Fall 2020 – CS Special Topics Courses

1. **CS 494* – Hummel (BioE/CS) – Wearable Technologies**
   - **This class/section is for CS undergrads only**
2. **CS 494 – Miao – Introduction to Cryptography**
3. **CS 594 – Asudeh – Responsible Data Science and Algorithmic Fairness**
4. **CS 594** – Caragea – Deep Learning for Natural Language Processing**
5. **CS 594 – Ravi – Learning and Optimization in Vision**

*CS Undergraduate students must submit a modification of major to use the class as a technical elective.

**CS 594 taught by Cornelia Caragea in Fall 2020 will be converted to CS 533. So, the class will count as a regular CS 5xx coursework and not as a special topics CS 594 for graduation requirements.

Back to List

Fall 2020 – CS Special Topics Courses

CS 494– Wearable Technologies

- Instructor: Joe Hummel (CS) and Hananeh Esmailbeigi (BioE)
- Meeting time: T 2-4:50pm
- CRN: 43965

Course Description:
This course, taught primarily by BioE, is a lab-based course in the design and construction of wearable devices. Students work in teams of 4, 2 from CS and 2 from BioE, and use HW and SW skills to design, construct, and analyze wearable devices and the collected data. The class is run like a true lab, where class meets once/week for 3 hours, and class time is devoted mostly to project work. If necessary, projects are completed outside of class time.

Coursework:
3 assigned projects + a final project of team's own design. Each project involves HW and SW integration, some circuit design, data analysis, and a team report. Team work required.

Prerequisites:
CS 251 and 261, and CS 362 if possible.

Back to List

Fall 2020 – CS Special Topics Courses

CS 494– Introduction to Cryptography

- Instructor: Peihan Miao
- Meeting time: T 2-3:15pm
- CRN: 45500/45501

Course Description:

This course is an introduction to modern cryptography from both foundational and practical perspectives.

We will present definitions and constructions of various foundational cryptographic primitives and algorithms, including encryption schemes, pseudorandom generators, hash functions, message authentication codes, block ciphers, digital signature schemes, and so on. We will see what security guarantees are desirable, how to properly define these security guarantees, and how to design cryptographic objects that satisfy them.

Besides the theoretical foundations, we will also study a variety of practical cryptosystems, such as DES, RC4, SHA, El Gamal, Diffie-Hellman, and so on. We will see precisely how these practical cryptosystems are built, and what security level they achieve in various scenarios.

Time permitting, more advanced topics will also be covered, such as zero-knowledge proofs, secure computation, blockchain and cryptocurrency.

Assignments and Grading:

There will be weekly homework/programming assignments, two (2) in-class midterms and one (1) final exam.

Prerequisites:

(a) CS 251, and (b) STAT 381 or IE 342 or STAT 401. CS 401 is highly recommended.

Specifically, we will assume familiarity with basic (discrete) probability and modular arithmetic. Students enrolled are also expected to have had some exposure to algorithms, mainly to be comfortable reading pseudocode and to be familiar with big-O notation.

Fall 2020 – CS Special Topics Courses

CS 594– Responsible Data Science and Algorithmic Fairness

- Instructor: Abol Asudeh
- Meeting time: TR 11-12:15 pm
- CRN: 43136

Course Description:
This course views data-driven and algorithmic decision making through the lens of data ethics and societal impacts. It shall cover the important aspects of the timely research area of responsible data science. The course will empower the graduate students with tools to start exploring/conducting research in this area.

Main Topics:
The course aims at exploring different aspects of responsible data science and societal impacts of data-driven and algorithmic decision making. While a major focus of the course will be on Algorithmic Fairness, other aspects including transparency, accountability, equity, and stability will also be considered. The course will cover the breadth from different aspects:

- Introduction and background: Course outline, aspects of responsibility in data science through recent examples.
- Data Ethics Terms: Fairness, Transparency, Accountability, Stability, Equity, and Diversity, etc.
- Fairness Definitions: Individual v.s. Group fairness, Fairness based on Model Independence v.s. Separation, Intersectional Fairness, Diversity as Fairness
- Impossibility results
- Ecosystem of Data, the pipeline of big data
- Bias in Social Data
- Data Preparation for achieving Fairness by Preprocess techniques
- Post-process techniques for achieving Fairness
- Fairness in Classification
- Fairness beyond Classification: Assignment, Human-designed models, Ranking, and Recommendation, Clustering, etc.
- Fairness and causality
- Stable Decision Making
- Data Investigation: Coverage, Data Profiling, Provenance, and Nutritional Labels

Evaluation/Grading:

- Active Class Participation: 30%
- Presentation: 35%
- Final Project: 35%

Back to List

Fall 2020 – CS Special Topics Courses

CS 594– Deep Learning for Natural Language Processing

- Instructor: Cornelia Caragea
- Meeting time: M 3-5:50pm
- CRN: 43915

Course Objectives:
Natural language processing (NLP) is one of the most important technologies today due to the large and growing amount of online text that needs to be understood in order to get the enormous value out of it. Although many machine learning models have been developed for NLP applications, recently, deep learning approaches have achieved remarkable results across many NLP tasks. The course provides an introduction to research in deep learning applied to NLP. We will cover topics such as word vector representations, convolutional neural networks, recurrent neural networks, attention mechanisms, Transformers, and capsule networks. We will also cover tools and software available for building and training deep neural networks. Through lectures and programming and reading assignments students will learn the necessary skills for applying and designing neural networks for practical NLP problems.

Course Work and Evaluation:
Students will be evaluated based on reading and programming assignments, paper presentations, and a class project. Students are encouraged to attend every lecture and to participate in class discussion.

Prerequisites:
Linear algebra and calculus, machine learning, natural language processing.

Fall 2020 – CS Special Topics Courses

CS 594– Learning and Optimization in Vision

- Instructor: Sathya Ravi
- Meeting time: TR 2:00-3:15pm
- CRN: 27441

Course Objectives:

This class explores the foundations and applications of optimization and learning algorithms that are used to train state of the art vision models. We will focus primarily on models that serve as building blocks in large scale vision applications. Being able to learn meaningful representations of data has been a key factor towards achieving high performance and reliability in vision tasks. Owing to the inherent ill-posedness of the vision tasks, a number of recent studies has shown that understanding the implicit regularization/bias of the models produced by the *optimization/learning* procedure may be vital to enable turn key applications. The course will be divided into two parts. In the first part of the course, we will go over the foundations of optimization algorithms that are commonly used in *discriminative* models that are used in tasks such as object detection, semantic segmentation, and optical flow estimation. In the second half of the course, we will then explore techniques that are used in *generative* models that are used in tasks such as image inpainting, adversarial networks, and adversarial training.

Course Work and Evaluation:

Students will be evaluated based on homework, paper presentations, and a class project. Students are strongly encouraged to attend every lecture, and to participate active in class proceedings.

Prerequisites:

1. CS 401 (Computer Algorithms I) or equivalent; and
2. CS 415 (Computer Vision) **or** CS 412 (Machine Learning) or equivalent.

Familiarity with the optimization material covered in CS 512 (Advanced Machine Learning) will be a plus, but we will introduce them as and when required. Basic background in Linear algebra, calculus, machine learning, general mathematical maturity (for example, intuition for a random variable) will be assumed. Note that CS 515 (Advanced Computer Vision) is **NOT** a prerequisite

Fall 2020 – CS Special Topics Courses

for this class.

Targeted audience

Graduate students from Computer Science, Electrical Engineering, Statistics, Mathematics, and related areas.

Attendance

Students are required to participate in the discussions during lectures. Hence, attendance is essential, and thus is expected.

Course Outline/Readings

We will cover topics at the intersection of Optimization, Computer Vision, and Machine Learning. For each topic, we will cover a representative vision application that benefits the approach directly. The tentative schedule is given below:

Week	Start Date	Topics
1	08/25/2020	Introduction to Continuous Optimization
2	09/01/2020	Representational power; Vision Architectures: CNNs, AlexNet, ResNet
3	09/08/2020	Gradient methods for Unconstrained and Constrained Optimization
4	09/15/2020	Stochastic Optimization and Duality
5	09/22/2020	Backpropagation and Method of Adjoints
6	09/29/2020	Mode Connectivity, and Lottery ticket hypothesis
7	10/06/2020	Generalization ability of models, Algorithmic stability
8	10/13/2020	Double descent curves, shallow vs flat minima
9	10/20/2020	A brief tour of unsupervised learning, and Curse of Dimensionality
10	10/27/2020	Introduction to generative modeling
11	11/03/2020	Convolutional Sparse Coding Model
12	11/10/2020	Wasserstein Distance and Generative Adversarial Networks
13	11/17/2020	A game theoretic perspective on GAN Dynamics
14	11/24/2020	Efficient Adversarially Robust Training
15	12/01/2020	Project Presentations
16	12/08/2020	Final week (project report due)

Note: In this class, more stress is given to the efficiency of the optimization/learning algorithms used. For example, we will discuss algorithms to learn Generative Adversarial Networks using polynomial computational/statistical resources and connections to image inpainting, and background subtraction.