EECS 498/598: Applied Machine Learning for Affective Computing

Winter 2021
Emily Mower Provost

Everywhere we look, machine learning is uncovering new ways of sensing and modeling human behavior. But, how does this work? Does this even work? The course will cover current practices in measuring and sensing human behavior via machine learning.

Course Description

This course covers the concepts and techniques that underlie machine learning of human behavior across multiple interaction modalities. Topics include: speech/text/gestural behavior recognition through applications of machine learning, including deep learning. The course will also include discussions of the cybersecurity challenges associated with this domain. Fluency in a standard object-oriented programming language is assumed. Prior experience with speech or other data modeling is neither required nor assumed.

Prerequisites:

Students should have taken EECS 281 and (MATH 214 or MATH 217 or MATH 296 or MATH 417) or graduate standing.

Time and Location:

Monday, Wednesday 3:00-4:30 (Remote)

Learning Objectives

1. Understand the value of affective computing in industry and research.
2. Understanding for the common signals used to measure behavior (speech, text, face/vision).
3. Understand the cybersecurity implications of these technologies.
4. Learn machine learning methods in affective computing.
5. Demonstrate an understanding of the concepts by building systems that sense and interpret human behavior.
6. Demonstrate an understanding of the limitations of the technologies *critically* interpreting the newest advances in human-centered technologies.

Course Evaluation

The evaluation of this course will include homework, a midterm, a final, and a final project.

Lecture-by-Lecture

Section 1: Intro and overview

Objective: Provide an overview of the field

1. Introduction ---------- Slides: Set 1

Section 2: Machine learning and Behavior

Objective: Cover background in machine learning
Objective 2: Explain how behavior is measured and the implication of different measurement practices
Objective 3: Tie measurement into theory

2. Extracting behavior 1: language
3. Machine learning 1: Linear regression
4. Machine learning 1: Logistic regression and research spotlight
5. Probability prep and Gaussian Mixture Models
6. Machine learning 2: Gaussian Mixture Models and research spotlight
7. Affective Science 1: Emotion overview
8. Extracting behavior 2: Audio
9. Machine learning 4: Hidden Markov Models 1
10. Machine learning 5: Hidden Markov Models 2
11. Machine learning 6: Hidden Markov Models 3
12. Machine learning 7: Hidden Markov Models 4
13. Data 1: How do we set up data collections and gather data?
14. How do we assess system performance?
15. Hyperparameter optimization and regularization
16. Machine learning 8: Neural networks
17. Machine learning 9: Neural networks
18. Machine learning 10: Convolutional Neural Networks (CNN)
19. Machine learning 11: CNN and Recurrent Neural Networks (RNN) + Demo
20. Machine learning 12: Bias, privacy, and interpretability
21. Machine learning 13: RNNs and research spotlight + Demo
22. Machine learning 14: Autoencoders + Demo
23. Machine learning 15: Representation Learning and Regularization

Section 3: Real world challenges

Objective: Address challenges (both practical and ethical) in this space and discuss how/if these challenges can be mitigated

24. Sources of variation and methods to counter: audio and text ---------- Slides: Set 14
25. What should we do? Ethics, security, privacy (spotlight: HARPA) ---------- Slides: Set 14
26. Project presentations

Grading

Course Evaluation	+1%
Homework	25%
Participation	5%
Midterm	20%
Final	20%
Final Project	30%

2% draft proposal
4% proposal
4% milestone check-in
10% final presentation
10% final report

The values are subject to adjustments based on the discretion of the instructor.

Project

The course includes a semester-long group project. The project will be open ended. It will require you identify a domain that would benefit from continuous sensing, methods to collect and analyze the data, and visualizations to present the results to a user population. Projects should take advantage the knowledge gained in this class and other courses in EECS.