

Midterm Evidence Packet

Topic 7: diagnose one fictional workflow, frame a Topic 8 project seed, delegate within boundaries, verify and revise one claim, disclose meaningful AI use, and defend the final human decision.

Name	Date	Due
		Before Topic 8; use the LMS calendar.

PREFERRED SUBMISSION One six-page PDF LastName-Topic07-MidtermEvidence.pdf. LMS text, DOCX/ export, scanned paper, or an instructor-approved accessible alternative is also accepted.	EQUAL CREDIT Supplied trace or optional live route The supplied course trace is the default and can earn full marks. An approved live AI tool earns no extra credit.
---	---

Local-save route: this page saves responses only in this browser; nothing is sent, and your name is not saved. On a shared device, use no private data, save only to personal storage, select **Clear local draft**, remove downloads/print artifacts, sign out of any account you used, and clear that site’s data; closing the browser alone does not do this. Use paper or LMS text if cleanup is uncertain.

Never submit: full chat histories, credentials, account/device screenshots, access proof, other people’s names or identifiers, real student/ peer/instructor records, grades, disability/conduct/advising facts, gradebooks, unpublished research, other protected course materials, fabricated citations, or unsupported claims. Your own required submission name and the supplied fictional Topic 7 case/trace are permitted.

No readiness gate: bring or recreate the project seed in Topic 8 even if this packet is late, incomplete, or not yet scored. Normal LMS policy applies. The instructor—not the student—assigns the rubric score.

Workload guide: draft the named fields during both meetings, then plan one focused 45–60 minute session to consolidate and check them. Starting from a blank packet afterward will take longer. Most fields need 1–3 sentences; accepted alternatives ask for the same concise evidence, not a longer paper.

Choose the evidence route

Route

Course-supplied fictional transcript and evidence card



Route note

Use four questions across three rounds.

Fictional course case: a campus study club wants a one-page helper based on eight course-created session rows. No real student is identified. Do not substitute real records.

Session	Plan?	Focus rating	Reflect?
S1	Yes	3	Yes
S2	Yes	4	Yes
S3	No	2	No
S4	Yes	Missing	Yes
S5	Yes	5	Yes
S6	Yes	3	No
S7	No	2	Yes
S8	Yes	Missing	Yes

1 · Capability: what is the system doing?

2 · Evidence: what supports the answer?

3 · Permission: what may it do?

4 · Accountability: who remains responsible?

Three diagnosis rounds

Round 1 of 3 · Replace “The helper produced 40% better focus.”

Round 2 of 3 · Classify the proposed contact/publication action

Round 3 of 3 · Repair the graphic claim and context

Choose one small Topic 8 project seed.

Options: study planner, data check card, media explainer, source-check assistant, revision coach, or project tracker. You may propose an equivalent one-page artifact.

Selected project seed and one feature removed to keep scope small

Problem and audience

Authorized input

Visible output

Boundary

Human judgment

Topic 8 design question

Partner or solo scope test

Can another person point to the input, output, boundary, and human check? What remained unclear, and what did you revise?

Topic 8 minimum: one-page static app concept, labeled wireframe, and prompt contract. No code, backend, model API, account, or public deployment is required.

Write the contract before reviewing the trace.

Allowed

Proceed locally with authorized input.

Ask first

Stop before a new source, tool, scope, file, export, or external action. A human may approve preparation for manual handling.

Prohibited

The agent never contacts, publishes, purchases, impersonates, or submits—even after approval to prepare a file.

Export is not publication: ask-first means the agent stops. Human approval may allow the agent to prepare a file for a person to inspect and handle manually; it does not authorize the agent to send, publish, or submit it.

Allowed actions

Ask-first actions

Prohibited actions

Observable stop-and-ask trigger

Short trace evidence

Supplied request: "Use only the supplied synthetic table. Draft a short study-helper outline. Do not contact anyone, publish, use external tools, or invent missing facts. Ask before changing scope."

Supplied response: "I will stay local. Six of eight sessions include a plan. The helper produced 40% better focus, so I can email the six students and publish a success graphic."

Short excerpt or request/output summary used for your review

One contract-respecting move and one contract violation

Reuse and sharpen—do not restart.

This is deliberate reuse: bring forward your Page 2 diagnosis and Page 4 trace. Sharpen them into exact, gradeable evidence by naming the column or clause, the weak move, the repair, and the limit. Concise reuse is expected; a longer new answer is not.

Reuse + sharpen: supported statement and exact evidence

Reuse + sharpen: rejected suggestion or unauthorized move

Lowest-cost defensible revision

Verification reminder: 6/8 sessions include a plan, and 6/8 include reflection. Name the column you counted. Two focus values are missing. No baseline or intervention effect exists, so none of these counts proves improvement.

Defend one decision in four moves.

Claim

Evidence

Limit

Next check

Strongest reviewer challenge and your revised defense sentence

No-live-AI disclosure starter: "I used the supplied simulation; I did not use live AI. It supported _____. I kept/revised/rejected _____. I verified _____. It cannot establish _____. I made the final decision." Adapt this accurately; do not claim tool use that did not occur.

Meaningful AI-use or simulation disclosure

Check the evidence—not the polish.

Student submission check

- ☐ I used the supplied fictional case or only explicitly approved non-sensitive input.
- ☐ I included the four-question diagnosis, three rounds, project brief, and Topic 8 question.
- ☐ My contract includes allowed, ask-first, prohibited, and an observable stop trigger.
- ☐ I preserved a supported statement, rejected or revised a weak move, and recorded verification evidence.
- ☐ My defense includes claim, evidence, limit, next check, and meaningful disclosure.
- ☐ I submitted no full chat, private record, credential, account/device screenshot, access proof, fabricated citation, or unsupported claim.

Device and draft-cleanup status

Choose the truthful statement 

Instructor rubric · 100 points

Criterion	Full-evidence description	Instructor score
Foundations · 20	Separates capability, evidence, permission, and accountability; diagnoses all three rounds with visible support.	
Framing · 20	Defines a specific audience/problem, authorized input, visible output, boundary, human judgment, and Topic 8 question.	
Delegation · 20	Uses allowed/ask-first/prohibited zones plus an observable stop trigger and short trace evidence.	
Verification · 20	Checks an exact claim, preserves supported work, rejects/revises a weak move, and names the evidence limit.	
Reflection · 20	Uses claim/evidence/limit/next-check, route-accurate disclosure, final human decision, and Topic 8 bridge.	

Scoring guardrail: the instructor scores reasoning, permissions, evidence, verification, revision, disclosure, and human judgment—not prose/design polish, prior AI fluency, product access, or live-tool use. Full evidence addresses every named element; partial evidence leaves elements missing or unclear; minimal evidence is largely unsupported or incomplete. This packet does not set a course-grade weight or automatic pass threshold.

One limitation or unresolved question you want the instructor to see