

Draft Guidelines for AI Procurement – Office for Artificial Intelligence

Copied by [Team Defence Information](#) from <https://www.gov.uk/government/publications/draft-guidelines-for-ai-procurement/draft-guidelines-for-ai-procurement-on-22/10/2019> Go to the source for current version.

Introduction

Artificial Intelligence is a technology that has the potential to greatly improve our public services by reducing costs, enhancing quality, and freeing up valuable time of frontline staff.

Recognising this, the UK Government published the Data Ethics Framework and a Guide to using AI in the Public Sector (also referred to as the AI Guide) to enable public bodies adopt AI systems in a way that works for everyone in society.

These new procurement Guidelines will help inform and empower buyers in the public sector, helping them to evaluate suppliers, then confidently and responsibly procure AI technologies for the benefit of citizens.

The Guidelines follow the highest standards as set in the Data Ethics Framework, which outlines principles to guide the design of appropriate data use in the public sector. Teams developing AI systems in-house should refer to the AI Guide to using Artificial Intelligence in the Public Sector. All technology projects and programmes should follow the Technology Code of Practice.

Why guidelines for public procurement of AI-systems are needed

To use AI to better manage or run public services, first, you will need to consider how to procure the technology if you chose not to build the system entirely in-house. You should define your purchasing strategy in the same way as you would for any other technology, but acquiring AI systems raises various new challenges that will need to be addressed.

Due to uncertainty of the technology, the market being fast-moving and immature, and the lack of standards, it can be difficult to choose from existing procurement routes. Access to markets and suppliers as well as draft detailed contracts, that encourage innovation and mitigate risks are also difficult to find.

Failure to promote ethical and technically robust considerations, diversity and openness through AI procurement may also lead to poor procurement decisions for AI systems that limit accountability, undermine social values, entrench the market power of large businesses and decrease public trust.

Innovative procurement approaches have the potential to foster innovation, create competitive markets for AI systems and uphold public trust in public sector adoption of AI.

How to use the guidance

These Guidelines aim to provide a starting point for planning and delivering public procurement approaches concerning AI solutions and services - once you have determined that an acquisition of an AI-driven system is needed to improve a business function and/or public service. The Guidelines are not intended to be a silver-bullet for solving all public sector AI adoption challenges related to procurement,

but to provide a path towards a more effective and responsible procurement approaches for AI and data driven systems.

The Guidelines have been structured around the main commercial stages of a typical procurement lifecycle:

- Preparation and Planning
- Publication
- Selection Evaluation and Award
- Contract Implementation

Many of the important decisions take place at the Preparation and Planning stage, before publishing a contract notice.

The Guidelines assume the application of the Public Contracts Regulations 2015 and the Public Sector Equality Duty under the Equality Act 2010. For more information on these legislations go to GOV.UK.

Who this guidance is for

These Guidelines are for multidisciplinary teams involved in public procurement decisions relating to AI projects:

- Policy officials and organisation leads considering an AI-based solution and/or planning and delivering AI projects
- Procurement officials and commercial teams responsible for the planning and delivery of AI projects
- Analysts, data scientists and digital, data and technology experts who are developing project-specific requirements and evaluating, using and maintaining AI systems
- Chief Data, Information, Technology and Innovation Officers considering planning and delivering AI projects
- Suppliers who want to better understand the best practice processes, technical and ethical expectations for AI projects, and to tailor their offerings appropriately

Overview

1. Explore procurement processes that focus on the challenge rather than a specific solution

- Explore the use of innovative procurement processes to acquire AI systems
- Focus on developing a clear problem statement, rather than on detailing specifications of a solution
- Support an iterative approach to product development and reflect this accordingly in the invitation to tender

Read more about point 1

2. Define the public benefit of using AI while assessing risks

Apply Data Ethics Framework Principle 1: Start with clear user need and public benefit:

- Set out clearly in your invitation to tender why you consider AI to be relevant to the problem and be open to alternative solutions
- Explain in your invitation to tender that the public benefit is a main driver of your decision-making process when assessing proposals
- Conduct initial AI impact assessments even before you are starting the procurement process, ensure that your interim findings inform the invitation to tender and revisit the assessments at decision points

Read more about point 2

3. Include your procurement within a strategy for AI adoption

- Consider aligning your work with relevant AI or AI-related initiatives from across government
- Establish networks across your organisation and across the civil service to share insights and learn from each other

Read more about point 3

4. Incorporate references to legislation and codes of practice in the invitation to tender

- Apply Data Ethics Framework Principle 2: Be aware of relevant legislation and code of practice

Read more about point 4

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

Apply Data Ethics Framework Principle 3: Use data that is proportionate to the user need:

- Ensure having data governance mechanisms in place from the start of the procurement process
- Assess whether relevant data will be available for the project
- Define if and how you will share data for the procurement initiative and the subsequent project

Read more about point 5

6. Develop a strategy to address technical and ethical limitations of using training data

Apply Data Ethics Framework Principle 4: Understand the limitations of the data:

- Highlight known limitations of the data in your invitation to tender and require tenderers to describe their strategies on how to address these shortcomings, as well as to have a plan for addressing relevant limitations you may have missed
- Make ethical considerations part of your evaluation criteria for proposals

Read more about point 6

7. Conduct procurement with diverse multidisciplinary teams

Apply Data Ethics Framework Principle 5: Ensure robust practices and work within your skillset:

- Seek to develop ideas and make decisions throughout the procurement process in a multi-disciplinary team
- Require the successful suppliers(s) to assemble a team with the right skillset

Read more about point 7

8. Focus on mechanisms of accountability and transparency throughout procurement

Apply Data Ethics Framework Principle 6: Make your work transparent and be accountable:

- Promote a culture of accountability of AI-powered solutions through your procurement practices
Maximise transparency in AI decision-making
- Consider the need for interpretability of the algorithms and make this one of your decision criteria during the procurement process

Read more about point 8

9. Consider the life-cycle management of the AI system

Apply Data Ethics Framework Principle 7: Embed data use responsibly:

- 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
- Ensure that knowledge transfer and training is part of the engagement
- Highlight the need for insights on how to manage the appropriate use of the application by non-specialists

Read more about point 9

10. Create a level and fair playing field for suppliers

- Engage AI system suppliers early and frequently throughout the process
- Reach out in various ways to a wide variety of AI solution suppliers
- Encourage an open environment that supports competition in the AI ecosystem

Read more about point 10

Guidelines for AI procurement – read more sections

1. Explore procurement processes that focus on the challenge rather than a specific solution

Why is it important?

To acquire the AI systems which best meets the challenges you want to address and foster responsible innovation.

Explore the development and use of innovative procurement processes to acquire AI systems

Innovation-oriented procurement procedures provide opportunities to accelerate the adoption of new technologies, to promote innovation and the ethical development for AI.

Innovative public procurement processes that include practices: detailing challenging problems, organising technology contests, providing opportunities for demonstrators, and giving newly established suppliers the opportunity to compete for public sector contracts, have the potential to boost innovation and help new companies grow.

There is more than one route to market for AI services depending on which kind of challenge you aim to address and you should aim to strategically develop your procurement approach. Frameworks such as G-Cloud, the Digital Outcomes and Specialist framework and Spark DPS are useful starting points.

Examples of procurement processes that support the development and deployment of AI systems include, but are not limited to:

- Agile procurement processes that allow you to go to market at 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. Examples include the GovTech Catalyst or the Scottish Government-run CivTech[®] accelerator programme.
- Innovation Partnerships that enable the procurement of technologies that cannot be delivered by the current options available to the market
- AI Procurement Frameworks that prescribe the terms and conditions applying to any subsequent contract and allow to pre-vet suppliers against a set of predefined criteria that can include ethical requirements for example, the the as the Government of Canada AI suppliers list

When making use of innovative approaches of procuring emerging technologies you should also focus on best practices that have evidently increased the supplier base of smaller and innovative suppliers, which is important for fast developing markets. 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 suppliers

Focus on developing a clear problem statement, rather than detailed specifications for a solution

AI technologies are rapidly developing, with new technologies and products constantly making it to market. Therefore, you should focus on describing the challenges and opportunities that you are facing and leverage technology partners to decipher what technology makes the most sense for the issue at hand. Output-based requirements in the invitation to tender should this approach.

Early market engagement can help you to determine the scope and feasibility of the requirements and, in turn, the most appropriate way to design and structure the requirements to allow the optimum response. Refer to Guideline 10 for more considerations on market engagement.

Support an iterative approach to product development and reflect this accordingly in the invitation to tender

The AI Guide provides information on how to plan and prepare for AI implementation AI-powered solutions. AI systems differ significantly from other technology tools in their unique ability to learn and adapt through ongoing, periodic training with new data. Therefore, 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 the first phase of procurement, with only the winner going on to develop the full solution. This is encouraged by initiatives, for instance, Small Business Research Initiative (SBRI). Building and testing the AI systems and its application in phases facilitates informed decision-making, innovation, and transparency, providing you with relevant information to conduct meaningful impact assessments and evaluate risks.

2. Define the public benefit of using AI while assessing risks

Apply Data Ethics Framework Principle 1: Start with clear user need and public benefit

Why is this important?

The Data Ethics Framework Principle 1 states that before you start working with data, you must consider the user need and expected public benefit. This is particularly important for the acquisition and deployment of AI systems since many applications are novel and ethical issues can arise through their use in the public sector. 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 specific risks which must be identified and managed early in the procurement phase.

Set out clearly in your invitation to tender why you consider AI to be relevant to the problem, and be open to alternative solutions

In most circumstances, you should only refer to the need for an AI solution in your invitation to tender if there is strong indication that the technology will address the problem that you are trying to solve. You should identify whether you need an AI system with an analysis of policy challenges and alternatives, and by comparing against other potential courses of action when the AI project does not have a clear research and innovation focus. The AI Guide provides guidance on understanding AI and how AI can help. If during the evaluation of the tender responses it becomes evident that another solution that does not incorporate AI is better able to address the problem, you should consider following this alternative delivery path.

Before starting the procurement process, assess whether AI could be part of a solution to your problem. If you lack the capabilities in your team to carry out this assessment, you should consider seeking these from elsewhere in your organisation or relevant professional network (for example, academia, trusted vendors). You can also consider consulting and possibly collaborating with appropriate stakeholders. For this assessment, it is important to consider engaging a multi-stakeholder community to define and test a clear policy problem statement and reflect the findings in the invitation to tender.

Explain in your invitation to tender that the public benefit is a main driver of your decision-making process when assessing proposals

When setting out the requirements in the invitation to tender, you should explicitly refer to the public benefit. For example, specify success and failure criteria in the context of public benefit: what do you expect a system to achieve and be capable of, and what are the types of failure and harm that must be avoided. The Data Ethics Framework describes this in more detail in Principle 1.

Conduct initial AI impact assessments at the start of the procurement process, ensure that your interim findings inform the invitation to tender and revisit the assessments at decision points

To improve your understanding of the potential impacts of the use of AI and to mitigate the risks, it is important that you start an assessment in a systematic manner before you acquire an AI system, and make sure that the findings also inform your commercial strategy. There will be different considerations depending 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 complete assessment.

An initial assessment should outline user needs and affected communities, in addition to potential issues such as inaccuracy and bias of the AI system. It should also consider scenarios of 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 data sets. Associated risks and their respective mitigation strategies must be recognised by a suitable risk owner with seniority to have decision-making power, and should include a go/no go decision.

You should also consider asking the potential suppliers to identify risks, and how they would mitigate them. This can give you valuable information regarding each tenderer's attitude towards, and awareness of potential risks. Where you identified significant risks in your initial assessments, you should require tenderers to specify out how they would address them.

Data protection impact assessments and equality impact assessments can provide a useful starting point for assessing potential unintended consequences. For 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.

3. Include your procurement within a strategy for AI adoption

Why is it important?

To take advantage of economies of scale and 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.

Consider aligning your work with the work of teams across government leading on relevant AI or AI-related initiatives

Linking up with teams across government will help you to make use of expertise and has the potential to make your procurement effort more effective and more likely to be successful. It might allow 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 strategic AI initiatives is that there may be special support for initiatives that align to the strategy, including access to additional experts.

Examples for teams and organisations that might be useful to consult are the Office for AI, the Government Digital Service, the Centre for Data Ethics and Innovation or teams and organisations with specific domain knowledge. You can also find examples of AI use cases in the AI Guide.

Establish networks across your organisation and across the civil service to share insights and learn from each other

To improve your practices you should actively seek out collaboration across departments and fields of expertise. You should also share knowledge and feedback via expert communities the {Digital Buying Community}(https://www.gov.uk/service-manual/communities/digital-buying-community), the Data Science Community of Interest or other similar networks in your region.

Within your organisation 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.

4. Incorporate references to legislation and codes of practice in the invitation to tender

Apply Data Ethics Framework Principle 2: Be aware of relevant legislation and code of practice

Why is this important?

To comply with existing laws and regulations, to support the standardisation of norms, and enables better policy making in a dynamic innovation technology space.

It is essential that compliance with all relevant legislation, codes of practice, regulation, and guidance related to data use and AI is one of the tender requirements.

Relevant legislation and codes of practice are but are not limited to:

Some relevant legislation and codes of practice are:

- Data Ethics Framework
- Technology Code of Practice
- Code of conduct for data-driven health and care technology
- Government Service Design Manual

- Digital Service Standard Guidelines
- The Resources for Chief Technology Officers
- Data Protection law and ICO guidance
- The Digital Economy Act
- Freedom of Information Act

The Data Ethics Framework sets out further important legislation that typically applies to using data. For more information on this topic refer also to the section in the guide on using artificial intelligence in the public sector that deals with ensuring your use of AI is compliant with data protection laws.

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

Apply Data Ethics Framework Principle 3: Use data that is proportionate to the user need

Why is this important?

Data is the basis of the majority of AI powered solutions. 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). The recognition that data is a sensitive asset that must be handled with care, especially when it is personal data, should be considered integral to the solution.

The Data Ethics Framework Principle 3 sets out considerations that need to be considered to ensure that the proposed data collection, storage or analysis is proportionate. Your invitation to tender should provide details on these proportionate data requirements.

Ensure having data governance mechanisms in place from the start of the procurement process

You should aim to set out a data governance approach from the start of the procurement process and reflect this in the invitation to tender where relevant. Governance needs to cover all data activities related to the proposed project, for example, granting data access to project members, storing data in other locations for analysis, and reviewing data consent (purposes we are authorised to use the data for).

The AI Guide also highlights guidance for governance when running your AI project. A good example of being transparent about what data is being shared and for what purpose is with the Public Service Delivery powers within Part 5 of the Digital Economy Act which require all data shared using this legal gateway to be published in an open register.

Assess whether relevant data will be available for the project

You should determine, at a high level, data availability before starting your procurement process. This entails gaining an initial awareness of data sources needed to ensure a successful project.

In cases where data is not available for the use case in mind, you may be able to find data through third parties; vendors, partners, or data brokers. If data is not available through any channel, engage data

scientists (for example, through vendors or the Data Science Community of Interest) to assess whether the use case can be addressed at all in a data driven manner.

This should help you to ensure that you have access to relevant data to address your problem and consent to use it. If this is not the case, you can establish the data requirements and ask for collection and collation of the relevant data as part of the invitation to tender.

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 releasing data to suppliers in line with Data Protection law during the procurement process. This can allow suppliers to craft a better response to the invitation to tender that is tailored to your needs, with assumptions, timelines, and fees that match your situation as closely as possible. This can improve the quality of invitation to tender responses you receive and increase the success of the procurement project.

Depending on the type of data, this could be done by opening the data to the public or by releasing a data 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 data set. Otherwise vendors might make erroneous assumptions that can impact the quality of bids and subsequently of the project.

Create and document the appropriate data sharing conditions aligned with data protection requirements such as:

- Minimum requirements for the environment where the vendor will host the data (for example, enterprise laptop that meets the vendor's standards for their sensitive data)
- Data consent form signed by the vendor's lead for the pursuit (for example, stating that the data will be used exclusively for the pursuit and for no other purpose)
- Date for data deletion (for example, immediately upon submission of the vendor's invitation to tender response)
- Confirmation of deletion (for example, written confirmation of deletion signed and submitted by the vendor's lead for the pursuit)

You can also consider the use of anonymisation techniques to help safeguard data privacy, including data aggregation, masking, and synthetic data. Keep in mind however that you must manage anonymised data as carefully as the original data, since it may inadvertently expose important insights. Invitation to tenders should encourage innovative technological approaches that make less intrusive use of data or which achieve the same or similar outcomes with less sensitive datasets.

Certain vendors may have data that is highly complementary to the initiative, and it is in your best interest to consider using these data. It is thus important to have a framework to decide under what circumstances to accept data from a vendor. Decision criteria could include:

- Vendor: some vendors could be pre-qualified as accepted data suppliers, be more trust-worthy from years of working together, or have a strong reputation related to their data assets

- Domain: some domains; health, justice, and immigration, are very sensitive. Use of third-party data in these domains should require careful thought before going through

- Data precedence and integrity: before using any third-party data, you 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 commercialise it

6. Develop a strategy to address technical and ethical limitations of using training data

Apply Data Ethics Framework Principle 4: Understand the limitations of the data

Why is this important?

Though available, legal to use, and proportionate to need, there may be limitations to data (for example, data bias) that make an AI system unreliable or misleading and therefore unfit for purpose. This is highlighted by Data Ethics Framework Principle 4 and the Guidance for understanding AI ethics and safety.

Highlight known limitations of the data in your invitation to tender and require tenders to describe their strategies on how to address these shortcomings, as well as to have a plan for addressing relevant limitations you may have missed.

Principle 4 in the Data Ethics Framework specifically points out what to consider when deciding if a source of data is suitable, including:

- Representativeness (whether the data accurately represents the segment of the population in scope for the AI solution)¹
- Provenance (for example, how and why the data was collected)
- Gaps in data quality (for example, many values missing from a data element)
- Errors in the data
- Bias present in the data (data is not representative of the population the algorithm will be applied to)
- Lack of clarity in metadata (for example, confusing or vague data element names)

Be aware of the importance of data quality for AI and, when possible, check completeness, diversity, and accuracy of the data before starting the procurement process. Articulating your data quality observations and the apparent limitations of the data potentially used are crucial. The supplier of the AI system must be aware of these considerations, if not during the procurement, then at the latest after the contract has been awarded and design of the system has begun.

If you do not have the right skills or means to comprehensively check for possible limitations of your data, provide suppliers with guiding insights into the state of the data and its origin, so that they can draft adequate proposals. Also, ensure the data requirements include data quality assessment, and, if required, development of mitigation strategies for low quality data.

If a thorough assessment of the data has not been made, make it a requirement in the invitation to tender that work should be preceded by a comprehensive check of the data the AI system will use to train.

Make ethical considerations part of your evaluation criteria for proposals

There are key robust practices you can ask for suppliers to demonstrate when providing AI solutions. The Guidance for understanding AI ethics and safety provides a useful framework to identify those. Besides having an ethical framework within their company, robust practices include:

- Having an internal AI ethics approach, with examples of how it has been applied to design, develop, and deploy AI solutions
- Processes to ensure accountability over outputs of algorithms
- Avoiding outputs of analysis which could result in unfair decision making
- Designing for reproducibility
- Testing the model under a range of conditions
- Defining acceptable model performance

7. Conduct procurement with diverse multidisciplinary teams

Apply Data Ethics Framework Principle 5: Ensure robust practices and work within your skillset

Why is this important?

Developing, evaluating and delivering AI invitation-to-tenders will be more effective with diverse teams that understand the interdependent disciplines AI covers, including: domain expertise (for example, healthcare, transportation), systems and data engineering, model development (for example, deep learning), and visualisation/information design, among others.

Principle 5 for the Data Ethics Framework highlights the need for robust and consistent practices to make best use of data. This is particularly true for AI applications because actions or decisions may follow from that data. Therefore, it is key to address them in the invitation to tender by making decision criteria an explicit requirement.

Seek to make procurement decisions using multidisciplinary teams

Develop an understanding of the skills needed to effectively acquire and maintain an AI-powered solution before the start of the procurement process.

Consider assembling multidisciplinary teams that specialise in designing, procuring, and operationalising AI systems. These multidisciplinary teams should include expertise in: policy from the domain where the AI solution will be applied (for example, justice), machine learning/data science, data engineering, technology (software and hardware), procurement, data ethics, and human rights².

You should also be aware of the importance of diverse teams for AI development and deployment. Consider different genders, ethnicities, socioeconomic backgrounds, disabilities, and sexualities. Encouraging a mix of perspectives helps ensure solutions to problems are inclusive, and can mitigate the

potential for bias in data and AI systems. The Public Sector Equality duty supports this by encouraging diverse recruitment.

If your team lacks expertise, you could reach out to professional networks within your organisation, or across the public sector, to gather important feedback at all stages of procurement.

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

Require the successful supplier(s) to assemble a team with the right skillset

As part of the invitation to tender, require suppliers to consider the skills, qualifications and diversity of the team that will develop and deploy the AI system. This should be one of your decision criteria for evaluating proposals.

8. Focus on mechanisms of accountability and transparency throughout procurement

Apply Data Ethics Framework Principle 6: Make your work transparent and be accountable

Why is this important?

To build public trust in the legitimacy of AI systems, procurement should be transparent, explain how the proposed solution works, and define accountability for its deployment, allowing the solution to be evaluated independently. You should make it a requirement for the supplier to be transparent about:

- The available toolkit, including the list of software tools the supplier proposes to use
- The origin and nature of any data the supplier plans on bringing to the project
- Data used to train algorithms the supplier will bring to the project
- The algorithms used

In addition, the service supplier should have transparent processes and practices, including documenting the algorithm-building process, and being able to clearly explain how data for algorithm training is selected. You should also ensure that you meet the Open Contracting Data Standard where possible, and openly publish documents at all stages of the contracting process to ensure transparency.

Promote a culture of responsibility for AI-powered solutions

Public institutions will run into difficulties relying on black-box algorithms to make decisions that could impact citizens in a significant way, especially within the context of growing awareness about algorithmic bias and potential discriminatory effects (areas the Centre for Data Ethics and Innovation is investigating). You will need to apply different considerations depending on your use case and how you plan to apply AI. You should work with the supplier to explain the AI application to promote transparency, and ensuring your approach can withstand scrutiny. You should link to the impact assessment described in Guideline 2.

In some cases, you could consider a requirement for suppliers to allow independent audit(s) of their solutions. This can help prevent or mitigate unintended outcomes. Part of this would be to consider how

applicable accountability requirements in law, for instance, freedom of information legislation and data protection requirements, will be implemented throughout the project lifecycle.

Due to the nature of AI technologies, liability for certain areas will reside with the government department, particularly around the application of the AI solution as well as data access and transfer. However, liability may also need to sit with the supplier, including technical and quality assurance.

Maximise transparency in AI decision-making

Encourage transparency of AI decision making (i.e. the decisions taken and/or insights generated by AI) through the requirements in the tender. One way to do this is to encourage explainable AI. You can also make it a requirement for suppliers to include training and knowledge transfer to your team. Finally, you can ask for documentation that provides information about the algorithms used (for example, the data used for training algorithms, whether the model is based on supervised, unsupervised, or reinforcement learning, and any limitations).

You should ask for documentation on the proposed solution. This is especially important when the algorithm is a pre-packaged solution, as opposed to an algorithm that will be built and/or customised as part of the project. You can also ask suppliers to provide information on their model building methodology, including how they select variables, build samples (where applicable), and validate the model. You should be aware that algorithm building is an iterative process, and that it depends on creativity as much as it does on science.

You should be aware that this documentation will give you an indication of a supplier's general approach. You should not expect a fully detailed plan, as this will invariably change as the project progresses.

Consider the explainability of algorithms and make this one of your decision criteria

In some domains, it is less important to know exactly how a machine learning model has arrived at a result if one can demonstrate logical steps to achieving the outcome. In other words, the ability to know how and why a model performed the way it did is a more appropriate means of evaluating transparency in the context of AI. 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, with respect to intrinsic attributes (i.e. considerations of fairness and mitigation of bias). This should be included in the documentation supplied by your supplier.

It is also important to consider the potential for tension between explainability and accuracy when acquiring AI solutions. Classic statistical techniques, for example, decision-tree models are easier to explain but might have less predictive power, whereas more complex models, for example, neural networks, have high predictive power but are black boxes due to their complexity. However, in various instances, particularly from a policy perspective it is extremely valuable to understand how decisions are made by algorithms, for example, in the case of predictive analytics. You should also consider the need for transparency for decision making systems where citizens might expect transparency.

9. Consider the lifecycle management of the AI system

Apply Data Ethics Framework Principle 7: Embed data use responsibly

Why is this important?

The functionality and consequences of AI systems may not be apparent in the procurement process and often only become evident during deployment, requiring extended communication and information sharing between the buyer and supplier.

For AI-powered solutions in the public sector, implementation plans, sustainable and ongoing evaluation methods, and mechanisms to feed back into the data model are crucial to ensure ethical use. For more information, refer to the Guidance for understanding AI ethics and safety. You should make clear in your invitation to tender that such considerations by the supplier count and will be discussed during procurement.

Consider during AI procurement that lifespan testing, not a one-time decision is required

An AI solution may need support throughout its lifecycle. Knowing where to go for support is vital. 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 AI solution.

Consider implementing process-based governance frameworks as suggested in the Guidance for understanding AI ethics and safety. This provides basis to integrate norms, values, and principles informing procedures and protocols that define the project workflow.

Testing the model on an ongoing basis is necessary to maintain its accuracy. An inaccurate model can result in erroneous decisions that negatively impact citizens. Therefore, you should establish with the supplier how the efficacy of the model will be monitored once deployed. The National Cyber Security Centre guidance for assessing intelligent tools for cyber security also highlights the importance of these considerations.

Ensure that knowledge transfer and training is part of the engagement

Consider making knowledge transfer a requirement in the invitation to tender. 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 finalised.

Set out clearly your expectations for project documentation.

Highlight the need for insights into use of the application by non-specialists

Operational or service staff must have enough knowledge of or training on the solution to understand how to use it and act on its outputs. You should address the need for staff training and support to avoid the misuse of AI applications with the AI supplier. The application must have an easy way to report any suspected unauthorised behaviour to relevant authorities within or outside the organisation.

Enable end-to-end auditability by implementing process logs that gather the data across the modelling, training, testing, verifying, and implementation phases of the project lifecycle. Such a log should allow for the variable accessibility and presentation of information with different users in mind to achieve interpretable and justifiable AI.

10. Create a level and fair playing field for suppliers

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.

Engage AI system supplier early and frequently throughout the process

Early supplier engagement can help to determine the scope and feasibility of the requirements and, in turn, the most appropriate way to design and structure the requirements, increasing the likelihood that the winning supplier will meet the needs of the public entity at a competitive cost.

Ways to engage vendors early include:

- Having vendors provide inputs on possible evaluation criteria for the invitation to tender
- Hosting vendors to walk them through the invitation to tender

Early market engagement helps to identify new potential suppliers and/or solutions, leverage the playing field, build capacity in the market, and makes you aware of new developments.

To mitigate any risks that could be associated with market engagement (for example, commercial confidentiality, protection of IP, fettering discretion of tender process) make sure to broadly advertise the engagement opportunity, allowing all interested parties to participate, ensure that there is adequate time for responses and reasonable time for supplier selection, and where appropriate that responses can be marked as confidential.

Reach out in various ways to a wide variety of AI solution suppliers

Due to the rapidly developing landscape of AI service suppliers, in great part filled by start-ups, it is helpful to consider non-traditional ways of market engagement to attract AI solution suppliers. For example, you can consider in-person presentations to explain the needs that lead to the proposal, attend information sessions at co-working spaces, use webinars or social media.

Consider reaching out to non-traditional stakeholders, research institutes and academia. In some cases, these might have the right skills to be part of an AI implementation, and in all cases, they can act as advisory partners. Keep in mind that successfully designing and deploying AI in organisations as big and complex as public agencies requires much more than technical expertise. It requires experience in change management, familiarity with public organisations, and the ability to manage complex projects.

Encourage an open environment that supports competition in the AI ecosystem

Consider strategies to avoid vendor-lock in, particularly in relation to black-box algorithms. These practices could involve leveraging 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 (for example, re-trained through your data). It is therefore useful to consider whether your organisation should own that IP. The arrangements should be mutually beneficial and fair.

1. For more information on fairness during data selection, refer to: Leslie, D. (2019). Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector. The Alan Turing Institute. 

2. 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. (2018). 