

1 Course description

Mathematical optimisation is one of the cornerstones of fields such as Machine Learning, Artificial Intelligence, and Operations Research. Most decision support methods have, at some level, a mathematical optimisation technique at its core, and it is precisely these methods that we will learn in this course.

Linear Optimization is a powerful framework in which one seeks to represent systems by means of linear objective functions and constraints. Using the analogy that variables represent decisions or parameters to be defined, constraints represent rules that a valid configuration or a plan of action for the system, and the function is a measure of performance, one can use that framework to support decision making in a wide range of applications, from planning industrial chemical plants to training models that learn from data.

In this course, the students will learn the basics of linear optimisation theory as well as advanced algorithms available and how they can be applied to solve challenging real-world inspired optimisation problems. Throughout the course, we will also look into practical and research applications of linear optimisation.

2 Learning outcomes

Upon completing this course, the student should be able to

- understand how several important problems arising from diverse fields can be cast and solved as linear optimisation problems;
- know the main techniques for solving linear optimisation problems and how to apply them in practice;
- know how to use optimisation software for implementing and solving linear optimisation problems.

3 Teaching methods

This course uses a combination of lectures and exercise sessions that introduce the material. The assessment is continuous throughout the course by means of homework assignments. This means that the student's learning depends on their initiative in attending the lectures and exercise sessions, formulating questions and discussing with colleagues.

The course will be taught by a composition of the following methods:

- lectures;
- exercise sessions;
- guided self-study;
- theoretical and practical exercises;
- homework assignments and feedback.

All lectures will be delivered in person. Voice-over recordings will be provided after the lectures as well. As preparation for the lectures, the students will be requested to read through the lecture notes provided and formulate any questions to be answered during the lectures.

The exercise sessions will take place as in-person sessions. These sessions will consist of practical exercises that will be done in smaller groups. These will focus on illustrating and developing concepts using computational tools and provide support for doing the homework assignments. The teaching assistants are also available to address questions related to content and homework assignments.

4 Assessment

The final grade of the course is composed of two components:

H: 5 homework assignments, each worth 20 points;

P: Participation component, worth 10 extra points;

The conversion scale to the 1-5 scale is as follows.

1-5	0-100
Fail	0-50
1	51-60
2	61-70
3	71-80
4	81-90
5	91-100

Table 1: Conversion from 0-100 to 1-5 scale

4.1 Homework assignments

A total of 5 homework assignments will be handed out. Each homework is worth 20 points, adding to a total of 100 points. The homework for each fortnight (fortnight = two weeks) will be available on Monday from MyCourses and will have as the deadline the Wednesday evening two or three weeks after the release. The submission of the solutions must be made through the course MyCourses website. Homework submissions after the deadlines will be penalised by 5 points per day after the deadline.

The homework will be composed of theoretical and computational exercises. The computational skills required to solve the exercises will be introduced in the exercise sessions, but it is expected that the students learn and practise the language on their own. Supporting material for that will be provided. The programming language that will be used in this course is Julia (julialang.org).

4.2 Participation bonus

Students with attendance at lectures and exercise sessions greater or equal to 10 (out of 11 sessions each) earn 10 extra points. Otherwise, points will be given proportionally to attendance.

5 Course material

Main study material: lecture notes, exercises, homework assignments.

The lecture notes are available at github.com/gamma-opt/optimisation-notes. These notes are open source and may be updated as needed. Please make sure to use an up-to-date version.

6 Course schedule

A tentative schedule for the course is given. The content of each class may be adapted according to the pace of the classes.

Week	Lecture	Content
2	1	Introduction - modelling linear problems
3	2	Linear algebra, polyhedral sets
4	3	Basis, extreme points and optimality of LP
5	4	The simplex method
6	5	Duality I - the dual simplex
7	6	Duality II - other applications
8	–	<i>Break between Periods III and IV</i>
9	7	Decomposition methods: Benders decomposition
10	8	Mixed-integer programming formulations
11	9	Branch-and-bound
12	10	Cutting planes method
13	11	Branch-and-cut: additional methods
14	–	<i>Easter holiday</i>
15	12	Computational and practical points

Table 2: Schedule of classes