Master Course Description for EE-445

Title: Fundamentals of Optimization and Machine Learning

Credits: 4

Course Catalog Entry:

EE 445: Fundamentals of Optimization and Machine Learning is an introduction to optimization and 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 their properties, convex modeling, duality, linear and quadratic programming, with emphasis on usage in machine learning problems including regularized linear regression and classification.

Coordinators:

- Lillian Ratliff, Assistant Professor, Electrical and Computer Engineering
- Maryam Fazel, Professor, Electrical and Computer Engineering

Goals: To give ECE students the foundational mathematical concepts and theory that underpins modern optimization and machine learning algorithms. Provide a background in mathematical reasoning, and convex problem modeling and solving. Develop a mathematical understanding of how convex optimization tools are used in the design, and analysis of machine learning algorithms and optimization problems used in various ECE application domains including data science, decision-making and control, communication, and signal processing.

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

- 1. Identify and characterize convex sets, functions, and optimization problems.
- 2. Develop skills to model applied problems as convex optimization problems.
- 3. Gain experience with the modeling environment CVX, implement simple optimization methods such as gradient descent in Python.
- 4. Model basic machine learning algorithms using the language of convex optimization.

Textbook:

- Main: Optimization Models in Engineering (Giuseppe Calafiore and Laurent El Ghaoui)
- Supplementary: Convex Optimization (Stephen Boyd and Lieven Vandenberghe)
- Supplementary: Introduction to Applied Linear Algebra: Vectors, Matrices, and Least Squares (Stephen Boyd, Lieven Vandenberghe)

Prerequisites by Topic:

- 1. Calculus sequence: Math 224 or Math 324
- 2. Linear Algebra: Math 208 or Math 308 or Math 136 or AMATH 352
- 3. Python: EE 241 or EE 235 or CSE 163

Course Structure Overview:

1. Module-1: Review of Mathematical Foundations

Module-1 concentrates on the review of the primary mathematical tools used in convex optimization and machine learning as relevant for this course. Topics include:

- Vectors, function properties, and norms
- Matrices, eigenvalue decomposition
- Symmetric matrices, positive semi-definite matrices, singular value decomposition, and principal component analysis (PCA)

Applications will be used to introduce the concepts above and reinforce their importance:

Examples of applications include:

- Machine learning: learning from data, discovering patterns and structure in data, dimensionality reduction
- Control and Signal Processing with Applications in Neuroscience: image compression and facial recognition, identifying neurons from the shape of its action potential, spike-triggered covariance analysis
- Quantitative Finance: modeling and analysis of the shape of a yield curve, portfolio asset analysis, interest rate modeling

2. Module-2: Least Squares Regression in Machine Learning

Module-2 concentrates on one of the quintessential tools in machine learning, namely least square regression. Topics include:

- Linear equations and least squares
- Least squares data fitting
- Regularized least squares, Ridge regression, Kernel methods
- Generalization and cross validation

As in Module-1, applications will be used throughout the module to motivate the different methods and topics. Examples of applications include:

- Machine Learning: regression, classification, ranking, feature selection
- Control and Signal Processing: linear prediction, smoothing, estimating missing data, filter design

• Quantitative Finance: forecasting portfolio returns, portfolio asset management 3. Module-3: Introduction to Convex Analysis and Optimization

Optimization is at the core of every machine learning model. Module-3 concentrations on convex analysis, modeling and optimization with connections to its use in machine learning. In particular, it will be demonstrated that machine learning problems and algorithms arising in different domains can be modeled using the language of convex optimization. Topics include:

- Convex sets and functions, convex conjugate duality
- Convex optimization problems, gradient descent
- Lagrangian, duality in convex optimization (weak/strong duality) optimality conditions including Karush-Kuhn-Tucker conditions
- Linear and quadratic programs

Applications will be used throughout this module to convey concepts. Examples of applications include:

- Machine Learning: solving least squares via gradient descent, Maximum A Posteriori (MAP) inference via linear programming and duality
- Control and Signal Processing with Applications in Game Theory: finding Nash equilibria in matrix games via linear (zero-sum) and quadratic (general sum) programming, system identification for ARX/ARMA models.
- Quantitative Finance: risk assessment via linear programming, meanvariance analysis in portfolio selection and asset pricing

4. Module-4: Applications

Module-4 combines concepts from the previous three modules by revisiting the example applications from Modules 1-3 in greater detail. The applications will be the primary focus and connections will be drawn to different aspects of the modeling, data analysis and solutions (optimization problem or algorithm) as they relate to the concepts from Modules 1-3. Example applications include:

- Machine Learning: LASSO and kernel methods
- Control and Signal Processing: sparse signal reconstruction, linear quadratic control design
- Quantitative Finance: hedging interest rate sensitivity of a portfolio, asset allocation, interest rate simulation using Maximum Likelihood Estimation (MLE)

Course Structure: The class meets for two 1 hour 20-minute lectures and one 2 hour discussion section per week. The latter is administered by teaching

assistants. Homework (with theoretical and computational components) is assigned weekly. One exam is given nominally at the end of the 5th week, and a comprehensive final exam is given at the end of the quarter.

Computer Resources: The course uses Python for the computational components of the homeworks, and some use of the modeling environments CVX (which can be called within Matlab) or CVXPY (which uses Python). Students are expected to use their personal computers.

Laboratory Resources: None.

Grading: Approximate distribution: Homework 35%, Midterm Exam 25%, Final Exam 40%. The grading scheme in any particular offering is the prerogative of the instructor.

ABET Student Outcome Coverage: This course addresses the following outcomes:

H = high relevance, M = medium relevance, L = low relevance to course.

(1) An ability to identify, formulate, and solve complex engineering problems by applying principles of engineering, science, and mathematics (H) The homework and exams require direct application of mathematical knowledge to engineering problems, and require students to model engineering problems in the language of convex optimization.

(3) An ability to communicate effectively with a range of audiences (L) Students will learn and apply techniques to rigorously and formally apply and communicate theoretical concepts.

(7) An ability to acquire and apply new knowledge as needed, using appropriate learning strategies.

Religious Accommodation Policy:

Washington state law requires that UW develop a policy for accommodation of student absences or significant hardship due to reasons of faith or conscience, or for organized religious activities. The UW's policy, including more information about how to request an accommodation, is available at Religious Accommodations Policy

(https://registrar.washington.edu/staffandfaculty/religious-accommodations-policy/).

Accommodations must be requested within the first two weeks of this course using the Religious Accommodations Request form (https://registrar.washington.edu/students/religiousaccommodation request/).

Prepared By: Lillian Ratliff, Maryam Fazel Last Revised: 2/21/21

Textbook table of contents

- 1. Introduction
 - 1. Motivating examples
 - 2. Optimization problems
 - 3. Important classes of optimization problems
 - 4. History
- I Linear Algebra Models
 - 2. Vectors and functions
 - 1. Vector basics
 - 2. Norms and inner products
 - 3. Projections onto subspaces
 - 4. Functions
 - 5. Exercises
 - 3. Matrices
 - 1. Matrix basics
 - 2. Matrices as linear maps
 - 3. Determinants, eigenvalues, and eigenvectors
 - 4. Matrices with special structure and properties
 - 5. Matrix factorizations
 - 6. Matrix norms
 - 7. Matrix functions
 - 8. Exercises
 - 4. Symmetric Matrices
 - 1. Basics
 - 2. The spectral theorem
 - 3. Spectral decomposition and optimization
 - 4. Positive semidefinite matrices
 - 5. Exercises
 - 5. Singular Value Decomposition
 - 1. Singular value decomposition

- 2. Matrix properties via SVD
- 3. SVD and optimization
- 4. Exercises
- 6. Linear equations and least squares
 - 1. Motivations and examples
 - 2. The set of solutions of linear equations
 - 3. Least-squares and minimum-norm solutions
 - 4. Solving systems of linear equations and LS problems
 - 5. Sensitivity of solutions
 - 6. Direct and inverse mapping of a unit ball
 - 7. Variants of the least squares problem
 - 8. Exercises
- 7. Matrix Algorithms
 - 1. Computing eigenvalues and eigenvectors
 - 2. Solving square systems of linear equations
 - 3. QR factorization
 - 4. Exercises
- **II** Convex Optimization Models
 - 8. Convexity
 - 1. Convex sets
 - 2. Convex functions
 - 3. Convex problems
 - 4. Optimality Conditions
 - 5. Duality
 - 6. Exercises
 - 9. Linear, quadratic, and geometric models
 - 1. Unconstrained minimization of quadratic functions
 - 2. Geometry of linear and convex quadratic inequalities
 - 3. Linear programs
 - 4. Quadratic programs
 - 5. Modeling with LP and QP

- 6. LS-related quadratic programs
- 7. Geometric programs
- 8. Exercises
- 10. Second-Order cone and robust models
 - 1. Second-order cone programs
 - 2. SOCP-representable problems and examples
 - 3. Robust optimization models
 - 4. Exercises
- 11. Semidefinite Models
 - 1. From linear to conic models
 - 2. Linear matrix inequalities
 - 3. Semidefinite programs
 - 4. Examples of SDP models
 - 5. Exercises
- 12. Introduction to Algorithms
 - 1. Technical preliminaries
 - 2. Algorithms for smooth unconstrained minimization
 - 3. Algorithms for smooth convex constrained minimization
 - 4. Algorithms for non-smooth convex optimization
 - 5. Coordinate descent methods
 - 6. Decentralized optimization methods
 - 7. Exercises

III Applications

- 13. Learning from Data
 - 1. Overview of supervised learning
 - 2. Least-squares prediction via a polynomial model
 - 3. Binary classification
 - 4. A generic supervised learning problem
 - 5. Unsupervised learning
 - 6. Exercises
- 14. Computational Finance

- 1. Single-period portfolio optimization
- 2. Robust portfolio optimization
- 3. Multi-period portfolio allocation
- 4. Sparse index tracking
- 5. Exercises
- 15. Control Problems
 - 1. Continuous and discrete time models
 - 2. Optimization-based control synthesis
 - 3. Optimization for analysis and controller design
 - 4. Exercises
- 16. Engineering Design
 - 1. Digital filter design
 - 2. Antenna array design
 - 3. Digital circuit design
 - 4. Aircraft design
 - 5. Supply chain management
 - 6. Exercises