CS109 Lecture 11: LLM Learning Guide

Inference

Summer 2026

Lecture 11: LLM Learning Guide

Before lecture: read the Lecture 11 slides (inference: updating belief distributions with Bayes' theorem). Then open your favorite LLM and work through the concepts below **in order**. Each concept has a *Learn* prompt and a *Test me* prompt. Attempt the "Test me" questions yourself before asking the LLM to grade you.

How to get the most out of this:

- Tell the LLM your level: "I'm a student in Stanford's CS109 (probability for computer scientists). Explain at that level."
- After any answer, ask "why?" at least once and ask it to show each step.
- When it states a formula, ask it to define every symbol before continuing.
- Attempt practice yourself first, then say: "Here is my reasoning: _____. Where exactly am I wrong, if anywhere?"

Concept 1: Inference as belief updating

Learn: > "Explain inference as updating one's belief about a hidden random variable given observations, using Bayes' theorem: posterior $\propto$ prior $\times$ likelihood. Contrast a prior belief (before data) with a posterior belief (after data). Use a concrete example like updating belief in a patient's vision ability after they read (or miss) a letter."

Test me: > "Give me a hidden random variable with a prior and an observation model, and make me explain which piece is the prior, which is the likelihood, and which is the posterior. Then check my labeling."

Concept 2: The update loop over a belief table

Learn: > "Explain how inference over a discrete random variable is implemented as a loop: for each value a, multiply the prior $P(A = a)$ by the likelihood $P(\text{obs} \mid a)$ to get an unnormalized posterior, then normalize. Represent the belief as a dictionary/table (non-parametric). Work a small numeric example by hand."

Test me: > "Give me a discrete prior (as a small table) and a likelihood for one observation, and make me compute the unnormalized posterior and then normalize. Check each product and the normalization."

Concept 3: Normalization

Learn: > "Explain the normalization constant in Bayes' theorem: it equals $P(\text{observation}) = \sum_a P(A = a)P(\text{obs} \mid a)$, the same for every value of a, and it rescales the posterior to sum to 1. Explain when you can skip it (finding only the mode / argmax).”

Test me: > “Give me a set of unnormalized posterior values and make me (1) normalize them and (2) identify the mode without normalizing. Check both.”

Concept 4: Multiple observations

Learn: > “Explain how to incorporate several conditionally independent observations: multiply the prior by the product of the individual likelihoods (or, equivalently, update sequentially where each posterior becomes the next prior). Show that both routes give the same posterior.”

Test me: > “Give me a prior and two observations, and make me update using both methods (sequential, and all-at-once product of likelihoods). Check that I get the same posterior both ways.”

Concept 5: Discretizing continuous variables

Learn: > “Explain how a continuous hidden variable (like an age or a distance) is handled in inference by discretizing it into a dictionary of closely spaced values, then running the same prior-times-likelihood-then-normalize loop. Use Bayesian carbon dating as the example, where the likelihood comes from a decay model.”

Test me: > “Give me a continuous inference scenario and make me describe how to discretize the hidden variable, what the likelihood term is, and how the update loop works. Check my description.”

Concept 6: Reading likelihoods off a model

Learn: > “Explain how to write the likelihood $P(\text{obs} \mid \text{hypothesis})$ from a described data-generating model, including cases where the likelihood is an indicator (0 or 1), a PMF (binomial/Poisson count), or a PDF (normal/exponential observation). Work one example of each.”

Test me: > “Give me three short scenarios and make me write the correct likelihood for each (indicator, PMF, and PDF). Check that I picked the right kind of likelihood and set it up correctly.”

Wrap-up prompt (after all six concepts)

> “Give me one multi-part CS109-style inference problem where I must (1) set up a prior over a discretized hidden variable, (2) write the likelihood from a described model, (3) update the belief and normalize, and (4) incorporate a second observation. Let me solve every part, then grade each part, tell me which concept it tested, and tell me the single concept I should review most before the next lecture.”

Note: next lecture generalizes inference to models with many random variables (Bayesian networks) and introduces sampling-based inference (rejection sampling, MCMC) for when the joint is too large to handle directly.

Bring any sticking points to lecture; the TAs and instructor will help you work through them on the worksheet.