Machine Learning for Computational Economics

EDHEC Business School — Spring 2026

Instructor: Dejanir Silva **Email:** dejanir@purdue.edu

Location: EDHEC (E-learning room) | Course site: `https://dejanirsilva.github.io/mlce`

Course schedule: Thu 01/15 09:00–12:15 & 14:45–18:00; Fri 01/16 09:00–12:15 & 13:45–18:00; Sat 01/17 09:00–12:15.

Course Description

This course combines classical numerical methods in economics and finance with modern machine-learning approaches for solving and estimating dynamic models. Students learn to connect discrete- and continuous-time dynamic programming with neural-network approximations and gradient-based optimization. The course culminates with the *Deep Policy Iteration (DPI)* algorithm for solving high-dimensional continuous-time models. Each module blends theory, implementation, and coding in Julia/Pluto.

Learning Outcomes

By the end of the five modules, you will be able to:

- Solve dynamic optimization problems in discrete and continuous time.
- Implement stable numerical schemes (finite differences, collocation) and verify convergence.
- Build, train, and interpret shallow/deep neural networks; choose and tune SGD variants (Momentum, RMSProp, Adam/AdamW).
- Apply *Deep Policy Iteration* (DPI) to solve high-dimensional continuous-time models (policy/value nets, drift estimation, control updates without closed-form FOCs).

By the end of the course, students will be able to solve and estimate complex dynamic models that arise in macroeconomics, asset pricing, and corporate finance using modern ML methods.

Prerequisites

Graduate-level macroeconomics or asset pricing. Familiarity with dynamic optimization and basic statistics is expected. All computations are done in Julia. Prior experience is helpful but not required. To prepare, review the "Getting Started with Julia" section of the QuantEcon tutorial (`https://julia.quantecon.org/intro.html`).

Texts & References

The main reference are the lecture notes **"Machine Learning for Computational Economics"** developed for this course. The notes go through the theory and implementation of the methods in detail. It includes a lot of examples and guided implementations of the different methods in Julia.

Additional reading: There are several useful references that complement the lecture notes. The handbook chapter on numerical methods by Fernández-Villaverde, Rubio-Ramírez and Schorfheide (2016) is a great summary of the classical methods for solving dynamic models. For the fundamentals of machine learning, the classic books by Hastie, Tibshirani and Friedman (2009) and Goodfellow, Bengio and Courville (2016) are great references. My discussion of fundamentals of Machine Learning is closer to Prince (2023). The material on the DPI method is based on *"Machine Learning Methods for Continuous-Time Finance"* by Duarte, Duarte and Silva (2024), and the discussion of estimation of high-dimensional diffusion models is based on Duarte, Duarte and Silva (2025).

High-Level Schedule (5 x 3h)

Module	**Theme**	**Core Topics & In-Class Activities**
01	Introduction	Introduction to the course and the main topics. The three curses of dimensionality. Overview of the course.
02	Discrete-Time Methods	Bellman equation; value function iteration; endogenous gridpoint method (EGM); interpolation; *hands-on:* consumption–savings baseline, Tauchen & policy updates.
03	Continuous-Time Methods	From Bellman to HJB; diffusions & jumps; stationary HJB; boundary conditions; FD vs. collocation; viscosity solutions; stability/consistency/monotonicity; *hands-on:* Solving continuous-time consumption–savings problem and option pricing problem.
04	Fundamentals of ML	Supervised pipeline; linear models as baseline; SNN/DNN; activation functions; SGD, Momentum, RMSProp, Adam/AdamW; Lux.jl model building; *hands-on:* fitting DNN to high-dimensional polynomial, breakpoint adaptivity.
05	Deep Policy Iteration (DPI) for CT Finance	Three curses and DPI: (i) drift evaluation in high-d; (ii) NN approximations (value/policy); (iii) control update without closed-form FOCs; implementation loop; *applications:* Lucas orchard, Hennessy–Whited, high-d portfolio choice.

Detailed Module Plan

Module 01: Introduction

- Introduction to the course.
- The three curses of dimensionality.
- Julia language and tools.

Recommended reading: Course notes on *Introduction*.

Module 02: Discrete-Time Methods

- Bellman equation; fixed-point arguments; value vs. policy iteration.
- Approximation: grids, Markov-chain approximation, interpolation.

- Endogenous Grid Method (EGM).
- *Lab:* implement consumption–savings; compare VFI vs. EGM.

Recommended reading: Course notes on *Discrete-Time Methods*.

Module 03: Continuous-Time Methods

- From discrete time to HJB; diffusions and jumps.
- Numerics: FD schemes, stability/consistency, viscosity notion, collocation.
- *Lab:* explicit and implicit schemes, option pricing problem.

Recommended reading: Course notes on *Continuous-Time Methods*.

Module 04: Fundamentals of ML

- SNN/DNN expressivity; universal approximation (width/depth versions).
- Optimization: SGD family (Momentum, RMSProp, Adam/AdamW); regularization.
- Lux.jl basics; training loops; monitoring; generalization checks.
- *Lab:* fit DNN on polynomial target.

Recommended reading: Course notes on *ML Fundamentals*.

Module 05: Deep Policy Iteration (DPI)

- DPI loop: value/policy networks; drift evaluation; argmax without closed-form controls.
- Practicalities: scaling, batching, AD details, stability, and stopping.
- *Applications:* Lucas orchard, Hennessy–Whited corporate model, high-d portfolio choice.
- *Lab:* solve two trees model using DPI method.

Recommended reading: Course notes on *DPI Method* and Duarte et al. (2024).

This syllabus may be updated to reflect pacing and student feedback; any changes will be announced in class and on the course site.

References

Duarte, Victor, Diogo Duarte, and Dejanir H. Silva, "Machine Learning for Continuous-Time Finance," *The Review of Financial Studies*, 2024, *37* (11), 3217–3271.

_ , _ , and _ , "Estimation of High-Dimensional Diffusion Models: A Hyper-Dual Approach," 2025.

Fernández-Villaverde, Jesús, Juan Francisco Rubio-Ramírez, and Frank Schorfheide, "Solution and estimation methods for DSGE models," in "Handbook of Macroeconomics," Vol. 2, Elsevier, 2016, pp. 527–724.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville, *Deep Learning*, MIT Press, 2016.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2 ed., New York, NY: Springer, 2009.

Prince, Simon J. D., *Understanding Deep Learning*, Cambridge, UK: Cambridge University Press, 2023.