CS 4333 – Current Directions in Artificial Intelligence (4-0)

Syllabus:

Week	Topics	Applications	Readings, assignments, and tests
1	*Unit : AI in Commercial Enterprise Applications* • Goals of big data and analytics (2h) • Data collection, lifecycle, governance, and infrastructure (2h) • Data lakes • Map/reduce	• Analytics • Optimization ▪ Logistics ▪ Production ▪ Human resources • Customer targeting	Selected handouts
2	*Unit : Advanced Neural Networks* • Implementation of convolutional neural networks (2h) • Implementation of sequential neural networks (LSTM, RNN) (2h)	• Image/object classification • Natural-language processing	Selected handouts; lab on using a neural network implementation
3	*Unit : Adversarial Machine Learning and AI Security* • Data integrity and threats to data integrity (2h) • “Evasion” attacks and SPAM (2h)	• Transferability of adversarial perturbations	Selected handouts
4	*Unit : Adversarial Machine Learning and AI Security* • Adversarial perturbations and attacks on vision-based systems (4h)	• Physical adversarial attacks	Selected handouts; test on the material so far
5	*Unit : Explainability, Interpretability, and Model Debugging* • Salience maps and other gradient-based approaches (e.g., LIME) (2h) • The function of and requirements for explanations to developers, end consumers (1h) • Human-machine teaming • Epistemic trouble with explanations (30 min) • End of theory vs. knowing things (30 min)	• Image classification and class activation maps (Grad-CAM)	Selected handouts; lab on using explanation capabilities of a neural network
6	*Unit : Testing and Verification of AI Systems* • Performance evaluation (2h) • Accuracy and related	Many	Selected handouts

	statistics(1h) • Confusion matrices and ROC curves(1h)		
7	*Unit : Testing and Verification of AI Systems* • Development lifecycles (2h) • Formal verification. (2h) • Lightweight • Heavyweight	Many	Selected handouts; lab on attesting of a neural network
8	*Unit : Measurement, Bias, Causality, Ethics* • Explicit measurement modeling and other frameworks for validity (1h) • Proxies, selection bias, and risks of harmful bias in data-driven systems (1h) • Confounding and problems of causality (1h) • Ethics and ethics principles (1h) • DoD AI Ethics Principles	• Bias in image classification with neural networks	Selected handouts; test on the material so far
9	*Unit : Privacy and AI* • Mosaic theory problems & proxies (1h) • Privacy-preserving approaches (2h) • Differential privacy and the fundamental need for noise • Legal regimes and attendant compliance issues (1h)	Many	Selected handouts; label on manipulation of data to assess privacy
10	*Unit : AI Legal Issues* • Privacy and data protection (2h) • Data handling issues encountered in practice • Intellectual property (2h) • Data licensing issues encountered in practice	Many	Selected handouts
11	*Summary and final exam*	No new material	Test on the whole course

Catalog Description

A survey of current important topics in artificial intelligence for students who are not computer-science majors. Topics include big-data management, advanced topics in neural networks, adversarial machine learning, explainability, testing and verification of artificial-intelligence systems, privacy issues in artificial intelligence, and other legal issues in artificial intelligence.

Prerequisites

Either CS3331 or CS3310, and either CS3332 or CS3315

Course and Learning Objectives

This course covers the practical use of AI technologies in solving problems within enterprises. Starting from a view of what problems are well solved with data-driven analysis, the course covers modern methods (including deep neural networks) and a host of practical issues that users of these technologies should be aware of. These include security and adversarial examples; privacy, data management, and data governance; testing and verification of system performance and safety; explainability, interpretability, and debugging; data issues such as measurement, preprocessing/cleaning, data bias, and causality; and legal and ethical issues with AI technologies.

Course Plan

Lecture hours will be split 75% synchronous lectures, recorded for asynchronous consumption, and 25% discussion and question-answer sessions to discuss topics and applications covered in lecture or of interest to students. Discussions will enable a deeper exploration of course topics, especially topics like performance management and ethics, where applications often entail tradeoffs.

Grading and Assessment

Problem sets/labs - 40%

Two tests and a comprehensive final exam - 60%

Teaching materials

Papers to be assigned in class, including these or similar papers:

- C. Dao-Duc, H. Xiaohui, and O. Morere, Maritime Vessel Images Classification Using Deep Convolutional Neural Networks, Proc. 6th intl. Symp. on Information and Communication Technology, December 2015, pp. 276-281
- M. Benaddy et al., Recurrent Neural Network for Software Failure Prediction, Proc. 4th Intl. Conf. on Engineering and MIS, June 2018, article 16
- V. Zantedeschi, M.-I. Nicolae, and A. Rawat, Efficient Defenses against Adversarial Attacks, Prof. 10th ACM Workshop on AI and Security, November 2017, pp. 39-49
- T. Ha, S. Lee, and S. Kim, Designing Explainability of an Artificial Intelligence System, Proc. of the Technology, Mind, and Society Conference, April 2018, Article 14
- D. Marijan, A. Gotlieb, and M. Ahuja, Challenges of Testing Machine Learning Based Systems, Prof. IEEE Intl. Conf. on Artificial Intelligence Testing, April 2019
- R. Al-Shabandar, and G. Lightbody, The Application of Artificial Intelligence in Financial Compliance Management, Proc. Intl. Conf. on Artificial Intelligence and Advanced Manufacturing, October 2019, Article 8
- L. Esquivel, E. Barrantes, and F. Darlington, Measuring Data Privacy Preserving and Machine Learning, Prof. 7th Intl. Conf. on Software Process Improvement, October 2018
- O. Erdelyi, and J. Goldsmith, Regulating Artificial Intelligence: Proposal for a Global Solution, Proc. AAAI/ACM Conf. on AI, Ethics, and Society, December 2018, pp. 95-101
- D. Roman, and P. Natalia, Artificial Intelligence Legal Policy: Limits of Use of Some Kinds of AI, Prov. 8th Intl. Conf. on Software and Computer Applications, February 2019, pp. 343-346

Other information:

Scheduling:

Three quarters after CS4000, the first course in the certificate sequence.

Duplication:

CS3333 shares some material with CS3310 and CS3315 but is mostly advanced material that has not been covered before: implementation of neural networks, adversarial methods, and trust in artificial-intelligence systems.

MOVES offers three courses with some overlap (MV3025, MV4025, and MV4100) but those courses are focused on AI for simulations as the titles say. A new ME course (number unknown) will focus on AI for autonomous vehicles, a more specialized topic. OA4321 covers some AI topics but its focus is on a broader class of decision-support systems. OS4118 covers statistical and machine learning at an advanced level, and has four prerequisites, something impossible for students in our certificate program to fulfill.

Resources:

No additional faculty or equipment resources are required.

Learning Objectives:

- Students can identify in applications the key current challenges to applying artificial intelligence to military applications, including obtaining data, getting sufficient processing speed, verifying correctness, anticipating adversary manipulations, testing, providing explanations, encouraging trust in methods, and dealing with legal issues.
- Students can trace cases exemplifying the challenges and use them to explain the challenges.
- Students can rate different artificial-intelligence techniques as to their suitability in addressing the challenges and make intelligent recommendations to military organizations for methods to use.