Session 1 Handout: What AI Is Doing When It Answers You

FAST-01 Applied AI Literacy at Work · 7formations · Instructor: Wadood Chaudhary · All examples fictional

Why this session matters

Most people meet generative AI as a box that answers questions. It writes quickly, it sounds sure of itself, and it is often right, which is exactly why it is easy to trust too much. This session gives you a working mental model of what the tool is doing, shows you where AI already sits inside the software you use every day, and ends with the first draft of your Personal AI Use Map: a task-by-task plan for where AI fits in your job.

Key concepts

1. Generative AI predicts; it does not look things up

A large language model (LLM), the engine inside Microsoft 365 Copilot, ChatGPT, Gemini, and Claude, was trained on a huge amount of text to do one thing very well: predict which words are likely to come next. When you ask a question, it builds an answer a small piece at a time, each piece chosen because it fits the pattern of the words before it.

That has two consequences you will use every day:

- **Fluent is not the same as correct.** The model is excellent at producing text that *sounds* like a good answer. Whether the facts are right depends on what it learned, what you supplied, and some luck.
- **It sounds just as sure when it is guessing.** There is no built-in warning light. A made-up shuttle time is written in the same confident tone as a correct one.

Some tools now search the web or your company files before answering ("grounding"). That helps a great deal, but the reply is still generated text. The tool can misread a source, blend two sources, or fill a gap with something plausible. Grounding reduces errors; it does not remove the need to check.

2. Why AI answers vary

Ask the same question twice and you may get two different answers. The tool picks among several likely next words, with a little randomness, so the wording, and sometimes the facts, can change. For writing a welcome note, variety is useful. For a refund rule, it is a reminder to check every time.

3. AI is already inside your tools

You may be using AI without opening a chat window:

Where	What the AI does	What it can see
Email (Outlook, Gmail)	Suggested replies, thread summaries, tone rewrites	The messages you open
Phones and voicemail	Transcribes calls and voicemail, autocorrects	What callers say, including numbers read aloud

Where	What the AI does	What it can see
Documents and spreadsheets	Copilot or Gemini side panels that draft, summarize, or analyze	The file, and with Copilot, other files you have access to
Customer chat and help desks	Suggested answers, auto-summaries of tickets	Customer messages and history
Scheduling tools	Suggested meeting times and follow-ups	Calendars and attendees

The same data rules apply to built-in features as to chat tools. If it would be wrong to paste a customer's card number into a chatbot, it is wrong to have an AI feature summarize the voicemail where they read it out, unless your employer's policy explicitly covers that use.

4. The DOL AI Literacy Framework is the floor

In February 2026 the U.S. Department of Labor issued its AI Literacy Framework (Training and Employment Notice 07-25). It describes basic content areas: understanding AI principles, exploring uses, directing AI, evaluating outputs, and using AI responsibly. The Massachusetts FAST AI program treats that as the baseline. This course goes past it: you apply each idea to your own tasks, in your employer's approved tool, and prove it with a checked work product.

5. Three piles: good fit, use with care, no AI

Every task you do can be sorted into one of three piles. Ask three questions:

1. **What data does it need?** If the task requires restricted data (personal information, PHI, confidential business data), it is "no AI" unless policy and the specific tool allow it.
2. **Can I check the result?** If you have no source to check against and no way to tell right from wrong, it is "no AI."
3. **Who is harmed if it's wrong?** If a customer, patient, or coworker will rely on it, it is at least "use with care."

Pile	Typical tasks	What it requires
Good fit	Turning your own notes into a checklist; brainstorming; summarizing a public document for yourself	Normal reading before use
Use with care	Anything a customer will read; summaries others rely on; safety reminders from an SOP	A named check against a named source, and a named reviewer for high-stakes items
No AI	Tasks needing restricted data; final decisions about people, money, or safety; anything policy forbids; anything you can't check	A person does it; AI may help with a sanitized piece if policy allows

Worked example: rating front-desk tasks at Cape Harbor Inn

Cape Harbor Inn (fictional, Chatham) has a front-desk associate who spends six hours a week answering guest emails about check-in, parking, pets, and the ferry shuttle.

- **Data:** the answers come from the inn's approved FAQ. No card numbers are needed.
- **Checkable?** Yes: every fact is in the FAQ.
- **Harm if wrong?** Medium. A wrong shuttle time means a missed ferry; a wrong pet rule means a dog turned away at the desk.

Rating: use with care. Check: compare every fact with the FAQ before sending.

Compare "take card payments over the phone for deposits." It needs card numbers, which are personal information under Massachusetts law when combined with a name. **Rating: no AI.**

And "write the weekly 'what's on in Chatham' note" uses public event information and is low harm. **Rating: good fit.**

Checklist: before you rate a task

- ☐ I described the task in general terms (no names or numbers).
- ☐ I know what data the task needs.
- ☐ I know what I would check the result against.
- ☐ I know who relies on the result and what happens if it's wrong.
- ☐ My reason mentions data, checkability, or harm, not just "saves time."

Do and don't

Do	Don't
Treat every AI answer as a draft until a person checks it	Treat a confident tone as evidence
Rate tasks one by one, even within the same job	Decide "AI is fine for my job" or "AI is useless for my job" overall
Use your employer's approved tool with your work account	Use a personal account or a free tool for work content
Notice the AI already built into your email, phone, and documents	Assume built-in features don't count
Ask "what data, what check, what harm?"	Rate a task as a good fit only because it saves time

Try this at work

1. Write down five tasks from your last workday, in general terms.
2. For each, answer the three questions: data, checkability, harm.
3. Put each in a pile. Show the list to a coworker who does the same job and ask which rating they would change.
4. Pick one "good fit" task and try it in your approved tool this week with non-sensitive text. Note one thing it did well and one thing you had to fix.

Glossary

Term	Plain-language meaning
Generative AI	Software that creates new text, images, or audio based on patterns learned from examples
Large language model (LLM)	The type of generative AI behind chat assistants; predicts the next likely words
Prompt	What you type or paste to tell the AI what to do

Term	Plain-language meaning
Output	What the AI produces in response
Grounding	Giving the AI a specific source (a document, a search result) to base its answer on
Hallucination	A confident, plausible statement that is false or not supported by the source
Approved tool	An AI tool your employer has reviewed and permitted for work, used with a work account
AI Use Map	Your task-by-task plan rating each task good fit, use with care, or no AI

Key takeaways

- AI writes what is likely, not what it has verified.
- AI already sits in many of your tools; the same rules apply.
- Rate tasks one at a time using data, checkability, and harm.

Naming note (October 2026): Microsoft's documentation now titles Microsoft 365 Copilot and Copilot Chat as "Microsoft Copilot" and "Microsoft Copilot Chat"; your screens may show either name. Tool names are trademarks of their owners; no endorsement implied. Tool features change frequently; descriptions reflect vendor documentation as of October 2026.