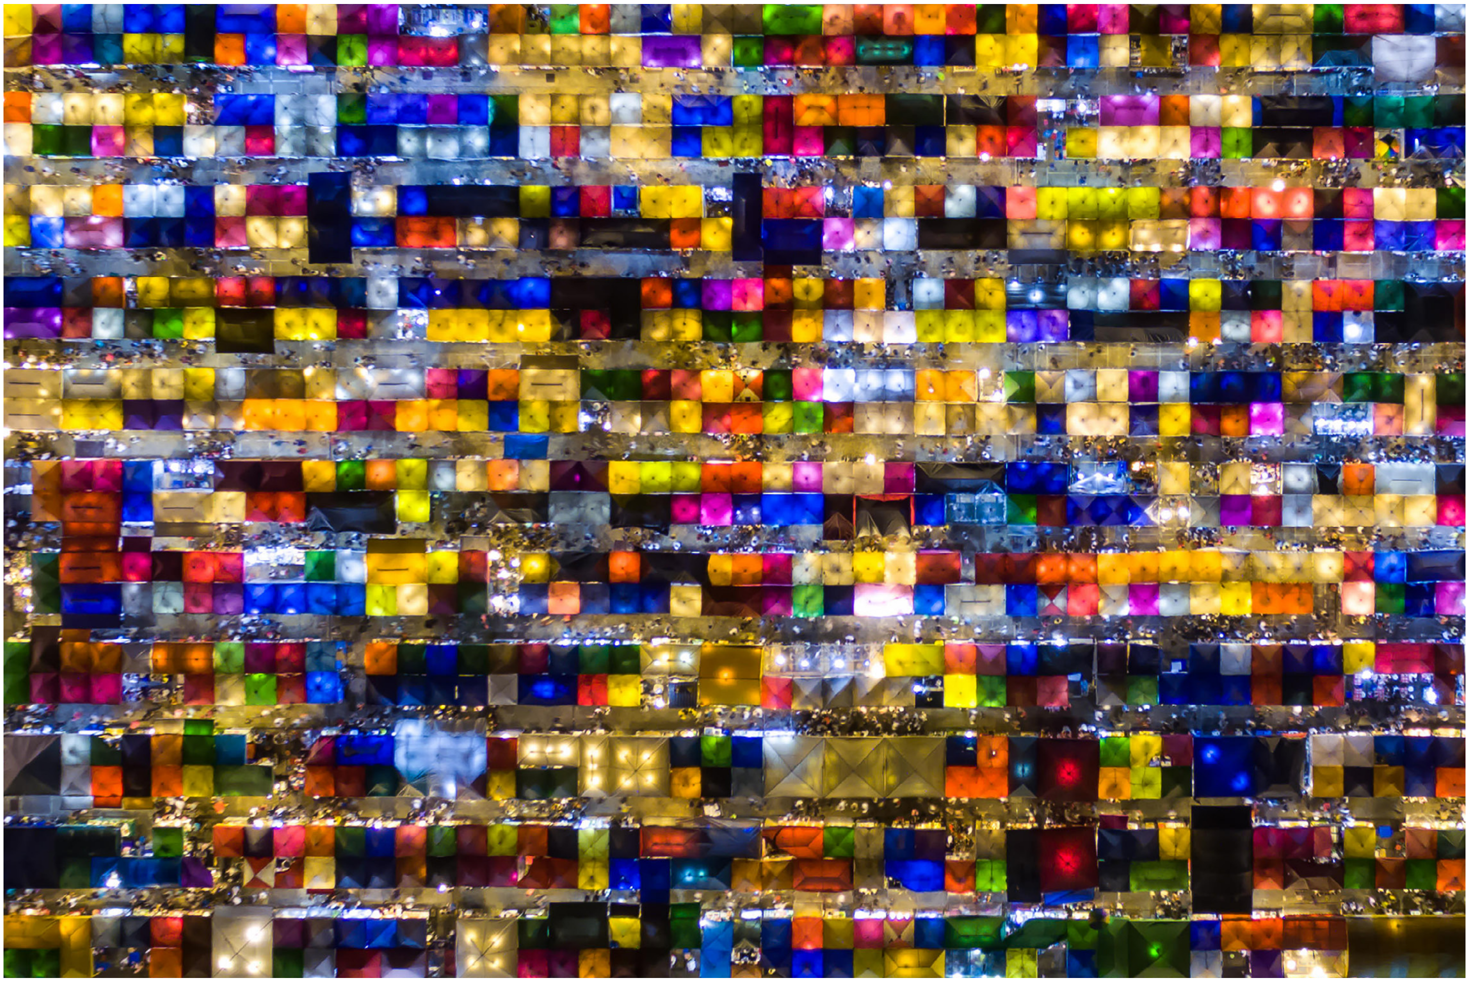

10 things governments should know about responsible AI

Organizations that use AI are increasingly faced with an emerging problem: They are struggling to govern the creation, deployment and management of AI-based services. Governments are no exception. This white paper explores the realities of how government organizations use AI, how they can be important stewards for the responsible implementation of this technology, and how they can access the right tools to ensure proper AI governance.

10 things governments should know about responsible AI

1. AI is highly relevant in the government sector
2. Government is a data billionaire
3. Data flows through the public sector on a wave of trust
4. Humans are at the heart of responsible AI
5. Responsible AI has 5 characteristics
6. Trusted data must be the foundation
7. Differentiate the risks of responsible AI
8. Responsible AI requires AI governance
9. Tools embed responsibility in AI
10. Governments can start to build responsible AI now

1.
AI is highly relevant in the government sector

Digitalization of government is key to increasing the efficiency of its programs and the effectiveness of its policies. Two elements critical to digital transformation are data and artificial intelligence.

Data is the fuel of the digital age. It enables new ways of improving and automating processes as well as the services offered to citizens, and it provides key insights to allow governments to predict and react appropriately to circumstances.

AI is equally important. It plays a crucial role in unlocking the value of data and providing a deeper understanding of the information. Together, data and AI—including more recent developments in generative AI and foundation models—can help drive efficiency, enhance quality and support human-friendly automation, helping governments make better and faster decisions.

Today, many countries around the globe recognize the importance of data and AI. Public administrations in Germany, the UK, Japan and beyond are already on the path of digitalization with huge investments.

The hype surrounding ChatGPT in 2023 and the emergence of large language models (LLMs) have finally contributed to the widespread consideration and application of AI.

Governments are now intensely driving discussions regarding governance of AI—not just within government itself, but across industries, in recognition of the need to implement AI in a responsible way. The EU’s AI Act and United States’ executive order on safe and secure AI are examples of such actions.[1, 2, 3]

Based on current trends, there is little doubt that AI will soon become a technology cornerstone for the public sector. However, despite these initiatives, governments are largely still in the exploration phase, merely looking at the enormous potential of AI.

2.
Government is a data billionaire

If the variety and volume of data owned by an organization indicate the potential value AI can deliver, then no sector stands to gain more than the public sector.

Government agencies have many databases at their disposal—citizen records with details of permits and family relations, for instance, or inventories that contain information about private and public assets, such as properties, houses, streets, traffic signs and trees. Any data related to businesses, migration, public safety, education and more are stored in public databases. It is easy to conclude that governments are data billionaires.

Today, many governments and their agencies are slowly beginning to derive value from such data, which may not be available in digital form, or may not be easily accessible or comparable. Though most citizen welfare services are application-based services, many have not yet migrated away from paper-based forms, and even those that have digital offerings continue to offer paper forms. This is to ensure that certain demographics are not excluded from accessing welfare services.

Forms—or the data they capture—are necessary for many situations, whether claiming social benefits, acquiring funds to implement energy-reducing measures in a building, or settling legal matters. However, this leads to a mass of unstructured data stored in written text form.

Government agencies also have vast quantities of records related to available geospatial information that they use to direct city traffic, manage major public events, or inform urban development or public infrastructure investments. These records are also important to nonpublic organizations, such as those in logistics and automotive industries, and startups innovating with new data-enabled business models.

Whatever the format, much of the data is siloed and disconnected. While this is by regulatory design in some circumstances, there is an opportunity to gain more value from the government-held data by consolidating or integrating it while also ensuring robust governance measures are in place.

At the same time, AI's natural language processing (NLP) capabilities can help public servants to be faster and more efficient while performing administrative work—which often includes reading, checking, understanding and writing information. NLP can help capture insights from all forms of unstructured data, including images, and accelerate decision making related to document processing.

Overall, combining AI with data can enable governments to improve their citizen experiences, increase their administrative efficiency, and inform their decisions with crucial insights.

3.
Data flows through the public sector on a wave of trust

The public sector has a prominent role as a trustee in handling data. Along with healthcare organizations and financial services firms, government agencies should be seen as the most trusted institutions in relation to data.

However, with the increasing use of digital services, the misuse of personal data has become a concern for citizens. While government has legitimate reasons for storing the data of citizens, unfortunate behaviors in other industries have led to trust-related issues. An OECD study shows that 41% of citizens do not trust their national government.[4]

Government agencies understand the importance of collecting and connecting personal data to deliver enhanced digital services to citizens. For example, if citizens need to inform their social security agency of a change in circumstances, it becomes much simpler for them if the government makes this information available across the agency. It helps the government initiate service bundles tailored to the individual needs of citizens. On the other hand, if the data fails to flow within or across agencies, the confidence of citizens in their government is reduced.

The full potential of personalized services cannot become a reality unless governments enable data sharing across their agencies and organizations. In some cases, a lack of sharing is enforced due to strict data regulations. Also, implementing an effective and governed data-sharing capability between governmental departments is complex. The reluctance to share and poor-quality data can make these undertakings costly and lengthy.

As a result, the concept of trust between citizens and the government is becoming much more complex. Although the digital literacy of citizens varies, they are demanding privacy and data rights.

This is happening at a time when government agencies wish to introduce more data-driven services enabled by AI technologies. Governments can better address trust-related issues by transparently communicating the way citizens' data will be acquired, stored and used, and by telling compelling stories of the benefits for the individual.

4.
Humans are at the heart of responsible AI

The IBM Institute for Business Value reported in 2022 that one of the five urgent activities governments must address is the development of ethical data and technology practices.[5] But when it comes to AI, humans are part of both the problem and the solution.

Today, almost all public services are supported by a backbone of technology. These systems, in general, are based on clearly defined laws and must be regulated and governed accordingly. And public servants are heavily involved in defining the policies and laws, designing and building the IT applications, and making decisions supported by the applications.

Human bias in AI

All humans possess cognitive biases. These are unconscious errors in thinking that make humans misinterpret information they perceive in the world. This condition is often attributed to *heuristics*, which refers to the mental shortcuts humans take to make decisions and judgments.

In relation to this context, the UK's Department for Work and Pensions commissioned a study in 2009. The objective was to collect factual evidence regarding bias and discrimination related to ethnic minorities in the labor market. Identical applications were created with different names—some that would be perceived to be typically White and some that would be perceived to belong to specific Black, Asian or other ethnic minority groups. The study revealed that 68% of White applicants received positive responses, while only 39% of the ethnic minority applicants fared similarly.[6]

We can even display a bias toward computers. Automation bias is the human inclination to trust automated systems and their decisions, and to ignore contradictory, often correct, information.

Different types of human biases can often end up in technological solutions. They can penetrate the designs we create, the data we capture or the hands-on training we perform for AI-enabled solutions.

As a result, there is always a risk that we—including governments—will deliver solutions that are not fair and inclusive. Hence, complete reliance on automated aids and decision support systems should be avoided.

Responsible development of AI requires mitigating the potential for human bias to influence this process. There must be clarity about who trains the AI system, what data is used and what goes into the algorithm's recommendations.

Human involvement at every stage of AI

To tackle the risk of bias, governments must include humans in decision making. Furthermore, Article 22 of the General Data Protection Regulation (GDPR) states that when a solely automated decision is made, resulting in a legal or similarly significant event, individuals have the right not to be subjected to it.[7] Although some caveats authorize such automated decision making, having a human in the process is preferable.

IBM's Trust and Transparency Principles, which guide its approach to data and AI, echo this thought. The first principle states the following:

"The purpose of AI is to augment human intelligence. This means that we do not seek to replace human intelligence with AI, but support it. Since every new technological innovation involves changes to the supply and demand of particular job roles, IBM is committed to supporting workers in this transition by investing in global initiatives to promote skills training around this technology."[8]

Diverse teams for better AI outcomes

Governments must ensure teams have diverse characteristics—such as gender, ethnicity, socioeconomic status and disability—so that they can better represent society.

According to the IBM Policy Lab, "Having a diverse design team broadens the understanding of user habits, enabling greater exploration of use cases, both the positive and the negative."[9]

Government role in the digital gap

For the responsible implementation of AI, citizens and other government constituents must also be included, to better shape how services can be developed more effectively and with greater impact. However, when it comes to digital services, there is an increasing concern regarding two factors: access and skills.

For example, in the UK in 2021, Lloyds Bank reported that 21% of the population was digitally disadvantaged, and 36% of working adults lacked the necessary digital skills.[10] In such a situation, it becomes essential for governments to work toward closing the gaps. Enabling a higher percentage of the population to utilize digital services can reduce the burden on the government. Furthermore, digitally savvy citizens can be included in the design of AI-based services.

5. Responsible AI has 5 characteristics

Responsible AI can be defined according to five characteristics: transparency, explainability, fairness, robustness and privacy. These traits apply across the AI lifecycle, which includes design, development, use and maintenance.

Transparency

The OECD Principles on AI state that there should be transparency and responsible disclosure around AI systems to ensure that people understand AI-based outcomes.

Transparent systems should disclose the purpose and use of AI. They should share information on data storage and privacy policies, confidence measures, levels of procedural regularity, and error analysis.

In 2021, the UK's Centre for Data Ethics and Innovation, or CDEI (since renamed the Responsible Technology Adoption Unit), conducted research that revealed awareness about the use of algorithms in the public sector was a priority for study participants.[11] The participants wanted transparency regarding the description of the algorithm, usage of the algorithm, human oversight, potential risks and technicalities of the algorithm.

The CDEI also recounted a tension between simplicity and transparency. It found that providing tiers of information—depending on the expertise of the stakeholder researching the algorithm-enabled service—was a satisfactory method of reducing that tension.

Explainability

People have a right to know if they are interacting with an AI system and how it has come to a certain decision or recommendation. This should be explained in nontechnical terms so that it is comprehensible to all. In 2019, the European Parliamentary Research Service (EPRS) stated the following: "Explainability must address both the technical processes of an AI system and the related human decisions taken in accordance with the EU guidelines."[12]

Fairness

Fairness relates to treating citizens, or groups of citizens, equitably. AI can assist humans in making fairer choices, countering biases and promoting inclusivity.

Governments should prevent discrimination related to protected characteristics, such as gender, race, age and veteran status. They should try to ensure that their AI-enabled systems are examined for fairness.

Robustness

IBM Research® summarizes robustness as making AI "hack-proof."[13] Responsible AI-enabled systems should resist adversarial attacks and exceptional conditions without causing unintentional harm. They must withstand intentional interference, such as active sabotage with poisoned data and incidental interference with data corruption.

According to the EPRS, "Trustworthy AI requires algorithms to be secure, reliable and robust enough to deal with errors or inconsistencies during all phases of the system."

Responsible AI systems should be robust not only in the ideal conditions of testing but also in the imperfect circumstances of real life.

Privacy

To honor the privacy of individuals, responsible AI must fully disclose what data is collected, how it will be used and stored, and who has access to it. This characteristic features prominently in Australia's Artificial Intelligence Ethics Framework: "AI systems should respect and uphold privacy rights and data protection and ensure the security of data."[14]

The AI designed, developed and used by governments must exhibit all five characteristics to be considered responsible.

6.
Trusted data must be the foundation

The issue of trust and responsibility does not lie in AI models alone. Often AI failures happen due to problems in data preparation and data organization. Examples include racial bias built into public safety algorithms and gender bias related to the language used by job applicants in their resumes.[15]

Today, government data repositories are spread across multiple clouds, applications, locations, environments and vendors, leaving public servants inhibited by data silos. And the problem is not just the distribution and scale of all these different data sources; it's also in the increased variety of forms that the data takes. The reality that government organizations find themselves in today is one where they are drowning in various data sources. Eighty percent of some data team members' time is spent on data cleaning, integration and preparation,[16] with only the remaining 20% on actually putting the data to use.

To get the full benefit of the huge quantities of data that governments hold—and to take advantage of the latest capabilities provided by LLMs and generative AI—those governments must turn that available data into knowledge at a much faster pace than is typically available today in their organizations.

Nonetheless, governments must base their AI models on a responsible foundation of data, not sacrificing privacy and governance for speed, but catering for all as new solutions are created to deliver better outcomes for government.

Integrated data management capabilities, such as virtualization, discovery, governance, curation and orchestration, can help governments develop confidence in the data on which their AI models are built and applied. Additionally, those capabilities will enable governments to unlock data from silos and turn it into actionable insights at a faster pace.

A government agency's maturity in the four phases of collecting, organizing, analyzing and infusing AI into processes and services will likely vary across its organization. It is better to focus on delivering one common business outcome that could improve maturity in each of those phases instead of attempting to build maturity across the whole agency one phase at a time. The latter is a never-ending cycle that can struggle to demonstrate value; the former is a flexible approach wherein teams can learn along the way and implement these lessons while delivering the next set of outcomes.

7. Differentiate the risks of responsible AI

Some governments are developing or have already developed frameworks to understand potential AI use cases and the related risks that can lead to a loss of trust. For example, the NIST AI Risk Management Framework in the US has been established, followed by a focus on federal implementation.[17] Singapore developed its Model AI Governance Framework in 2020,[18] and Japan is following suit.[19] At the end of 2023, the European Parliament and Council reached political agreement on the European Union's Artificial Intelligence Act ("EU AI Act").

Different risk levels will impact the measures and the toolbox necessary to manage responsible AI. A framework should enable a practical discussion about the value of AI in public services and the related challenges and pitfalls of using AI in a government context. The following three examples demonstrate how risk profiles can vary.

1. Chatbot

Chatbots, sometimes known as virtual assistants, are one of the most common starting points for governments. They're employed to manage large volumes of citizens' questions or to help citizens navigate complex legislation.

A chatbot uses artificial intelligence and natural language processing to understand human questions and automate responses, simulating and creating meaningful conversations. IBM watsonx Assistant, for example, is a capability leveraging generative AI, upon which many such chatbots are built.

A successful virtual assistant is trained for specific interaction styles and outcomes, with nondiscriminatory dialogues. Sharing public services data through AI-enabled chatbots could be classed with an easily manageable risk profile.

2. Document analysis

Government work often requires the reading, analysis and writing of documentation. Examples include benefit and grant applications, the filing of taxes, witness statements, and ministerial and public correspondence.

AI-enabled solutions can help public servants find relevant information, compare documents, structure content, summarize issues, generate responses and more. This risk profile requires the ability to trace every step taken by artificial intelligence in categorizing, summarizing and comparing documents, and sourcing additional relevant information in the process.

AI can be employed in support of the judiciary. In Germany, courts have faced delays in proceedings, hearings and outcomes, with a backlog of more than 10,000 cases at the Stuttgart Higher Regional Court. A transparent, traceable AI solution equipped with natural language understanding, called OLGA20, helps extract metadata and categorize cases. This helps judges and clerks process vast quantities of documentation at a faster pace, contextualize comprehensive information and resolve cases more quickly.[20]

3. Enablement of decision making

Lastly, there are higher risk levels when AI has an influence on decision-making processes. One example is the application of AI in child welfare services, where social workers often face complex and sensitive decisions related to child protection, family support and placement. AI systems can analyze vast data sets, including family histories, socioeconomic factors and prior case outcomes, to assist social workers in making more informed and timely decisions. For instance, predictive analytics models can identify patterns that may indicate a higher risk of child abuse, enabling early intervention and support for vulnerable families.

However, the use of AI in social services decision making is not without its challenges: There are concerns about biases in the algorithms, potentially leading to unfair treatment or reinforcing existing disparities. Additionally, the ethical implications of relying on machine-driven decisions in sensitive and nuanced social contexts must be carefully considered.[21]

Based on these three examples, we learn that the level of risk depends on the specific AI use case. Each area of the public sector (tax, social services, justice, etc.) is likely to have the full range of low-, medium- and high-risk profiles. Governments should introduce AI frameworks for their operations and use them for guidance. The decision to use responsible AI should be made on a case-by-case basis.

8. Responsible AI requires AI governance

Government organizations must also consider the governance of AI and its use in government. As AI technology continues to advance, concerns about its potential negative consequences have increased. Recent developments in generative AI have raised concerns in the public consciousness. Government reputation is again at risk, with the threat also of litigation and fines related to misuse of data and compliance errors.

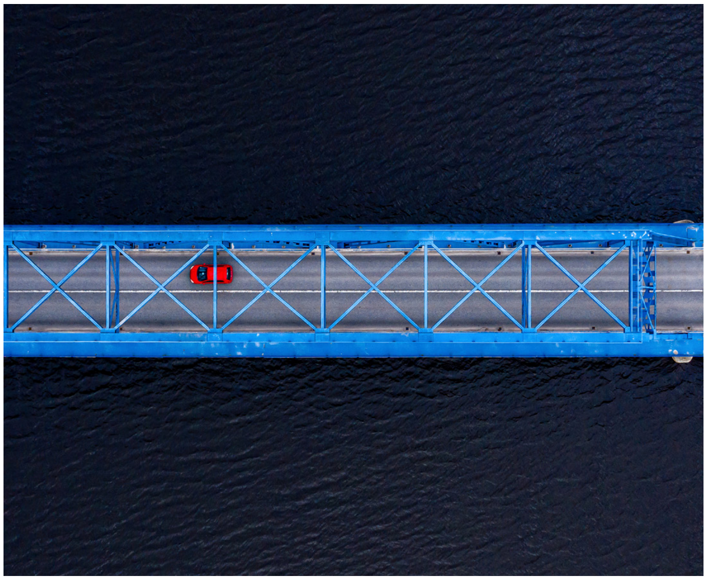

Establishing AI governance

Government organizations must proactively create an AI governance strategy that will address the distinct new risks of foundation models—including LLMs—as well as the better-understood risks from more established predictive machine learning.

AI governance—the ability to direct, manage and monitor the AI activities of an organization—provides a level of organizational rigor and human oversight into how AI models are created and deployed. It can complement any existing MLOps processes with activities intended to strike the appropriate balance between the benefits and risks of AI.

The three areas that governance must address are: compliance to satisfy regulations; risk management to proactively detect and mitigate risk; and lifecycle governance to manage, monitor and govern AI models.

Implementing AI in parallel

As governments begin to develop AI-enabled solutions, they can use them both to establish and refine the necessary AI governance measures, and to ensure that the right processes and tools are employed to deliver them.

9.
Tools embed responsibility in AI

Manual tools and processes for building, deploying, managing and monitoring models may work well when building a proof of concept, but they are not sufficient when scaling AI across the enterprise and into production. Governments must take advantage of the automated tools and processes necessary to shorten the model lifecycle, driving visibility and transparency.

Ideally, government organizations will automate AI governance activities to streamline processes and use a single set of governance policies and workflows across platforms and applications.

Embarking on AI-enabled solutions is an opportunity to enhance suboptimal tooling, automate and consolidate. Governments should seek out tooling that drives visibility across the organization with automated collaborative tools, customizable dashboards and reports.

As for data governance, governments must also adopt tooling that ensures only the right users see the right data at the right time.

IBM AI governance

The IBM watsonx™ AI and data platform includes three core components and a set of AI assistants designed to help scale and accelerate the impact of AI with trusted data across organizations. The three core components, outlined below, are watsonx.ai, watsonx.data and watsonx.governance.

watsonx.ai

This is a next-generation enterprise studio for AI builders to train, validate, tune and deploy both traditional machine learning and new generative AI capabilities powered by foundation models. It enables governments to build AI applications in a fraction of the time with a fraction of the data.

IBM's library of foundation models includes both IBM-developed as well as third-party and open-source models, including some that have been rebuilt on IBM's data sets to reduce model sizes and mitigate problems (such as bias) and other common risks; these models can be plugged into extra applications.

watsonx.data

This fit-for-purpose data store is built on an open lakehouse architecture, supported by querying, governance and open data formats to access and share data.

It allows government organizations to access all their data through a single point of entry across all clouds and on-premises environments.

watsonx.governance

This is an end-to-end toolkit to govern generative AI and traditional machine learning models across the entire AI lifecycle and to enable responsible, transparent and explainable AI workflows.

Watsonx.governance helps clients govern the training data they use and the AI they deploy so that they can operate, scale and succeed with trust as insurance. It offers the ability to improve adherence to AI regulations, such as the proposed EU AI Act, and internal compliance standards.

Scaling AI with trust

The watsonx platform includes four key characteristics that allow governments to scale trusted AI:

- **Open:** Watsonx is based on available open technologies, providing model variety to cover the breadth of government use cases and compliance requirements.

- **Trusted:** Watsonx is designed with principles of transparency, responsibility and governance. Governments can use watsonx to build AI models trained on their own trusted data to help manage legal, regulatory, ethical and inaccuracy concerns.

- **Targeted:** Watsonx is designed for enterprise and targeted at business domains. It unlocks new value for government-related use cases.

- **Empowering:** Watsonx provides a platform that empowers enterprises to bring their own data to tune, train and deploy generative AI models. It also supports diverse sovereignty requirements that may be government requirements.

10.
Governments can start to build responsible AI now

Transforming into a government enabled by responsible AI can be daunting. While technology is the enabler of this transformation, humans are at the heart of this process.

Before building responsible AI, government agencies must determine the following factors:

- Services and citizen engagements that can benefit from the transformation

- Areas that can show quick results to prove value and justify the transformation

- Internal users who can gain an advantage from AI

- Skills necessary to execute the AI project

- Gaps in capabilities that need to be filled

Government agencies should invest in an end-to-end model for accelerating digital transformation. With the help of the model, they can:

- Generate innovative ideas and bring together experts, practices and technologies to implement those ideas.

- Focus on the pain points of users—citizens and public servants—to ensure value is first delivered in priority areas.

- Use the right ecosystem of partners and stakeholders to increase relevance and build productive, meaningful relationships.

IBM Garage as an approach

IBM Garage™ is an end-to-end model built on a collaborative exchange of ideas, knowledge, experimentation, expertise and support. It's based on three stages:

Co-create: Defines your vision
Co-execute: Proves the value of your vision
Co-operate: Scales your solution and your team's capabilities

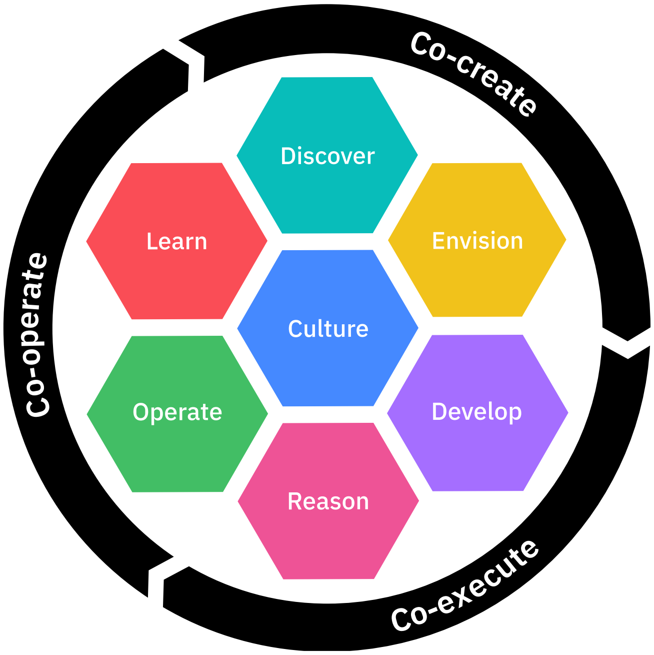

IBM Garage is designed to meet today's demands for modernization, transformation and growth. It can help government agencies accelerate their AI transformation journey and deliver value iteratively and rapidly. And it can work with diverse, multidisciplinary teams within the government agency, enabling them to innovate and solve problems faster.

During the pandemic, for example, a COVID-19 response unit was formed in the U.S. state of Rhode Island. An interagency organization tasked with mitigating and reducing the virus's spread within the state, the unit generated surveillance data, analytics and insights to be shared internally and with the public. However, the team relied on cumbersome, paper-based processes to merge data sets into a single, trusted source of information. As a result, the unit needed a way to synthesize data sets to produce reports and generate insights for leaders and citizens.

The state's COVID-19 response unit participated in a design thinking and analytics insights usability workshop with IBM to improve its data and analytics operations. The unit used Enterprise Design Thinking® and IBM Garage methodologies to assess "as-is" data and analysis processes, map its desired "to-be" state and be more prepared to respond to the pandemic.

Building responsible AI is an ongoing process

Responsible AI is vital to optimizing government operations. It can help governments deliver services their citizens need and effectively deal with issues like global warming. This paper has shown how it can also help build citizens' confidence in their governments.[22]

However, responsible AI is not just a status at a point in time. Building it is an ongoing process in which several factors play a crucial role:

- The principles that define it
- The data on which it is based
- The humans that are included in the lifecycle
- The framework that supports the assessment of need and relevance
- The methods that build accountability and governance into development and application
- The tools that assess the trustworthiness and identify areas of concern

Get started now

AI proof-of-concept projects have struggled to evolve and generate real value for organizations. This makes it all the more important for governments to adopt a proven approach like IBM Garage that can chart their transformation journey with an iterative framework—guiding them from ideation to building to scale. It is never too late—or too early—to start building trust in digital services. Begin now.

About the authors

Eckard Schindler
Strategy Advisor
Global Government Industry
IBM Technology
schindler@de.ibm.com

Eckard works with clients and partners worldwide on the digital transformation of public administration. He has many years of experience in strategy consulting, sales and business development in the public sector. Using AI to bring value to governments is one of his business passions. Eckard lives in Germany.

Sharon Moore
CTO
Global Government Industry
IBM Technology
sharon.moore@uk.ibm.com

Sharon's personal mission is transforming public service with technology. She helps governments around the world do better for citizens and plays an active role in driving change for good in the tech industry. Responsible AI is one of her focus areas. Sharon lives in Scotland, UK.

1. “Digital-by-default: A new concept in Germany’s development co-operation,” OECD, 17 December 2021.
2. “Roadmap for digital and data, 2022 to 2025,” GOV.UK, 9 June 2022.
3. “Japan launches Digital Agency to push ahead with long-overdue reforms,” Kazuaki Nagata, The Japan Times, 1 September 2021.
4. “Trust in Government,” OECD, 2021.
5. “Government transformation in tumultuous times,” IBM Institute for Business Value, 13 April 2022.
6. “Undercover job hunters reveal huge race bias in Britain’s workplaces,” The Guardian, 17 October 2009.
7. “Automated individual decision-making, including profiling,” article 22, General Data Protection Regulation (GDPR), 27 April 2016.
8. “What is AI ethics?” IBM, March 2021.
9. “Five Technology Design Principles to Combat Domestic Abuse,” IBM, 11 November 2020.
10. “Essential Digital Skills Report 2021,” Lloyds Bank, September 2021.
11. “BritainThinks: Complete transparency, complete simplicity,” Centre for Data Ethics and Innovation (CDEI), 21 June 2021.
12. “EU guidelines on ethics in artificial intelligence: Context and implementation,” Tambiama Madiega, briefing, European Parliamentary Research Service, September 2019.
13. “Securing AI systems with adversarial robustness,” IBM, 15 December 2021.
14. “Australia’s AI Ethics Principles,” Australian Government, Department of Industry, Science and Resources, 2022.
15. “What Do We Do About the Biases in AI?” James Manyika, Jake Silberg, and Brittany Presten, Harvard Business Review, October 2019.
16. “Data Integrity Trends: Chief Data Officer Perspectives in 2021,” Corinium, June 2021.
17. “House Dems call on White House to make agencies adopt NIST AI framework,” FedScoop, 21 July 2023.
18. “Singapore’s Approach to AI Governance,” Personal Data Protection Commission Singapore, January 2020.
19. “Governance Guidelines for Implementation of AI Principles Ver. 1.1,” Ministry of Economy, Trade and Industry, January 2022.
20. “Judicial systems are turning to AI to help manage vast quantities of data and expedite case resolution,” IBM, January 2024.
21. “Child welfare algorithm faces Justice Department scrutiny,” Sally Ho and Garance Burke, Associated Press, 31 January 2023.
22. “Rhode Island navigates an unpredictable pandemic,” IBM, March 2022.

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the
United States of America
April 2024

The content in this document (including currency OR pricing references, which exclude applicable taxes) is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY

OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

The client is responsible for ensuring compliance with laws and regulations applicable to it. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the client is in compliance with any law or regulation.