

AI and Deep Learning

Week 2 – Activities and empirical work

Student and instructor route

Four days; offline, CPU-only, reproducible

Remember

Each day moves from an individual concept commitment to a dataset-based comparison and ends with the narrowest claim supported by frozen evidence. The questions test reasoning about models and systems, not memory of software function names.

How each day works

1. **Before class.** Complete the pre-class questions and record one reason for each choice.
2. **Concept studio.** Vote silently, discuss with a partner, revote, inspect the reveal, and repair one overclaim.
3. **Empirical studio.** Read the dataset card before opening code. Record the unit, information boundary, split roles, baseline, and metric.
4. **Final reveal.** Open final or stress evidence only after the system and evaluation rule are frozen.
5. **Exit ticket.** State the result, the evidence, its scope, and one important limitation.

Important

Do not open a later label file, instructor notebook, answer key, or expected metric early. The order is part of the evidence design.

Daily empirical catalogue

Day	Local dataset	Main comparison and decision
1	1,800 synthetic tabular rows	prior baseline, logistic model, small MLP, then a deliberately invalid leaked model; select threshold on validation only
2	1,200 synthetic shape images and 180 stress images	majority, flat MLP, small CNN; compare clean performance, parameter count, and shortcut robustness
3	15 course documents and 30 controlled questions	three transparent TF-IDF retrievers; tune retrieval and abstention on development only
4	16 policy or untrusted documents and 32 tasks	vulnerable retrieve-first assistant versus a guarded controller with calculator, refusal, abstention, and human gates

Evidence rules

1. Fit learned preprocessing on training or development evidence only.
2. Choose thresholds, architecture settings, and retrieval settings before the final or stress reveal.

3. Compare systems on the same rows and report the relevant denominators.
4. Keep prediction, retrieval, support, action, and causal claims separate.
5. Treat retrieved instructions as data, never as higher-priority policy.
6. No Day 4 classroom path performs an external or high-impact action.

Compact 90-minute route

Use 22 minutes for the concept studio, 58 minutes for the empirical notebook, and 10 minutes for debrief and exit tickets. In the empirical block preserve the contract, baseline, freeze, final reveal, and failure audit; shorten only optional model exploration.