16:332:579 Machine Learning for IoT [Fall 2023]

Soon to be: Multimodal Learning for Sensing Systems
(Graduate Course 2024)

Jorge Ortiz
Rutgers University

July 3, 2023

1 Abstract

The "Machine Learning for IoT" course, scheduled for Fall 2023 at Rutgers University, provides a comprehensive exploration of multimodal learning and its applications in IoT and sensor systems. The course, taught by Jorge Ortiz, delves into the fundamentals of sensor data and multimodal learning, advanced feature learning, fusion techniques, and practical applications of multimodal learning. It also covers using neural architectures, sequence transformers, and graph neural networks in multimodal learning. The course integrates theoretical learning with hands-on activities and case studies, culminating in a project presentation. It's designed to equip students with the knowledge and skills to apply machine learning techniques to IoT and sensor data, with a particular focus on multimodal learning.

2 Course Synopsis

The course covers a variety of topics related to multimodal learning and sensor data. Here's a brief synopsis based on the schedule:

- Week 1: Introduction to Multimodal Learning and Sensor Data.
- Week 2: Sensor Data Processing and Time Series Analysis.
- Week 3: Continuation of Time Series Analysis and Feature Learning from Sensor Data.
- Week 4: Advanced Feature Learning and Sequence Alignment and Multimodal Learning.
- Week 5: Midterm Exam 1 and Advanced Multimodal Fusion Techniques.
- Week 6: Advanced Multimodal Fusion Techniques and Feature Learning for Multimodal Data.

- Week 7: Early and Late Fusions in Multimodal Learning and Multimodal Learning in Distributed Sensing Systems.
- Week 8: Complex Fusion Techniques in Multimodal Learning and Practical Application of Multimodal Learning: Case Study 1.
- Week 9: Neural Architectures for Multimodal Learning and hands-on activity.
- Week 10: Sequence Transformers in Multimodal Learning and Midterm Exam 2.
- Week 11: Graph Neural Networks (GNNs) in Multimodal Learning and hands-on activity.
- Week 12: Multimodal Learning in Reinforcement Learning.
- Week 13: Transfer Learning and Domain Adaptation in Multimodal Learning and hands-on activity.
- Week 14: Transformers and Attention Mechanisms in Multimodal Learning and hands-on activity.
- Week 15: Review and Recap, Course Project Presentations and Wrap-up.

The course includes a mix of theoretical sessions, practical case studies, and hands-on activities. The course concludes with project presentations and a wrap-up session in the final week.

Detailed Teaching Schedule (Preliminary)

- **Week 1**
  - **Session 1:** Introduction to Multimodal Learning – Overview, History, Importance
  - **Session 2:** Introduction to Sensor Data – Overview, Types of Sensors, Importance
- **Week 2**
  - **Session 1:** Sensor Data Processing – Data Cleaning, Noise Removal, Missing Data Imputation
  - **Session 2:** Time Series Analysis – Overview, ARIMA
- **Week 3**
  - **Session 1:** Time Series Analysis (Continued) – Fourier and Wavelet Transforms

  - **Session 2:** Feature Learning from Sensor Data – Introduction, PCA, ICA, LDA

- **Week 4**
  - **Session 1:** Advanced Feature Learning – t-SNE, Autoencoders
  - **Session 2:** Sequence Alignment and Multimodal Learning – Dynamic Time Warping, Early, Late, and Hybrid Fusion

- **Week 5**
  - **Session 1:** Midterm Exam 1 (Coverage: Weeks 1-4)
  - **Session 2:** Advanced Multimodal Fusion Techniques – Canonical Correlation Analysis, Multimodal Factor Analysis
  - Hands-on Activity: Applying Advanced Multimodal Fusion Techniques to Sensor Data

- **Week 6**
  - **Session 1:** Advanced Multimodal Fusion Techniques – Additive Fusion, Multiplicative Fusion
  - **Session 2:** Feature Learning for Multimodal Data – Multi-view Learning, Multiset Canonical Correlations

- **Week 7**
  - **Session 1:** Early and Late Fusions in Multimodal Learning – Overview, Techniques, Applications
  - **Session 2:** Multimodal Learning in Distributed Sensing Systems – Sensor Network Architectures, Data Fusion in Sensor Networks

- **Week 8**
  - **Session 1:** Complex Fusion Techniques in Multimodal Learning – Complex-Valued Neural Networks, Quaternion Networks
  - **Session 2:** Practical Application of Multimodal Learning: Case Study 1 – Context, Problem Definition, Data Acquisition, Multimodal Learning Techniques (including various fusion techniques), Results, and Conclusion

- **Week 9**
  - **Session 1:** Neural Architectures for Multimodal Learning – CNNs, RNNs
  - **Session 2:** Hands-on Activity: Applying Neural Architectures to Multimodal Data

- **Week 10**

- **Session 1:** Sequence Transformers in Multimodal Learning – Transformer-XL, XLNet
- **Session 2:** Midterm Exam 2 (Coverage: Weeks 5-10)

- **Week 11**
  - **Session 1:** Graph Neural Networks (GNNs) in Multimodal Learning – Introduction, Graph-based Fusion
  - **Session 2:** Hands-on Activity: Applying GNNs to Multimodal Data

- **Week 12**
  - **Session 1:** Multimodal Learning in Reinforcement Learning – Integration of Multimodal Data and RL
  - **Session 2:** Guest Lecture: Applications of Multimodal Learning in Sensing Systems

- **Week 13**
  - **Session 1:** Transfer Learning and Domain Adaptation in Multimodal Learning
  - **Session 2:** Hands-on Activity: Transfer Learning and Domain Adaptation for Multimodal Data

- **Week 14**
  - **Session 1:** Transformers and Attention Mechanisms in Multimodal Learning
  - **Session 2:** Hands-on Activity: Applying Transformers and Attention Mechanisms to Multimodal Data

- **Week 15**
  - **Session 1:** Review and Recap
  - **Session 2:** Course Project Presentations and Wrap-up