that one door where they keep all the chemicals. AI could help build a briefing form about every facility.

Apart from prioritizing inspections, generative AI can also help in documenting infractions, sending notices, and generating first drafts of inspection reports. Saving time on paperwork can add up. Based on an analysis of the US federal government workforce conducted in 2017, it's estimated that automating manual tasks through AI could free up tens of millions of regulatory staff-hours. For example, up to 60 million hours a year could be saved for activities related to compliance and enforcement operations alone. Another 26 million hours annually of inspectors' time could be freed up as well.²²

In the six years since this analysis was done, AI capabilities have vastly improved and, with the advent of generative AI, time savings likely would now be even higher. AI can now take on an even greater number of low-value tasks, while inspectors could work on high-value, cognitive tasks.

AI augmentation can allow inspectors to spend more time doing what they're supposed to—inspect high-risk operations.





Getting started

When thoughtfully paired with complementary AI tools and human judgment, generative AI can open up transformative possibilities for regulation development and enforcement. The path to that transformation features some significant challenges, but three considerations could help regulators improve their chances of success as they embark on their AI journeys:

Become familiar with underlying technology

Understanding the basic capabilities of different types of machine learning and generative AI can help leaders decide which technologies are appropriate for which tasks. The relative transparency of a traditional machine learning setup might be better for finding patterns in incident reports rather than a generative AI tool. Detailed knowledge about how the AI models work can also help mitigate challenges with privacy and security. For example, using “knowledge embedding” to restrict the scope of documents that an LLM will review can help improve accuracy and help keep queries out of the public domain. Knowing AI’s capabilities lets one imagine what is possible.

Put users at the center


Reimagined regulatory processes should work like any well-designed tech project: put the users at the center. Depending on the process that AI is augmenting, those users may be citizens, businesses, or even the regulators themselves. Centering around the user can help shape

good design. The basic principles of human-centered design apply here. First, focus on the outcome that users desire and discover what factors may be currently hindering that progress. Next, develop progressively more refined solutions that can bring technology to help reduce those challenges. Finally, find metrics that correlate to serving the mission and continually measure them.

Invest in developing trust in AI

While economies evolve, and regulators should evolve with them, regulations are often a slow turning ship. Balancing trust and stability with the benefits of a fast-moving technology can be a challenge for regulators.

Trust is important to regulators. When it comes to trusting AI, a central tenet of Deloitte’s trustworthy AI framework is transparency. Developing trustworthy AI will require transparency in two ways. First, transparency of the AI models themselves. The opacity of black box AI models, where inputs enter and a result emerges with no one certain what happened in the middle, can be concerning especially if your life or livelihood is on the line. Public trust will likely become stronger when there is an oversight mechanism that stakeholders can understand. So, if regulators are using AI to augment decisions, they should be sure to choose an algorithm with explainable features. If AI is used to find hidden trends in data, then features of the data used to train the model should be public, if not the training data itself. If AI automates a process, the results should be consistent. AI in government is expected to be held to a higher standard of trust than experimental models in the private sector.



The second dimension of transparency is about the purpose of AI models. AI programs should be impartial, transparent, accountable, consistent, secure, and should protect privacy. Leaders should not only build in all the right governance safeguards from the beginning, but also be clear about when AI models are used and why.

This clarity can help show the public and workforce alike that AI is not being used to infringe on privacy or replace workers but is designed to bring a real benefit to the public. Our research has shown that communicating benefits is **among the surest ways to generate trust in government.**²³

Toward a future of better, faster regulatory processes

AI makes big changes possible. The prospect of reengineering a process to create a better one incorporating machine learning can sound daunting, but it doesn't need to be. Consider starting small if you're new to AI. New York City built the original version of their AI-driven inspection tool by interviewing veteran firefighters about which factors, anecdotally, correlated with structure fires.²⁴

Within months, the first tool had collected enough data to help build a second tool. Both tools were an improvement on the old system of a card catalog, siloed by station.²⁵ Test new software. Automate a process. See if generative AI can reliably produce a few anodyne compliance reports. Find a real problem, with a clear, measurable outcome, and measure it. As regulators know, sometimes to understand a system, you have to inspect it yourself.

Endnotes

1. The Boeing 78 Dreamliner contains an estimated 6.5 million lines of code, but only 2.3 million physical parts. See: Jeff Desjardins, “How many millions of lines of code does it take?,” *Visual Capitalist*, February 8, 2017; Boeing, “World class supplier quality,” accessed October 31, 2023.
2. Devin Coldewey, “How ‘The Mandalorian’ and ILM invisibly reinvented film and TV production,” *TechCrunch*, February 21, 2020.
3. Dan Maloney, “Putting the brakes on high-frequency trading with physics,” *Hackaday*, February 26, 2019.
4. Deloitte, “Regulatory Intelligence,” homepage, accessed October 31, 2023.
5. Brian Perron, “Generative AI for social work students: Part I,” *Medium*, March 20, 2023.
6. Angie Heise, “Generative AI and public sector,” Microsoft’s Public Sector Center of Expertise, accessed October 31, 2023.
7. Editorial, “AI to boost regulator in Denmark?,” XBRL, February 17, 2017.
8. Alison Banney, “You can now file your tax return with help from a virtual chatbot,” *Finder*, June 18, 2018.
9. Edward Van Buren, William D. Eggers, Tasha Austin, Joe Mariani, and Pankaj Kamleshkumar Kishnani, *Scaling AI in government: How to reach the heights of enterprise-wide adoption of AI*, Deloitte Insights, December 13, 2021.
10. Scott W. Bauguess, “The role of big data, machine learning, and AI in assessing risks: A regulatory perspective,” champagne keynote address, US Securities and Exchange Commission, June 21, 2017.
11. Issie Lapowsky, “How bots broke the Federal Communications Commission’s (FCC’s) public comment system,” *Wired*, November 28, 2017; David Shepardson, “US broadband industry accused in ‘fake’ net neutrality comments,” *Reuters*, May 6, 2021.
12. Jeff Kao, “More than a million pro-repeal net neutrality comments were likely faked,” *HackerNoon*, November 23, 2017.
13. Ibid.
14. Emprata, *FCC restoring internet freedom docket 17-108: Comments analysis*, August 30, 2017.
15. Lorenzo Franceschi-Bicchierai, “More than 80% of all net neutrality comments were sent by bots, researchers say,” *Vice*, October 4, 2017.
16. Deloitte, “Trustworthy AI: Bridging the ethics gap surrounding AI,” accessed October 31, 2023.
17. Adam Sadilek, Stephanie Caty, Lauren DiPrete, Raed Mansour, Tom Schenk Jr., Mark Bergtholdt, Ashish Jha, Prem Ramaswami, and Evgeniy Gabrilovich, “Machine-learned epidemiology: Real-time detection of foodborne illness at scale,” *npj Digital Medicine* 1, 2018.
18. Kristen M. Altenburger, “Artificial intelligence and food safety: Hype vs. reality,” *Food Safety Magazine*, December 16, 2019.
19. William D. Eggers and Mike Turley, “Future of regulation: Case studies,” Deloitte Center for Government Insights, accessed October 31, 2023.
20. National Science Foundation, “Fighting food poisoning in Las Vegas with machine learning,” news release, March 7, 2016.
21. Jesse Roman, “In pursuit of smart,” *National Fire Protection Association Journal*, November/December 2014.
22. William D. Eggers, Mike Turley, and Pankaj Kamleshkumar Kishnani, *The regulator’s new toolkit: Technologies and tactics for tomorrow’s regulator*, Deloitte Insights, October 18, 2018.
23. Deloitte’s Center for Government Insights, “The digital citizen survey,” homepage, Deloitte Insights, accessed October 31, 2023.
24. Roman, “In pursuit of smart.”
25. Ibid.

About the authors

Matthew Gracie

magracie@deloitte.com

Matthew Gracie is a managing director in the Strategy & Analytics team at Deloitte Consulting LLP. He leads Deloitte's Regulatory Intelligence portfolio and is a thought leader with global and national experience in strategy, analytics, marketing, and consulting.

William D. Eggers

weggers@deloitte.com

William D. Eggers is the executive director of Deloitte's Center for Government Insights where he is responsible for the firm's public sector thought leadership.

Sam Walsh

sjwalsh@deloitte.co.uk

Sam Walsh leads and delivers complex transformation across the public sector. She works extensively with regulators and Arms Length Bodies, helping them to improve user experience, exploit digital technologies, and embrace innovation.

Scott Streiner

sstreiner@deloitte.ca

Scott Streiner is a senior adviser with Deloitte whose areas of focus include strategy and governance, regulatory modernization, transportation, and space. Before Deloitte, he was with the Government of Canada for 31 years. He has also served as vice chair of the OECD's Network of Economic Regulators and the Council of Canadian Administrative Tribunals.

Pankaj Kamleshkumar Kishnani

pkamleshkumarkish@deloitte.com

Pankaj Kishnani is a research specialist with the Deloitte Center for Government Insights. His research focuses on identifying trends in emerging technologies and innovations that can transform public sector. He also closely tracks digital transformation, regulation of emerging technologies, and service delivery innovations in government.

Acknowledgments

The authors would like to thank Kavita Majumdar, Abrar Khan, and Shambhavi Shah from Deloitte Insights for their editorial support and inputs.

About the Deloitte Center for Government Insights

The Deloitte Center for Government Insights shares inspiring stories of government innovation, looking at what's behind the adoption of new technologies and management practices. We produce cutting-edge research that guides public officials without burying them in jargon and minutiae, crystalizing essential insights in an easy-to-absorb format. Through research, forums, and immersive workshops, our goal is to provide public officials, policy professionals, and members of the media with fresh insights that advance an understanding of what is possible in government transformation.

Continue the conversation

Industry leadership

Matthew Gracie

Managing director | Strategy & Analytics | Deloitte Consulting LLP
+1 571 766 327 | magracie@deloitte.com

Matthew Gracie leads Deloitte's Regulatory Intelligence portfolio and has the privilege of working with governments around the world on their regulatory opportunities and challenges with an emphasis on the use of AI.

William D. Eggers

Executive director | Center for Government Insights | Deloitte Services LP
+1 571 882 6585 | weggers@deloitte.com

William Eggers is the executive director of Deloitte's Center for Government Insights where he is responsible for the firm's public sector thought leadership.

Contributors

Editorial: Abrar Khan, Shambhavi Shah, Emma Downey, and
Debashree Mandal

Creative: Molly Piersol and Pooja N Lnu

Cover artwork: Sofia Sergi and Sonya Vasilieff



Published in collaboration with Deloitte Insights.

About this publication

This publication contains general information only, and none of Deloitte Touche Tohmatsu Limited, its member firms, or its and their affiliates are, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your finances or your business. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser. None of Deloitte Touche Tohmatsu Limited, its member firms, or its and their respective affiliates shall be responsible for any loss whatsoever sustained by any person who relies on this publication.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee (“DTTL”), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as “Deloitte Global”) does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the “Deloitte” name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.