CDSS 94 (UC Berkeley): Posttraining

Final Project Checkpoint B

Instructors: Karina Nguyen & Kevin Miao • Spring 2026

Checkpoint: Final Project Checkpoint (B)
Due on Gradescope: March 9, 2026 at 4:59pm

Deliverables (Final Project B)

1. **Write-up:** PDF write-up covering the sections below.
2. **Architecture diagram:** A clear visual showing how your system's components connect (hand-drawn is fine).
3. **Project Progress Meeting:** Schedule a progress meeting for 3/9 using this link: https://tinyurl.com/cdss94-proj-b.

Goal of this checkpoint (Final Project B)

This is a **technical blueprint and project plan**. In Checkpoint A you identified a problem worth solving. Now you need to answer: *what exactly are you building, how does it work, how will you know it is working, and what does the path to a finished project look like week by week?*

The emphasis is on **thoughtful design choices**. We do not expect everyone to train a model from scratch. You may be building prompt pipelines, RAG systems, evaluation frameworks, agent workflows, fine-tuning infrastructure, or something else entirely. What matters is that every technical decision is deliberate and justified.

Note. Your architecture will change. That is fine. The goal here is to force precise thinking *now* so that the changes you make later are informed, not panicked.

Guiding questions

You do **not** need to use a rigid section structure. Treat the questions below as a checklist to make sure your write-up is complete.

1. **What are you building?** Describe the system end to end: inputs, outputs, and the user-facing experience. Include 1–2 concrete usage scenarios that walk through how a real user interacts with your tool from start to finish.
2. **How does the system work?** Break down your architecture into components (models, APIs, databases, pipelines, UI, etc.). For each component, explain:
   - What it does and why it exists.
   - What alternatives you considered and why you chose this approach.
   - What off-the-shelf pieces you are using vs. what you are building yourself.

   Your architecture diagram should make the data flow and component relationships clear at a glance.

3. **Where are the thoughtful AI design decisions?** Where in your system do you need to think carefully about how AI behaves? This could include prompt design, output validation, guardrails, human-in-the-loop steps, calibration, failure modes, bias mitigation, privacy boundaries, or anything else that separates a thoughtful system from a naive one. Be specific: name the decision, explain the tradeoff, and describe your current approach.

4. **What is your data strategy?** Where does your data come from (existing datasets, user-generated, synthetic, scraped, manually curated)? How much do you need? What are the quality, licensing, and privacy constraints? If you are fine-tuning or doing any form of training, describe the pipeline. If you are not training, explain what data your system consumes at inference time and how you ensure its quality.

5. **What are your KPIs?** Define 2–4 measurable objectives you are optimizing for. For each KPI:
   - State the metric precisely (e.g., "task completion rate," "user preference win rate vs. baseline $> 60\%$").
   - Identify a baseline (current state, naive approach, or existing tool).
   - Set a target that defines your minimum viable win.

6. **How will you evaluate?** Describe your evaluation plan: what tests you will run, what datasets or scenarios you will use, and how you will distinguish between "the model is bad" and "the system around the model is bad." Explain how this evaluation feeds into the user study you will run in Final Project C.

7. **What is your week-over-week plan?** Provide a week-by-week timeline from now through Final Project D. For each week, state:
   - The milestone or deliverable for that week.
   - Who on the team owns it.
   - What the "done" criteria are (how you know the week was successful).