

Deloitte.

The Government & Public Services AI Dossier

A selection of high-impact use cases



About the Deloitte AI Institute

The Deloitte AI Institute™ helps organizations connect all the different dimensions of the robust, highly dynamic, and rapidly evolving Artificial Intelligence ecosystem. The AI Institute leads conversations on applied AI innovation across industries, with cutting-edge insights, to promote human-machine collaboration in the “Age of With.”

The Deloitte AI Institute aims to promote the dialogue and development of AI, stimulate innovation, and examine challenges to AI implementation and ways to address them. The AI Institute collaborates with an ecosystem composed of academic research groups, start-ups, entrepreneurs, innovators, mature AI product leaders, and AI visionaries to explore key areas of artificial intelligence including risks, policies, ethics, the future of work and talent, and applied AI use cases. Combined with Deloitte’s deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result, delivers impactful perspectives to help organizations succeed by making informed AI decisions.

No matter what stage of the AI journey you are in: whether you are a board member or a C-Suite leader driving strategy for your organization—or a hands-on data scientist bringing an AI strategy to life—the Deloitte AI Institute can help you learn more about how enterprises across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute for a full body of our work, subscribe to our podcasts and newsletter, and join us at our meet-ups and live events. Let’s explore the future of AI together.

www.deloitte.com/us/AllInstitute



Foreword

Artificial intelligence (AI) continues to advance by leaps and bounds, delivering breathtaking capabilities once thought to be far off in the future. With a remarkable capacity to understand complex inputs and generate valuable outputs—and the rapidly emerging ability to execute real-world actions through intelligent agents and physical AI—AI is opening the door to innovations and new ways of working that were almost unthinkable just a few years ago.

As the AI landscape evolves, so does this compendium. Our latest edition features 130 of the most compelling use cases for AI across six major industries:



Consumer



Energy, Resources & Industrials



Financial Services



Government & Public Services



Life Sciences & Health Care



Technology, Media & Telecommunications

For each of these industries, we explore innovative uses for AI that can address enterprise challenges in new ways, expand and improve capabilities in every business function, and deliver advantages in efficiency, speed, scale, and capacity. To further provide context and clarity, each case specifies the primary business function it supports and whether agentic and/or physical AI is used. These labels are presented for informational purposes, helping you quickly grasp the intention and scope of each case.

Of course, every powerful tool presents potential risks, and AI is no exception. To help you better understand and manage the risks associated with AI, we use Deloitte’s Trustworthy AI™ framework throughout this compendium to illuminate factors that contribute to trust and ethics in AI deployments, and to offer practical steps for strengthening governance and risk mitigation. The specific objective of our Trustworthy AI™ framework is to help organizations create AI systems that are (1)

fair and impartial, (2) robust and reliable, (3) transparent and explainable, (4) safe and secure, (5) responsible and accountable, and (6) private.

Given AI’s rapidly expanding scope and reach, this compendium offers just a glimpse of what the technology can do. Our goal is to convey what AI is currently capable of, and even more important, to inspire the next wave of AI-driven innovation. As AI technology continues to improve and organizations increasingly embrace it, we anticipate even more impressive and compelling use cases in the future—including those that have yet to be imagined.

We hope the use cases highlighted here will spark new ideas, provide a foundation for successful deployments, and set organizations on a path to harness the maximum value from this powerful new technology.



Nitin Mittal
Global Head of Generative AI
Deloitte Consulting LLP



Beena Ammanath
Global Deloitte AI Institute Leader
Deloitte Consulting LLP



Jim Rowan
US Head of AI
Deloitte Consulting LLP

The Government & Public Services AI Dossier

AI is becoming a critical tool for driving mission outcomes in the government and public services industry. From frontline service delivery to policymaking and infrastructure management, AI systems are reshaping how government and public services organizations operate, deliver services, and engage with constituents.

In an environment often constrained by legacy systems, budget pressures, and complex regulatory mandates, AI offers the ability to modernize processes, improve responsiveness, and unlock new levels of efficiency and insight.

AI is already being deployed for a wide range of use cases across the sector, such as drafting policy documents, analyzing and summarizing global legislation and policies, simulating urban planning scenarios, hyper-personalizing education, and engaging with the public through virtual public servants.

Agentic AI enables more autonomous support across many areas—including regulatory oversight, permitting, benefits, intelligence, and health documentation—augmenting, not replacing, human expertise. Similarly, physical AI offers governments powerful tools to protect people, deliver services, and manage enormous physical domains (roads, utilities, public buildings, transit systems, public safety infrastructure, etc.) that are chronically under-resourced given the maintenance and oversight they require.

AI innovations, when paired with secure data integration and worker oversight, can reduce administrative burdens, streamline operational workflows, and enhance the precision and speed of government decision-making.

As government and public sector use cases mature and the regulatory landscape evolves, AI is poised to become not just a tool for modernization, but a catalyst for more adaptive and data-driven governance. However, the adoption of AI in government and public services also brings unique challenges.

Public institutions must meet high standards for transparency, fairness, and accountability—often in environments where privacy concerns, procurement complexity, and uneven digital maturity can limit deployment. As such, success depends not just on technical innovation, but on cross-agency collaboration, trust-building with the public, and alignment with legal and ethical frameworks.

AI systems are reshaping how government and public services organizations operate, deliver services, and engage with constituents.

Note: The tags below each use case indicate its primary business function and whether Agentic or Physical AI is used.

Tags

Primary business function

Agentic AI

Physical AI

AI-supported regulatory examinations and inspections

Augmenting regulators with multi-agent systems for fast, more accurate review

Agentic AI systems can support regulators with specialized agents that analyze data, review large volumes of documents, and uncover insights to streamline examinations and inspections.

ISSUE/OPPORTUNITY

Regulatory examinations and inspections are essential to ensuring compliance in industries such as financial services, food safety, and transportation. These processes require regulators to review vast amounts of historical and current market data, pull from multiple document sources, and interpret complex regulations. The work is labor-intensive, time-consuming, and prone to bottlenecks, which can delay the detection of compliance risks.

As regulatory responsibilities expand and the business environment becomes increasingly complex, regulators need tools that can streamline and accelerate analysis without sacrificing quality or rigor. Agentic AI offers the ability to automate routine examination tasks, highlight anomalies, and provide real-time support to examiners.

Tags

Compliance & Risk

Agentic AI

HOW AI CAN HELP

Quantitative analysis

AI agents can extract formulas, run calculations, and compare data against regulatory benchmarks to identify areas of potential concern.

Document review

The systems can process large volumes of reports, filings, and other documents, uncovering key insights and flagging areas that warrant attention from human examiners.

Dynamic examiner support

A conversational interface allows regulators to interact with the system, request clarifications, and guide the direction of analysis, ensuring a human stays in the loop.

Cross-domain applicability

The framework currently being used for financial services examinations can be adapted to other types of regulatory inspections, including food safety, tax compliance, and transportation.

AI-supported regulatory examinations and inspections

MANAGING RISK AND PROMOTING TRUST



Private

Because examinations can involve sensitive institutional and personal data, agentic systems should employ strict data protection protocols and comply with all relevant privacy regulations.



Robust and reliable

Flawed analyses can lead to missed risks or incorrect enforcement actions. To avoid such problems, AI agents should be validated against known regulatory outcomes and tested continuously for accuracy.



Responsible and accountable

Since regulatory judgments can have serious legal and economic consequences, AI agents should be used primarily as decision-support tools, with human regulators retaining final authority over important decisions and actions.



Fair and impartial

To ensure equitable oversight, agents should be designed to apply standards consistently across institutions and sectors, minimizing the risk of bias in how examinations are prioritized or conducted.

POTENTIAL BENEFITS

Faster examinations at lower cost

Automated analysis and document review can shorten the time required to complete inspections, helping regulators do their jobs efficiently and effectively.

Improved risk detection

By examining vast amounts of quantitative and qualitative data, AI agents are more likely to detect problems or anomalies that could go unnoticed.

Scalable oversight

With AI agents handling routine analysis, regulators can expand their oversight capacity without a corresponding increase in staff, helping to maintain compliance in complex and high-volume domains.

AI-driven permitting

Accelerating and simplifying government permitting processes

Agentic AI systems can streamline permitting workflows by using specialized agents to scan applications, extract key data, check compliance, and provide real-time feedback to applicants and reviewers.

ISSUE/OPPORTUNITY

Permitting is a critical function for government agencies at the federal, state, and local level, covering areas such as land use, energy, waterways, and housing. These processes can be slow, opaque, and labor-intensive, requiring reviewers to manually assess large volumes of documentation against complex and sometimes overlapping regulations. Small errors or missing information in

applications can create lengthy delays, increasing costs for both applicants and agencies. With government entities facing larger workloads and limited staffing, applicants and regulators alike are seeking opportunities for efficiencies. New approaches can help accelerate permitting while ensuring compliance with local, state, and federal requirements.

HOW AI CAN HELP

Automated document review

AI agents can scan and extract relevant data from submitted applications, checking for completeness and identifying missing or inconsistent information.

Compliance validation

Specialized agents can evaluate applications against permitting requirements, standard operating procedures (SOPs), and applicable regulations, flagging potential compliance issues early in the process.

Applicant feedback and support

Using AI and other advanced technologies, agents can provide real-time feedback to applicants, helping them correct errors or supply missing documentation before delays occur.

Reviewer assistance

AI agents can support human reviewers by summarizing key findings, verifying that evaluations align with regulatory standards, and highlighting areas requiring human attention.

AI-driven permitting

MANAGING RISK AND PROMOTING TRUST



Fair and impartial

To maintain public trust, agents should be designed to apply requirements consistently across all applicants.



Robust and reliable

Errors in interpreting regulations can lead to incorrect approvals or major delays. To improve consistency and reliability, agents should be trained on verified SOPs and continuously updated as requirements evolve.



Private

Since applications often contain sensitive personal and corporate data, agentic AI systems should be designed with strong data protection measures and full compliance with privacy regulations at all three levels of government.



Responsible and accountable

Permitting decisions can have important legal and financial ramifications. To avoid errors in accuracy or judgment, agentic AI systems should be primarily used for decision support, with final approval authority retained by human reviewers.



POTENTIAL BENEFITS



Reduced costs and permitting timelines

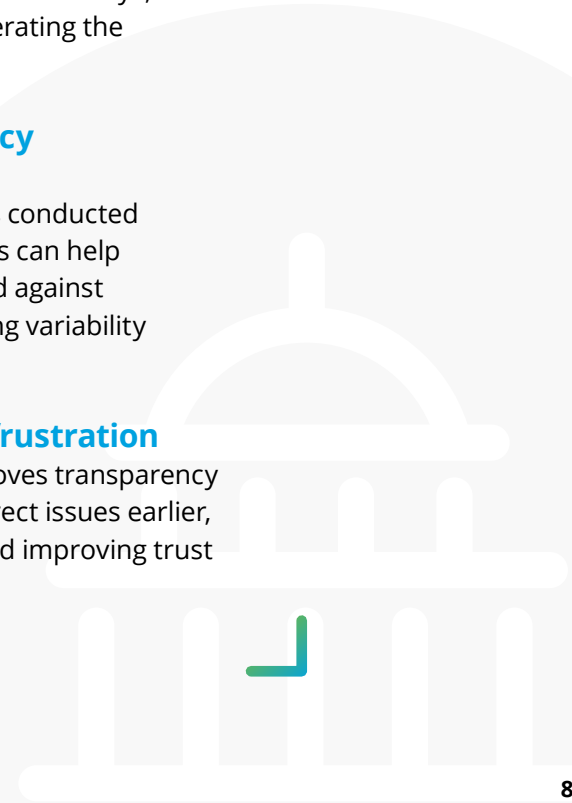
By automating document review and compliance checks, processing times can be reduced from months to days, lowering costs and accelerating the permitting process.

Improved consistency and fairness

Standardized evaluations conducted or supported by AI agents can help all applicants be assessed against the same criteria, reducing variability across reviewers.

Reduced applicant frustration

Real-time feedback improves transparency and helps applicants correct issues earlier, reducing costly delays and improving trust in the process.



AI-enhanced benefits eligibility

Improving access to government benefits programs

AI systems can streamline the application and determination processes for government benefits programs and public assistance.

ISSUE/OPPORTUNITY

Determining eligibility for public benefits such as health services, food stamps, unemployment insurance, workers comp, and childcare can be a time-consuming process for applicants and government agencies alike. Traditional public assistance portals may have redundant data entry, unclear requirements, and complex navigation, which can frustrate applicants (especially those with limited digital literacy) and delay access to benefits.

Meanwhile, government caseworkers manage lengthy manual intake processes and may have a lack of support in interview preparation and execution. These challenges increase agency workloads and can lead to incomplete applications, inconsistent reviews, and a higher likelihood of errors.

Tags

Customer Service

Agentic AI

HOW AI CAN HELP

Conversational intake

Conversational agents can replace static application forms with guided, human-like interactions that clarify questions and adapt dynamically to applicant responses.

Automated document intelligence

Document analysis agents can extract and validate data from submitted records, reducing manual effort and maintaining accuracy in eligibility assessments.

Interview assistance

AI agents can support human caseworkers by preparing interview outlines, validating applicant responses in real time, and highlighting missing or inconsistent information.

Onboarding and on-the-job assistance

AI can accelerate new caseworker onboarding by generating tailored training materials and insights from historical cases, helping agencies address workforce challenges more efficiently. AI agents can also support human caseworkers by preparing interview outlines, validating applicant responses in real time, and highlighting missing or inconsistent information.

End-to-end applicant support

The specialized agents can work together across the intake and review processes to provide applicants with a seamless experience that is less intimidating and confusing, helping to ensure applications are complete and compliant.

AI-enhanced benefits eligibility

MANAGING RISK AND PROMOTING TRUST



Transparent and explainable

In order for applicants and caseworkers to trust AI-driven recommendations, agents should provide clear explanations of eligibility determinations, including references to applicable rules and policies.



Robust and reliable

Conversational agents might misunderstand user inputs, especially in cases of language barriers or low digital literacy. Agents should be trained on multiple languages and continuously monitored for performance.

POTENTIAL BENEFITS

Faster access to benefits

Automating intake and eligibility can reduce processing times, helping applicants receive support more quickly.

Cost savings and improved productivity

Using AI agents to streamline processes and automate manual tasks can reduce costs and enable caseworkers to focus more attention on complex cases and applicant support.

Improved applicant experience

Guided interactions, real-time feedback, and simplified workflows can reduce frustration and confusion, especially for applicants with limited digital literacy.

Smarter intelligence

Improving intelligence reporting via AI agents and multimodal data integration

Agentic AI systems can autonomously collect, synthesize, and cross-reference intelligence data from various sources, producing faster, more comprehensive insights.

ISSUE/OPPORTUNITY

Intelligence analysts, including those working with open-source intelligence (OSINT), face the daunting task of searching, collecting, and interpreting vast amounts of multimodal data—from social media and public records to satellite imagery, communications, and financial reports. This manual process is time-consuming and limited by the capacity of human analysts to track patterns across large datasets.

As threats evolve and the volume of public and classified data continues to grow at breakneck speed, agencies need tools that can help accelerate the production of intelligence reports while improving accuracy, adaptability, and depth of insight. Agentic AI offers the potential to automate core intelligence gathering activities so human analysts can focus on validation and higher-level reasoning.

HOW AI CAN HELP

Automated data discovery

AI agents can continuously scan for and retrieve information from multiple structured and unstructured data sources, both classified and public—including briefings, news media, and other reports—reducing the need for manual searches and data collection.

Contextual linking of insights

If analysis of an old report produces new leads—such as a location or personal association—AI agents can automatically launch related queries to build a more complete picture of personal networks and activities.

Multimodal data integration

Specialized AI agents can process different types of inputs (e.g., text, imagery, sensor data) and combine the results into comprehensive, integrated insights.

Adaptive intelligence reporting

As the various agents collaboratively draft and refine reports, they can highlight key findings and update them dynamically as new data becomes available.

Smarter intelligence

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Automated agents might misinterpret ambiguous or misleading data, potentially affecting the accuracy of intelligence outputs. This is especially true for open-source information, which can be intentionally misleading or fake. AI agent actions should be validated against verified sources and continuously monitored to reduce errors.



Safe and secure

Given the sensitive nature of intelligence queries, special care must be taken to prevent adversaries from influencing the model (or gathering their own intelligence from what is queried).



Transparent and explainable

Providing adequate transparency in autonomous information exchange can pose significant technical and operational hurdles, especially given the complexity of integrating data from multiple sources. AI agents should document their reasoning and data sources so the outputs are easier to trust and audit.



Responsible and accountable

Given the life-and-death consequences of intelligence information, AI agents should be viewed as decision-support tools, with human analysts and leaders continuing to have final responsibility for decisions and actions.

POTENTIAL BENEFITS

Faster examinations at lower cost

Automating data collection and integration accelerates intelligence reporting and reduces manual effort, which can enable agencies to respond more quickly and efficiently to emerging threats.

Deeper insights

Cross-modal, autonomous querying and analysis can uncover hidden connections that might be missed with traditional manual processes, providing a more comprehensive and dynamic understanding of complex domains.

Improved analyst capacity

By offloading routine search and synthesis tasks, AI agents allow human analysts to focus on higher-level reasoning, strategic decision-making, and collaboration with colleagues.

Automated clinical documentation

Reducing provider burden by automating patient visit notes

Agentic AI systems can automatically generate structured clinical documentation from patient visits, reducing providers' workload, improving record quality, and enabling them to spend more time with patients.

ISSUE/OPPORTUNITY

For providers in large government health care systems, documenting patient visits can be an administrative burden. Traditional processes require physicians to spend hours after each visit manually writing SOAP (Subjective, Objective, Assessment, Plan) notes and updating electronic health records (EHRs). This

not only disrupts the flow of care but often requires after-hours "pajama time" spent on paperwork, which contributes to provider burnout. Also, incomplete or inconsistent documentation creates risks for care quality, regulatory compliance, and continuity across providers.

HOW AI CAN HELP

Automated transcription and generation

A generator agent listens to the audio of patient visits, converting it into structured SOAP notes and EHR-ready data, which minimizes the need for manual entry.

Validation and compliance

A validator agent reviews the generated notes for clinical accuracy, completeness, and regulatory compliance, confirming that records meet organizational and legal standards.

Workflow orchestration

An orchestrator agent manages the entire documentation process, routing tasks among the specialized agents and providing smooth, efficient workflows.

Seamless integration

The agentic AI system formats the validated documentation into common formats (e.g., DOCX, HTML) for direct integration into EHR systems, reducing friction and manual effort for providers and staff.

Automated clinical documentation

MANAGING RISK AND PROMOTING TRUST



Fair and impartial

Documentation agents should be designed and tested to work effectively across varied patient populations and care contexts.



Robust and reliable

Because errors in transcription or orchestration could compromise documentation quality, AI agents should be validated regularly and continuously updated to reflect evolving medical and compliance requirements.



Transparent and explainable

For providers to trust that AI-generated documentation accurately reflects patient encounters, AI agents must be able to explain how notes were created and why specific items were flagged for correction.



Private

Some patients might be uncomfortable having their medical visits recorded and then analyzed by AI. Every patient should be given the choice to opt out. Also, AI agents must strictly comply with HIPAA and other relevant regulations, encrypt audio and text records, and minimize retention of personally identifiable information.



POTENTIAL BENEFITS

Lower costs and reduced provider burnout

Automating documentation can reduce clinical staff workload, reducing costs and freeing providers from after-hours paperwork, which can improve their work-life balance.

Improved record quality and compliance

Standardized, AI-validated notes can improve the completeness and consistency of patient records, enhancing care continuity and compliance with regulations.

Data-driven insights

Structured documentation enables advanced analytics and quality improvement initiatives, helping government health systems improve care delivery at scale.



Virtual public servant

Citizen engagement

AI can enable virtual assistants that provide personalized responses to citizen questions about public services.

ISSUE/OPPORTUNITY

Government organizations perform a range of functions, from supporting public health to promoting tourism. Data about government and public services, however, is often stored in a variety of formats and locations (e.g., on-prem, cloud), challenging interoperability. When citizens contact agencies to inquire about

services and resources, human agents are challenged to rapidly access and summarize information to satisfy citizen questions. This is a time-consuming, labor-intensive endeavor for the organization, and it may not meet citizen expectations for fruitful engagement.

HOW AI CAN HELP

A digital agent for engagement

An AI-enabled virtual assistant can serve as the interface between citizens and government information, helping with questions and transactions via empathetic, natural language.

Reaching across datasets

The virtual assistant can distill and summarize information from myriad sources on a variety of topics to answer questions in a multitude of languages regarding service requirements and appointment options.

Virtual public servant

MANAGING RISK AND PROMOTING TRUST



Responsible and accountable

While virtual assistants may be valuable for providing information, they may not be suited to providing true insight and advice. Agencies need to guard against over-reliance on AI solutions and the potential for citizens to take some action based on faulty or improper AI output.



Safe and secure

A model tasked with providing accurate information may be a target for cyber criminals seeking to access sensitive information or manipulate the model and its underlying data. Many government agencies contend with cybersecurity regulations and standards, making model security a priority.



Robust and reliable

Model accuracy and timeliness depends in part on the data sources it can access. If information is outdated or incorrect, it creates a risk of erroneous outputs. Human stakeholders responsible for updating information have a direct impact on model reliability and user trust.



POTENTIAL BENEFITS

Promoting citizen engagement

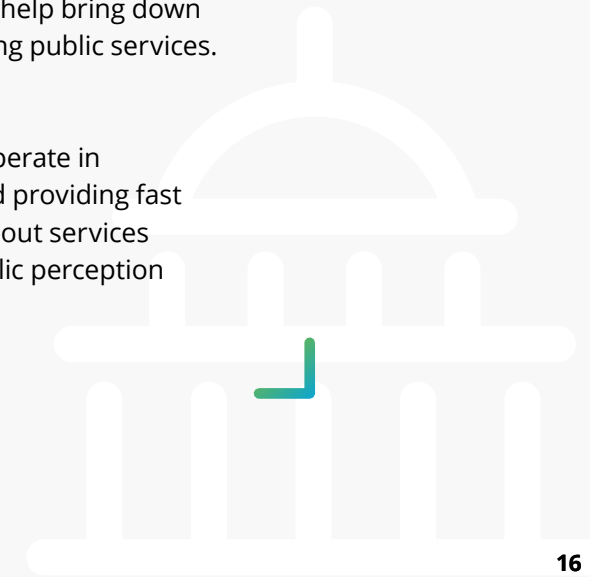
When public services are more accessible due to more efficient and robust technology, user engagement and citizen satisfaction in government offerings increase.

Increasing accessibility

A virtual assistant powered by AI can interact with citizens in their preferred language and ultimately help bring down social barriers to engaging public services.

Citizen satisfaction

Government agencies operate in service to the public, and providing fast access to information about services promotes a positive public perception of government function.



Insights for all

Knowledge management

AI can serve as an interface to help public sector organizations become insight-driven by making data more accessible.

ISSUE/OPPORTUNITY

From the census to transportation and procurement, government agencies collect and release huge amounts of open datasets. By encouraging the use, reuse, and distribution of open datasets, government organizations can promote data-driven innovation and

citizen-centric services if combined with an agency's internal datasets. For public industry stakeholders to become truly insight-driven, they require the means to interrogate all relevant data, even if they lack a technical background in data science or related fields.

Tags

Cross-functional

HOW AI CAN HELP

Greater accessibility

AI can provide a natural language interface that allows non-technical users to access and understand data that might otherwise only be accessible to technical users.

Democratizing insights

Rather than placing all of the burden for data analysis, interpretation, and visualization on a technical team, an AI interface reduces the effort by allowing more stakeholders to work with the data and derive their own insights.

Insights for all

MANAGING RISK AND PROMOTING TRUST



Safe and Secure

An AI model that taps into a variety of datasets can make it difficult to control which data is accessed by which stakeholders in which organization, raising important considerations for model security and governance.



Private

When dealing with sensitive and proprietary information that is subject to varying laws and regulations across jurisdictions, organizations are called on to ensure the AI model does not leak, inadvertently divulge, or inappropriately access sensitive or restricted data.



Transparent and explainable

To accurately interpret data and AI outputs, the end user needs to understand which data was referenced for the output, which could not be accessed, and the potential biases in the available data.



POTENTIAL BENEFITS

Scaling data access

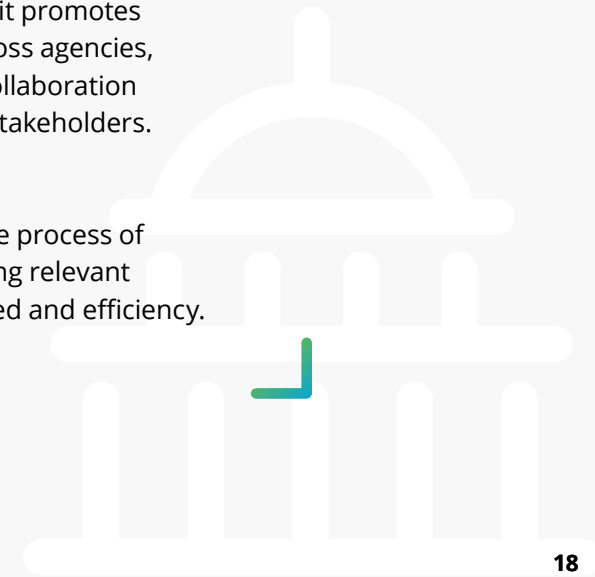
An AI solution that can access a variety of datasets and data types allows public servants to draw conclusions from a broader set of knowledge and information.

Fostering collaboration

When more public servants can access insights and knowledge, it promotes insight-driven action across agencies, helping to fuel greater collaboration between a larger set of stakeholders.

Faster insights

AI can help accelerate the process of identifying and consuming relevant information, driving speed and efficiency.



Simulating urban planning scenarios

Urban planning/future of cities

AI can be used to help urban planners in the ideation and design of novel urban concepts.

ISSUE/OPPORTUNITY

More than 56 percent of the world's population—4.4 billion people—lives in cities.¹ By 2050, the urban population is expected to double, with upwards of 70 percent of people living in cities. The scale and speed of urbanization brings a host of challenges, such as lack of affordable housing, overburdened transportation

systems, traffic congestion, lack of drinking water, sanitization issues, and environmental quality. The challenge for city officials and urban planners is to imagine the future of cities by overcoming creative hurdles and developing city designs that are resilient, sustainable, and human centric.

HOW AI CAN HELP

Generating 3D city models

Using AI, thousands of 3D images can be rapidly created to help guide and refine a city design. Such 3D images form part of the design brief for urban planners and the master city plan.

Simulate natural disasters

AI can simulate natural disasters like earthquakes, floods, or hurricanes to evaluate the vulnerability of city infrastructure and plan for resilient urban infrastructure.

Planning for the future

By simulating population growth and demographic trends, AI can develop scenarios for urban expansion and plan for adequate infrastructure, housing, transportation, and public services that accommodate urban growth.

Simulating urban planning scenarios

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

While AI models may create interesting or attractive designs, they require human review and validation to ensure they meet urban planning requirements and can be feasibly built in the real world.



Transparent and explainable

A lack of contextual knowledge of urban planning may lead AI to develop improbable scenarios, and analysts need to be able to understand how and why the model produced an output in order to confirm and validate it.



POTENTIAL BENEFITS

Super-charge creativity

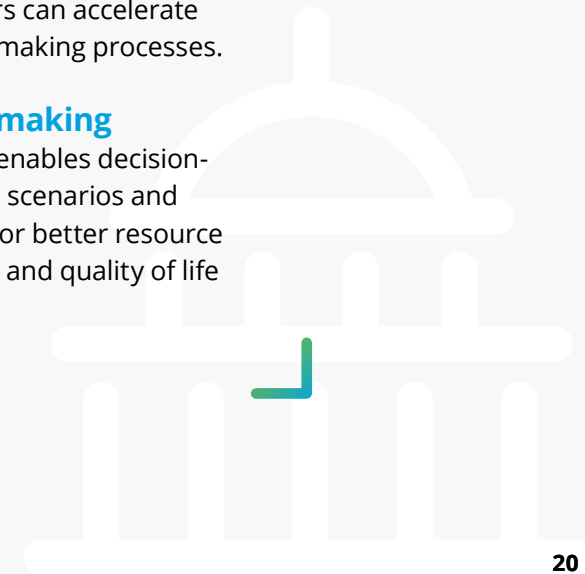
Using AI to rapidly create a plethora of designs and scenarios helps city officials imagine the future of cities and plan for upcoming challenges.

Faster ideation and iteration

With a faster method to create design iterations, urban planners can accelerate the design and decision-making processes.

Improved decision-making

Using AI in city planning enables decision-makers to model various scenarios and optimize urban designs for better resource utilization, sustainability, and quality of life for residents.



Education 2.0

Hyper-personalized education

AI can be used to hyper-personalize digital teachers that can adapt to student learning needs and curricula.

ISSUE/OPPORTUNITY

The demand for schoolteachers can often exceed supply. While the available teachers contend with larger class sizes, they also need to accommodate students with different learning styles and educational needs. Yet, because of

the one-to-many nature of traditional schools, teachers are challenged to deliver the kind of personalized learning support and instruction that students need to be successful.

HOW AI CAN HELP

A digital, adaptive teacher

AI can serve as a virtual instructor, drawing from resources and lesson plans to hyper-personalize the learning experience. The model can check the student's work and comprehension and adapt lessons and learning strategies according to the student's individual weaknesses, strengths, and preferences.

A force multiplier for teachers

When personalized digital teachers can work with students one-on-one to master new skills and knowledge, the human instructor can focus on higher level planning, interacting with students, evaluation, and student support.

Education 2.0

MANAGING RISK AND PROMOTING TRUST



Responsible and accountable

While digital teachers can offer valuable advantages in adaptive learning, the model should not be expected to satisfy all of the important lessons teachers impart, such as social lessons around collaboration, conflict resolution, and empathy. The human element in teaching is essential, and educational institutions need to take a responsible approach to integrating AI-enabled teachers.



Robust and reliable

Because AI is susceptible to outputting inaccuracies and hallucinations, there is a risk that a virtual teacher could teach incorrect facts or produce poor learning strategies.



Private

Student data is subject to education regulations, making model security and data privacy a priority when deploying digital teachers.



POTENTIAL BENEFITS

Catering to the student

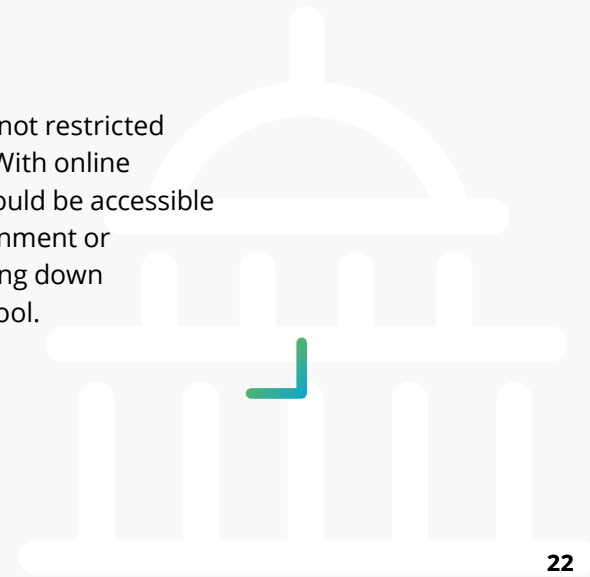
Employing adaptive learning with AI can promote knowledge retention and understanding by tailoring teaching approaches to the student's learning style.

Remedy the talent gap

Leveraging AI helps overcome teacher shortages, allowing more students to access quality education.

Removing barriers

An AI-enabled teacher is not restricted to a physical classroom. With online access, digital teachers could be accessible to students in any environment or geography, helping to bring down barriers to attending school.



Digitizing policymaking

Policy creation assistant

AI can be used to search large volumes of policy documents and output natural language responses to user queries in complex policy environments.

ISSUE/OPPORTUNITY

Because the data relevant to government agencies is stored in different locations and formats, it can be difficult for analysts and policymakers to effectively query datasets and retrieve relevant information in a timely manner. With nomenclature issues,

it can also be challenging to identify associated data topics and types. The result is a diminished ability to digitize policymaking and discussion, as well as increasingly complicated interactions around policy matters.

HOW AI CAN HELP

AI assistant

AI can identify data related to common themes and topics and then summarize that information in response to user queries, helping to identify policy differences, conflicts, and gaps.

Citizen engagement in policymaking

Using AI, governments can create interactive platforms and chatbots that encourage citizens to participate in policymaking discussions. The AI-driven interface can gather public opinions and feedback on policies, making it easier for citizens to voice their views.

Digitizing policymaking

MANAGING RISK AND PROMOTING TRUST



Private

Some of the data relevant to policy issues may be sensitive or restricted, and the AI model may require controls to limit which users can access which datasets.



Fair and impartial

Various stakeholders aim to influence policymaking. AI might be biased in giving higher weight to comments and input coming from some sources over others. This has the potential to produce biased policies that favor certain businesses or segments of society.

POTENTIAL BENEFITS

Data query at scale

By reviewing large volumes of policy documents, AI can help users accelerate information gathering and increase their capacity and efficiency in querying disparate datasets.

Participatory policymaking

Using AI to better identify and incorporate a varied views and stakeholders supports more robust and complete representation in policy matters.

Drafting contracts and SoWs

Procurement

AI can analyze offerings from existing vendors, match them to organizational needs, generate requests for proposals, and analyze the responses.

ISSUE/OPPORTUNITY

Governments procure billions of dollars in goods and services annually.² Traditionally, government procurement requires significant volumes of paperwork, which can lead to delays. Many government procurement contracts are highly detailed and often incorporate a range of clauses

and requirements from payment terms to export controls to wage and workforce requirements. Drafting requests for proposals (RFPs) and contracts and then generating statements of work (SoWs) requires significant time and resource investments.

Tags

Procurement/Sourcing & Supply Chain

HOW AI CAN HELP

Automated drafting

AI can automate the RFP and SoW writing processes by generating the initial drafts based on templates, historical documents, or specific prompts provided by procurement officials.

Extracting information

AI's advanced Natural Language Processing (NLP) capabilities can help extract relevant clauses and requirements from existing contracts, SoWs, and legal documents. Such information can be used to either create new contracts or assess the risks posed by existing contracts.

Drafting contracts and SoWs

MANAGING RISK AND PROMOTING TRUST



Transparent and explainable

AI may not be able to explain why certain clauses are added to a contract while others are excluded, which is vital information for the human user validating the outputs.



Private

Ingesting existing and historical contract data may pose data privacy and legal hurdles. Model governance is necessary to ensure the AI model, as well as the underlying data, meet privacy rules, regulations, and standards.

POTENTIAL BENEFITS

Time savings

Creating initial document drafts with AI can expedite the writing process and lead to significant time savings, compared to manually creating each RFP or SoW from scratch.

Improved consistency

AI can develop drafts while adhering to predefined guidelines in prompts, which supports a greater level of consistency across report writing.

Multilingual citizen services

Service delivery

AI can help with language translation to support the delivery of services to citizens.

ISSUE/OPPORTUNITY

Many governments around the world serve populations with varying language proficiency and linguistic backgrounds. This challenges agencies to develop multilingual websites, translate official documents, and support frontline workers with translation tools so they can better communicate with all citizens.

Tags

Customer Service

HOW AI CAN HELP

Aiding frontline workers

AI can be used to create real-time audio and text messages in different languages as frontline workers interact with residents around a variety of services, such as social care, health care, and emergency response.

Announcement and website translation

Government websites and public information (e.g., health and travel advisories) can be translated quickly to make essential information more accessible to the population.

Translating official documents

Government agencies often handle the publication of official documents, laws, regulations, and policies. AI can streamline the translation process and help produce accurate and consistent translations.

Multilingual citizen services

MANAGING RISK AND PROMOTING TRUST



Fair and impartial

The data used to train an AI model for use in translation may not be consistently accurate or robust across all languages, which could in turn lead to poorer translations and less access to citizen services for some language speakers than for others.



Private

The translating model may be exposed to sensitive information, necessitating steps to ensure the model does not mishandle or inappropriately divulge protected data and thus violate data privacy regulations.

POTENTIAL BENEFITS

Real-time translation

When audio or text can be translated into a multitude of languages in real time, it enables more seamless and conversational interactions with different language speakers.

Translation at scale

AI can handle large volumes of document translation, giving an agency capacity to ensure government information and services are accessible to a greater audience.

Summarizing legislative documents

Legislative administration

AI can help legislative staff more rapidly transcribe and summarize hearings, legislation, documents, and official announcements.

ISSUE/OPPORTUNITY

Legislative offices are expected to hold hearings on important topics, respond to constituents, and make public announcements in the form of press releases. Manually transcribing hearings and meetings is a time-consuming task.

Further, developing new legislation (where staff play a pivotal role in research) requires sifting through voluminous policy proposals and research published by experts.

HOW AI CAN HELP

Summarizing official documents

Auto-generating transcripts of hours-long committee hearings and summarizing important bills and hearings can significantly reduce the administrative burden on staffers.

Process and summarize policy proposals and research

Legislative staff review a large volume of policy proposals and recommendations published by experts. AI can quickly summarize the documents for them, so staffers can spend more time on higher level policy analysis and decision-making.

Summarizing legislative documents

MANAGING RISK AND PROMOTING TRUST



Fair and impartial

AI may perpetuate latent biases based on its training set and generate skewed summaries that are partisan and favor certain ideologies.



Private

Ingesting internal policy proposals can expose sensitive information, requiring organizations to take measures that protect the confidentiality of internal documents.


POTENTIAL BENEFITS

Reducing burdens

Generating summaries of official hearings can reduce administrative burdens on legislative staff so they can focus on more complex tasks.

Saving time

AI can quickly retrieve information and summarize it, saving legislators and staff time when reviewing lengthy, complex, or detailed documents.



Global policy tracking

Automated tracking and analysis of public policy

Organizations in the public and private sector can use AI to monitor, interpret, and analyze public policy developments in real time across hundreds of countries.

ISSUE/OPPORTUNITY

Tracking national policy developments on a global scale is a resource-intensive and highly fragmented process. Policy documents vary in language, structure, formatting, and accessibility, making it

difficult for international organizations, governments, and advocacy groups to maintain a coherent and timely view of global policy trends.

HOW AI CAN HELP

Data collection and analysis

AI can automatically gather, structure, and analyze vast volumes of policy documents from government websites and public sources around the world. The technology performs multilingual data extraction, applies natural language processing to categorize and summarize policies, and synthesizes insights into structured outputs that can be validated by subject-matter experts.

Global policy tracking

MANAGING RISK AND PROMOTING TRUST



Fair and impartial

The AI model should be designed to avoid reinforcing systemic biases. Human experts from different backgrounds and regions are embedded in the feedback loop to validate model outputs, helping to ensure representation across geographies and policy contexts.



Robust and reliable

Automated systems should undergo rigorous, iterative testing to help ensure the reliability of outputs. Policy insights are continuously benchmarked against human analysis and real-world policy documents to maintain a high level of accuracy and dependability, especially in politically sensitive or under-reported regions.

POTENTIAL BENEFITS

Greater accuracy, scalability and knowledge sharing

AI can enable ongoing, real-time monitoring and analysis of thousands of policies across hundreds of countries without an exponential increase in manual effort. It can also enable local entities and other stakeholders to identify global policy trends, compare regional approaches, and uncover best practices.

Improved efficiency

Combining data collection and policy analysis with AI-powered automation can save organizations significant time and resources.

Improved decision-making

Use of AI for policy tracking and analysis can provide organizations with more timely, structured, and reliable data to support effective decision-making and planning.

Autonomous patrolling and threat detection

Robotic patrolling for government facilities

Autonomous mobile robots (AMRs) patrol buildings and facilities, continuously monitoring environments, detecting potential threats, and escalating incidents based on severity. Systems interpret scenes locally and alert security personnel when intervention is required.

ISSUE/OPPORTUNITY

Continuous physical security coverage is costly and difficult to staff with humans alone. Fixed cameras lack mobility and context. Security guards patrol facilities on fixed routes or remain stationed at checkpoints, creating predictable gaps in coverage that leave areas unmonitored for extended periods.

Fatigue and attention degradation affect human guards during overnight shifts or long patrol routes, reducing vigilance when monitoring repetitive environments where incidents are infrequent. Guards cannot be in multiple locations simultaneously, forcing

facilities to choose between comprehensive coverage requiring large security teams or accepting blind spots during routine patrols. Fixed cameras provide static views but cannot investigate suspicious activity, follow subjects through facilities, or provide contextual assessment of situations that require mobile perspective.

Physical AI robots can provide persistent, mobile monitoring while reserving human attention for verified incidents requiring judgment, de-escalation, or physical intervention.

HOW AI CAN HELP

Autonomous navigation

Robots move through buildings independently, following patrol routes, navigating around obstacles, using elevators, and adapting paths when areas are blocked or access is restricted.

Threat escalation logic

Events are classified by severity, with minor anomalies logged for review while urgent threats trigger immediate alerts with live video feeds to security personnel.

Visual scene interpretation

AI identifies abnormal situations such as unauthorized personnel in restricted areas, doors left open after hours, unattended packages, or environmental hazards like water leaks or smoke.

Continuous operation

Robots patrol without fatigue, maintaining consistent coverage during overnight hours, weekends, and holidays when staffing human security teams is most expensive and challenging.

Human alerting

Security teams are notified when needed, receiving situational context, robot location, and video evidence to inform their response without requiring constant monitoring of robot feeds.

Autonomous patrolling and threat detection

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Robots that misclassify routine activity as threats create alert fatigue that causes security teams to discount warnings; those that miss genuine incidents defeat the purpose of deployment. Both failure modes are serious. Reliability across varying lighting, occupancy levels, and facility configurations is essential for the system to deliver sustained operational value.



Private

Autonomous robots continuously recording video as they patrol occupied facilities capture detailed data about individuals' movements and behaviors throughout the working day. Organizations must define clear limits on what is recorded, how long footage is retained, who can access it, and what prevents patrol data from being used for employee monitoring beyond legitimate security purposes.



Fair and impartial

Threat classification models trained on specific facility types or populations may generate disproportionate false positives for individuals whose appearance or behavior differs from training data norms. In government facilities serving the public, biased threat detection creates civil rights exposure and undermines the institutional trust on which legitimate security operations depend.

POTENTIAL BENEFITS

Improved coverage

Facilities are monitored continuously as robots provide round-the-clock patrols without breaks, shift changes, or attention lapses that create security gaps.

Faster detection

Threats are identified earlier through constant mobile monitoring that catches developing situations before they escalate, rather than discovering incidents only during periodic guard rounds.

Security consistency

Coverage does not degrade over time as robots maintain identical patrol patterns and alertness levels regardless of hour, day, or duration of deployment.

Lower labor dependence

Fewer routine guard tasks as robots handle repetitive patrols, allowing security personnel to focus on response, investigation, and situations requiring human judgment.

Audit-ready physical decision trails

Sensor data, movement paths, and escalation decisions can be logged for post-incident review, public accountability, and legal scrutiny.

Smart city operations and urban infrastructure modernization

City-scale intelligence for monitoring, service delivery, and resilient operations

Physical AI combines urban sensors, drones, and autonomous robots to monitor and manage city infrastructure in real time. Edge intelligence detects events, sending data to cloud platforms for analysis and decision-making, while robots perform tasks like inspection and maintenance, linking sensing and action across the city.

ISSUE/OPPORTUNITY

Smart city initiatives aim to deliver dense, sustainable urban environments with high service quality and operational efficiency. In practice, however, both city infrastructure and large public complexes still rely heavily on fragmented data, manual inspections, and labor intensive operations. Urban systems—transportation, utilities, public safety, and environment—operate in silos with limited real time, street level visibility, making it difficult to detect issues early, coordinate responses, or prioritize maintenance. At the same time, large residential and commercial developments depend on sizable human workforces for cleaning, maintenance, and deliveries, resulting in high operating costs, uneven service quality, and limited scalability as demand grows.

The opportunity is to deploy physical AI as foundational urban infrastructure. Edge integrated sensors, drones, and autonomous mobile robots enable continuous, real time monitoring and verification across cities and facilities. By detecting hazards at the edge, streaming curated events to the cloud, and unifying insights across systems, physical AI enables coordinated dispatch, proactive intervention, and demand driven resource allocation—delivering more consistent services, improved resident experience, and more resilient urban operations without relying on static schedules or manual oversight.

Tags

Operations

Physical AI

HOW AI CAN HELP

City-scale sensor and edge integration

AI aggregates data from transportation, utilities, and environmental sensors with edge compute to detect anomalies locally (e.g., hazards, leaks, congestion) and publish event alerts with minimal latency.

Robotic service execution

AMRs handle cleaning, waste handling, internal logistics, and routine maintenance across public buildings, transit hubs, campuses, and residential complexes. Service schedules adapt dynamically based on actual usage patterns.

Infrastructure health monitoring

Detects stress and failure in urban systems by analyzing diverse data such as bridge vibration patterns, water flow irregularities, or grid fluctuations, identifying maintenance needs before failures disrupt service.

Decision support for operators

AI informs planners without exerting autonomous control, providing recommendations that human operators review and approve before implementation.

Human governance and oversight

Public authorities retain control over policy decisions and operational changes, with AI systems providing analysis and recommendations.

Mobility and crowd analysis

Improves traffic and pedestrian flow by identifying congestion patterns, predicting rush-hour bottlenecks, monitoring crowd density at public events, or detecting unusual movement patterns.

Smart city operations and urban infrastructure modernization

MANAGING RISK AND PROMOTING TRUST



Private

City-scale sensors, drones, and robots operating continuously across public spaces and residential areas create a mass surveillance capability covering entire urban populations that have no practical ability to opt out. Governments must define strict, publicly documented limits on what is collected, how long it is retained, and what prevents repurposing beyond the infrastructure management purposes that justified deployment.



Responsible and accountable

The use case preserves human governance by design; however, it is easy for AI recommendations to become de facto decisions as alert volumes grow and operator capacity is stretched. AI recommendations must be clearly distinguished from human decisions in operational records, and accountability frameworks must be actively maintained, not assumed to hold just because that was the original design intent.



Fair and impartial

Governments have fairness obligations that private sector entities do not. AI-driven urban resource allocation that systematically directs maintenance and services toward higher-density or more commercially active areas is not just an operational shortfall—it is a failure of the public duty that justifies the deployment. Monitoring for equitable outcomes across the full jurisdiction must be built into governance from the outset.

POTENTIAL BENEFITS

Reduced operating costs

Replaces large human teams performing repetitive tasks with autonomous systems, significantly lowering facility management expenses while enabling smart city developments to achieve financial sustainability.

24/7 facility operations

Enables round-the-clock service delivery—cleaning, deliveries, maintenance—without shift-based labor costs or coordination complexity, supporting the always-on nature of smart city infrastructure.

Greater urban resilience

Cities respond more effectively to stress and incidents, as cross-system visibility enables faster detection of cascading failures, coordinated response across departments, and proactive mitigation of predictable disruptions.

Consistent service quality

Eliminates variability in service delivery based on staffing availability, ensuring all residents receive uniform maintenance and support as envisioned in smart city livability standards.

Better-informed planning

Decision-makers gain clearer visibility into city operations, enabling evidence-based infrastructure investments, more accurate service demand forecasts, and better understanding of how proposed changes might affect interconnected systems.

Contacts



Beena Ammanath
Global Deloitte AI Institute
Leader, Deloitte AI Institute
United States, Lead
Deloitte Consulting LLP
bammanath@deloitte.com



Tiago Durao
Deloitte AI Institute
Portugal, Lead
Deloitte Portugal
tdurao@deloitte.com



Geert Hallemeesch
Deloitte AI Institute
Belgium, Lead
Deloitte Belgium
ghallemeesch@deloitte.com



Chris Lewin
Deloitte AI Institute
Asia Pacific Lead
Deloitte Singapore
chrislewin@deloitte.com



Audrey Ancion
Deloitte AI Institute
Canada, Lead
Deloitte Canada
knuttall@deloitte.com



Richard Eudes
Deloitte AI Institute
France, Lead
Deloitte France
reudes@deloitte.fr



Jan Hejtmanek
Deloitte AI Institute
Central Europe, Lead
Deloitte Central Europe
jhejtmanek@deloitte.com



Wessel Oosthuizen
Deloitte AI Institute
Africa, Lead
Deloitte Africa
woosthuizen@deloitte.com



Naser Bakhshi
Deloitte AI Institute
Netherlands, Lead
Deloitte Netherlands
nbakhshi@deloitte.nl



Roman Fan
Deloitte AI Institute
China, Lead
Deloitte China
rfan@deloitte.com



Prashanth Kaddi
Deloitte AI Institute
India, Lead
Deloitte India
pkaddi@deloitte.com



Sulabh Soral
Deloitte AI Institute
United Kingdom, Lead
Deloitte United Kingdom
ssoral@deloitte.com



Dr. Bjoern Bringmann
Deloitte AI Institute
Germany, Lead
Deloitte Germany
bbringmann@deloitte.com



Tomas Meca Figueres
Deloitte AI Institute
Spain, Lead
Deloitte Spain
tomecafigueras@deloitte.es



Sultanbek Khunkaev
Deloitte AI Institute
Middle East, Lead
Deloitte United Arab Emirates
sukhunkaev@deloitte.com



Dr. Elea Wurth
Deloitte AI Institute
Australia, Lead
Deloitte Australia
ewurth@deloitte.com.au



Martin Cabrera
Deloitte AI Institute
Chile, Lead
Deloitte
mcabreraa@deloitte.com



Alfredo Maria Garibaldi
Deloitte AI Institute
Italy, Lead
Deloitte Italy
agaribaldi@deloitte.it



Tomotake Kozu
Deloitte AI Institute
Japan, Lead
Deloitte Japan
tomotake.kozu@tohatsu.co.jp



Jefferson Denti
Deloitte AI Institute
Brazil, Lead
Deloitte Brazil
jdenti@deloitte.com



Nicolas Griedlich
Deloitte AI Institute
Luxembourg, Lead
Deloitte Luxembourg
ngriedlich@deloitte.lu



Carlos Labanda
Deloitte AI Institute
South-Latin America, Lead
Deloitte Colombia
clabanda@deloitte.com



This communication contains general information only, and none of Deloitte Touche Tohmatsu Limited, its member firms, or their related entities (collectively, the “Deloitte Network”) is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser. No entity in the Deloitte Network shall be responsible for any loss whatsoever sustained by any person who relies on this communication.

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.

Copyright © 2026 Deloitte Development LLC. All rights reserved.

Endnotes

1. “Overview,” Urban Development, The World Bank, accessed September 7, 2023.
2. US Government Accountability Office (GAO) analysis of Federal Procurement Data System, June 2025.