Optimisation algorithms

The aim of this course is to introduce different optimisation algorithms that can be used to optimise any partial differential equation. No particular application is targeted. Instead, the steps of each algorithm will be discussed for the optimisation of generic functions. Nevertheless, the methods presented in this course are state of the art in many aeronautical and mechanical engineering applications.

The following topics are covered:

- An introduction of the lingo and the terminology used in the community: how to formalise an optimisation problem and how to extract exact and approximate gradients of the underlying PDE.
- Descent-based optimisation algorithm: 1st order algorithms, line-search algorithms and second order methods.
- Constrained optimisation: introductions of the Lagrangian and the Lagrange multiplier (adjoint equations)
- Derivative-free methods: Population and stochastic methods.
- Strategies for reducing the cost of function evaluations, with application to stochastic-based methods.

The participants are meant to implement and compare different algorithms using python codes and notebooks.

Prerequisite: Skills in applied math (linear algebra, taylor-expansions, series, integration by parts, …) and scientific computing (python programming).

Bibliography:

- M.J. Kochenderfer & T.A. Wheeler. Algorithms for optimisation. The MIT Press. 2019
- M.D. Gunzburger. Perspectives in flow control and optimisation. SIAM 2003.

Timing: The Course is offered in the first part of the M2 year.

Credits: 1.5 ECTS

Hours: 18 hours.