

TRAINING MANUAL

Course on
**ARTIFICIAL INTELLIGENCE 101 FOR THE
PUBLIC SECTOR**

JANUARY, 2026

Message by the Minister for Planning, Development & Special Initiatives



I welcome this initiative aimed at enhancing understanding of Artificial Intelligence within the public sector. In an era of accelerating technological change and increasing demand for data-driven governance, strengthening the digital and analytical capacity of our civil service is essential for effective and informed public administration.

Pakistan's development strategy, encapsulated in the URAAN Pakistan, 5Es-based Five-Year National Economic Transformation Plan: Exports, E-Pakistan, Environment, Energy, and Equity, provides a roadmap for inclusive and sustainable growth. This framework highlights the importance of digital transformation and human capital development as critical components of socioeconomic progress, and it aligns with ongoing efforts to modernize public service delivery and policy implementation.

The training material presented in this manual will support civil servants in understanding the practical applications of AI, prompt engineering, and emerging technologies in governance. I encourage officers to engage deeply with this resource and apply these insights thoughtfully in their day-to-day roles, improving outcomes for citizens across Pakistan.

I commend the Civil Services Academy and its partners for this timely contribution to capacity building and look forward to the positive impact it will have on our collective efforts to advance responsive and forward-looking public administration.

Ahsan Iqbal Chaudhry

Minister of Planning, Development
& Special Initiatives
Government of Pakistan

Message by the Minister for Information Technology & Telecom



The formulation of Pakistan's National Artificial Intelligence Policy marks a significant milestone in our national journey toward responsible and inclusive digital transformation. The policy recognizes that while technology will shape the future of governance, real and lasting impact depends on the capacity of our institutions and the preparedness of our people.

A key pillar of the National AI Policy is therefore the upskilling of the public sector—ensuring that government officers at all levels are equipped to understand, use, and govern Artificial Intelligence with confidence and responsibility. Without this human and institutional readiness, even the most advanced technologies cannot deliver meaningful outcomes for citizens.

This training program at the Civil Services Academy represents a practical step in translating that policy vision into action. By introducing probationary officers to the fundamentals of AI, its applications in public administration, and the principles of ethical and responsible use, we are embedding digital competence at the very foundation of civil service training.

The objective is clear: to build AI-literate public leaders—officers who can make informed decisions, safeguard public interest, and leverage emerging technologies to improve governance, efficiency, and service delivery.

I commend the Civil Services Academy and atomcamp for institutionalizing this initiative and for setting a national example of how policy commitments can be transformed into sustainable institutional capacity.

Shaza Fatima Khawaja

Minister for Information Technology &
Telecom
Government of Pakistan

Message by Director General Civil Service Academy (CSA)



It is with great pride that I announce the introduction of our new Artificial Intelligence (AI) module at the institute, an important stride in advancing the capabilities of our public sector workforce. In an era where data-driven governance and technological innovation are reshaping the way governments serve their citizens, understanding and applying AI is no longer optional; it is essential.

This module has been carefully designed to provide public officials with the knowledge, analytical tools, and ethical frameworks required to harness AI for evidence-based policymaking, improved service delivery, and enhanced administrative efficiency.

Through this initiative, we aim to strengthen institutional capacity and support the national vision for a digitally empowered and citizen-centric public service. We are living through a decisive shift. Artificial Intelligence is no longer a distant concept but a force already reshaping how economies function, how institutions operate, and how citizens experience the state.

As public servants, you stand at the forefront of transformative change. I urge all participants to approach this program not merely as a training exercise, but as an opportunity to reimagine the future of governance through innovation, collaboration, and integrity. The responsible use of AI can help us anticipate challenges, design inclusive solutions, and ensure that technology serves the greater public good. Let this module inspire you to lead with foresight and purpose, shaping a public sector that is adaptive, transparent, and responsive to the evolving needs of our society.

Farhan Aziz Khawaja

Director General
Civil Service Academy (CSA)

Message by the Secretary, Ministry of Planning, Development & Special Initiatives



The Government of Pakistan remains committed to strengthening the capacity of public sector institutions through the responsible adoption of emerging technologies that enhance efficiency, transparency, and quality of service delivery. Artificial Intelligence is no longer a future concept; it is a present capability that must be understood, governed, and applied with care.

This training manual represents a collaborative initiative of the Ministry of Information Technology and Telecommunication, the Civil Services Academy, and atomcamp, developed to introduce Civil Service Academy probationers to the foundational principles and practical implications of Artificial Intelligence in governance.

The focus of this initiative is not experimentation, but structured and informed adoption—positioning AI as a decision-support tool that operates within established governance frameworks, institutional protocols, and principles of public accountability. This document is the first in a planned series aimed at building progressive AI literacy across the civil service. It establishes a baseline understanding of AI and prompt engineering, with emphasis on real administrative use cases, responsible deployment, and the necessity of human oversight.

It is my expectation that officers engaging with this material will approach it with professionalism and critical judgment, recognizing that while AI can enhance productivity and analytical capacity, responsibility for decisions and outcomes always rests with the public servant. This initiative marks an important step toward developing a more capable, digitally empowered, and future-ready civil service.

Awais Manzur Sumra

Secretary Ministry of Planning,
Development & Special Initiatives
Government of Pakistan

Message by Secretary, Ministry of IT & Telecommunication



The Government of Pakistan is committed to strengthening the capacity of its public sector institutions by responsibly adopting emerging technologies that enhance efficiency, transparency, and service delivery. Artificial Intelligence is no longer a future consideration; it is a present capability that must be understood, governed, and applied with care.

This training manual represents a collaborative initiative of the Ministry of Information Technology & Telecommunication (MoITT), the Civil Services Academy (CSA), and atomcamp.

The focus of this initiative is not experimentation, but structured adoption—ensuring that AI is used as a decision-support tool aligned with governance standards, institutional protocols, and public accountability.

This document is the first in a planned series designed to build progressive AI literacy across the civil service. It lays the foundation by introducing core concepts of Artificial Intelligence and prompt engineering, with an emphasis on real administrative use cases, responsible usage, and human oversight. Subsequent materials will continue to deepen practical capability while addressing policy, security, and implementation considerations.

It is my expectation that officers engaging with this material will approach it with professionalism and critical judgment, recognizing that while AI can enhance productivity, responsibility for decisions and outcomes always rests with the public servant.

This initiative marks an important step toward a more capable, digitally empowered, and future-ready civil service.

Zarrar Hasham Khan

Secretary Ministry of Information
Technology & Telecommunication
Government of Pakistan

Message by Co-Founder atomcamp



It is a great pleasure to join hands with the Civil Service Academy, Lahore, and the Ministry of Information Technology & Telecom in this important initiative to upskill the next generation of civil servants and decision-makers.

We are living through a decisive shift. Artificial Intelligence is no longer a distant concept but a force already reshaping how economies function, how institutions operate, and how citizens experience the state.

From predictive systems in public health and agriculture to AI-assisted policymaking and data-driven service delivery, emerging technologies are redefining what effective, transparent, and accountable governance looks like. In this era, data is a strategic resource and AI is the capability that turns that data into insight, foresight, and action across every sector of our national life.

This moment demands public leaders who understand AI not as a buzzword, but as a governance priority. They must know how algorithms influence decisions, how to regulate and deploy these technologies responsibly, and how to ensure that AI becomes a driver of inclusive development rather than inequality. Through this collaboration, we are laying the foundation for an empowered, AI-literate leadership that can ask the right questions, make informed decisions, and guide responsible innovation within government. The journey toward a smarter, more resilient, and more sustainable future begins now - and I am honoured to be part of it.

Naveed Iftikhar

Co-founder atomcamp

Acknowledgement



I am pleased to formally extend my profound appreciation to our respected Director General, CSA, for his visionary leadership and strategic foresight in introducing the Artificial Intelligence (AI) module into our training framework. This initiative marks a significant step toward modernizing our curriculum and ensuring that our probationers are equipped with the skills required to navigate an increasingly technology-driven environment.

I would also like to acknowledge the faculty of the Program Wing, CB Wing, and the IT Team for their dedicated efforts, technical competence, and collaborative spirit in supporting the development and implementation of this module. Your contributions reflect commendable professionalism and a strong commitment to academic excellence.

Additionally, I wish to offer special thanks to Dr. Muqem ul Islam, Director General (KIMS), for his invaluable support, and to Dr. Naveed Iftikhar, Co-Founder atomcamp, for his expert guidance and partnership in strengthening the AI learning experience for our trainees. Their contributions have added immense value to this initiative and further enriched the quality of the program.

Collectively, these efforts demonstrate our institution's commitment to innovation, capacity-building, and continuous improvement in public service delivery and serving the nation. I extend my sincere appreciation to all individuals and departments involved for their unwavering dedication and exemplary teamwork.

Syed Shabbir Akbar Zaidi

Director (CTP/CB) Civil
Service Academy (CSA)
Walton, Lahore

Table of Contents

➤	COURSE OVERVIEW	-----	10
➤	MODULE 1 : INTRODUCTION TO AI	-----	13
➤	MODULE 2 : INTRODUCTION TO LARGE LANGUAGE MODELS (LLMS)	-----	30
➤	MODULE 3 : INTRODUCTION TO PROMPT ENGINEERING	-----	46
➤	MODULE 4 : BEGINNER HANDS-ON AI TOOLS FOR PRACTICAL APPLICATION	-----	51
➤	Annexure	-----	58
➤	Contact us	-----	62

Course Overview

This two-day beginner-level instructional program provides probationary officers at the Civil Services Academy with a structured introduction to Artificial Intelligence (AI) and prompt engineering, with a specific emphasis on their application in governance and public sector management. While primarily designed for civil servants graduating from the Civil Services Academy, the program is equally applicable for public sector officials across ministries, departments, public sector training academies, and government organizations seeking a foundational understanding of AI for administrative and policy-related work.

The curriculum covers the principles, uses, and limitations of AI and trains officers to apply no-code tools and prompt engineering strategies to generate precise, context-aware, and administratively appropriate outputs. The program emphasizes the responsible and informed use of AI, ensuring that officers understand both the capabilities and constraints of such systems in official settings.

The accompanying manual serves as a guide to support officers and public sector officials in adopting AI concepts and tools in a manner consistent with professional standards, ethical considerations, data confidentiality requirements, and established government protocols.

Duration: 02 days

Target Audience: Probationary officers at the Civil Services Academy, civil servants involved in future policy making, administration, office management, and public service delivery, as well as public sector officials from ministries, departments, public sector training academies, and autonomous bodies seeking to build foundational capacity in Artificial Intelligence.

Prerequisites: Basic computer literacy and familiarity with text-based interfaces

Note on the AI Training Series

This manual is the first volume in a structured, multi-stage national training series on Artificial Intelligence for the public sector. We will publish the next three manuals soon.

The series has been designed as a progressive capability-building pathway—moving from foundational awareness to advanced institutional readiness. Civil servants develop not only understanding, but practical and strategic competence in AI for governance.

Structure of the Series

The training framework is organized into four levels:

AI 101 — Foundations of Artificial Intelligence for the Public Sector

AI 201 — Applied AI Systems, Governance, and Secure Deployment in Government

AI 301 — Government-Grade Models, Fine-Tuning, and AI Infrastructure

AI 401 — National AI Infrastructure, Data Centers, and Digital Sovereignty

Learning Objectives:

01

AI Understanding & Application

02

Master Prompt Engineering Techniques

03

Execute Government-Specific Tasks Using AI

04

Ensure Ethical, Responsible, and Secure Use

05

Apply Learning Through Practical Exercises

MODULE 1 : INTRODUCTION TO ARTIFICIAL- INTELLIGENCE

What is AI?

Artificial intelligence (AI) is a system that learns from data and uses that learning to make predictions or recommendations or actions. Instead of following fixed rules, AI improves its output by finding patterns in past information and applying them to new situations.

Data is the foundation of AI. This data can be text, images, audio, video, numbers, or records of human activity. AI systems convert this data into mathematical signals and learn which patterns are significant and which are not.

After training, an AI system can classify information, predict outcomes, undertake actions, or support decisions. For example, it can estimate the demand of a commodity, detect unusual behavior in a public space, or suggest the next best action regarding a legal issue based on previous cases.

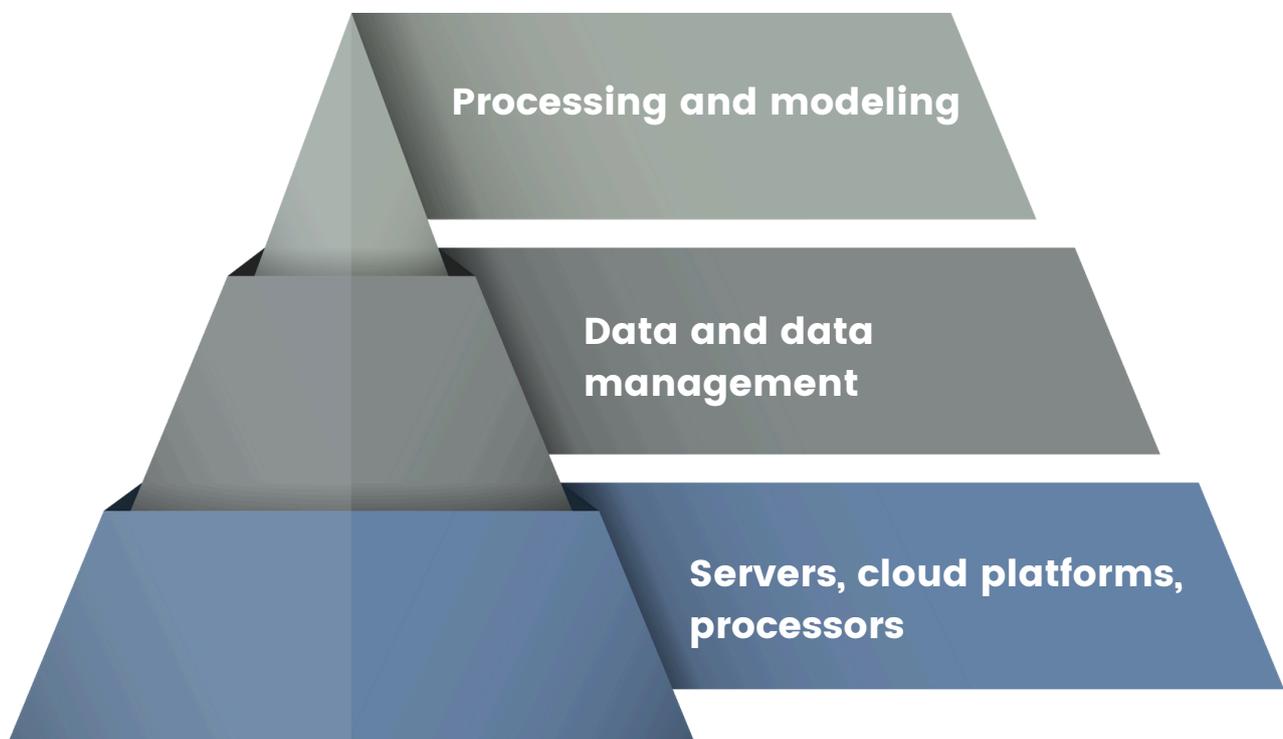
For public service use, it is important to understand that AI does not think or exercise judgment. It produces outputs based on the data it has seen, the assumptions built into its design, and the objectives it is given. This means AI can support officials by improving speed, consistency, and analysis, but responsibility for decisions, accountability, and public outcomes always remains with the human authority.

AI is therefore best used as a decision-support tool, not a decision-maker. Its value depends on the quality of data, clarity of objectives, and the controls put in place to ensure lawful, fair, and transparent use.

AI functions through a small set of interdependent layers. The first is computing infrastructure, which includes servers, cloud platforms, and processors that provide the raw power needed to train and run AI systems. Without adequate computing capacity, AI models cannot be trained efficiently or deployed at scale.

The second layer is data and data management. This covers how data is collected, cleaned, stored, and governed. High-quality, well-managed data is essential because AI outputs are only as reliable as the data they learn from.

The third layer is processing and modeling, where algorithms analyze data to produce predictions or recommendations. This layer turns raw data into usable outputs, with humans responsible for defining objectives, validating results, and applying judgment in real-world contexts.



AI is not new

AI did not start in 2022 or 2023. It is not something that suddenly came out of nowhere.

AI stands on hundreds of years of work in:

- Mathematics
- Statistics
- Programming
- Logic and reasoning

In 1955, John McCarthy gave the formal term “Artificial Intelligence.” Since then, many researchers have developed ideas about how machines can learn, solve problems, and make decisions.

The real jump happened in the 1990s, when the internet, the web, and computing power grew rapidly. After that:

- Better hardware
- Faster chips
- Cheaper storage
- Cloud computing

All these things made it possible to build AI systems at a large scale. So, AI is the result of decades of work.

What is new today is not the idea of AI, but its accessibility and practical use. Systems that once existed only in research labs are now embedded in everyday tools used by governments, businesses, and citizens, making AI a direct operational factor in policy, regulation, service delivery, and public accountability.

AI will reshape how people work, how services are delivered, and how businesses compete by changing costs, speed, and decision-making across the economy. Policymakers must therefore understand AI to deliberately maximize public value while managing risks related to jobs, equity, safety, and trust.

Digitilisation and AI



Digitisation

Converting analog information into digital form.



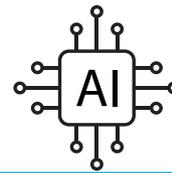
Digitalisation

Using digital technologies to improve existing processes.



Digital Transformation

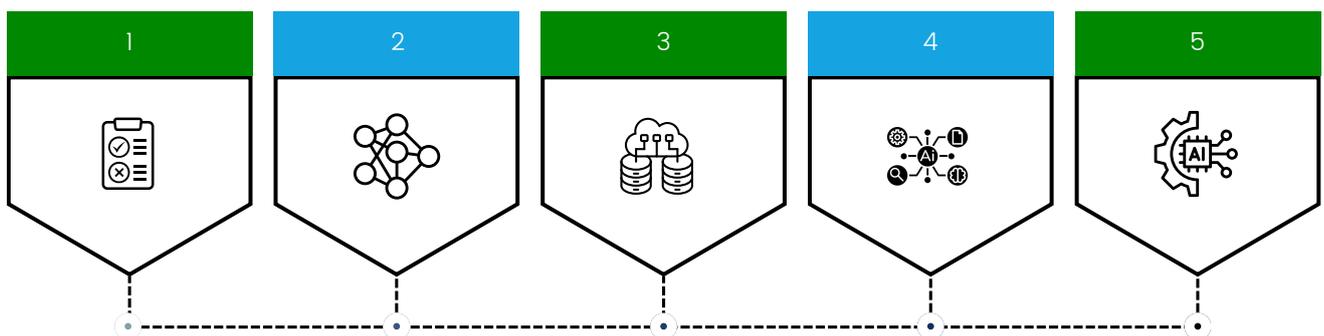
Reinventing how an organisation operates using digital tools – changing culture, processes, and citizen services.



Artificial Intelligence

Machines that can learn, reason, and perform tasks that normally require human intelligence.

Evolution of AI



Rule-Based Systems

- Humans wrote explicit “if-then” rules for machines.
- Examples: **Expert systems for medical diagnosis.**

Machine Learning Emerges

- Instead of hand-coding rules, computers learn patterns from data.
- Examples: **Spam email filters.**

Deep Learning & Big Data

- Neural networks with many layers (“deep”) learn from massive datasets.
- Global Examples:
- (**Google Photos, self-driving cars.**)

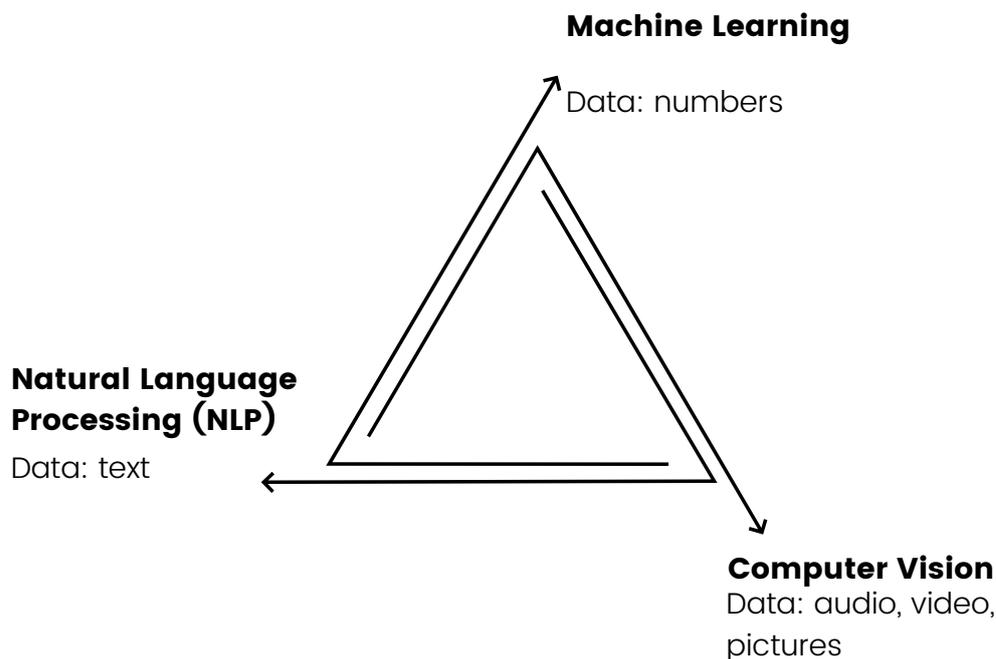
Generative AI

- AI can now create text, images, audio, and even code.
- Global Examples: **ChatGPT, DALL-E, Google Gemini.**

Agentic AI

- AI systems that not only answer but act on your behalf – connecting across tools, automating workflows.
- Global Examples: **AutoGPT, Microsoft Copilot, Lindy.ai.**

Subfields of Artificial Intelligence



Different AI sub-fields primarily work with different forms of data. Machine learning often relies on structured numerical data, natural language processing focuses mainly on text and speech, and computer vision works with images and video.

Foundation models & LLMs

Advanced AI systems are built on Foundation Models (FMs). These are large, general-purpose models trained at scale and adapted across many tasks, marking a shift away from narrow, task-specific machine learning.

Large Language Models (LLMs) are one prominent class of foundation models, specialized in understanding and generating human language.

Other equally critical foundation model classes include vision models for interpreting images and video, multimodal models that combine text, images, and video, and scientific models used for climate forecasting, genomics, and materials discovery. Together, these models constitute the current frontier of AI capability.

Why AI Matters for Government, Policy, and Society

Artificial Intelligence is not only changing how tasks are performed; it is reshaping how work is organized, how public services are delivered, and how societies function.

Across government operations, AI is increasingly used to support decision-making, improve service delivery, manage large volumes of information, and increase administrative efficiency. Beyond government, its influence extends to healthcare, education, finance, industry, culture, and the future of work.

For policymakers and civil servants, understanding AI is no longer optional. Effective governance now requires a clear understanding of how AI systems work in practice, where their limitations and risks lie, and how quickly technological change can outpace existing laws, regulations, and institutional processes.

Policymakers need to understand AI's societal and sector-specific impacts, helping participants anticipate change, assess policy consequences, and adapt regulations in a way that protects public interest while remaining practical and future-oriented.

As governments begin to use AI at scale, the role of the state expands beyond adoption to stewardship. Public institutions must decide where AI can assist human judgment, where it should be limited, and where it should not be used at all. These choices affect public trust, accountability, fairness, and long-term state capacity. Clear policy direction is therefore essential to ensure that AI strengthens institutions rather than undermines them.

Why Does Data Matter?

Everything in AI depends on data quality, this theme appears throughout the document. AI is only as good as the information it learns from.

Example: If you train a model on messy customer data, the model will make messy predictions.

The three most important skills for AI (Large Language Models) users

These three skills make anyone effective with AI - no coding required:

a. Asking Clear Questions (Prompting)

Clear instructions lead to clear answers

Example:

Instead of: "Explain finance."

Try: "Explain basic personal finance to a beginner in 5 bullet points."

b. Checking and Verifying Responses

AI sometimes produces inaccurate or fabricated details. Users should verify important information.

Example:

If AI generates a medical explanation, cross-check it with a trusted source.

c. Using AI with Your Own Context

Context lets AI read your documents so answers are accurate and grounded.

Example: Upload your company policies → ask questions → AI responds using the exact document.

Everyday Examples of AI

Here are simple, practical examples you can relate to:

- Email Assistance: AI drafts replies, but you review tone and correctness.
- Document Summaries: AI summarizes PDFs, but you confirm important details.
- Language Help: AI rewrites text in simple English for clarity.
- Data Lookup: AI answers questions using your uploaded files

AI Safety Basics

AI literacy also includes safe usage:

- Don't paste confidential data into public models.
- Always double-check critical information (legal, medical, financial).
- Understand model limitations - AI is not a final authority.

Never upload confidential summaries, meeting minutes and documents on LLMs like ChatGPT, Gemini or Deepseek.

How AI Connects to Machine Learning and Deep Learning

Now that you understand AI at a high level, it's important to see how the different parts of AI fit together. Artificial Intelligence is a broad field. Under it, we have Machine Learning (ML) and Deep Learning (DL), which are simply different ways of building intelligent systems.

Think of it like this:

- AI → The big goal: making computers think, learn, and make decisions.
- Machine Learning → A major method inside AI that allows computers to learn patterns from data.
- Deep Learning → A more advanced type of ML that uses layered neural networks to understand complex patterns like images, speech, and language.

Why Machine Learning and Deep Learning Matter for AI Today

AI used to rely on rules written by humans. But as you saw in earlier sections, modern AI works because systems can learn instead of being manually programmed. That learning ability comes from ML and DL.

What is Machine Learning and Deep Learning?

To make it simple:

Machine Learning

Machine Learning means the computer learns patterns from data instead of us telling it every single rule.

Examples:

- Predicting budget spend
- Predicting population growth

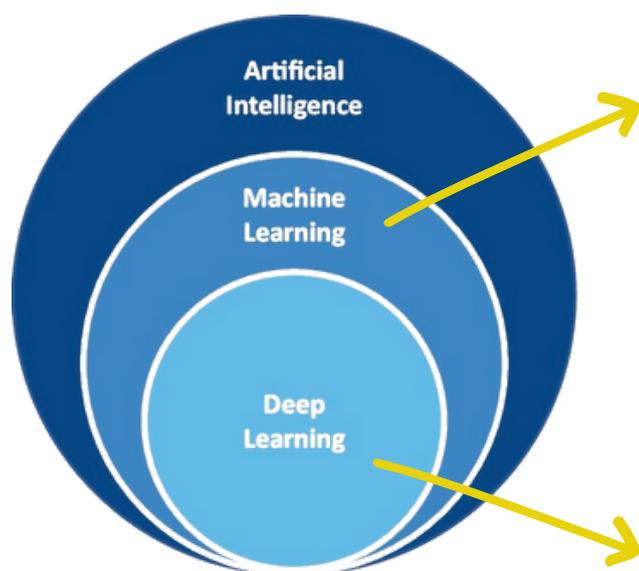
The machine looks at past data and finds patterns.

Deep Learning

Deep Learning is a special type of Machine Learning that uses many layers of computation. These layers help the system understand complex things like:

- Images
- Speech
- Videos
- Natural human language

Deep Learning is why we have modern AI today.



Machine learning (ML) is a type of artificial intelligence (AI) that trains algorithms on data to find patterns, make predictions, and improve performance on tasks without being explicitly programmed for every scenario, enabling systems to learn from experience

Deep learning is a subset of machine learning that uses multi-layered artificial neural networks, inspired by the human brain, to learn complex patterns from vast amounts of data, enabling computers to perform tasks like image recognition, speech transcription, and language understanding with minimal human intervention, powering modern AI applications.

AI Capabilities for Different Use Cases

AI Capability	What Makes It Work
ChatGPT answering questions	Deep Learning + LLMs
Predicting budget spend	Machine Learning
Face detection in safe city cameras	Deep Learning
Email spam detection	Machine Learning
Image generation	Deep Learning

One Simple Example to Connect Them All

Imagine you want a system that identifies whether a photo contains an authentic ID

- Using AI: You want the computer to “recognize” the person using his ID card photo
- Using Machine Learning: You gather thousands of labeled images and teach the system patterns of what a authentic ID card looks like.
- Using Deep Learning: A neural network automatically learns much more complex details -text positions, photo details- without you explicitly programming any of it.

Practical Use of Artificial Intelligence in Government Work

This section explains where and how AI can be practically used in government operations, with a focus on tasks civil servants already perform.

Drafting and Documentation

Civil servants spend a large part of their time preparing:

- Office notes
- Briefs for seniors
- Speeches and statements
- Summaries of meetings

How AI helps

- Producing first drafts of notes and summaries
- Converting bullet points into formal language
- Rewriting documents for clarity and structure

Good practice

- Use AI for the first draft only
- Always review, verify, and finalize yourself
- Never submit AI-generated text without human approval

Research and Policy Support

Officers frequently need:

- Background information on issues
- Comparisons of international practices
- Summaries of laws, reports, and studies

How AI helps

- Summarizing lengthy documents
- Structuring policy briefs
- Listing pros and cons of options
- Generating issue briefs for meetings

Good practice

- Cross-check all facts
- Use official sources for final validation
- Treat AI output as a research assistant, not a source of record

Administrative Efficiency

Routine work such as:

- Email drafting
- Scheduling support
- Preparing standard replies
- Formatting documents

How AI helps

- Drafting routine correspondence
- Creating templates for repetitive tasks
- Structuring reports and presentations
- More time to focus on: analysis, oversight, field engagement, and decision-making

Using AI in Service Delivery

AI can support—not replace—citizen-facing services.

Examples

- Chatbots for basic information
- Automated responses for common queries
- Ticket classification in grievance systems
- Support for helplines and facilitation desks

What Must Always Remain Human

- Final decisions on cases
- Complaint resolution authority
- Disciplinary and legal actions
- Any matter affecting rights, benefits, or penalties

Rule:

AI can assist service delivery.

It must never replace accountability.

From Perception AI → Generative AI → Agentic AI

1. Perception AI (The “Seeing and Hearing” Era)

AI has evolved in stages. Each stage added new abilities and allowed machines to work in more intelligent ways. Here’s a simple breakdown to help you understand how we reached today’s modern AI systems.

What it can do:

- See objects in images
- Recognize faces
- Detect suspicious activities from video feeds.
- Hear and transcribe speech
- Classify documents or sounds

Examples:

- Phone unlock with face recognition
- Google Photos detecting “cat,” “mountain,” or “birthday”
- Security cameras detecting motion
- Speech-to-text apps convert voice into written words

Perception AI understands, but it does not create.

2. Generative AI (The “Creating” Era)

This is the second wave. Generative AI can write, create, and generate new content based on patterns it learned.

What it can do:

- Write essays, emails, and summaries
- Generate images, videos, and voice
- Produce code
- Answer questions
- Translate languages
- Draft reports or documents

Examples:

- ChatGPT is writing a lesson plan or summary
- Midjourney is generating an artwork
- GitHub Copilot writes code automatically
- AI tools are creating marketing posts or product descriptions

Generative AI can produce new content, not just classify it.

3. Agentic AI (The “Doing” Era)

Agentic AI is the newest stage. These systems don't just create content - they can take actions, use tools, and complete tasks on your behalf.

What it can do:

- Book meetings
- Search the internet and gather information
- Send emails
- Update databases
- Run workflows or automations
- Follow multi-step goals

Examples:

- An AI assistant that reads your email, drafts responses, and sends them when approved
- An AI agent helping/answering calls citizens on motorway and highways.
- An AI agents that routes citizen complaints to relevant departments ,follow up on actions and reports back with progress.
- An AI finance agent that categorizes expenses and generates monthly reports
- AI that takes a goal like “Prepare a 5-slide summary from this PDF” and performs multiple steps automatically

Agentic AI can plan, reason, and act, not just see or create.

AI vs Rule-Based Automation

Many people confuse AI with automation. Rule-based automation follows fixed, human-written instructions: if X happens, do Y. It cannot learn or improve - everything depends on predefined rules.

AI, on the other hand, learns patterns from data instead of relying on hard-coded logic. It adapts, improves with more data, and can handle complex, unpredictable situations where manual rules would fail.

Rule-Based Automation

- You write fixed rules
- If X happens, do Y
- No learning
- No improvement
- Everything depends on human-defined steps

Example - Rule based vs AI:

The rules are already described in government systems to transfer payroll on monthly basis. A simple automation executes this workflow.

AI system is when you use the historical pattern to detect fraud. Over time the system learns how much is transferred from government accounts, how much should have been transferred. By learning this patterns, AI can detect fraudulent activities.

Similarly in Auditing, an AI agent/system can scans huge data by quickly reading millions of records at once, spotting patterns and highlighting anything that looks unusual compared to normal behavior. Humans must review and judge these alerts to confirm real fraud and ensure fairness.

Module 2 : Introduction to Large Language Models (LLMs)

Large Language Models are AI systems that understand and generate human language. They are trained on large amounts of text, such as books, articles, websites, and code. Through this training, they learn how language works in context, which allows them to write, summarize, and answer questions in natural language.

LLMs work by predicting the next word based on previous words. At scale, this enables useful language-based assistance, but it does not mean the model truly understands facts or intent.

How LLMs Are Developed

- Large volumes of text data are collected and cleaned
- The model is trained to predict the next word
- Deep learning layers capture language patterns
- Training requires high-performance computing (GPUs)
- Models are then fine-tuned for safety and quality
- Deployed on cloud or on-premise systems

What Policymakers Need to Understand

LLMs learn patterns in language, not verified facts. They are powerful support tools, but their outputs must be reviewed and governed, especially in public sector use.

Simplified LLM Training Workflow

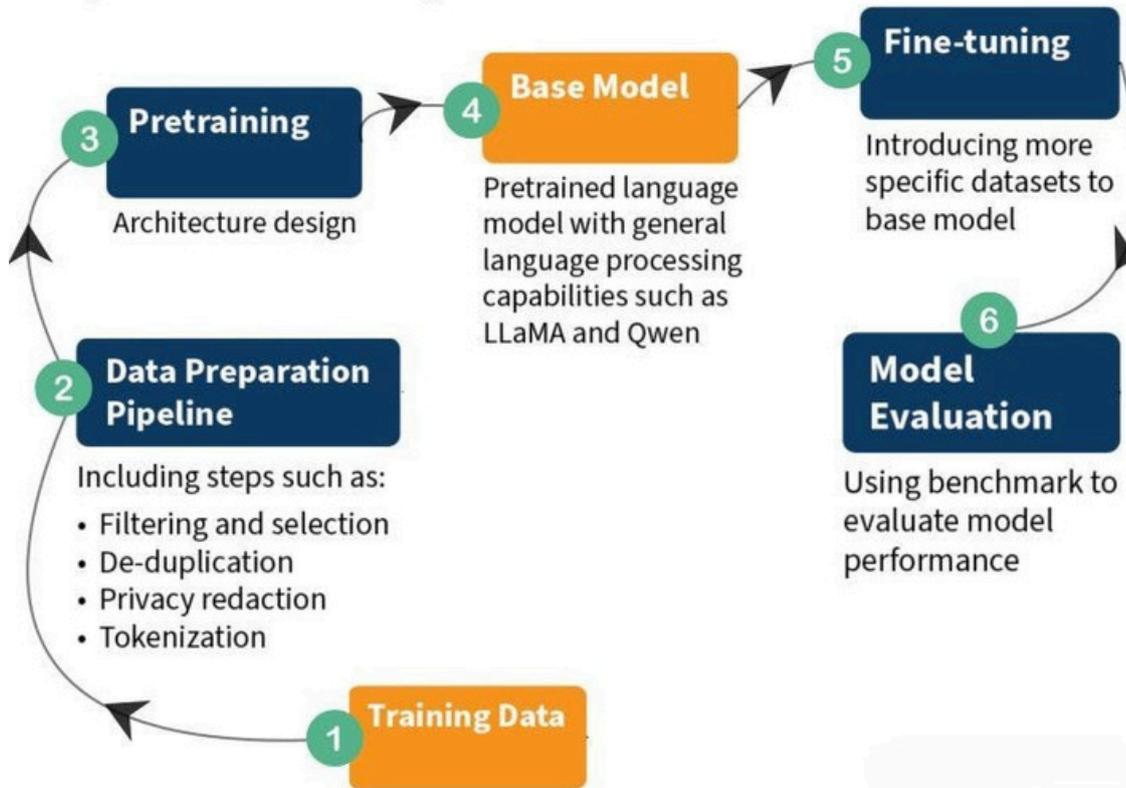


Figure 2 : How LLMs are Trained

Use Case	Key Benefit
Policy Drafting	Faster document creation with a consistent official tone
Legal & Official Letters	Improved grammar accuracy and standardized formatting
Citizen Services	Chatbots and multilingual support for public assistance
Data Analysis	Rapid summarization and trend extraction from large datasets
Training & Upskilling	Personalized learning and skill development content
Digital Transformation	Automation of repetitive documentation and workflows

How Governments Should Use LLMs: Simple Maturity Path

LLMs are trained on generic, global data. Government value comes from how they are used, not from building new models on day one. Risks (hallucination, bias, leakage) reduce as usage maturity increases.

LLM Usage Maturity (Low to High)

1. Better Prompting – Fastest Gains

Clear, structured prompts reduce errors and hallucinations

2. RAG (Context Injection) – Controlled Outputs

Uses official laws, policies, and documents as context. It can improve accuracy and reduce hallucination.

3. Fine-Tuning – Domain Specialization

Trains models on curated government data. It requires more financing and compute. Government departments will have to move towards fine tuning for optimal use of LLMs.

4. Sovereign LLMs – Strategic Endgame

Full control over data and compliance. It requires long-term investment in compute and talent.

Policy takeaway

Start with usage discipline, progress to context control, and pursue model ownership only when readiness exists. This is the strategic path that the government of Pakistan and provincial governments should adopt.

Designing AI Systems for Government Use

Retrieval-Augmented Generation (RAG)

Retrieval-Augmented Generation (RAG) is a method where an AI system retrieves information from approved documents or databases before generating a response.

Instead of relying on general knowledge, the AI:

1. Searches trusted government data sources
2. Retrieves relevant sections
3. Uses only that information to generate its response

This ensures outputs are grounded in official, up-to-date, and verifiable data.

Government Use Examples

- Querying service rules, acts, or regulations
- Answering internal questions using departmental SOPs
- Supporting officers with policy look-ups
- Creating AI assistants for ministries that only reference approved documents

Key Takeaway

RAG ensures AI systems do not “invent” answers, but instead respond based on government-approved knowledge.

Fine-Tuning

Fine-tuning is the process of training an AI model further on domain-specific data so that it consistently reflects a particular style, structure, or subject area.

While general AI models are trained on broad public data, fine-tuning adapts them to:

- Government language
- Sector-specific terminology
- Standard formats used in official work

It shapes how the AI responds, not what policy decisions are made.

Government Use Examples

- Generating standardized notices, letters, or reports
- Assisting with routine documentation in ministries
- Supporting sector-specific functions (health, finance, education)
- Creating internal AI tools trained on historical government documents

Key Takeaway

Fine-tuning helps AI systems sound and behave like a government assistant, not a generic chatbot.

Data Sovereignty, Confidentiality, and In-House AI Systems for Government

Why Data Sovereignty Matters in Government AI

Data sovereignty refers to the principle that government data must remain subject to national laws, regulatory oversight, and institutional control, regardless of how it is processed or analyzed.

For government entities, this includes:

- Citizen records
- Policy drafts and cabinet summaries
- Financial, security, and administrative data
- Internal communications and reports

When such data is processed by externally hosted AI systems, control over storage, usage, and retention may be lost.

Key Risks of Using Foreign-Hosted AI Tools

Using online, externally developed AI platforms for government work can introduce:

- Loss of control over sensitive data
- Unclear data retention and reuse policies
- Jurisdictional risks (data stored outside national boundaries)
- Exposure to model training on confidential inputs
- Non-compliance with national security and privacy laws

Even when tools claim not to store data, verification and auditability are limited.

Principle of In-House and Organization-Specific AI Systems

In-house AI systems are:

- Deployed on government-controlled infrastructure
- Accessed only by authorized personnel
- Trained or configured using approved internal documents
- Governed by existing IT, security, and data policies

Such systems function as internal digital assistants, not public AI services.

Core Design Principles

Government AI systems should follow these principles:

- Data Never Leaves the Organization
 - All documents, queries, and outputs remain within official systems.
- Use Only Approved Internal Sources
 - AI systems must reference:
 - Acts, rules, notifications
 - Departmental SOPs
 - Official reports and datasets
- No Training on External or Public Data by Default
 - The system is contextualized using internal documents only.
- Role-Based Access Control
 - Different officers see only the data relevant to their mandate.
- Human Oversight by Design
 - AI supports officers; final authority remains human.

Building Confidential AI Agents Using Internal Government Data

Step 1: Identify Approved Knowledge Sources

Before any AI system is built, departments must define:

- Which documents are approved for AI use
- Classification levels (public, internal, confidential)
- Update and version-control mechanisms

This ensures data governance precedes AI deployment.

Step 2: Use Retrieval-Based Approaches (Not Public Training)

Instead of uploading documents to public AI tools, organizations should:

- Store documents in secure internal repositories
- Use retrieval-based systems where AI fetches relevant information at runtime
- Prevent models from learning or memorizing confidential data

This approach keeps data contextual but not exposed.

Step 3: Develop Organization-Specific AI Agents

Each department can develop AI agents tailored to its function:

- Policy drafting assistants
- Legal and regulatory lookup systems
- Service delivery support tools
- Internal reporting and briefing assistants

Step 4: Developing Pakistan's Own Local and Domain-Specific LLMs

For true data and linguistic sovereignty, Pakistan should invest in the development of local large language models (LLMs) that:

- Are trained and hosted entirely within Pakistan
- Support local languages and administrative terminology
- Reflect national legal, regulatory, and cultural context
- Are aligned with public-sector use cases rather than consumer applications

Such models would not replace global research efforts but would provide trusted, sovereign AI foundations for government use.

Understanding AI Models for Government and Official Work

Modern governance increasingly integrates Large Language Models (LLMs) to support drafting, policy analysis, research synthesis, and structured communication. This section provides an overview of key AI models - ChatGPT, Grok, Gemini, Claude, and DeepSeek - outlining their core capabilities, practical applications, and known limitations in government and administrative settings.

The objective is to enable officers to make informed, task-specific choices when using AI tools, while maintaining appropriate human oversight and adhering to principles of accuracy, transparency, data security, and ethical responsibility. Used correctly, these models can enhance efficiency and decision support without replacing professional judgment or accountability.

Understanding Limits of LLMs : Bias, Context, and Role of Human Judgment

For Pakistan's public sector, it is essential to recognize that most LLMs are trained predominantly on English-language data and Western knowledge systems. As a result, their outputs may conflict with local culture, societal norms, legal frameworks, or state interests.

Techniques such as better prompt design, RAG, and fine-tuning can reduce risk but do not eliminate underlying bias or epistemic limitations. LLMs are probabilistic systems, not sources of truth; verification and human judgment must remain central and be explicitly embedded in policy design across education, health, and other critical sectors.

Overview of Different LLMs for Government Use

1. ChatGPT (OpenAI)

Overview:

ChatGPT is a widely used LLM developed by OpenAI, built on the GPT architecture. It is optimized for conversational interaction, analytical reasoning, and structured drafting tasks which are commonly required in government offices.

Capabilities:

- Drafting official reports, summaries, memos, and policy briefs
- Converting raw notes into structured Minutes of Meetings (MoMs)
- Preparing citizen-facing communication in multiple languages
- Analyzing text, extracting patterns, and generating recommendations

Limitations:

- May occasionally produce factual inaccuracies (“hallucinations”)
- Output quality depends heavily on clear and well-structured prompts
- Cannot access confidential or internal government databases
- Token limits and knowledge cut-offs may restrict long or updated content

Government Context Examples:

- Drafting policy briefs for ministries and attached departments
- Preparing disaster response notices in Urdu and English
- Structuring office orders, meeting agendas, and summaries
- Reviewing lengthy documents to extract key insights

2. Grok (xAI)

Overview:

Grok is an LLM developed by xAI, designed for speed, real-time information processing, and concise reasoning. It is suitable for fast-moving administrative and communication tasks.

Capabilities:

- Rapid summarization of lengthy documents
- Quick reasoning support for contingency planning
- Generating real-time insights for press statements or talking points
- Interpreting trends in fast-evolving situations

Limitations:

- Tone may be less formal than required for official documentation
- Limited suitability for legal or highly detailed policy drafting
- Not ideal for sensitive or neutral regulatory content

Government Context Examples:

- Preparing rapid talking points for senior leadership
- Summarizing district or field reports
- Drafting short public safety alerts (heatwaves, smog, floods)
- Supporting crisis communication units with quick situational updates

3. Gemini (Google)

Overview:

Gemini is Google's LLM designed for reasoning, drafting, and seamless integration with Google Workspace tools such as Docs, Sheets, and Gmail, making it useful for routine office workflows.

Capabilities:

- Drafting emails, memos, meeting notes, and briefs
- Assisting with data analysis and summaries using Google Sheets
- Translating or summarizing government notifications
- Comparing policy options and drafting evaluations

Limitations:

- Free versions offer limited output length and depth
- Requires internet connectivity and Workspace access
- Less suitable for very long or complex administrative documents

Government Context Examples:

- Drafting inter-departmental correspondence
- Summarizing spreadsheets (budgets, ADP progress, complaint logs)
- Preparing working papers for meetings
- Converting field reports into structured summaries

4. Claude (Anthropic)

Overview:

Claude is designed for accuracy, neutrality, and long-document handling, making it well-suited for detailed administrative, regulatory, and policy-related work.

Capabilities:

- Reading and summarizing long reports (50+ pages)
- Drafting structured inquiry reports and policy options
- Maintaining a formal, neutral, bureaucratic tone
- Supporting multi-step policy and regulatory analysis

Limitations:

- Slower than Grok for quick-response tasks
- May provide overly cautious responses
- Free versions may limit document uploads

Government Context Examples:

- Reviewing lengthy PC-I and PC-II documents
- Drafting inquiry and investigation reports
- Preparing legislative or regulatory briefs
- Conducting comparative policy analysis

5. DeepSeek (China/Global Open-Source Ecosystem)

Overview:

DeepSeek is a cost-efficient LLM designed for structured outputs, data processing, and large-scale usage, making it suitable for routine administrative automation.

Capabilities:

- Processing structured data and tables
- Generating standardized forms, templates, and checklists
- Automating repetitive drafting tasks (notices, certificates, formats)
- Producing rule-based and consistent outputs

Limitations:

- Less creative or nuanced than ChatGPT or Claude
- Requires carefully structured prompts
- Not suitable for high-level policy or strategic drafting

Government Context Examples:

- Creating standardized formats for office orders and certificates
- Processing large datasets (complaints, attendance, monitoring logs)
- Drafting routine public notices
- Developing templates for administrative workflows

Summary Table: AI Models for Government Use

Model	Strengths	Limitations	Best Government Uses
ChatGPT	Strong drafting ability, versatile, multi-lingual	Risk of factual errors; needs good prompts	Policy drafting, summaries, MoM, public notices
Grok	Fast reasoning, real-time insights	Tone less formal	Rapid updates, crisis comms, quick summaries
Gemini	Google Workspace integration	Free version limited	Email drafts, data summaries, workflow automation
Claude	Excellent with long/technical documents	Sometimes too cautious	Inquiry reports, legislative drafts, briefings
DeepSeek	Efficient, structured outputs	Less nuanced	Templates, notices, formats, dataset processing

MODULE 3 : INTRODUCTION TO PROMPT ENGINEERING

A prompt is the instruction or request given to an AI system to guide its response. Prompt engineering is the practice of designing these instructions in a structured and intentional way to achieve accurate, relevant, and reliable outputs.

Effective use of AI depends largely on how clearly this prompt is written. Prompt engineering refers to structuring instructions so the AI understands its role, the context of the task, and the expected output. For government officers, good prompting enables AI systems to produce precise, relevant, and context-appropriate results, especially when used for analysis, drafting, and decision support. This is especially useful for complex tasks such as policy drafting, analytical reviews, stakeholder communication, and structured decision support, where clarity and consistency are critical.

Large Language Models do not reason in the human sense. They generate responses by predicting likely patterns based on training data, which means they can produce outputs that sound confident but are incorrect or incomplete (hallucinations). Thoughtful prompting can reduce this risk by narrowing the scope, adding context, and requesting structured outputs, but it cannot eliminate hallucinations.

Effective use of AI depends on the quality of prompts. A clear definition of the AI's role, relevant context, expected output, and task clarity significantly improve results, especially when using AI agents in government work. While many prompt frameworks exist, no single framework will remain sufficient as LLMs continue to improve. Understanding the principles of good prompting matters more. The example provided on the next pages can be adapted and refined to improve productivity over time.

Anatomy of an Effective AI Prompt

Act as a senior policy analyst in the Ministry of Commerce.

Role

Develop a policy note for the Commerce Minister outlining the key advantages and disadvantages of bilateral free trade agreements.

Task

The policy note is required for internal briefing purposes ahead of a cabinet-level discussion. The intended audience is senior decision-makers who require a clear, balanced overview rather than technical economic detail.

Context

Explain the purpose of bilateral free trade agreements, outline three potential benefits for the national economy, and highlight three key risks or sectoral challenges, using a neutral, evidence-informed approach.

Reasoning

Maximum 300 words, formal and objective in tone, with a short introduction, bullet-pointed benefits, bullet-pointed risks and challenges, and a brief conclusion.

Output Format

Do not advocate for or oppose any agreement, avoid naming specific countries unless requested, and end with a summary paragraph.

Stop Conditions

SAMPLE PROMPT

You are a public health policy analyst supporting a district administration in Pakistan.

Use only authentic Government of Pakistan data, including population and demographic data from official sources such as the Pakistan Bureau of Statistics (PBS) and other publicly available government datasets. Do not rely on private surveys or estimates unless clearly stated as assumptions.

Analyze the district Mardan's population demographics with a focus on children under five years of age, including population size, gender distribution, urban–rural split, and any relevant socio-economic indicators available in government data.

Based on this analysis, identify priority areas and population segments for a child vaccination campaign. Highlight gaps, risks, and practical considerations for outreach.

Present the output as a short policy brief with:

1. Key demographic insights
2. Implications for vaccination coverage
3. Actionable recommendations for campaign design

Keep the language clear and suitable for senior government decision-makers.

KEY LESSONS FOR PROMPT ENGINEERING

A key lesson in prompt engineering is that clarity matters more than complexity. AI systems respond best when instructions are direct, specific, and free of ambiguity. Long or complicated prompts are not necessarily better; what matters is whether the task, purpose, and expected outcome are clearly stated.

One widely used approach is to define the role the AI should play, such as an analyst, policy drafter, or administrative assistant. Assigning a role helps the AI align its response with the user's intent and reduces irrelevant or unfocused output, which is especially important in government contexts.

Another important element is context. Providing background information, constraints, or reference material allows the AI to tailor its response to the policy, sector, or institutional setting. Without context, the AI may produce generic answers that are less useful for public-sector decision-making.

Clearly specifying the expected output is also critical. Stating whether the response should be a summary, a set of options, a comparison, or a draft policy note improves both accuracy and usability. This is particularly valuable when AI is used to support reviews, briefings, or internal reports.

Finally, users should recognize that prompt frameworks are guides, not fixed rules. As AI models continue to improve, rigid formulas will become less important than understanding core principles: clear intent, sufficient context, and iterative refinement. Effective prompting is an ongoing practice that improves with experimentation and feedback.

Ten Government-Centric Practice Exercises

1. Rewrite a vague prompt to address a specific economic issue in Punjab (e.g., agriculture productivity or SME development).
2. Draft a contextualized smog-related public safety alert as the Assistant Commissioner of Gujranwala.
3. Create a few-shot prompt to classify citizen complaints (Sanitation, Water Supply, Encroachment, Revenue, Health, Education).
4. Prepare a structured 150-word bullet-point summary of the Annual School Census Report with three actionable recommendations.
5. Refine a summary of polio campaign progress for a briefing to the Minister for Health.
6. Develop a chain-of-thought prompt outlining monsoon flood preparedness steps for District Khairpur.
7. Use role-based prompting to propose a three-point municipal waste management plan as a Deputy Commissioner.
8. Design a prompt chain for digitizing land records in Khyber Pakhtunkhwa, including a policy brief, press release summary, and bilingual awareness post.
9. Draft a neutral 100-word summary of a land dispute in South Punjab for submission to the Commissioner.
10. Complete a capstone by designing an AI-assisted workflow for a government task (e.g., PC-I summary, MoM, divisional review, or CMO briefing note).

MODULE 4 : BEGINNER HANDS-ON AI TOOLS FOR PRACTICAL APPLICATION

This module provides probationary officers with guided, beginner-friendly hands-on exposure to selected AI tools that demonstrate how Large Language Models can be applied in practical, low-risk, non-technical workflows.

The focus of this module is experiential learning—allowing participants to see, build, and interact with AI-powered systems without coding, while reinforcing responsible use, human oversight, and relevance to administrative work.

Module Objectives

- Understand how modern AI tools are used beyond chat interfaces
- Create structured documents using Claude Artifacts
- Build a simple no-code chatbot for informational use
- Design a basic automation for repetitive administrative tasks
- Understand the concept of local / on-premise AI models using Ollama

This module equips probationary officers with practical exposure to modern AI tools while reinforcing governance-first thinking.

Participants leave with a clear understanding of:

- What AI tools can do
- What they should not be used for
- How AI fits into structured administrative workflows rather than replacing them

Creating Structured Documents Using Claude Artifacts

To demonstrate how AI can be used to generate well-structured official documents such as briefs, notes, checklists, and summaries in a controlled format.

Tool Overview

Claude Artifacts allow users to generate and edit long-form structured outputs (documents, tables, templates) in a separate workspace, rather than conversational chat.

Hands-On Activity

Task: Create a structured briefing note using Claude Artifacts.

Participants will:

1. Open Claude and select the Artifact feature
2. Use a role-based prompt such as:
3. “Act as a Section Officer. Create a one-page briefing note on improving complaint redressal timelines in district offices.”
4. Generate an artifact structured as:
 - Title
 - Background
 - Key Issues
 - Recommendations
5. Edit and refine the document within the artifact workspace

Building a No-Code Chatbot Using Botsonic

To demonstrate how AI-powered chatbots can be created without coding for basic informational or internal support use cases.

Tool Overview

Botsonic is a no-code platform that allows users to build chatbots trained on uploaded documents or predefined information.

Hands-On Activity

Task: Build a simple informational chatbot.

Participants will:

1. Create a Botsonic account
2. Upload a small sample document (e.g., office procedures, FAQs, training notes)
3. Configure a chatbot to answer basic questions such as:
 - a. Office timings
 - b. Required documents
 - c. Process explanations
4. Test the chatbot by asking sample questions

Governance Note

- Not to upload confidential or sensitive data
- That such chatbots are suitable only for non-sensitive, informational use

Creating a Simple Automation Using Make.com

To introduce the concept of workflow automation, showing how AI and automation reduce repetitive administrative effort.

Tool Overview

Make.com is a no-code automation platform that connects apps and triggers actions based on events.

Hands-On Activity

Task: Build a basic automation workflow.

Participants will:

1. Create a simple scenario in Make.com
2. Connect two basic modules (trigger + action)
 - a. Trigger: New entry added to a Google Sheet (e.g., complaints log)
 - b. Action: Automatically generate a formatted summary using AI
 - c. Output: Send the summary via email or store it in a document
3. Run and test the automation

Understanding Local AI Models Using Ollama for On-Prem Use

To introduce the concept of local/on-premise AI, especially relevant for government environments with data sensitivity concerns.

Tool Overview

Ollama allows users to download and run small language models locally on their computers without internet dependency.

Hands-On Activity

Task: Download and run a small language model locally.

Participants will:

1. Observe a live demonstration of installing Ollama
2. Download a lightweight model (e.g., LLaMA-based model)
3. Run a basic prompt locally such as:
 - a. Summarize this paragraph in formal language.”
4. Compare local model behavior with cloud-based AI tools

Creating Custom AI Gems for Professional Productivity

This session introduces participants to the concept of custom AI Gems designed to support specific professional and administrative tasks.

Tool Overview

Google Gems are custom-configured AI assistants designed to perform specific, recurring professional tasks with consistency and structure.

Hands-On Activity

Task: Creating a “Meeting Minutes (MoM) Assistant” Gem

Participants will:

1. Open Google Gemini using the link <https://gemini.google.com/app>
2. Create a new Gem with the purpose of converting raw meeting notes into formal Minutes of Meeting.
3. Set the Gem’s role as a government documentation assistant using a neutral and factual tone.
4. Define a fixed structure for output: meeting details, discussion summary, decisions, and action items.
5. Test the Gem by pasting rough notes and review the generated MoMs before final use.

Data Analysis Using Julius

To demonstrate how officers can use a no-code AI data analysis tool to quickly extract insights from data without writing formulas, code, or queries.

Tool Overview

Julius is an AI-powered, no-code data analysis tool that allows users to upload datasets (CSV or Excel) and interact with them using plain language questions.

Hands-On Activity

Task: Analyze a small dataset and generate insights for a briefing note.

Participants will:

1. Upload a simple dataset (e.g., complaints data, attendance records, service delivery statistics) into Julius.
2. Ask Julius to explain what the dataset contains and identify key variables.
3. Ask questions such as:
 - a. “What are the main trends in this data?”
 - b. “Which category has the highest volume?”
 - c. “Are there any noticeable changes over time?”
4. Generate one chart or visual summary using natural language.
5. Ask Julius to produce a short, plain-language insight summary suitable for an internal briefing.

Annexure-General Suggestions

1. Optimize Prompt Structure:

- Use numbered lists or bullet points in prompts to enforce structured outputs (e.g.,
- “include 1) objective, 2) strategies”).
- Incorporate specific constraints (e.g., “use budget data”) to avoid generic responses.
- Experiment with temperature settings (if available, e.g., via API) to balance creativity and precision (low temperature for CoT, higher for creative summaries).

2. Maximize Efficiency:

- Minimize token usage by avoiding redundant context in chained prompts (e.g., reference prior outputs instead of repeating).
- Use batch testing to compare multiple prompt variations quickly (e.g., test CoT with/without step labels).
- Leverage Grok’s free access on grok.com or x.com for rapid prototyping, noting usage limits.

3. Enhance Effectiveness:

- Combine techniques (e.g., CoT within role-based prompts) for complex tasks like policy analysis.
- Include few-shot examples in prompts to guide output style (e.g., “format like this: [sample policy summary]”).
- Validate outputs against real-world data (e.g., cross-check with health or tourism reports).

1. Iterative Refinement:

- Analyze LLM outputs for biases or inaccuracies, especially for sensitive issues (e.g., community disputes).
- Use feedback loops in prompt chains to refine outputs iteratively (e.g., “revise based on stakeholder feedback”).
- Log prompt-output pairs to identify patterns and optimize future prompts.

2. Ethical Considerations:

- Ensure prompts avoid generating biased or harmful content, critical for diverse communities.
- Verify factual accuracy in outputs, especially for legal or financial recommendations, using official data sources.

Specialised AI Tools, Apps, and Websites for Different Professions

No	Profession / Domain	Representative AI Tools & Primary Use
1	Advertising & Marketing	Jasper (marketing content), Surfer SEO (content optimisation), Albert.ai (automated media buying)
2	Art & Design	Midjourney (text-to-image), Adobe Firefly (image & text effects), Invoke (custom model training)
3	Content Creation & Writing	Grammarly (grammar & style), Jasper (idea generation), Hemingway App (readability analysis)
4	Education	Google AI Essentials (education tools), Khanmigo (AI tutoring), Slides AI (presentation generation)
5	Engineering	GitHub Copilot (code suggestions), Lovable (coding assistant), Adept (AI software task)
6	Finance & Banking	IBM Watsonx (risk & fraud detection), Daloopa (financial data extraction), Tesseract (investment)
7	Healthcare	Google AI (Breast Cancer Detection), Babylon Health (symptom analysis), Tempus (genomic)
8	Human Resources	Paradox (AI screening), Jobscan (resume optimisation), Talentsoft (employee learning)
9	Legal	Harvey AI (legal research), Casetext (case & statute search), Evisort (contract review)
10	Logistics & Supply Chain	FedEx AI (routing), Freightos (freight pricing), Locus (route optimisation)
11	Manufacturing	Siemens AI (predictive maintenance), Augury (failure prediction), Seebo (quality prevention)
12	Media & Entertainment	Descript (text-based audio/video editing), Runway (AI video), ElevenLabs (voiceovers)
13	Project Management	Asana Intelligence (task management), Motion (auto scheduling), Reclaim AI (focus optimisation)
14	Real Estate	Zillow AI (home recommendations), Realtor.com AI (market insights), Restb.ai (property image)

No.	Profession / Domain	Representative AI Tools & Primary Use
15	Retail & E-commerce	Amazon AI (product recommendations), Dynamic Yield (personalisation), Criteo (targeted ads)
16	Sales	Relevance AI (lead generation), Salesforce Einstein (predictive scoring), Gong (sales-call analytics)
17	Science & Research	Perplexity (cited answers), Scite (paper evaluation), Elicit (research summarisation)
18	Software Development	GitHub Copilot, Amazon CodeWhisperer, Tabnine (code completion)
19	Customer Service	Intercom (chatbots), Zendesk (ticket routing), Drift (visitor engagement)
20	Data Science & Analytics	Numerous.ai (spreadsheet analytics), Tableau AI (visualisation), DataRobot (AutoML)
21	Photography	PhotoRoom (background removal), Luminar Neo (AI photo enhancement), Remove.bg (background removal)
22	Audio Production	Descript (dialogue editing), LALAL.AI (vocal separation), Altered Studio (voice modification)
23	Human Resources & Recruiting	Jobscan (resume analysis), Eightfold AI (talent matching), HiredScore (applicant ranking)
24	Administrative & Productivity	Notion AI (summaries & organisation), Microsoft Copilot (M365 assistance), Zapier (workflow automation)



The Artificial Intelligence 101 for Public Sector program is a joint initiative by the Ministry of Information Technology and Telecommunication, Ministry of Planning, Development and Special Initiatives, Civil Services Academy, and atomcamp, aimed at equipping Civil Service Academy probationers with essential AI awareness for effective governance and policy-making.

For further information, please contact the Civil Services Academy, Walter Lahore