Optimisation and Operations Research

Lecture 1: Introduction and Course Summary

Matthew Roughan
<matthew.roughan@adelaide.edu.au>
http://www.maths.adelaide.edu.au/matthew.roughan/notes/OORII/

School of Mathematical Sciences,
University of Adelaide

August 13, 2019

Section 1

Course Introduction

Who is teaching this course

Course Coordinator: Prof Matthew Roughan
Email: matthew.roughan@adelaide.edu.au
Office: Ingkarni Wardli, Room 6.17

Course Tutors: Sarah James and Max Wurm

Administrative Enquiries: School of Mathematical Sciences Office,
Level 6, Ingkarni Wardli

Course philosophy

You can see the Uni's graduate attributes on the web.
I think more in terms of the AMSI Industry working groups findings regarding what they NEED in applied mathematicians:

- logical/critical thinking
- practical ability to do problem structuring and solving, often beginning with messy data, or a messy problem description
- data analysis skills, ability to work effectively with data
- ability to work collaboratively, in multi-disciplinary teams
- communication skills
- ability to code/program

These are the things I am aiming to teach in this course
+ some optimisation of course

Where do you need to be, and when?

Lectures

The plan is to record lectures, but I make no guarantee to record every lecture. Sometimes machinery doesn't work properly.

Practicals

One session every week will be in a lab. In alternate weeks, it will a normal practical, or be focused on projects.

It is unlikely to be possible/useful to record these. Students should aim to attend in person. If unable to attend, students MUST go through the material/exercises by the due date (and submit them through Cody).

Tutorials

Students are required to attend one session every two weeks.

They take place fortnightly starting in Week 2.

You must attend the tutorial you have enrolled in via MyUni, exceptions to be agreed with the lecturer on a one off basis.

Tutorials will provide practise exam questions. You are expected to have at least attempted the problems before attending your tutorial.

Schedules and Outline

Note that these are still being tidied up.

Recommended resources

1. Taha. An introduction to operations research. Prentice Hall, 2007.
2. Nemhauser and Wolsey. Integer and combinatorial optimisation. Wiley, 1988.
3. Papadimitriou and Steiglitz. Combinatorial optimization – Algorithms and complexity. Prentice Hall, 1982.
4. Ahuja, Magnanti and Orlin. Network flows: theory, algorithms, and applications. Prentice Hall, 1993.
5. Wright. Primal-dual interior point methods. SIAM, 1997.

Lecture notes

Almost all materials for this course will be available via MyUni, or through the course web page at http://www.maths.adelaide.edu.au/matthew.roughan/notes/OORII/index.html.

- notes are still being tidied up, and final versions of some tasks added
- lecture recordings should be on MyUni as usual

The lecture slides cover all examinable material in this course.

Lecture notes are based on notes originally prepared by Liz Cousins and edited by Josh Ross and David Green, as well as myself.

Assessment

1. 10% for assignments (6 in total)
2. 20% for the project
3. 70% for the final exam

with the following additional requirements:

- An aggregate score of at least 50%
- A score of at least 40% on the Project.

Note also that

1. All written assignments are to be submitted via the appropriate handin box on level 6 of Ingkarni Wardli (in front of lifts), with a signed cover sheet attached.
2. Late assignments will not be accepted, except by prior arrangement (for a good reason).

The Project

The Project is a group project based upon a real world problem encountered in South Australia. This type of problem actually exists in many places, but an actual case study from S.A. will form the basis of the problem description.

More details will be provided in Week 2.

You will work in teams of 4-5.

You will need to choose a group from your practical session.

Please start to to think about your groups now! Groups WILL be finalised in the practical session of week 2. If you don't form a group, then you will be allocated to one.

Policy on plagiarism

Students are advised that the University of Adelaide has an official policy on academic honesty and plagiarism, which they are advised to read at
http://www.adelaide.edu.au/policies/230/

Students are encouraged to work together in order to enhance their understanding of the subject matter. To this end they may work together on assignments.

However, students are required to have, and may be required to demonstrate, a complete understanding of their submitted assignments. Failure to demonstrate a complete understanding of a submitted assignment may be interpreted as evidence of plagiarism.

Hint: One easy way to work together and yet avoid any issues with plagiarism is to plan how to do the assignment together, but then work on the assignment itself separately. This way, you are able to learn together, but are **not** tempted to copy each other's actual assignment.

Programming

- This course will require some computer programming
  - expectation is that you can write simple Matlab programs
  - help getting started
    - notes www.maths.adelaide.edu.au/matthew.roughan/Lecture_notes/matlab_notes.pdf
    - links
      http://au.mathworks.com/academia/student_center/tutorials/launchpad.html
      http://people.duke.edu/~hpgavin/matlab.html
- You must write your own code!
  - plagiarism applies to writing code as well as other handins

What are we trying to teach

- Translate real-world problem into maths
  - may involve some approximation
  - have to deal with tradeoff between accuracy and simplicity
- Linear (integer) programming
  - how to find optimum solutions
- Algorithms
  - how particular algorithms work
  - general algorithmic strategies
  - what is important in design and implementation
    - complexity
  - danger points
- Proofs
  - how can you show you definitely have the right answer?

All in the context of linear algebra and optimisation.

Section 2

Introduction to Optimisation and Operations Research

What is OOR?

Optimisation: The action or process of making the best of something; (also) the action or process of rendering optimal; the state or condition of being optimal. (Oxford English Dictionary)

Operations research: Operations research, also known as operational research, is an interdisciplinary branch of applied mathematics and formal science that uses advanced analytical methods such as mathematical modelling, statistical analysis, and mathematical optimization to arrive at optimal or near-optimal solutions to complex decision-making problems. (Wikipedia 28/04/10)

Some real world problems

Foam cutting:

A company makes various grades of foam rubber, for use in car seats, beds, etc. The foam is made in "buns", approximately 3.5m $\times$ 3.5m $\times$1m. From these, various smaller rectangular prisms are cut to satisfy orders that vary from client to client, and day by day.

Cuts are made by a guillotine, and so cuts go right through the block being cut. This can lead to wastage: some blocks (off-cuts) are too small or the wrong shape to be used in any order.

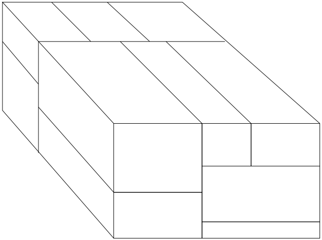

http://www.trc.adelaide.edu.au/aor/case_studies/dunlop.html (5/07/11).

Some real world problems

Foam cutting:

The firm wishes to minimise the cost of wastage.

This is often referred to as a *bin-packing* problem, and can arise in cargo handling as well as other places.

"Adelaide Operations Research undertook research into various algorithms that could be applied to the optimal foam cutting problem. Using historical data, and keeping track of off-cuts from each order fulfilled, we determined that the wasted material over a three month period could be reduced by approximately one-third – a saving of potentially $800,000 per year in raw materials."

http://www.trc.adelaide.edu.au/aor/case_studies/dunlop.html (5/07/11).

Some real world problems

Reserve design:

A problem facing conservationists and government agencies is which parcels of land (reserves) should be put aside as national parks, etc., in order to preserve as many species of plants and animals as possible. For genetic diversity and greater chance of survival of species, there must be multiple reserves containing each species. Contiguous areas might be preferable to widely spaced, smaller sites.

However, the purchasing of land, and then its maintenance and management, can be prohibitively expensive.

So the question is: which parcels of land should be reserved for the protection of our wildlife, to minimise costs (land acquisition or total), while maintaining sufficient representation of each species?

Some real world problems

Dial-a-ride problem:

Taken from

Cordeau (2006)
A Branch-and-Cut Algorithm for the Dial-a-Ride Problem,
Operations Research, 54 (3), 573-586.

"In the dial-a-ride problem, users formulate requests for transportation from a specific origin to a specific destination. Transportation is carried out by vehicles providing a shared service. The problem consists of designing a set of minimum-cost vehicle routes satisfying capacity, duration, time window, pairing, precedence, and ride-time constraints."

Some real world problems

Automotive manufacturing:

"Major automobile manufacturers often produce a wide range of vehicles on a single assembly line. The time taken to fit various options depends on their complexity and the type of car. There are rules concerning how closely two cars requiring certain options may be spaced on the assembly line. As a result, determining the number of cars that can be built on any given day, given the types of cars that have been ordered, can be difficult to determine. If assembly line resources are not used efficiently the maximum capacity may not be achieved, with subsequent impact on the potential revenue that can be generated. It is crucial to know when capacity limitations may occur so that actions can be taken to address the situation. "

http://www.trc.adelaide.edu.au/aor/case_studies/auto.html (5/07/11).

Some real world problems

Automotive manufacturing:

"We developed software that analyses the end-to-end manufacturing process of one of the plants of a major automobile manufacturer. The software uses vehicle orders, assembly line rules and production information to determine when artificial limitations in capacity might occur. When capacity issues are identified, the software can automatically determine which rules are causing the pressure on capacity. Knowing this, production managers can then take action to loosen certain rules, for example by using extra staff at key assembly line stations, in order to increase overall capacity and thus ensure the most efficient operation of the plant overall. "

http://www.trc.adelaide.edu.au/aor/case_studies/auto.html (5/07/11).

Section 3

Our First Problem

Think of an optimisation problem...

- What is it that you ideally wish to achieve?
  - *e.g.*, maximise profit
  - *e.g.*, minimise risk

This is the *Objective*

- What is it that you have control over?
  - *e.g.*, **who** should do W?
  - *e.g.*, **how** many X should I make?
  - *e.g.*, **where** should I make Y?
  - *e.g.*, **when** should I do Z?
  - *e.g.*, **what** model of V should I choose?

These are the *Variables*

- Are there any restrictions?
  - *e.g.*, we can't make negative numbers of items
  - *e.g.*, we have a limited resources or time

These are the *Constraints*

Our first problem

Example (A scheduling problem)

A manufacturing company makes three types of items: desks, chairs, and bed-frames, using metal and wood.

A single desk requires

- 2 hours of labour,
- 1 unit of metal and
- 3 units of wood.

One chair requires

- 1 hour of labour,
- 1 unit of metal and
- 3 units of wood.

Making one bed frame requires

- 2 hours of labour,
- 1 unit of metal and
- 4 units of wood.

Our first problem

Example (A scheduling problem: (cont))

In a given time period, there are

- 225 hours of labour available,
- 117 units of metal, and
- 420 units of wood.

The profit on

- one desk is \$13,
- one chair is \$12, and
- one bed frame is \$17.

Our first problem

Example (A scheduling problem: (cont))

The company wants a manufacturing schedule designed, which maximises profits without violating any constraint on resource availability during that time period.

So what are the steps in Formulation?

1. Read the question right through! Look carefully.
2. Tabulate any data
3. Identify the variables
4. Identify the objective
5. Formulate the constraints, *i.e.*, define the feasible region
6. Write down the LP (Linear Program)

Our first problem

2. Tabulate the *data:*

	Labour (hrs)	Metal	Wood	Profit
Desk	2	1	3	13
Chair	1	1	3	12
Bedframe	2	1	4	17
Total	225	117	420	

Our first problem

3. Define the *decision variables:*

   x_1 to be the number of desks;

   x_2 to be the number of chairs; and,

   x_3 to be the number of bedframes,

made per time period.

Our first problem

4 Identify (formulate) the *objective function:*

Here we wish to maximise the profit, z, (in dollars), so we have:

$$\max \; z = 13x_1 + 12x_2 + 17x_3.$$

We might sometimes write this as

$$\underset{\mathbf{x}}{\operatorname{argmax}} \; z = 13x_1 + 12x_2 + 17x_3.$$

which means find the *argument* $\mathbf{x}$ that maximises z.

Our first problem

5. Formulate the *constraints*, including any non-negativity constraints

$$
\begin{array}{ccccccccl}
2x_1 & + & x_2 & + & 2x_3 & \leq & 225 & & \text{(labour)} \\
x_1 & + & x_2 & + & x_3 & \leq & 117 & & \text{(metal)} \\
3x_1 & + & 3x_2 & + & 4x_3 & \leq & 420 & & \text{(wood)} \\
x_1 \geq 0, & & x_2 \geq 0, & & x_3 \geq 0 & & & & \text{(non-negativity)}
\end{array}
$$

Question:

Should we include any other constraints on the x_i's?

Our first problem

5. Formulate the *constraints*, including any non-negativity constraints

$$
\begin{array}{ccccccccc}
2x_1 & + & x_2 & + & 2x_3 & \leq & 225 & & \text{(labour)} \\
x_1 & + & x_2 & + & x_3 & \leq & 117 & & \text{(metal)} \\
3x_1 & + & 3x_2 & + & 4x_3 & \leq & 420 & & \text{(wood)} \\
x_1 \geq 0, & & x_2 \geq 0, & & x_3 \geq 0 & & & & \text{(non-negativity)}
\end{array}
$$

Question:

Should we include any other constraints on the x_i's?

What about integer constraints?

Feasible region

$$
\begin{array}{ccccccccc}
2x_1 & + & x_2 & + & 2x_3 & \leq & 225 & & \text{(labour)} \\
x_1 & + & x_2 & + & x_3 & \leq & 117 & & \text{(metal)} \\
3x_1 & + & 3x_2 & + & 4x_3 & \leq & 420 & & \text{(wood)} \\
x_1 \geq 0, & & x_2 \geq 0, & & x_3 \geq 0 & & & & \text{(non-negativity)}
\end{array}
$$

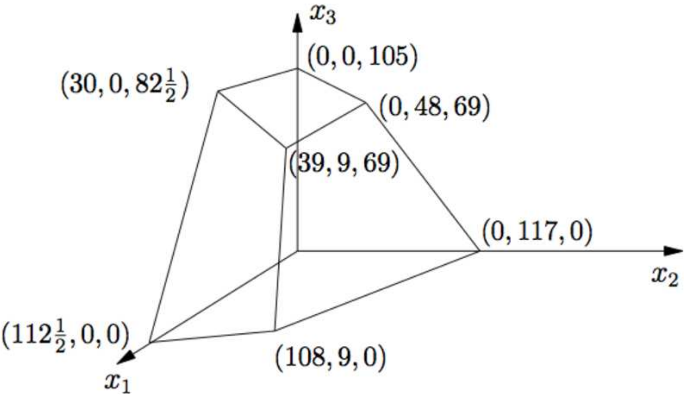

Our first problem

6. Write down the *linear program* (LP)

Our LP:

(objective function)

$$\max \quad z = 13x_1 + 12x_2 + 17x_3$$

subject to

$$\begin{array}{rcrcrcl} 2x_1 & + & x_2 & + & 2x_3 & \leq & 225 \\ x_1 & + & x_2 & + & x_3 & \leq & 117 \\ 3x_1 & + & 3x_2 & + & 4x_3 & \leq & 420 \end{array}$$

(non-negativity)

$$x_1 \geq 0, \qquad x_2 \geq 0, \qquad x_3 \geq 0$$

Note:
If the integer constraints were added, it would be an Integer Linear Program (ILP).

Our first problem

Matrix form of the problem:

$$\begin{array}{rrcl} \max & z & = & \mathbf{c}^T \mathbf{x} + z_0 \\ \text{such that} & A\mathbf{x} & \leq & \mathbf{b} \\ \text{and} & \mathbf{x} & \geq & 0 \end{array}$$

where vector inequalities mean every element of the vector must satisfy the inequality.

Takeaways

- Course Admin
- Optimisation problems have three parts
  - variables
  - objective
  - constraints
- We will deal with LPs (Linear Programs)
  - objectives and constraints are linear

Further reading I