OR 604: Practical Optimization George Mason University

Department of Systems Engineering and Operations Research

Fall 2017

Time: Mondays 7:20 PM to 10:00 PM

Classroom: Nguyen Engineering Building, Room 1110

Professor: Dr. Steven Charbonneau

Phone: (202) 418-4076 (work) from 08:30 AM to 5:30 PM

(703) 550-5006 (home) from 6:30 PM to 9:00 PM

Email: scharbo2@gmu.edu

Office hours: By Appointment only

Text:

Required: Optimization in Operations Research, 2nd Edition, Rardin, R.L., Prentice Hall, 2017 (0134384555) (https://www.amazon.com/Optimization-Operations-Research-Ronald-Rardin/dp/0134384555/ref=sr_1_1?s=books&ie=UTF8&qid=1484419441&sr=1-1&keywords=rardin+operations+research)

Course Description: This course builds on the techniques learned in Analytics and Decision Analysis (OR 531). As data of huge sizes becomes ubiquitous, analysts must learn to set up, formulate, solve, and interpret prescriptive models of unprecedented size. This course describes optimization techniques and codes capable of working in the Big Data setting. The course will use state-of-the-art optimization packages coupled with a programming language and databases suitable for big data analyses. The course covers model formulation, convexity, linear programming, integer programming, and heuristic methods. *You will be writing and running your own code for each homework assignment. Do not be confused, it is a very large part of this course. If you do not know how to write code, you will be provided the resources and adequate opportunity to learn through application and frustration. This means you will allocate more time to this course in the beginning of the semester than your peers. If you have concerns, start the tutorials that are in the course blackboard site before the course starts.*

Course Objectives: This course focuses on developing student skills in the formulation and implementation of large-scale optimization models. At the end of this course, students will be able to:

- Formulate large-scale optimization models by taking advantage of the regular structure of sets of constraints
- Solve large-scale optimization models by employing advanced techniques and heuristics that allow generally intractable problems to solve
- Take advantage of the advanced features of commercial optimization engines

- Set up, read from, and write to databases using SQL statements run from a programming language
- Create and run routines in Python that can be used to formulate and solve optimization problems.

Course Schedule: The course schedule is subject to change as the course progresses. Modifications will be posted on blackboard.

Lecture	Date	Topic	Prep Work
Lesson 01	AUG 28, 2017	**Model Formulation I** • Course Introduction • Software Stack Overview • Mathematical Programming Overview	Text: Chapter 1, 2, 3
Labor Day – No Class			
Lesson 02	SEPT 11, 2017	**Model Formulation II** • Resource Allocation Models • Blending models • Formulating models in Gurobi • Checking models in Gurobi	Text: Chapter 4 Sections 1 and 2 Papers: Summation Notation Doc 1 Summation Notation Doc 2
Lesson 03	SEPT 18, 2017	**Model Formulation III** • Scheduling models • Inventory models • Reading and understanding the Gurobi output and LP files	Text: Chapter 4 Sections 4 and 5 Papers: Converting your paper models to Python/Gurobi models
Lesson 04	SEPT 25, 2017	**Model Formulation IV** • Assignment models • Transportation models • Dominos case study overview	Text: Chapter 10 Dominos Case Study
Lesson 05	OCT 02, 2017	**Model Formulation V** • Integer Programming • Binary Programming • Converting Dominos to a Binary Programming Problem	Text: Chapter 12 Dominos Case Study

Lesson 06	OCT 10, 2017 (Tuesday Class)	**Model Formulation VI** • Facility location problem • Vehicle routing problem • Sweep algorithm	Text: Chapter 11 Dominos Case Study Papers: A Tutorial on Column Generation and Branch and Price for Vehicle Routing Problems (read to understand pages 1-3, scan the rest of the paper) A Sweep Algorithm for the Mix Fleet Vehicle Routing Problem
Lesson 07	OCT 16, 2017	**Model Formulation VI** • Introduction to the NFL Case Study • Formulating Core Constraints	NFL Case Study Paper: Sports Scheduling: Algorithms and Applications
Lesson 08	OCT 23, 2017	**Model Formulation VII** • Formulating more complex constraints	NFL Case Study
Lesson 09	OCT 30, 2017	**Model Formulation VIII** • Soft Constraints • Goal Programming	Text: Chapter 8 Paper: Hard and Soft Constraints in Linear Programming NFL Case Study
Lesson 10	NOV 06, 2017	**Heuristics I** • Probing the NFL Model • Reading in a model • Recreating the modeling environment from a file	AWS Tutorial TBP: Advanced Gurobi Features I
Lesson 11	NOV 13, 2017	**Heuristics II** • Multi-processing - turning the probing heuristic into a multi-processed routine	AWS Tutorial TBP: Multi-processing in Python
Lesson 12	NOV 20, 2017	**Heuristics III** • Reduced Models • Adding feasibility cuts	TBP: Advanced Gurobi Features II

		• Adding solution cuts	
Lesson 13	NOV 27, 2017	**Heuristics IV** • Distributed programming • Sleeping routines • SSH	TBP: Distributed Programming TBP: SSH in Python
Lesson 14	DEC 04,2017	**Heuristics V** • Wrapping it all up	No New Reading
Reading Day	DEC 11, 2017	**Do not make travel plans**	We'll know by NOV 13

Grading Scheme:

Homework: 80%

Final/Class Participation: 20% (ask me about this – please)

Coursework & Grading: Unless otherwise indicated, you are expected to work individually on homework assignments. You must submit homework directly to me via email at scharbo2@gmu.edu or through blackboard.

Academic Integrity: GMU is an honor code university; please see the University Catalog for a full description of the code and the honor committee process. The principle of academic integrity is taken very seriously and violations are treated gravely. What does academic integrity mean in this course? Essentially this: when you are responsible for a task, you will perform that task on your own. When you rely on someone else's work in an aspect of the performance of that task, you will give full credit, in writing, as a cover document to your homework submission. Another aspect of academic integrity is the free play of ideas. Vigorous discussion and debate are encouraged in this course, with the firm expectation that all aspects of the class will be conducted with civility and respect for differing ideas, perspectives, and traditions. When in doubt (of any kind) please ask for guidance and clarification.

GMU Email Accounts: Students must use their Mason email accounts to receive important University information, including messages related to this class. See http://masonlive.gmu.edu for more information.

Additional Notes:

I will make every effort to use Blackboard to post homework, assignments, lecture notes, and grades. I will send out email notices each time I have uploaded new information to blackboard.

Failure to turn in homework on the due date will result in a 0% for that submission.

Best way to contact the professor is by email.

The following is a great source for getting started on Python and SQL: The GMU library has an academic subscription to Safari Books Online; a repository of technical books on programming, big data, math, and all sorts of IT subjects. These are books (from publishers like Wrox, Sams, and O'Reilly) you generally would spend between $50 and $150 if you wanted to buy your own copy. As an academic license, it is limited in its features (cannot bookmark or highlight pages) but it is an excellent resource. Get to it from the following URL and use your GMU credentials to log on:
https://login.gmu.edu/login?service=https%3a%2f%2flogin.mutex.gmu.edu%2flogin%3fqurl%3dezp.2aHR0cDovL3Byb3F1ZXN0LnNhZmFyaWJvb2tzb25saW5lLmNvbS8.dWljb2RlPXZpdmE-

Hardware:
You will need to bring a laptop to class each lesson. If you have both a Mac and a PC, I encourage you to bring the PC. If you were looking for a reason to buy a new laptop and you think this course is your reason – good on you. I strongly encourage you to get a PC based laptop running Windows with an Intel processor (as opposed to AMD).

1. *Why Intel:* For some reason, Intel processors are better designed for math intensive operations whereas AMD has its strengths in other areas (*I have no idea what that would be*). Case in point, two former students bought brand new AMD laptops with four cores, 3.5+ GHz speed, with 4 to 8 GB RAM specifically for this course. Their laptops could not get a feasible solution in less than 10 hours. Running their models on my 5 year old Intel i3 duo core laptop with 4 GB RAM and 2.1 GHz speed, I could get a feasible solution in about 10 minutes. Your call on how you want to spend your time.
2. *Why Windows:* If you have a Windows laptop your life will be much easier (*in this course at least*) than if you try to use a Mac laptop. I didn't write the Gurobi software, I'm just telling you the cold hard facts. For you Windows haters out there, sometimes you just have to acknowledge that Mac doesn't always mean easier. If you insist on using a Mac, once (if) you get your environment up and running, you should be okay. The latest release of Gurobi may have fixed this situation. In the past year, I have had no problems with students running Mac's. Gurobi has gotten much better with their apple support, but Anaconda in Mac and integrating it with Gurobi is a little iffy. That being said, last semester students were successful with Macs *(WARNING: You will be on your own to get your stack configured. I have never used a Mac in my life and am too old to start).* At the end of the semester you Mac users can tell me if this is still a valid statement. If you have a Chrome book running Linux, good luck. You are on your own. I have no idea if it even works.

Software:

You will need to download and configure software for this course. Before downloading and installing all of the software listed in this syllabus, scan the **Gurobi Quick Start Guide** to figure out the order and sequence of software to be loaded. ***If you do not reference the quick start guide for your OS, you will undoubtedly run into issues. Please read it first.***

NOTE: If you are using Windows 8.0 or higher as your OS, you MUST use the 64 bit versions of the software. You don't get a choice. Well, actually you do, but it won't work.

Gurobi – We will use Gurobi as our optimization library. I recommend you use the latest version of Gurobi (currently 7.5). If you have some other version loaded, it will be well worth it to upgrade your license to 7.5. There is no reason not to upgrade since it is free for academic users (you are an academic user if you are taking this course). You will need to register on the Gurobi website to download the software. Go to www.gurobi.com to register and download your academic license.

Before you can use Gurobi, the software must validate your academic standing. It does so by checking your IP address and ensuring it is a legitimate academic domain. Once validated, your license is good for a year. You can validate your license in one of two ways:

1) *Easy but inconvenient:* Go to the GMU campus, log onto the network as a student, and then follow the directions in the Gurobi Quick Start Guide for validating Gurobi. This should be your preferred method for getting your laptop version of Gurobi registered.
2) *Convenient but potentially frustrating:*
    a. Download and install Cisco AnyConnect for your computer operating system (https://itservices.gmu.edu/downloads/index.cfm), once you have installed it, you will be able to skip this portion of the instructions.
    b. Start Cisco AnyConnect
    c. Once the connection is made, open an internet browser to any page to confirm your connection is valid, and navigate to the Gurobi website and follow the links to the Quick Start Guide for your OS.
    d. Follow the instructions in the **Gurobi Quick Start Guide** for academic validation.

Python –Gurobi supports the Anaconda distribution of Python 2.7 and 3.5 (does not support the 3.6 release). This statement is true for both Windows and OSX. Your instructor only supports (that means provides assistance in) Python 2.7 and 3.5. I strongly encourage you to use the Anaconda distribution for 2.7. Gurobi also works with other IDE's (PyCharm, PyScripter, …), but it will not be guaranteed to work. Go to Continuum Analytics download page (https://www.continuum.io/downloads) to get the correct version of Python (did I mention you want 2.7?). Follow the instructions to the letter on getting Gurobi to run in Anaconda. Most

students who come to class without having this working ended up thinking they were smarter than the instructions and did it their own way.

Not that it matters (actually, it does for you), but I like Python 2.7 over 3.5 because the number of libraries available to 2.7 is vastly greater than the number of libraries available under 3.5. Python 3.X series underwent major revisions from 2.X that caused the two versions to not be compatible (e.g., the print command is not compatible across the two languages). Developers of all Python packages had to (still have to) update their libraries. I suspect in about 5 more years, I'll change the syllabus to state that 3.5 is the preferred language, but until then, its 2.7 or suffer the consequences. For what it's worth, load Anaconda with 2.7 and then build a 3.5 environment that you can switch to. If you don't understand what I just said – then do 2.7 and leave it at that.

SQLite Studio – This is an optional piece of software. If you are inexperienced with Structured Query Language (SQL) you may find this a useful tool. It is a graphical user interface and database management tool for SQLite (the database library that ships with Python). By using this software you will be able to test your SQL commands and statements before you run them in your Python routines. To download a copy of SQLite Studio go to http://sqlitestudio.pl

You will use SQLite in this course. Python interfaces very nicely with PostgreSQL, MySQL, and MS SQL Server. That being said, I only have SQL Server and SQLite loaded on my computer. The configuration requirements for SQL Server are specific to each computer and environment. The configurations for SQLite are common across all systems since it is included with Python. If you would like your homework to be successfully graded, you will use SQLite.