Module 5 Summary: The Automation Problem

Prof. Dr. Joerg Osterrieder

Digital Finance — BSc Course

Key Concepts by Lesson

L1: ML Foundations

- Classification ("what kind?") vs. regression ("how much?")
- Bias-variance trade-off: underfitting vs. overfitting
- Train/validation/test split; cross-validation
- Feature engineering for financial time series

L2: Generative AI & LLMs

- Transformer architecture and attention mechanism
- LLMs: next-token prediction, not "understanding"
- Financial use cases: document processing, compliance review
- RAG pipelines for financial document Q&A

L3: Limits of Prediction

- Stationarity: mean, variance, autocorrelation must be constant
- Regime changes break models trained on historical data
- Pitfalls: data snooping, look-ahead bias, survivorship bias

L4: MLOps

- MLOps = DevOps principles for ML lifecycle
- Drift types: concept drift, data drift, covariate shift
- Model monitoring, versioning, and governance
- Production gap: notebook $\neq$ production system

Module 5 answers: Can machines make better financial decisions than humans — and at what cost?

Key Formulas and Definitions

Bias-Variance Decomposition

$$\text{Expected Error} = \text{Bias}^2 + \text{Variance} + \text{Irreducible Noise}$$

High bias = underfitting (model too simple). High variance = overfitting (model too complex).

Stationarity Requirement

$$E[y_t] = \mu, \quad \text{Var}(y_t) = \sigma^2, \quad \text{Cov}(y_t, y_{t+k}) = \gamma(k) \quad \forall\, t$$

Financial data almost always violates stationarity — regimes shift, volatility clusters, trends emerge.

Overfitting Diagnostic

If training error $\ll$ validation error $\rightarrow$ overfitting. If both are high $\rightarrow$ underfitting. If converging at low level $\rightarrow$ good fit. Plot learning curves (error vs. data size) to diagnose.

Drift Detection

Concept drift: $P(Y|X)$ changes. **Data drift:** $P(X)$ changes. **Covariate shift:** $P(X)$ changes but $P(Y|X)$ stays the same. All require retraining or model replacement.

The hardest part of financial ML is not building the model — it is knowing when the model has silently stopped working.

Connections to Other Modules

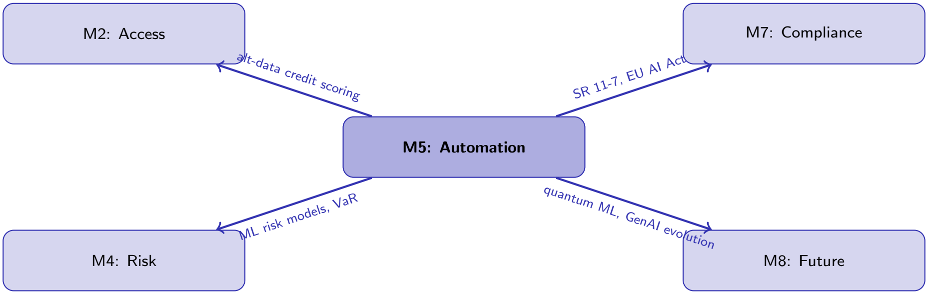

- **Automation → Access (M2):** ML powers alternative credit scoring that includes the "credit invisible" — but introduces algorithmic bias
- **Automation → Risk (M4):** ML models estimate VaR, detect fraud, and predict default — model risk is the cost of automation
- **Automation → Compliance (M7):** SR 11-7 governs model risk; the EU AI Act classifies credit scoring as "high risk" requiring human oversight
- **Automation → Future (M8):** Quantum computing may accelerate ML training; GenAI capabilities will continue to expand

ML is the engine; data is the fuel; governance is the brake. All three must work together.