

CLEARED FOR AI

Module 1 Workbook

Ground School: What AI Is and Isn't

About 86 minutes · 7 lessons · exercises · logbook · quiz key

Name: _____ Team: _____ Date: _____

Prepared with AI assistance and reviewed by Nils Roed, airline captain. Educational material only. No specific productivity, savings or financial results are promised. Use only non-sensitive data in exercises. Sources: Mollick, Anthropic, OpenAI, NIST, Prosci (see the module page). © 2026 Cleared for AI.

Lesson 1: The Automation Paradox (~12 min)

Key points

- Capable automation shifts the human job from doing to briefing, monitoring and intervening.
- Complacency is a predictable response to a usually-right system. Plan for it rather than blaming people.
- Brief it, monitor it, be ready to take over.
- Accountability stays with the person and company that send the work.

Exercise

Exercise (about 8 minutes). Open Grok or any AI chat tool. Paste in a short, made-up set of notes for a small job (or use this: *"Replace kitchen faucet, customer supplies fixture, also check slow drain in bathroom sink, 2 hours labor at \$95/hr, trip charge \$45"*). Ask the AI to turn it into a customer-ready estimate. Then play pilot monitoring: compare every line with your notes. Did it add anything you didn't say? Leave anything out? Do the maths? Write down one thing you'd have missed if you hadn't checked.

My notes and results

Green / Amber / Red?

☐ Green ☐ Amber ☐ Red Why: _____

Lesson 2: What AI Actually Is (and Isn't) (~10 min)

Key points

- Language models predict likely text. They're fluent drafters, not sources of truth.
- Fluency isn't accuracy. Confident wording proves nothing.
- Give the AI your real information and tell it to use only that.
- Chat, then workflow, then agent: more autonomy means more guardrails.

Exercise

Exercise (about 6 minutes). In your AI tool, ask: *"What are the business hours and return policy for [your company name]?"* Notice whether it admits it doesn't know or produces something plausible. Next, paste in a few real facts (or made-up ones for practice) and ask again, adding: *"Use only the information I provided. If something isn't covered, say so."* Compare the two answers and note how context changed the output.

My notes and results

Green / Amber / Red?

☐ Green ☐ Amber ☐ Red Why: _____

Lesson 3: The Jagged Frontier (~11 min)

Key points

- AI's capability edge is jagged: excellent and poor tasks can look identical.
- Outside the frontier, AI fails convincingly, and people tend to accept persuasive wrong answers.
- Map your team's tasks green, amber or red with real examples, and revisit the map often.
- The higher the stakes, the stronger the check.

Exercise

Exercise (about 7 minutes). Pick three small tasks from your real work and run each through your AI tool using non-sensitive data where you know the right answer. Suggestions: (1) rewrite a terse email in a warmer tone; (2) total a list of 12 numbers and apply 8% tax; (3) answer a factual question about a recent event in your industry. Grade each green, amber or red. Did any red result *look* convincing? That's the frontier.

My notes and results

Green / Amber / Red?

☐ Green ☐ Amber ☐ Red Why: _____

Lesson 4: Confident and Wrong: Hallucinations (~11 min)

Key points

- Hallucinations are plausible but false outputs, including invented sources, wrong specifics and quiet omissions.
- Check every specific, open every link, and compare important summaries with the source.
- Brief to reduce errors: give source material, allow 'I don't know,' and break tasks into checkable steps.
- A second AI isn't independent verification. Check against a trusted source.

Exercise

Exercise (about 7 minutes). Ask your AI tool: *"Give me three statistics, with sources, about small-business adoption of AI."* Then verify each one: open the source, find the number. Record how many you could fully confirm. Next, ask the same question but add: *"Only include statistics you can link to a specific, real source. If you're unsure, say so."* Compare the results. Whatever happens, the point stands: the check is yours to do.

My notes and results

Green / Amber / Red?

☐ Green ☐ Amber ☐ Red Why: _____

Lesson 5: Learning Your Aircraft: Hands-On Hours (~14 min)

Key points

- Competence comes from hands-on hours on real tasks, not from reading.
- Knowledge isn't ability (Prosci ADKAR). Practice closes the gap.
- Use a ground → simulator → supervised line → debrief structure over about a month.
- Converse with the AI rather than asking one-shot questions, and keep a team logbook.

Exercise

Exercise (about 10 minutes). Start your AI logbook now. Make a simple table with columns: *Date · Task · What I asked · Result · Green/Amber/Red · Better prompt next time*. Pick one recurring task from your job and try it in your AI tool. Then hold a short conversation: ask for two alternatives, give feedback ("more concise," "friendlier," "add a call to action"), and ask the AI to critique its own final draft. Log what changed between the first answer and the last.

My notes and results

Green / Amber / Red?

☐ Green ☐ Amber ☐ Red Why: _____

Lesson 6: Where Your Data Goes (~10 min)

Key points

- Data shared with an outside tool can't be unshared. Treat it as cargo you must account for.
- Read the terms for your specific tool and plan. Business plans often differ from consumer ones.
- Information traffic lights: green is fine, amber is approved tools only, red is don't paste. When in doubt, strip it out.
- Prompt injection: keep humans approving any AI action that sends, pays, deletes or shares.

Exercise

Exercise (about 6 minutes). Open the settings and the privacy or data-use page of the AI tool you use most. Find out: (1) whether your conversations are stored and for how long; (2) whether they may be used to train models, and how to opt out; (3) what's connected (email, drive, calendar, browsing). Write the answers in your logbook. Then take a realistic work prompt and rewrite it with every identifying detail stripped out. Does the AI still give a useful answer?

My notes and results

Green / Amber / Red?

☐ Green ☐ Amber ☐ Red Why: _____

Lesson 7: Train Judgment, Not Buttons (~10 min)

Key points

- Assume today's AI is the weakest you'll use. Plan for constant change.
- Tool-specific skills fade in months (Prosci). Invest in durable judgment.
- Durable skills: task sorting, briefing, verification, data discipline, escalation.
- Schedule recurrent training and share near misses without blame. Leaders reinforce the habits.

Exercise

Exercise (about 6 minutes). Write your team's first "near miss" report, real or hypothetical, in four lines: *What was the task? What did the AI produce? What almost went wrong? What check caught it (or would have)?* Then ask your AI tool to help you design a 15-minute monthly team segment for sharing AI wins and near misses. Edit its suggestion into something you'd actually run, and note what you changed and why.

My notes and results

Green / Amber / Red?

☐ Green ☐ Amber ☐ Red Why: _____

[illegible]

Quiz answer key

1. **Why does capable automation make the human supervisor more important, not less?**
Answer: Because the human's job shifts to briefing, monitoring and stepping in, and the rare failures matter most. *The automation paradox: capability concentrates judgment. The rare unexpected moment is exactly when an alert crew matters most.*
2. **A large language model answers questions mainly by...**
Answer: Predicting likely text based on patterns from training. *Language models generate likely-sounding text. That's why they're fluent, and why fluency isn't proof of accuracy.*
3. **What does the 'jagged frontier' describe?**
Answer: AI's strengths and weaknesses are irregular and hard to predict, even between similar tasks. *Mollick's term: two tasks can look alike and sit on opposite sides of what AI does well.*
4. **In the consulting-firm study, what happened on the task outside AI's capabilities?**
Answer: AI users were more likely to reach the wrong answer. *Persuasive wrong output led many AI users astray. AI fails convincingly, not obviously.*
5. **Which AI error is usually hardest to catch?**
Answer: A quiet omission of an important detail. *Omissions only show up when you compare the output with the source, which is why important summaries get compared line by line.*
6. **Asking a second AI to check the first AI's work is...**
Answer: Useful as a brainstorm, but not independent. Check against a trusted source. *Like cross-checking instruments fed by the same faulty sensor: both AIs can share blind spots.*
7. **In Prosci's ADKAR model, which gap do hands-on hours mainly close?**
Answer: Knowledge → Ability. *Understanding how isn't the same as being able to do it in real work. Practice builds ability.*
8. **Under the information traffic-light rule, a customer's account number is...**
Answer: Red: don't paste unless the tool is specifically approved for it. *Personal and financial identifiers are red. When in doubt, strip it out.*
9. **What is prompt injection?**
Answer: Hidden instructions in outside content that try to make an AI do something unintended. *That's why humans should approve any AI action that sends, pays, deletes or shares, especially when the AI reads outside content.*
10. **Why train 'judgment, not buttons'?**
Answer: Tool-specific skills go stale in months, while briefing, verification and data discipline transfer between tools. *Prosci describes a 3–4 month half-life for tool-specific AI skills. Durable judgment survives tool changes.*