

Machine Learning for Business

FALL 2019

Instructor: Helen Mou & Nick Ross

Time and Location: 101 Howard, Monday from 6-9 PM starting October 21st through Monday December 2nd.

Contact Information: helenmou@gmail.com ncross@usfca.edu

Overview and Objectives

Data is at the core of product management. Today's product managers are increasingly responsible for shipping machine learning-driven product features, making critical product decisions based on machine learning techniques, and developing strong partnerships with data science counterparts. Our Machine Learning for Business certificate provides the foundational understanding of data science and machine learning necessary for extracting business value from the data produced by your product and your organization.

The certificate covers supervised and unsupervised machine learning topics, including classification, clustering, recommendation systems, and natural language processing. Implementation considerations, ethical standards and business strategies around these topics will also be covered. Technical topics are taught with a focus on applications, algorithm design, and product tradeoffs.

This program is geared towards product and business managers. While there are no formal technical requirements, we find previous experience in product or business management and data analysis to be helpful.

Machine Learning for Business is taught by Nicholas Ross and Helen Mou. Professor Ross is currently Practicum Director and Assistant Professor for USF's Master of Science in Data Science program, previously working in the video game industry. Helen Mou works as Product Lead at Shopify where she leads cross-functional teams to build Kit, a virtual employee for Shopify's 800,000+ entrepreneurs.

Learning Objectives

This certificate has the following learning objectives:

1. Define the difference between analytic and production systems.
2. Understand common use cases for different data science techniques (clustering, classification, recommendation systems, regression and NLP).
3. For each of the above techniques, identify data constraints and potential fitting issues. Describe how parameter choice effects each of these techniques.
4. Critically assess data science techniques from an ethical perspective.
5. Analyze data science problems and make optimal decisions with regards to tool choices.

Course Format

Generally, speaking the course format will be one and half hours of lectures, followed by a break and then a guest speaker.

Guest speakers are chosen based on their background and knowledge and will give a short (20-30 minute presentation) followed by a Q&A.

While no NDA has been signed to participate in this course, please use discretion when repeating information from guest speakers. They are actively giving their time and effort in order to help you learn from their experiences and would appreciate you using your judgement before sharing.

Textbook

There is one required text for this course, *Data Science for Business: What You Need to Know about Data Mining and Data-Analytic Thinking* which will be provided to all students. Additional readings will be sent out before each class.

Schedule

The certificate will follow the following schedule:

- October 2: Class #1
 - Introduction: Chapters #1 & #2
 - Guest Instructor: Vincent Chio, Shopify
 - Topic: Personalized marketing recommendations at Shopify (in combination with Helen: Intelligent Kit)
- October 28: Class #2
 - Introduction to ML Models
 - Chapter #3: Basic Classification
 - Guest Instructor #1: Connor Ashenbrucker (Head of Customer Success) and Christa Simone (UX Researcher), No Red Ink.
 - Guest Instructor #2: Michael Gorman (Senior Advisor) + Huanjin Chen (Chief Technology Officer), ShareThis
- November 4: Class #3
 - Recommendation Systems
 - Similarity, Neighbors and Clusters (Chapter #6)
 - Guest Instructor: Kabir Bedi, Amazon
- November 11: Class #4

- Organization, Team Roles and Management
 - Data Science Tools
 - Unsupervised Learning (Chapter #6 cont.)
 - Guest Instructor: Emma Kelsey and Ting Ting Liu, Branch International
- November 18: Class #5:
 - Final lesson on Classification
 - Overfitting, AUC, ROC (Chapter #5 & #7)
 - Guest Instructor: Kevin Hu and Kate Adamson, importance of interpretability and examples from Plaid. Coming with Kate Adamson, who was also previously head of mortgage, and has experience with sales + growth, and perspective on how ML interacts with business/client needs.
 - November 25: Class #6
 - NLP
 - Guest Instructor: Deven Dharm, Deloitte
 - December 2: Class #7
 - Deep Learning
 - Ethics
 - Guest Instructor: Gio Gujarati, Valimail

Receiving the certificate

In order to receive this certificate, students must attend all but one lecture as well as actively and positively participate in the course. If a student either misses more than one course *or* is disruptive, negative or aggressive to any of the guest speakers they will be asked to leave and not receive the certificate.