FAST-17 Lab Guide: Responsible AI for All Staff

Provider: 7formations | Course: FAST-17 | Instructor: Wadood Chaudhary | Three labs (about 45 minutes each), one per session Audience: trainees (steps, worksheets) and instructors (setup, answer notes, timing). Instructor-only notes are marked **Instructor**.

All companies, people, and numbers are fictional. Labs run on paper or in the employer's approved assistant (Microsoft 365 Copilot Chat, ChatGPT Business or Enterprise, Gemini for Google Workspace, or Claude for Work) or the 7formations sandbox. This course explains Massachusetts rules so employees know they exist; it is not legal advice.

Suggested reading: Laurence Lars Svekis, *AI Literacy for All Employees: Understanding the Role of Artificial Intelligence in the Modern Workplace* (Kindle edition, 2026), chapters 8 to 12 and 14. The labs are original to 7formations.

0. Lab environment and safety rules (read before Lab 1)

Tools. Work account in the approved tool only. Lab 2 step 4 is the only step that requires an AI tool; it has a paper version.

Data

- Fictional data only. Under G.L. c. 93H and 201 CMR 17.00, a Massachusetts resident's name combined with a Social Security number, driver's license or state ID number, or financial account or card number is personal information. Never paste that combination into any AI tool.
- The Lab 2 "overshared prompt" contains fictional identifiers so you can practice removing them. Do not run the overshared version in any tool. Run only your rewritten version.
- If you paste something you shouldn't have, stop and tell the instructor. Reporting is the habit we practice.

Fictional organizations

- **North Shore Builders**: a 120-person residential builder in Beverly, MA.
- **Back Bay Accounting Partners**: a 25-person accounting firm in Boston.
- **Bayside Precision Parts**: a 60-person CNC machine shop in Worcester County, MA.

Completion. W1 to W3 complete / incomplete against each lab's criteria; W4 is the final (Appendix B). One resubmission is allowed.

Lab 1 (Session 1): Ethical risk check and bias detective

Scenario. North Shore Builders started using an AI shift scheduler that assigns weekend overtime across 14 crews. After two months, a crew lead notices that workers with a "childcare" availability flag have not been offered weekend overtime once. Overtime pays 1.5 times the regular rate.

Goal. Complete W1 Ethical risk check on the scheduler, and match eight bias examples to the four bias types.

The four principles: fairness, accountability, transparency, privacy. **The five-step ethical check:** (1) stakeholders: who is affected? (2) risks: what could go wrong, for whom? (3) transparency: do affected people know AI is involved and how to question it? (4) oversight: who reviews and can override? (5) compliance: which laws or policies apply?

Steps

1. **Read the brief (5 min).** One page: how the scheduler works, which fields it uses (crew, certifications, past hours, availability flags), and who approves the schedule.
2. **Run the five-step check (15 min).** In pairs, complete W1 for the scheduler.
3. **Bias detective (12 min).** Match the eight examples below to historical, representation, measurement, or evaluation bias, and write who you would tell.
4. **Accountability (5 min).** The site manager says "the computer decides, not us." Write two sentences in reply.
5. **Debrief (8 min).** What should North Shore do this week?

Bias examples (fictional)

1. A résumé model trained on ten years of mostly male hires ranks women lower.
2. A voice tool misunderstands strong regional accents.
3. The scheduler treats badge-in hours as effort, so workers who log site visits by phone look less productive.
4. A safety-glasses detector was tested only on the day shift under bright lights.
5. An image tool shows only men for "construction manager."
6. A credit model scores rural customers poorly because few rural loans were in its data.
7. A call-quality model rates calls by length, so short, efficient calls score as poor.
8. A translation tool was approved after testing only English and Spanish, then used with Portuguese-speaking customers.

Deliverable: W1 Ethical risk check

Step	North Shore scheduler
Stakeholders	
Risks (who could be harmed, how)	
Transparency (do workers know? can they question it?)	
Oversight (who reviews, who can override)	
Compliance (law or policy that applies)	
Principle most at stake, and why	
What I would report, to whom	

Completion criteria. All five steps completed with specifics; at least six of eight bias examples matched correctly; a report path named.

Lab 2 (Session 2): Redact and rewrite

Scenario. A staff accountant at Back Bay Accounting Partners wants help rewriting a client email and pastes this into a free personal chatbot:

> "Rewrite this email to Maria Lopez, SSN 012-34-5678, about her 2025 return showing $84,210 income and a $2,140 refund, account 4417 at Eastern Bank. She lives at 18 Marlborough St, Boston."

(All details are fictional.) Then the office receives a phone call in a voice that sounds exactly like the office manager, asking for a client's tax documents to be sent to a new email address immediately.

Goal. Complete W2 Data-safe prompt (before and after) and record the right response to the call.

Three kinds of sensitive information: personal information (about identifiable people), confidential business information (pricing, bids, client lists, unreleased financials), and regulated information (covered by specific law, such as the c. 93H combinations, health information, or tax data).

Steps

1. **Classify (8 min).** Circle every item in the overshared prompt and label it personal, confidential, or regulated.
2. **What went wrong (5 min).** List the two mistakes: the tool (personal and unapproved) and the data (identifiers the task did not need).
3. **Rewrite (10 min).** Write a data-safe version that gets the same help: remove identifiers, use placeholders such as [CLIENT NAME] and [AMOUNT], keep only what the task needs, and direct the client to the secure portal.
4. **Run the safe version (5 min).** Run only your rewritten prompt in the approved tool. Check that the output still does the job.
5. **If it already happened (5 min).** Write what the accountant should do now: report through the firm's incident channel the same day, and do not delete the chat to hide it.
6. **The call (8 min).** Work the branching scenario: pause, verify through a known number or in person, report to IT or security. The FBI has warned that criminals use AI-generated voice and video to impersonate trusted people.
7. **Settings check (4 min).** Find your approved tool's memory and data-training settings (the instructor shows where).

Deliverable: W2 Data-safe prompt

Before (overshared)	Items removed and why (personal / confidential / regulated)	After (data-safe prompt)	Tool used	Output still useful?

Plus: three-line incident response and three-line call response.

Completion criteria. Every identifier removed; placeholders used; task still achievable; incident and call responses name a verification or report step.

Lab 3 (Session 3): Governance clinic and personal commitment plan

Scenario. Bayside Precision Parts has a one-page AI use policy (sample in the lab packet): approved tools are Copilot Chat with work accounts; no customer drawings, personal information, or controlled technical data in any AI tool; high-stakes uses (hiring, discipline, pay, safety sign-off) need HR or management approval; AI-assisted customer documents carry a short note; concerns go to the operations manager or the anonymous form.

Goal. Complete W3 Governance scenarios and AI-use note, and W4 Responsible-AI commitments and 90-day learning plan (final).

Steps

1. **Four scenarios (15 min).** In groups, decide for each: allowed, ask first, or not allowed. Cite the policy line and say what you would do. - A. A machinist wants to use a free AI app to summarize a defense customer's drawing notes. - B. The shipping lead wants AI to draft a customer delay notice. - C. A supervisor asks you to use AI to rank employees for a layoff. - D. Your team has quietly started using an unapproved AI add-on that uploads customer files.
2. **AI-use note (8 min).** Write a plain-language note telling customers how Bayside uses AI in quotes and delivery notices (three or four sentences: what AI does, what a person checks, how to reach a person).
3. **Future of your role (7 min).** List two tasks in your job AI is likely to change in the next two years, including any that agents could do, and one skill to build for each. Use the WEF Future of Jobs 2025 skills list as a prompt (analytical thinking, resilience and flexibility, AI and big data literacy, among others).
4. **Commitments and plan (12 min).** Write W4.
5. **Pair review (3 min).** A partner checks that your plan has a date and a reviewer.

Deliverable: W3 Governance scenarios and AI-use note

Scenario	Allowed / Ask first / Not allowed	Policy line	What I would do

Plus the customer AI-use note.

Deliverable: W4 Responsible-AI commitments and 90-day learning plan (final)

Part	My plan
Three commitments (tools, data, speaking up)	
Two tasks AI may change in my role	
Skill 1 and how I will practice it (low-risk task)	
Skill 2 and how I will practice it	
Check-in date (within 90 days)	
Reviewer (role)	

Completion criteria. W3: all four scenarios decided with a policy line; customer note in plain language. W4: see Appendix B.

Appendix B: Rubric for W4 (Not yet / Meets / Strong)

Criterion	Not yet	Meets	Strong
Commitments	Generic ("be careful")	Three specific commitments covering tools, data, and speaking up	Each names the tool, data rule, or channel by name
Future of role	Missing or "AI will take my job"	Two realistic task changes	Includes an agent or automation example and the human check that remains
Skills and practice	No practice task	Two skills, each with a low-risk practice task	Practice tasks are checkable and use approved tools
Accountability	No date or reviewer	Date within 90 days and a reviewer role	Includes how progress will be shown
Data and ethics	Plan would use restricted data	No restricted data; principles respected	Names a principle at stake in one task