Macroeconomic Forecasting and Analysis in the Machine Learning Era

Philippe Goulet Coulombe, Ph.D.

OeNB Summer School · Joint Vienna Institute, Vienna · August 24–28, 2026

Overview

Successful deployment of Machine Learning (ML) across the sciences calls for both a solid grasp of algorithms and a deep understanding of the application domain. This course brings those elements together in the context of empirical macroeconomic analysis. It pursues two objectives: to give researchers and practitioners the practical skills needed to work with modern ML methods on macroeconomic data, and to cultivate the mindset required to adapt or design algorithms for the unique challenges posed by macroeconomic applications.

We begin with core ML concepts, then progress through classical and modern techniques—from regularized linear models to the architecture of large language models. Emphasis is placed not just on how these tools operate, but on the contexts in which they are most effective. The course devotes particular attention to recent innovations tailored to macroeconomics, including forecasting, nowcasting, latent state extraction, interpretability, density prediction, and other central tasks in applied time-series econometrics.

Prerequisites: Basic econometrics; working knowledge of `R` or `Python`.

Tentative Schedule

Day	Topics
Monday	Introduction to ML; Regularized Linear Models; Cross-validation
Tuesday	Kernel Methods; Tree-Based Models (Random Forests, Boosting)
Wednesday	Neural Networks; Advanced NN Topics (LLMs, Double Descent); Text as Data
Thursday	Macroeconomic Applications; Interpretability; Uncertainty Quantification
Friday (AM)	Causality & Impulse Response Functions; Nowcasting; Wrap-up

Topics

1. **Introduction to Machine Learning**
    - Bias–variance trade-off; generalization

- Prediction vs. causality

2. **Regularized Linear Models**
   - Ridge regression; Lasso; Elastic Net
   - Cross-validation and time-series strategies (blocked CV, validation sets)
3. **Kernel Methods**
   - Feature engineering and basis expansions
   - Kernel Ridge Regression and the kernel trick
4. **Tree-Based Models**
   - Decision trees; bagging; Random Forests; boosting
   - Why tree ensembles work so well
5. **Basic Neural Networks**
   - Structure: layers, neurons, activations
   - Shallow vs. deep networks
   - Training: gradient descent, Adam; regularization
6. **Uncertainty Quantification**
   - Conformal prediction
   - Time-series considerations
7. **Interpretability**
   - Variable importance (permutation); partial dependence; Shapley values
   - Limits of post-hoc methods; inherently interpretable models
   - Dual interpretation via proximity weights (forecasts and causal estimates)
   - Implied factors
8. **Macroeconomic Applications**
   - Macro Random Forest
   - Hemisphere Neural Network (Neural Phillips Curve)
   - Hemisphere Neural Network (Density Forecasting)
   - RF-based adaptive moving average for macro monitoring
   - Deep Time-Varying Parameter VARs
   - Maximally Forward-Looking Core Inflation (generalized ridge)
   - Supervised Detrending Networks (tentative)
   - Assemblage VAR (tentative)
9. **Advanced Neural Networks Topics**

- Why (and why not) use NNs; double descent

10. **Demystifying Large Language Models**
    - Word embeddings and representation learning
    - Attention as regression: connecting transformers to econometric intuition
    - The transformer architecture and its variants
11. **Text as Data for Macroeconomics**
    - Sentiment extraction from central bank communication
    - Text-based indicators for nowcasting and forecasting
    - From bag-of-words to embeddings in economic applications
12. **Causality and Impulse Response Functions**
    - ML-based local projections; extracting/communicating nonlinearities
    - Proximity-weight interpretability
    - Empirical pitfalls: implicit shrinkage; shock overfitting