Statistical Inference and Learning for Engineers (Syllabus)

This is a course about statistical inference (inferring the values of some quantity from noisy observations) aimed at graduate students in electrical and computer engineering. Although the traditional applications for this material are in signal detection and estimation in communications and DSP applications (at least for ECEs), the class is better thought of as a course in theoretical statistics and (some) machine learning.

1 Administration

The instructor for this class is:

Prof. Anand Sarwate
CoRE 517
anand.sarwate@rutgers.edu
https://adsarwate.github.io/

The best method of contact is email or in person before or after class.

1.1 Prerequisites

There is one major prerequisites for the course:

- ECE 541 (Stochastic Signals and Systems)

A solid background in probability is required for this course. Students would also benefit from a background in linear algebra and optimization.

Students who are not in the ECE department are welcome to attend assuming they have the right background in probability and linear algebra.

1.2 Statement on inclusiveness

We must treat every individual with respect. We are diverse in many ways, and this diversity is fundamental to building and maintaining an equitable and inclusive campus community. Diversity can refer to multiple ways that we identify ourselves, including but not limited to race, color, national origin, language, sex, disability, age, sexual orientation, gender identity, religion, creed, ancestry, belief, veteran status, or genetic information. Each of these diverse identities, along with many others not mentioned here, shape the perspectives our students, faculty, and staff bring to our campus.

In engineering, accounting for and incorporating these diverse perspectives fuels excellence and innovation. In turn, our goal is to develop the tools to let us build technologies to support and sustain a more just and equitable future for all. Accommodations

Rutgers University welcomes students with disabilities into all of the University's educational programs. In order to receive consideration for reasonable accommodations, a student with a disability should contact the Office of Disability Services (ODS, `https://ods.rutgers.edu/`), participate in an intake interview, and provide documentation. ODS may then provide you with a Letter of Accommodations. Share this letter as soon as you can with the instructors for your classes and discuss the accommodations if needed.

2 Course material

2.1 Learning outcomes

In this course, students should develop knowledge and skills that are integral to reading and writing theoretical and algorithmic research that uses probability theory. In particular, by the end of this course, students should be able

- to understand theoretical papers which use the mathematical tools and techniques we will learn and discuss,
- to apply key results and mathematical tools to analyze algorithms and experiments in a variety of applications,
- to write a clear and rigorous mathematical proof using these tools (as part of the exercises),
- to present/teach the key mathematical ideas and arguments used in a research paper to their peers,
- to work effectively with peers from different labs/research groups by learning a common mathematical "language", and
- to practice methods for investigating and explaining to themselves and others advanced mathematical tools which they may find useful in research.

3 Textbook

The textbook for the course will be listed as SIEDS for short:

- P. Moulin and V.V. Veeravalli, Statistical Inference for Engineers and Data Scientists, Cambridge University Press, 2019. [1]

There are other (older) books which are also useful references:

- Harry L. Van Trees, Detection, Estimation, and Modulation Theory (Part I), Wiley-Interscience, 2001.
- Bernard C. Levy, Principles of Signal Detection and Parameter Estimation, Springer 2008.
- Peter J. Bickel and Kjell A. Doksum, Mathematical Statistics, Basic Ideas and Selected Topics, Vol. 1, (2nd Edition), Pearson, 2006.
- H. Vincent Poor, An Introduction to Signal Detection and Estimation, 2nd Edition, Springer, 1998.
- Steven M. Kay, Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory, Prentice Hall, 1993.
- Steven M. Kay, Fundamentals of Statistical Signal Processing, Volume II: Detection Theory, Prentice Hall, 1998.

This class is offered in a semi "flipped" format. Videos recorded from 2021 are available on YouTube.

3.1 Schedule of topics

This is a tentative list of topics to be covered in the semester.

Week	Topics	Reading
1	Statistical decision theory, likelihood, risk	SIEDS 1.1-1.6
2	Intro to hypothesis testing	SIEDS 2.1-2.2, 1.7, 2.3-2.3.1, 2.4
3	Neyman-Pearson, asymptotics	SIEDS 2.4, 3.1-3.2, 6.1, 7.1-7.4, 5.1-5.7
4	Extensions of hypothesis testing	SIEDS 3.1-3.2, 3.5, 4.1-4.2, 4.5, 5.5
5	FDR and sequential detection	SIEDS 3.5, 9.1
6	MMSE Estimation	SIEDS 11.1-11.7
7	MMAE and exponential families	SIEDS 11.7.4 - 11.8
8	Frequentist estimation	SIEDS Chapter 12
9	Rao-Blackwell and bias-variance tradeoff	SIEDS Chapter 12
10	Cramér-Rao lower bound	SIEDS Chapter 13
11	MLE asymptotics	SIEDS 13.5, 14.1-14.8
12	EM Algorithm and sequential estimation	SIEDS 14.11, 14.12
13	Predictive models and classical machine learning	SIEDS 14.11, 14.12
14	Final project presentations	

4 Assessment

This is a graduate course and we expect students to engage actively and meaningfully with the material.

4.1 Components of the grade

This course will be graded on:

- 15% homeworks (5)
- 30% quizzes (5)
- 25% final project
- 30% take-home final exam

The homeworks are intended to be challenging and are "formative" assessments. The quizzes test the knowledge from the previous homework and are "summative" assessments.

4.2 Final project

The final project for this class (or term project) can take one of many different forms. In the end, however, you will produce a document of around 10 pages in length (allowing for equations, including references) or so. That is not a hard-and-fast number, but more to indicate that it should be a substantive project, not just a homework assignment. The report should be able to be read and understood by another student in the class. Think of it as writing lecture notes that could be used in next year's class. A prototypical report structure:

1. Introduction to the problem and its context (what is the application)
2. Description of the technical/mathematical problem/challenge as it relates to the class.
3. A solution of the problem (or an approach to solving the problem)
4. Worked out analytical examples or numerical examples to illustrate the problem.

Some example approaches to a project include:

- Learn about an application of the material in this class to teach it to other students in the class. The important components of this type of project are that you find a specific application, describe how a topic or topics from the class can be applied to the problem, and then give an application/example worked out analytically and maybe numerically as well.

- Apply the ideas in the class to some aspect of your own research. You shouldn't just present your research, but instead find a related problem that you can analyze and explain to other students in the class through your report.
- Find a recent (or not so recent) research paper on a topic related to the class and dig a deeper. For example, if it is an algorithm, implement the algorithm and apply it to data that is different from what appears in the paper. Write a report that explain the problem, the solution, and the experiment, and write it at a level that can be understood by other students in the class.
- As a last resort, if you cannot find a project report that looks like one of these, you can write a literature review around a technical problem related to the class. However, this literature review will have to be comprehensive and will involve you downloading and reading (not the full paper but at least up to the main results) at least 20-40 papers. Of these options, this will be the most work.

The project in total is worth 25

- The proposal is worth 5
- The interim report is worth 5
- The final project report is worth 15

4.3 Academic integrity

Please be sure to follow the Academic Integrity guidelines. These are very important, especially for the project: be sure to cite your sources and do not use any figures or other materials without explicit written permission. Plagiarism will be grounds for failing the course.

Students should familiarize themselves with the Academic Integrity Policy, available online (http://nbacademicintegrity.rutgers.edu) Quoting from these guidelines:

The principles of academic integrity require that a student:

- make sure that all work submitted in a course, academic research, or other activity is the student's own and created without the aid of impermissible technologies, materials, or collaborations.
- properly acknowledge and cite all use of the ideas, results, images, or words of others.
- properly acknowledge all contributors to a given piece of work.
- obtain all data or results by ethical means and report them accurately without suppressing any results inconsistent with the student's interpretation or conclusions.
- treat all other students ethically, respecting their integrity and right to pursue their educational goals without interference. This principle requires that a student neither facilitate academic dishonesty by others nor obstruct their academic progress.
- uphold the ethical standards and professional code of conduct in the field for which the student is preparing.

Adherence to these principles is necessary to ensure that:

- proper credit for ideas, words, images, results, and other scholarly work, no matter the form or media, is attributed to the appropriate individual(s).
- all student research and work are fairly evaluated, and no student has an inappropriate advantage over others.
- the academic and ethical development of all students is fostered.
- the reputation of the University for integrity, ethics, scholarship, and professionalism is maintained and enhanced.

Any violations to this policy will be reported to Office of Student Conduct (New Brunswick). Violations of academic integrity will be treated in accordance with university policy, and sanctions for violations may range from no credit for the assignment, to a failing course grade to (for the most severe violations) dismissal from the university.

References

[1] P. Moulin and V. V. Veeravalli. *Statistical Inference for Engineers and Data Scientists*. Cambridge, UK: Cambridge University Press, 2019.