WORLD ECONOMIC FORUM

*Unlocking Public Sector AI*

# **AI Procurement in a Box:** AI Government Procurement Guidelines

**TOOLKIT**

JUNE 2020

**Cover credit:** Paweł Czerwiński

# Contents



# 1 What is artificial intelligence (AI)?

AI has been formally defined as “technologies [that] aim to reproduce or surpass abilities (in computational systems) that would require ‘intelligence’ if humans were to perform them. These include: learning and adaptation; sensory understanding and interaction; reasoning and planning; optimization of procedures and parameters; autonomy; and creativity.”[1]

New AI approaches developed in the past decade, particularly the use of deep-learning neural networks, have dramatically advanced the capability of AI to recognize complex patterns, optimize for specific outcomes and make automated decisions. Doing this requires massive amounts of relevant data, a strong algorithm, a narrow domain and a concrete goal, and can result in dramatic improvements in reliability, efficiency and productivity.

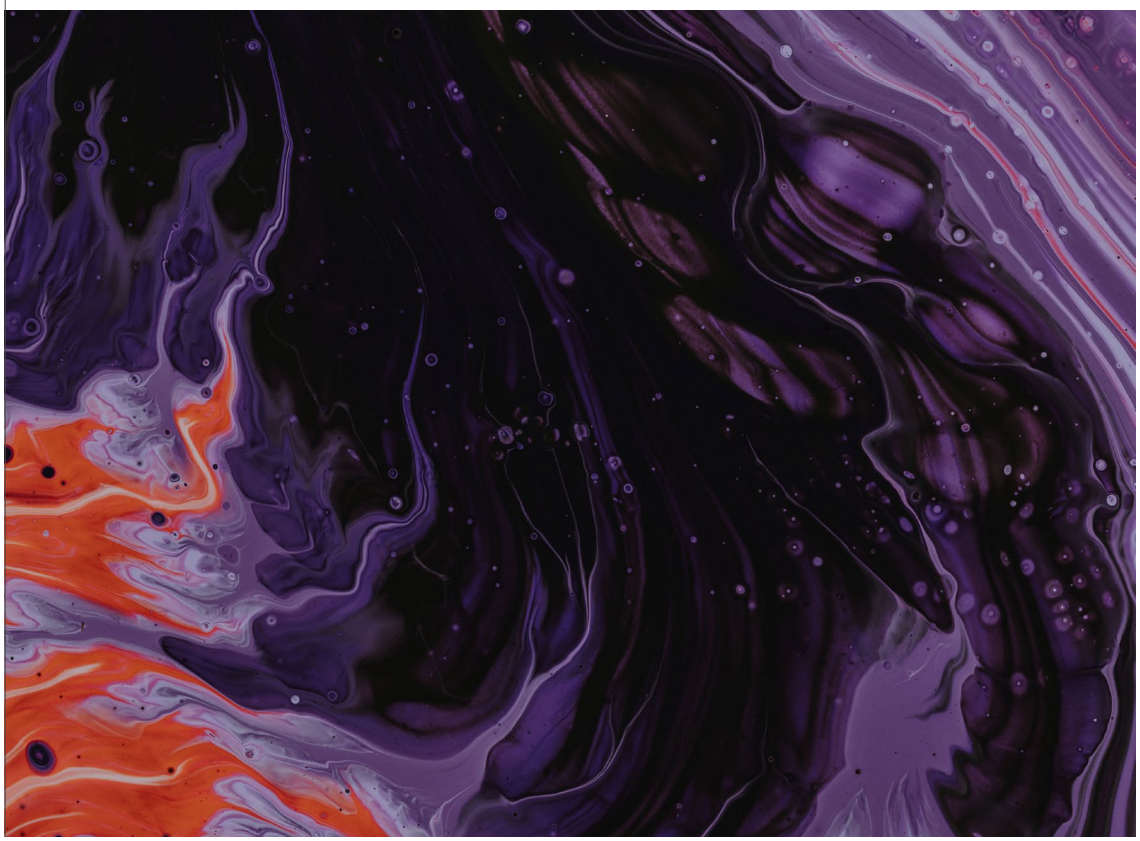

# 2 Why do we need guidelines for public procurement of AI?

Governments are increasingly seeking to capture the opportunities offered by AI to improve public-sector productivity and the provision of services to the public, and to stimulate the economy. AI holds the potential to vastly improve government operations and meet the needs of citizens in new ways, ranging from traffic management to healthcare delivery to processing tax forms. However, governments often lack experience in acquiring modern AI solutions and many public institutions are cautious about harnessing this powerful technology. Guidelines for public procurement can help in a number of ways.

> New examples are emerging of negative consequences arising from the use of AI in areas such as criminal sentencing, law enforcement and even employment opportunities.

First, government and the general public have justified concerns over bias, privacy, accountability, transparency and overall complexity. New examples are emerging of negative consequences arising from the use of AI in areas such as criminal sentencing, law enforcement and even employment opportunities. As citizens increasingly demand the same level of service from their governments as they do from innovative private-sector companies, public officials will be required not only to identify the specific benefits AI can bring, but also to understand the negative outcomes that can be generated.

Governments do not have the latitude of using the inscrutable "black box" algorithms that increasingly characterize AI deployed by industry. Without clear guidance on how to ensure accountability, transparency and explainability, governments may fail in their responsibility to meet public expectations of both expert and democratic oversight of algorithmic decision-making and may inadvertently create new risks or harms.

Governments rely on the expertise, and previously developed models, of technology providers and may lack the necessary skills to fully understand or trace algorithmic causality. Technology providers understand these challenges and look to governments to create clarity and predictability about how to manage them, starting in the procurement process. While companies are generally wary of stricter guidelines for government procurement, common-sense frameworks can help governments overcome reluctance to procure complex new technologies and actually open new markets for companies. Transparent guidelines will permit both established companies and new entrants to the AI space to compete on a level playing field for government contracts.

Second, AI procurement can build on a foundation of previous efforts to improve the effectiveness and efficiency of government technology procurement or be integrated into existing efforts. These may include legislation or policy measures such as frameworks or model contracts.

Established principles of good government technology procurement may take on added significance in AI procurement. For example, many governments already ensure that procurement efforts are run by multidisciplinary teams. Experience has shown that a lack of diversity in AI teams and positions of leadership has correlated with inadvertent harms or discrimination to vulnerable minority groups and protected classes. Given government's role in upholding inclusion, an added emphasis on a multidisciplinary approach and diversity may be necessary in AI procurement.

Some of the elements highlighted in the guidelines might already be evaluated in existing governance approaches but are not brought together holistically for decision-making. Closer working relationships between different teams should simplify the review of governance processes of AI systems even if they happen throughout different governance bodies and should integrate them in a strategy for AI adoption.

Third, as noted, AI has advanced rapidly in recent years, spurring further research and applications. New uses of AI that are of interest to governments will continue to emerge and will bring with them both benefits and risks. It is important that governments prepare for this future now by investing in building responsible practices for how they procure AI.

Finally, government procurement rules and purchasing practices often have a strong influence on markets, particularly in their early stages of development. As industry debates setting its own standards on these technologies, the government's moral authority and credibility can help set a baseline for these discussions.

Overall, the guidelines aim to guide all parties involved in the procurement life cycle – policy officials, procurement officers, data scientists, technology providers and their leaders – towards the overarching goal of safeguarding public benefit and well-being.

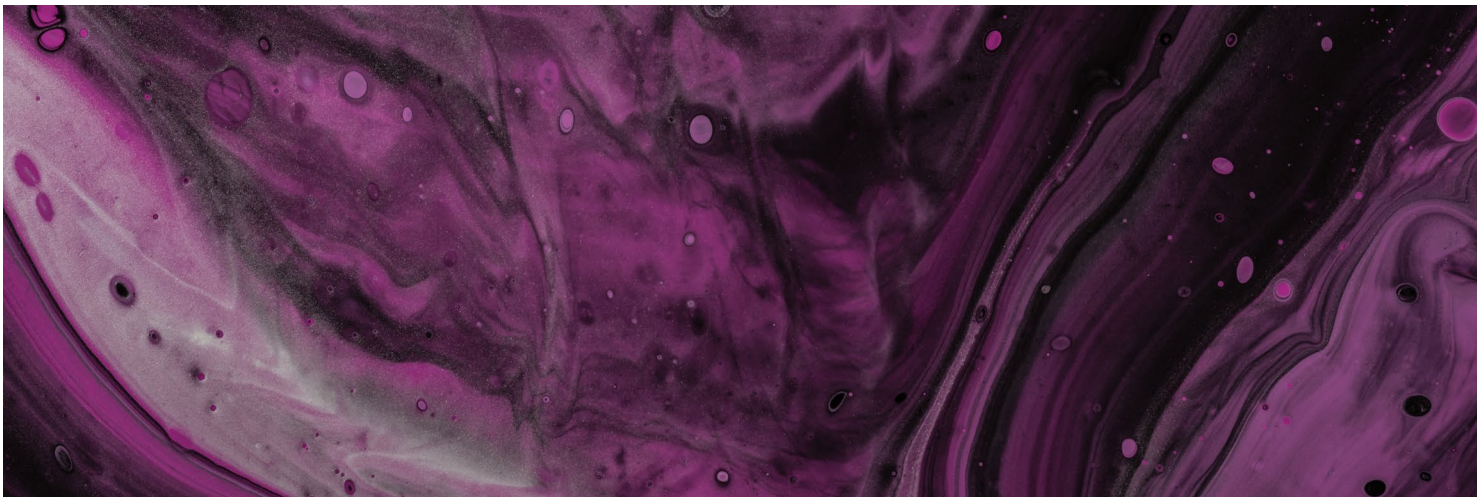

# 3 How were these guidelines developed?

The guidelines were developed by the World Economic Forum Centre for the Fourth Industrial Revolution, in consultation with a multistakeholder community. Project fellows from the UK Government's Office for AI, Deloitte and Salesforce worked with Forum staff, and in partnership with Splunk-convened workshops with appropriate representatives from government, academia, civil society and the private sector to explore key issues and co-design responses.

# 4 How to use the guidelines

The guidelines provide fundamental considerations that a government should address before acquiring and deploying AI solutions and services. They apply once it has been determined that the solution needed for a problem could be AI-driven. The guidelines are not intended as a silver bullet for solving all public sector AI-adoption challenges, but by influencing how new AI solutions are procured, they can set government use and adoption of AI on a better path.

**Specifically, the guidelines will help:**

- Policy officials to accelerate attainment of their policy goals
- Procurement officials and commercial teams to develop AI-related requests for proposals and to manage procurement processes
- Data practitioners (e.g. statisticians, data scientists, digital and technology experts) to safeguard public benefit and identify and manage potential risks
- AI-solutions providers to better understand the core expectations for government AI projects and to align their proposals with emerging standards for public procurement

The guidelines consist of 10 high-level recommendations, ordered roughly sequentially in terms of their relevance to the cumulative process of procurement, each containing:

- Multiple principles relating to each guideline
- Explanatory text elaborating on the thinking and substance underlying each principle

BOX 1 **Taking a risk-based approach to AI procurement**

> This increases the risks and sensitivities about AI deployment in many use cases.

It is important to approach AI procurement proportionality and not all guidelines may apply to the same extent to all procurement decisions. This is also why it is crucial to conduct an initial AI impact assessment and then act appropriately and proportional.

Important issues that can drive your decision whether to add additional ethical criteria to consider within your procurement approach, can fall within the following categories, many of which are closely interlinked. Note that this is not an exhaustive list of issues that need to be considered nor does it give you the answers whether your AI project might be more or less risky but it highlights key areas that need to be investigated further, particular in a public sector context.

## Key variables to consider in a risk assessment:

**Data:**

- **Data sensitivity** – The more sensitive the data that you are using within the AI system is, the more checks you should be building in. You need to closely consider if the data could be re-identified or give away any personal information.
- **Data quality** – The less sure you are about the quality of your data, the better it is to build in additional assurances to avoid bias and de-risk the project. Ensuring the representativeness of the data set might be difficult to ensure and qualitative measures might need to be taken. It is important to consider specific societal bias that could be reflected in the data for public sector use cases.
- **Data consent** – If meaningful personal data consent in the context that you are planning to use an AI-driven solution is not clear, the project is considered riskier. Also ensure that you are not inferring consent to a certain use of the data that does not comply with the original use case.

**Field of use:**

- **Public scrutiny** – If the project is within a sector of intense public scrutiny because of privacy concerns, legal concerns, interest and/or frequent litigation, the stakes are also higher. Fields, among others, such as health, social assistance, employment, financial services, insurance, the criminal justice systems, immigration, access and mobility, or decisions about permits and licences are examples of areas of applications that demand further considerations.

**Socioeconomic impact:**

- **Stakeholders involved** – The higher the impact on individuals, businesses, and communities, the more important it gets to thoroughly consider AI ethics and scrutinize the application of AI.
- **Scope of impact** – It is important to consider factors such as how many people are impacted; how high the impact is and how high the likelihood of impact is. The risk also increases when decisions of the systems are linked to groups of people that are particularly vulnerable.

**Financial consequences for agency and individuals:**

- **Scope of financial impact** – The higher the potential financial consequences, the more you should address all areas linked to AI specific considerations.
- **Types of financial impact** – The financial consequences can be diverse and include monetary aspects as well as the access to credit, economic opportunities, schooling or training, insurance and certifications.

**Impact of the AI system on your processes, employees and core businesses:**

- **Core functions impact** – If the AI system is central to the core function of the agency, you should take on a more mandated approach to not only mitigate technical risks but also for reputational risk. The more tech dependence you create the riskier.
- **Business functions impact** – Consider whether you are replacing a business function rather than just improving and adding to the status quo, this might also impact your decision on how much to scrutinize the procurement process.
- **Job loss** – the more processes are automated, the more job losses can be expected. This increases the risks and sensitivities about AI deployment in many use cases.
- **Human in the loop** – The less checks and balances you have in place, the more risk. You should focus on adding explainability, interpretability and mindful friction to your AI deployment.

## Example of tools that are already used within the public sector and the risks attached to this adoption:

- **Spam-filters in email programs** – designed to detect and block unwanted emails. Have the least risk prone use of AI in the public sector but can lead to discrimination if certain email addresses are blocked. However, “human in the loop” is usually included at various junctions so that the program isn’t

making decisions completely on its own, thus easily mitigating risk.

- **AI in cybersecurity solutions** – designed to protect networks, programs, and data from attack, damage, or unauthorized access. At first sight less prone to risks related to AI ethics, but we need to closely consider how the system is used in practice. If AI is used to better predict threats or identify cyber security risks, thus in a supporting function rather than making specific decisions, this use case seems to have a lower risk profile and thus would demand a less stringent approach to the implementation of all parts of the guidelines.
- **Chatbots** – designed to converse with people via voice interfaces or text messages. If they mainly provide information back to you and make it easier to sort through a large amount of data, rather than driving decisions, the use case seems to be less risk prone. But since they will likely be built into key processes and will have citizens interact with them, it is advised to follow the AI procurement guidelines to support those purchases.
- **Fraud detection** – designed to detect, prevent and manage fraudulent patterns in the data. Well tested use case of AI in the public sector, allows departments to make more effective enforcement decisions but the risk can be high if data quality is poor and if vulnerable groups are over proportionally targeted. False positive can also have high financial consequences and data sensitivity can be high depending on the use case. Hence, AI procurement guidelines should be followed.
- **AI in policing or social services** – designed to support and/or drive decisions in fields such as law enforcement, crime prevention, public safety, children welfare, social programs. The use of AI in those fields involves large risks as policy decisions are built into those systems and socioeconomic impacts are high. These use cases need to be put under particular scrutiny and procurement decisions need to follow very clear rules that include system testing, ethical considerations and a great focus on data governance. AI procurement guidelines should be closely followed.
- **AI in HR** – designed to take on key HR tasks including hiring, retaining talent, training, benefits and employee satisfaction. Employment decisions have high stakes with critical consequences for individuals, organizations and society. Algorithms can make predictions in ways that disadvantage certain groups. Hence, concerns about AI algorithms bias and discrimination are particularly heightened, further complicated by labour and anti-discrimination laws. Finally, unique aspects of the human resources setting, including small datasets, complex social interactions, data privacy concerns and the need for accountability pose challenges and require close procurement guidelines governance.

**Examples on how to do this:**

1. AI risk assessment tool: The tool aims to help you decide on a proportional approach to AI procurement. It sets out examples for decision criteria to include in a risk assessment of any potential solutions that contain AI capabilities. The tool outlines some of key questions you should consider when deciding your procurement strategy, considering what questions to ask in your RFP and assessing a solution.

2. Alan Turing Institute, Understanding artificial intelligence ethics and safety: This guide is an end-to-end guidance on how to apply principles of AI ethics and safety to the design and implementation of algorithmic systems in the public sector. The ethical platform includes; a list of values that orient you in deliberating about the ethical permissibility and impact of a prospective AI project; a set of principles that all members of your project delivery team should be well-acquainted with and a framework that operationalizes these values and principles in an end-to-end workflow governance model.

3. Canadian directive on automated decision-making: The Canadian government has developed a risk-based approach to AI adoption in the public sector which divides the AI systems in different levels. The four factors used to determine the risk-level are impact on: the rights of individuals or communities, the health or well-being of individuals or communities, the economic interests of individuals, entities, or communities and the ongoing sustainability of an ecosystem. Based on the risk-level, the guide provides insights on how to best approach AI procurement from a proportionality view and to what extent each requirement should be applied. Please refer to figure 1.

FIGURE 1 **Canadian Directive on Automated Decision-Making**

| Level | Description |
|---|---|
| 01 | The decision will likely have little to no impact on:<br>– The rights of individuals or communities.<br>– The health or well-being of individuals or communities.<br>– The economic interests of individuals, entities, or communities.<br>– The ongoing sustainability of an ecosystem.<br>Level 01 decisions will often lead to impacts that are reversible and brief. |
| 02 | The decision will likely have moderate impacts on:<br>– The rights of individuals or communities.<br>– The health or well-being of individuals or communities.<br>– The economic interests of individuals, entities, or communities.<br>– The ongoing sustainability of an ecosystem.<br>Level 02 decisions will often lead to impacts that are likely reversible and short-term. |
| 03 | The decision will likely have high impacts on:<br>– The rights of individuals or communities.<br>– The health or well-being of individuals or communities.<br>– The economic interests of individuals, entities, or communities.<br>– The ongoing sustainability of an ecosystem.<br>Level 03 decisions will often lead to impacts that can be difficult to reverse, and are ongoing. |
| 04 | The decision will likely have very high impacts on:<br>– The rights of individuals or communities.<br>– The health or well-being of individuals or communities.<br>– The economic interests of individuals, entities, or communities.<br>– The ongoing sustainability of an ecosystem.<br>Level 04 decisions will often lead to impacts that are irreversible, and are perpetual. |

As the technological sophistication and government use of AI evolves, the guidelines should be updated to reflect new learning and leading practices. This is a living document that is intended to integrate feedback from practitioners over time. Much of that feedback will come from two sources: the project's community of subject matter experts, and the pilots to be held with the UK, the United Arab Emirates, Colombia and other partner governments. We also welcome feedback from other stakeholders and the general public. If you wish to provide feedback, please share via email: AI@weforum.org.

Ultimately, the goal is that these guidelines will enable governments and international bodies to set the right policies, protocols and perhaps even standards to facilitate effective, responsible and ethical public use of AI.

# 5 Guidelines overview

What are the key considerations when starting a procurement process, writing a request for proposal (RFP), and evaluating RFP responses?

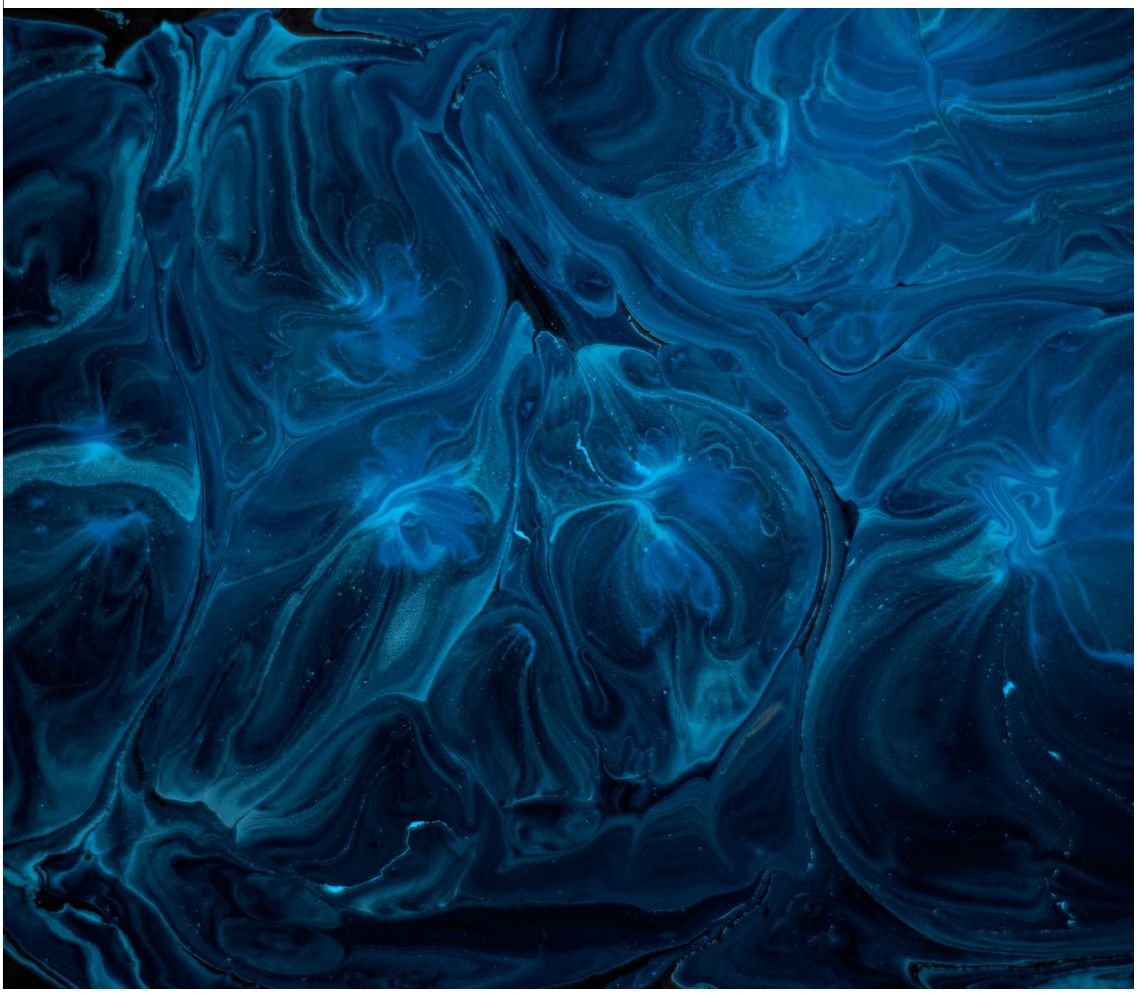

| | Guideline | Principles |
|---|---|---|
| 01 | Use procurement processes that focus not on prescribing a specific solution but rather on outlining problems and opportunities, and allow room for iteration. | **a.** Make use of innovative procurement processes to acquire AI systems.<br>**b.** Focus on developing a clear problem statement, rather than on detailing specifications of a solution.<br>**c.** Support an iterative approach to product development. |
| 02 | Define the public benefit of using AI while assessing risks. | **a.** Set out clearly in your RFP why you consider AI to be relevant to the problem and be open to alternative technical solutions.<br>**b.** Explain in your RFP that public benefit is a main driver of your decision- making process when assessing proposals.<br>**c.** Conduct an initial AI risk and impact assessment before starting the procurement process, ensure that your interim findings inform the RFP, and revisit the assessment at decision points. |
| 03 | Align your procurement with relevant existing governmental strategies and contribute to their further improvement. | **a.** Consult relevant governmental initiatives such as AI national strategies, innovation and/or industrial strategies, and guidance documents informing public policy about emerging technologies.<br>**b.** Collaborate with other relevant government bodies and institutions to share insights and learn from each other. |
| 04 | Incorporate potentially relevant legislation and codes of practice in your RFP. | **a.** Conduct a review of relevant legislation, rights, administrative rules and other relevant norms that govern the types of data and kinds of applications in scope for the project and reference them in the RFP.<br>**b.** Take into consideration the appropriate confidentiality, trade-secret protection, and data-privacy best practices that may be relevant to the deployment of the AI systems. |
| 05 | Articulate the technical and administrative feasibility of accessing relevant data. | **a.** Ensure that you have proper data governance mechanisms in place from the start of the procurement process.<br>**b.** Assess whether relevant data will be available for the project.<br>**c.** Define if and how you will share data with the vendor(s) for the procurement initiative and the subsequent project.<br>**d.** Ensure that you have the required access to data used and produced by the vendor(s) solution. |

| | Guideline | Principles |
|---|---|---|
| 06 | Highlight the technical and ethical limitations of intended uses of data to avoid issues such as historical data bias. | **a.** Consider the susceptibility of data that could be in scope and if usage of the data is fair.<br><br>**b.** Highlight known limitations (e.g. quality) of the data in the RFP and require tenderers to describe their strategies on how to address these shortcomings. Have a plan for addressing relevant limitations that you may have missed. |
| 07 | Work with a diverse, multidisciplinary team. | **a.** Develop ideas and make decisions throughout the procurement process in a multidisciplinary team.<br><br>**b.** Require the successful bidder(s) to assemble a team with the right skill set. |
| 08 | Focus throughout the procurement process on mechanisms of algorithmic accountability and of transparency norms. | **a.** Promote a culture of accountability across AI-powered solutions.<br><br>**b.** Ensure that AI decision-making is as transparent as possible.<br><br>**c.** Explore mechanisms to enable interpretability of the algorithms internally and externally as a means of establishing accountability and contestability. |
| 09 | Implement a process for the continued engagement of the AI provider with the acquiring entity for knowledge transfer and long-term risk assessment. | **a.** Consider during the procurement process that acquiring a tool that includes AI is not a one-time decision; testing the application over its lifespan is crucial.<br><br>**b.** Ask the AI provider to ensure that knowledge transfer and training are part of the engagement.<br><br>**c.** Ask the AI provider for insights on how to manage the appropriate use of the application by non-specialists. |
| 10 | Create the conditions for a level and fair playing field among AI solution providers. | **a.** Reach out in various ways to a wide variety of AI solution providers.<br><br>**b.** Engage vendors early and frequently throughout the process.<br><br>**c.** Ensure interoperability of AI solutions and require open licensing terms to avoid vendor lock-in. |

# 6 Detailed explanation of guidelines

## 6.1 Use procurement processes that focus not on prescribing a specific solution, but rather on outlining problems and opportunities and allow room for iteration.

**Why is it important?**

To acquire the AI systems that best address the challenge you want to address and encourage responsible innovation.

**a. Make use of innovative procurement processes to acquire AI systems.**

- Innovation-oriented procurement procedures provide opportunities to accelerate the adoption of new technologies such as AI systems, to promote innovation and to support secondary policy criteria such as support for small and medium-sized enterprises and the ethical development of AI.
- For example, these processes support early market engagement, enable you to go to market in different stages and can include the use of proofs of concept. These provide the opportunity to test the technologies on your problem area before making a final buying decision. Innovative public procurement processes that include practices such as detailing challenging problems, organizing technology contests, providing opportunities for demonstrators, and giving newly established providers the opportunity to compete for public-sector contracts, have the potential to boost innovation and help new companies become established. This market-making role also encourages small enterprises with new ideas and reduces the risks for new technology start-ups.
- By strategically choosing the procurement approach depending on the nature of the challenge that you mean to address, these processes could include, for example:
  - Agile procurement processes that allow you to go to market in different stages and can include proofs of concept to test the technologies before the final purchase.
  - Challenge-based procurement processes that have vendors compete against each other based on their AI skills and include an evaluation of the technologies applied to the challenges they mean to address.
- Innovation partnerships that enable the procurement of technologies that cannot be delivered by the current options available to the market.
- Dynamic purchasing systems – procedures currently used mainly for products commonly available on the market – can accelerate uptake of technologies that are rapidly developing. As

a procurement tool, it is similar in some ways to an electronic framework agreement but, as new suppliers can join at any time, this allows newly established firms to participate in the framework agreements when they meet the set criteria.

- AI procurement frameworks that prescribe the terms and conditions applying to any subsequent contract and allow the pre-vetting of providers against a set of predefined criteria that can include ethical requirements.

- When making use of novel approaches to procuring emerging technologies you should also focus on best practices that have been shown to increase the supplier base of smaller and innovative suppliers, which is important for fast-developing markets such as AI. These practices include, but are not limited to:

  - Setting out and following a detailed procurement timeline at the start of the campaign.

  - Breaking down large proposals into smaller work components.

  - Encouraging collaboration between different bidders.

**b. Focus on developing a clear problem statement, rather than on detailing the specifications of a solution.**

- AI technologies are developing rapidly, with new technologies and products constantly being introduced to the market. By focusing on describing the challenges and/ or opportunities that you want to address and drawing on the expertise of technology partners, you can better decipher what technology is most appropriate for the issue at hand. By focusing on the challenge and/or opportunity, you might also discover a higher-priority issue, or realize you were focusing on a symptom rather than the root cause.

- Beyond playing to each stakeholder's strength, this approach has two added benefits. First, it demands and promotes early market engagement, which we explain in further detail in Guideline 10. Second, it makes it easier for newer AI service providers (such as start-ups) to participate, as the government will not be focused on a specific product. Nurturing an emerging AI ecosystem is a key economic investment in the future.

**c. Support an iterative approach to product development.**

- AI-powered solutions differ significantly from other technology tools in their unique ability to learn and adapt through ongoing, periodic training with new data. Therefore, the procurement process should allow room for iteration, while ensuring a robust, fair and transparent evaluation and decision process.

- For example, a phased challenge-based procurement could serve to evaluate different competitors' minimum viable products (MVPs) during phase one of procurement, with only the winner going on to develop the full solution. This building and testing in phases within the procurement cycle facilitates informed decision- making, innovation and transparency. It also provides you with relevant information to conduct meaningful impact assessments and evaluate risks.

Encouraging collaboration between different bidders.

FIGURE 2 **Visual to depict the challenge-based procurement process used by the UK GovTech Catalyst challenge**

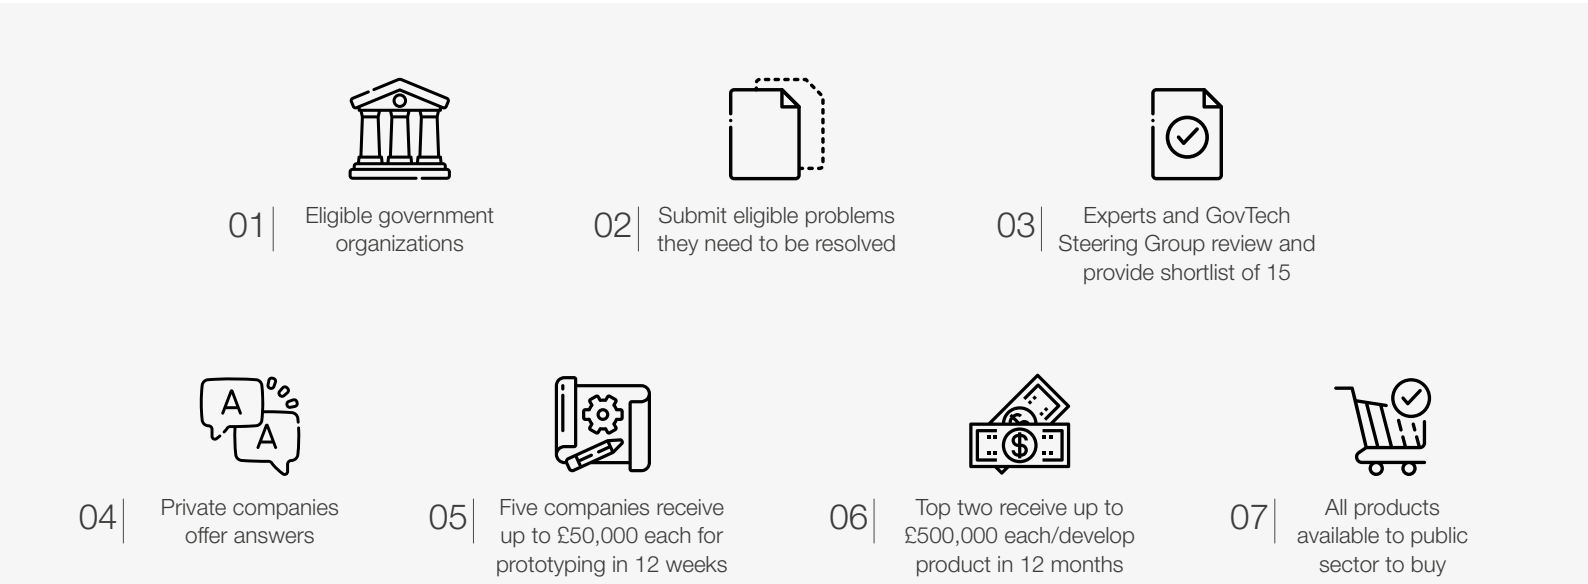

## 6.2 Define the public benefit of using AI while assessing risks.

> What do you expect such a system to achieve and be capable of, and what are the types of failure and harm that must be avoided?

**Why is it important?**

Defining the public benefit goal provides an anchor for the overall project and procurement process that the AI is intended to achieve. AI also brings new and specific risks that must be identified and managed as early as the procurement phase of the project.

**a. Set out clearly in your RFP why you consider AI to be relevant to the problem and be open to alternative technical solutions.**

- In most circumstances, you should refer to the need for an AI solution in your invitation to tender only if there is strong indication that the technology will address the problem that you are trying to solve. A need for the acquisition of an AI system should arise through analysis of policy challenges and alternatives, and be compared to other potential courses of action when the AI project does not have a clear research and innovation focus. If, during the evaluation of the tender responses, it becomes evident that another solution that doesn't incorporate AI is better able to address the problem, you should make the decision to follow this alternative delivery path.

- Assess whether AI could be part of a solution to your problem, before starting the procurement process. If you lack the capabilities in your team to carry out this assessment, you should seek these from elsewhere in your organization or relevant professional network (e.g. academia, trusted vendors) and make the consultation and collaboration with appropriate stakeholders a priority. For this assessment, it is crucial to engage a multistakeholder community to define and test a clear policy problem statement and reflect the findings in the RFP.

- Pre-market engagement is also often essential in helping you to describe your problem and narrow down the tasks that AI may be able to assist with. This will help you better communicate to potential suppliers what you are asking for and why, as well as highlighting where the gaps are. Documenting user need and challenges to the best of your ability is crucial, since the success of the project also depends on how well AI system providers understand the problem.

**b. Explain in your RFP that public benefit is a main driver of your decision-making process when assessing proposals.**

- When setting out the requirements in the RFP, you should consider explicitly referring to public benefit as well as user needs. When determining user needs, public servants should be confident that they are acting in the public benefit. With regard to AI systems, the public benefit extends beyond value for money and also includes considerations about transparency of the decision-making process and other factors that are included in these guidelines.

- In practice this requires you, for example, to specify success and failure criteria in the context of public benefit: What do you expect such a system to achieve and be capable of, and what are the types of failure and harm that must be avoided? Conducting an impact assessment will help you to set these issues out. Refer to Guideline 7 for additional information on adding ethical requirements to the RFP.

**c. Conduct an initial AI risk and impact assessment even before starting the procurement process, ensure that your interim findings inform the RFP, and revisit the assessment at decision points.**

- To better understand the potential impacts of the use of AI and to mitigate the risks, it is important to start an assessment in a systematic manner before the acquisition of an AI system and to make sure that the findings also inform your commercial strategy. There will be different considerations depending on which policy challenges you are trying to solve and which potential application of AI could help to address this challenge. Without knowing which AI system you will acquire, it is not possible to conduct a whole assessment.

- An initial assessment should outline user needs and affected communities, as well as potential risks such as inaccuracy and bias of the AI system. It should also include some consideration of scenarios involving unintended consequences. The initial assessment should make you think about strategies to address these potential impacts, including but not limited to citizen panels, transparency reports and testing on differentially private or synthetic datasets. Associated risks and their respective mitigation strategies must be recognized by a suitable risk owner with decision-making power and should include a go/no-go decision.

- In your invitation to tender, you should consider asking potential suppliers to identify risks and explain how they would mitigate them. This can give you valuable information regarding how careful each tenderer is and how aware they are of potential risks. Where you identified significant risks in your initial assessments, you should specifically require tenderers to set out how they would address those.

– Data protection impact assessments and equality impact assessments can provide a useful starting point for assessing potential unintended consequences. In assessing these, you should consider how the use of these systems, such as semi-automated or solely automated decisions, interact with mechanisms of oversight, review and other safeguards. We developed a high-level risk assessment, which allows you to make a more informed decision about your approach, and introduced the concept of a proportional approach to AI procurement. See the AI risk assessment tool in the workbook. For other examples of risk assessment questionnaires for automated decision-making, refer to the government of Canada's Directive on Automated Decision Making, and the framework on Algorithmic Impact Assessments from AI Now.

– In addition to the above, there should be systematic and continuous risk monitoring during every stage of the AI solution's life cycle, from design to post-implementation maintenance. AI solution providers can do this by identifying, drafting mitigations for and reporting risks through a project management function, which is central to the implementation (refer to Guideline 9 for more information on how to consider life-cycle management during the procurement process). The impact assessment should be revisited where necessary (e.g. in the event of significant changes to the opportunity statement).

BOX 2 **Example of human rights assessment from Google Cloud**

Google Cloud launched a Celebrity Recognition tool to a select set of media and entertainment customers to help them identify and label celebrities in professionally produced content, such as movies and sporting events. From the beginning of the product development process, they engaged in a human rights impact assessment (HRIA) and internal AI principles reviews. In partnership with BSR, a human rights non-profit organization, and using the UN's Guiding Principles on Business and Human Rights as a framework, the team considered potential impacts throughout numerous dimensions including privacy, discrimination, freedom of expression and many others. Aspects such as consultation with potentially affected stakeholders, dialogue with independent expert resources and paying particular attention to those at heightened risk of vulnerability or marginalization were part of the methodology. Their input played an essential role in shaping the API's capabilities and the policies established around them.[2]

**Some mitigation strategies adopted after this initial human rights risk assessment:**

– Creation of "Service Specific Terms" that customers need to agree with. These limit the range of content upon which the API can be used and that address issues such as copyright, hate speech, child rights, surveillance and censorship.
– Adoption of a narrow definition of celebrity that respects the principle of informed consent by only including those that have actively and deliberately sought a role in public life.
– Creation of an "opt-out" option for celebrities not wanting to be included in Google's celebrity database.

FIGURE 3 **Visual of the SDLC stages, with sample AI risk assessment question for each stage.**

| SDLC stage | Sample AI risk mitigation considerations |
|---|---|
| **01** Requirements gathering and analysis | – Is the use of AI/ML necessary for the desired outcome?<br>– Should AI/ML even be discussed at this stage? |
| **02** Design | – Do we have consent to use the data sources required by the solution?<br>– Do we fully understand the implications of using external data, models or solutions? |
| **03** Implementation and coding | – Do we have the right skills or domain expertise to develop the solution?<br>– Does the development process protect data confidentiality and integrity? |
| **04** Testing | – What level and type of bias is acceptable in the solution?<br>– Do the acceptance criteria set appropriate levels of accuracy to ensure the model performance is satisfactory? |
| **05** Deployment | – Have users received adequate training to ensure they understand the output of the system?<br>– Is it transparent to users how the solution is deriving an output? |
| **06** Maintenance | – Do the system administrators know what metrics to examine to validate that models are operating as expected?<br>– Is there a clear process for updating or refining models using new data? |

## 6.3 Aim to include your procurement within a strategy for AI adoption across government and learn from others.

**Why is it important?**

To ensure that you use procurement strategically to support efforts on AI development and deployment, and to spread the knowledge of the public application of an emerging technology.

**a. Consult relevant AI national strategy initiatives and guidance documents from ministries and departments informing public policy of emerging technologies.**

– Many countries are currently in the process of drafting and releasing national AI strategies, and some have already published theirs. Prior to commencing an AI rollout, evaluate how your pursuit of acquiring an AI system aligns to your country's overall strategy.

– This allows you to include secondary policy aims in your strategic procurement and potentially make use of economies of scale by pooling the demand for AI systems. An added benefit of aligning to a national AI strategy is that there may be special support for initiatives that align to the strategy, such as access to additional experts.

**b. Consult with government agencies that have looked into procuring AI solutions, irrespective of the outcome of the procurement efforts.**

– To improve your practices and share your experiences, you could actively seek out collaboration across departments and fields of expertise. You could also share knowledge and feedback via expert communities, such as the digital-buying community, professional networks or meet-ups.

– Within your department it can be helpful to set up platforms and networks that allow for the exchange of information, experiences and best practices about the purchasing of AI-powered solutions.

> Many countries are currently in the process of drafting and releasing national AI strategies, and some have already published theirs.

## 6.4 Ensure that legislation and codes of practice are incorporated in the RFP.

**Why is this important?**

Conforming with existing laws and regulations ensures compliance; incorporating codes of practices supports the standardization of norms; and surveying the relevant rules enables better policy-making in a dynamic innovation technology ecosystem.

**a. Conduct a review of relevant legislation, rights, administrative rules and other relevant norms that govern the types of data and kinds of applications in scope for the project.**

– Conduct a review of relevant legislation, human and civil rights, administrative rules, and other relevant norms that govern the types of data and kinds of applications connected to the problem being addressed and solutions being proposed. Clarify the appropriate adjudicative and administrative jurisdictions within the domestic government in relation to conflicts of laws concerning the data. Depending on the problem being addressed in the invitation to tender, existing laws and regulations relevant to that government function may already have some rules on the use, processing, transfer etc. of data. Incorporate those rules and norms into the RFP by referring to the originating laws and regulations.

– When identifying the relevant rules, sources should include not only formal law, but also industry best practices, trade organization consensus guidelines and other forms of norm-setting mechanisms of soft law. For example, freedom of information laws[3] establish rules about what needs to be made available to the public, and data ethics frameworks guide the design of appropriate data use in government and the wider public sector.

**b. Take into consideration the appropriate confidentiality, trade secret protection and data privacy best practices that may be relevant to the deployment of the AI solutions.**

– To meaningfully evaluate proposed AI solutions, government officials must strike the right balance between preserving accountability through transparency and reassuring vendors that the trade secrets associated with their products and services, as well as their business confidentiality, will not be compromised. Information about government processes should be open by default, with the limits of disclosure justified in exceptional circumstances such as export controls, national security or ongoing criminal investigations.

– In those circumstances where confidentiality and trade- secrecy protection can be justified in light of public-interest considerations, investigate the possibilities of facilitating transparency through partial disclosure, limited review options and other means of enhancing public trust.

## 6.5 Articulate the technical feasibility and governance considerations of obtaining relevant data.

**Why is this important?**

Availability of relevant data is a prerequisite for any AI solution, so time should not be spent discussing AI procurement if no data will be available. In addition, access to data should be granted only after careful consideration by the data-governing party(ies).

**a. Ensure that you have proper data-governance mechanisms in place from the start of the procurement process.**

– Set out a data-governance approach from the start of the procurement process. Given the importance and complexity of data governance, it is almost mandatory to implement sound data-governance processes before engaging in transformative AI projects. Governance needs to cover all data activities related to the proposed project, such as granting data access to project members, moving/storing data in other locations for analysis, and reviewing data consent (the purposes for which we are authorized to use the data).

– Data governance, and all other aspects of an AI initiative, require a multidisciplinary team. Refer to Guideline 7 for more information on multidisciplinary teams.

– In the absence of a data-governance framework, ensure that it is clear who is accountable (who is responsible for data management during the procurement process and the subsequent project).

**b. Assess whether relevant data will be available for the project.**

> Data is crucial for modern-day AI tools.

– Data is crucial for modern-day AI tools. You should determine, at a high level, data availability before starting your procurement process. This entails developing an understanding of what data might be required for the project. The idea is not to assess all possible data sources, but to build general awareness of data sources of potential interest. Data documentation, using data dictionaries, for example, is helpful when trying to build a high-level understanding of data assets.

– In cases where data is not available for the use case in mind, you may be able to find data through third parties, for example, vendors, partners or data brokers. If data is not available through any channel, engage skilled data scientists (for example, through vendors) to assess whether the use case can be addressed at all in a data- driven manner.

**c. Define if and how you will share data with the vendor(s) for the procurement initiative and the subsequent project.**

– Depending on the sensitivity of your project and data, it is worth considering the release of data to providers during procurement so that bidders can craft a response to the RFP that is tailored to your needs, with assumptions, timelines and fees that match your situation as closely as possible. This improves the quality of RFP responses you receive.

– If you are releasing data that is sensitive and not meant for public consumption, consider releasing only a sample, so that vendors have a clear idea of what the data enables them to do without having access to all of it. When doing this, make sure that you provide a sample that is representative of the overall dataset. Otherwise, vendors might make erroneous assumptions that can impact the quality of bids and consequently the integrity of the project.

– Create and document the appropriate data-sharing conditions. For example:

– Minimum requirements for the environment where the vendor will host the data (e.g. enterprise laptop that meets the vendor's standards for their sensitive data).

  – Data consent form signed by the vendor's lead for the project, stating that the data will be used exclusively for the pursuit and for no other purpose. It should be clear to vendors that while in possession of the data they are not allowed to use the data for any purpose other than that specified in the RFP.

  – Date for data deletion (e.g. immediately upon submission of the vendor's RFP response). In

no circumstances should governments allow vendors to keep data after the procurement process, or after the conclusion of the project for successful bidders.

  - Confirmation of deletion of all data (e.g. written confirmation of deletion signed and submitted by the vendor's lead for the project).

- There are many anonymization techniques to help safeguard data privacy, including data aggregation, masking and synthetic data.[4] Keep in mind, however, that you must manage anonymized data as carefully as the original data, since it may inadvertently expose important insights. RFPs should encourage innovative technological approaches, such as those mentioned above, that make less intrusive use of data or that achieve the same or similar outcomes with less sensitive datasets.

- Certain vendors may have data that is complementary to the initiative, and it is in your best interest to consider using this data. It is important to have a framework that gives guidance regarding the circumstances under which it is reasonable to accept data from a vendor. Decision criteria could include:

  - Vendor: some vendors could be pre-qualified as accepted data providers, be considered more trustworthy as a result of their previous track record as existing suppliers or have a strong reputation related to their data assets.

  - Domain: some domains – such as health, justice and immigration – are very sensitive. Use of third-party data in these domains requires careful scrutiny before it is accepted.

  - Data precedence and integrity: before using any third-party data, the government should have a clear understanding of how the data was collected, the governance processes employed to ensure its integrity, and whether the third party offering the data is legally allowed to commercialize it for the RFP.

**d. Ensure that you have the required access to data used and produced by the vendor(s) solution.**

1. Access and control of data used and produced by AI models is critical in monitoring, assessing and rectifying performance.

2. You must ensure that you have access to raw input, processed/combined and enriched data produced by the vendor(s) AI models. This should also include third party and open source data, particularly if there is the chance that these will not be available/maintained on a long-term basis.

3. Dependent on the solution(s) proposed the vendor(s) may not be willing or able to provide full access to all data (e.g. to protect IP for SaaS or COTS solution):

- Access to data should be provided with as wide a scope as possible. The supplier should be able to clearly articulate the reason for restricted sharing and this should be limited to only relevant areas not a blanket justification (e.g. commercially sensitive training sets do not preclude sharing enriched model outputs).

- You should ensure that, where restricted access is justified, the supplier provides relevant, up-to-date and representative sampled data sets. Ideally these with be constructed from operational/live data.

4. Data ownership should be clearly articulated by the supplier:

   4a. You should aim for contractual ownership of the data on a persistent basis.

   4b. As a minimum enriched data produced by the AI model(s) should be under "shared ownership" with access rights to all remaining data.

   4c. Ideally key data sets should be available for your internal teams to use learn and develop enhanced/new systems and approaches.

FIGURE 4 **Sample data governance framework**

Deloitte's data governance framework enables organizations to be specific in terms of what goals will be prioritized, what capabilities will be deployed and what results are expected

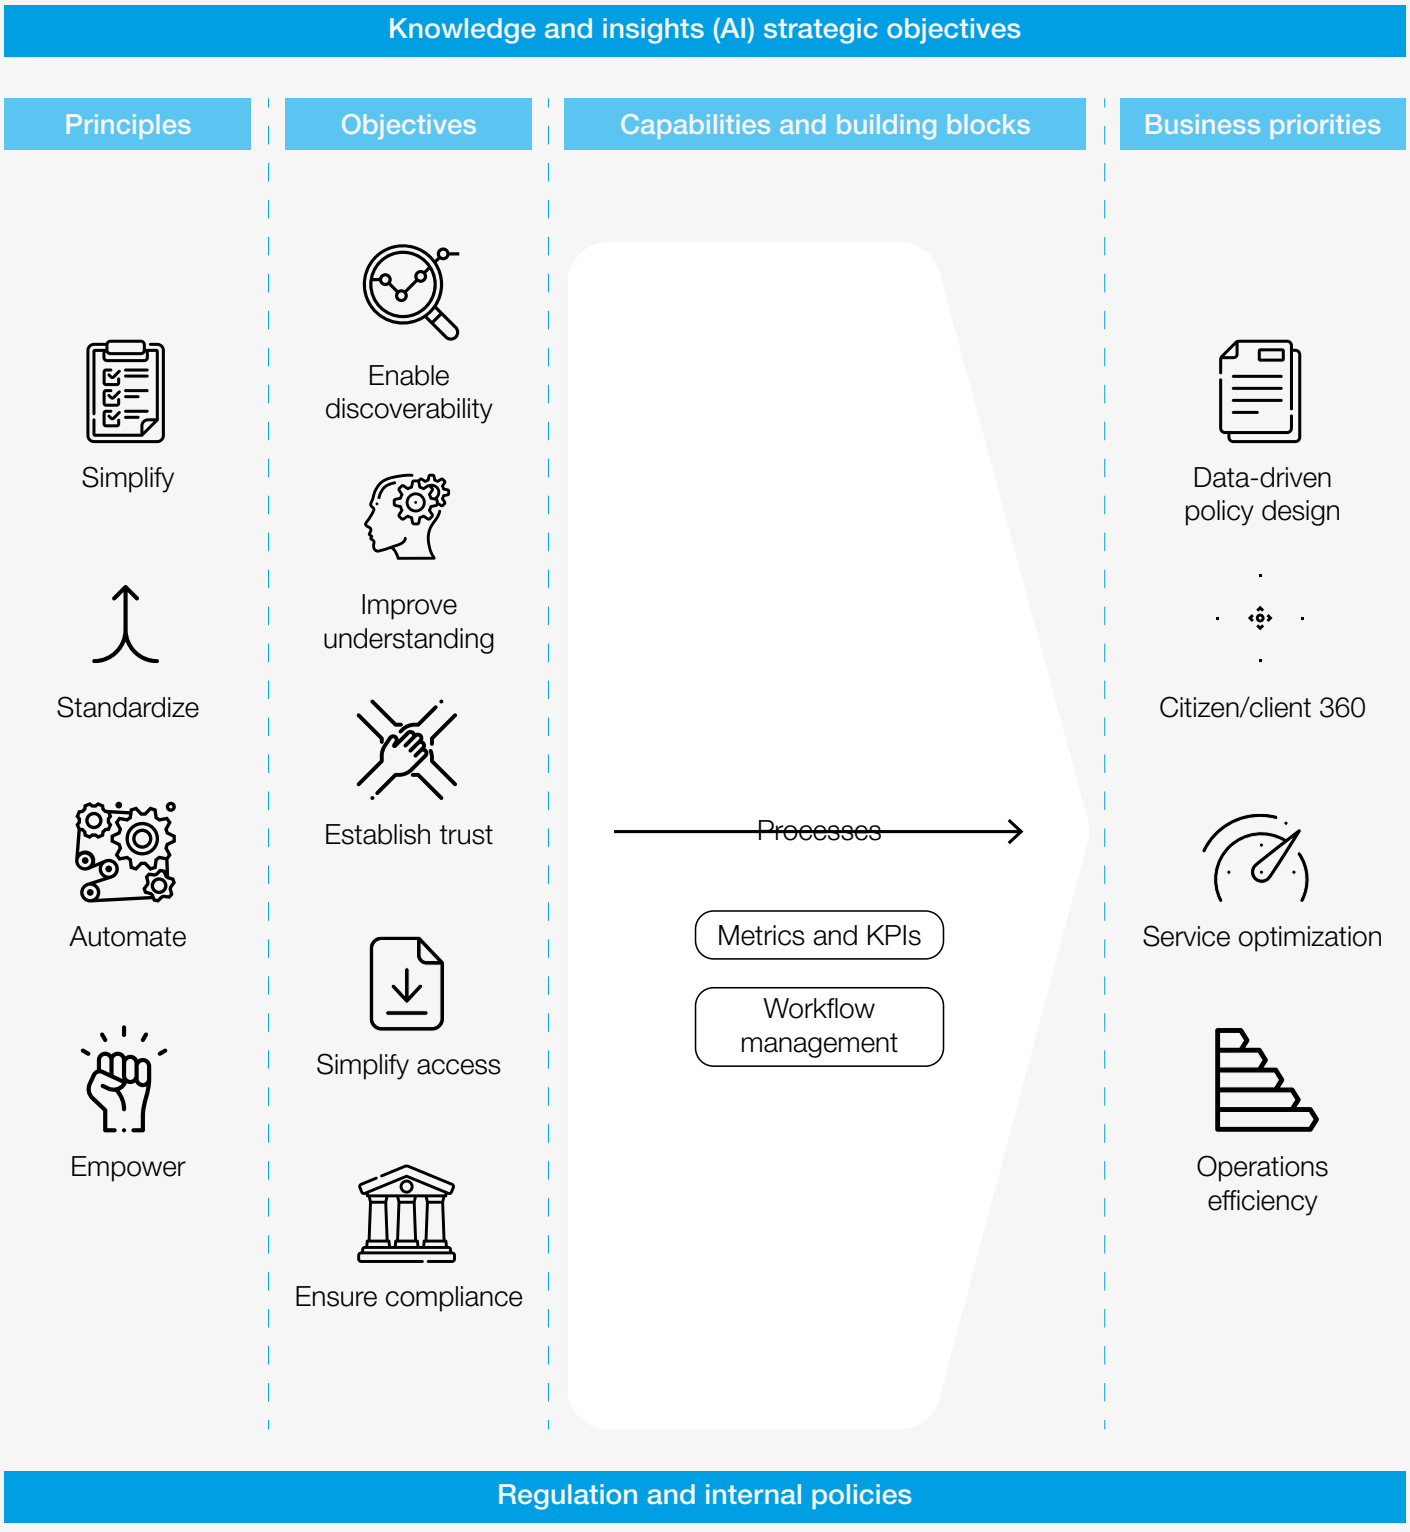

## 6.6 Highlight the technical and ethical limitations of using the data to avoid issues such as bias.

> Though available, legal to use and proportionate to need, there may be limitations to data (e.g. data bias) that make an AI approach inappropriate, unreliable or misleading.

**Why is this important?**

Though available, legal to use and proportionate to need, there may be limitations to data (e.g. data bias) that make an AI approach inappropriate, unreliable or misleading.

**a. Consider the susceptibility of data that could be in scope and whether usage of the data is fair.**

- As important as data protection is, not all data is sensitive (e.g. open-government data is freely accessible online). All data, sensitive or not, must have its integrity safeguarded, but it is not necessary to keep non-sensitive data behind closed doors. It is important to assess the privacy needs of different datasets to determine the right level of protection. Normally, personally identifiable information (PII), such as financial and health data, is considered extremely sensitive. The RFP needs to reflect data governance requirements for both the procurement process and the project that are in accordance with the nature of the data.
- Select data that fits criteria of fairness. For example, the data should be representative of the population that the AI solution will address, as well as being reasonably recent.[5]

**b. Highlight known limitations of the data (e.g. quality) in your RFP and require tenderers to describe their strategies on how to address these shortcomings. Have a plan for addressing relevant limitations that you may have missed.**

Considerations when deciding if a source of data is suitable include:[6]

- Representativeness (whether the data accurately represents the segment of the population in scope for the AI solution)
- Provenance (including how and why the data was collected)
- Gaps in data quality (e.g. many values missing from a particular data element)
- Bias present in the data (if it is not representative of the population to which the algorithm will be applied)
- Lack of clarity in metadata (for example, confusing or vague data element names)

- Check data completeness, representativeness and accuracy of potential sources before starting the procurement process. Articulate data quality observations and the apparent limitations and, if possible, share those insights through the RFP. Bidders must be aware of these data considerations during the procurement process or, in cases where data is sensitive, the selected provider(s) must be made aware after the contract has been awarded.
- If you do not have the right skills or means to comprehensively check for possible limitations of your data, provide vendors with guiding insights into the high-level state of the data and its origin,[7] so that they can draft adequate proposals. Also, ensure the RFP's data requirements include performing a comprehensive data quality assessment and, if required, development of mitigation strategies for low-quality data.

## 6.7 Work with a diverse, multidisciplinary team.

**Why is this important?**

Developing and fulfilling a proper AI RFP will require a diverse team that understands the interdependent disciplines that AI covers, including: domain expertise (e.g. healthcare, transportation), systems and data engineering, model development (e.g. deep learning) and visualization/information design, among others.

**a. Develop ideas and make decisions throughout the procurement process in a diverse and multidisciplinary team.**

- Develop an understanding of the skills that are needed to effectively acquire and maintain an AI-powered solution, before starting the procurement process.
- Assemble multidisciplinary teams that specialize in designing, procuring, evaluating and operationalizing AI systems. These multidisciplinary teams should include expertise in: policy from the domain (e.g. justice) in which the AI solution will be applied, machine learning/data science, data engineering, technology (software and hardware), procurement, ethics and human rights.[8]

– Ensure that you have a diverse team. This should include people from different genders, ethnicities, socioeconomic backgrounds, disabilities and sexualities. You should also make sure that there is a mix of perspectives and points of view. This ensures that problems and solutions are tackled from different angles and helps to mitigate bias.

– This is important when it comes to evaluating tender responses. You need to be certain that you have the right expertise in your team to compare AI-driven solutions. Technical, business as well as legal and ethical experts are needed to score the different bids. You can integrate processes in your procurement decision to ensure that a multidisciplinary evaluation is mandatory. If expertise is lacking within your team, you can reach out to pools or professional networks within your organization or across the civil service.

Note that many value-laden decisions will likely be made during development (i.e. post-procurement), and it is essential that your team maintains the skills, or at the very least access to expertise, to ensure that all important decisions and trade-offs are made or overseen internally, rather than exclusively by a contractor or vendor.

**b. Require the successful bidder(s) to assemble a team with the right skill set.**

– As part of your requirements, ensure bidders provide evidence of the skills and qualifications of the proposed project resources who will develop and deploy the AI solution.[9] This should be part of the RFP response and it should be one of your decision criteria when evaluating the proposals.

## 6.8 Focus throughout the procurement process on mechanisms of accountability and transparency norms.

**Why is this important?**

To build public trust in the legitimacy of the AI system, the procurement process should enable accountability in understanding how the AI solution works, so that it can be evaluated independently and thus promote a culture of responsibility over the AI solution life cycle.

**a. Promote a culture of accountability across AI-powered solutions.**

– Public institutions cannot rely on black-box algorithms to justify decisions that affect individual and collective citizens' rights, especially with the increased understanding about algorithmic bias and its discriminatory effects on access to public resources. There will be different considerations depending on the use case and application of AI that you are aiming to acquire, and you should plan to work with the supplier to explain the application for external scrutiny, ensuring your approach can be held to account. These considerations should link to the risk and impact assessment described in Guideline 2. Under certain scenarios, you could consider making it a requirement for providers to allow independent audit(s) of their solutions. This can help prevent or mitigate unintended outcomes.

– Providers and public officials should incorporate risk analysis for the unexpected and unintended effects of AI-powered solutions, within the limits prescribed by the law, and specify their respective responsibilities in the contract. Note that the laws and standards for assigning accountability may differ according to jurisdiction. For example, the Canadian federal government's Directive on Automated Decision-Making requires the associate deputy minister of the respective federal entity to sign off on an algorithmic impact assessment (AIA) as part of an AI project.

– Consider how applicable accountability requirements in law, such as freedom of information legislation and data- protection logging requirements, will be implemented throughout the project life cycle.

**b. Ensure that AI decision-making is as transparent as possible.**

– Encourage transparency of AI decision-making (i.e. the decisions and/or insights generated by AI). One way to do this is to encourage the use of explainable AI. You can also make it a requirement for the bidder to provide the required training and knowledge transfer to your team, even making your team part of the AI-implementation journey. Finally, you can ask for documentation that provides information about the algorithm (e.g. data used for training, whether the model is based on supervised, unsupervised or reinforcement learning, or any known biases).

– Documentation is especially important when the algorithm is a pre-packaged solution that the bidder will bring to the project, as opposed to an algorithm that will be built and/or customized as part of the upcoming project. Finally, you can also ask bidders to provide information on their model-building methodology, including how they select variables, build samples (where applicable) and validate the model. Be aware, however, that algorithm-building is an iterative

process and that it depends on creativity as much as it does on science.

– Documentation provided by a bidder will give you directional awareness of their practices and methods; it will not give you a step-by-step guide that details exactly what would be done during the project, as the exact process will invariably shift from project to project to meet the needs of each scenario.

**c. Explore mechanisms to enable interpretability of the algorithms internally and externally as a means of establishing accountability and contestability.**

– With AI solutions that make decisions affecting people's rights and benefits, it is less important to know exactly how a machine-learning model has arrived at a result if we can show logical steps to achieving the outcome. In other words, the ability to know how and why a model performed in the way it did is a more appropriate means of evaluating transparency in the context of AI. For example, this might include what training data was used, which variables have contributed most to a result, and the types of audit and assurance the model went through in relation to systemic issues such as discrimination and fairness. This should be set out as documentation needed by your supplier.

– It is also important to consider the potential tension between explainability and accuracy of AI when acquiring AI solutions. Classic statistical techniques such as decision-tree models are easier to explain but might have less predictive power, whereas more complex models,such as neural networks, have high predictive power but are considered to be black boxes. Given these challenges you should think carefully about.

FIGURE 5 **Diagram to explain what is meant by a "black box" algorithm and why they're an issue**

In a traditional model where a service provider interfaces with the service recipient, the recipient can communicate back and forth with the service provider regarding an outcome and/or determination. The recipient can ask questions regarding a decision and challenge an outcome.

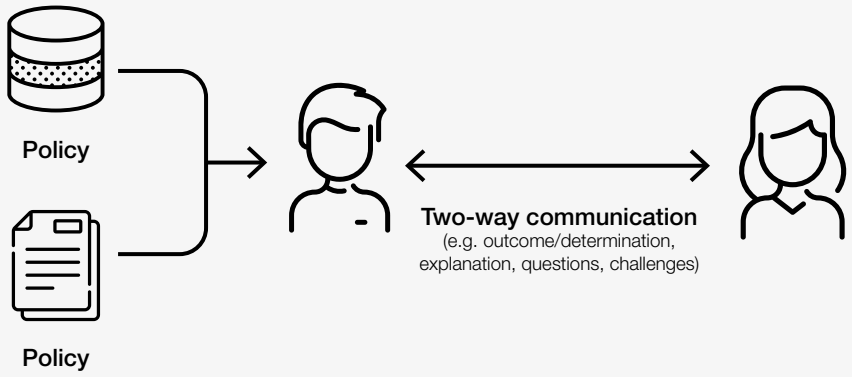


With a fully automated system that uses a technique such as neural networks, the service recipient cannot expect to understand the outcome. This is because certain algorithms, such as neural networks, are very accurate but do not explain their path to a decision.

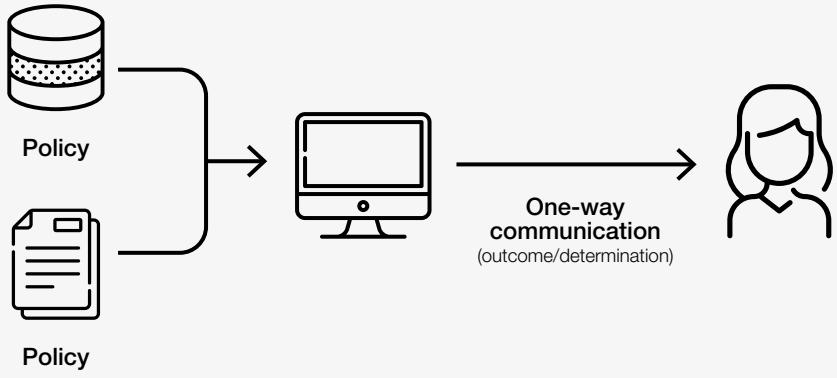

BOX 3 | **Sample type of documentation to ask for: Google Model Cards**

Machine learning models are often distributed without a clear understanding of how they function. For example, under what conditions does the model perform best and most consistently? Does it have blind spots? If so, where? Model cards address that issue by providing information about a model's performance and limitations. These "cards" are short documents accompanying trained machine learning models that provide benchmarked evaluation in a variety of conditions. They are aimed at experts and non-experts alike. Developers can use them to make better decisions about what models to use for what purpose and how to deploy AI responsibly. For journalists and industry analysts, they might provide insights that make it easier to explain complex technology to a general audience and they might even help advocacy groups better understand the impact of AI on their communities.

FIGURE 6 | **Model card example**

| Cards content | Sample information provided |
|---|---|
| **01** Overview of the model | – Simple description of the model<br>– Input data and output of the model<br>– Model architecture used |
| **02** Limitations | – Factors that might degrade the model's performance<br>– Situations in which the model might perform less than optimally |
| **03** Performance | – Model's performance on various evaluation datasets drawn from different sources than the training data |

BOX 4 | **Solution to address explainability: example from Google Cloud AI Explainability**

The most useful models are often the most explainable, as they are the most trusted. Cloud AI Explanations help developers and enterprises understand why their AI model made the decisions it did by quantifying how each data factor contributes to the output. They can use this information to improve the models or share useful insights with their end users. The What-If tool, an interactive visual interface, also allows users to investigate model behaviour by using dataset visualization to explain performance. AI Platform users can develop a deeper understanding of how their models work under different scenarios and build rich visualizations to explain model performance to business users and other stakeholders.

## 6.9 Implement a process for the continued engagement of the AI provider with the acquiring entity for knowledge transfer and long-term risk assessment.

**Why is this important?**

The functionality and consequences of AI systems may not be apparent in the procurement process and often become evident only over the duration of its application, requiring extended communication and information-sharing between the procuring entity and the system developer.

For AI systems in the public sector, sustainable and ongoing evaluation methods and means of providing feedback on the data model are crucial to ensure that the tool's use remains ethical. You should make clear in your RFP that this should be considered by the provider and discussed as part of the procurement process.

**a. Consider during the procurement process that acquiring a tool that includes AI is not a one-time decision; testing the application over its lifespan is crucial.**

> "The tool will need support during its life cycle."

- The tool will need support during its life cycle. Knowing where to go for that support and how much support is available will be vital for getting the most out of any tool. Accepting the potential impact of any support gaps or employing outside expertise both come at a cost. This should be factored in when purchasing an intelligent tool.

- Consider the implementation of a process-based governance framework that provides a template for the integration of the norms, values and principles that inform the procedures and protocols organizing the project workflow.

> "Testing the model on an ongoing basis is necessary to maintain its accuracy."

- Testing the model on an ongoing basis is necessary to maintain its accuracy. An inaccurate model can result in erroneous decisions and affect users of public services.

  Therefore, you should establish with the provider how the efficacy of the model will be monitored once deployed.

**b. Ask the AI provider for knowledge transfer and training to be part of the engagement.**

- Make knowledge transfer a requirement under the RFP. Evaluate the thoroughness and logic of the knowledge- transfer plan to ensure that government resources will be able to use the tool appropriately on their own once the project is finalized.

- Set out clearly your expectations for project documentation. Ensure that maintenance and auditing of the AI solution would be possible by entities independent of the vendor.

**c. Ask the AI provider for insights on how to manage the appropriate use of the application by non-specialists.**

- Operational or service staff must have enough knowledge or training on the solution to understand how to use it and successfully exploit its outputs. You should address the need for enough training and support to avoid the misuse of AI applications with the AI provider. The application must make it easy to report any suspected unauthorized behaviour to the relevant authority(ies) within and/or outside the organization. Enable end-to-end auditability with a process log that gathers the data across the modelling, training, testing, verifying and implementation phases of the project life cycle. Such a log will allow for the variable accessibility and presentation of information with different users in mind to achieve interpretable and justifiable AI.

**d. Make ethical considerations part of your evaluation criteria for proposals.**

- There are robust ethical practices that you should require suppliers to demonstrate when providing AI solutions. Leading AI-solution providers have begun to create internal frameworks for the ethical design, development and deployment of AI, which cover processes to ensure accountability over algorithms, avoiding outputs of analysis that could result in unfair and/or biased decision- making, designing for reproducibility, testing the model under a range of conditions and defining acceptable model performance. Bidders should be able not only to describe their approach to the above, but also to provide examples of projects, complete with client references, where these considerations have been followed.[10]

- Make comprehensive, transparent algorithm assessment one of the requirements in the proposal and, if applicable, state minimum performance metrics that the model must meet. If possible, work with bidders to determine what the thresholds should be. As part of testing the model, you should work with the provider to establish how often you need to update the model with new data. Testing over the lifespan of the model for suitability and accuracy is highly important, especially when the AI is supporting critical services.

## 6.10 Create the conditions for a level and fair playing field among AI solution providers.

**Why is this important?**

Government spending can be used to create a fair, competitive market, which leads to better AI systems. In addition, early engagement with AI vendors can result in more relevant responses, increasing the probability of success for the procurement and the subsequent project.

While AI systems generate new challenges that you need to reflect within the requirements and procurement approach, you must be proportionate in your approach and not impose any unnecessary burdens that would deter a wide diversity of suppliers, including small and medium sized enterprises (SMEs), Voluntary, Community and Social Enterprise (VCSE) suppliers and those owned by under-represented groups, from competing for public contracts.

**a. Reach out in various ways to a wide variety of AI solution providers.**

– Given the rapidly developing landscape of AI service providers, largely comprising smaller enterprises such as start-ups, consider non-traditional methods of market engagement to attract AI solution providers. For example, explain the needs that lead to the proposal through in-person presentations, webinars, information sessions at co-working spaces and/ or online platforms such as LinkedIn or Twitter.

– Consider reaching out to non-traditional stakeholders, such as research institutes and academia. In some cases, these may have the right skills to be part of an AI implementation, and in all cases, they can act as advisers.[11]

– You should ensure that you have taken action to attract a wide diversity of suppliers to bid such SMEs, VCSEs and other under-represented businesses. You should test your approach to ensure it will not deter bidders or create unnecessary burdens on them either during the bidding process or during contract delivery. You must be proportionate in your approach.

– Keep in mind that successfully designing and deploying AI in organizations as big and complex as public agencies requires much more than technical expertise. It requires experience in change management, familiarity with public organizations, and the ability to manage complex projects.

**b. Engage vendors early and frequently throughout the process.**

– Market engagement is a process; it takes place prior to procurement and aims to identify potential bidders and/ or solutions, build capacity in the market to address challenges and opportunities, and inform the design of the procurement and contract.

– Early engagement between government and industry is vital to a successful AI purchasing campaign. Early supplier engagement can help to determine the scope and feasibility of the RFP and, in turn, the most appropriate way to design and structure the requirements, increasing the likelihood that the winning bidder will meet your needs at a competitive cost. Ways to engage vendors early include having vendors provide inputs on possible evaluation criteria for the RFP, and hosting vendors to walk them through the RFP. Approaches like this are already being deployed in Canada, for example, and greatly help government and the private sector increase the effectiveness of procurement.

– To mitigate any risks that could be associated with market engagement (e.g. commercial confidentially, protection of intellectual property [IP], fettering discretion of tender process), be sure to broadly advertise the engagement opportunity, allow all interested parties to participate, ensure that there is adequate time for responses and reasonable time for bidder selection and, where appropriate, that RFP responses can be marked as confidential.

**c. Ensure interoperability of AI solutions and require open licencing terms to avoid vendor lock-in.**

– Consider strategies to avoid vendor lock-in, particularly in relation to black-box algorithms. These practices could involve the use of open standards, royalty-free licensing and public domain publication terms.

– During the design and deployment of the AI solution, it is likely that either a new algorithm will be designed, or an existing one will be tailored (e.g. retrained through your data). It is therefore useful to consider whether your department should own that IP and how it would control it. The arrangements should be mutually beneficial and fair, and require royalty-free licencing when adopting a system that includes IP controlled by a vendor.

– In order to preserve access to systems that become obsolete, ensure the ability to reverse-engineer the system to allow for maintenance of the AI solution independent of the vendor.

# 7 Acknowledgements

The World Economic Forum's Unlocking Public Sector Artificial Intelligence project, in collaboration with the Government of the United Kingdom, Deloitte Consulting and Splunk is a global, multistakeholder and cross-disciplinary initiative intended to help shape the public sector's adoption of AI, and emerging technologies in general, around the world. The project has engaged leaders from private companies, governments, civil society organizations and academia to understand public-sector procurement of AI technology, identify challenges and define principles to guide responsible and ethical procurement. The opinions expressed herein may not correspond with the opinions of all members and organizations involved in the project.

## Lead authors:

**Sabine Gerdon**
Artificial Intelligence and Machine Learning Fellow, World Economic Forum, Seconded from the Office for Artificial Intelligence, Government of the United Kingdom

**Eddan Katz**
Project Lead, World Economic Forum

**Emilie LeGrand**
McGill University Integrated Management Student Fellow

**Gordon Morrison**
Director of EMEA Government Affairs, Splunk Inc.

**Julián Torres Santeli**
Artificial Intelligence and Machine Learning Fellow, World Economic Forum, Seconded from Deloitte Canada's AI practice

We would like to thank our Unlocking Public-Sector AI project community as well as the following contributors for their insights:

**Rashid Alahmedi**
Senior Specialist Technolgy and Solutions, Dubai Electricity and Water Authority

**Greg Ainslie-Malik**
Machine Learning Architect, Splunk Inc.

**Jesus Alvarez-Pinera**
Head of Data, Food Standards Agency

**Shelby Austin**
Managing Partner, Growth and Investments and Omnia AI, Deloitte

**Yousef Al-Barkawie**
Partner, Analytics and Cognitive Middle East Leader, Deloitte

**Neil Barlow**
Head of Vehicle Policy and Engineering, Driver and Vehicle Standards Agency

**Kathy Baxter**
Architect, Ethical AI Practice, Salesforce

**Lorena Cano**
Digital Trade Fellow, World Economic Forum from Inter-American Development Bank

**Ashley Casovan**
Executive Director, AI Global

**Michael Costigan**
Artificial Intelligence and Machine Learning Fellow, World Economic Forum from Salesforce

**Sue Daley**
Associate Director, techUK

**Nihar Dalmia**
Government and Public Sector AI leader for Deloitte Canada, Deloitte

**Gourav Dhiman**
Business Development Manager, XLPAT

**Cosmina Dorobantu**
Deputy Director of Public Policy Programme, The Alan Turing Institute

**Leslie Harper**
Senior Sector Specialist, Inter-American Development Bank

**James Hodge**
Chief Technical Adviser, Splunk Inc.

**Hamad Karam**
Senior Specialist Artificial Intelligence, Dubai Electricity and Water Authority

**Andrew Kim**
Head of AI Policy, Google Cloud

**Steven Knight**
AI Lead, Food Standards Agency

**Benjamin Leich**
Economic Adviser, Better Regulation Executive

**Katherine Mayes**
Programme Manager, techUK

**Maha Mofeez**
Chief Corporate Officer,
Bahrain Economic Development Board

**Valesca Molinari**
Automotive and Autonomous Mobility Fellow,
World Economic Forum from Baker McKenzie

**Mariam Al Muhairi**
Head, Centre for the Fourth Industrial Revolution
United Arab Emirates

**Khalid Al Mutawa**
Director, Bahrain Information
and eGovernment Authority

**Brandie Nonnecke**
Founding Director, CITRIS Policy Lab

**Arwa Al Qassim**
AI Lead, Centre for the Fourth Industrial Revolution
United Arab Emirates

**Ana Rollan**
Artificial Intelligence and Machine Learning Fellow,
World Economic Forum from BBVA

**Nada Al-Saeed**
Data Policy Fellow, World Economic Forum from
Bahrain Economic Development Board

**Komal Sharma Talwar**
Director, XLPAT and TT Consultants

**Leonard Stein**
Senior Strategic Adviser, Splunk Inc.

**Jitin Talwar**
Founder, XLPAT and TT Consultants

**Sandeep Singh Kohli**
Co-founder, XLPAT

**Ahmad Al Tawallbeh**
Specialist Artificial Intelligence, Dubai Electricity
and Water Authority

**Abbey Thornhill**
Assistant Economist, Better Regulation Executive

**Adrian Weller**
Programme Director for AI, The Alan Turing Institute

**Mark Woods**
Director, Technology and Innovation, Splunk Inc.

**Tim Woodbury**
Director of State and Local Government Affairs,
Splunk Inc.

Thank you also to the teams in the UK from the Defence Science and Technology Laboratory, the Department for Transport, the Home Office Accelerated Capability Environment and local governments that supported the user testing and piloting. The steering and working group from the Department of Digital, Culture, Media and Sport, the Government Digital Service, the Cabinet Office, the Crown Commercial Service and the Centre for Data Ethics and Innovation has been instrumental to progressing this work, in particular:

**Sue Bateman**
Deputy Director for Policy and Innovation,
Government Digital Service

**Oliver Buckley**
Executive Director, Centre for Data Ethics and
Innovation

**Stephen Hennigan**
Deputy Head of Office for Artificial Intelligence,
United Kingdom Government

**Sana Khareghani**
Head of Office for Artificial Intelligence, United
Kingdom Government

Thank you to everyone who contributed through interviews, workshops and discussions in the last 18 months in Dalian, Dubai, London, Manama, San Francisco, Tianjin, Toronto and Washington DC.

# 8 Endnotes

1. Definition from the Engineering and Physical Science Research Council, a UK government research funding body.
2. For the complete results, see https://services.google.com/fh/files/blogs/bsr-google-cr-api-hria-executive-summary.pdf
3. For an up-to-date list of freedom of information laws around the world, see https://en.wikipedia.org/wiki/Freedom_of_information_laws_by_country (link as of 29.05.2020).
4. For more information on data anonymization, refer to: "Guide to basic data anonymisation techniques", Personal Data Protection Commission, Singapore. 25 January 2018.
5. For more information on fairness during data selection, refer to: "Understanding artificial intelligence ethics and safety. A guide for the responsible design and implementation of AI systems in the public sector", section "Data fairness", David Leslie, the Alan Turing Institute.
6. For more information on data selection criteria, refer to: "Understanding artificial intelligence ethics and safety. A guide for the responsible design and implementation of AI systems in the public sector", section "Data fairness", David Leslie, the Alan Turing Institute.
7. For example, summary statistics such as number of rows present, number of missing values for each data field, description of how the data is collected and processed.
8. For more information on the domain and technical skills required to deliver an AI engagement, refer to: "Searching for superstars isn't the answer. How organizations can build world-class analytics teams that deliver results", Deloitte.
9. ibid.
10. AI ethics is a deep and evolving field, and there are various publications on the matter, including those listed below. Refer to these sources for a full background on the topic.

    – "OECD principles on artificial intelligence", Organizations for Economic Co-operation and Development

    – "Ethics guidelines for trustworthy AI", Independent High-Level Expert Group on Artificial Intelligence set up by the European Commission

    – Understanding artificial intelligence ethics and safety. A guide for the responsible design and implementation of AI systems in the public sector", section "Data fairness", David Leslie, the Alan Turing Institute.

    – "For a meaningful artificial intelligence. Towards a French and European strategy", Cédric Villani
11. Examples of organizations include the Alan Turing Institute in the UK and the Vector Institute, MILA, and the Alberta Machine Intelligence Institute in Canada.

WORLD ECONOMIC FORUM

COMMITTED TO IMPROVING THE STATE OF THE WORLD

The World Economic Forum, committed to improving the state of the world, is the International Organization for Public-Private Cooperation.

The Forum engages the foremost political, business and other leaders of society to shape global, regional and industry agendas.

World Economic Forum
91–93 route de la Capite
CH-1223 Cologny/Geneva
Switzerland

Tel.: +41 (0) 22 869 1212
Fax: +41 (0) 22 786 2744
contact@weforum.org
www.weforum.org