



AI 101: Beyond the Chatbot





Four waves made today's agents possible



Reference milestones: Dartmouth workshop, 1956; deep-learning breakthrough, 2012; Transformer, 2017; public chat assistants, 2022.



AI is not magic. It is a stack we can inspect.



Once we can name the parts, we can test them, set boundaries, and make better decisions.



A neural network learns by adjusting many tiny connections

Inputs

Pixels
Sensor readings
Words



Outputs

Defect probability
Failure risk
Next token

Training


Show examples, measure the error, and adjust numerical weights until results improve.

Using the model

Freeze the learned weights, provide a new input, and calculate an output.



A large language model predicts the next piece of text



What you send in counts as input tokens; what the model writes back counts as output tokens. Most services charge per token, and output usually costs more.

Why it feels intelligent

At enormous scale, prediction captures language, concepts, patterns, and useful procedures.

Why fluency can mislead

Its basic job is a plausible continuation, not a built-in guarantee of truth or intent.



The harness turns a model into an agent

Model

The brain: proposes what to do next.

Memory

Keeps useful facts and task state.

Tools

The hands: reads, writes, searches, or operates.

Loop

Observe, plan, act, check.

Guardrails

The leash: permissions, approvals, limits, and stops.

The model suggests. The harness decides what it can see, do, remember, and when it must stop.



The loop repeats four simple moves



Finish

The result meets the stated goal.

Ask a human

The evidence is weak, the case is unusual, or approval is required.

Fail safely

A guardrail stops the action before harm spreads.



Agentic AI is a dial, not an on-off switch

Suggest

A human performs every action.

Prepare

The agent drafts; a human approves.

Execute within bounds

The agent acts under policy and limits.

Operate and escalate

The agent handles routine cases and stops on exceptions.

Increase autonomy only as evidence, guardrails, and reversibility increase.



Start where errors are reversible and evidence is available

Good first use

Frequent task, measurable result, trusted data, limited harm if it fails, clear human owner.

Bad first use

Irreversible decision, controlled-data ambiguity, safety-critical action, no audit trail, unclear accountability.

Manufacturing value starts in bounded workflows



Good starting points

- Quality evidence triage
- Predictive maintenance support
- Engineering knowledge retrieval
- Supplier and schedule risk analysis
- Documentation and compliance preparation

**Budget more time for organizing your data than for the AI itself.
Pilots die on messy data.**



There are three sensible ways to adopt AI

1

Approved software service

Cost and effort: lower setup, faster start.

Require enterprise data protections, access controls, and vendor review.

2

Private deployment on your own systems

Cost and effort: higher.

Offers more control for sensitive data, but needs internal skills, security, and maintenance.

3

Deliberate waiting

Cost and effort: low technology spend.

Build literacy, clean data, and policy now so waiting does not become falling behind.



CMMC makes the data boundary an executive decision

Decision rule

AI tooling that touches CUI must live in CMMC-compliant, contract-approved environments.

1 Classify first

Know whether the data is CUI, export-controlled, proprietary, personal, or public.

2 Approve the environment

Match the tool, storage, model, access, and logs to contract requirements.

Pasting controlled technical data into a public chatbot can jeopardize contract eligibility.

CMMC means Cybersecurity Maturity Model Certification. Use current contract, cybersecurity, export-control, and records requirements. This class is not legal or compliance advice.



Three boundaries to set this week



No CUI or export-controlled data in public AI tools.



No agent connected to production, ERP, or payment systems without written approval.



No AI-drafted compliance or contract submissions released without human verification.



Hallucination is a system-design problem

Why it happens

- Plausibility is not truth
- Context may be missing or conflicting
- The model may guess instead of saying "I don't know"
- Confident wording does not mean the answer is right

How to contain it

- Ground answers in approved sources
- Require citations and repeatable checks
- Test on real failure cases
- Allow "I don't know"
- Keep humans on consequential decisions



AI makes old fraud faster and more convincing

Instructions can be attacked

Malicious text can try to override an agent's instructions or make it leak data.

Executives can be impersonated

A cloned CEO voice authorizes a wire. A convincing fake supplier email changes payment details.

Connected tools raise the stakes

Broad permissions can turn one bad instruction into a real operational or financial action.

Control: verify every payment or credential change through a separate trusted channel.

Also use least privilege, approvals, isolation, logs, limits, and a kill switch. Security framing: OWASP Top 10 for LLM Applications 2025.

Watch this 2014 warning: we are the horses

1915



2014

CGP Grey, “Humans Need Not Apply”

Mechanical muscle reduced the economic role of horses. Mechanical minds can reduce demand for human cognitive labor. The warning has held up because capability keeps getting cheaper, broader, and easier to deploy.

<https://www.youtube.com/watch?v=CMFj75kBQ1U>



Scan to watch

The lesson is not panic. It is to adapt while we still have choices.



Make adaptation practical for workers and small manufacturers

For employers

- Map tasks before eliminating roles
- Train domain experts to supervise AI
- Create mobility paths before disruption
- Measure safety, quality, autonomy, and who benefits

For public leaders

- Fund shared training, secure test environments, and cyber readiness
- Ask companies for data boundaries, human accountability, and a workforce plan
- Keep small manufacturers included through pooled expertise and practical pilots
- Build retraining through local workforce agencies and community colleges



At your table: name one workflow and one boundary

One workflow

Where could AI remove friction or improve a measurable result?

One boundary

What must the AI never see, decide, or do on its own?

The full design worksheet belongs in the take-home handout.



Open Q&A

Manufacturing. Engineering. Defense. Policy. Security. Workforce.

What decision do you need to make next?

One commitment per table: one workflow, one boundary

Workflow

Where will you explore value?

Boundary

What will you protect from the start?

We must be proactive rather than reactive, or jobs will be lost and the economy will recess.

Keep learning: AI is moving very fast. What applies today may not apply a year from now. Regular follow-up courses can keep this community current.

Questions, concerns, or help writing AI policy? Reach out anytime:

Tristen Pierson, Axiovex Systems

start@axiovetxsystems.com

<https://axiovetxsystems.com>

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

