FAST-05 Projects: Applied Data Science and Machine Learning for Business Users

Provider: 7formations | Course: FAST-05 | Instructor: Wadood Chaudhary | One capstone and two mini projects

Projects apply the course to models and forecasts trainees meet at work, without building any. Scoring uses **Not yet / Meets / Strong**. Completion requires Meets on every capstone criterion; mini projects are formative unless the employer asks for scoring. One resubmission allowed. Use fictional, public, or approved aggregated data only; never customer personal information or PHI.

Capstone: Decision Memo on a Model or Vendor AI Feature

Brief. Choose a model, score, forecast, or vendor AI feature used (or proposed) in your role: CRM lead scoring, ERP demand forecast, claim or fraud triage, quality flags, scheduling predictions. If you have none, use ClaimGuard (Lab 5) or the clinic no-show model (Lab 3). Write a 2-page decision memo for a non-expert manager evaluating whether and how to use it, and present it in three minutes.

Deliverables

1. Prediction map (W1 format).
2. Metrics in business terms: what errors cost in each direction, and the threshold or override approach you recommend.
3. Data and fairness review: 3+ data questions or problems, who could be harmed.
4. One verified analysis: an AI-assisted finding confirmed by manual calculation (fictional or approved data).
5. Vendor or owner questions (6+) and a monitoring plan (owner, measure, trigger, action).
6. 100-word explain-it-back for an affected customer or coworker.
7. Recommendation: adopt / adopt with conditions / pilot / don't adopt, with reasons.

Step plan

Step	When	What you do
1	After session 1	Choose the model; W1 map
2	After session 2	Costs and threshold reasoning
3	After session 3	Data and fairness review
4	After session 4	Verified analysis
5	Session 5	Questions, monitoring, explain-it-back; present

Rubric

Criterion	Not yet	Meets	Strong
Metric interpretation	Accuracy only or no costs	Metrics tied to business cost both directions	Threshold or override grounded in cost and capacity
Data judgment	Fewer than 3 issues	3+ real data or fairness issues, specific	Issues sized; affected groups and mitigations named
Verification	No manual check	One AI finding confirmed manually	Differences explained and reconciled
Communication	Jargon; recommendation unclear	Non-expert understands decision and reasons	Explain-it-back scored clear by peers; next steps owned
Governance	No monitoring	Owner, measure, trigger, action	Linked to relevant rules (e.g., MA DOI Bulletin 2024-10, employer policy)

Strong vs weak example (recommendation paragraph)

Weak	Strong
"ClaimGuard is 95% accurate, so we should use it. AI is the future and our competitors use it."	"Pilot ClaimGuard for 90 days on claims from policies older than 2 years only, at a threshold that keeps flags within adjuster capacity (about 80/month). The vendor's 95% accuracy is on Midwest auto claims and doesn't tell us how many frauds it would catch for us; we've asked for homeowners precision and recall and for ZIP-code proxy testing. Claims supervisor reviews flagged-vs-confirmed monthly; if fewer than 15% of flags are confirmed for two months, we pause."

Mini project 1: Metric Translator

Brief. Pick three metrics from a dashboard, vendor sheet, or report you see (or use the Lab 2 table). Translate each into a plain sentence with business consequences for your team.

Deliverables. Three translations; for each: the metric, its value, a plain-language sentence, what a wrong decision costs, one question to ask the owner.

Step plan. (1) Pick metrics; (2) define each in one sentence; (3) attach costs; (4) test the sentences on a coworker; (5) revise.

Rubric

Criterion	Not yet	Meets	Strong
Accuracy	Definitions wrong	Definitions correct	Includes base rate context
Plain language	Jargon remains	A coworker understood all three	Uses real counts ("14 of 20")
Business link	No consequence	Consequence and question per metric	Suggests a better metric where needed

Strong vs weak example. Weak: "Precision is 28%, which is the positive predictive value." Strong: "Of every 50 claims the model flags, about 14 are fraud. So adjusters spend most review time on legitimate claims; ask whether a 0.6

threshold catches more fraud within our capacity."

Mini project 2: Forecast Review Routine

Brief. Design a monthly forecast-vs-actuals review for a forecast your team uses (or Bayside's BP-2207). Include error tracking, out-of-range alerts, override logging, and override review.

Deliverables. A one-page routine; a template sheet (columns: month, forecast, low, high, actual, abs error, outside range?, override, reason, approver, override error); three months of fictional or approved example entries; a rule for when to request model retraining.

Step plan. (1) Pick the forecast; (2) build the template; (3) fill examples; (4) verify one AI-computed error by hand; (5) write the retraining and escalation rule.

Rubric

Criterion	Not yet	Meets	Strong
Error tracking	No error measure	Abs error and MAE by period	Separates periods before/after known changes
Override discipline	Overrides unlogged	Reason, source, approver, review	Override accuracy compared to original forecast
Verification	AI numbers unchecked	One checked manually	Template formulas documented

Strong vs weak example. Weak: "Check the forecast monthly." Strong: "First Monday each month the planner logs actuals; any month outside the 80% range for two months triggers a review with sales; overrides are logged with source and compared to the original forecast after actuals arrive; if override error is lower three months running, request retraining."