Data Driven Approaches in (Multi-objective) Bayesian Optimisation - An Introductory Course

Tinkle Chugh

Department of Computer Science
University of Exeter, UK

Course page: `https://tichugh.github.io/essai-bo-course/`

In this course

You will learn

- Why multi-objective optimisation
- What are existing multi-objective optimisation methods and their shortcomings
- Introduction to Bayesian machine learning - Gaussian Processes
- Why Bayesian (multi-objective) optimisation
- How to utilise Bayesian inference to alleviate computation time
- Examples of real-world optimisation problems

Course Structure

Day 1

- Introduction to multi-objective optimisation
  - Definition of multi-objective optimisation problems (MOPs)
  - Pareto dominance and Pareto optimality

Course Structure

Day 2

- Foundation of Gaussian Processes
  - Hyperparameter Optimisation
  - Uncertainty Quantification

Course Structure

Day 3

- Bayesian Optimisation
  - Exploitation vs Exploration
  - Acquisition functions

Course Structure

Day 4

- Multi-objective Bayesian Optimisation
    - Mono- and multi-surrogate approaches
    - Summary

Any questions on the module structure?

Examples of Optimisation

Engineering Design[1]

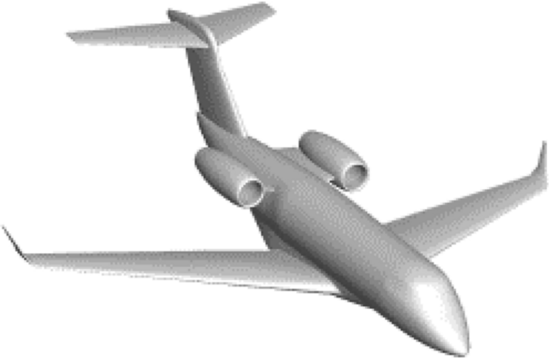

[1] Forrester, A. I. J., Sobester, A., & Keane, A. J. (2008). Engineering Design via Surrogate Modelling: A Practical Guide. Wiley

Hyperparameter tuning in Neural Networks[2]

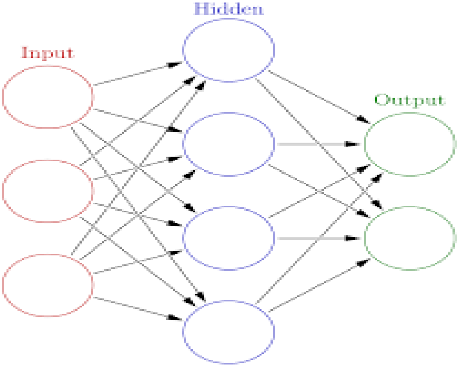

[2] Wu et al. Hyperparameter Optimization for Machine Learning Models Based on Bayesian Optimization, Journal of Electronic Science and Technology, 17, 2019, 26-40.

Examples of Optimisation

Drug Discovery[3]

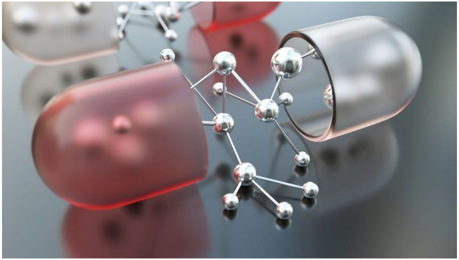

[3] Colliandre L, Muller C. Bayesian Optimization in Drug Discovery. Methods Mol Biol. 2024;2716:101-136.

Examples of Optimisation

- Material Science[4]
- Financial Modelling and Portfolio Optimisation[5]
- Reinforcement learning[6]

[4]Xue, D., Balachandran, P., Hogden, J. et al. Accelerated search for materials with targeted properties by adaptive design. Nat Commun 7, 11241 (2016).

[5]Gupta, A. A Bayesian Approach to Financial Model Calibration, Uncertainty Measures and Optimal Hedging. Oxford University, UK, 2010.

[6]Paul et al. , JMLR, Robust Reinforcement Learning with Bayesian Optimisation and Quadrature, JMLR, 21, 1-31, 2020.

Terminology

Decision variables or parameters

x is a decision vector (design), $\mathrm{x} = (x_1, x_2, \ldots, x_D)^\top$

$\mathcal{X}$ is the feasible decision space

Here, without loss of generality, we will be illustrating on continuous space, $\mathcal{X} \subset \Re^D$

Objective function

Using the objective function (or fitness function), we assess the quality of a decision vector, $y = f(\mathrm{x})$

We will be assuming objective function(s) to be minimised

Optimisation Problem

Single-objective formulation

$$\mathrm{x}^* = \underset{\mathrm{x} \in \mathcal{X}}{\operatorname{argmin}} f(\mathrm{x})$$

Solving Single Objective Optimisation Task

We can put an order on $\mathcal{X}$:

- if $f(\mathrm{x}) < f(\mathrm{x}')$, then x is better than x'
- if $f(\mathrm{x}) > f(\mathrm{x}')$, then x is worse than x'
- if $f(\mathrm{x}) = f(\mathrm{x}')$, then x is equivalent to x'

This may seem trivial; the situation is not as clean-cut in the multi-objective setting

Examples Single Objective Optimisation Task

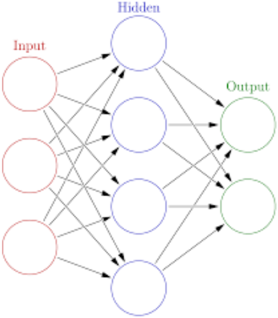

Training a neural network requires optimising a loss function:
$L(\mathrm{y}, \hat{\mathrm{y}}) = \frac{1}{N} \sum_{i=1}^{N} (y_i - \hat{y}_i)^2$

Examples Single Objective Optimisation Task

Travelling Salesman Problem: Given a list of cities and the distances between each pair of cities, what is the shortest possible route that visits each city exactly once and returns to the origin city?

Examples Single Objective Optimisation Task

Benchmark problems - test of algorithms

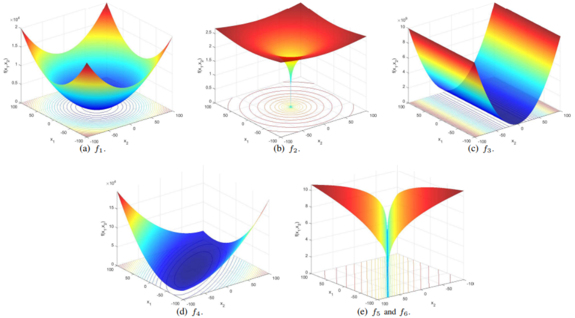

Source; GNBG-Generated Test Suite for Box-Constrained Numerical Global Optimization `https://arxiv.org/pdf/2312.07034`

Examples Single Objective Optimisation Task

Benchmark problems - test of algorithms

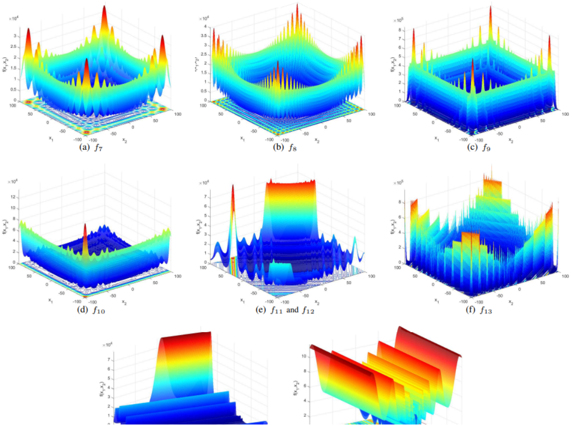

Source; GNBG-Generated Test Suite for Box-Constrained Numerical Global Optimization `https://arxiv.org/pdf/2312.07034`

Multi-objective Optimisation

Multiple objective functions:

$$\text{minimise} \quad \mathrm{f}(\mathrm{x}) = (f_1(\mathrm{x}), \ldots, f_M(\mathrm{x})) \qquad \text{subject to} \quad \mathrm{x} \in \mathcal{X}$$

Example

- Minimising manufacturing cost and maximising the quality of product
- Maximising the fuel efficiency and minimising the soot emission and minimising NOx

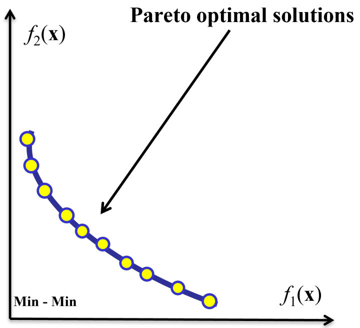

A Toy Example

Multi-objective Optimisation

- In multi-objective optimisation, we have $M > 1$ objectives
- The **objective vector** y is returned as $\mathrm{y} = \mathrm{f}(\mathrm{x})$
- A putative solution is the tuple of a valid design and its evaluation: $\mathrm{s} = (\mathrm{y}, \mathrm{x})$

Optimisation Problem

Multi-objective formulation

$$\mathrm{x}^* = \operatorname*{argmin}_{\mathrm{x} \in \mathcal{X}} (f_1(\mathrm{x}), f_2(\mathrm{x}), \ldots, f_M(\mathrm{x}))$$

Examples Multi-objective Optimisation

Bias and Variance Trade off

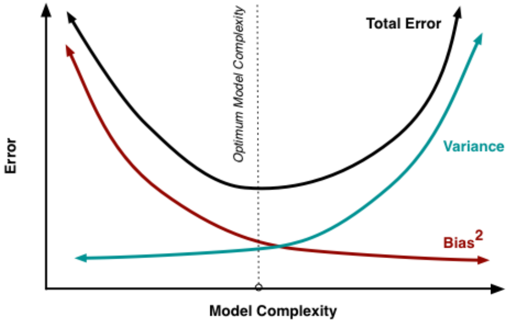

$$\text{Bias: } E(\hat{y}) - y$$
$$\text{Variance: } E[(\hat{y} - E[\hat{y}])^2]$$

Chemical Engineering: Synthesis gas production

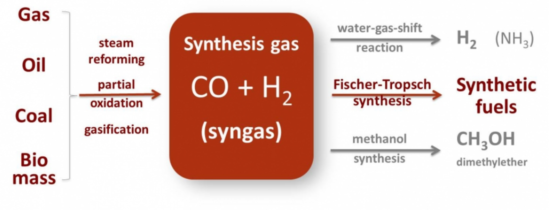

Objectives: 1. Maximise CH_4 rate 2. Maximise CO rate, and 3. Minimise H to CO ratio

Examples Multi-objective Optimisation

Pump shape optimisation

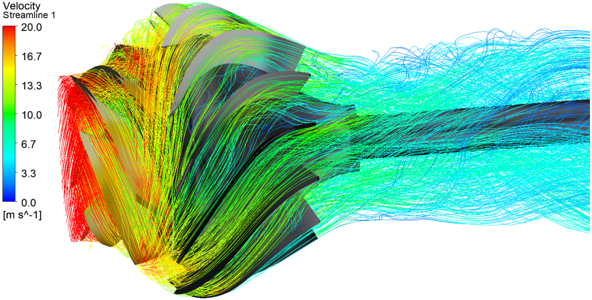

Objectives: Maximising efficiencies at different flow rates

Black Box Optimisation Problem

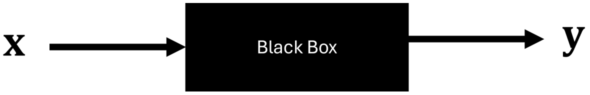

No analytical expression for objective functions are available

Comparing Solutions

a dominates b

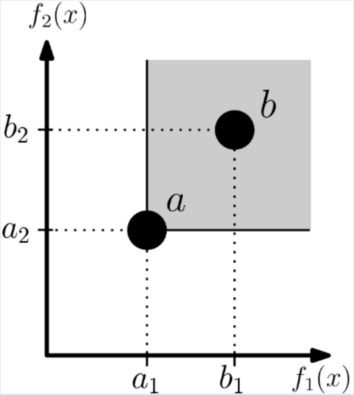

a and b nondominated

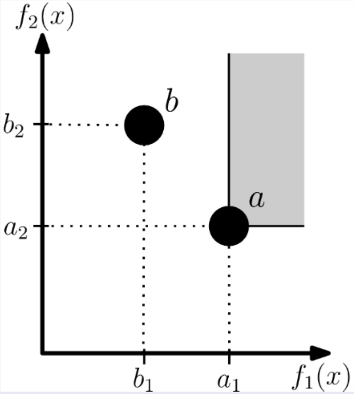

Pareto Optimality

Vilfredo Pareto

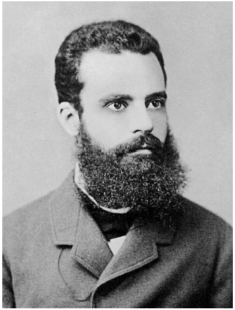

Unknown photo authorship (pre 1926)

Italian Polymath, Engineer and Academic

A decision vector $\mathbf{x}$ is Pareto optimal iff $\nexists \mathbf{x}' \in \mathcal{X}$ where $\mathbf{x}' \prec \mathbf{x}$

$\mathbf{x}$ is therefore optimal in the sense that it is impossible to find an alternative solution that improves performance in any objective, without having to also degrade performance in at least one other objective

Pareto Set and Pareto Front

From Pareto optimality we can define the optimal set of of solutions to the multi-objective optimisation task as:
$\mathcal{P} = \{\mathbf{x} \in \mathcal{X}]|\nexists \mathbf{x}' \in \mathcal{X}, \mathbf{x}' \prec \mathbf{x}\}$

Note it is possible that the Pareto Set $\mathcal{P}$ for a problem to have infinite membership

The image of $\mathcal{P}$ under $\mathbf{f}$ is known as the Pareto front, $\mathcal{F}$

Pareto Set and Pareto Front

$$
\begin{aligned}
\text{Minimize} \quad & f_1(x_1, x_2) = x_1^2 + x_2^2 \\
\text{Minimize} \quad & f_2(x_1, x_2) = (x_1 - 1)^2 + x_2^2 \\
\text{Subject to} \quad & 0 \leq x_1 \leq 1 \\
& x_2 = 0
\end{aligned}
$$

A Toy Example

With NSGA-II

Summary So Far

- In multi-objective optimisation, we seek to optimise more than one objective simultaneously
- We are concerned with an approximation to the Pareto front and Pareto set - to define our trade-off solutions