


Human-AI-Human

A workflow for keeping human judgment and accountability explicit.



Before use: a human makes the decision

Use only after review. If a concern remains unresolved, pause, verify, ask the responsible person, or reject the output.

Three safeguards travel together

Privacy: protect sensitive information. Accuracy: verify facts and claims. Judgment: keep meaningful decisions and accountability with people.

The final human step is not ceremonial. It is where professional responsibility lives.

A short implementation guide

Use the workflow for one bounded task, then make the review repeatable.

1. Define the work and name its owner

State the purpose, audience, desired result, and constraints. Name the person responsible for the result and any required approval. Ask whether AI adds value and whether local policy permits the proposed use.

2. Set the input boundary

Use a tool approved for this task and these inputs. Keep patron, employee, credential, and confidential organizational data out of unvetted systems. Prefer public or fictional material for practice. If approval or data handling is unclear, pause and ask the responsible person.

3. Ask for bounded assistance

Provide the allowed source material, requested format, and limits. Ask AI to flag missing information rather than invent it. Use it for drafts, options, or summaries; keep the decision with the named human owner.

4. Review against evidence and consequences

Check facts, citations, calculations, and claims against appropriate sources. Review privacy, bias, tone, accessibility, and likely consequences. Consider whose needs or perspective are missing. A second AI answer does not replace checking the evidence.

5. Decide, document, and learn

Accept, revise, verify, or reject. Follow local approval and disclosure requirements before use. Record the owner, sources checked, important changes, and final decision where appropriate to the work. Use recurring problems to improve procedures and training.

A useful review record: Task | Owner | Allowed inputs/tool | Evidence checked | Changes | Decision/approval

Workshop: make judgment visible

A reusable 30-minute exercise for library staff and supervisors.

Purpose and preparation

Participants will define a bounded task, identify unsupported claims in a sample draft, and explain who owns the final decision. Give each pair pages 1, 4, and 5; keep the worked example on page 5 covered until the debrief. Bring pens or an accessible editable copy. No AI account is required.

0-5 minutes | Introduce the model

Walk through Define, Assist, and Review and own. Ask: What has the first human decided before AI starts? What must the final human be able to explain?

5-10 minutes | Define

Pairs read the fictional source card and fill in the first section of the worksheet. Ask them to specify allowed inputs, the owner, and a useful prompt. For this exercise, assume the tool is approved for fictional training data.

10-18 minutes | Review the sample

Reveal the simulated AI draft. Pairs mark each unsupported claim, identify who might be excluded or misled, and write a corrected announcement. They must choose a disposition: accept, revise, verify, or reject.

18-25 minutes | Compare decisions

Use the worked example to discuss evidence and consequences. Ask: Which changes were factual? Which required judgment? Who may release the announcement? What would you do if registration or access details could not be confirmed?

25-30 minutes | Transfer and check learning

Each participant names one real task, one input boundary, one review check, and its accountable owner. Success means they can explain their decision with evidence, avoid invented details, and identify a human owner. Do not collect sensitive examples.

Participant worksheet

Work in pairs. Use the fictional source card on page 5.

Define

Purpose and audience? Who owns the result? What inputs and tool are allowed?

Assist

Write a bounded prompt. What may AI help with? What must it leave for a person?

Review and own

List unsupported claims and source checks. Check privacy, bias, tone, accessibility, and consequences.

Decide

Accept, revise, verify, or reject? Write your corrected announcement and explain the decision. Who authorizes use?

Transfer

Choose one real task. Name its input boundary, review check, and accountable owner.

Practice case and worked example

Fictional training material. The sample draft is deliberately flawed.

Source card: facts supplied by the program owner

A library is offering an introductory spreadsheet workshop for adults. It takes place November 12, 2026, from 2 to 3 p.m. in the meeting room. Admission is free. Registration is required, with 12 places available. Participants should bring a laptop. No registration method, loaner equipment, or accessibility arrangements have been supplied.

Simulated AI draft: review before use

"Drop in for our free spreadsheet workshop on November 12, 2026, from 2 to 3 p.m.! All ages welcome, no registration needed. Laptops are provided, and everyone will leave an expert. The workshop is fully accessible."

Example prompt

Draft a short announcement for adults using only the source-card facts. Do not invent logistics, equipment, access arrangements, or outcomes. Flag missing details separately for the program owner to confirm.

Worked review

Correct drop-in/no-registration language and the audience. Restore the location, capacity, and bring-a-laptop requirement. Remove the unsupported equipment, expertise, and accessibility claims. Confirm the registration route and access information with the program owner; consider how the laptop requirement affects participation.

Corrected draft: pending owner review

Join our introductory spreadsheet workshop for adults on November 12, 2026, from 2 to 3 p.m. in the library meeting room. Admission is free. Registration is required, and 12 places are available. Please bring a laptop.

Decision and ownership

Revise, then verify the missing registration and access details. The program owner reviews the announcement and follows the library's publication approval process. Do not release this exercise draft as a real announcement.

Framework source: Robin Hastings, AI Policy Implementation Case Study, "The operating model," and Presentation Excerpts, slide 3 (HQ).
Exercise and worksheet newly drafted for this package.