

# **OIDD 255X: AI, Business, and Society (1 CU)**

**Instructor: Prasanna (Sonny) Tambe**

*Last updated 10/22/2021*

## **SAMPLE SYLLABUS, ALL ELEMENTS ARE SUBJECT TO CHANGE.**

### **Course Objectives**

The progression of AI-based technologies promises to transform many aspects of business, labor, and even society. The goal of this course is to provide students with an understanding of the capabilities of modern AI technologies, with an emphasis on being able to critically assess where they can provide societal value, and where they may create new societal challenges. The course is not intended to provide a deep-dive into the workings of these technologies in the same way as a computer science course might. Rather, business and policy decision-makers will be confronted with a number of important issues as AI becomes integrated into the social decision-making fabric. This course is intended to provide a framework for people who may have to confront these legal, ethical, and economic challenges. In doing so, a course objective is to ensure that students who complete the course are comfortable enough in the inner-workings of these technologies to think critically across many AI contexts as well as different domains ranging from public policy, to criminal justice, to health inspections, HR, and marketing.

This 1 CU course is oriented around hands-on critical written assessments, labs, exams, and a presentation. Broadly, data rich firms in finance, tech, management, marketing, and other industries are increasingly adopting AI as a tool to accelerate and improve decision-making. It is important for modern managers to understand the opportunities and challenges introduced by data and AI so that they can credibly communicate about these issues with others in the firm. We will cover many of these issues, so that you will be able to think about the opportunities and challenges that arise when firms try to use AI to solve business problems.

Labs will reinforce your learning of how AI works, and how it is being used to solve business problems. We will focus in the labs on gaining experience with introductory machine learning concepts. Students will spend time inside and outside of the classroom combining data and code to provide a foundation for understanding the deep challenges that this will bring to organizations.

### **Course Overview**

Within the last decade, there has been a dramatic rise in interest in the use of AI technologies applied to many domains, including finance, HR, policy, marketing, and strategy. As a result, the modern "digital leader" increasingly requires the use of technology, statistics, and data analysis skills to facilitate business analysis as well as a grasp of where AI technologies can perform well, and where they may fail. This includes knowing how to a) effectively frame data-driven questions, b) use AI-driven algorithms, and c) understand the capabilities of this new generation of tools that are becoming available to automate decision-making from data.

The class includes readings and critical assessments due before class and in-class discussions of topics related to the application of AI technologies, including ethics, bias, and the potential for AI to fuel gains in productivity. The class also requires the completion of two labs to ensure that

participants have a deep understanding of how algorithms are applied for decision-making and what the constraints are of these approach. A background in coding is not required in this class. Labs will be based in “no-code” tools such as MS Azure ML Studio.

Throughout the semester, we will cover the applications of AI to a number of industries (e.g. medicine, journalism, criminal justice), including inviting guest speakers to hear about applications in specialized domains. A learning goal of this course is exposure to how AI is changing decision-making in different industry contexts, and how organizations are reacting to these changes.

Class time in most weeks will be dedicated to lecture and discussion of some of the most important topics facing the AI community. Class time in some weeks will be devoted to supervised work on projects that are meant to underscore AI-based challenges. Through these exercises and discussions, students are expected to become proficient at applying data to business decisions and at effectively analyzing big data sets to inform decisions about business problems using data analysis tools.

## Course web site

We will be using Canvas to submit all assignments and receive grades. All course information will be posted on the Canvas website.

## Required textbooks and software

There is no required textbook. There are frequent readings that will consist of selected online content which will be posted on the course site. As part of your homework, you may be expected to download and install some open source software.

## Deliverables and grading

During this course, you will be assigned a number of hands on data projects which you will spend time on both in class and out of class. You are expected to participate in classroom discussions (there is more information below about participation). The breakdown of points is as follows:

Deliverable	Weight	Points
AI presentation	5%	25
Labs	5% each (x 3)	75
Response writeup	5% each (x 2)	50
Midterm	15%	75
Participation	15%	75
Final Exam	25%	125
AI pitch deck	15%	75
<b>TOTAL</b>	<b>100%</b>	<b>500</b>

With each project, you will be provided with a set of guidelines. Deliverables may include short, informal analyses and an accompanying recommendation. Group projects will be completed in small groups (two to three students, no more than three). You *may* also be asked to evaluate the contribution of each of your team members after the group project.

The classroom presentation and discussion presents a unique opportunity for you to develop and enhance your confidence and skills in articulating a personal position, sharing your knowledge, and reacting to new ideas. All of you have personal experience that can enhance our understanding of this subject, and we want to encourage you to share that experience.

## **Participation and Professionalism**

This course, like many other courses at Wharton, uses learning methods that require active involvement (e.g. attendance, participation in discussions, and in-class exercises). Not only is this the best way to learn, but it also develops your communication and presentation skills. Regular attendance, participation, presentations, and in general, presenting yourself professionally are all very important, and are an important part of your grade.

Active participation requires good preparation—thoughtful completion of homework before class is essential. We recognize that expressing viewpoints in a group is difficult, but it is an important skill for you to develop. We will do what we can to make this as easy as possible. Remember though that only regular and insightful contributions will be rewarded.

The grade we assign for your class participation and attendance is a careful, subjective assessment of the value of your input to classroom learning. We keep careful track of attendance, your contributions towards each class session, and these contributions can include (but are not restricted to) raising questions that make your classmates think, providing imaginative yet relevant analysis of a situation, contributing background or a perspective on a classroom topic that enhances its discussion, providing thoughtful feedback on the presentations of other students, and simply answering questions raised in class. A lack of preparation, negative classroom comments, or improper behavior (such as talking to each other, sleeping in the classroom or walking in and out of the class while the lecture is in process) can lower this grade.

In particular, because this class emphasizes projects and learning-by-doing, attendance is expected. Missing an excessive amount of class without justification will negatively impact your final grade.

Participation grades will include attendance as well as contributions to the class discussion. This is not a “binary” score. Please note that the resulting score is not a simple high/low measure. The distribution of participation scores in past semesters has routinely been bell-shaped, with particularly high scores being awarded only to students that have distinguished themselves in terms of participation.

## **Grading Guidelines**

At Wharton, we strive to create courses that challenge students intellectually and that meet the Wharton standards of academic excellence. If you believe that an assignment or project grade you received was unjustified, you can appeal the grade. To appeal the grade you must write a one-page explanation as to the reason for your appeal and hand it along with your graded assignment back to the TA responsible for that assignment.

Please think twice before appealing a grade: the TA will completely re-grade the assignment, which may increase your grade, but may also lower it (e.g., if the TA catches more mistakes the second time around). If after re-grading you feel that your grade was again unjustified, you can appeal the grade with the instructor.

Points will be deducted from late assignments, labs, or projects at the rate of a 20% penalty for each day the submission is late.

## Schedule

Session	Date	Topic	Due
1	xxx	<b>Course Intro</b>	
2	xxx	<b>AI: The past to the present</b>	
3	xxx	<b>AI Applications: Presentations</b>	Presentations
4	xxx	<b>AI Applications: Presentations</b>	
5	xxx	<b>Understanding Digital Data</b>	
6	xxx	<b>How AI works: ML</b>	
7	xxx	<b>How AI works: Loss functions</b>	
8	xxx	<b>In class ML Lab</b>	
9	xxx	<b>How AI works: Deep learning</b>	ML Lab
10	xxx	<b>In class Deep Learning Lab</b>	
11	xxx	<b>Quiz</b>	Quiz
12	xxx	<b>Ghost Work: How Data Gets Labeled</b>	
13	xxx	<b>Reinforcement Learning</b>	
14	xxx	<b>Generative AI</b>	
15	xxx	<b>AI Competition: Software, Skills, Data, and Computation</b>	DL Lab
16	xxx	<b>AI Competition: Networks and Regulation</b>	
17	xxx	<b>AI Barriers: Bias</b>	Auto ML Lab
18	xxx	<b>AI Barriers: Explainability</b>	
19	xxx	<b>Industry analysis: HR + Guest</b>	
20	xxx	<b>AI Barriers: Ethics &amp; Law</b>	Bias response
21	xxx	<b>AI Barriers: Privacy</b>	
22	xxx	<b>Industry analysis: Healthcare + Guest</b>	
23	xxx	<b>AI Culture and Governance</b>	Ethics response
24	xxx	<b>AI &amp; Jobs I</b>	
25	xxx	<b>AI &amp; Jobs II: Scenario exercise</b>	
26	xxx	<b>Exam</b>	Final exam
27	xxx	<b>Final project presentations 1</b>	
28	xxx	<b>Final project presentations 2</b>	Final project

## Individual Session Descriptions

### Session 1: Course Introduction

*Session Overview:* Introduction to course essentials, including topics and concepts covered, session structure, and grading. Group-led discussion of critical issues in AI.

### Session 2: AI: The past to the present

*Session Overview:* Surveys a “history” of artificial intelligence, ranging from the 1800’s to modern day. Discusses the Turing Test, the AI winter and spring, expert systems, and the rise of modern statistical machine learning.

### **Session 3: Presentations - AI applications**

*Session Overview:* Groups make short presentations on modern applications on AI issues.

### **Session 4: Presentations - AI applications**

*Session Overview:* Groups make short presentations on modern applications on AI issues.

### **Session 5: Understanding digital data**

*Session Overview:* Covers the foundational technical knowledge required for understanding how raw digitized data is connected to progress in AI. Includes bits and the digitization of information, and the role of transistors and Moore's Law.

### **Session 6: How AI works - Machine Learning**

*Session Overview:* An introduction to ML techniques, with comparisons with pattern recognition and regression. Introduction to how prediction works in the ML context. This session covers how to think about different popular ML algorithms (logistic regression, trees, etc.).

### **Session 7: How AI works - Loss functions**

*Session Overview:* Discussion of how to evaluate AI models and further discussion of the costs and benefits of different AI models. Explanation of ROC curves, false positives and false negatives, and confusion matrices. Discussion of how the need for well labeled output limits which tasks are amenable to machine learning/automation. Will also cover platforms such as Kaggle, and how ML competitions are "specified", including prediction quality measures. Finally, we will walk through an example of how no-code tools are used for machine learning.

### **Session 8: In class Machine Learning lab**

*Session Overview:* Use of the open source [WEKA data mining tool](#) to complete an empirical case study on the use of machine learning to predict HR attrition in the IBM employee flight data set.

### **Session 9: How AI Works: Deep Learning**

*Session Overview:* Introduction to deep learning and how it works. Contrasts the benefits of hand-generating ML features vs. using a deep learning engine. Covers feed forward networks and back propagation.

## **Session 10: Deep Learning Lab**

*Session Overview:* Uses [Teachable Machine](#) to illustrate key concepts related to deep learning and to underscore the changing economics of deep learning. Students use deep learning to classify lung xrays as being diseased or healthy.

## **Session 11: Quiz**

*Session Overview:* Quiz on concepts from Sessions 2 through 9

## **Session 12: Labeling data for ML systems**

*Session Overview:* Discusses various approaches firms use to label data. Also covers the changing economics of the industry forming around large scale data labeling. Introduces the concept of returns to data scale in the AI economy.

## **Session 13: Reinforcement Learning**

*Session Overview:* Covers reinforcement learning. Discusses the q-learning approach to solving reinforcement learning problems. Includes discussion of applications where reinforcement learning has been successful, such as games and ride sharing.

## **Session 14: Generative AI & GANs**

*Session Overview:* Covers applications of AI to content generation, including GPT-3. Works through examples related to text, music, art, and code. Explains the tradeoffs and limitations of generative models, and the use of GANs to develop deepfakes and other high-quality artificial content. Discusses some of the controversies around the use of AI for content generation.

## **Session 15: AI Competition - The Economics of AI**

*Session Overview:* Discusses key inputs to AI - labor (e.g. developers), capital (e.g. the cloud), software (e.g. TensorFlow) - and the implications for AI-based competition. Also covers AutoML and what it implies for AI deployment.

## **Session 16: AI Competition - Networks and the Cloud**

*Session Overview:* We cover the foundational knowledge required for understanding how the ownership and governance of networks affects AI outcomes. Discussions of key initiatives like network neutrality, GDPR/CCPA, and the ACCESS act.

## **Session 17: AI Barriers: Bias**

*Session Overview:* Covers bias in ML algorithms. Included discussions of tradeoffs between accuracy and bias and questions around high profile cases such as the OPTUM system and the ProPublica/Northpoint application to criminal justice.

## **Session 18: AI Barriers: Explainability**

*Session Overview:* Discusses problems related to explainable AI and why explainable AI has become a major focus of AI development. What are the trade offs in organizational adoption? What new technologies are emerging to make AI decisions more explainable?

## **Session 19: Industry Analysis - HR**

*Session Overview:* Walks through a deep case analysis of an application of AI to HR applications, with an emphasis on challenges related to deployment, bias, interpretation, and legal issues that arise when applying AI to HR. Likely to involve a guest visit to provide industry context.

## **Session 20: AI, Ethics, & the Law**

*Session Overview:* Discusses the challenging issues around AI ethics. This includes the question of AI morality and integrates examples such as the Trolley problem. Half of the class session is dedicated towards an in-class debate.

## **Session 21: AI & Privacy**

*Session Overview:* Covers challenges at the intersection of AI and privacy, including the privacy paradox, legislating whether consumers must be required to “opt-in”, relevant case law, and differences across countries. Also covers differences in state-led and market-led approaches to managing information privacy as well as implications for competition in the technology industry.

## **Session 22: Industry Analysis - Healthcare**

*Session Overview:* Walks through a deep case analysis of an application of AI to healthcare applications, with an emphasis on challenges related to deployment, bias, interpretation, and legal issues that arise when applying AI to healthcare. Likely to involve a guest visit to provide industry context.

## **Session 23: AI Management & Governance**

*Session Overview:* Covers how organizations must change to integrate AI-supported decisions. Includes issues around governance (e.g. the role of AI councils), challenges with AI-led management, and how management may be different in organizations that rely heavily on data-driven decision-making. Discusses data-driven culture.

## **Session 24: AI & Jobs I**

*Session Overview:* Discussion of the impact of AI on the future of work. Includes job displacement, robots, universal basic income, and other related topics. Discusses how the workforce and organizations are changing to adapt to AI. Also discusses what types of managers are needed to govern AI driven decisions, including the emergence of new governance vehicles, such as AI councils and Chief Data Officers.

## **Session 25: AI & Jobs II**

*Session Overview:* Conducts a scenario planning exercise. Students are divided into groups and asked to “game out” scenarios and policy responses under different assumptions about the rate and direction of AI-led job displacement. In the last twenty minutes, students present back to the class their policy conclusions under the different scenarios.

## **Session 26: Final Exam**

*Session Overview:* Final exam covering all concepts covered through the course, with an emphasis on those covered in the second half.

## **Session 27: Final Project Presentations I**

*Session Overview:* Students present their “pitch deck” for a new AI driven venture.

## **Session 28: Final Project Presentations II**

*Session Overview:* Students present their “pitch deck” for a new AI driven venture.