

Course title: Fundamental limits of statistical machine learning algorithms

Instructor: Prof. Raj Rao Nadakuditi

Credits: 3

Class Times: Monday, Wednesday 10:30 AM - 12 PM

Course description: This course explores the theoretical and practical limitations of machine learning algorithms, such as computational complexity, data quality and quantity. The course covers both classical and modern results in algorithmic learning theory, as well as recent advances and challenges in deep learning. The course also discusses the implications of these limitations for the design and deployment of machine learning systems in various domains.

Course objectives: By the end of this course, students will be able to:

- Devise computational experiments that will allow them to bring into sharp(er) focus the trade-offs of different machine learning algorithms, such as neural networks, support vector machines, decision trees, and boosting with respect to underlying parameter(s) of interest such as data quality, quantity or noisiness.
- Map the computational experiments into a theoretical program that can facilitate the analysis of this method
- Understand and predict what theoretical tool is needed to be able to link the theory to the observed experimental or empirical phenomenon

Course prerequisites: Students should have a solid background in mathematics (linear algebra, calculus, probability, and statistics) and computer science (algorithms, data structures, and programming). Familiarity with basic machine learning concepts and methods is also expected.

Course format: The course consists of lectures, readings and multiple projects to emphasize master of the learning objectives.

The course will have an experiential format with the students and the instructor co-creating the syllabus based on collaborative discussions.