Harvard University
School of Engineering and Applied Sciences

ENG-SCI 201/APMTH 231: Decision Theory

Spring 2026
Course Information

Course information

Instructor:
Demba Ba. (demba@seas.harvard.edu).
150 Western Ave, Room SEC 3.308
Boston, MA 02134
(617) 495-1228 [Voice]
Office Hours: T 9:00-10:00 AM, 4:00-5:00 PM.

Teaching Fellows:
Shubham Choudhary. (shubham_choudhary@g.harvard.edu).
Sumedh Hindupur. (shindupurravindra@g.harvard.edu).
Section: W 5:00–6:00 PM, MD 119.
Office Hours: W 6:00–7:00 PM, MD 119; F 1:00–2:00 PM, SEC 3.314.

Grader:
Mozes Jacobs. (mozesjacobs@g.harvard.edu).
Office Hours: Th 2:00–3:00 PM, SEC 3.314.

Lecture:
T/Th 11:15 AM–12:30 PM, SEC 1.413.

Course Administrator:
Sarah Gayer. (sgayer@g.harvard.edu).

Course Website:
`https://canvas.harvard.edu/courses/163605/`

Course overview

ES 201/AM 231 is a course in statistical inference and estimation from a signal processing perspective. The course will emphasize the entire pipeline from writing a model, estimating its parameters and performing inference utilizing real data. The first part of the course will focus on linear and nonlinear probabilistic generative/regression models (e.g., linear, logistic, Poisson regression), and algorithms for optimization (ML/MAP estimation) and Bayesian inference in these models. We will pay particular attention to sparsity-induced regression models, because of their relation to artificial neural networks, the topic of the second part of the course. The second part of the course will introduce students to the nascent and exciting research area of model-based deep learning and sparse auto-encoders. We will see, for instance, how neural-networks with ReLU nonlinearities arise from sparse probabilistic generative models introduced in the first part of the course. This will form the basis for a rigorous recipe we will teach you to build interpretable deep neural networks, from the ground up. More broadly, model-based deep learning and sparse auto-encoders have become popular approaches to reverse-engineer intelligence in both biological and artificial settings: in each case, we are able to train these systems to perform complicated tasks, but our understanding of how they do so remains opaque.

Reverse engineering intelligence—whether in the brain or in artificial neural networks—means systematically probing what features, patterns, or concepts high-performing systems are actually representing and how. For example: what aspects of "the world" does a population of neurons encode? How does an LLM organize concepts? We will use statistical, optimization, and computational tools to shed light on these questions, moving beyond just performance to true interpretability. Importantly, this quest is not only about intellectual curiosity: greater understanding brings transparency, control, and trust to real-world systems—enabling safer brain-machine interfaces, fairer decision-making, and more socially responsible deployment of large-scale AI.

We will invite an exciting lineup of speakers. We encourage you to pursue a final project that could lead to prototype or production solutions to challenges businesses face around AI adoption due to lack of transparency.

What makes this course unique?

- **Fundamentals, fundamentals, fundamentals:** Linear and nonlinear regression, optimization, kernel methods, SVMs, and regularization—the core concepts underpinning statistical inference: you will get challenged mathematically!
- **Frontiers:** Model-based deep learning, sparse autoencoders, and recent advances in interpretable AI, bridging theory and emerging technologies.
- **Application/Prototype:** Final projects often go beyond theory to tackle real-world data and transparency challenges, including the opportunity to develop production-level solutions and product prototypes.

Prerequisites

The official prerequisites for this course are APPLIED MATH 21a or MATH 21a, and STATS 110 or equivalents. This is a highly interdisciplinary graduate-level course that will involve a combination of theory *and* computational modeling that are both motivated by data analysis problems. The key requirements are intellectual curiosity and a desire to learn to think about data in new ways. A certain level of mathematical maturity is assumed. In particular, prior exposure to abstract linear algebra and real analysis will deepen your understanding of the materials.

Textbook

Machine Learning: A Bayesian and Optimization Perspective, by Sergios Theodoridis, Academmic Press, 2015.

Policy on collaboration

To get the most out of this course, you are encouraged to struggle with the course assignments on your own and reach out to the course staff during Office Hours. **You are allowed to *discuss* the *content* of assignments with fellow students but *not* their *solutions***. Your write-up of assignments must entirely be your own. Moreover, at the top of every assignment, you are kindly asked to acknowledge the students you have discussed said assignment with. We also ask you to acknowledge the use of books, articles, websites, lectures, discussions, etc., that you have consulted to complete your assignments.

Policy on Generative AI

I encourage you to use generative AI tools to, for instance, clarify your understanding of topics from class. In addition, if you find them useful, I encourage you to experiment with them to work on your problem sets. The course assignments have as a goal to teach you certain thinking processes/ways of tackling problems. In that sense, the actual solution to the problem itself matters very little to us. That's why, if you do use generative AI tools for your problem sets, we ask that you, for every problem, to submit the sequence of prompts and answers from the AI tool that led you to a solution.

Grading information

The final grade for this course will be based on your performance on problem sets, a midterm examinations and a final project/paper. There will not be a final exam in this class.

1. **Problem sets**: There will be **4** (roughly) bi-weekly **problem sets** that will count towards **45% of your final grade**. Problem sets will be due at the

beginning of class on the date stated in the course calendar. Problem sets will consist of "pen-and-paper assignments" and/or computational assignments, which you will submit electronically.

2. **Midterms**: There will be **1 midterm examination** that will count towards **30% of your final grade**.

3. **Final project/paper**: You can choose to work on a **final project** or **final paper**, which will count towards **15% of your final grade**.

   - The final project can either be performed alone or in a group of up to three people. All members of a group will get the same grade on the final project. The bigger the group, the higher the expected output of the project. The goal of the final project is three-fold: we ask that you (a) formulate an interesting question that involves real data, (b) gather the data required to answer this question and (c) used concepts from the course to tell an interesting story, and (d) give a (10 to 15-minute) presentation. The grade will be based on your ability to apply the concepts and tools taught in class, and your ability to integrate (a), (b), (c) and (d). We strongly encourage that you consult with the course staff as early as possible. You can choose to pursue one of two tracks

   Track 1 : Traditional Research
     - Answers/explores a research-style question
     - Can use external data sources and libraries
     - Focus on methodology and analysis

   Track 2 : Prototype Development
     - Builds a solution to a transparency challenge faced by a business
     - Creates something pitchable/demonstrable
     - Can use external data sources or APIs

     For this track, your goal is to create a compelling demonstration—a proof-of-concept or prototype is sufficient to show your approach in action. The course staff will be happy to advise on possible frameworks or technologies, so students from all backgrounds can participate meaningfully in this applied dimension.

   - We will treat the final paper as a single-person assignment. We ask that you formulate a question based on a data-driven problem of your choice, that can utilize the tools for statistical inference we will learn in the class. This can either come from a problem related to your undergrad/grad thesis, or simply a problem that interests you. We ask that you explain the problem/question, survey the current literature on approaches towards solving the problem, suggest new ways to solve the problem based on (a) tools you will learn in the first part of the class, and (b) using model-based deep learning approaches from the second part of the class. We ask that you submit the paper in 4-page, two-column, IEEE conference format. You can use an additional page for references, for a total of 5 pages.

4. **Project/paper proposal**: We ask that you put together a short one-page project/paper proposal to be discussed with the course staff following Spring Recess (see course calendar). The project/paper proposal will count towards **10% of your final grade**.

Policy on Late Assignments

Except for the exams and the final-project report, you get to turn in one assignment up to three days late, no questions asked.

Disclaimer

While the above weights are used for computing the final grade, I reserve my right to scale the grades based on the performance of the entire class.

Table 1: **Course Calendar**

	Date	Topic	Assignments
		Regression, Classification and Optimization	
T	01/27	Course overview	Lec. 1 slides
Th	01/29	Maximum likelihood estimation	**Pset 1 out**, Lec. 2 notes
T	02/03	Linear regression	Lec. 3 notes
Th	02/05	Intro to Convex Optimization I	Lec. 4 notes
T	02/10	Intro to Convex Optimization II	Lec. 5 notes
Th	02/12	Exp. family and weighted least-squares	**Pset 1 due, Pset 2 out** Lec. 6 notes
T	02/17		
Th	02/19	SVMs: classification meets optimization	Lec. 7 notes
T	02/24		
Th	02/26	Features space, kernels, RKHS	Lec. 8/9 notes
T	03/03		**Pset 2 due, Pset 3 out**
		Regularization, hierarchical models and Bayesian thinking	
Th	03/05	MAP estimation/regularized regression	Lec. 10 notes **Pset 2 due, Pset 3 out**
T	03/10	Cross-validation/parameter selection	Lec. 11 notes
Th	03/12	**Midterm 1**	
Spring Break: Saturday March 14 – Sunday March 22			
T	03/24	Perceptron, ANNs and learned feature extraction	**Pset 3 due, Pset 4 out** Lec. 12 notes
Th	03/26	Optimizing ANNs: back-propagation algorithm	
T	03/31	Model-based networks: unrolling, (D. Ba) optimization layers, sparse auto-encoders	**Project proposal due** Reading assignment
Th	04/02	Sparse auto-encoders in neuroscience (B. Tolooshams)	**Reading assignment**
T	04/07	Sparse auto-encoders in AI (T. Fel)	**Reading assignment**
Th	04/09	State-space models and Kalman filtering	**Pset 4 due** Lec. 15/16 notes
T	04/14	Kalman filtering, RNNs	
Th	04/16	LLMs, transformers and RNNs (TBD)	
T	04/21	Generative modeling: VAEs, diffusion models, norm. flows (TBD)	
Th	04/23		**Reading assignment**
T	04/28	**Project presentations**	
Th	04/30	**Project presentations**	
Mon	05/04	**Final project due**	