Engineering Faculty Document
No.: EFD 74-25
November 19, 2024

TO: The Engineering Faculty

FROM: The Faculty of the Agricultural and Biological Engineering Department

RE: New graduate course – **ABE 54200 –** FROM CHIPS TO CLOUD: MACHINE LEARNING IN IOT AND COMPUTER SYSTEMS

The Faculty of the Agricultural and Biological Engineering Department has approved the following new graduate course. This action is now submitted to the Engineering Faculty with a recommendation for approval.

FROM (IF ALREADY OFFERED WITH TEMPORARY NUMBER): **ABE 59100**

Spring; 3 total credits; Lecture

Pre-requisites: STAT 30100 OR ABE 20500 OR CHE 32000 OR Graduate standing

Enrollment for last 3 spring offerings was between 10 and 15 students. The first spring it was taught as a 1-credit course and there were 25 students enrolled.

TO: ABE 54200 - FROM CHIPS TO CLOUD: MACHINE LEARNING IN IOT AND COMPUTER SYSTEMS

Spring; 3 total credits; LEC/75/2/16

Pre-requisites: STAT 30100 OR ABE 20500 OR CHE 32000 OR Graduate standing

This course provides a comprehensive introduction to the theory, foundations, and latest advances in applying machine learning to the design and optimization of computer systems, particularly those with computational constraints, such as Internet of Things (IoT) devices. Students will learn to develop and optimize machine learning algorithms for computer vision and IoT applications, while managing model complexity and performance. The course covers current developments, including decentralized learning and federated learning, as well as real-time and streaming algorithms for IoT. Additionally, students will gain foundational knowledge in cloud computing and database technologies, essential for deploying machine learning solutions across various cloud platforms.

<u>Learning Outcomes</u>:

- **Design and Evaluate Machine Learning Algorithms**:
Students will **design** and **evaluate** supervised and unsupervised machine learning algorithms, particularly for computer vision tasks and IoT applications, evaluating their performance using appropriate metrics.

- **Analyze and Optimize Model Performance**:
Students will **analyze** machine learning models using techniques like cross-validation, regularization, and quantization, and **optimize** model performance by balancing accuracy and computational complexity.

- **Apply Decentralized Learning and Streaming Algorithms in IoT**:
Students will **apply** concepts of decentralized learning, including federated learning and real-time data processing, to **design** streaming algorithms on IoT devices.

- **Understand and Evaluate Cloud Computing Fundamentals and Database Technology**:
Students will **explain** cloud computing fundamentals and **differentiate** between structured (SQL) and unstructured (NoSQL) databases, with an emphasis on how these systems are **deployed** and **evaluated** on cloud platforms.

RATIONALE:

This course will serve as a building block for students wishing to explore further applications of Machine Learning in applications. Theory, foundations, and latest advances in applying machine learning to building computer systems, especially systems under constrained computational resources as in Internet of Things (IoT) devices. There have been 4 offerings of the course – the first one was a one-credit offering and the other three were three credit offerings. There were 25 students in the first offering and between 10 and 15 in the others.

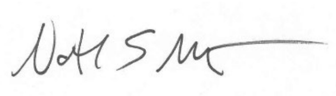

Nathan Mosier, Head of the Agricultural and Biological Engineering Department
Link to Curriculog entry: https://purdue.curriculog.com/proposal:30687/form

Syllabus for the Spring 2024 cohort

1/2 of lectures = data science/data engineering/ML concepts and intuition
1/4 of lectures = computer vision and IoT algorithms
1/8 of lectures = reading technical papers and presenting
1/8 of lectures = guest lectures and discussion

Example lectures: https://schaterji.io/teaching/

Course description: This course will cover data science and data engineering algorithms for use in computer vision tasks and IoT devices. This will allow students to develop an algorithmic understanding of the tools and techniques in this field. Further, the course will also include knowledge of commercial software by cloud computing and machine intelligence software vendors such as Google, Microsoft, and Amazon. The course will also cover cloud computing and serverless infrastructure, data engineering best practices, and machine learning innovations targeted at making large networks function on diverse hardware. Finally, the course will also introduce decentralized learning, computational complexity, and ethics to be able to understand the nuances of deploying algorithms in the real world.

Syllabus – link: https://docs.google.com/document/d/1l3YKx-woN6y79vmrnJ4l89VEMt0xH1uNsIXuZPpZLs/edit

Week 1: Introduction to Machine Learning and IoT

Overview of IoT and Computer Vision Applications
Introduction to ML Algorithms
Supervised and Unsupervised Learning
IoT data-driven ML
Datasets and Train/Validation/Test Sets
Example Datasets (MNIST, CIFAR, etc.)
Overview of Data in IoT and CV contexts

Week 2: Training ML Models and Error Analysis

Training an ML Model
Overfitting, Underfitting, Bias/Variance Trade-off
Cross-validation Methods (k-fold, leave-one-out)
Model Evaluation
Regularization (Ridge, Lasso, ElasticNet)
Hyperparameter Tuning

Week 3: Support Vector Machine and Deep Learning Basics

Introduction to Support Vector Machines (SVM)
Theory and Applications in IoT and CV
Deep Learning Overview
Perceptron, Evolution of Neural Networks
Deep Learning for Computer Vision and IoT

Week 4: Activation Functions and Gradient Descent

Activation Functions
Sigmoid, ReLU, Softmax: Use Cases
Gradient Descent
Batch Gradient Descent vs. Stochastic Gradient Descent
Learning Rates and Convergence

Week 5: Information Theory and Model Complexity

Information Theory in ML
Cross-Entropy, KL-Divergence, Entropy of a Distribution
Model Complexity
AIC (Akaike Information Criterion), BIC (Bayesian Information Criterion)

Week 6: Autoencoders and Dimensionality Reduction

Autoencoders
Introduction to Autoencoders for Dimensionality Reduction
Visualization Techniques
Dimensionality Reduction
PCA, t-SNE, UMAP
Regularization and Sparsity

Week 7: Computer Vision Applications (Part 1)

Examples of CV Models
Object Detection, Segmentation, Classification
Metrics for CV Performance
Accuracy, Precision, Recall, F1 Score, IoU (Intersection over Union)

Week 8: Computer Vision Applications (Part 2)

Convolutional Neural Networks (ConvNets)
Architecture and Use in CV
Examples with IoT (e.g., Smart Cameras)

Residual Networks (ResNets)
Deep Networks and Performance Improvement

Week 9: Machine Learning for Streaming Data

Streaming Algorithms for IoT Data
Processing Real-time Data in ML Systems
Challenges in Streaming Data
Continuous Learning, Data Drift, and Adaptation

Week 10: Diffusion Models

Introduction to Diffusion Models
Theory and Applications in ML
Recent Advances in Generative Models
Use of Diffusion Models in IoT and CV

Week 11: Cloud Computing Fundamentals

Cloud Computing for ML and IoT
Cloud Architecture for Large-scale ML
Database Technology
Structured SQL and NoSQL Databases for IoT

Week 12: Serverless Computing and Cloud Infrastructure

Serverless Computing
ML Workflows in Cloud Environments (e.g., AWS Lambda)
Introduction to AWS Infrastructure
Optional Lecture: Infrastructure for Cloud and IoT

Week 13: ML Ethics and Data Hygiene (Part 1)

ML Ethics: Data Privacy
Overview of Privacy Issues in IoT and CV
GDPR and Data Collection Regulations
Model Fairness and Bias
Addressing Bias in ML Models

Week 14: ML Ethics and Data Hygiene (Part 2)

Data Provenance and Model Versioning
Tracking Changes in Models and Datasets
Data Preprocessing and Postprocessing
Best Practices in Data Engineering for ML

Week 15: Guest Lectures and Review

Guest Lecture on Entrepreneurship (Disruptive Innovation)
Guest Lecture on IoT in the Wild
Course Review and Final Discussion

Assigned readings (on class website) – see examples here - https://schaterji.io/1PAW/awesome/

Top ML and CV conference papers and high-profile journals for vision papers
Podcasts – examples - https://schaterji.io/podcast.html

Evaluation

- **Class participation, pop quizzes, and assessment: 15 points**
- **Homework: 30 points [method of continuous assessment - e.g., writing summaries, presenting papers and concept, multiple choice/etc]**
- **Mid-term: 30 points [open-book, open-notes, you can use the same Google doc you create for HWs].**

The main purpose of the mid-term will be for you to look up and assimilate concepts, recapitulate the concepts taught and discussed in class, map them to what you learned in class, and then be able to write out the answers formally, uploading your document by the deadline. This is so you get good at technical writing, communicating well to an ML/CS-y audience, capturing the tone of the class. I want this course to whet your appetite to think algorithmically, and to know what to apply, where. This is especially constructive if you are an ML beginner, but also helpful if you are an intermediate (someone who has applied ML packages without a clear intuition) wanting to know the advances in ML at a more algorithmic level. The mid-term exam (and assessment in general) will further this goal.

- **End-term: 30 points [open-book, open-notes, open-laptop, needs to be electronically submitted].**

I will not accept a hand-written exam to encourage formal writing.

- **Instructor feedback: 5 points bonus [by honor code]**

This is for the Purdue Qualtrics survey.