E6772: Machine Learning for Computer and Communications Networks

Anwar Walid

Course Description: Machine learning has emerged as an important tool for the design, performance optimization and management of computer and communications networks. This is driven by (i) the demand for new higher performance and more reliable network services, (ii) increasing network and systems complexities, (iii) and, importantly, the availability of massive measurement data and higher computing resources. The data-driven machine learning solutions can be applied to augment or replace traditional approaches, that are based on fixed models or human-crafted heuristics, in various functions and layers of the network to learn new and deeper insights, predict and enhance performance, and enable adaptive service automation and optimization. Computer and communications networks also bring new challenges and nuances to the application of machine learning, such as the distributed nature of the data (data collected in multiple places and may be infeasible/costly to move it to one location) and need for explainability. This course covers machine learning methods and introduces their applications in fundamental problems in networking ranging from packet-level to connection-level and application-level in wired, wireless and datacenter networks. Students gain practical experience and preparation for research in this field, and gain understanding of concepts and tools of broader applications. For each application domain, we investigate how machine learning is applied to address limitations of traditional approaches and provide new advantages in emerging networking paradigms. We will also explore open problems and future directions.

Prerequisites: basic knowledge in networking or computer systems, and basic understanding of statistical learning, machine learning or equivalent.

Topics:

- Machine Learning Methods:
  - Supervised learning
  - Unsupervised learning
  - Reinforcement learning
  - Federated learning
  - Transfer and meta learning
  - Generative models
  - Model interpretability, explainability, and robustness
- Machine Learning applications in:
  - Traffic classification and prediction
  - Packet scheduling
  - QoS routing
  - Congestion control for the Internet, wireless and datacenter networks
  - Network topology design
  - Network security, predictive anomaly and intrusion detection.
  - Network and resource management including emerging 5G network slicing and automation.
  - WiFi channel access
  - Self-organizing networks (SONs)
  - Network caching and probabilistic data structures
  - Job scheduling in distributed clusters
  - Content streaming applications

Course requirements: Term Project, Paper Presentation, Homework.

Textbook: No textbook required. Course material will be provided, including references and online resources.