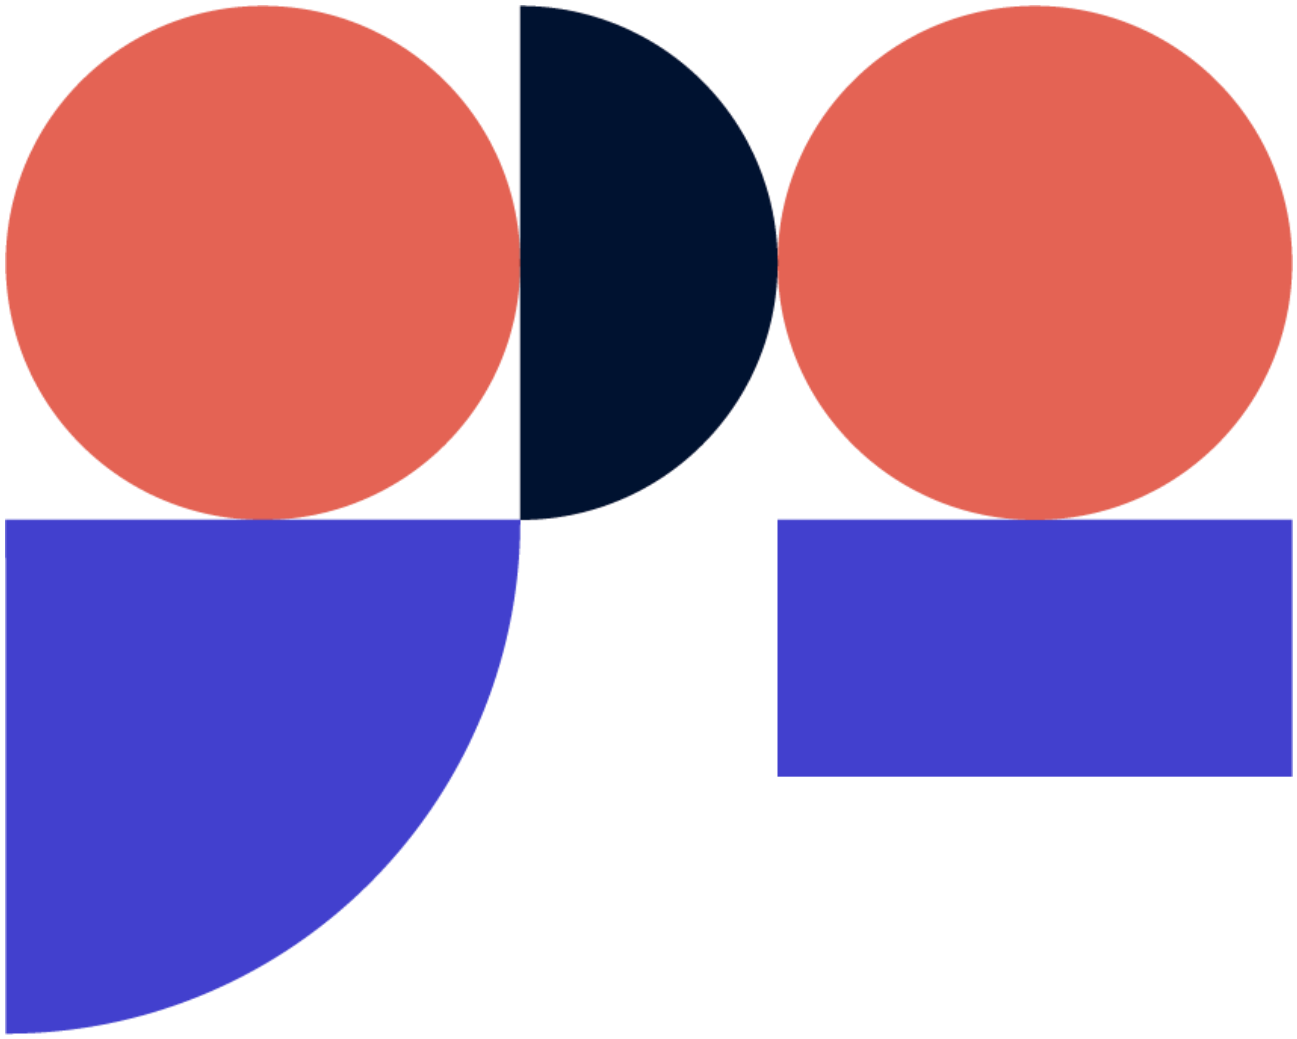

AI use guide for civil servants

oxford insights

Table of contents

Introduction

Artificial Intelligence (AI) has enormous potential to transform how governments operate. AI can help governments:

- Make sense of their data, leading to more informed policies that better respond to citizens' needs.
- Improve service delivery, helping citizens access the right services faster and more easily.
- Streamline operations, cutting bureaucracy and allowing civil servants to focus on the tasks that matter most.

The impact of AI across these three areas is real. For instance, the Alan Turing Institute estimates that AI could automate up to 84% of repetitive public service transactions in the United Kingdom, saving the equivalent of 1,200 person-years of work each year. This is time that can be used to deliver better policies and services that improve outcomes for citizens.

Uzbekistan is committed to leading this change, and harnessing the benefits of AI in the public sector. Uzbekistan has set an ambitious vision to become a leader in AI and is making bold investments to reach its goal, building AI infrastructure, promoting AI uptake across sectors, and strengthening AI skills and literacy across the general population and the workforce.

Civil servants play a key role in supporting Uzbekistan's vision. As a civil servant, you are both a recipient and a driver of this transformation. As Uzbekistan integrates AI in the public sector, you will increasingly encounter it in the tools you use every day and in your projects and services. At the same time, you have the opportunity to drive this change by learning to use AI to improve your work, and spot opportunities to improve your rojects, and services.

This guide is designed to help you understand what AI is, how it can be applied in government, and how to use it in your role, giving you the knowledge and tools to use it responsibly, ethically, and effectively.

Who is this guide for?

This guide is designed as a practical resource for all public officers, with a particular focus on non-technical staff.

What does it cover?

This guide is divided into two parts.

The first part focuses on **learning about AI**, providing essential background knowledge to be able to understand and talk about AI. It covers:

- Understanding AI.
- AI in Uzbekistan.
- AI in the public sector.

The second part provides **practical guidance** on how to use AI in your role, covering:

- How to use generative AI in your job.
- Implementing AI in your project or service.

Each section also has additional guidance boxes, with suggestions for relevant readings and resources to gain a more in-depth understanding of each topic.

What will I learn?

By the end of this guide, you will have met 6 learning objectives:

1. You know what AI is and what its limitations are.
2. You are familiar with key policies and legislation in Uzbekistan and know how to consult them.
3. You know how to build the skills and expertise needed to implement and use AI.
4. You know how to use AI responsibly, ethically, and securely in your job.
5. You know how to assess whether AI is the right tool to support your project or service.
6. You understand the AI life cycle and know how to start an AI project in your organisation.

Understanding AI

What is AI

The UN defines Artificial Intelligence (AI) as 'the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings'. Often, this means systems with the intellectual processes characteristic of humans, such as the ability to reason, discover meaning, generalise, or learn from past experience'.

You have probably come into contact with AI. In particular, generative AI systems such as ChatGPT have become very popular among the general public since 2022. However, long before this breakthrough of AI into popular use, researchers attempted to build systems that would be able to perform tasks for which would usually need human intelligence. This section provides a brief overview of AI, including key terms, a short history, and the main types of AI systems in use today.

The birth of AI

The idea of machines that can think like humans has been part of human's imagination for centuries. However, it was not until the 1950s that this idea started to take real shape. This period is often seen as the birth of Artificial Intelligence (AI) as we know it today.

Two key figures played a central role in the early development of AI: Alan Turing and John McCarthy.

Alan Turing, an English mathematician and computer scientist, is often called the founder of modern computing and the father of AI. After making major contributions to computer science, in 1950 he published a famous paper called "Computing Machinery and Intelligence". In it, he described what later became known as the Turing Test, a way to test whether a machine could demonstrate human-like intelligence.

The term 'artificial intelligence' was introduced a few years later by **John McCarthy** in 1955. McCarthy, a professor at Dartmouth College in the United States, organised a summer workshop to explore ideas about 'thinking machines'. This event is considered the official start of AI as a research field.

After the Dartmouth workshop, interest in AI grew quickly. Researchers began developing systems that could perform tasks seen as intelligent, such as playing games like chess or checkers. By the mid-1960s, AI research was being funded by governments, and dedicated AI laboratories were being set up around the world.

Rule-based AI

From the 1950s to the 1980s, most AI research focused on rule-based systems, also called **symbolic AI**. These systems are programmed with explicit rules to perform tasks, rather than learning from data. Rule-based AI works using a series of "if-then" statements: if a condition is true, then the system takes a specific action. For example, a rule might be: '*if* the customer asks about recent transactions, *then* retrieve and display the last five transactions'. These systems can then combine many rules, making more complex decisions.

However, rule-based systems also have clear limitations. These systems cannot learn on their own, and require humans to program every rule. Therefore, these systems struggle with complex situations, where it is impossible to program rules for every possible scenario. Despite these limitations, rule-based AI is still in use today, particularly for tasks with small amounts of data and simple rules, such as robotic process automation.

By the 1980s, rule-based AI was reaching its limits, as the computing resources necessary to power AI were still low, and many systems failed to deliver the results promised during the previous years of AI excitement. These limitations led to a period of reduced funding and interest in AI research, known as the first AI winter. Things began to change only in the 1990s, with the rise of machine learning. Although machine learning had existed since the 1950s, it had been limited by the lack of large datasets and computing power. As technology advanced and more data became available, machine learning approaches became more popular. Instead of relying on rules, AI began to learn from examples.

Machine Learning, neural networks, and deep learning

Machine Learning (ML) is the basis of modern AI systems. Machine learning involves training a statistical model by processing a large amount of data and generalising to unseen data. In plain language, this means teaching a computer to recognise patterns by analysing a large amount of data ('training'), to then become able to make predictions about new information it has not seen before.

After a training process, an ML model is able to provide outputs based on inputs it has never encountered. For instance, a ML model trained on thousands of photos of cats and dogs learns to recognise the characteristics (patterns) that distinguish them. Once trained, the AI model can identify whether a new photo it has never seen before shows a cat or a dog. We will talk more about the training process in the section below [see 'Training AI'].

Neural networks are one specific way of making a machine learning system. Researchers realised that the best way to create an AI system would be to replicate the human brain. In essence, the brain is a

system of neurons connected in complicated ways. Researchers attempted to mimic this structure by building a digital system made up of artificial neurons: ***a neural network***. Just like in a human brain, a neural network is made up of connected neurons. The connections between the neurons store information about the training data. During training, the strength of these connections is modified through a complex statistical process, so that some connections become stronger, and others become weaker. Over time, these adjustments help the network recognise patterns and make better predictions.

Deep learning is a branch of machine learning where a neural network has multiple layers of neurons. These extra layers allow the network to recognise more complex patterns and relationships between inputs and outputs. The learning process works the same way as in simpler neural networks: connections between neurons are adjusted, with some becoming stronger and others weaker, so the network gradually improves at the task.

Training AI: supervised, unsupervised, and reinforcement learning

Machine learning systems are trained over prolonged periods of time, using large amounts of data. As mentioned, the principle underlying machine learning is that the system learns to generalise after learning from many examples. However, there are multiple ways through which this learning process can take place.

In the case of **supervised learning**, the model is trained using labelled examples. Each example in the training data includes both the input and the correct output. For instance, if we want to use supervised learning to train an AI system to identify animals in pictures, we provide it with thousands of images labelled as ‘cat’, ‘dog’, or ‘cow’. By analysing these examples, the system gradually learns the visual characteristics of each animal and can correctly identify new images it has never seen before.

In the case of **unsupervised learning**, the training data is unlabeled. The model is not given the correct answers, but it is instead asked to find similarities between the training data. For example, with animal pictures, an unsupervised model could be instructed to group images that look similar together. The system might create groups corresponding to cats, dogs, and cows, even though it does not know the actual names of the animals.

In the case of **reinforcement learning**, the system learns by trial and error. The system is asked to complete a task and receive feedback in the form of rewards or penalty based on how well it has completed the task. This feedback can be assigned by a human, or automated. The system then adjusts

its behaviour to maximise the rewards, and, over time, it uses the feedback to become better at performing its task. For example, a model trained to identify animals could receive points for correctly identifying an animal and lose points for mistakes.

Supervised, unsupervised, and reinforcement learning are used for different kinds of tasks.

- Supervised learning is suitable when the correct answers are already known, and the system needs to learn how to reproduce them.
- Unsupervised learning works best when the goal is to explore data and find natural groupings or patterns.
- Reinforcement learning is useful for tasks where it is clear what success looks like, but not how to get there.

Types of AI systems

The training methods discussed above can be used to train different types of AI systems.

Predictive AI

Before generative AI became widely available in 2022 with tools like ChatGPT or Gemini, AI typically referred to **predictive AI systems**. The purpose of predictive AI is to use machine learning methods to predict outcomes based on large amounts of data, identifying hidden patterns that would not be found with more traditional statistical methods. This approach is more effective than other statistical methods when:

- there is a lot of data, and;
- the relationships between data and the outcomes are particularly complex and not easily captured by traditional statistical methods.

Natural language processing, computer vision, and speech recognition

These applications use machine learning to create an AI system that is able to recognise and sometimes manipulate language, images or videos, and speech. Some examples are translation software, apps that can be used to recognise plants, and speech recognition in smartphones.

Generative AI

Following the introduction of transformer models in 2017, AI systems gained the ability to process entire sequences of data, such as sentences, at once, rather than one word at a time. Transformer models are a type of neural network that learns the context of sequential data and can generate new data out of it. This breakthrough forms the basis of all modern **generative AI** systems

Explained simply, a generative AI system is one that learns to generate more objects that look like the data it was trained on.The most common generative AI systems are **large language models** (LLMs). LLMs are trained using a large amount of text from the internet, and produce text as output. The first popular LLM was ChatGPT, but other examples include Gemini, Claude, Deepseek, and Mistral. Other generative AI systems can **generate audiovisual outputs**, creating images, videos, or sound from prompts.

The principle behind generative AI is simple: users provide a task in the form of a written instruction, called a prompt, and the system generates an output based on that instruction. These models are easy to use because they do not require technical skills and can take instructions and respond in everyday language, known as 'natural language'. However, generative AI tools do not actually understand what they are creating: they are purely predicting what looks or sounds right based on past examples. This means that they can produce outputs that sound or look convincing, but are inaccurate or misleading We discuss risks of generative AI tools more in depth in the following sections [See 'Risks of generative AI tools']

Agentic AI

Agentic AI refers to systems that combine multiple generative AI capabilities to perform complex, multi-step tasks autonomously with minimal human intervention.
For example, an agentic AI research tool can design a research strategy, perform online desk research, and compile findings in a report with minimal human guidance.

Additional guidance

- BBC, 'AI: 15 key moments in the story of artificial intelligence'.
- Coursera, 'The History of AI: A Timeline of Artificial Intelligence'.
- NewScientist, 'Alan Turing'.
- Forbes, 'Are We Heading Into Another AI Winter?'.
- MIT, 'Machine learning, explained'.
- IBM, 'AI vs. machine learning vs. deep learning vs. neural networks'.
- Oliver Wyman, 'The Risks Of Generative AI To Business And Government'.
- OECD, 'Generative AI: Risks and Unknown'.

AI in Uzbekistan

Before starting or contributing to any AI project, it's important to understand how Uzbekistan is approaching artificial intelligence as a country. Knowing the national vision, policies, and governance structures will help you make sure your work fits within the government's priorities and follows existing rules.

This section gives an overview of Uzbekistan's AI ecosystem, including the organisation involved in overseeing this technology, as well as the most relevant laws and governance frameworks. Understanding this context will help you design and implement AI initiatives that are responsible, effective, and aligned with national goals.

National Strategy for the Development of Artificial Intelligence

As a civil servant, it is important to understand the national vision for artificial intelligence and how your work can contribute to it. The National Strategy for the Development of Artificial Intelligence Technologies until 2030 sets out how Uzbekistan plans to use AI to drive innovation, improve public services, and strengthen the country's digital economy.

Approved by Presidential Decree No. PF-358 on 14 October 2024, Uzbekistan's Strategy for the Development of Artificial Intelligence Technologies until 2030 sets out a **comprehensive roadmap to accelerate the adoption and development of AI across sectors**. The Strategy seeks to establish the legal, technological, and economic foundations needed to integrate AI into everyday government work, business activity, and social life.

For you, this means that any AI project you design or manage should align with the goals of the Strategy, helping ensure that efforts across government contribute to a shared national direction. The Strategy is based on five core objectives aimed at advancing the safe, inclusive, and effective use of AI in Uzbekistan:

1. **Building a strong legal foundation**: Developing national legislation in line with international best practices, promoting innovation, and establishing ethical and security standards for AI use.
2. **Accelerating adoption across sectors**: Promoting the application of AI technologies in social and economic spheres to create high-value products and services.
3. **Strengthening standards and international cooperation:** Aligning national standards with global norms and deepening collaboration with international organisations, research institutions, and private sector partners.

4. **Expanding technical and data infrastructure:** Developing the infrastructure required to train, test, and deploy AI systems, including open data and big data resources.
5. **Enhancing knowledge and skills:** Raising public awareness and improving AI-related competencies and skills across the population to ensure that citizens and professionals can participate in and benefit from AI-driven transformation.

Each of these objectives is supported by specific actions with defined timelines and clearly assigned responsibilities across relevant government entities.

The Strategy also sets out targets to track progress toward its 2030 vision. These include, among others, advancing Uzbekistan into the top 50 countries in the Oxford Insights Government AI Readiness Index, increasing the number and value of AI-based software products developed domestically, and expanding the use of AI technologies across the civil service. Together, these indicators provide a clear framework for monitoring implementation and evaluating the country's overall progress in building AI capacity.

You can find the full strategy here https://lex.uz/en/docs/7159258

Who oversees AI governance in Uzbekistan?

If you're working on or planning an AI project, it's important to know who is responsible for what. Understanding the main institutions involved in AI governance will help you find the right partners, get approvals faster, and ensure your work aligns with national priorities.

AI governance in Uzbekistan is led by the Ministry of Digital Technologies, which acts as the main coordinating body for the country's Strategy for the Development of Artificial Intelligence Technologies. The Ministry makes sure that all ministries and agencies work toward shared goals and that AI-related activities across government remain consistent and aligned with the national strategy.

The Ministry also identifies key AI projects in collaboration with other government bodies and submits them to the Cabinet of Ministers for approval. This ensures that important initiatives receive both political support and the resources needed to move forward.

To support the implementation of the Strategy, the Ministry has created three key institutions:

- **The Center for the Development of Artificial Intelligence and the Digital Economy**, which leads research and supports the use of AI technologies across sectors.
- **The Research Institute for the Development of Digital Technologies and Artificial Intelligence**, which focuses on advancing AI research, promoting innovation, and connecting government, academia, and the private sector.
- **The Artificial Intelligence Alliance**, a new non-governmental, non-profit organization designed to strengthen coordination and collaboration across Uzbekistan's AI ecosystem. The Alliance brings together representatives from government, the private sector, academia, and civil society to support national AI development and promote Uzbekistan's presence in the global AI community.

If you're developing or managing an AI project, these institutions can be valuable partners. They can offer guidance, technical expertise, and opportunities for collaboration.

Relevant governance frameworks

Ethical AI Principles

When designing or using AI systems, civil servants must ensure these technologies uphold fairness, transparency, and accountability. The Government's **Ethical AI Principles** serve as a practical guide to achieve this. Developed in line with international standards such as the UNESCO Recommendation on the Ethics of Artificial Intelligence and the OECD AI Principles, these guidelines reflect Uzbekistan's commitment to ensuring that AI is developed and applied responsibly.

Understanding and applying these principles in your daily work will help you design and use AI systems responsibly. By following them, you can make sure that any AI project you are involved in not only improves efficiency and supports innovation, but also respects citizens' rights, protects their data, and strengthens public trust in how the government uses technology.

The **Ethical Principles for Artificial Intelligence** adopted by the Government of Uzbekistan establish clear guidelines to ensure that AI development and use are responsible and aligned with international standards. The principles can be summarised as follows:

Protecting Personal Data

If you work with AI in government, you will almost certainly work with data, and often with personal information. Understanding Uzbekistan's data protection laws is key to making sure the projects you design or manage respect citizens' rights and comply with legal requirements.

Because AI systems rely on large amounts of data, including personal or sensitive information, they come with important legal and ethical responsibilities. Knowing the national data protection framework will help you make the right decisions when developing or using AI tools in your organisation.

By following Uzbekistan's data protection laws, you help safeguard citizens' privacy, prevent data misuse, and strengthen public trust in digital government. It's important to understand how these rules apply at every stage of an AI project—from collecting and storing data, to processing, sharing, and eventually deleting it—so that AI is used safely and responsibly.

Law on Personal Data Protection

The main legislation on this matter is the **Law on Personal Data Protection**. This law is guided by five core principles that ensure the responsible handling of personal information:

- Respect for individual rights and freedoms: all personal data must be processed for legitimate purposes and in accordance with the law.
- Accuracy and reliability: information must be correct and up to date to avoid harm or misuse.
- Confidentiality and security: personal data must be protected from unauthorised access or disclosure.
- Equality of rights: everyone involved in data processing (citizens, data owners, and operators) has defined and fair rights and responsibilities.
- Protection of society and the state: data protection supports not just individuals, but also national security and public trust in digital systems.

As a civil servant, it's your responsibility to ensure that any AI project you contribute to complies with these principles. Doing so helps create systems that citizens can trust and that align with Uzbekistan's broader goals for responsible digital transformation.

The Personalisation Agency under the Ministry of Justice is the national authority responsible for enforcing these rules. It maintains the State Register of Personal Databases, issues registration certificates for data controllers, and ensures compliance across government institutions. The Agency

also develops procedures for data protection, issues corrective instructions when violations occur, and represents Uzbekistan in international discussions to align national practices with global standards.

Law on Cybersecurity

Another relevant law is the Law on Cybersecurity, which provides a legal framework for protecting information systems and digital assets across the public sector. Since AI systems often rely on complex digital infrastructure, civil servants need to be aware of this law to ensure that government AI deployments protect critical systems, respond appropriately to cybersecurity incidents, and maintain system stability. Compliance with the Law on Cybersecurity contributes to safeguarding data integrity, system availability, and public trust in digital services.

You can find them here:

- **Law on Personal Data:** https://lex.uz/docs/4831939
- **Law on Cybersecurity:** https://lex.uz/en/docs/5960609#

Improving AI literacy in the public sector

As a civil servant, you play a key role in how Uzbekistan adopts and uses artificial intelligence. Whether you work on policy, IT, public service delivery, or administration, understanding AI will help you make better decisions and deliver greater value to citizens.

AI is becoming increasingly common in government work, from automating routine tasks to analysing large amounts of data for better policy insights. To use these technologies effectively and responsibly, it's important to understand what AI can do, what its main risks are, and how to make sure it is used safely and effectively.

Building your AI literacy will help you:

- Recognise where AI could make your work more efficient or effective.
- Understand how to protect citizens' privacy while using AI tools.
- Collaborate with technical experts and external partners.
- Align your projects with Uzbekistan's national AI priorities and ethical principles.

To support you in this journey, the Government of Uzbekistan is investing in initiatives that help public servants develop the skills and understanding needed to use AI effectively and responsibly.

In line with this vision, Presidential Decree No. PQ-189 of 22 October 2025 introduces a comprehensive plan to build AI skills across all levels of education and government. The decree approves a Plan of Measures for developing qualified specialists in artificial intelligence technologies and increasing youth engagement. It includes the launch of the National "AI Day" Programme, designed to raise awareness and improve AI knowledge among school students through dedicated lessons, textbooks, and teacher training materials.

Starting from the 2026/2027 academic year, AI fundamentals and basic skills will also be integrated into the national school curriculum under the subject "Informatics and Information Technologies." This initiative ensures that future generations — as well as today's public servants — develop the foundational understanding needed to use AI confidently, ethically, and effectively in support of Uzbekistan's national digital transformation goals.

To start learning today, initiatives such as AI Study and Five Million Prompters provide accessible ways to learn about AI — what it is, how it works, and how you can use it in your daily work. These programmes are a useful resource for any civil servant who wants to use technology confidently and responsibly.

Additional guidance

- Ministry of Digital Technologies, 'On the approval of the Strategy for the Development of Artificial Intelligence Technologies until 2030'.
- OECD, 'AI Principles'.
- UNESCO, 'Recommendation on the Ethics of Artificial Intelligence'.
- ISO, 'Building a responsible AI: How to manage the AI ethics debate'.

AI in the Public Sector

The relationship between Government and AI

In today's world, how governments use artificial intelligence matters more than ever. For civil servants, understanding the relationship between AI and government is critical, not only to enhance operations and service delivery, but also to foster public trust and align new technologies with national priorities. This section outlines how governments engage with AI.

There are three main ways in which government interacts with AI:

- **Using AI:** governments deploy and adopt AI systems to improve operations, service delivery, and policy-making. For example, AI chatbots can help citizens access government information or complete applications online.
- **Enabling AI:** governments act as a facilitator, creating the right environment for innovation by developing foundational infrastructure (e.g. 5G network or data centers), funding research, developing skills, and supporting collaboration between businesses, universities, and public institutions. For instance, by supporting training programmes, the government helps develop a new generation of AI professionals.
- **Governing AI:** government establishes laws, policies and oversight mechanisms to ensure AI is aligned with national priorities and is developed and deployed effectively and responsibly. Governments play a key role in steering AI in a way that it benefits their citizens and minimises risks. For example, the role of the national government in setting rules to prevent non-consensual use of personal data for training AI systems.

Civil servants in Uzbekistan may engage with AI across all three dimensions, by using it in their daily work, contributing to initiatives that enable its broader adoption, or helping shape governance and oversight mechanisms.

However, **this guide is focused primarily on the first dimension: government as a user of AI**. Unlike the other two dimensions, adoption directly involves individual civil servants and their day-to-day interaction with AI systems, while enabling and governing AI largely occurs at the institutional or policy level.

Government as a user of AI

Governments around the world are increasingly harnessing AI to develop better policies, improve internal processes and make public services more accessible and effective. AI can be used in many areas of government. For example, it can automate repetitive tasks such as processing documents or analysing large datasets, saving time to public employees. It can also help citizens interact more easily with public services through chatbots or virtual assistants, or support policy teams with better data analysis and forecasting.

According to the OECD's report 'Governing with Artificial Intelligence', governments are increasingly adopting AI, leading to benefits in a wide range of areas, including citizen participation, tax administration, and forecasting natural disasters.

To help illustrate what this looks like in practice, Singapore's Public Sector AI Playbook identifies two main ways AI can support your work as a civil servant:

- **Automating repetitive tasks:** using AI to handle routine or time-consuming processes, such as document classification or information extraction, so that public servants can focus on higher-value work.
- **Augmenting human capabilities:** using AI to complement human judgement. For instance, by identifying patterns, predicting outcomes, or personalising services for citizens.

Here you can find an overview of the main applications that fall under these two categories. Please note that the examples below are not exhaustive, but provide an overview of typical areas where governments are applying AI.

Table 1: Common applications: automating repetitive tasks in the public sector[1]

Automating repetitive tasks	
Recognition	
Video analytics Analysing footage to identify elements and their behaviour, count elements or ensure compliance with certain rules.	**Example of use case** Monitoring traffic flow through AI-enabled cameras to manage congestion in cities.
Information extraction and summarisation Automatically processing and summarising large volumes of text to extract relevant information for analysis or decision-making.	**Example of use case** Extracting key insights from thousands of citizen feedback forms to inform service design.
Conversation systems	
Chatbot AI-powered assistants that handle conversations through messaging applications, websites, mobile apps or over the phone.	**Example of use case** Chatbots answering citizens' questions about the documentation required to get a passport or a driving license.
Speech and language tools Converting speech to text, translating languages, or generating written summaries to support efficiency.	**Example of use case** Transcribing parliamentary sessions and providing instant translation to a different language.

[1] This table is based on information from Singapore's Public Sector AI Playbook.

Table 2: Common applications: augmenting human capabilities in the public sector[2]

Automating repetitive tasks	
Personalisation	
Personalised services Using AI to tailor information, recommendations, or services to specific user needs and contexts.	**Example of use case** Recommending relevant training courses to civil servants based on their role and experience.
Pattern and anomaly detection	
Sentiment analysis An AI application that detects sentiments and identifies issues of concern reflected in text documents, helping governments understand public attitudes and emerging topics.	**Example of use case** Analysing survey responses using AI to understand how people feel about new government programmes.
Feedback analysis An AI tool that reads and organises citizen feedback to find common problems or suggestions.	**Example of use case** Reviewing online forms or call centre notes to identify frequent complaints about a public service.
Fraud analytics Identifying irregular transactions or behaviours that may indicate fraudulent activity.	**Example of use case** Checking benefit payment records to find duplicate or suspicious claims.
Forecasting and decision support	
Predictive modelling Using AI to make informed predictions about what happens in the future, based on variables such as past data.	**Example of use case** Forecasting which areas are most at risk of flooding or where demand for healthcare services will increase.

[2] This table is based on information from Singapore's Public Sector AI Playbook.

Additional guidance

- OECD, 'Governing with Artificial Intelligence, The State of Play and Way Forward in Core Government Functions'.
- Singapore Government Developer Portal, 'Public Sector AI Playbook'.
- Oxford Insights, 'Government AI Readiness Index'.
- Singapore Smart Nation and Digital Government Office, 'Public Sector AI Playbook: Annex'.
- UK Government Digital Services, 'Artificial Intelligence Playbook for the UK Government: Appendix: example AI use cases in the public sector'.

How to use generative AI in your job

Using generative AI tools in your daily job

While there are several types of AI systems [see 'Understanding AI'], most people will only use generative AI in their day-to-day work. For instance, you might have heard of or already used generative AI tools such as ChatGpt, Claude, Gemini, Deepseek, Midjourney, DALL-E, or Runway ML.

These tools are easy to use, do not require technical experience, and are often available to the public for free. This Playbook provides an overview to help civil servants better understand how such technologies work and where they may be relevant to government tasks and addressing real-life situations. However, it is fundamental that you understand its **risks and limitations**, and know how to use these tools **ethically, safely and responsibly**.

This section is about using generative AI tools in your day-to-day work. It covers the main risks of generative AI tools, and how to mitigate them by learning to use them ethically, safely, and responsibly. It also discusses some common applications of generative AI that you might explore in your day-to-day work, and use cases to avoid.

Risks of generative AI tools

Generative AI tools present two broad categories of risks:

1. Data privacy: To use a generative AI tool, you need to provide information in the form of a prompt. This could be a few lines of text or even a whole document. Anything you share with a generative AI model could be revealed intentionally or unintentionally:

Your information is typically stored by the service provider. Storing information makes it open to the risk of **data leaks** (where the information is unintentionally revealed because of errors in the provider's system) or **data breaches** (where the information is intentionally by malicious actors, often through a cyber attack).

This information might also be accessed by the service provider, its partners and sub-contractors. For instance, many publicly available generative AI systems use the information provided by their users to **train or improve their models**. This means that your information might become part of the model's data set, and future outputs could reproduce or share part of that information to other users.

2. Biased, inconsistent, or inaccurate outputs. Generative AI tools create outputs such as text, images, or videos, based on patterns they have learned from huge amounts of existing information. These tools

do not actually understand what they are creating, they are purely predicting what looks or sounds right based on past examples. For instance, when a generative AI tool writes text, it chooses each word by predicting what usually comes next in similar sentences. It does not know what the words mean, or whether what it is writing is true.

This means that generative AI tools can produce answers that sound confident and credible, but can be **wrong or misleading**. Generative AI tools might **hallucinate**, inventing facts, statistics, or references that do not actually exist. It is also important to recognise that generative AI outputs are not neutral: their algorithms tend to reproduce the assumptions, and biases that exist within the training dataset.

Generative AI tools also often tend to adjust their response to agree with the user's opinion and expectations. For instance, a generative AI tool might agree with an incorrect statement from its user, repeat a user's mistakes, or change an accurate answer when the user disagrees. This behaviour is called ***sycophancy***. Sycophancy can make generative AI systems less reliable, reinforcing false information and giving users a false sense of confidence, which might lead to poor decision-making.

Using generative AI tools ethically, safely, and responsibly

There are good practices for using generative AI that may apply in your everyday life. As a civil servant, however, you need to be particularly careful about how you use generative AI because:

- You might have access to unpublished or sensitive government information;
- You might have access to personal information;
- Using generative AI incorrectly might lead to poor outputs that directly affect government decisions.

As a general rule, before using any generative AI tools, you should always **ask yourself three questions:**

1. How will this generative AI tool use my information?

You should know **where this information is stored**, **how long the tool's provider will keep it**, and **how it will be used**. This is important to ensure that your information is handled in compliance with data protection legislation and any additional guidance or policy of your organisation. You should always check that the provider's practices align with your national and organisational policies.

These practices can vary depending on the tool and whether you are using a free version. For instance, some tools might allow you to prevent your information from being used for training,

or automatically prevent it for paid accounts. You should always check whether your data might be used for training, and opt out whenever possible.

2. Can I share this information?

- You should never input information that is **classified, sensitive or might reveal the intent of the government** outside of what is already public. For instance, you should never share policy decisions or plans that are not yet publicly available.
- You should never put **personal or identifiable data**, such as a person's name, email, address, and any login credentials and passwords.
- You should not enter any information that is **covered by data protection law**, or an obligation of **confidentiality** owed to a third party.

You can anonymise information by removing personal identifiers (such as names or email addresses) or swapping details for generic names (e.g. "Project X"). However, **anonymisation is not a perfect solution**, as it may still be possible to identify the anonymised information from context.

As a general rule, never enter information that you would not want to be publicly available.

3. How can this output be misleading?

Generated output can look convincing, but it might be wrong. You should always **review any generated output** with a critical eye and verify any sources and links. You should **never rely on AI-generated output as your only source of information**. As a civil servant, it is especially important that you carefully assess generated content as your work directly impacts government policy. Using incorrect output might lead to misguided policy decisions, with negative impacts on citizens' lives.

Common generative AI applications for civil servants

With the right safeguards, generative AI can be a helpful support in your day-to-day tasks. Below, we present some example uses of generative AI, with their risks and limitations. You might already be familiar with some of these, or you might want to experiment with them.

Table 3: Common generative AI applications for civil servants

Application	Risk	Mitigation
Conduct desk research	Generate biased or incorrect content; hallucinate sources and references.	Always review links and references to ensure their existence, credibility, and pertinence; cross-check facts with credible sources; do not rely solely on AI outputs.
Review and summarise documents	Generate biased or incorrect content; expose sensitive information such as government's intent.	Avoid AI-summaries for confidential documents or text; supplement with human review.
Translating text or documents	Generate incorrect translations or missing the original nuances; expose sensitive information.	Avoid AI-translation for confidential documents; be aware that AI translations can limit the full understanding of the original text.
Writing emails	Share personal data; expose sensitive information.	Remove any personal data; avoid confidential input.
Summarising meetings	Share personal data; expose sensitive information.	Consider the appropriateness of AI summaries before each meeting; always review output; remove any personal data.
Writing and creating presentations	Generate biased or incorrect content; expose sensitive information; losing audience trust if the output is evidently AI-generated.	Always review and edit output; avoid confidential input.
Creating images, videos, or sound	Generate biased or incorrect content; misusing or sharing personal characteristics without the person's consent (such as voice, appearance); expose sensitive information; losing audience trust if the output is evidently AI-generated.	Ensure subjects' consent before creating outputs; avoid confidential input.
Writing code	Introduce security vulnerabilities; generate incorrect or unsafe code.	Review all code thoroughly; do not input proprietary algorithms.

Brainstorming	Reinforce biases; expose sensitive information.	Treat AI as a suggestion tool critically assessing ideas; avoid confidential input .

Use cases to avoid

You should never use generative AI tools to fully **write government policy or policy-supporting documents.** This would mean sharing sensitive government information that is not currently public and any incorrect, misleading, or biased AI-generated output would directly influence government action.

This guidance applies to publicly available tools. Your organisation may have paid subscriptions to specific generative AI models, or its own internal models, which might handle and store information differently. Check with your organisation to understand what tools are in use and how they manage data.

Additional guidance

- University of Kansas, 'Helping students understand the biases in generative AI'
- The Atlantic, 'AI Is Not Your Friend'
- UK Government,'Guidance to civil servants on use of generative AI'
- Harvard University, 'Ideas for experimenting with Generative AI: Use cases and things to keep in mind'

Checklist: before using an AI tool

Identifying the tool

Does your organisation have a preferred generative AI tool for staff to use?

Do you understand how the tool you are using will store and process your information?

Will the tool handle your information in line with data protection laws and your organisation's policies?

Are there steps you can take to limit how your information is reused, for example, preventing it from being used to train the AI?

Reviewing the input

Is the information you want to input publicly accessible?

If not:

Do you have permission to share it?

Does any part of this information fall under data protection laws or confidentiality obligations?

Have you removed any personal or identifying information?

Is the information classified, or does it reveal undisclosed information about the intent of the government?

Reviewing the output

Have you reviewed the AI-generated output to ensure it fits your needs and context?

Could the AI-generated output contain misleading or false information?

Have you reviewed all sources of the AI-generated output?

Have you cross-checked the AI-generated output with other reliable sources?

Implementing AI in your project or service

Do you need AI?

AI is a tool, and like any tool, it can help solve some problems, but not all. A hammer is useful for hanging a picture, but it won't help you cook an egg. In the same way, before you decide to use an AI solution, you should first consider whether **AI is the right tool for the problem you are trying to solve**.

To do this, first you need to:

1. Define the problem you are trying to solve;
2. Understand whether AI is the right tool to address it;
3. Understand whether you have the right data for an AI solution.

1. Defining your problem

Defining your problem might mean identifying a current issue in your project or service, or an area where performance could be improved. There might be a step in your current processes that takes a considerable amount of time, is particularly expensive, or where performance could be enhanced. For instance, caseworkers needing to review high amounts of PDF documents may be causing bottlenecks in a citizen-facing service.

To identify a problem, you can:

- **Talk to the people involved in the process**: such as analysts, service workers, or the admin team. Those closer to the process are often the most knowledgeable about the pain points and opportunities for improvement. For instance, a caseworker will know which steps of the process take them the most time to complete, or are the most repetitive.
  You should always conduct some in-depth interviews with your team to understand their perspectives. In addition, you might want to conduct some workshops, focus groups, or surveys to engage a broader number of people.

- **Engage the users of your service:** you can think of users as anyone outside of your organisations who uses or interacts with your project or service. If you are working on a government service, this means talking to the citizens or businesses who use it. Users experience the process from a different perspective, and can provide valuable insights into which steps are slow or could be improved. User feedback complements the perspectives of your team and can help identify opportunities for meaningful improvements. For instance, a citizen using a government service to apply for a passport can say which steps they found most time-consuming, highlighting steps that might be optimised or simplified.

- **Use metrics, benchmarks, or KPIs:** quantitative measures can help you identify the blockers and bottlenecks in process. For instance, the percentage of unsolved cases, or a high number of man hours spent on a task might indicate possible performance gains. You should still speak to the teams involved to understand the reason between such metrics, and what improvements might help them improve performance.

Once you have identified a problem, it can be helpful to **define it using a clear problem statement**. At a minimum, a problem statement should cover:

- **A brief description of the problem:** *e.g. caseworkers in the Department of Pensions spend 10 hours per week reviewing citizen-submitted forms*
- **The impact of the problem:** *e.g. this causes delays for the customer service team and frustration for citizens, and results in 15% of key information being missed*
- **What success would look like:** *e.g. we want a faster, more accurate way to extract this information, reducing the time spent by caseworkers by 50%, cutting errors so that less than 5% of key information is missed, and maintaining full compliance with data privacy requirements.*

2. Understanding whether AI is the right tool to address your problem

Once you have identified and defined your problem, you need to determine whether AI is the right tool to address it. AI works best for particular types of tasks. AI may be suitable if your problem involves tasks that are:

- **data-driven**: with large amounts of data for AI to learn from, such as text, images, audio, or numerical data.
- **repetitive or high-volume**: with large amounts of information that would require a lot of time to process manually
- **digital**: using digital inputs, such as PDFs rather than physical documents, so that the AI tool can integrate easily with existing systems.

Good tasks might also be:

- **pattern-based:** focused on finding patterns, trends, or unusual cases, such as recognising fraud indicators.

- **rule-based:** following clear or learnable rules, such as how certain text indicates a particular claim type.

For instance, AI might be helpful to assist with:

- Reviewing large volumes of text, audio, or video to identify and extract key information.
- Analysing video to detect, identify, or count objects of interest.
- Helping citizens or businesses navigate large amounts of information, such as legal documents or government guidance.
- Transcribing meetings or recordings.
- Translating written or spoken content.
- Detecting patterns in data, such as identifying unusual claims or recurring themes, for example, in fraud detection or analytics.

The 'AI in the Public Sector' section describes common use cases for AI tools in the public service. These examples may also help you consider whether similar solutions could work in your context.

Use cases to avoid

AI systems should never be used for **fully automated decision making**. Fully automated decision making means that the AI system is making decisions without any human input or oversight. This violates the principle of Human Oversight in Uzbekistan's Ethical AI Principles.

This means that a human should always review the AI's outputs before any decisions are made.For example, an AI system might flag a benefit claim as potentially fraudulent. While the AI can help identify risks, a human reviewer should always assess the case before taking any action, such as stopping a payment. This is because AI automated decision making limits:

- **Accountability**: AI systems involve many people and steps. When something goes wrong, it can be difficult to know who is responsible. For instance, responsibility could fall with the developers who built the system, the team who provided the training data, or the civil servant using the tool. This reduces accountability and makes it harder for the government to take clear responsibility for its decisions, which is fundamental for delivering services.
- **Transparency**: AI systems do not always show how they reach a decision (this is called operating as a 'black box'). As a result, neither the people making the decision nor those affected by it can easily understand why the AI produced a certain outcome. This is especially important in government, where decisions may need to be explained clearly to citizens.

- **Fairness**: AI can reproduce biases present in the data or the algorithm, potentially leading to unfair outcomes.
- **Accuracy**: AI systems can make mistakes that a human reviewer might find and prevent.

AI tools are also **not suitable on their own for use cases where rules and conditions change frequently**. As discussed in previous sections [see 'Understanding AI' and 'Understanding whether AI is the right tool to address your problem'] AI performs best with tasks that are based on **clear rules or patterns** from which it can learn. In cases with complex nuances and where rules change often, human oversight is particularly essential to ensure decisions are accurate and appropriate.

Additional considerations

When deciding whether AI is appropriate, consider how it **fits within existing organisational workflows**. Consulting the people involved in the process can help identify where automation is feasible, where it could bring the most benefits, and where it should not be attempted, for example due to complexity, risk, or ethical considerations.

You should also always assess whether the **AI system could create legal or ethical risks.** These might include data protection constraints and cybersecurity concerns. You may want to consult with your organisation's legal or compliance teams as needed.

3. Understand whether you have the right data for an AI solution

Once you have defined your pain points, and assessed that AI could be a viable solution, you need to understand whether you have the right conditions in place for AI to properly work in your project or service.

To be able to use AI in your project or service, you will need **good quality data** that **informs the problem you are trying to solve.**

- AI systems rely on data to learn patterns, make predictions, or generate content. Poor or incomplete data can lead to inaccurate results and unreliable outcomes. Good quality data should be accurate, complete, consistent, and up to date.
- In order to instruct your AI tool, you will need to feed it data on the task you want it to perform. For example, if you want to train a risk model for earthquakes, you must have reliable and

detailed historical earthquake data. If you want an AI model to extract information from PDFs, you must have access to those documents in a consistent, readable format.

The amount of data you need, and the overall complexity of your AI project, will vary greatly depending on whether you are building and training a new model or using a pre-trained one.

- A **pre-trained model** has already been trained on a large dataset for a particular task, and can be reused or fine-tuned for a different but related task. While you might still need data to fine-tune your model, this approach requires far less data, time, expertise, and computing power than training a model from scratch.
- **Building your own model** from scratch means training it entirely on your own data. This gives you full control and flexibility, but it also demands much larger volumes of high-quality data for training, substantial computing resources, and expertise to design, train, and test the model effectively.

Before starting an AI project, you must **assess what data your project will need**. You will need to:

- Identify the data or datasets required for the project.
- Determine the availability of these datasets and how they can be accessed. For example, they may already be collated within your organisation, held by other organisations that you will need to contact, or may need to be compiled by merging multiple sources.
- Ensure that the datasets are in a usable condition, such as being annotated, de-sensitised, or cleaned, and ready for analysis.

Your Chief Data Officer can help you assess your data needs and start the data mapping process.

Additional guidance

- Mural, 'How to Write Problem Statements You'll Actually Use'.
- Singapore Smart Nation and Digital Government Office, 'Public Sector AI Playbook: Annex'.
- UK Government Digital Services, 'Artificial Intelligence Playbook for the UK Government: Appendix: example AI use cases in the public sector'.
- United Kingdom's Information Commissioner Office, 'What is automated individual decision-making and profiling?'.
- Nebius, 'Understanding pre-trained AI models and their applications'.

Checklist: do you need AI?

Defining the problem

- Have you engaged those working on the process you are reviewing?
- Have you engaged the users of your project or service?
- Have you reviewed any quantitative metrics, benchmark, or KPIs?
- Have you drafted a clear problem statement?

Understanding whether AI is the right tool to address your problem

- Have you assessed whether AI capabilities are appropriate for this problem?
- Have you considered how AI would fit with existing workflows?
- Have you considered whether using AI could create legal or ethical risks?
- Have you engaged your organisation's legal, compliance, or data protection teams to review potential risks?

Understand whether you have the right data for an AI solution

- Have you identified the data you would need to inform an AI solution?
- Have you verified the availability, quality, and format of this data for AI use?
- Do you know who owns this data, and if it is not your organisation, do you know how to access it?
- Have you engaged your Chief data officer to discuss data needs and data mapping?

Next steps: building and deploying AI

Once you have understood whether you need AI, you will need to think about how to create, deploy, and maintain the AI system.

When designing an AI project, it is best to start small and increasingly expand on what works, rather than trying to create a large, complex system from the start. This approach is called **iterative design.** Iteration means improving your AI system one step at a time. In an iterative process, you develop and test small versions of your AI system, learn from the results, and make gradual improvements. Each round, or iteration, builds on what you learned before. This helps you refine how the system works and respond to new data or feedback early, without wasting time or resources on a full system that might not work as intended. Iteration works particularly well with AI because models often need fine-tuning: they rarely perform perfectly on the first try, and small changes in the data, algorithms, or settings can have a big impact on results.

There are several ways to create an AI solution. You can:

- **Re-use what already exists:** many governments develop their own AI systems that can be applied across departments. You should check whether the government of Uzbekistan already possesses an AI system that might work for your task.
- **Build it in-house:** your team or organisation might be able to build the solution using internal experts, such as developers or data scientists. This might work well for large departments with sufficient expertise, or for particularly sensitive cases.
- **Build it with external support:** you can hire external contractors to build the AI system. This might be helpful if you do not have the right skills in house, but it requires clear instructions and oversight. Choose a contractor through public procurement procedures with experience specifically in AI projects and an understanding of your domain.
- **Buy it:** depending on the task you need, you might be able to purchase an off-the-shelf product from a third party supplier.

How AI systems are developed and deployed: the AI lifecycle

The AI lifecycle is the process of moving from a business problem to an AI solution that solves that problem. As mentioned above, the AI lifecycle is not a linear process. You are likely to need to take often go back and repeat steps, revisiting each of the steps in the life cycle is revisited many times throughout the design, development, and deployment phases.

Different sources define the stages of the AI lifecycle in slightly different ways, but most cover five main phases:

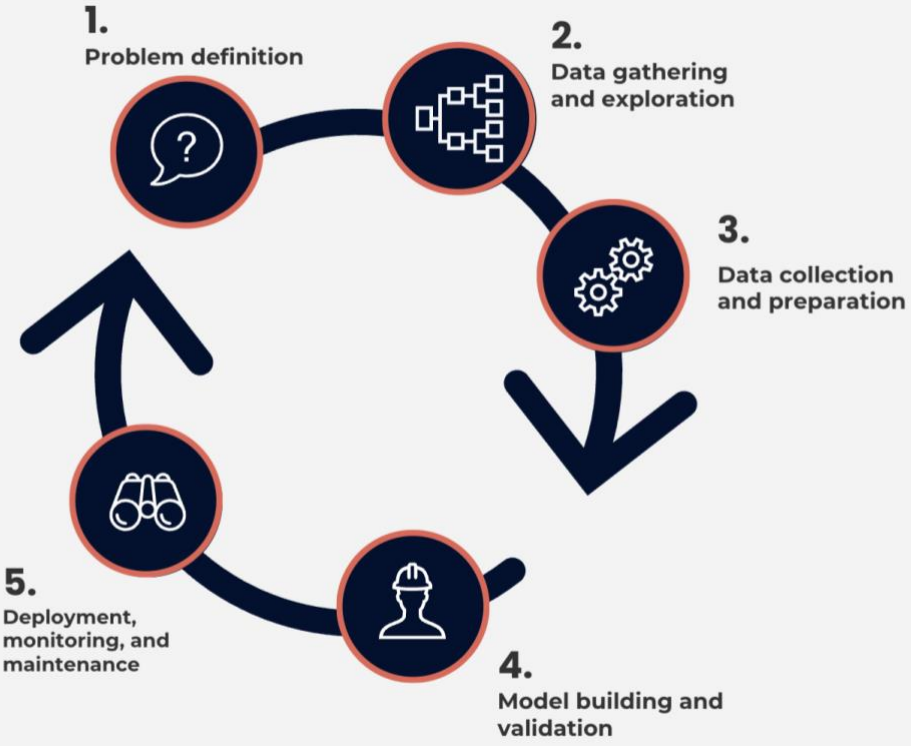

You should keep the AI lifecycle in mind when planning how to develop an AI solution. For example, you will need to consider any training for staff who will use the system, as well as how to embed monitoring and maintenance into your processes.

Regular monitoring and maintenance are particularly important as **AI systems can lose accuracy over time**, a problem called model drift. Model drift can happen for several reasons. For example, the data the model was trained on may become outdated or no longer reflect current conditions, or the model may not be designed to handle changes in the data over time. For this reason, it is important that you include maintenance and monitoring plans from the very beginning of your AI project, especially when specifying requirements in a contract for a purchased AI solution.

Starting your AI project

Which option you choose, and how you proceed, will depend on several factors, including internal skills and expertise, budget, timeframes, data availability, task complexity, and regulatory or security requirements.

The Ministry of Digital Technologies can help you assess these factors and define the next steps. You can contact them through the Ministry's official channels — via the website digital.gov.uz or by email at info@digital.uz.

Additional guidance

- Inter-American Development Bank, 'Responsible use of AI for public policy'.
- University of Michigan-Dearborn, 'AI's mysterious 'black box' problem, explained'
- Datacamp, 'Understanding Data Drift and Model Drift'
- United Kingdom's Information Commissioner Office, ' Annex A: Fairness in the AI lifecycle'