CHAPTER 07 / TRAINING

Same lineage. New signal.

Training changes weights. Request context does not.

Training changes weights. Request context does not.

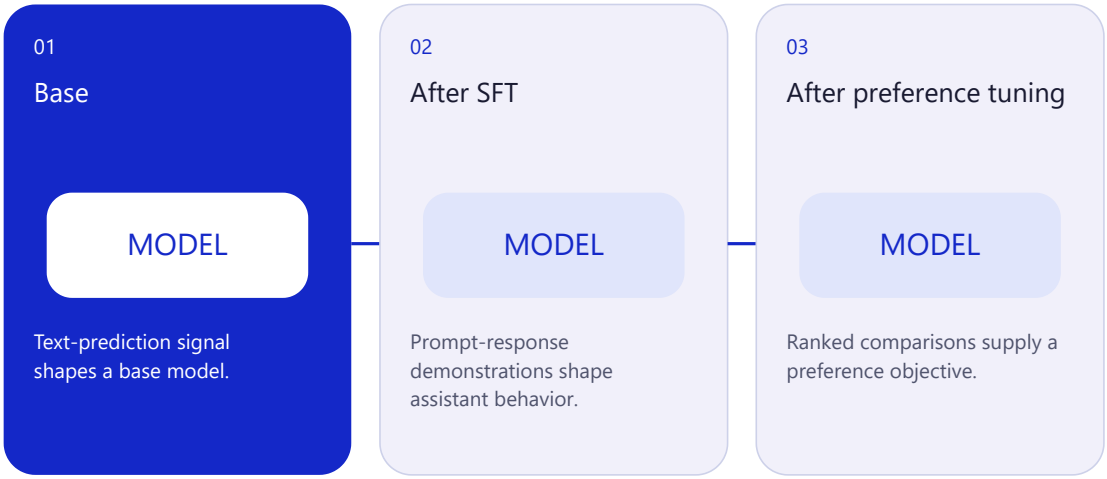

One lineage: each stage starts from the previous weights.

One common recipe, not every model. Scripted samples, not observed outputs. Selecting a preference only records data; it does not train a model.

Pre-training

Text prediction fits the initial model parameters.

Demonstrations

Supervised fine-tuning uses examples of desired responses.

Preferences

Comparisons supply a proxy objective; preference is not proof of truth.

↗ **Text prediction fits the initial model parameters.**

PUT IT TO WORK

Rank two made-up security answers, name your criteria, then compare where colleagues might disagree.

CLOSE THE PAGE. EXPLAIN IT.

Which training signal could shape a confident answer, and what would you still verify?