

# FALL 2022 – CS SPECIAL TOPICS

1. CS 294\* – Hayes – Web App Development
2. CS 494\*\* – Soni – Natural User Interactions
3. CS 594 – Cheng – Socially Responsible AI: Theories and Practice
4. CS 594 – Rooshenas – Deep Generative Models
5. CS 594 – Ziebart – Imitation Learning

\*Course counts as a free elective for CS Undergraduate students

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

## CS 294 – Web App Development

- Instructor: David Hayes
- Meeting time: MWF 5-5:50pm
- CRN: 47562

# Build Apps

Learn the tools to build installable mobile (and desktop) PWA apps. The Progressive Web App architecture is used in many popular apps: Spotify, Pinterest, Starbucks, Twitter, Forbes, Alibaba, Lancome and more. You'll be able to build and share apps that work on Android and Apple phones and on Windows.



# Build Your Portfolio

At the end of the semester, you'll have several apps in your GitHub portfolio; these will be useful when you are interviewing. Your apps will include a variety of features: data manipulation, geolocation, camera input, remote data, Google Maps, visualization, Google Machine Learning APIs, text-to-speech and more.

# Build Your Knowledge and Skills

You'll learn a little about a lot of things: we need to learn HTML, CSS and frameworks to build user interfaces (UI); JavaScript to manipulate the UI and interact with web servers and device features; visualization APIs like Google Maps and Google Charts; animation and gaming in the browser; working with data APIs; Chrome Developer Tools to build and test our work; GitHub for version control and to host our apps; and more.

# Who should take the course?

You! If you want to understand how the web works, enjoy learning a little about a lot of things, have the desire to build things, and have some programming experience beyond an introductory course. If you've taken CS211 and CS251, you're ready for the course. Students with less experience but a lot of motivation can succeed in the course but should talk with Professor Hayes before enrolling.

Scan the QR code or navigate to <https://bit.ly/CS294-PWA> for more info.



# FALL 2022 – CS SPECIAL TOPICS

## CS 494 – Natural User Interactions

- Instructor: Nikita Soni
- Meeting time: TR 11-12:15pm
- CRN: 45500/45501

### Course Description

Rapid advances in computer technology have increased the opportunities for developing advanced user interfaces that allow users to interact in *natural ways* (i.e., in the same way as people might interact with other people and physical objects around them). However, what does it mean for an interface to be natural? Also, how do you design natural user interfaces? This class will introduce students to the design, development, and evaluation of a range of Natural User Interaction (NUI) technologies. Unlike keyboard and mouse-based interaction technologies, NUI allows users to interact with technology through a range of human abilities such as touch, gestures, voice, or whole-body movements. Some examples of NUI are touchscreen interactions with smartphones, mid-air gestures with augmented or virtual reality (AR/VR) environments, whole-body interactions with Microsoft Kinect's console, and voice interactions with voice-based interfaces such as Apple's Siri. Beyond traditional keyboard and mouse interfaces, NUI interfaces have added a new level of complexity to the interface design process. There is an increasing demand for researchers and programmers with skills to design advanced user input technologies. The projects students will do in the class will strengthen their portfolios to showcase skills in developing advanced user interfaces that leverage users' natural interaction capabilities.

### Coursework (subject to change)

In this course, students will engage with the learning material through reading relevant literature, participating in group discussions and class presentations, working on hands-on individual and group projects, as well as attending guest lectures. In-class lectures will include discussions of hardware-to-software NUI pipeline, NUI design principles, and key considerations when developing NUI software, including existing platforms, toolkits, and APIs. Students will apply their knowledge in a project of their choice where they will design, implement, and evaluate a NUI prototype.

### Prerequisite

CS 342 Software Design. Experience with one or more of the listed programming languages is recommended for this course: Java, Python, C#, or Objective C. Students who have taken CS422 (User Interface Design and Programming) or CS522 (Human-Computer Interaction) will feel more comfortable with this course, but it is not a pre-requisite.

### Course Outline (Tentative)\*

Week	Topics / In-Class Activities
Week 1	Syllabus, HCI, NUI introduction Designing NUI: Human Factors and Interaction Design Methods

## FALL 2022 – CS SPECIAL TOPICS

Week 2	Developing NUI: Recognition-Based Interaction And Current NUI Frameworks And Algorithms In-Class Activity: Project Discussion
Week 3	In-Class Demo For Current NUI Frameworks Touch/Gesture, Voice, And Whole-Body Interaction Interfaces. In-Class Presentations: Group Project Outlines
Week 4	Evaluating NUI: Live User Studies and Offline Testing In-Class Activity: Project Updates + Group Working / Meeting Time
Week 5	Touch and Gesture Interaction I: NUI Principles Of Designing Touch Interactions In-Class Activity: Project Updates + Group Working / Meeting Time
Week 6	Touch and Gesture Interaction II: Toolkits and APIs for Designing Touchscreen Software Applications In-Class Activity: Project Updates + Group Working / Meeting Time
Week 7	Working with Users and the IRB In-Class Activity: Project Updates + Group Working / Meeting Time
Week 8	Motion and Whole-Body Interaction I: NUI Principles Of Designing Whole-Body Interactions In-Class Activity: Project Prototype Video Demos (8-10 min per group)
Week 9	Motion and Whole-Body Interaction II: Toolkits and APIs for Designing Whole-Body Interaction Software Applications In-Class Activity: Project Updates + Group Working / Meeting Time
Week 10	Speech and Voice Interaction I: NUI Principles Of Designing Whole-Body Interactions In-Class Activity: Project Updates + Group Working / Meeting Time
Week 11	Speech and Voice Interaction II: Toolkits and APIs for Designing Voice-Based Software Applications In-Class Activity: User Study Report Presentations (8-10 min per group)
Week 12	NUI for Feedback: Haptics, Auditory Interfaces, and more In-Class Activity: Project Updates + Group Working / Meeting Time
Week 13	Guest lecture on NUIs for AR/VR Research Topic and Opportunities: NUIs and Children / Skeptical Views of NUIs
Week 14	Guest lecture In-Class Activity: Project Presentations & Final Demos
Week 15	In-Class Activity: Project Presentations & Final Demos

# FALL 2022 – CS SPECIAL TOPICS

## CS 594 – Socially Responsible AI: Theories and Practice

- Instructor: Lu Cheng
- Meeting time: TR 11:00 – 12:15
- CRN: 27441

**Course Description:** In the current era, people and society have grown increasingly reliant on artificial intelligence (AI) technologies. AI has the potential to drive us towards a future in which all of humanity flourishes. Yet, if not used *responsibly* – in accordance with ethical and moral norms, and legal and policy considerations -- it also comes with substantial risks for oppression and calamity.

This course centers around the social responsibility of AI: what we *should* and *should not* do. It will provide students with a systematic framework and a unified perspective of Socially Responsible AI (SRAI) aimed to bridge from the conceptual AI principles to responsible AI practice. We start with real-world examples as motivation, followed by an inclusive definition of SRAI. We then explore the broader impact of AI on modern society and the principles of fairness, accountability, and transparency as we gain a deeper understanding of the importance of a shared set of social values. The majority of content in this course will focus on how to develop algorithmic solutions that align with these principles (some representative topics include fair AI, explainable AI, robust AI etc.) and how to use AI for social good (some representative topics include misinformation and hate speech). This course will serve as a convenient entry point for students to think about the effects of AI algorithms have on individuals, population groups, and on society at large, to understand the problems and challenges of SRAI, and to identify how their developed skills can contribute to making AI more socially responsible.

### Course Work

1. Reading, discussion, presentation of research papers
2. 2 assignments to implement models and try out existing tools
3. Course project

**Note 1:** Each student will be asked to present from 2 papers, and to be the discussant for 2 other papers – this means writing a short-written critique for the paper in question and be ready to participate in discussion. Exact workload will depend on the class size.

**Note 2:** Attendance is required. The instructor will start taking attendance after the add/drop period.

### Prerequisite

Students are required to have taken and received an 'A' or 'B' in one of the following courses: 412 Introduction to Machine Learning, 421 Natural Language Processing, 521 statistical NLP, 533 Deep learning for NLP, 583 Data Mining and Text Mining, 559 Neural Networks.

**Course Outline (Tentative)**

Week	Topic
	<b>Weeks 1-9: Theories</b>
Week 1	Background on AI Responsibility
Weeks 2-4	Fairness in Machine Learning
Weeks 4-6	Explainable/Interpretable AI
Weeks 6-7	Privacy
Weeks 7-9	Distribution Shift and Robustness
	<b>Weeks 10-14: Responsible Practice</b>
Week 10	Misinformation and Echo Chamber
Week 11	Hate Speech, Cyberbullying
Week 12-13	Bias in NLP
Week 14	Privacy-Preserving Applications
Week 15	Project Presentations

## CS 594 – Deep Generative Models

- Instructor: Pedram Rooshenas
- Meeting time: MW 3:00-4:15pm
- CRN: 29095

**Course Description:** Generative modeling is one of the leading directions in machine learning that addresses the problem of modeling the data distribution. Knowing data distribution has many applications, including detecting out-of-distribution data, creating robust classifiers, and creating realistic artificial samples. Probabilistic graphical models were traditionally used for generative modeling, however, deep neural networks enables the introduction of different categories of deep generative models such as autoregressive models, variational autoencoders, energy-based models, flow-based models, and generative adversarial networks. Deep generative models are an important part of unsupervised learning, and has many application in physics and engineering as well as high-dimensional text and image generation.

**Goal:** The main goal of this course is to familiarize the students with generative modeling and deep generative models, and prepare them to conduct research in this area or apply these techniques in their research projects. After successfully completing this course, the students will be able to understand the probabilistic approach to machine learning, understand the generative models and open research problems in the field, and apply deep generative models in their research projects.

**Student Deliverables:** The deliverables include assignments, project, quizzes and exam.

- Project: The students work on individual projects involving one (or more) deep generative models. The goal of the project is to study and improve or apply these models to their research problems. The outcome is a project report and presentation. The project's proposal must be submitted by Week 11.
- Assignment: The students are provided with four programming assignments.
- Quiz and Exam: The course include three in-class quizzes and an exam to evaluate the understanding of the students.

### Tentative Syllabus:

---

Week 1	Introduction
Week 2	Probabilistic Graphical Models
Week 3	Inference and Sampling
Week 4	Parameter Estimation and Evaluation
Week 5	Deep Learning
Week 6	Autoregressive Models
Week 7	Variational AutoEncoder
Week 8	Normalizing Flow
Week 9	Generative Adversarial Networks
Week 10	Energy Based Models
Week 11	Review and Exam
Week 12	Latent Space EBMs
Week 13	Denoising Diffusion Probabilistic Models
Week 14	Hybrid Models
Week 15	Student presentation
Week 16	Project Report

---

Evaluation: The course grading structure includes the following parts:

Programming Assignments	40%
Exam	20%
Quiz and participation	10%
Presentation	10%
Project report	20%

---

Textbook:

Main textbook: Kevin P. Murphy, Probabilistic Machine Learning: Advanced Topics. MIT Press, 2022. Other reference books:

Tomczak, J. M., Deep Generative Modeling. Springer Nature, 2022.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville, Deep Learning (DL), MIT Press, 2015

Christopher M. Bishop, Pattern Recognition and Machine Learning, Springer

We also use a set of tutorials and papers.

**Pre-requisites:** Students are required to have taken and received an 'A' or 'B'; or obtain the consent of the instructor. Students must be proficient in Probability, Linear Algebra, and Calculus. The students must be familiar with Python and Tensorflow/PyTorch.

## CS 594 – Imitation Learning

- Instructor: Brian Ziebart
- Meeting time: MW 9:30-10:45
- CRN: 43915

**Course Description:** Enabling artificial intelligence (AI) systems to behave like humans is an increasingly important challenge for applications ranging from self-driving vehicles to surgical robotics. Imitation learning seeks to address this challenge by using machine learning techniques to enable the AI system to learn from human demonstrations. Two main approaches have emerged: behavioral cloning directly mimics the policy of the human (i.e., the action taken in different situations) to sequence together behavior; and inverse reinforcement learning (also known as inverse optimal control) seeks to learn a reward or cost function that rationalizes demonstrated behavior and induces optimal behavior that is similar to that of the demonstrator.

**Goal:** The primary goal of this course is to prepare students to conduct research in imitation learning--either in the development of improved imitation learning methods or the application of imitation learning techniques to new domains. Students will gain familiarity with the key technical challenges surrounding imitation learning and state-of-the-art imitation learning methods.

### Student Deliverables/Evaluation:

- **Quizzes / Final Exam.** Students will complete a combination of in-class and out-of-class quizzes and a final exam to assess understanding of lecture materials.
- **Research Paper Readings.** Students will be required to read all papers. Each student will present at least once, and students will be required to participate in paper discussion. To ensure that students read all assigned papers, they will be required to turn in a summary for each paper before class. Students will be graded based on their presentation as well as their participation during paper discussions.
- **Research.** Students will submit research proposals at the 3rd/ 4<sup>th</sup> week, a project status update at the end of the 8th week, and a final write-up towards the end of the course (structured as a research paper). Each team will also provide a 20-minute presentation of their project in class.

**Prerequisites:** Students are required to have taken and received an 'A' or 'B' in CS 411 and CS 412; or taken and received an 'A' in CS 594 (Reinforcement Learning).

**Exams:** A comprehensive final exam assess understanding of lecture materials.

# FALL 2022 – CS SPECIAL TOPICS

## Tentative Schedule of Topics:

Week	Topic
1	Overview
2	Optimal Control: MDPs, POMDPs, Influence Diagrams, LQR, iLQR, MPC
3	Reinforcement Learning: Model-based; Model-free
4	Behavioral Cloning, Neural Networks
5	On-Policy vs. Off-Policy Imitation
6	Preference Elicitation; Active Learning
7	Probabilistic Model-based Inverse Reinforcement Learning
8	Discriminative Model-based Inverse Reinforcement Learning
9	Model-Free Inverse Reinforcement Learning
10	Deep Inverse Reinforcement Learning
11	Adversarial Inverse Reinforcement Learning
12	Shared Autonomy
13	Value Alignment & AI Safety
14	Project Presentations
15	Project Presentations / Recap for Final Exam