EECS 498/598: Applied Machine Learning for Affective Computing

Winter 2020
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. 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:

Tuesday, Thursday 1:30-3 (EECS 1005)

Learning Objectives

1. Understand the value of affective computing in industry and research.
2. Develop an understanding for the common signals used to measure behavior (speech, text, face/vision).
3. Learn machine learning methods in affective computing.
4. Demonstrate an understanding of the concepts by building systems that sense and interpret human behavior.
5. 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
2. Affective computing in research and industry

Section 2: Machine learning and Behavior Measurements

Objective 1: Cover background in machine learning
Objective 2: Explain how behavior is measured and the implication of different measurement practices

3. Probability and linear algebra prep
4. Extracting behavior 1: Text
5. Machine learning 1: Linear and logistic regression
6. Machine learning 2: Gaussian Mixture Models
7. Extracting behavior 2: Audio
8. Machine learning 3: Hidden Markov Models 1
9. Machine learning 4: Hidden Markov Models 2
10. Hyperparameter optimization and regularization
11. Extracting behavior 3: Video
12. Deep learning 1: Feed-forward NN
13. Deep learning 2: CNN
14. Deep learning 3: LSTM/GRU
15. Deep learning 4: Adversarial methods
16. Anomaly detection
17. How do we measure behavior? Designing datasets for affective computing.
18. Midterm

Section 3: Real world challenges

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

19. What can we do? Challenges and limitations
20. Sources of variation and methods to counter: audio
21. Sources of variation and methods to counter: text
22. Federated learning
23. Student-teacher models
24. Building models of a user
25. Detecting deviations from normal patterns
26. Visualizations
27. What should we do? Ethics, security, privacy (spotlight: HARPA)
28. 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 slight 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.