

Business Analytics Syllabus

Instructor

Prof. Ciamac Moallemi

Decision, Risk & Operations

Columbia Business School

Office: Uris Hall, room 416

E-mail: ciamac@gsb.columbia.edu

Office hours: Mondays, 12:30pm–2pm (or by appointment)

Teaching Assistants

Nikhil Kumar

E-mail: nkumar15@gsb.columbia.edu

Office hours: Wednesdays, 12:30pm–2pm (or by appointment), location TBA

Kai Yuan

E-mail: kyuan17@gsb.columbia.edu

Office hours: Tuesdays, 5:00pm–6:30pm (or by appointment), Uris 4th floor cubicle 4J

Course description

Business analytics refers to the ways in which enterprises such as businesses, non-profits, and governments can use data to gain insights and make better decisions. Business analytics is applied in operations, marketing, finance, and strategic planning among other functions. The ability to use data effectively to drive rapid, precise and profitable decisions has been a critical strategic advantage for companies as diverse as WalMart, Google, Capital One, and Disney. For example, Capital One uses sophisticated analytic capabilities to match credit card offerings to customers more accurately than their competition. WalMart uses analytics to monitor and update its inventory in a way that allows it to serve its customers at an exceptionally low cost. In addition, many current and recent startups such as Palantir and Splunk are based on the application of analytics to large data bases. With the increasing availability of broad and deep sources of information — so-called “Big Data” — business analytics are becoming an even more critical capability for enterprises of all types and all sizes.

In this course, you will learn to identify, evaluate, and capture business analytic opportunities that create value. Toward this end, you will learn basic analytic methods and analyze case studies on organizations that successfully deployed these techniques. In the first part of the course, we focus on how to use data to develop insights and predictive capabilities using machine learning, data mining and forecasting techniques. In the second part, we focus on the use of optimization to support decision-making in the presence of a large number of alternatives and business constraints. Finally, throughout the course, we explore the challenges that can arise in implementing analytical approaches within an organization.

The course emphasizes that business analytics is not a theoretical discipline: these techniques are only interesting and important to the extent that they can be used to provide real insights and improve the speed, reliability, and quality of decisions. The concepts learned in this class should help you identify opportunities in which business analytics can be used to improve performance and support important decisions. It should make you alert to the ways that analytics can be used — and misused — within an organization.

We have three goals in this course. The first is to help you think critically about data and the analyses based on those data — whether conducted by you or someone else. The second is to enable you to identify opportunities for creating value using business analytics. The third is to help you estimate the value created using business analytics to address an opportunity. Business analytics is an integral part of modern management — this course should provide you with the foundation you need to understand and apply these methods to drive value.

Course materials

There is no required textbook for the class. There will be cases, articles as well as slides that we will distribute in each class, and post on canvas.

Requirements and Grading

There will be three homework assignments to hand in, preparation assignments for sessions, and a final. In addition, there will be periodic “concept checks” — short quizzes to be done individually to check your understanding of the material. If you understand the material, it should take you no more than 10 minutes to do each concept check.

Your grade will be determined as follows:

- Final exam: 50%
- Participation: 10%
- Concept checks: 10%
- Homeworks: 30%

You are expected to come to class prepared, and ready to discuss the pre-class reading, case or assignment questions.

Homework: Homework assignments will need to be submitted on Canvas. There is one individual assignment and two group assignments (Type A, to be done by learning teams).

- Homework 1 (individual assignment) is due (i.e., should receive a Canvas time stamp) before the start of session 4.
- Homework 2 (group assignment) is due (i.e., should receive a Canvas time stamp) before the start of session 7.
- Homework 3 (group assignment) is due (i.e., should receive a Canvas time stamp) before the start of session 11.

Software: This course will require the use of Windows Excel and we will provide a business analytics excel add-in. This add-in only runs under Windows, so Mac users will need to be able to boot into Windows.


Syllabus at a Glance

Session 1: Introduction

Predictive Analytics

- **Session 2:** Predicting outcomes I / lending analytics
- **Session 3:** Predicting outcomes II / recommendation analytics
- **Session 4:** Quality of predictions I / healthcare analytics
- **Session 5:** Quality of predictions II / financial analytics
- **Session 6:** Predictions and skill versus luck / sports analytics

Prescriptive Analytics

- **Session 7:** Testing / retail analytics
- **Session 8:** Simulating the future / pension analytics 
- **Session 9:** Optimizing complex decisions / salesforce analytics
- **Session 10:** Optimizing with multiple objectives / portfolio analytics

Implementation

- **Session 11:** Decision-support systems -from concept to deployment- / supply chain analytics
 - **Session 12:** Review
-

Other Resources

Industry speakers on campus:

- April 16, 2014: Ricardo Sanchez, Technical Director for North America, Cintra (visit <http://www7.gsb.columbia.edu/cprm/events> for more information)
- The Institute for Data Sciences and Engineering sponsors a number of talks that are relevant to the topic:
(visit <http://idse.columbia.edu/events/calendar> for more information)

Optional recommended readings:

While there is no required textbook for the class and there is no definitive and comprehensive reference on the subject, there are several books we recommend in case you would like to learn more and follow up beyond what we will cover in class.

- For data mining and predictive analytics, the following are introductory textbooks:
Data Science for Business, Provost and Fawcett: O'Reilly
Data Mining for Business Intelligence, Concepts, Techniques and Applications, Shmueli, Patel, and Bruce: Wiley
For Excel modeling and optimization, the following is a good textbook:
Management Science: The Art of Modeling with Spreadsheets, Powell and Baker: Wiley
The three books above are on reserve at the Watson library.
- For scientific retail management, the following is an interesting account of various approaches that may be adopted to leverage analytics:
The New Science of Retailing, Fisher and Raman: Harvard Business Press
- An interesting recent industry report on big data and the role that analytics may play in industry is:
Big data: The next frontier for innovation, competition, and productivity
and can be found at:
http://www.mckinsey.com/insights/mgi/research/technology_and_innovation/big_data_the_next_frontier_for_innovation

Optional Software:

- Tableau: this is a data visualization tool, which should already be installed on your computers. If not it may be downloaded at <http://www.tableausoftware.com/products/desktop/download>
You may use the key: TD0L-6BD7-9C80-6C79-0484 to activate it.
- Rattle: this is an open source R-based data analysis tool.
<http://rattle.togaware.com/>
An introductory book to Rattle
Data Mining with Rattle and R: The Art of Excavating Data for Knowledge Discovery, Williams: Springer
is on reserve at the Watson library.
- Orange: this is an open source data-mining tool which is intuitive and workflow oriented.
<http://orange.biolab.si/>