SAP Institute for Digital Government
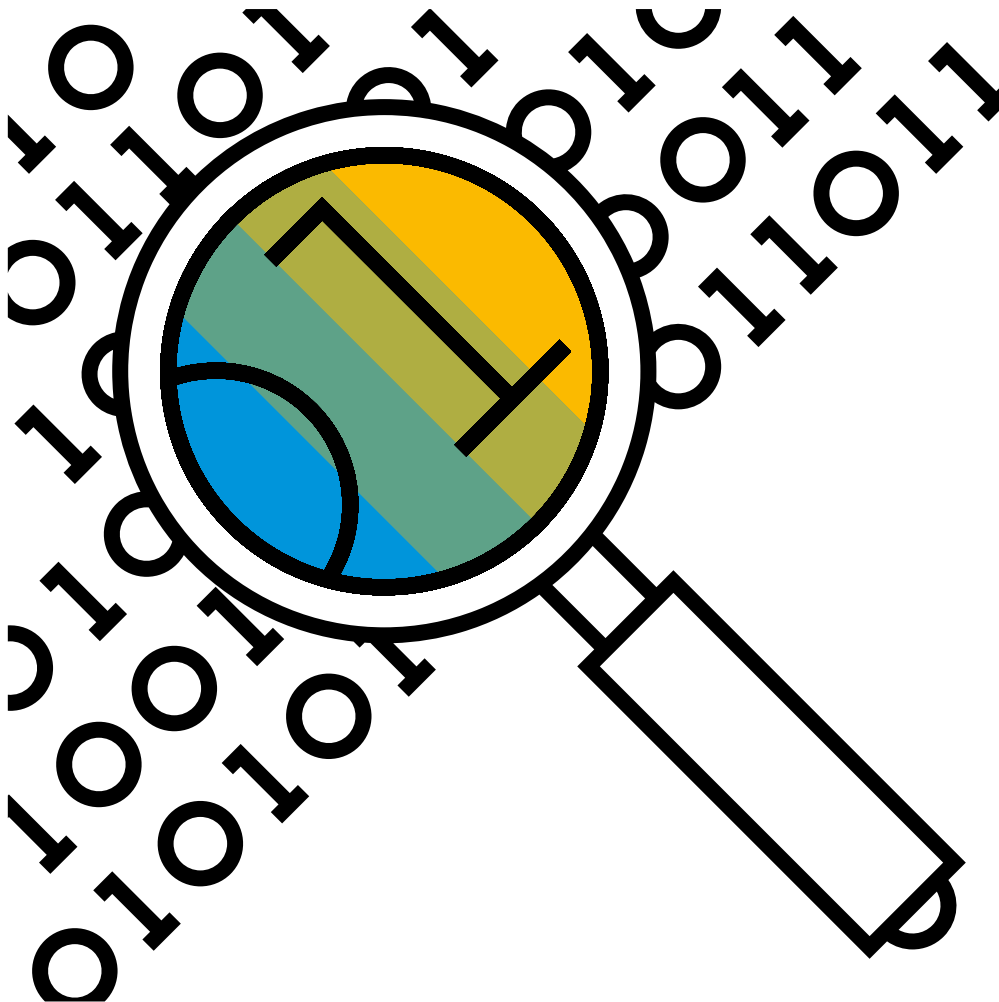
SAP Thought Leadership Paper | PUBLIC
Public Sector

Building Artificial Intelligence Capability in the Public Sector
Accelerating the Adoption of Artificial Intelligence in Government Organizations

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

THE BEST RUN SAP

Table of Contents

Authors
Doctor Tapani Rinta-Kahila,
UQ Business School, The University of Queensland
Doctor Ida Someh,
UQ Business School, The University of Queensland
Professor Marta Indulska,
UQ Business School, The University of Queensland
Ian Ryan,
SAP SE

Government bodies need to make a shift towards becoming more results-oriented organizations with flexibility, entrepreneurialism, and accountability for performance. While COVID-19 has accelerated this shift, artificial intelligence (AI), when deployed and managed correctly, has a key part to play in the transition. The technology holds a promise of revolutionizing the efficiency and flexibility of public services.

But public organizations can be slow to make this change. Various rules, policies, and legislation constrain their ability to drive transformative change.

So how do we build public sector AI capability?

Challenges facing the Public Sector

Public organizations are expected to adopt transformative technologies, increase efficiency, and reduce costs, whilst fulfilling their main roles and functions in serving citizens, such as providing healthcare, safeguarding consumer rights, and protecting national borders. These complex demands present them with unique challenges for developing and implementing AI systems. AI is not a magic bullet: it can yield quick wins but achieving its true potential requires a long-term commitment with assured investments. You can build a road, but without maintenance that road will soon fall into disrepair and need to be closed.

In an earlier report, we introduced 'capability development' as part of a multi-pronged approach to deliver successful AI programs in public sector[1]. In this thought leadership paper, we delve deeper into capability development. Drawing on our insights from studying several government AI projects, we elaborate on five specific AI capabilities public sector organizations could develop, and we also provide related recommendations.

OUR RESEARCH

To understand what capabilities are needed to successfully develop, deploy, and maintain AI systems in the public sector, we studied three AI projects undertaken by Australian public agencies. In each project, we interviewed key stakeholders, including data scientists, business-domain managers, and business architects. Table 1 summarizes the context, background, and aims of each project we drew upon in this thought leadership paper.

Table 1: The AI Projects Studied for this Paper

Project title	Public sector organization	Goal	Stage
Health AI	A state hospital's emergency department	To save lives by identifying patients at risk of developing sepsis while awaiting treatment	Proof of concept developed; yet to be implemented
Police AI	A state police department	To make police investigation data accessible and usable by enabling digitized keyword search	Implemented and in use
Tax AI	A state tax revenue management office	To improve taxpayer services and overall tax debt collection rates by identifying taxpayers at a risk of becoming chronic debtors	Implemented and in use

1. Rinta-Kahila, T., Someh, I., Gillespie, N., Indulska, M., Gregor, S., Van Leent, R., & Ryan, I., (2020). '**Delivering AI Programs in the Public Sector: Guidelines for Government Leaders**', University of Queensland and SAP SE Thought Leadership Report.
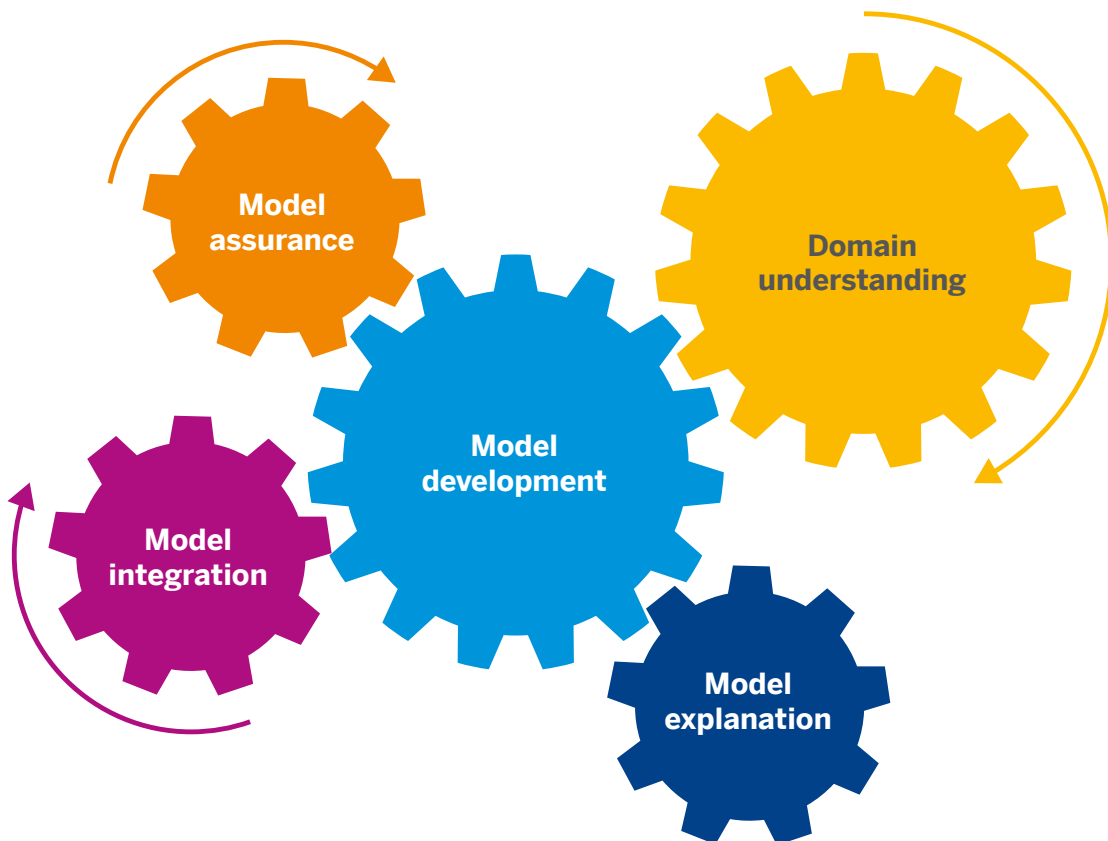
Five AI Capability Areas

Analysis of our case study data pointed to five distinct AI capability areas: model development, domain understanding, model explanation, model integration, and model assurance. Each capability is discussed as follows: first, we define the capability; then we describe which aspects the capability involves, especially in terms of resources and activities; and finally, we discuss how public sector organizations can go about building the capability.

The capability areas we propose reflect the nature of AI systems, in that at their core sits a machine learning (ML) model. An ML model is trained using existing data to recognize certain patterns in real-world problem domains. The learning process can be supervised (a specific learning goal is assigned) or unsupervised (the algorithm freely explores relationships between variables without human tagging or labelling). Once the model is trained, it can reason over new instances of data, select actions, and perform tasks and routines. A unique feature of AI models is their 'online learning', which ensures continued model training and evolution based on streaming data after they are deployed.

The five capability areas are inter-related and mastering all is the key for successful and meaningful use of AI in public sector.
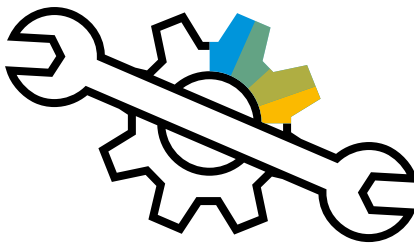
Capability Area 1: Model Development

DEFINITION
An organization's ability to train a set of accurate and robust ML models to support various organizational tasks.

KEY ASPECTS
To begin building AI models, public organizations must access or create suitable training data sets, test different competing ML algorithms on the data, and assess the models' performance against various criteria. These assessments need to include ethical considerations during model development to ensure any negative unintended consequences, such as biases and discrimination, are identified and rectified before model is put in use. This requires both technical and human resources. Our case analysis pointed to three key aspects that model development capability was nurtured:

DATA CURATION
(accessing legacy data, generating new data, and sharing data across the organization)

MODEL CO-PRODUCTION
(partnering with universities, vendors and other public organizations, sharing data and models outside of the organization)

HUMAN TALENT ACQUISITION
(hiring data scientists, 'ethicists', 'guardrail managers', training internal staff)

Role requirement: this capability is typically driven by data scientists and supported by data engineers and domain experts.

- Data curation
- Model co-production
- Human talent acquisition

HOW TO BUILD THE CAPABILITY

The ability to build robust ML models depends on the availability of, and access to, high-quality data. Many government organizations run legacy systems that lack features, scalability, and integration possibilities. Data is scattered across different departments and siloed in their proprietary systems instead of being shared and easily accessible across the organization. Therefore, creating a model development capability in the public sector hinges on investments in cloud platforms that not only provide ready access to novel ML tools, but also help break down data silos.

Health AI and Police AI opted to use a commercially-available cloud platform because it gave them an assortment of analytics products to choose from, based on their specific development needs, and allowed them to retain the IP of the developed models. To enable widespread data use, the available data had to be catalogued, typically using proprietary tools, to ensure shared understanding of data across various employee groups.

Sometimes the existing data is not of sufficiently high quality or lacks important attributes. This means that new data needs to be gathered. In Health AI, significant datasets were generated to establish an accurate 'ground truth' for predicting sepsis. The clinical director reported: *"We cleared about half a million patient records… it took me five years to do (the first) 1,200 patients."* The team then iterated between different models, such as logistic regression and deep learning, to identify models that deliver the best performance.

Model development capability also requires human talent, including data scientists who can work with complex datasets and models. However, according to the managers we interviewed, private sector offers higher salaries, which can lead to a high attrition rate for data scientists in the public sector. The lack of financial support for AI projects was noted by many of our interviewees: risk-averse public organizations are not easily willing to invest significantly in hiring highly paid experts to work on uncertain AI projects. Organizations we studied had coped with this challenge by partnering with other public organizations. Co-production is common in delivering public services especially in dynamic environments, where a single organization may not be able to keep up with changing needs and aspirations of the public.

Collaborating with other hospitals helped the Health AI team to validate and improve their model by giving them access to larger datasets: *"We shared the code and that's going across everywhere"* (clinical director). They also partnered with universities, as these institutions are host to experts in data science and user experience design and are motivated to collaborate for the greater good and access to research data.

In health, collaboration with private sector partners and other public sector bodies is **crucial to helping fund** the development of AI systems.

Capability Area 2: Domain Understanding

DEFINITION
An organization's ability to develop new or enhanced knowledge of business domains.

KEY ASPECTS
AI promises to deliver new insights about the domain of implementation, which can help stimulate innovation and new ways of working[2]. This requires a curious and critical mindset from the implementers. Indeed, across the projects we studied, AI development entailed a critical assessment and scrutiny of various taken-for-granted assumptions within the business domain. Three key aspects of developing a domain understanding capability are:

CODIFYING CURRENT IMPLICIT KNOWLEDGE
(teaching algorithms through tagging or labelling data points with insights and experiences of domain experts about decision attributes and outcomes)

UPSKILLING DOMAIN EXPERTS
(AI concepts and tools, data value, ethical issues, AI-based process management)

CHANGING INCUMBENT PARADIGMS
(challenging existing decision assumptions, generating new domain-specific knowledge, embracing algorithm-driven insights)

Role requirement: this capability is typically driven by domain experts and end users and supported by business managers and data scientists.

- **Codifying current implicit knowledge**
- **Upskilling domain experts**
- **Changing incumbent paradigms**

---

2. Someh, I., Wixom, B., & Zutavern, A. (2020). 'Overcoming Organizational Obstacles to Artificial Intelligence Value Creation: Propositions for Research', in Proceedings of the 53rd Hawaii International Conference on System Sciences, 5809–5818.

HOW TO BUILD THE CAPABILITY

In developing ML models, particularly when supervised algorithms are used, domain experts need to teach the algorithms the knowledge and rules of the application domain. As machines learn from codified knowledge (for example, outcomes that humans tag or label), domain experts need to be freed from their routine work to create the necessary data points for the machine to learn from. As the algorithms learn from the data, they will shed light on new insights that can potentially change incumbent ways of working and paradigms.

Changing incumbent paradigms of thinking is difficult and requires substantial organizational flexibility and openness to learning. Our case organizations achieved augmented domain understanding through active collaboration between domain experts and data scientists. Typically, the process was facilitated by creating a space for domain experts and data scientists to come together and develop ML models iteratively. As data scientists exposed insights from the data to domain experts, domain experts became more cognizant of factors that they had previously overlooked.

Doing this successfully requires coordination, which was made possible by creating positions that bridge the two groups. For instance, in both Health AI and Tax AI, one manager who had an adequate degree of understanding of both topic domain and data science was put in charge to orchestrate the interaction between the groups and also act as a 'translator'. The clinical director of Health AI commented: *"My job is just to sit back and make sure it's safe and validate things they come up with".*

A product owner of Tax AI explained how carrying out some initial descriptive data analyses on tax-debt recovery prompted them to question some of the business's incumbent assumptions early on. They discovered that a cohort of young citizens were failing to pay not because they had financial issues but because they were being contacted by physical letters, a medium that this cohort overlooked. This finding prompted process changes toward digital communication channels.

"We obviously have to **change what we're doing** if we want the taxpayer to change" – Tax AI product owner.

Capability Area 3: Model Explanation

DEFINITION
An organization's ability to communicate an ML model's goals, logic, and outputs to various stakeholders.

KEY ASPECTS
Considering that transparency of operations is a critical requirement for public organizations, managers need to develop capability to explain the goals, mechanics, and outputs of their AI models for different stakeholders[3]. This is a real challenge because many complex AI models operate in an inscrutable manner that is not understandable to humans[4]. Therefore, model explanation capability needs to build on a unique combination of technical tools and human skills that reveal machine insights to human decision-makers and help them make sense of the insights[5]. It covers the following aspects:
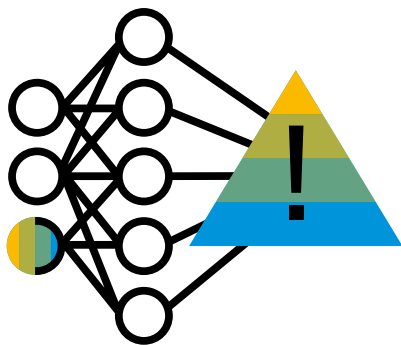
- **Model explainability**
- **Traceability**
- **Visualization**

EXPLAINING MODEL LOGIC TO STAKEHOLDERS
(communicating AI decisions, providing interpretability and guidance on the decisions including limitations associated with probabilistic nature of AI)

TRACING OUTPUTS INTO INPUTS
(mechanism for technical transparency and validation of model decisions, identifying sources of bias)

VISUALIZING MODEL OUTPUTS TO DOMAIN EXPERTS
(building explanatory system interfaces that are user friendly and intuitive)

Role requirement: this capability is typically driven by algorithm explainers and supported by data scientists and domain experts.

3. Someh, I., Wixom, B.H., Beath, C., & Zutavern, A., (2022) 'Building an Artificial Intelligence Explanation Capability', Forthcoming in MIS Quarterly Executive.
4. Asatiani, A., Malo, P., Nagbøl, P.R., Penttinen, E., Rinta-Kahila, T., and Salovaara, A., (2020). 'Challenges of Explaining the Behavior of Black-Box AI Systems'. MIS Quarterly Executive, 19, 4, 259–278.
5. Rinta-Kahila, T., Someh, I., Indulska, M., Ryan, I., and Van Leent, R., (2021). 'Building Explainability into Public-Sector Artificial Intelligence', The University of Queensland and SAP SE.

HOW TO BUILD THE CAPABILITY

In the development phase, data scientists need to explain the logic behind the ML models to domain experts to ensure user acceptance and effective use of the models. When building sophisticated ML models, such as those leveraging deep learning algorithms, full transparency of the models might not be feasible. Such complex models learn autonomously from data and update their knowledge across layers, making it impossible to track how inputs are transformed into outputs.

In these cases, public organizations must ensure that they can provide sufficient explanations for the models' behaviour and how it has reached to its particular conclusions (i.e., how inputs are transformed into outputs). When full transparency of the decision-making process is not feasible, there is a need for humans to qualitatively review and validate AI decisions. This can occur by building visualizations that show or expose AI decisions to the relevant domain experts.

Integrating explanations in the AI system's domain use requires the input of system developers whose job is to wrap ML models into consumable user interfaces. The Health AI project team built an interface that communicated the model's outputs, decision-making features, and confidence levels associated with each decision. The team leveraged human-centered design approaches to ensure the interface offered high usability for doctors and nurses. Partnerships with universities provided vital input to user-interface development. For example, the developers collaborated with a design school: *"They have that strong human-centred design expertise […] before we invested heavily in a big ICT integration… working with those stakeholders to absolutely confirm that what we're doing has every possible chance of succeeding"* (data-analytics director). Using colours and other cues was key to creating interfaces that clearly communicate relevant information to the domain expert.

When developing the Health AI technology, the development team **created a simple yet informative user interface** by collaborating with a design school.
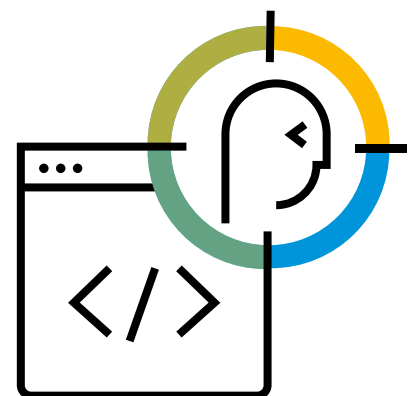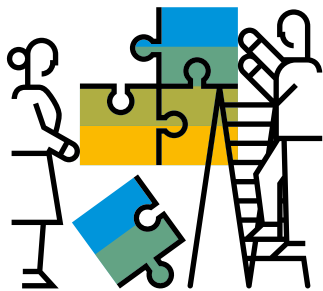
Capability Area 4: Model Integration

DEFINITION
An organization's ability to integrate ML models into existing technical interfaces, work processes, products, platforms, and structures.

KEY ASPECTS
Once developed, proof-of-concept AI models must be turned into usable products or services that can be deployed to support public service provision. The ML models need to be integrated into service platforms, various enterprise software systems, and existing workflows. The process involves clarifying the role of AI in performing work (for example, online learning, autonomous action, and automation versus augmentation), changes to existing product or service design, connectivity with other systems and services (for example, through APIs), and creating a positive user experience for domain experts or other users. Thus, it encompasses the following aspects:

CHANGE MANAGEMENT
(ensuring end-user acceptance, process impact and changes to metrics, internal and external user experience management)

TECHNICAL INTEGRATION
(building an AI interface into an existing workflow, establish pipelines for data flow, IT infrastructure)

WORK REDESIGN
(changing processes to align with and support data-driven decision-making, automation versus augmentation, providing guardrails for AI use)

Role requirement: this capability is typically driven by IT developers and UX designers, supported by data scientists and domain experts.

- **Change management**
- **Technical integration**
- **Work redesign**

HOW TO BUILD THE CAPABILITY

The Police AI project provides an illuminating example on how technical integration can be achieved by using AI systems in conjunction with flexible cloud platforms. The police department's data-analytics manager explained how their various databases hosted significant amounts of digitized case data from police investigations. The data consisted of scanned images of police officers' handwritten notes that had been saved into the electronic systems. To use this data, one would have to know exactly what to look for and where.

Consequently, the manager led the development of an image recognition AI model that was taught to interpret handwritten text. This integration of AI into the data-management process armed the department with a novel search capability that made data easily retrievable and usable by typing keywords into a search engine.

However, technical integration is not the most challenging part – it is the integration into work processes as organizational structures in the public sector emphasize stability and predictability over adaptability and risk-taking. Overall, AI entails a shift from deterministic thinking (for example, blind heed to thresholds) toward a more holistic orientation where AI models inform and augment (but not override) domain experts.

To achieve successful augmentation, an ability to operationalize new learnings is required both at the individual level (domain experts, managers) and the organizational level (processes, priorities, hierarchies) – work roles may have to change, and people must be willing to learn new skills and unlearn old assumptions. This is not easy, as employees used to existing job descriptions may struggle to expand their skillset and may be reluctant to adopt new ways of working. Failure to successfully augment, however, is a stumbling block for many AI implementations, including Tax AI: although the system was successfully introduced and its effectiveness proven, the department struggled to integrate it into daily work. The project manager reflected: *"That was a real lesson for us, that just by putting it in there and then thinking somebody would be able to use it and get value from it is probably unrealistic".*

In Health AI, the development team spent significant effort in consulting with triage nurses on what they would require from the system to be able to effectively use it. During interactions with data scientists and system developers, the nurses were educated about the system's classifications and how they should be considered in daily work. The challenge of integration like this can be seen as an opportunity to adopt a data-driven mindset and by doing so rise to a new level of capability and performance.

Considering these challenges with social integration, organizations need to invest significantly in **education and change management** for domain experts and managers alike.

Capability Area 5: Model Assurance

DEFINITION
An organization's ability to continuously supervise, assess, and improve the performance of an ML model, and control its effects on stakeholders

KEY ASPECTS
Managing AI systems does not end in development and productization – they are not a 'set and forget' technology and require ongoing monitoring to function effectively. For instance, 'model drift' becomes a concern when there are changes in the input data that the ML model cannot handle (for example, some taxpayers' payment behaviour changes due to unforeseen disruptions in a particular industry). Furthermore, changes in regulations and stakeholders' needs may lead to incumbent ML models' performance deteriorating, calling for intervention. Finally, it is possible that bias can creep into the models over time due to imbalanced distribution of input data (for example, if certain ethnic group grows its share among the customers of a public service, but the model has been trained with data where this group's presence has been lower). Such changes could create discrimination as an unintended consequence, compromising their ethicality of AI models. This means that organizations need capabilities and processes to ensure models' outputs are continuously reviewed in response to changes. Due to AI systems' online learning features, a specific focus on monitoring outputs is important. We thus distinguish three distinct aspects of model assurance capability:

SUPERVISING MODEL PERFORMANCE
(governance mechanisms, cadence of model review and monitoring, model drift management)

ASSESSING AND MEASURING MODEL VALUE FOR STAKEHOLDERS
(metrics and methodology for value assessment, tracking and communicating outcomes over time)

PREVENTING UNINTENDED CONSEQUENCES
(early warning, impact assessment, mechanism for instigating changes to the model, managing model learning and user engagement)

Role requirement: this capability is typically driven by algorithm supervisors and supported by data scientists, domain experts, and business managers.

- Model supervision
- Measuring model value
- Unintended consequences

HOW TO BUILD THE CAPABILITY
The models may need to be retrained with fresh training data to ensure they are able to serve their purpose adequately even when conditions change. With Health AI, the plan is to largely hand over the model management to the organization's IT service delivery team who manages the ongoing operation of all the department's digital systems. To ensure that the service team has appropriate understanding of how the model works, they have been looped into the development cycle of Health AI with regular meetings. Still, collaboration between the data science team and service delivery team is expected to continue and prompts for their interactions have been built into the system (automatic 'model health checks').

As the data-analytics director of Health AI explained: *"[A] pitfall for projects like this is that you'll find a specialist data science team who goes out and builds something, but then they get called to work on the next high priority strategic project, and there's no one left to maintain the product".*

While the online learning feature of AI was not utilized in our three projects, there was still a need to build assurance mechanisms into the model to maintain their accuracy and relevance over time. The data-analytics director elaborates: *"There are some components of support that that team won't be familiar with [...] they're not going to be data scientists themselves. [...] if you start getting bias in certain cohorts over time, or the predictive accuracy of the model begins to fail over time, we've got pre-defined alerts within these reporting solutions that would effectively tell that support team to tap us on the shoulder..."*

In addition to investing resources to ongoing management and coordination between teams, model assurance capability requires the organization to consider and prepare for a possibility of unintended consequences for stakeholders affected by the model. These efforts can include proactive analyses of alternative scenarios done in workshops attended by various stakeholder groups.

Involving different stakeholder groups into both **model development and ongoing management** will help to mitigate potential undesired issues.

Assembling Expertise to Build AI Capability

Building AI capability within public sector organization hinges on assembling teams that can carefully integrate various expertise needed for successful AI implementation. The capability areas we introduce in this report point to five roles (summarized in Table 2) that can contribute expertise to build and use AI safely. Some roles, such as data scientists, have become established in other industries, but less so in public agencies. Others entail entirely novel types of roles such as algorithm explainers and supervisors that play a key role in meeting ethical and legal considerations of the models.

Algorithm explainers will convey meaningful information regarding the logic, mechanics, and outputs of the model to various stakeholders (for example, customers, executives, and domain experts). This role has become important as the European General Data Protection Regulation (GDPR) has provisioned 'the right for explanation' for all European citizens[6]. When the models are operationalized, it is the role of the algorithm supervisor to constantly review model decisions, assessing the technical performance and the effects on stakeholders. This is key to minimizing unintended consequences.

Table 2: Public Sector AI Capability Areas

AI capability areas	Description	Who will drive it
1: Model development	An organization's ability to train a set of accurate and robust ML models to support various organizational tasks	Data scientists
2: Domain understanding	An organization's ability to develop new or enhanced knowledge of business domains	Domain experts and end users
3: Model explanation	An organization's ability to communicate ML model's goals, logic, and outputs to various stakeholders	Algorithm explainers
4: Model integration	An organization's ability to integrate ML models into existing technical interfaces, work processes, products, platforms, and structures	UX designers and developers
5: Model assurance	An organization's ability to continuously supervise, assess, and improve the performance of an ML model, and control its effects on stakeholders	Algorithm supervisors

6. Someh, I., Wixom, B.H., Beath, C., & Zutavern, A., (2022) 'Building an Artificial Intelligence Explanation Capability', Forthcoming in MIS Quarterly Executive.

Looking Ahead

Public organizations face unique challenges in adopting AI technology. The five capabilities we have discussed, point toward recommendations for nurturing successful AI-driven change in the public sector:

- **Jump-start AI development with pre-trained solutions:** budget-constrained agencies can jump-start their AI development by investing in externally developed solutions for carefully selected, narrowly specified use cases. Consider investing in pre-trained solutions and pre-built models to avoid high initial costs of developing AI from scratch.
- **Grow your internal expertise as you mature:** avoid running AI systems that you do not understand. While it is relatively easy to source technical expertise from AI vendors, it is your responsibility to ensure that your people fully understand the system and its limitations, and can be accountable for its decisions (for example, identify, understand, and control for risks that come with AI systems). Take ownership of these models by growing your internal data-science expertise over time. Develop complementary skills in ethics, law, and AI governance to safeguard the use of models.

- **Scrutinize legacy systems and data:** breaking free from silos created by old systems. Invest resources for overhauling your IT architecture. Compared to traditional analytics systems that focused on reporting or simple statistics, systematically building AI in public sector needs an array of contemporary platforms and systems that enable big data capture in the first place that is of high quality necessary for model training and permits sharing and collaboration. This is essential for establishing an agile data-sharing and AI development environment.
- **Leverage academic talent and join forces with other public agencies:** universities have the skills and the curiosity for projects using cutting edge technologies and they can also gather data that you may not be able to. So, collaborate with them and make them a part of your team. Partnering with your peer organizations is another avenue for extending your capabilities. Sharing data and models with them will help you to speed up and scale up AI development.
- **Focus on learning and continuous evolution:** start a cultural change. Achieving a successful technical installation is only one side of the coin – the organization's social system needs to transform at the same time, so lead that change from the very onset.

We advise public managers to ensure there is **organizational commitment to, as well as investment in,** AI to develop systems that deliver benefits while minimising risks.

About the SAP Institute for Digital Government

We live in an increasingly disrupted world and are witnessing an unprecedented transformation of how governments, businesses, and citizens operate and interact.

This transformation is readily evident in the changing role of government as it addresses this disruption: increasing expectations of citizens in how they engage with government services; the ability of government operations to effectively and safely utilise the valuable data within and across the ministries; and creating secure and economically sustainable environments and delivering the mission of government in helping drive nation-building.

SAP has been a key enabler of government services and processes for over 30 years. As a global company, we have first-hand experience partnering with governments. In 2014, along with several academic and government institutions, SAP created the SAP Institute for Digital Government to support governments in responding to these challenges. The institute facilitates a forum for exchange of ideas and thought-leadership demonstrating the public value of digital government to tackle real-world, complex issues.

To learn more about the SAP Institute For Digital Government, please visit https://discover.sap.com/sap-institute-digital-gov/en-us/index.html.

Follow us

www.sap.com/contactsap

THE BEST RUN SAP