Proposal: New 3xx ML Courses

Vasilis Charisopoulos, Lillian Ratliff, Maryam Fazel, Nathan Kutz, Banghua Zhu

Context

- There is high demand for machine learning courses in the undergraduate ECE curriculum, especially at the junior level.
- Existing ECE courses are all at the senior level, and an extensive amount of time is focused on covering common introductory topics which could more efficiently and effectively be covered in junior level offerings.
- ECE can also use these courses to attract non-ECE majors
- To address this gap, we are proposing two new complementary 300 level courses to begin Fall 2025.
    - 344: Data-Driven Modeling & Machine Learning
    - 345: Introduction to Foundations of Machine Learning (with 445 as the follow-up course in the sequence)

Integration with ML Pathway

Pre-requisites:
- Calc sequence (Math 124-126, or honors equivalent)
- Python: EE241 or CSE 163 (EE235 also suffices)

345: Introduction to Foundations of ML

- **Summary**: foundations of core machine learning techniques such as regression and classification, and linear algebra methods for feature analysis and dimensionality reduction.
- **Focus**: theoretical underpinnings of ML
- Intended for ML/DS pathway majors; open to all ECE and non-ECE majors

Offered once per year; enrollment cap 80

Recommended to take simultaneously or in same year

344: Data-Driven Modeling & ML

- **Summary**: provide students with a working knowledge of applied ML concepts and techniques via hands-on data science
- **Focus:** computation and implementation of DS/ML methods
- Intended for any ECE major; open to non-ECE majors

Offered multiple times per year; enrollment cap 80

Pre-req. for DS/ML pathway[3]

Current Advanced Course Offerings[2]:

- 443: Machine Learning for Signal Processing
- 445: Foundations of Optimization and Machine Learning[1]
- 467: Machine Learning for Cybersecurity

[1]We are revising the MCD for this course to emphasize more advanced machine learning methods as well as convex optimization, modeling and algorithms.
[2]New courses such as 'Robotics and Reinforcement Learning', 'Introduction to Foundation Models', etc. to be added eventually
[3]This is more of a suggestion as it would require changing the other MCDs. It will be a prerequisite for 445; changes to this MCD are forthcoming.

Justification: 345 Introduction to Foundations of Machine Learning

The course (345) “Introduction to Foundations of ML” is intended to be an introductory course for the ML pathway, a course the ECE curriculum currently lacks. It aims to provide students with the requisite theoretical foundations of ML as needed for more advanced courses. The concept for 345 emerged out of prior instantiations of 445 “Foundation of Machine Learning and Optimization,” the curriculum of which aims to cover both core ML methods *and* convex modeling as well as more advanced ML methods. Given students preparation upon entering 445, we learned that the agenda is extremely ambitious and an ideal offering would consist of a introductory course and an advanced course.[1]

Content & Outcomes Overview: As the introductory course, 345 will focus on foundations of core ML techniques (such as regression, classification, and clustering) and linear algebra methods for feature analysis and dimensionality reduction. This course will serve as an entry point to the DS/ML pathway by providing the students with the primary skills needed for more advanced courses. Students will be able to 1) understand and utilize linear algebra techniques, including matrix decomposition, for efficient feature extraction and analysis, and dimensionality reduction, 2) demonstrate an understanding of the theoretical underpinnings of quintessential ML methods including linear regression and classification, and simple clustering techniques, and 3) identify and reason about the appropriate use of supervised versus unsupervised ML approaches in real-world data science tasks. The class will be grounded in ECE and ECE-adjacent application domains including control and robotics, power systems, neuroscience, and signal processing.

Demand for ML: There is a high demand for ML classes in ECE at the undergraduate level. At present, all the ECE ML/DS classes are at the senior level. In these courses, an extensive amount of time is focused on covering common introductory topics which could more efficiently and effectively be covered in junior level offerings. This course is aimed at addressing demand and making the ML/DS more cohesive and efficient.

Relationship to 344: This course is complementary to 344 “Data-Driven Modeling and Machine Learning,” which focuses on computational aspects of ML via brief introductions to practical implementation of a broad array of ML methods. Students are encouraged to take both courses either in parallel or sequentially in their Junior year. This course is primarily aimed at ML pathway ECE majors but is open to all ECE and non-ECE majors.

[1]The MCD for 445 is under revision to reflect these changes as well.

Justification: Data-driven modeling and ML (344)

The proposed course, Data-Driven Modeling and Machine Learning, is aimed at providing students with working knowledge of applied ML concepts and techniques. It complements the proposed EE-345, as the latter focuses on theoretical underpinnings of ML methods. The concept emerged after a successful prior offering (as a special topics course) by J. Nathan Kutz, combined with increasing demand for ML classes at the undergraduate level.

Focus: This course focuses on computational methods as “tools” for different engineering tasks, rather than their theoretical underpinnings. It will offer students hands-on experience implementing and evaluating ML methods in Python and guided opportunities to develop software literacy with tools such as git/github, test-driven software development, LLM-based coding assistants etc. In addition to regular homeworks, students will **(i)** develop a personal portfolio page hosted on Github pages or a similar static website hosting service and **(ii)** complete a coding-based final project (in lieu of a “traditional” final exam). This project will be an opportunity for a deeper dive into a particular ML method, task, or dataset.

Outcomes: Ideally, students completing the course will be able to incorporate advanced ML methods and modern software development skills into the broader UW ECE curriculum, including machine learning courses offered at the senior level, capstone projects, etc. Beyond coursework, students interested in research experience - i.e., working with UW ECE faculty that employ ML methods in their own research - can even use the final coding project as an opportunity to explore datasets or methods used by the particular lab they seek to join.

Finally, there is a high demand for ML classes in ECE at the undergraduate level. At present, all the ECE ML/DS classes are at the senior level. In these courses, an extensive amount of time is focused on covering common introductory topics which could more efficiently and effectively be covered in junior level offerings. This course is aimed at addressing demand and making the ML/DS more cohesive and efficient.

Justification: Foundations of Optimization and Machine Learning (445)

This is a revised version of the existing EE 445 course, now as a second course in the 345-445 sequence. The topic of convex optimization models is new in this course, building upon the introductory models studied in 345 and discussing more advanced machine learning problems.

In addition to covering the mathematical foundations, we will build convex models using the modeling tool CVX/CVXPY and solve practical convex optimization problems with black-box solvers (hidden “under the hood”), while implementing selected algorithms in Python (such as gradient descent and variations).

Overview: EE 445 is an introduction to optimization and advanced machine learning models motivated by their application in areas including statistics, decision-making and control, and communication and signal processing. Topics include convex sets and functions, convex optimization problems and convex modeling, duality, linear and quadratic programming, and basic algorithms such as gradient descent, with emphasis on supervised and unsupervised machine learning problems (e.g., regularized regression, robust classification, clustering, and logistic regression), building upon EE 345 concepts.

Relationship to 344 and 345: This course follows the 3xx courses. Only 345 is a prerequisite; taking 344 before is helpful but not required.