

Unlocking Public Sector AI

AI Procurement in a Box:

Pilot case studies from
the United Kingdom

TOOLKIT
JUNE 2020

Contents

3	1. Department for Business, Energy & Industrial Strategy
4	Overview
4	Procurement process
5	1.1 Few government teams have procured AI solutions. What proved to be the opportunities and challenges?
6	1.2 Why were you keen to use the guidelines for AI Procurement?
6	1.3 How did you go about consulting the guidelines?
6	1.4 What made it easier for you to implement the guidelines?
6	1.5 You used a challenge-based procurement process. Do you have any insights that you'd like to share?
7	1.6 Best practice
8	2. Food Standards Agency
9	Overview
9	2.1 Few government teams have procured AI solutions. What proved to be the opportunities and challenges?
10	2.2 Why were you keen to use the guidelines for AI Procurement?
10	2.3 How did you go about consulting the guidelines?
10	2.4 How did you go about the procurement process?
14	Acknowledgements



The UK government used the AI Procurement in a Box toolkit to design their own [guidance](#) for AI Procurement in government. In this document two teams are sharing their experiences of applying this guidance in practice.

1

Department for Business, Energy & Industrial Strategy



Overview

What is the challenge that you are trying to solve with AI?

The Better Regulation Executive (BRE) aims to develop a regulatory system that is simpler for businesses to navigate, while maintaining important protections for citizens and the environment. The historical and incremental build-up of regulations can lead to disproportionate burdens on business, particularly if obligations are poorly co-ordinated. Digital innovations give policy-makers the opportunity, for the first time, to tackle this problem by looking at the UK's regulatory environment as a whole.

The BRE and the department for Business, Energy and Industrial Strategy (BEIS) sought technological solutions to help analyse the cumulative effect of different regulations on business. For example, it was looking for solutions that examine the complexity of regulations, how often they change and the interactions between them. Successful solutions could help the government prioritize future regulatory reform.

The BRE was mainly seeking technology solutions to the following problems:

- How can we analyse the stock of existing regulation and identify which requirements apply to different businesses and sectors?

- How can we assess how challenging individual regulatory requirements are for different businesses and sectors to comply with?
- How can we assess how challenging the cumulative stock of regulation is for different businesses and sectors to comply with?

From the outset, the BRE highlighted that from its experiences there are several things that make regulations more complicated and challenging for businesses to comply with. These include, for example, the length, complexity and scope of regulations, frequent changes to regulations, overlaps between different regulations and the different types and amount of regulation. The data for the project was available online as open data. The majority of the stock of UK regulation is stored at www.legislation.gov.uk. It's free for anyone to use and republish and is updated regularly.

The government team welcomed different approaches to tackling this challenge and wanted to ensure that the solution is accessible and usable for a range of government policy-makers.

Procurement process:

Which vehicle or framework were you using? How long did you tender for?

The project is part of the GovTech Catalyst challenge fund. The Catalyst uses a £20 million fund to help solve public-sector problems (called “challenges”) using innovative digital technology.

The overall programme is delivered in two phases. A decision to proceed with phase two was made after the outcomes from the first phase were evaluated.

Phase 1: Technical feasibility – The first phase involved a feasibility study and research and development (R&D) contracts being awarded up to £50,000 (including value-added tax (VAT)). This is for

each project to demonstrate the technical feasibility of the proposed solution.

Phase 2: Prototype development and evaluation – The second phase involved up to two R&D contracts being awarded to businesses chosen from the successful first phase applicants. Up to £500,000 (including VAT) was allocated for each contract to develop a prototype and undertake testing for up to 12 months.

More information on the procurement process can be found [here](#).

1.1 Few government teams have procured AI solutions. What proved to be the opportunities and challenges?

We had the notion at the beginning of the project and before the start of the procurement that emerging technologies such as AI could help us to make sense of all the data we had. Nevertheless, we consciously did not provide details on what techniques to use in our invitation to tender.

Instead, we shared questions with the suppliers that needed answers and we shared the data that we had access to and thought could be useful in tackling the challenge. This way of working had a lot of benefits for us since the successful suppliers proposed a mix of different techniques and approaches, which we did not think of in the beginning and the project also evolved substantially over time.

From the start of the project, we saw two major opportunities. First, the opportunity to tackle a challenge that we weren't able to tackle previously and second, the opportunity to learn more about AI in government by delivering a project ourselves.

AI enabled us to answer a question that we previously weren't able to answer. The scale and the complexity of the question were too large to be addressed with traditional methods. It would have been too resource intensive and we are sure we would not have gained the same in-depth insights without using machine learning and natural language processing. The opportunity to use these novel techniques and methods allowed us to do something that no one has ever done before as well as come to conclusions that no one was able to reach previously.

The other opportunity was to experience a new way of working. Our team usually focusses on regulating new technologies as well as encouraging regulators to adopt new technologies to improve how they regulate. Therefore, developing an AI solution in partnership with industry allowed us to learn about

the pros and cons of AI adoption in the public sector. It also let us experience first-hand what it takes to overcome AI adoption challenges and how to make the most of the benefits of these new technologies.

Of course, there were challenges that we needed to overcome to deliver a successful AI procurement and project. The overriding question during the whole procurement process was how accurate will the AI-driven solution be at the end? There was a lot of excitement about the novelty of the approach and the potential to answer questions that nobody has answered before, but it was first and foremost key for us that the results of the AI model were helpful and future-proof. Therefore, the results needed to be accurate and interpretable. We focused on procuring an AI solution that added value to our overall processes and that we were able to integrate in our decision-making processes. We did not want to conduct an interesting data science experiment that did not lead to process optimization and that we could not use effectively.

Another challenge was the fact that the BRE, before the start of the project, had little technical expertise in-house. To make effective decisions during the procurement process we partnered with BEIS Digital and the experts at the National Archives. As a result, what at the start seemed to be a challenge actually became an opportunity. We were able to expand our network throughout government, got the right experts on board and reached out to cross-government networks to ask questions and share our lessons learned. We found that there is really a need for multidisciplinary teams in AI procurement. Breaking down organizational silos not only allows you to tap into the knowledge needed to be successful, but also raises awareness of what you are doing and what you are learning throughout your organization.

1.2 Why were you keen to use the guidelines for AI Procurement?

Adoption of AI-driven tools was new for us as a team, but also for the government as a whole. There is the need for skills and expertise to get AI-driven projects right, in particular the procurement of these projects. When we found out about the guidelines, we thought that it is great to have a central repository that provided a guide to what to think about during the different phases of the procurement process. We were

in need of guidance on best-practices and ideas and therefore the guidelines were a useful tool for us.

The framework of the guidelines gave us the structure needed to check our understanding of the requirements, develop a plan for our procurement and the project more widely and cross-check our approach. Most importantly, the guidelines helped us

when writing the invitation to tender and preparing the interviews with suppliers.

The two aspects that we found particularly helpful and that changed the way we were thinking about the procurement were:

1. Explainable AI – Ways and methods to ensure that the results of the AI-driven system are interpretable
2. Audits – The use of third-party audits that we did not consider previously

1.3 How did you go about consulting the guidelines?

We used the checklist to help us structure our approach. We shortlisted the key questions in the guidelines that were particularly important to our project. We did this while drafting the request for proposal (RFP), which helped us to add relevant questions and requirements right away. The guidelines helped us not only to clarify our approach, but also gave us specific ideas about what to prioritize and what further expertise we needed in our team.

The guidelines gave us the confidence and the skills that we needed to run the process and confidently engage with suppliers.

We also got feedback from the suppliers who responded to the RFP; those that consulted the guidelines found them helpful, providing a better understanding of our procurement approach and why we were asking the questions that we did.

Generally, the guidelines were most helpful right at the start of the procurement. It helped most at the planning and preparing stage since they provide an overview of issues that might come up and that users should be aware of. We “skilled up” our team during the process and got experts to join to help shape and guide the process.

1.4 What made it easier for you to implement the guidelines?

Usability of the guidance and accessibility of the concepts and processes discussed are important. It was helpful that the guidelines outlined the procurement process that you need to go through and provided guidance at each step. The checklist helped with preparations at the beginning and was useful to refer to at every stage of the procurement process. It is also important to include in any type of guidance not only what issues to consider, but also why this is important along with ideas on how to address them during your process.

It is important to use plain English because most of our team using the guidelines did not have much

experience with data science and AI techniques, and were learning along the way. In the future, case studies like this one and examples of use cases will make the guidelines even more accessible for a broad audience.

In a project like ours we learned that if the scope of the project is quite broad and you are working with a lot of unstructured data, it is really important to build in assurance processes and mitigation strategies, such as focusing on explainable AI and using third-party audits.

1.5 You used a challenge-based procurement process. Do you have any insights that you'd like to share?

The challenge-based procurement process is useful for projects with a high level of innovation. For us, the approach was helpful because it gave us a lot of flexibility in the procurement process. It allowed us to trial different approaches, conduct a feasibility

study and learn what works for us. It ensured that we scoped out the project extensively and encouraged us to challenge our previous assumptions at different points of time, all of which led to an excellent AI-driven solution that we can use in government.

1.6 Best practice

Take a very agile, user-centred approach to AI procurement. In traditional procurement processes you are sometimes prone to almost design a solution before you go out to tender. As such, it is important to keep questions in an RFP at a high level, the requirements not too onerous and to continue to have an open mind about what responses you might receive even if you started out with an AI project in mind.

Don't underestimate the benefits of working in a multidisciplinary team during an AI procurement. Consider bringing in the right skills as and when you need them and plan in advance so your processes are not slowed down because of resource constraints. Contact experts in the field and make use of cross-government networks to gather insights and share experiences.

Be aware of the potential risks of your project and actively mitigate them through your approach. One of the challenges in our project was the vast amounts of unstructured data that we worked with. We needed to adapt the data source, restructure and enrich that data with other data sources while at the same time ensuring that the AI model results are interpretable.

Make the guidelines your tool and ensure that you use them for the relevant parts of your project. Not everything in the guidelines will necessarily be relevant to your AI project. It is important to acknowledge this and focus on the areas that are of greatest relevance to you. Ethics are an important consideration for our project, but we focused more on aspects of "explainability" and transparency rather than bias concerning individuals because we are not using sensitive data.

Use the guidelines iteratively and return to them during different phases of the procurement process, using them as a way to challenge your assumptions and test your approach.

There is not much else out there on the procurement of AI-driven solutions in terms of best practices or lessons. If you would have to do this without any resources this could be a pretty daunting process. The more you can make use of the guidelines during your procurement process, the better. Using the guidelines definitely helped us and improved our approach and we strongly encourage teams across the UK Government, but also globally to use the guidelines.

BOX 1 The supplier's view

All participating potential suppliers had already heard about the UK's data ethics framework before the procurement process.

Most suppliers found it important that they were asked to describe how their approach to AI development and deployment met government digital service and ethics standards. They supported the suggestion to include ethical considerations in the proposal evaluation.

Some of the suppliers highlighted that questions considering ethical standards can be quite vague and general. "There are many different standards and their requirements are also a bit overlapping, and not every standard is as relevant to this specific project." Therefore, suppliers suggested including questions only related to the most relevant

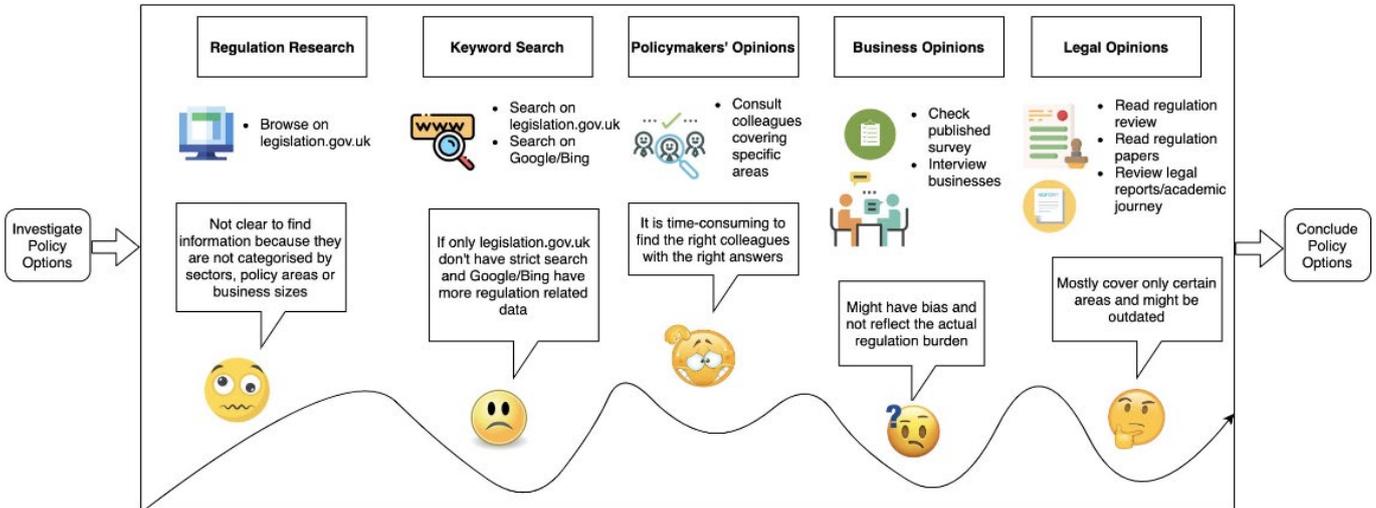
standard or more specific questions on how certain standards will be met during the AI development.

All suppliers stated that they would agree to a third-part audit of the AI system that they developed. All suppliers agreed that it is important to ensure trust in the AI-driven solution.

"Given the need to build user and public understanding and confidence in AI solutions, as well as to detect and address emerging issues with the use of the technology, an audit requirement could be a useful measure. But the detail would matter – the specification, who the auditors were, whether they (and the audit protocols) could cope with the variety of applications and technologies in use, such that it would be a meaningful process that did add value."

FIGURE 1 | User journey maps from a supplier's perspective
Presentation by Qualimental Technologies

User Journey Map - As Is



User Journey Map - To Be

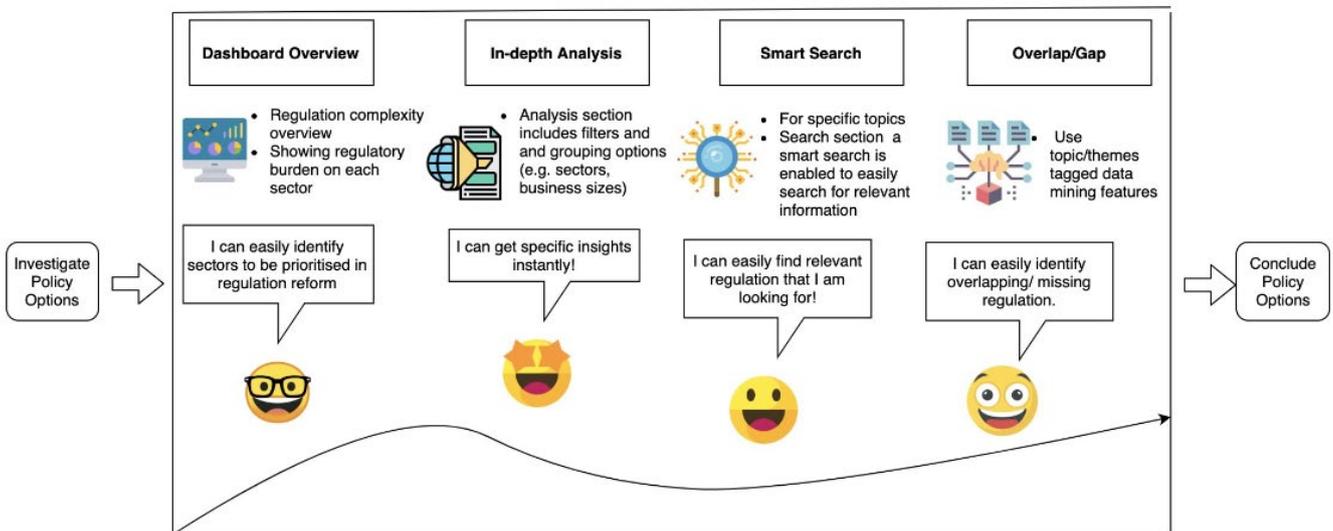


FIGURE 2 | **Explainable AI was a priority**
 Presentation by Qualimental Technologies

Explainability

User Interface



Info Icon on front-end application to show more information

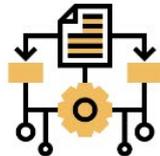


Channels for users to challenge the AI results and provide feedback

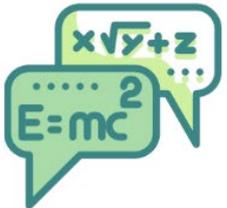
Explainable AI



Explain training datasets used & Nature of the data available in the training database



Explain type and structure of the model, the processing algorithms and extracted features



Specify the formula, variables, factors and explain function nature in simple technical language

2

Food Standards Agency



Overview

The United Kingdom's Food Standards Agency (FSA)'s aim is to be able to protect consumers now and in the future. Therefore, it is important for the organization to forecast potential risks and take action in a timely manner. As the competent authority in charge of regulating the food and feed sector in the UK, the FSA needs to be aware of risks affecting UK consumers related to safety and authenticity in the sectors. Over the last two years, the agency has extended its use of emerging technologies to identify risks.

The focus on the project that the agency most recently procured was to develop further predictive capability within the agency to mitigate against food and feed safety risks. The FSA wants to develop an overarching artificial intelligence (AI)-based system that while remaining agile and decentralized, enables additional data and intelligence sharing, the re-use of technical solutions throughout government and clearly ties into actions taken by the remainder of the FSA.

2.1 Few government teams have procured AI solutions. What proved to be the opportunities and challenges?

The FSA expects many benefits from using AI applications. AI techniques enable us to build more complex models, allow us to make better predictions and help us to better identify risks. The use of AI systems enables us to more effectively allocate resources and as a result, saves the organization costs. We have also found that we were able to reuse our AI toolbox in other use cases and share the findings as well as the tools themselves with other organizations, such as local authorities.

Traditional approaches like statistics are very powerful, but have their limits. AI-based applications are well placed to consolidate large amounts of information that are stored in multiple data sets throughout the world, to analyse this information, identify patterns and to provide us with actionable insights. AI-driven systems allow us to gather insights in minutes that otherwise would take hours, days and probably in some cases, months to compile.

For example, analysts in the FSA have been using AI-based models and real-time weather data to predict toxin contamination produced by mould, which can be prevalent in food commodities like fruits and nuts. There is a very clear correlation between the weather conditions during harvest and storage of those commodities and the levels of toxins at the point of consumption by people. This smart and data-driven approach allowed us to target inspections and controls. We used those trained models and also applied them in similar use cases.

Access to standardized data and data quality are key challenges for AI adoption for the FSA as well as the whole public sector. Currently, data is essential for any AI project. Exploring the potential of data trusts to share data and intelligence throughout organizations, which the Office for AI has done in the UK is a good start, but needs broader collaboration throughout government and industry.

The UK has also advanced significantly with regard to the open data agenda in recent years, but even when the data is open and readily accessible, there are still challenges when you operate with different data sets that aren't standardized. There is much more that the government could and should be doing to make data sets accessible and usable. For example, when it comes to the quality of the data and the interoperability of different data sets. A government-wide open data catalogue or a data dictionary with all relevant information including data quality would be a tool that could boost AI and analytics uptake considerably.

Another challenge for AI adoption in the public sector is the concern around ethical considerations. In particular, how to operationalize ethical principles, how to practically apply guidance and how to best implement "explainable AI" (meaning AI systems that allow us to interpret the outcomes of the algorithms). There is currently little centralized support in the UK public sector and every team has to develop an approach from scratch that works for them, which in turn, creates inefficiencies as well as uncertainties. So, the FSA is trying to proactively address this gap by facilitating conversations throughout government about AI ethics and linking up with institutions such as the Centre for Data Ethics and Innovation.

Last, but not least, in-house capabilities and skills are a key challenge, but also opportunity.

To raise awareness and build expertise the FSA started with identifying proof of concepts that really show the value of AI to decision-makers. The agency also focused on business engagement and partnered with experts in the field on proof of concepts and pilots. Since their first AI-driven projects the agency has also been successful in recruiting and upskilling their teams.

2.2 Why were you keen to use the guidelines for AI Procurement?

We were proactively looking to support our AI procurement efforts with specialist knowledge and best-practice throughout government departments. We found it challenging to find examples of best-practice for AI procurement within government departments. There are more experienced delivery teams in various departments, but there is not yet a centralized approach to sharing knowledge. We found that the AI Procurement Guidelines were a good repository of key considerations and they helped us to structure our approach to the procurement of the new AI-based system effectively. It was easy to see what considerations

to focus on and the guidelines served as a basis for discussion as to how this related to our project.

Not all of the issues mentioned in the guidelines were relevant to our project, but having an overview of all possible considerations and ideas on how to best learn from others who have done similar in the past was very helpful. What we found particularly useful were the insights on how to encode legislation and standards into our work. It also helped us think about public benefits and potential impacts of our work beyond the immediate business need.

2.3 How did you go about consulting the guidelines?

The guidelines were particularly useful at the specifications stage and helped us ensure that we are going to a market with a clear problem statement asking for the right things. The guidelines built on the [Tech Code of Practice](#), which we are already using to design, build and buy technology. This made it easy to integrate them in our processes. We used the guidelines to include signposting for ethical requirements, technical feasibility and capability building in our

invitation to tender. The guidelines worked for us as a checklist to ensure that we addressed the important issues that are peculiar for AI procurement and our project.

We found working with the guidelines effective. After we issued our invitation to tender, we received fewer supplier questions and more responses to our invitation to tender than the previous time we procured for AI capabilities.

2.4 How did you go about the procurement process?

We found the Digital Outcomes and Skills (DOS) framework fit for purpose for our procurement approach. The framework enables a flexible approach to IT procurement since it is set out to highlight the challenge that you want to address rather than focus on one specific technology. This framework often helps the public sector buy, design, build and deliver digital outcomes by finding appropriate specialists to deliver agile software development. The DOS framework is also often used for the procurement of digital teams or individual contributors, to work alongside in-house

delivery teams. This procurement approach aligns with the principles set out in the AI procurement guidelines since it asks for a focus on the challenge rather than the specific solution.

A key aspect of our procurement is that we tendered for a call-off contract allowing work packages to be agreed and awarded throughout the term of the agreement. This supports the agile delivery of the projects and provides us with the flexibility to react to new findings and mitigate risks with delivery timescales and project alterations.

BOX 2 | The supplier's view

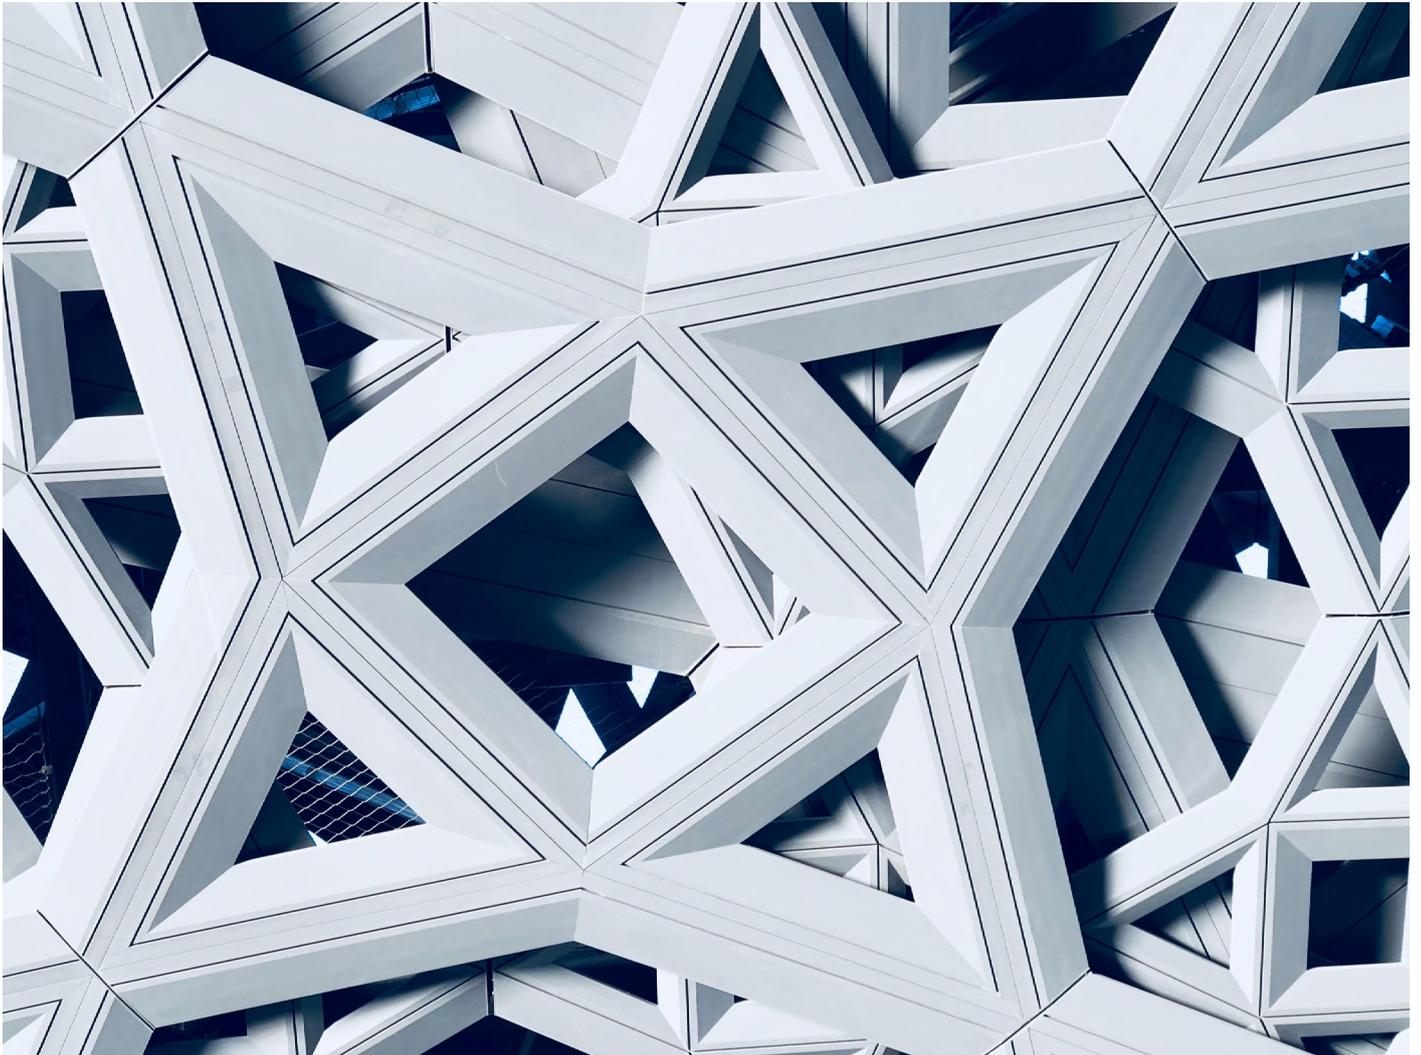
The suppliers highlighted that for someone new to this domain the UK AI procurement guidelines are a helpful and a good summary of UK Government procurement for AI.

Feedback included the importance of terms and references for non-technical staff in AI procurement teams in government as well as from the suppliers' perspective. "For non-technical experts who work in AI (e.g. procurement experts, business change consultants, leadership, Scrum Masters) it would be useful to have a short intro saying what AI is, what it isn't and how it relates to other terms such as statistics, economics, data science and machine learning. All of these definitions will be according to a UK Government definition as few of them (especially AI) have singular, agreed definitions."

Explaining why ethical considerations are key for AI procurement is also important. One supplier explained that, for example, "creating traditional predictive models involves coding a series of instructions that allows the predictive engine to appear to make decisions (like a human) based upon source data. The supplier further explains how the AI approach differs from the traditional approach: "Coding instructions is incredibly labour intensive

and as such forces the creator (developer) to think deeply about every instruction and the implications of including or excluding it. Machine learning uses algorithms that enable the machine to write these instructions automatically based purely upon source data. This saves time, but also removes a lot of "thinking time" from the predictive model creation. As such, it's easy for a machine learning algorithm to pick up biases in the data and codify them as instructions and it's hard for the creator to then spot these biases. Hence, properly executed AI projects must add this 'thinking time' back in."

Furthermore, feedback shows that more needs to be done in the future to ensure that it is easier to prove compliance with ethical frameworks in government. The communications from the government need to be clearer on the specific expectations for suppliers. Feedback was that often suppliers have ethical considerations embedded into their ongoing processes and it is difficult for them to provide evidence for standalone tasks because they see this as merely good practice. In their view this does not mean that they are not operating ethically, but that they are ethical by design and that those considerations are ingrained in everything that they do.



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 Technology 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.



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