BME|SIE 578: Artificial Intelligence for Health and Medicine - #AI4HAM

Instructor: Vignesh Subbian, PhD | E-mail: vsubbian@email.arizona.edu

Class Time: Tuesday and Thursday – TBD | Venue: TBD

Office Hours: 60 minutes after class or by appointment

Course Catalog Information: (3 credits) The practice of modern medicine in a highly regulated, complex, sociotechnical enterprise is a testament to the future healthcare system where the balance between human intelligence and artificial expertise will be at stake. The goal of this course is to introduce the underlying concepts, methods, and the potential of intelligent systems in medicine. We will explore foundational methods in artificial intelligence (AI) with greater emphasis on machine learning and knowledge representation and reasoning, and apply them to specific areas in medicine and healthcare including, but not limited to, clinical risk stratification, phenotype and biomarker discovery, time series analysis of physiological data, disease progression modeling, and patient outcome prediction. As a research and project-based course, student(s) will have opportunities to identify and specialize in particular AI methods, clinical/healthcare applications, and relevant tools.

Prerequisites: A basic foundation in linear algebra, discrete mathematics, probability and statistics, and data structures are recommended for this course. Additionally, students should be open to learning new computing tools and technologies. Tutorials or primers will be provided for most prerequisite topics and tools. This course is generally suitable for graduate students in applied math, biomedical engineering, computer engineering, mechanical engineering, statistics, systems and industrial engineering, computer science, information science, and management information systems. If you are a student from outside of these disciplines, please talk to the instructor to discuss your interests and obtain consent to enroll.

Course Reference Textbooks:

There are no required textbooks for this course. Reference and reading materials will be provided via D2L. A few general sources that we will use in the course are listed below.

- Stuart Russell and Peter Norvig. 2009. *Artificial Intelligence: A Modern Approach* (3rd ed.). Prentice Hall Press, Upper Saddle River, NJ, USA.
- Toby Segaran. 2007. *Programming Collective Intelligence* (First ed.). O'Reilly.
- Tony J. Cleophas and Aeilko H. Zwinderman. 2015. *Machine Learning in Medicine - a Complete Overview.* Springer.
- Sunila Gollapudi, S. 2016. *Practical Machine Learning*. Packt Publishing Ltd.
- Peter Harrington. 2012. *Machine Learning in Action*. Manning Publications Co., Greenwich, CT, USA.
- Selected seminal and contemporary readings from peer-reviewed literature such as Proceedings of Machine Learning in Healthcare, Artificial Intelligence in Medicine, IEEE Transactions on Biomedical and Health Informatics, and other relevant venues.

Course Topics:

This is a tentative list of topics. Topics may be revised on a yearly basis based on needs, trends, and updates in the methodological domain, application areas, and/or tools*. Also, it is important to note that topics in Part I (foundations), Part II (applications), and Part III (implementation & evaluation) will be "matched" and discussed in tandem.

Part I: Foundations

1. Introduction to Human and Artificial Intelligence: terminologies, computational models of intelligence; conceptual frameworks from cognitive and educational psychology, neuroscience, information theory, and linguistics; philosophical foundations of AI
2. Review of relevant mathematical and statistical concepts: logarithmic loss, cross entropy optimizing cost functions; linear and logistic regression.
3. Forms of Learning: supervised, semi-supervised, unsupervised, active, and transfer learning
4. Supervised Learning: (a) Decision trees, non-parametric methods for learning, support vector machines, (b) Bio-inspired Learning (from perceptron to deep learning): neural basis of computing, classical neural networks, deep neural networks, deep belief networks, recurrent neural networks, and convolutional neural networks.
5. Unsupervised Learning: basic and advanced clustering techniques, dimensionality reduction (feature selection and feature extraction)
6. Knowledge Representation and Reasoning: Prepositional logic, first-order logic, ontological engineering, probabilistic reasoning
7. Time-series analysis: temporal models (probabilistic reasoning over time)
8. Emerging paradigms and concepts in artificial social and emotional intelligence

Part II: Applications

9. Unique characteristics and challenges in medicine and healthcare; History and status quo of intelligent and expert systems in medicine.
10. Risk stratification, patient outcome prediction, disease progression modeling
11. Clinical decision-making and intelligent systems to support evidence-based medicine
12. Phenotype and clinical/bio-marker discovery, Relevance to personalized medicine
13. Analysis of tissue morphology and other medical imaging applications

Part III: Implementation and Evaluation

14. Tools and Technologies for implementing AI methods**
15. Model evaluation and performance metrics, cross-validation, model interpretability
16. Ethics of AI: bias, fairness, accountability, and transparency in machine learning; Ethical, Legal, and Social Issues of AI in medicine and healthcare

*Topics not listed above such as other methodological concepts such as problem-solving by search, humanistic intelligence, and biomedical applications such protein-protein interactions, optimization of treatment protocols are welcome as a part of research critiques or individual projects.

**The course will encourage the use of Python-based tools and libraries. Students are not required to have prior experience with Python and machine learning libraries, but are encouraged to get

accustomed with such tools during the course of the semester. Prior experience with computing tools such as MATLAB can be easily transferred/applied to this course (for example, see startup guide here).

Course Learning Outcomes:

At the end of this course, each trainee/student will be able to

(1) Understand models of human and artificial intelligence, specifically computational models of intelligence.
(2) Comprehend a collection of machine learning models (identified and covered in the course), and their applications in medicine and healthcare.
(3) Identify and apply appropriate intelligent system models and computational tools to specific problems in biomedicine and healthcare.
(4) Analyze the performance of specific models as applied to biomedical problems, and justify their use and limitations.
(5) Identify, understand, and interpret methods and evidence from artificial intelligence and other relevant literature.
(6) Effectively communicate and disseminate knowledge in any science or engineering domain in the context of computing, systems, and/or biomedical applications.

Assessment of Learning

Research Critiques (23%)

Research critiques are meant to allow for critical thinking around specific methods and their connection and relevance to applications. This will involve in-class demonstration and discussion sessions based on research and written critiques prior to the session.

Project (70%)

This semester-long project will involve identification of method(s) and a biomedical problem of interest, selection of appropriate tools and datasets, systematic review of pertinent literature, application and evaluation of the methodology, interpretation of results, and development of a comprehensive research manuscript. Preliminary work for the project will begin during the third week of the semester. Detailed milestones and guidelines will be available on D2L at the start of the semester.

Quizzes (7%)

There will be short in-class quizzes during the course of the semester. Tentative dates for the quizzes will available on the online (D2L) course calendar during the second week of the semester.

Assessment Scale:

The final letter grade will be based on UA's standard grading scale and assigned as follows. The minimum overall total (lower limit) may be lowered, but will not be raised.

A: 90 to 100 | B: 80 to 89.99 | C: 70 to 79.99

D: 60 to 69.99 | E: 0 to 59.99

Statement of Inclusion

As an interdisciplinary course with a focus on sociotechnical systems and implications to society and human health, the *diversity** of participants in this course is a valuable source of ideas, problem solving, and engineering creativity. If you feel that your contribution is not being valued for any reason, please speak with me privately. If you wish to communicate anonymously, you may do so in writing or speak with Linda Cramer or Trisha Stanley. As members of the UA community, it is our shared responsibility to cultivate a climate where all students/individuals are valued and where both they and their ideas are treated with respect.

**includes every participant's identity, personal and academic/professional background (includes technical/clinical experience, internship/co-op/research experience), interests, and expertise.*

Statement of Personal Challenges and Preferences

- It is the University's goal that learning experiences be as accessible as possible. If you anticipate or experience physical or academic barriers based on disability or pregnancy, please let me know immediately so that we can discuss options. You are also welcome to contact Disability Resources (520-621-3268) to establish reasonable accommodations.
- Also, I will gladly honor any request to address you by an alternate name or gender pronoun. Please advise me of this preference early in the semester so that I may make appropriate changes to my records.

Course Policies

- **Team Work:** For team projects, students are required to work on teams of 2 or 3. Students' background and prior experience will be used to form teams. Every member of the team is expected contribute equally to the project. The entire team will share the same score for each release; however, if the instructor/members identify others who have not contributed (or have contributed above and beyond), the scores will be varied accordingly.
- **Extra Credit:** Extra credit opportunities (if any) will be announced in class.
- **Submission Policies:** All project-related submissions are due at 11.59 pm on the day specified on the write-up (usually Saturdays). Late submissions will receive 85% credit after the due date. Late submissions beyond one week after the due date will receive 10% credit.
- **Make-up Policy:** Make-up quizzes may be given in case of illness or personal emergencies. In case of illness, a written and signed note from the physician is required.
- **Audit Policy:** Students auditing this course should work with one of the project teams, provide research and/or computing support, and submit a report summarizing their experience and preliminary results from the project.
- **Academic Integrity:** Every student/student team should strive for honest and responsible conduct. All forms of academic misconduct including, but not limited to, plagiarism, unauthorized collaboration, and fabrication are strictly prohibited. Failure to adhere to these guidelines will be dealt with as recommended in the UA Student Code of Conduct and the submission/activity under consideration will receive 1% credit.

- **Grievance Policy**: All course-related grievances should be communicated in writing. See *Statement of Inclusion* for related information. The graduate college grievance policy is here.
- **Statement of Changes to Syllabus**: Information contained in this course syllabus, other than the grade and absence policy, may be subject to change with advance notice, as deemed appropriate by the instructor.
- **General Policies**: All policies and codes approved by the dean of students apply to this course.
    - *Absence Policy:* Absences for any sincerely held religious belief, observance, or practice will be accommodated where reasonable. Absences pre-approved by the UA Dean of Students (or dean's designee) will be honored.
    - *Classroom Behavior Policy:* See Statement of Inclusion
    - *Threatening Behavior Policy:* The UA Threatening Behavior by Students Policy prohibits threats of physical harm to any member of the University community, including to oneself.
    - *Nondiscrimination and Anti-harassment policy:* UA is committed to creating and maintaining an environment free of discrimination. See statement of inclusion and further details here.