Sample syllabus from IMA Math-to-Industry Boot Camp II

Introduction

The following syllabus is available to support the development of graduate student training programs and internships in business, industry, and government (BIG).

The IMA Math-to-Industry Boot Camp is divided into two sections: technical skill building in weeks 1-3 and projects in weeks 4-6. Interspersed in the course are other activities, including:

- Discussions with industrial scientists
- Software demonstrations
- Resume preparation and other career-related, skills-building activities
- Team-building and social activities

Required preparation

- Bring your own laptop and have the appropriate software installed and working
- Software uploading for Python, R, RStudio, AMPL
- Course prep for Python segment
- Course prep for statistics segment
- Course prep for stochastic modeling segment
- Course prep for optimization segment
- Myers-Briggs Type Indicator (MBTI) assessment

Lecture schedule

The lecture topics covered include programming in Python, statistics, machine learning, stochastic modeling, and optimization. Below are the syllabi and software requirements used for each lecture topic.

Topic: Programming in Python

Instructor: Sarah Miracle (University of St. Thomas)

Programming in Python syllabus

Course description:

This tutorial is designed for students with no programming experience and will quickly cover the basics of procedural programming in Python. Topics will likely include: numeric data types, variables, assignment statements, functions, using modules, loops, nested loops, conditionals, Boolean expressions, accumulator patterns, strings, lists, testing, recursion and some algorithms.

Preparatory material:

There are many great online references for Python.

Read before class:

- Chapter 1 in *How To Think Like a Computer Scientist.*
- CS for All
- Python documentation

Course material:

Example files, project descriptions, lecture notes, and other course documentation are available in this Dropbox folder.

Software requirements:

Python 2.X and the IDLE integrated development environment (which installs with Python)

Visit the Python website to install software by clicking on downloads and selecting the latest 2.X version (likely 2.7.13). For installation help, the Techwalla website has an article on "How to Install Python IDL."

Topic: Statistics

Instructor: Alicia Johnson (Macalester College)

Statistics syllabus

Course description:

The goal of this tutorial is to build statistical, data, and computer literacy. Since all techniques of statistical modeling and inference techniques can't be covered, this tutorial will focus on multiple linear regression. In doing so, topics in traditional statistics introductions will be bypassed in order to establish foundational concepts that can be generalized to other modeling techniques, such as spatiotemporal models, mixed effects models, etc. This tutorial will favor statistical applications using real data over statistical theory so that students walk away with a sophisticated set of tools with real applications.

Preparatory material:

Read and watch all written and video content in this pre-boot camp online tutorial. Students are strongly encouraged to try out all example RStudio code on their own computers. Additionally, complete all homework exercises in the pre-boot camp online tutorial.

Course material:

Course notes, homework, and RStudio resources are available online.

Software requirements:

R and RStudio

Install R software at The R Project for Statistical Computing website. Install RStudio software at the RStudio website.

Topic: Machine learning

Instructor: Gilad Lerman (University of Minnesota, Twin Cities)

Machine learning syllabus

Course description:

The tutorial will introduce students to machine learning methods that are common in industry and that can be understood in a relatively short time. Topics include:

- K-nearest neighbors as a simple solution for some problems in machine learning
- Support vector machines: maximal margin classifier, support vector classifiers, non-linear decision boundaries and the kernel trick
- Tree-based methods: regression and classification trees, bagging, random forests, boosting
- Principal component analysis: quick review of singular value decomposition, statistical interpretation and applications.
- Two examples of nonlinear methods: SNE/tSNE and SOM
- Clustering techniques: k-means, mixture of Gaussians, hierarchical, DBSCAN, nearest neighbors, spectra

Preparatory material:

An *Introduction to Statistical Learning with Applications in R* by Gareth James, Daniela Witten, Trevor Hastie and Robert Tibshirani (June 2013).

Software requirements:

R

Topic: Stochastic modeling

Instructors: Jasmine Foo (University of Minnesota, Twin Cities)
Kevin Leder (University of Minnesota, Twin Cities)

Stochastic modeling syllabus

Course description:

Topics to be covered:

- Basic stochastic processes – Poisson processes, continuous time Markov Chains, birth-death processes, and renewal processes
- Queueing theory – Applications of queueing models, fundamental results of queueing theory, birth-death queueing models

Preparatory material:

Probability section of the Introduction to Probability and Statistics class available at MIT OpenCourseWare.

References:

- Introduction to Probability Models by Sheldon Ross, 11th edition.
- Probability and Random Processes by G. Grimmet and D. Stirzaker.
- Course notes

Software requirements:

None

Topic: Optimization

Instructor: [Qie He (University of Minnesota, Twin Cities)]

Optimization syllabus

Course description:

This short course introduces how to apply optimization models and methods to solve decision making problems arising in many industries such as logistics, transportation, manufacturing, finance, and healthcare. The ultimate goal is teach students how to solve large-scale optimization problems from the real world in a scientific way. Students will receive trainings in the following three aspects: modeling, algorithms, and software. In particular, students are expected to:

- Improve the ability of formulating complex decision problems into optimization models in concise mathematical forms.
- Learn how to transform an optimization model written by humans into computer programs compatible with optimization software, with high-level modeling languages such as AMPL;
- Understand the capability and limitation of a variety of optimization software on the market, and be able to select the right solver and fine-tune it to solve your model;
- Understand the basic principles used in many optimization algorithms, and be able to develop customized solutions for non-standard optimization problems.

Topics to be covered:

1) Deterministic optimization
   a) Optimization applications and hands-on AMPL experiences (the first 1.5 hours)
      i) Two real applications of optimization
      ii) Your first AMPL program: Steel Production
      iii) The second AMPL program: Where should Walmart build its warehouses?
      iv) The third AMPL program: Amazon's two-hour delivery
   b) An engineer's perspective on optimization (the second 1.5 hours)
      i) A typical decision making process in practice
      ii) An overview of optimization models and existing solvers
      iii) Advanced modeling techniques with examples (scheduling, cutting stock, matching, etc.)
2) Optimization under uncertainty
   a) The impact of uncertainty on decision making (the first 1.5 hours)
      i) The fourth AMPL program: a farming example
      ii) The classical newsvendor problem
   b) An engineer's approach to handle uncertainty in decision making (the second 1.5 hours)
      i) The static models and where you can use them
         (1) Expectation minimization
         (2) Chance constrained optimization
         (3) Robust optimization and distributionally robust optimization

ii) The dynamic models and the challenges

Preparatory material:

Read the AMPL installation guide. Download and install the AMPL package. AMPL will be used to code and solve the optimization models formulated in class.

Read the introduction chapter of the AMPL book.

Course material:

The AMPL package, lecture slides, and exercises can be downloaded from this Dropbox folder.

Resources:

Textbook

- Robert Fourer, David M. Gay, and Brian W. Kernighan. AMPL: A Modeling Language for Mathematical Programming (2nd edition).

Recommended readings

- Bertsimas, D., & Tsitsiklis, J. N. (1997). Introduction to linear optimization. Belmont, MA: Athena Scientific.
- Boyd, S., & Vandenberghe, L. (2004). Convex optimization. Cambridge University Press.
- Nocedal, J., & Wright, S. (2006). Numerical optimization (2nd ed.). Springer Science & Business Media.
- Nemhauser, G.L., & Wolsey, L. (1999). Integer and combinatorial optimization. John Wiley & Sons.
- Birge, J., & Louveaux, F. (2011). Introduction to stochastic programming (2nd ed.). Springer Science & Business Media.
- Shapiro, A., Dentcheva, D. & Ruszczynski, A. (2014). Lectures on stochastic programming: modeling and theory (2nd ed.). SIAM.

Software requirements:

AMPL

For installation help, use this AMPL installation guide.