

Scoping, risk and commissioning

Guidance on scoping and commissioning the use of artificial intelligence (AI), and the risks involved

First published 17 June 2025

11 mins read

Scoping

The project's scope should be grounded in a clear and evidenced problem statement. This is an explanation of why an artificial intelligence (AI) tool or system is needed, as well as the desired outcomes and benefits. This should be supported by user research, including the views and needs of the officers and staff who will use the solution, from business analytical expertise in your force.

Decision-makers in your force will be interested in the cost of the project and whether it will represent value for money in the long run. A full projection will include not just the cost of building or buying a solution, but also the following costs:

- any upgrades to data infrastructure – for example, computational resources needed to train the tool or system
- any uplift in the workforce skills needed to support the technology – in order to use it appropriately, officers will need to be equipped with the knowledge of how the tool works, its risks and limitations
- ongoing hosting and transferring data in and out of your systems, including costs associated with cloud services
- testing and evaluation, especially if procuring third-party expertise
- any ongoing post-deployment maintenance
- the human-in-the-loop who will make use of what the AI produces or advises
- actions required to satisfy the increased or accelerated influx of information. For example, if you are introducing a risk assessment tool that accurately and quickly scores the risk of cases, do you have the capacity to act on the increased pace and volume of assessments
- external funding sources that could offset costs to the force
- the cost of not using the tool and continuing with business-as-usual

Early advice and support

Many of the issues associated with AI can be navigated if they are identified early enough. If not, issues around data protection, compliance, hidden costs and data sufficiency can delay or even derail your project.

You should engage with the following sources of expertise in your force at the outset:

Legal

Identify the areas of law that AI is likely to engage, including intellectual property, contract, equality and public law.

Data and information management

Seek advice on data protection and set up data protection officers to carry out the data protection impact assessment. For more information on data protection considerations go to the data section of the National Police Chiefs' Council (NPCC) [AI Playbook](#).

Information technology

Establish:

- the quality and quantity of data you will need to support the proposed AI tool or system, as well as the steps required to make sure these are in place
- the capabilities that will be needed in terms of skills, processes and systems, as well as any gaps that will need to be bridged (and how this can be done)
- which systems the solution will need to interoperate with and whether there are interdependent systems, such as Niche or Athena, that will need to be notified in advance for trialling and implementing the tool or system

Information and cyber security

Understand the information and cybersecurity implications and requirements of the type of tool or system being considered. Initiate a risk assessment. Go to the [Risk section](#) of this guidance for more information

Procurement

Seek advice on potential routes to market and procurement strategies.

Analytical

Build evaluation into the design of your project. From the outset, think about what the key performance indicators are likely to be. Ask analysts about how to measure these indicators and how to establish the baseline for comparison (for example, business-as-usual). If it is not something you routinely measure, establish how you can start measuring it now.

Stakeholder engagement, culture and inclusion

Map and identify groups who may be disproportionately affected by the AI, in order to inform the equality impact assessment and stakeholder engagement strategy.

Stakeholders include any other criminal justice partners or wider public sector partners who will receive or rely on AI-enabled outputs. For example, you may be implementing AI-facilitated generation of documents, such as witness statements or safeguarding reports. In this instance, ensure that partners in Crown Prosecution Service (CPS) and social services are content to receive AI-assisted versions of these documents and that they align with the standards and requirements of those partners.

Risk

There are risks associated with the deployment of all technologies in operational policing, but some risks are heightened by and specific to AI.

Risks, limitations and vulnerabilities associated with AI generally include:

- bias, as algorithms can reproduce and amplify biases in underlying data
- inaccuracy, as AI produces outputs that it considers probable or plausible, rather than what is necessarily true
- enhanced risk of cyberattack
- enhanced risk of unlawful disclosure of sensitive force data, especially with large language models (LLMs)

The lack of transparency in the operation of AI tools and systems makes it difficult to assess and test their security. Further, the scalability and autonomy of AI tools and systems means that any harm can be more far-reaching and less visible. There are also reputational risks, as the public's views on the acceptability of AI in different contexts is under-researched.

Specific uses of AI also each carry their own risks and vulnerabilities, for example:

- a glitch associated with LLM is 'hallucination', where the model makes something up when it doesn't know the answer
- a predictive tool might experience 'concept drift', where the relationship between your input and what you want to predict changes
- a predictive tool might also experience 'data drift', where the underlying data distribution changes, which changes the relationship between the input and the output and reduces accuracy

AI can also expose weaknesses in other aspects of your organisation, such as:

- poor data quality
- weak access controls to sensitive data
- poor data literacy across the wider workforce

You should ask your information security team (IST) to research the proposed type of tool or system in detail, including its data architecture, how it works, its risks and limitations, and cybersecurity requirements. This knowledge should then form the basis of an IST-led risk assessment that identifies and quantifies the risks associated with the project. The Police Digital Service (PDS) is developing a standardised risk assessment for AI in policing. In the meantime, forces may find the following resources a helpful basis for understanding the risks associated with AI and suitable actions to mitigate them:

- [Open Worldwide Application Security Project \(OWASP\) top 10 for large language model applications](#) sets out risks associated with LLMs and strategies to mitigate these
- [OWASP top 10 for machine learning security](#)
- [National Institute of Standards and Technology \(NIST\) AI risk management framework \(RMF\)](#)
- [NIST AI risk management framework: Generative AI profile \(NIST AI 600-1\)](#).
- [National Cyber Security Centre \(NCSC\) Guidelines for secure AI system development](#)
- [The NCSC cloud security principles](#), in instances where AI is to be hosted by a third party
- [PDS artificial intelligence and LLM \(large language models\) standard](#)

To ensure a proportionate approach to risk, you could consider quantifying and scoring risks. This can form a basis for establishing consistent thresholds and policies around what risks you will accept, monitor, report, seek external expertise or assurance for, or avoid altogether. The [INTERPOL AI risk assessment](#) includes a suggested approach that broadly aligns with the data protection legislation. Each risk is scored according to the level of harm that would arise and the likelihood of it occurring.

Sources of support

For funding and advice:

- NPCC Science and Innovation Coordination Committee
- NPCC AI Board
- Home Office
- Office of the Police Chief Scientific Adviser (OPCSA)

For advice on evaluation design and whether tools are analytically sound:

- College of Policing

For advice on governance:

- Association of Police and Crime Commissioners

For advice on data quality, quantity and management:

- Centre for Data and Analytics in Policing
- PDS

For advice on information security:

- PDS

For advice on information security:

- PDS

Commissioning

Due diligence to avoid duplication

When deciding what new tool or system to commission, forces should try and build on what has been tried and tested elsewhere. Some customisation is likely to be necessary with whatever tool or system you use. However, adapting something that is already tried and tested will be likely to:

- reduce the amount your force will have to spend on testing it
- prevent the proliferation of inconsistent, non-interoperable technologies across the service

As the [Data-driven technologies APP](#) your starting point for due diligence should be a landscape review produced in partnership with a national body. In practice, that could involve your force's AI lead reaching out to one of the following:

- NPCC AI Co-ordinator
- OPCSA
- College of Policing
- PDS
- Home Office

These national bodies will be able to advise on any instances of your intended AI use case they are aware of in policing, as well as the extent and results of any testing behind them.

You can also consult your regional innovation lead for what is happening in your region and any other technology-based or innovation-based forums that your force attends.

Your due diligence should confirm whether there is a tool or system that is:

1. Proven to be both effective and scalable. This applies to:

- products on Blue Light Commercial frameworks
- G-Cloud
- PDS-supported programmes
- programmes of work signed off by the NPCC (such as Digital Public Contact)

2. Proven to be effective and is on a national partner-backed programme to establish scalability (for example, a programme comparable to the Home Office's robotic process automation programme).

3. Considered promising by a national partner, such as the College of Policing, which is funding replication and testing.

4. Being implemented and tested by one or more forces, with promising results shown in testing.

For the earlier items on the list above, the tool or system will require more testing and due diligence. Sometimes, there may be a compelling reason not to choose the most-tested option, or even to explore something untested or wholly new. These reasons could include the following:

- You have trialled the products that have been tested and found promising elsewhere, but they do not work in your force.
- The products found promising elsewhere are not compatible with your force's systems, data and cyber security arrangements, or budget. You are satisfied that this reveals a genuine gap in the market, rather than your force's unreadiness to support the technology, or AI being an inappropriate solution for your problem.
- There are no tools that are relevant to the problem, or no tools that testing have been found to be effective. However, you are still confident that the problem you are trying to solve is amenable to AI.
- You are concerned that the tools or systems in use elsewhere could be superseded by a more innovative solution.

Options for building AI-enabled tools and systems

If your force is looking to commission something new, there are broadly three paths that forces take:

Off-the-shelf

An off-the-shelf approach is where forces buy the whole tool or system (for example, an algorithm or model) from a supplier. This may be appealing where a force's own data engineering capability is limited. Be aware that the more ready-made the tool or system is, the less scope you will have to:

- tailor it to your force's needs
- train it on your own data
- modify it to reduce risks such as bias, if they arise

Co-build with a supplier

Rather than buy a product from a supplier, you could buy the supplier's time to develop a solution together. This will provide more opportunities for those designing the tool to work with the officers and staff who will use it and to design something bespoke to their needs. It is also likely to be easier for the force to modify the tool or system if it is not working as intended, and for the force to retain intellectual property in the final product. It could, however, take more time and money than buying something off-the-shelf.

Self-build

Forces may be able to build their own tools and systems, depending on the complexity of the task and the depth of their own data engineering capability. For example, many forces would have the data science capability to customise a generative pre-trained transformer (GPT) to perform a specific function internally within the force. Some forces have built their own, more complex, AI-assisted systems, bringing in third-party expertise where needed, such as AI engineers.

The skills required for self-building are data architecture, data engineering, data science and user research. If your force does not have these skills, you could partner with other forces in the region. A major advantage of self-building is that the force will be better able to explain how the tool or system works and address any issues with it how it operates than with a third-party proprietary model. This option can, however, be more complex and time-consuming. The resulting product may also be harder to scale and integrate into wider systems.

Tags

Artificial intelligence