Unlocking Public Sector AI

AI Procurement in a Box:
Challenges and opportunities during implementation

TOOLKIT

JUNE 2020

WORLD
ECONOMIC
FORUM

Contents

Background

The World Economic Forum's Centre for the Fourth Industrial Revolution launched the project Unlocking Public Sector in AI, which offers AI procurement guidelines for government and public-sector organizations in November 2018. The Forum's goal is to help officials better understand this rapidly developing technology and mitigate potential risks. The guidelines are being piloted with government departments and agencies around the world, beginning with Bahrain and the United Kingdom, along with the Dubai Electricity and Water Authority in the United Arab Emirates.

The pilots have centred on workshops that focus on analysing the guidelines in the light of a specific potential use case of AI to test their applicability and gather feedback. The main aims of the workshops have been to:

– Gather feedback on the guidelines

– Test the guidelines in an example scenario

– Move from theory to practice and develop a roadmap for action

What is the purpose of this document?

This document is an overview of the key findings from the workshops conducted in Bahrain, the United Arab Emirates (UAE) and the United Kingdom (UK), which focused on moving the guidelines from theory to practice.

The document provides insights, identified by workshop participants, of key themes and aspects to consider when implementing the guidelines.

Discussions: Moving the guidelines from theory to practice

Workshop participants explored various themes related to the government's use of AI and how procurement plays a role in government adoption of the technology. The themes identified as priorities were:

– Data

– People: Skills, culture and leadership

– Procurement processes

– Ethics: Accountability, liability, transparency

(1) Data

Participants identified data governance as a major challenge for many public-sector organizations looking to adopt AI. This related to the foundational technology as well as the process required to successfully manage and leverage data. Levels of data maturity vary throughout government agencies, with different agencies finding themselves at different points in the data journey.

The discussions identified the following as key initiatives to move from theory to practice and apply the guidelines:

1. Exploring the potential and dos and don'ts of data sharing as part of the procurement process: While the benefits of sharing data with vendors during the procurement process were recognized, it was acknowledged that it would be challenging given the current regulatory environment, but trials and pilots could be considered to explore the potential and share best practices. Participants also identified that more broadly, sharing data between government agencies is necessary to address complex policy problems, but is currently a significant challenge. In the UAE, Dubai's answer to this problem is provided by Smart Dubai, a government agency tasked with improving the experience of living and visiting Dubai. Smart Dubai is creating the infrastructure and processes that enable government agencies to share data. While the initiative is relatively young, it is an example of tangible government action aimed at increasing data sharing.

2. Mapping opportunities related to third-party data: More analysis and guidance is needed on how to obtain and utilize third party data, whether it is procured or acquired through partnerships.

3. Ensuring data integrity: Organizations need to be able to trust their data. This requires data management practices that ensure the integrity of data. Practices can be borrowed from red teaming (i.e. data monitoring by different individuals/systems) also known as penetration to stress-test data. The exercise would consist of having data-savvy resources (beyond than those that created the data) attempt to derive personally identifiable information (PII) from the merged data set.

4. Setting out a national data governance strategy and providing public entities with an adoption roadmap: Since data governance is a cross-cutting theme that is important for every government entity, and entities need to share data with one another, it is crucial that entities within a country share a common approach to data governance. The UK government is working on a National Data Strategy[2] and a national data standards framework. It is likely that departments and agencies' efforts to develop data strategies will increase once the standards are in place.

5. Developing and implementing data action plans within agencies and teams: Data action plans tailor national strategies for the needs of specific groups at different levels (e.g. agency, department, team, sector etc.). An example is the Joint Rail Data Action Plan published by the UK government in August 2018.[3] It is a framework for government and the rail industry to foster greater data openness.

❝ **In the UAE, Dubai's answer to this problem is provided by Smart Dubai, a government agency tasked with improving the experience of living and visiting Dubai.**

(2) People: Skills, culture and leadership

Leadership, skills and culture are key to the implementation of the guidelines. AI change management, skills development and partnership between the public and private sector were the key sub-themes that were discussed. It was highlighted that the skills gap is potentially much bigger for AI procurement than for conventional procurement because it not only calls for multidisciplinary teams, but also very specific technical expertise. Mindset shifts and a greater appetite for risk are also needed to drive innovation and trial emerging technologies in a public-sector context. AI change management also includes business process engineering. Processes need to be effectively analysed and studied so that they can be improved and streamlined for optimum efficiency in terms of performance and cost. Since most likely you are not starting from scratch with a blank sheet of paper it is crucial to evaluate the process, analyse data and gather evidence to ensure that new considerations are aligned with the overall approach.

The discussions identified the following as key initiatives to move from theory to practice and apply the guidelines:

– **Putting a greater focus on knowledge sharing throughout different professions and the public and private sectors**: Knowledge and best practice sharing could be a key addition to other traditional up-skilling measures. It was suggested that the establishment of a multistakeholder group that provides knowledge sharing of AI projects throughout government would support this.

– **Focusing on bringing in experts**: Teams throughout government need to attract, develop and retain the people and skills required to achieve government transformation.

– **Up-skilling senior civil service officials in AI**: Provide training to ensure that public servants are aware of the opportunities and challenges of AI in the public sector, and can shape change and make effective decisions. Participants suggested that given the importance of leadership support, training can start with senior officials.

– **Obtaining strategic direction and support from senior leadership**: It is important to train senior leaders to better understand AI and develop and communicate an AI vision for their organization. Leaders must support their teams in increasing the impact of AI technologies.

– **Ensuring that everyone in the organization understands the benefits and impact of AI adoption**: When focusing on AI adoption it is important that there is an initial focus on quick wins for AI deployment and that it helps staff with their daily work. This includes demonstrating the benefits and measuring the impact of models in use.

– **Working in partnership with private sector organization**: Public-private partnerships could ensure that change is managed effectively. This could be done through proof of concepts that demonstrate the validity and value of AI solutions, or through programmes such as fellowships, secondments, hackathons and workshops that showcase the ability of AI to enable government to accomplish its goals.

– **Bringing together the right skills to evaluate tenders with ethical considerations**: It is important to have ethics skills as part of the procurement team. If they are not available in-house they should be built or borrowed. For example, in the case for the UK bodies such as Centre for Data Ethics and Innovation[4] or the Alan Turing Institute[5] can assist.

③ Procurement processes

To move from theory to practice as well as apply the guidelines determined by the participants, the following should be considered:

The discussions identified the following as key initiatives to move from theory to practice and apply the guidelines:

– **Developing a new procurement vehicle for AI**: Current procurement frameworks do not actively enable the agile nature of AI development and deployment nor do they address ethical questions. It was suggested that the need for separate new AI vendor lists or frameworks be considered as well as whether all tech procurement approaches should include ethical evaluation criteria by design. Dynamic purchasing systems, such as the emerging technology marketplace SPARK[6] in the UK, could be used to increase the supplier base and pre-qualify ethical and innovative providers.

– **Building ethics-by-design into the procurement process**: It is important to define the ethical requirements of the tool. It is also helpful to determine at which points during the procurement process you should perform an ethics assessment. For example, this could be done as part of pre-qualification of suppliers for a procurement framework or as part of each tender.

– **Developing and including straightforward definitions and a glossary**: This could address the right terms to use for AI in commercial discussions (e.g. data versus IP rights, trained versus untrained models), as well as model contract examples.

– **Introducing ethical oversight into the procurement process**: A body composed of senior leaders should exist to provide ethical oversight of the applications of data. For example, in the UK the Office for National Statistics has an Ethics Board[7] that reviews and approves usage of data based on ethical principles.

– **Reshaping business cases to measure holistic benefits**: Traditional business cases make it difficult to justify spending money for pre-commercial activities and proof of concepts, and currently fail to sufficiently capture non-monetary gains.

– **Choosing the right opportunities for AI deployment**: AI is particularly adept at certain tasks and focusing on those will make it easier to reap the benefits faster. In the beginning of the AI deployment journey, departments and agencies should focus on the immediate opportunities and explore AI use cases that are very likely to succeed. For example, when it comes to case work it is easier to automate the simple cases, allowing case workers to spend more of their time on the more complex ones.

Knowledge sharing

– **Setting up institutional processes to have a single point of contact or a repository of AI knowledge and information**: To be able to include the procurement within a strategy for AI adoption it would be helpful to have an overview of different AI-related government initiatives. Currently it is difficult to access this information and a team or organization such as the UK's Office for AI[8] could curate and share this information as well as provide coordination throughout government.

– **Sharing best practices for AI procurement and deployment across government**: There is a need to share leading practices to ensure that teams can learn from the challenges and successes of other projects. For example, sharing best practice throughout government on agile procurement helps to show that it is possible to comply with procurement rules and regulations while at the same time accelerating the processes.

Pre-market engagement

When conducting pre-market engagement, the issues that public officials need to consider and raise with suppliers include:

– Data requirements: What data do you need?

– Agile AI system development: Should we start with a proof of concept? At what cost could that be delivered? Should we consider outsourcing the scoping of the project to ensure that the challenge is well defined?

– Human in the loop: What is the strategy for keeping the human in the loop in AI deployment and how can AI-systems be integrated into process so that they work effectively with delivery teams?

– Public value and impact of the AI system: How can the public benefit and social value be demonstrated? What are the risks? How can an initial impact assessment be conducted in collaboration?

– Skills, culture and readiness: What skills are available in-house? Is the department ready to adopt the solution and integrate it with its processes? What are the internal challenges to the adoption of AI? Is the contract scalable? What training is required to effectively collaborate?

Lifecycle management

Many of the lifecycle management issues discussed below are applicable for IT procurement in general. However, AI accentuates these issues, given an algorithm's constant state of change resulting from periodic updates to maintain its accuracy. The following were identified as important issues to address prior to any AI procurement project:

– Define your post-deployment support needs and validate them with vendors.

– Establish a holistic and long-term budget plan that includes maintenance costs, such as cloud storage, computing needs, the need for human-AI interaction, bias controls, model retraining and auditing.

– Avoid vendor lock in. Consider that you may want to shift services such as data management and hosting from one provider to another. Interoperability is critical to achieve this.

Intellectual Property

– Be open to innovative risk and pricing agreements: The government and the private sector should have frank conversations about the risk level each can agree to take on for a given AI project, and the respective pricing to match.

– Establish best practice for ownership and usage rights of the AI solution: Ownership and usage rights should be considered and assigned to the party that is best placed to deliver the desired economic and social outcome.

(4) Ethics: Accountability, liability and transparency

Ethical considerations need to be built into the end-to-end lifecycle processes of any technology solution and AI is no exception. This must be stated and ensured at procurement stage. AI developers and public-sector officials must understand the ethical considerations of AI solutions. A technology-centric focus that solely revolves around improving the capabilities of an intelligent system doesn't necessarily consider human needs. An ethical, human-centric approach must be central to any AI deployment.

A common criticism of certain AI applications is the opaqueness of data processing and decision-making. Transparency, interpretability and auditability are important considerations when using AI in the public sector. There are different ways of enabling transparency including: thorough documentation of the data, processes and algorithms, releasing the source code, or simple explanations of the logic of the system aimed at non-experts. In addition, when discussing transparency, it is important to distinguish between the transparency of the AI system and that of the wider organizational decision-making process, which includes the algorithm.

Challenges regarding AI transparency are particularly evident with advanced AI systems, such as deep neural networks. When deploying machine learning algorithms in public-sector organizations, particularly those that can have a significant impact on the lives of citizens (e.g. immigration, law enforcement) it is crucial to ensure that an acceptable level of transparency is designed into the system.

The discussions identified the following as key initiatives to move from theory to practice and apply the guidelines:

–   **Establishing standards for the audit of AI systems and their underlying data**: It is necessary to consider what an audit should entail and what the minimum acceptable outcome should be.

–   **Implementing performance measures and standards to evaluate performance against ethical requirements**: There is a need to find ways to measure performance against ethical principles so as to be able to make informed decisions about those tendering.

–   **Acknowledging data ethics as a shared responsibility between public buyers and AI service suppliers**: There is a need to define, in broad terms, which responsibilities correspond to the government and which to vendors. This could be done at the macro level, for example, through a directive or on a case-by-case basis (e.g. each RFP could be used to define what the expectations are for that specific project). One way to achieve this is to define a set of enforceable standards on the safe storage and usage of data. Participants identified existing standards, such as UK's National Cyber Security Centre (NCSC) principles for cloud security[9], which can serve as a model.

- Articulating and demonstrating the public benefit of AI usage in tangible terms and developing an approach considering the proportionality of the application of AI: It is crucial to make informed choices about AI deployment and consider ethical applications while balancing risk and public benefit. Citizens could also be involved in these choices through citizen juries or online panels. Deploy a framework for processes that explain how the AI model operates in practice as well as how to examine its decisions retrospectively.

- Exploring the value of formal ethics panels: These could be administered by an external and independent body that is empowered to question AI applications in the public sector. The ethics panel could monitor the AI systems and specific applications through life and have technical as well as ethical expertise.

- Designing AI systems with a focus on how humans will interact with them: For example, historic appeals data could be used together with upheld data to warn people when a decision is likely to be appealed. The system could highlight which details of the case would have to be changed for it to be approved. In this specific use case, controlling for bias plays a very important role as well.

- Ensuring that there is a human in the loop for decision-making with direct impact to constituents: A human should have oversight of a machine's decision making. There are two different scenarios. For relatively simple decision-making, a machine can help automate the process and a person can review after a statistically significant percentage of the outcomes. For complex decision-making, it is essential to have human review in place before any action is taken.

- Focusing on iterative system development, user testing and good practice for assessment: There is a need for user testing and formal trials (proportionate to the impact of the solution). Iterating prototypes with customer groups (government and civil society) on a regular basis is important for any digital service, but is even more important when it comes to AI deployment. The expectation that this approach be taken should be set early, during procurement.

Roadmap for guidelines and AI adoption

Government departments and teams need to develop roadmaps and ensure that legislation enables the full implementation of the guidelines.

Examples of initiatives that can be included in a roadmap:

– **Make the guidelines user-centric**

  – Adapting the language in the guidelines to reflect local public procurement vocabulary and provide a technical glossary for AI specific terms.

  – Tie the guidelines to the local procurement cycle to make them more intuitive. If possible, do this in a visual format.

– **Link the guidelines to real-world examples**

  – Link the guidelines to examples and case studies of AI procurement initiatives.

  – Consider the overlap between the different areas of action that are mentioned in the guidelines such as accountability and data governance.

  – Provide examples of innovative data applications within government and the value unlocked by these applications.

– **Provide additional guidance**

  – Highlight what is best practice for general IT procurement versus what's specifically important for AI.

  – Highlight roles and responsibilities for procurement teams as opposed to a project manager and a delivery manager in an AI project.

  – Develop standard terms and conditions for AI/ML projects.

  – Provide sample templates for standard procedures and processes (e.g. request for proposal (RFP), invitation to tender (ITT), request for information (RFI)).

  – Raise awareness and understanding of innovative procurement routes, such as innovation partnerships.

– **Disseminate knowledge throughout procurement teams**

  – Collect procurement teams' success stories and challenges related to the guidelines.

  – Appoint AI procurement champions to share knowledge.

  – Establish a community of practice for AI-interested procurement officials.

  – Provide references to training programmes for procurement officials.

  – Collect best practice on the flexible use of standard procurement procedures for different project phases of an agile delivery concept such as proof of concept, discovery, iterative delivery/testing and deployment.

Workshops in the UK

The Department for Transport partnered with the Office for AI to host the workshop. For the purpose of this pilot workshop a virtual multidisciplinary team of procurement officials, data scientists, digital delivery experts and policy and analytical experience was formed to provide feedback and develop an example case study for the potential application of AI in rail modelling. Participants from the Department for Transport were asked to work through the guidelines using a checklist before developing a fictional ITT to illustrate the application of the guidelines in a potentially real-world scenario. The example scenario was shared with all participants ahead of the workshop and provided the basis for the discussion.

A session with the Department for Transport was held on 4 October 2019 to gather feedback from a diverse set of stakeholders. The use of AI in public transport is a critical opportunity to unlock the value of data to improve the quality and efficiency for the public transport sector, especially in rail transport. AI is already playing a positive role in the rail industry. Analytic tools are helping customers plan and book journeys, and data is helping to provide more accurate real-time journey information. While developments in AI are exciting, there is still some way to go before its potential is unlocked and success depends on numerous elements working in harmony. In this specific case, the workshop focussed on the potential application for AI for transport modelling as an example of a procurement process.

A session with a focus on AI deployment in local government was held on 16 October 2019 and sought to gather feedback from a diverse set of stakeholders. AI tools promise to improve back office functions and deliver efficiency gains and services more effectively. Many councils appear wary of making the initial required investment when budgets are already strained and they lack resources to identify AI applications. Furthermore, they have concerns about ethical considerations. This is, however, slowly starting to change and there are several examples in the UK and abroad that show AI is considered not as hype, but as a genuine enabler of change. This sentiment is likely to increase once AI has successfully been implemented and the benefits are clear. Service demands are high and are unlikely to decline. This, combined with financial pressures, means that local councils and communities' might experience huge benefits through the smart use of technology, such as robotic process automation and AI in the future.

On 25 October 2019, the Forum organized a workshop in collaboration with the Defence Science and Innovation Laboratory (DSTL) and techUK to gather feedback from a diverse set of stakeholders. AI tools promise opportunities to improve back office services and deliver efficiency gains as well as to provide intelligence to address defence and security-related challenges. The application of AI has already been investigated in the defence and security sector as the utilization of large amounts of data to support decision-making has been common for many years. A good example use case is the military readiness assessment, which is defined as the condition of the armed forces and their constituent units, formations and platforms (ships, planes etc).This was described by Deloitte as: "Large data volumes, diverse sources of information, complex interactions, and the need for speed and accuracy make military readiness a problem tailormade for AI to tackle. And if AI can help tackle readiness, it can help the military tackle just about anything."[10]

DSTL's AI hub helps to improve the country's capabilities in the application of AI-related technologies. The laboratory conducts cutting edge research to support the UK's Ministry of Defence (MoD) with opportunities to help keep people safe from different future defence and security threats. DSTL has recently produced guidance on AI projects in the form of a Biscuit Book.[11] Explaining the definitions and differences, latest thinking and developments, The Biscuit Book is something to be dipped into and easily digested – "just like a biscuit with tea". DSTL is at the heart of innovations in this area, working with other government departments, academia and institutions for the defence and security of the UK.

A session with the UK's Home Office's Accelerated Capability Environment was held on 14 October 2019 to gather feedback from a diverse set of stakeholders. The use of AI in policing, law enforcement and immigration has already been widely debated. AI tools promise opportunities to improve back office services as well as delivering services more effectively, but also particularly in these areas of application bring along many challenging ethical questions. While developments in AI are promising, society is still some way from unlocking its potential, and success in these areas depends crucially on ethical questions being answered fully. For example, the UK government has recently pledged to spend more money on the child abuse image database to trial aspects of AI including voice analysis and age estimation to see whether they would help track down child abusers. A paper by the security policy think tank Royal United Services Institute (RUSI) recently contributed to the debate over the use of machine learning algorithms. The paper that focused on predictive crime mapping and individual risk assessment, found algorithms that are trained on police data may replicate – and in some cases exacerbate – the existing biases inherent in the dataset, such as over- or under-policing of certain communities.[12] The Centre for Data Ethics and Innovation therefore also aims to develop a code of practice for the trialling of the predictive analytical technology in policing.[13]

Workshops in the United Arab Emirates

Between 28-31 October, the Dubai Electricity and Water Authority (DEWA) joined the Centre for the Fourth Industrial Revolution and the Dubai Future Foundation to assess the applicability of the procurement guidelines in the Emirati public sector, focusing on DEWA, the first utility in the region to pilot the AI framework. Over the course of the four days the teams discussed the benefits and challenges of adopting the guidelines in DEWA. As part of the exercise, the Forum developed and reviewed an RFP for a chatbot application, which allows DEWA executives to quickly obtain answers to data-related questions. The application, which is a continuation of a customer-facing chatbot deployed by DEWA, highlights the benefits of using chatbots. The DEWA application allows users to retrieve information in considerably less time than if the user were to search manually, resulting in significant productive enhancements for users. The workshop provided DEWA with insights into leading procurement practices for AI. In turn, the Centre for the Fourth Industrial Revolution team learned about the steps taken by the United Arab Emirates to lead in AI. Smart Dubai, a government agency tasked to improve the city experience for Dubai residents and visitors, is working on a government platform for data exchange throughout agencies. Still early in the process, data submission is for now voluntary and most data submitted is aggregated data. Dubai has also worked on AI tools for medical diagnosis with the Dubai Health Authority, as well as citizen-facing chatbots. These early steps demonstrate the type of initiatives governments around the world are considering as they begin their AI journey. The key outcome from the session was a roadmap for adoption of the AI guidelines. Participants agreed to prioritize the creation of a new procurement vehicle for AI and other disruptive technologies.

Workshops in Bahrain

Vendors and government entities were invited to discuss the format of the RFP, as well as the tendering processes for AI solutions.

From 5-7 November 2019, the Centre for the Fourth Industrial Revolution co-hosted, along with the Bahrain Economic Development Board and the Information and eGovernment Authority, a workshop in Manama to explore the applicability of the procurement guidelines in the kingdom's public sector. More than 70 participants attended the event, representing 25 different institutions from government, the private sector, academia and civil society. The participants helped the project team to identify potential high-value AI applications in Bahrain's public sector and test the value of applying the procurement guidelines. Reflecting actual business needs expressed by the Information and eGovernment Authority, a fictional RFP was drafted for the purpose of simulating a potential target state of the procurement process of AI.

Vendors and government entities were invited to discuss the format of the RFP, as well as the tendering processes for AI solutions. The key outcome from the session was a roadmap for adoption of the AI guidelines. Participants agreed to prioritize the creation of a national data strategy, the development of a data sharing policy framework, providing an effective procurement process, upskilling government teams for AI work and creating an IP knowledge hub.

Acknowledgements

Endnotes

1. https://www.smartdubai.ae/

2. https://www.gov.uk/guidance/national-data-strategy

3. https://www.gov.uk/government/publications/joint-rail-data-action-plan/joint-rail-data-action-plan-addressing-barriers-to make-better-use-of-rail-data

4. https://www.gov.uk/government/organisations/centre-for-data-ethics-and-innovation

5. https://www.turing.ac.uk/

6. https://www.crowncommercial.gov.uk/agreements/RM6094

7. https://www.statisticsauthority.gov.uk/about-the-authority/committees/nsdec/

8. https://www.gov.uk/government/organisations/office-for-artificial-intelligence

9. https://www.ncsc.gov.uk/collection/cloud-security

10. https://www2.deloitte.com/us/en/insights/industry/public-sector/ai-military-readiness.html

11. https://www.gov.uk/government/publications/the-dstl-biscuit-book

12. https://rusi.org/publication/briefing-papers/data-analytics-and-algorithmic-bias-policing

13. https://www.gov.uk/government/publications/the-centre-for-data-ethics-and-innovation-calls-for-evidence-on-online-targeting-and-bias-in-algorithmic-decision-making/centre-for-data-ethics-and-innovation-review-on-bias-in-algorithmic-decision-making

WORLD
ECONOMIC
FORUM

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.