

Enterprise AI Strategy (B8152)

Course Syllabus

Summer Semester 2021 – A Term

Professors Daniel First & Daniel Guetta

In almost every corner of the economy, analytics has gone from a siloed, technical activity to a centerpiece of transformation. Traditional businesses are reinventing themselves by creating entire new business lines with their data assets, personalizing their services, shaping their strategy using data-driven insights, building predictive models to anticipate the future, and optimizing the efficiency of business processes. *Fortune's Fastest Growing Companies list* included 32 tech companies in 2018 and 2019, but only 20 in 2020 – “tech” has moved from being an industry of its own to a competency essential for today’s most competitive businesses.

Yet the road to analytics nirvana is paved with well-publicized failures. In finance, firms such as [Point72](#), [Blue Mountain Capital](#), and [BlackRock](#) have faced serious challenges trying to integrate AI and data science into their operations. IBM’s Watson program experienced a well-publicized [setback](#) in their work with the MD Anderson Cancer Center. Amazon [scraped](#) an application evaluation algorithm when it was later shown to be biased against women. SoFi’s attempt to issue loans without FICO scores [reportedly](#) led to an increase in default rate followed by a quiet re-integration of FICO data. The UK government’s [attempt](#) to use analytics in lieu of exams to automatically assign grades to high-school students at the height of the pandemic led to widespread outrage and a quick reversal. Gartner estimates [85%](#) of big data science projects fail. Why do companies find it so hard to capture the value analytics offers?

A major reason for these difficulties is that executives are ill-prepared to guide and enable data scientists to deliver meaningful and lasting impact in this new world of ubiquitous analytics.

This course is intended for future CxOs, heads of business units or functions, or anyone who will one day oversee analytics initiatives, directly or indirectly. The course will also be useful for anyone going into consulting who will support leaders in carrying out analytics transformations in their organizations.

This course will teach you how to be a business leader who can transform your organization using analytics. You will learn how to ideate business use cases from your data assets, build analytics roadmaps for your organization, manage data science projects, think critically about business recommendations based on data, report on analytical findings, and leverage data-driven recommendations to transform your business.

Pre-requisites

There are no pre-requisites for this class, though you may find synergies with Business Analytics, Business Analytics 2, Modern Econometrics for Business, and Sports Analytics. The class also serves as an excellent complement to Analytics in Action.

No knowledge of data science, AI, coding, or mathematics is needed for this class, and this course does not include any coding or mathematical exercises.

Course outline

We will structure our course around the lifecycle of a data science project, broadly divided into four steps: Understanding the current state of analytics at your organization, identifying high-value analytics use cases, managing an analytics project, and moving from analytics to business impact.

The course will be heavily based around real-world cases. Example cases may include (all specifics subject to change):

- A sports team undergoing an analytics transformation as they look to improve team performance and reduce injuries. The team has more than twenty different sources of data, ranging from blood tests, to heart rate monitors, to velocity and acceleration trackers. They also have a number of PhDs in their data science department, but they are underutilized and their work is poorly understood. We'll follow them as they learn to use this data to prevent and diagnose injuries in real time.
- A biotech engineering company where projects planned for two years last five on average. Executives are struggling to use the oceans of data available to them to understand why these delays happen and how to prevent them. Learn how executives begin to define and measure KPIs, predict and prevent delays using machine learning, and improve the efficiency and performance of their organization and equipment using data-driven insights.
- A New York restaurant chain, famed for its healthy nutritious menus and seasonal offerings, as they look to scale. They have no team members that specialize in data analytics, and their datasets are spread across a broad range of systems. As the number of restaurants, staff, and operations of the group grow, intuition is no longer enough to make decision. We will follow them as they start creating a modern data infrastructure, and move from intuition to data as their main decision-making tool.
- A credit card company with a simple legacy credit decisioning model, based on a single dataset and heavily reliant on FICO scores that wants to transition to a more complex, multi-faceted technique for deciding whom to issue credit cards to. We will discuss the challenges they face in their path to this improved model, and the surprising ethical and legal hurdles that arise in their journey.
- An app targeted at mom-and-pops restaurants, and giving them a low-cost option for accepting orders online. The company is run by a family of restaurateurs with extensive

restaurant experience but little analytics experience. They are looking to use analytics to improve customers' experience in the app, including product recommendations and targeted offers. We will discuss their initial steps in taking this amorphous problem and creating a roadmap to its solution.

Course schedule

(Session details and schedule is preliminary and subject to change)

Introduction		Course overview	Introduction, motivation, goals. Famous examples.
	Module 1	Big Data: A Crash Course for Executives	What you need to know to discuss Big Data and AI intelligently in the boardroom. Types of analytics (descriptive, predictive, explanatory, optimization) and techniques (causal inference, deep learning, reinforcement learning). KPIs.
		Understanding the current state of analytics at your organization	People (organizational structure, skillsets), Process (how is data science done), Technology (data storage, processing, analytics); analytics software providers and consulting companies
Building an analytics roadmap	Module 2	Identifying use cases	Understanding the archetypal analytics projects and the business value they provide. Brainstorming; what can analytics do for my organization? Process- and data-driven brainstorming.
		From use case to analytic solution	The elements of a data science project: Target variable, unit of analysis, features. Using these concepts to formulate and evaluate a data science project.
	Module 3	Prioritizing & evaluating analytics use cases	Evaluating use case for business impact, technical feasibility, data feasibility adjacency to capabilities, and business roadmap. Separating the hype from the unpolished gems.
	Module 4	Building an analytics roadmap	6-month, 1-year, and 3-year objectives. Where to start. Calculating the business impact of the project. Communicating the impact to executives to receive the go-ahead.
		Creating a project workplan	The five phases of an analytics project. Agile principles. Showing outputs at the outset. Dividing workstreams. Team composition. How analytics projects differ from traditional consulting projects. The art of failing early and pivoting.
Module 5	Running a project & generating insights	Data access and infrastructure, technical tools, productivity tools. Defining sprint goals, standups, and communications to leadership. Effective work norms with technical colleagues. Generating hypotheses & how SMEs can enhance	

			the data science process. Shadow decks. Thinking critically about data-driven conclusions.
		Communicating findings	Communicating analytical insights and ML models. Synthesis, storylining, BLUF, top-down communication. Effective communication of results. Creating insights that are actionable.
Module 6		Risks and Ethics of AI	The seven types of risk in a machine learning project. Ethical pitfalls in machine learning models. Addressing ethical concerns.
		<i>Time permitting:</i> Translating into recommendations for impact	Evaluating different types of impact, creating actionable recommendations. Experimenting and monitoring KPIs over time. Building trust and process re-design.

Requirements and Grading

Grades for the class will be determined as follows

- 25% : attendance class participation
- 25% : Assignments (*graded on effort*)
- 50%: Course project

Assignments will be based on class discussions, and will be group efforts, completed with the group with whom you had the discussion in class. Details will be provided in class.

In addition, there will be a final course project to be completed in a team of 5 or 6 students. As a team, pick any company of your choice and suppose the CEO has tasked you with coming up with a roadmap for how the company might better use analytics to drive value at the company. Prepare a presentation (maximum 15 slides) for the CEO. The presentation should include the following

- **Executive Summary**
- **Part 1: How might this company use analytics?** This section might touch upon KPIs the company could monitor, dashboards they may create, machine learning models they may introduce, or questions they might address using data-driven insights.
- **Part 2: What is this company's data landscape?** What data do you think this company has? Are there any data sources the company has, but that you think might be under-utilized? What data quality or integration issues do you think this company may face? How could they be collecting data better? Are there any external data sources they might want to acquire?
- **Part 3: Feasibility study.** Evaluate your ideas from part 1 in terms of their business impact, feasibility, level of effort required, and likelihood of delivering business impact. Refer to the data landscape in Part 2 as part of your study. Make a recommendation about which one to start with, and size the business impact of this top recommendation. Carry out a risk assessment for this top proposal – what could go wrong? Are there any ethical considerations?
- **Part 4: Project plan.** Create a project plan for your top recommendation, broken down into phases.

You will be evaluated on the relevance of the questions your report addresses, the creativity of the analytics projects you suggest, the feasibility of these problems, your ability to apply principles from this course to assess potential issues that arise with different types of projects, and the clarity of your communication.

Time permitting, some teams might be asked to present their idea in our last lecture.