UCLA Anderson School of Management

MGMTMFE431-0X – ADVANCED FINANCIAL DATA ANALYTICS AND APPLICATIONS OF AI

[Fall 2025 – 4 Units – In-Person]

Course Syllabus – tentative

REVISED 9/8/2025

Teaching Team Information

Faculty of Record:	Professor Lars A. Lochstoer
Email:	lars.lochstoer@anderson.ucla.edu
Office location:	C4.17 Entrepreneur's Hall
Office hours:	Thursday 10am-11am
Course Assistant:	Kibeom Lee
Email:	kibeom.lee.phd@anderson.ucla.edu
Office location:	A301
Office hours:	Friday 2pm-3pm

Course meeting times and location

Course Day & Time:	Tuesdays from 10/7 to 12/2. Make up classes for 9/30, 11/11 (Veterans Day), and 11/25 (Thanksgiving) will be on 10/2, 11/13, and 12/4, respectively. All classes are from 1pm – 3:50pm.
Course Location:	G304
Course Site:	bruinlearn.ucla.edu

Pre-requisites/Co-requisites

MGMTMFE 402 Econometrics. MGMTMFE 407 Empirical Methods in Finance. MGMTMFE 413 Data Analytics and Machine Learning.

Course Description

Catalog Description

This is an advanced course in data science and AI with applications to finance. Topics include state-space models, the Kalman filter, methods for handling missing data, Markov regime switching models, a variety of deep learning and neural network models, and the use of large language models for model building, coding, signal generation and data analysis. Letter grading.

Expanded Description

AI will change the way finance professionals work with data and code. This class will use LLMs actively for analysis to help students leverage these tools. Examples and homework will focus on finance applications. The first part of the class will consider prompt engineering for model generation and coding. We start with the Kalman filter, state-space models and hidden Markov regime switching models as benchmarks for prediction methods, using LLMs to develop and implement these models. We then move to an in-depth treatment of different classes of neural networks and deep learning, with applications to both prediction and pricing. Examples of the latter include modeling the implied volatility surface and solving value function iteration problems using parallel programming. Finally, the class will study how large language models are constructed and tuned, and how to use them for signal generation. All coding examples will be in Python. All students are assumed to be very familiar with linear algebra, classic econometric methods, and calculus.

Note: This syllabus is tentative and depends on how the class proceeds. New topics could be introduced, and some topics could be dropped.

Course Objectives

At the end of the course, students will be able to:

1. Understand prompt engineering and model building using LLMs.
2. Understand advanced state space models, learning about latent variables, the Kalman filter, and hidden Markov regime switching models
3. Understand the main varieties of neural networks, the attention/transformer mechanisms, and how to use them for both prediction and pricing. The latter includes fitting volatility surfaces and solving value function iteration problems using parallel processing.
4. Understand and leverage large language models for use in analytics and signal generation.
5. Seamlessly move between advanced, but classic, econometric models and state-of-the-art AI tools and choose the best tools for addressing a given task.

Course Materials

Required:

- Course slides and handouts, including academic articles listed in the course outline. Please find the links to these materials on the course site.

Optional:

- Links and references to optional readings will be posted on the course site.
- Additional references that may be useful include:
  - http://faculty.marshall.usc.edu/gareth-james/ISL/ISLR%20Seventh%20Printing.pdf
  - https://web.stanford.edu/~hastie/ElemStatLearn/printings/ESLII_print12.pdf
  - http://www.deeplearningbook.org/
  - J. Durbin and S. J. Koopman, Time Series Analysis by State Space Methods, Oxford Statistical Science Series 38.
  - J. Alammar, M. Grootendorst, Hands-On Large Language Models, O'Reilly publishing, 2024. *See also associated GitHub repository*.
  - A. Bahree, Generative AI in Action, Manning publishing, 2024. *See also associated GitHub repository*.

Course Outline – tentative

Module/ Week	Date	Weekly Title & Key Topics	Pre-Class Reading/Media	Assignments Due
1	10/2/25	Prompt engineering: building econometric models with LLMs – theory and coding o Latent states, learning, the Kalman filter, filtering, smoothing, forecasting o A first example: Estimating inflation dynamics using the Kalman filter o Deriving the Kalman filter, maximum likelihood estimation of state space models o Developing and implementing the model using LLMs	• Lecture Notes, references, articles, and code posted on BruinLearn	
2	10/7/25	State space models: a general benchmark for prediction models • General state-space model, "Everything" is a special case of a state space model • How to handle missing data • Non-linear models: the unscented Kalman filter, hidden Markov regime switching models • Sufficient statistics in prediction • Developing and implementing a regime-switching state-space model with LLMs	• Lecture Notes, references, articles, and code posted on BruinLearn	• Homework 1 due
3	10/14/25	Neural networks and deep learning: why feed forward? • Time-series dependence, sufficient statistics, and feed-forward networks • Can a feed forward network approximate a regime-switching state-space model? Assessment and implementation using LLMs • Recurrent neural networks, long short-term memory networks, autoencoders	• Lecture Notes, references, articles, and code posted on BruinLearn	• Homework 2 due
4	10/21/25	Using neural networks for pricing and optimization; other network structures • Derivative pricing and volatility surface modeling using neural networks • Portfolio choice: solving a value function problem using deep learning • Parallel (GPU) computing using JAX	• Lecture Notes, references, articles, and code posted on BruinLearn •	• Homework 3 due
5	10/28/25	Neural networks and transformers: attention is all you need • AI Asset Pricing Models • Large scale transformer networks and the stochastic discount factor • Risk and return using neural networks • Understanding and visualizing nonlinearities	• Lecture Notes, references, articles, and code posted on BruinLearn •	• Homework 4 due

Module/ Week	Date	Weekly Title & Key Topics	Pre-Class Reading/Media	Assignments Due
6	11/4/25	Understanding Large Language Models • Markov models and word prediction • Word embeddings • Using LLMs in your code, sentiment and summaries • Using LLMs in batch processing and signal generation	• Lecture Notes, references, articles, and code posted on BruinLearn	• Homework 5 due
7	11/13/25	Using LLMs in practice o Creating sentiment signal from earnings call transcripts o Developing code to apply existing LLMs for information gathering and signal generation o Tuning o Creating a ChatBot	• Lecture Notes, references, articles, and code posted on BruinLearn	• Homework 6 due
8	11/18/25	In class midterm	•	• Project outline due
9	12/2/25	The virtue of complexity? • Conventional wisdom: always bet on sparsity • Complexity says increasing nonlinearities a good thing! True? Gaussian kernels and random Fourier transforms • Double-descent and flooding-the-zone approach to prediction	• Lecture Notes, references, articles, and code posted on BruinLearn	
10	12/4/25	Student project presentations		• Project slides and code due
Finals (week)	##/##/##	None		

Evaluation and Grading

Required Assignments and Weighted Percentages

This course will be graded using the following weighted percentages for each of the assignments in the course.

Assignments	% of Grade
Class participation	8%
Homework (6 assignments)	42%
Class project and presentation	25%
Midterm	25%
Total	100%

Grades

Your overall course grade will be determined by how your performance on graded assignments ranks in comparison with other students in the class according to the grade distribution model at Anderson. Note that courses in which an overall grade of C is received must be offset by higher grades in the same term for students to remain in good academic standing at UCLA. A grade of C- or below in a core class must be retaken.

Assignment Descriptions

The following are descriptions of your required assignments for this course. Specific instructions, submission information, and any accompanying rubrics are detailed on the Course Site.

Class participation

Class attendance will be taken and count towards the class participation grade. In addition, every student is expected to contribute to in-class discussions and do their best to answer questions when called upon.

Homework

The homework should be done in groups of four students. The students have to form their own groups. There will be four large homework assignments, best described as cases. Homework will typically involve a combination of some derivation and Python computing. Each homework will count for 7% of the grade. You may ask the teaching assistant clarifying questions regarding the homework.

Class project and presentation

The homework group has to be the same as the project group. There cannot be collaboration across groups. The group project should use techniques covered in the class to analyze a relevant problem in finance. The groups must acquire data for their project using online resources that are available to them at Anderson. The project must leverage an LLM or neural network/deep learning for a finance related analytics application. Each group must have their project idea approved by the professor by Week 6. The groups will present their projects in class. Every student is required to do a part of the presentation, and students are required to attend their peers' project presentations.

Midterm

There will be an in-class midterm in week 8 of the class.

Anderson and Course Policies

AI Usage Policy

This course will actively use AI and it is allowed as a tool to help in any aspect of the class, with the exception of the midterm.

Netiquette

The written language has many advantages: more opportunity for reasoned thought, more ability to go in-depth, and more time to think through an issue before posting a comment. However, written communication also has certain disadvantages, such a lack of the face-to-face signaling that occurs through body language, intonation, pausing, facial expressions, and gestures. As a result, please be aware of the possibility of miscommunication and compose your comments in a positive, supportive, and constructive manner.

UCLA Policies

Code of Conduct

All participants in the course are bound by the **UCLA Student Conduct Code** (https://deanofstudents.ucla.edu/individual-student-code) and **UCLA Anderson Honor Code** (https://www.anderson.ucla.edu/documents/areas/adm/web/AndersonHonorCode.pdf).

Academic Integrity

UCLA is an institution of learning, research, and scholarship predicated on the existence of an environment of honesty and integrity. As members of the academic community, instructors, students, and administrative officials are all responsible for maintaining this environment. It is essential that all members of the academic community practice academic honesty and integrity and accept individual responsibility for their work. Academic misconduct is unacceptable and will not be tolerated in this course. Cheating, forgery, dishonest conduct, plagiarism, and collusion in academic misconduct erode the University's educational, research, and social roles.

Students who knowingly or intentionally conduct or help another student engage in acts that violate UCLA's expectations of academic integrity will be subject to disciplinary action and referred to the Dean of Students' Office.

Please familiarize yourself with **UCLA's Academic Integrity Policy:** https://www.deanofstudents.ucla.edu/Academic-Integrity. Speak to your instructor if you have any questions about what is and is not allowed in this course.

Integrity in Research

Integrity in research includes not just the avoidance of wrongdoing, but also the rigor, carefulness, and accountability that are hallmarks of good scholarship. All persons engaged in research at the University are responsible for adhering to the highest standards of intellectual honesty and integrity in research.

Please familiarize yourself with the **University of California Policy on Integrity in Research** (https://www.ucop.edu/academic-personnel-programs/_files/apm/apm-190-b.pdf)

Accessible Education & Inclusive Education

Disability Services

UCLA is committed to providing a barrier-free environment for persons with documented disabilities. If you are already registered with the Center for Accessible Education (CAE), please request your Letter of Accommodation in the Student Portal. If you are seeking registration with the CAE, please submit your request for accommodation via the CAE website. Students with disabilities requiring academic accommodations should submit their request for accommodations as soon as possible, as it may take up to two weeks to review the request. For more information, please visit the CAE website (www.cae.ucla.edu), visit the CAE at A255 Murphy Hall, contact CAE by phone at (310) 825-1501, or by telecommunication device for the deaf at (310) 206-6083.

Equity, Diversity, and Inclusion

Please familiarize yourself with UCLA Anderson's commitment to maintaining an equitable, diverse, and inclusive community:
(https://www.anderson.ucla.edu/about/equity-diversity-and-inclusion)