Lesson 5.1: Machine Learning Foundations for Finance – Practice Exercises

Digital Finance Course

Exercise 1: Classification vs. Regression

Classify each financial problem as Classification (C) or Regression (R), and propose an appropriate evaluation metric:

1. Predicting whether a credit card transaction is fraudulent
2. Estimating the expected loss on a mortgage portfolio
3. Deciding whether to approve or deny a personal loan application
4. Forecasting next month's Value at Risk (VaR) for a trading desk
5. Detecting whether a customer will churn within 90 days
6. Predicting the settlement price of a commodity futures contract

For each problem, also answer:

- Is the data likely to be imbalanced?
- What is the cost asymmetry (which type of error is more expensive)?

Exercise 2: Confusion Matrix Calculations

A fraud detection model processes 50,000 transactions. The results are:

	Predicted Legitimate	**Predicted Fraud**
Actual Legitimate	48,200	500
Actual Fraud	120	1,180

Calculate:

1. Precision for the fraud class
2. Recall for the fraud class
3. F1-score for the fraud class
4. Overall accuracy
5. If each missed fraud costs \$5,000 and each false alarm costs \$50, what is the total cost of model errors?

Exercise 3: Bias-Variance Diagnosis

Three credit scoring models are evaluated. Diagnose each and recommend an action:

Model	Train AUC	Test AUC	Diagnosis?
Logistic Regression	0.72	0.70	?
Decision Tree (depth=30)	0.99	0.65	?
XGBoost (tuned)	0.91	0.88	?

For each model:

1. Classify as underfitting, overfitting, or good fit
2. Explain the relationship between bias and variance for this model
3. Recommend one specific action to improve performance
4. Which model would you deploy in production? Justify your choice.

Exercise 4: Feature Engineering for Anti-Money Laundering (AML)

You have a dataset of bank transactions with these raw fields:

`transaction_id, account_id, timestamp, amount, currency, sender_country, receiver_country, transaction_type, account_age_days, account_balance`

Design 8 engineered features that could help detect money laundering. For each feature:

1. Name the feature
2. Write the formula or logic
3. Explain why it captures suspicious behavior
4. Categorize it: ratio, rolling window, lag, aggregation, or time-since-event

Bonus: Which of your features would be most informative for detecting "structuring" (breaking large transfers into many small ones to avoid reporting thresholds)?

Exercise 5: Model Comparison – Tree-Based Methods

You are building a credit default prediction model. Compare three approaches:

Criterion	Decision Tree	Random Forest	XGBoost
How trees are built	?	?	?
Overfitting risk	?	?	?
Interpretability	?	?	?
Handles missing data	?	?	?
Training speed	?	?	?
Typical AUC range	?	?	?

Questions:

1. Fill in the table with appropriate descriptions
2. If a regulator requires full model explainability, which model would you choose and why?
3. Under what conditions might random forest outperform XGBoost?

Exercise 6: ROC Curve Interpretation

Two credit scoring models are evaluated on the same test set. Their ROC curves yield:

- Model A (Logistic Regression): AUC = 0.78
- Model B (XGBoost): AUC = 0.91

Questions:

1. What does it mean that Model B's AUC is 0.91? Express this in plain language.
2. At a False Positive Rate (FPR) of 5%, Model A achieves a True Positive Rate (TPR) of 45% and Model B achieves 82%. What does this mean for the bank?
3. The bank processes 100,000 loan applications per year, of which 3,000 eventually default. Using Model B at a 5% FPR operating point:
   - How many defaults will it correctly identify?
   - How many non-defaulters will be incorrectly flagged?
   - What is the precision at this operating point?
4. Under what business scenario would you choose Model A despite its lower AUC?

Exercise 7: Anomaly Detection vs. Supervised Classification

A fintech company processes payments and wants to build a fraud detection system. They have:

- 5 million labeled transactions from 2022–2023 (0.3% confirmed fraud)
- Evidence that fraudsters changed tactics significantly in late 2023
- A compliance team that can investigate 500 alerts per day

Questions:

1. Would you use supervised classification, anomaly detection, or both? Justify your architecture.
2. If you choose anomaly detection (Isolation Forest), what contamination parameter would you set and why?
3. The compliance team capacity is 500 alerts/day. If the model processes 200,000 transactions/day, what maximum false positive rate is acceptable?
4. Design a two-stage system that combines both approaches. Draw the data flow and explain each stage.

Exercise 8: End-to-End ML Pipeline Design

Design a complete ML pipeline for predicting customer churn at a neobank.

Given:

- 2 million customers, 18 months of data, 8% annual churn rate
- Available data: transaction history, app usage, demographics, customer service contacts
- Requirement: model must be explainable (EU AI Act compliance)

For each pipeline stage, specify your choice and justify it:

1. **Problem framing:** Classification or regression? Target variable definition?
2. **Feature engineering:** List 5 features you would create (with formulas)
3. **Validation strategy:** How would you split the data? Why?
4. **Model selection:** Which model(s) would you train? Why?
5. **Evaluation metric:** Which metric is primary? Why not accuracy?
6. **Deployment:** How would you monitor for model drift?