AI, Business, and Society
Course Syllabus (May 2023)
Prof. Kartik Hosanagar

Overview: The course provides an overview of AI and its role in business transformation. The purpose of this course is to improve understanding of AI, discuss the many ways in which AI is being used in the industry, and provide a strategic framework for how to bring AI to the center of digital transformation efforts. In terms of AI overview, we will go over a brief technical overview for students who are not actively immersed in AI (topics covered include Big Data, data warehousing, datamining, machine learning, etc). In terms of business applications, we will consider applications of AI in Media, Finance, Healthcare, Retail, and other industries. Finally, we will consider how AI can be used as a source of competitive advantage. We will conclude with a discussion of ethical challenges and a governance framework for AI. No prior technical background is assumed but some interest in (and exposure to) technology is helpful. Every effort is made to build most of the lectures from the basics.

The primary purpose is to help you develop managerial understanding of AI and its applications. This will necessarily involve "looking inside the hood" of AI systems. That said, the course is not meant to teach you how to be an ML engineer. While we might "look" at code in a session or two, you will not be coding in this course.

Textbook: A Human's Guide to Machine Intelligence by Kartik Hosanagar

Pre-readings before start of day 2:

1. Book chapter readings: Many sessions have a reading that is usually a chapter from my book. I expect that you should be able to do these readings in an hour and will have ample time during our course (e.g. evenings) to keep up with readings. However, if you have other commitments outside of our classroom, I recommend that you try to finish the course readings in advance of the course. That will ensure you have your evenings for any other activities you may have planned.

Sessions

Please note that the entire course will be taught during the MBA opportunity week in May. There will up to 3 sessions in a day as follows (all sessions will be virtual).

Day 1 (Monday): 12.00 to 1.30 pm PT and 1.45 to 3.15 pm PT
Day 2: 10.15-11.45 am, 12.00 to 1.30 pm ET and 1.45 to 3.15 pm PT
Day 3: 10.15-11.45 am, 12.00 to 1.30 pm EST, and 1.45 to 3.15 pm PT
Day 4: 12.00 to 1.30 pm PT, and 1.45 to 3.15 pm PT
Day 5 (Friday): 12.00 to 1.30 pm PT and and 1.45 to 3.15 pm PT

Day	Session	Content	Readings/ Book Chapter
1	1	**Start of Module 01: Intro to Big Data and AI**	

		In session 1 we will discuss the transformative potential of AI. Specifically, we will discuss AI's promise as a General-Purpose Technology and look at data to suggest that AI is likely to change most industries, thereby requiring organizational change on the part of companies as well. We will discuss the basics of Big Data and data infrastructure in preparation for taking a deeper dive into these topics in the next session. Students will walk away from this session with an understanding that a data explosion is occurring and that new tools like machine learning can be applied to this data to extract intelligence and automate decision-making. Topics in session 1: ● AI for Business Course Intro ● Is AI a General Purpose Technology? ● Basics of Big Data and data infrastructure	
1	2	Session 2 will wrap up our discussion of Big Data. We will start with an overview of what Big Data is, how it is being generated, and why it matters. We will then go on to discuss some differences between traditional analytics and Big Data analytics, the general skillset required for Big Data analysis, and tools for working with Big Data. To this end, we will discuss tools for managing Big Data as well as tools for analyzing Big Data. With regards to managing Big Data, we will discuss data warehouses like Snowflake. Regarding analyzing Big Data, we will discuss data mining tools like clustering and association rule mining. Our discussion of machine learning will cover the three types of ML (supervised learning, unsupervised learning, and reinforcement learning). Finally, I will introduce a simulation on A/B testing that students will play as part of a later class session on reinforcement learning. Topics in session 2: ● Big Data Overview ● Big Data Analysis ● Data Infrastructure ● Data Mining ● Basics of AI: Supervised, Unsupervised, and Reinforcement Learning	
2	3	**Start of Module 02: Machine Learning** In session 3 we will discuss Artificial Intelligence and Machine Learning in greater depth. We will cover what AI is, types of AI, a brief history of AI, and expert systems as an early approach to building AI. We will then discuss machine learning, a newer approach to AI that addresses the limitations of expert systems. Our discussion of machine learning will cover the three types of ML (supervised learning, unsupervised learning, and reinforcement learning) as well as factors that influence accuracy in ML systems. We will discuss several specific ML methods in depth (logistic regression, decision trees, random forests, and neural networks) and introduce model selection. Topics in session 3: ● Artificial Intelligence Overview ● Machine Learning Overview + Types of ML ● Accuracy of ML models ● Specific ML Methods: A Deep Dive ● Model Selection and Validation ● Quiz 01 (covers sessions 1 & 2)	Chap 04, 05 of textbook

2	4	Generative AI: We will focus primarily on Large Language Models (LLMs) and Text to Image (T2I) models. We will discuss technical primitives (how do they work), their impact on business productivity, as well as some strategic questions around building and capturing value with Gen AI applications	ChatGPT and Search: HBR Article
2	5	Gen AI Lab Session Topics in session 5: ● RAG, finetuning ● Scaling models: number of parameters, neurons, training datasets ● Custom GPTs ● Prompt Engineering	
3	6	We will discuss challenges associated with model training, such as overfitting, in detail. We will also discuss validation strategies such as split training and testing, k-fold cross validation, and iterative cross validation. Additionally, we will look at a sample dataset and ML code (in Python) to understand how data preparation, model setup, and validation works in practice. The goal of this session is not necessarily to learn coding but instead to help students understand what the application of ML in practice looks like. Students will also appreciate how data preparation and interpretation of results are often more time consuming than coding of any ML model. Topics in session: ● Training and validation ● ML in practice ● Intro to Google Colab and Jupyter Notebook (Run code live) Other topics/agenda items: ● Quiz 02 (primary focus: sessions 2, 3, 4 & 5)	
3	7	In this session, we will also discuss ML Ops. Specifically, we will cover the practices and tools of traditional Dev Ops and how ML Ops differs from these. For example, we will discuss some of the unique challenges associated with ML Ops, especially with regards to testing and performance monitoring. We will also discuss existing tools for ML Ops as well as the concept of model interpretability as part of the ML Ops workflow. I'll also intro you to several WYSWYG tools for ML that you may find helpful as non-data scientists who might be users of Ml tools in your careers Topics in session: ● ML Ops ● AutoML (Google cloud) ● Teachable Machines ● TensorFlow Playground ● Gen AI Ops	Neural networks (20 minutes)
3	8	Our guest speaker, Ikkjin Ahn (founder/CEO of adtech company Moloco) will talk about building differentiation around AI in a crowded adtech market. The focus will be to both understand what it takes to deploy AI at scale and also how to capture value around it.	
4	9	**Module 03: Business Applications**	Chap 03 of textbook

		This session will concentrate on how AI augments personalization on the web, and will cover types of recommender systems and how they work, the impact of personalization on markets, and challenges associated with personalization systems. We will cover content-based recommenders, collaborative filters, and hybrid recommender systems in depth, and we will use several well-known music streaming services to illustrate how these various recommender systems are used in practice as well as the differences between them. Students will also come away with an understanding that personalization goes beyond product recommendations, instead encompassing customer interactions more broadly. Time permitting, we will also discuss Autonomous Vehicles as another application of AI. Topics in session: ● Recommender Systems ● Impact of recommenders on markets ● Other forms of personalization on the web ● Challenges with personalization ● ML in Finance: Fraud Detection ● ML in Finance: Additional applications ● Quiz 03 (primary focus: sessions 6,7,8)	
4	10	**Start of Module 04: AI Strategy and Governance** Session 10 will be the first of two class sessions on AI Strategy and Governance. Students will walk away with an understanding that, to create an edge with AI, companies need to not only use AI, but create the right strategic frameworks and organizational processes around AI. To this end, we will begin by discussing how AI can drive business transformation, with a particular focus on the idea that transformative technologies like AI are often difficult to apply initially, but companies that persevere tend to benefit. We will then move into discussing how companies can apply AI more effectively by taking a portfolio approach to AI projects. We will also discuss how the democratization of machine learning is lowering barriers to AI use from a technical standpoint. With regards to organizational structure, we will discuss five specific organizational strategies that can help companies generate value from AI. Topics in session 10: ● AI-Driven Business Transformation ● Developing a Portfolio of AI Projects ● Lowering Barriers for AI Use ● AI in the Organization Structure	
5	11	In session 11 we will continue our discussion around AI Strategy and Governance, focusing on the risks that AI technologies bring and how certain governance guidelines can help manage those risks. Specifically, students will come away with an understanding of how algorithms can make discriminatory decisions, a serious social risk. We will also discuss how social risks can lead to additional risks for firms, such as reputational, legal, and regulatory risks. We will discuss how providing some user control and calibrated transparency, as well as implementing audits, can all help to manage these risks. Topics in session 11: ● Risks with AI ● AI Governance	HBR Article Starting AI ventures without Data

		• Trends in AI • Course Takeaways	
5	12	Final session will include a discussion of building an investment and operating thesis around data and AI • Guest speaker 2 • Quiz 04 (primary focus: sessions 9,10, 11, 12)	Chapter 10 of textbook

Grading

1. In-class quizzes: 80 points (not all quizzes will have equal points)
2. AI simulation game (1-page writeup + in-class performance) = 15 points
3. Generative AI: 1 page writeup: 5 points
    a. Each submission is graded by 2 fellow students (students can assign a score of 0-5
    b. Each student grades at least 2 submissions.
    c. The TA and I can review the grades but we will assign the final grade which may be an average of peer grades but could be very different if we view the submission differently.
4. Participation (attendance + keeping up with readings) is a necessary for a Pass grade independent of the numerical grades in the above components

Guidelines for Project Presentations

1. Imagine you are a Product Manager for a company that makes a Generative AI productivity tool for a certain profession. In the homework, we looked at a writing assistant. You can stick to the context or consider other contexts such as using Gen AI to draft legal contracts, write code, transcribe doctor's notes, etc. You are trying to think through how best to design the AI-human collaboration (e.g. what should AI do versus humans, when should AI/humans step in, etc). For your project, I'd like you to think through the dimensions along which to organize human-AI collaboration. For example, one dimension is whether AI is a copilot for humans or an independent completer of tasks. Another dimension is delegation, i.e. does a human complete the tasks but delegate to AI when in doubt or does an AI complete the task and delegate to a human when in doubt.
    a. Your main goal for this project is to bring structure to this task and identify the key dimensions along which you can think through Generative AI's role as a productivity tool for humans. Such a structure will allow you to think through the design space and reason through the choices you need to make in a more logical fashion than a more haphazard set of choices you might otherwise make. You may benefit from thinking through how team-based collaboration is designed and how to borrow ideas from there for human-AI collaboration.
    b. Use the task context you picked and illustrate how these dimensions might play out in that setting.
2. Hint: try using some Gen AI assistants for various tasks like writing (Sudowrite, copy.ai, writesonic, Hyperwrite), coding (Copilot, Harness, Codesquire.ai), legal contracts (Evisort, Spellbook, Latchapp), etc. You will notice they make different choices for *how* to leverage AI in the workflow.
3. The final deliverable can be a set of slides. No required length. The reports will be judged on content. Please prepare slides that can be read and interpreted independently by TAs and myself (i.e. designed for independent consumption in your absence). We will evaluate the reports for the following:
    a. Quality of information gathered ("research")
    b. Creativity and structure you brought to the problem
    c. Coherence ("writing; structure; logical presentation")

d. Examples (illustrated the concept well with a clear case study)

4. On the last day, I will ask all/most teams to present their project (it is in your best interest to have 1 person present because time management is usually better achieved with one person presenting).
5. The goal of the presentation would be convey the basic idea in under 5 minutes. I realize this is insufficient time to showcase all the work you have done (☺), but the other option will be to have an additional class.
6. Additional details regarding the presentation slides and the presentation will be emailed later.

Grading Approach: All team members usually receive the same grade unless there is broad consensus among members of a team that effort has been disproportionate.