

AI 101: Beyond the Chatbot

A plain-language field guide for manufacturing, engineering, defense, and public leaders.



AI is not magic. It is software that learns patterns from data. A **neural network** learns those patterns from examples. A **large language model (LLM)** predicts useful sequences of words. An **agent** puts a model inside a **harness** so it can pursue a bounded goal.

Token: a word or word piece.

Your prompt becomes input tokens, the reply becomes output tokens, and both fill the context window.

How we arrived here: rules, machine learning, neural networks, LLMs, agents.

Five building blocks of a modern agent

1 The model (the brain)

Interprets input and proposes the next step.

2 Memory

Supplies approved facts and prior state.

3 Tools (the hands)

Search, calculate, call software, or act.

4 The loop

Observe, plan, act, check, then repeat or stop.

5 Guardrails (the leash)

Limit access, require approvals, log actions, and stop unsafe work.

Harness: the software around the model that connects memory, tools, the loop, and guardrails.

Useful first workflows

- Search approved manuals, procedures, and lessons learned.
- Draft non-final work instructions, reports, or supplier messages.
- Triage maintenance, quality, and supplier signals for expert review.
- Compare requirements while preserving sources and traceability.

Five questions before any pilot

- ☐ **Outcome:** What useful result will we measure?
- ☐ **Boundary:** What may the system see and do?
- ☐ **Evidence:** How will we know when it is wrong?
- ☐ **Authority:** Which decisions stay with a person?
- ☐ **Transition:** How will affected people shape the change?

Take-home exercise: Name one workflow that AI could improve and one boundary it must never cross.

The risks are real and manageable

Challenge	Practical control
Hallucinations	Use approved sources, citations, verification, abstention, and human approval.
Hacking and misuse	Use least privilege, isolation, attack testing, monitoring, logs, and fast shutdown.
Sensitive data	Keep Controlled Unclassified Information (CUI), export-controlled data, and IP in approved environments.
Bias and accountability	Name an owner, test real cases, preserve evidence, and allow appeal.
Job displacement	Map changing tasks, retrain early, redesign roles, and share gains deliberately.
Over-automation	Set action limits, stop conditions, escalation paths, and rollback.

Three boundaries to set this week

1. No CUI or export-controlled data in public AI tools.
2. No agent connected to production, enterprise resource planning (ERP), or payment systems without written approval.
3. No AI-drafted compliance or contract submission released without human verification.

CMMC decision rule: AI touching CUI belongs only in Cybersecurity Maturity Model Certification (CMMC)-compliant, contract-approved environments.

For public leaders: Make retraining concrete. Partner with your local workforce agency and community college to make it real.



Recommended viewing

"Humans Need Not Apply" by CGP Grey (2014). Scan the code or use the link. The video is a useful provocation about automation and work.

<https://www.youtube.com/watch?v=CMFj75kBJIU>

"We are the horses."

Technology need not replace every human ability to transform the labor market. Prepare people and institutions instead of assuming helplessness.

Discuss: Which of our tasks look most like the horses' jobs, and what would we do about it this quarter?

Humans adapt. Leaders can smooth the transition by protecting what matters, teaching what changes, and shaping AI before disruption shapes them.

AI is moving quickly. What applies today may not apply in a year. Schedule regular refresher courses and update policy and controls. Waiting raises the risk of avoidable job loss and economic recession.

Questions, concerns, or AI policy help?

Tristen Pierson, Axiovox Systems
start@axiovoxsystems.com

<https://axiovoxsystems.com>

<https://www.linkedin.com/company/axiovoxsystems>

Course materials



Course slides (PDF)

<https://axiovox.github.io/shared-documents/docs/ai-101-beyond-the-chatbot-slides.pdf>

This handout (PDF)

<https://axiovox.github.io/shared-documents/docs/ai-101-handout.pdf>

Plain-Language Glossary

Quick definitions for leaders working across manufacturing, engineering, defense, and public service.



Use these definitions as a shared starting point. In practice, the exact meaning may depend on the system, contract, or policy involved.

AI BASICS

Artificial intelligence (AI)

Software that performs tasks that normally require human judgment, such as recognizing patterns or creating text.

Machine learning

A way of building AI by training software on examples so it can make predictions or spot patterns in new data.

Neural network

A machine-learning system made of connected layers that adjust as it learns from examples.

Large language model (LLM)

An AI model trained on large amounts of text to understand and generate language.

Token

A small piece of text, roughly a word fragment. The model reads everything you send as input tokens and writes its reply as output tokens. Input tokens include your instructions and any documents you paste in. Output tokens are everything the model writes back. Most AI services charge per token, and output tokens usually cost more than input tokens, so long documents and long answers drive the bill.

Context (context window)

The information an AI can consider at one time. A context window is the limit on how much fits.

Prompt

The instruction and information given to an AI to guide its response.

RISKS AND CONTROLS

Hallucination

A confident-sounding answer that is false or unsupported.

Prompt injection

Hidden or misleading text that tries to make an AI ignore its instructions.

Least privilege

Give each person or system only the access needed for the task, and no more.

Pilot

A small, bounded trial of an AI workflow with a clear measure of success, used to learn safely before a wider rollout.

AGENTS

Agentic

Describes AI software that can act toward a goal with some independence, like an AI agent does.

AI agent

AI software that can pursue a goal through several steps, rather than only answer once.

Harness

The software around a model that provides memory, tools, rules, and approvals.

Memory (agent memory)

Saved information the agent can use later, such as approved facts or the current state of a task.

Tools (agent tools)

Approved ways an agent can search, calculate, read a system, or take an action.

The loop (observe, plan, act, check)

The agent observes, plans, acts, and checks the result. It repeats until done or told to stop.

Guardrails

Rules and technical limits that control what an AI may see, do, or send.

Abstention

When an AI says it does not have enough evidence instead of guessing.

DEFENSE AND COMPLIANCE

CUI (Controlled Unclassified Information)

Information the U.S. government requires organizations to safeguard even though it is not classified.

CMMC (Cybersecurity Maturity Model Certification)

A U.S. Department of Defense program that checks whether contractors use required cybersecurity practices to protect sensitive information.

Export-controlled data

Technical or defense information that law limits from being shared with certain people, organizations, or countries.

ERP (enterprise resource planning)

Business software that connects core operations such as orders, inventory, production, purchasing, and finance.

On-premises (on-prem)

Running software on computers your organization owns and controls, rather than on a vendor's cloud.