

Ourbit Master AI - Day 11: Responsible AI, Privacy, Security and Data Judgment

Purpose: Make proportionate decisions about data, permissions, consequences and human oversight.

Source/build status: New gap.

Learning outcomes

- Apply data minimization before sharing information with an AI system.
- Distinguish permission to access data from permission to disclose or process it elsewhere.
- Use least privilege and consequence-based human oversight.
- Recognize when uncertainty about authorization requires stopping rather than proceeding.

Core lesson - Responsible AI, Privacy, Security and Data Judgment

Responsible AI use begins before a prompt is sent. Ask what data is actually needed, whether you are authorized to use it, what system it is entering, what action may follow, and who remains accountable. Apply data minimization, least privilege and consequence-based human oversight. Convenience is not permission.

Worked example

An employee wants to paste a customer spreadsheet into an AI service for summarization. First classify the information and confirm the service is approved for that use. Remove unnecessary identifiers, share only what the task needs, restrict access, and keep an accountable human responsible for the result. If authorization is unclear, do not upload the data.

Teaching / video direction

Visually reduce a fictional customer dataset from many fields to only the fields needed for a defined task. Overlay permission and approval gates. Do not use real personal data. Close by showing that responsible research also requires traceable evidence-Day 12.

Hands-on practical exercise

- Use a fictional dataset only; list ten fields a customer record might contain.
- Define a simple summarization task and mark which fields are actually necessary.
- Remove unnecessary identifiers and sensitive fields from the proposed input.
- State what authorization and system approval would be required before real data could be used.
- Identify who remains accountable for checking the output and any action based on it.
- Write a stop condition for the case where authorization or data handling rules are unclear.

Knowledge check

- What does data minimization mean in practice?
- Why does access to information not automatically permit uploading it to another service?
- What is least privilege?
- When should human oversight become stronger?
- What should happen when authorization is unclear?
- How does responsible data judgment build on Day 10 evaluation?

Assessment and remediation contract

The scored daily quiz uses 10 questions; 7/10 (70%) is a pass. Missed concepts map to exact lesson anchors for targeted remediation, including on a passing attempt. Attempts 1-9 below threshold return to remediation and a fresh eligible question set; an unsuccessful tenth attempt restarts the day in a new learning cycle while preserving history. The 50-question validated bank is a learner-release gate and will be scaled only after the Day 1 runtime passes acceptance.

Teaching / video script - 6-8 minute target

Open: Start with the worked scenario, not a generic definition: An employee wants to paste a customer spreadsheet into an AI service to summarize it.

Teach: Demonstrate the method visibly: Classify the data first, minimize what is shared, use only approved systems, remove unnecessary identifiers, and keep human accountability. Pause at the judgment point and explain why the system should not simply continue on autopilot: Do not disclose sensitive or confidential data merely because a tool makes the workflow convenient. Least necessary data is the default.

Practice bridge: Ask the learner to reproduce the method on a different low-risk example, mark the verification/control step, and explain the reason for it.

Close: Day 12: build traceable research and knowledge-work practices.

Production / QA notes

- Use the worked scenario as the primary visual sequence; avoid generic AI-brain imagery.
- Keep terminology vendor-neutral and define unavoidable technical terms in context.
- Do not equate fluent output with verified truth or AI assistance with autonomous authority.
- Before learner release, fact-check the final lesson and validate every scored item against explicit teaching.

Learning comes before credentials. Completing this reading alone does not issue a certificate. Programme assessment and practical requirements apply.