THE UNIVERSITY OF SYDNEY

# Statistical Optimisation under Uncertainty

*Gilad Francis*
*Fabio Ramos*
School of Information Technologies
FACULTY OF ENGINEERING & INFORMATION TECHNOLOGIES

## Overview

### Black-Box Optimisation

Many engineering problems revolve around finding an optimal set point. In most cases, it can be expressed as problem of finding the minimum or maximum of an unknown and complex function.

Optimisation is an expensive process that requires many experiments in order to assess the underlying function and obtain the best outcome.

We propose a strategy that will guide the optimisation process, i.e. finding the extreme of a black-box function with a set of unknown constraints. We present a possible application for such a strategy in robotic autonomous exploration.

## Contributions

### Active Bayesian Sampling

Develop Bayesian Optimisation (BO) framework that will generate new observation points with a set of unknown constraints while balancing the costs (energy, time, etc.).

## Bayesian Optimisation (BO)

### What is BO?

Bayesian Optimisation (BO) is an effective framework for finding the global extreme of a black-box noisy objective function, :

As in many machine learning problems, the only information about is obtained by observing its values in specific locations. However, each such observation carries a cost to system resources. Hence, it is desirable to minimize number of observation. BO is designed exactly for that purpose, find global extreme by selecting appropriate observation points [1].

The Bayesian Optimisation algorithm has 2 steps.

- **Objective function evaluation:** Based on the current set of observations. Forms a surrogate function that holds the updated prediction of . This is done efficiently with a Gaussian Process (GP) prior.
- **Acquisition function - Search for the extreme**: This process is guided by an acquisition function that estimates the utility of sampling in each location.

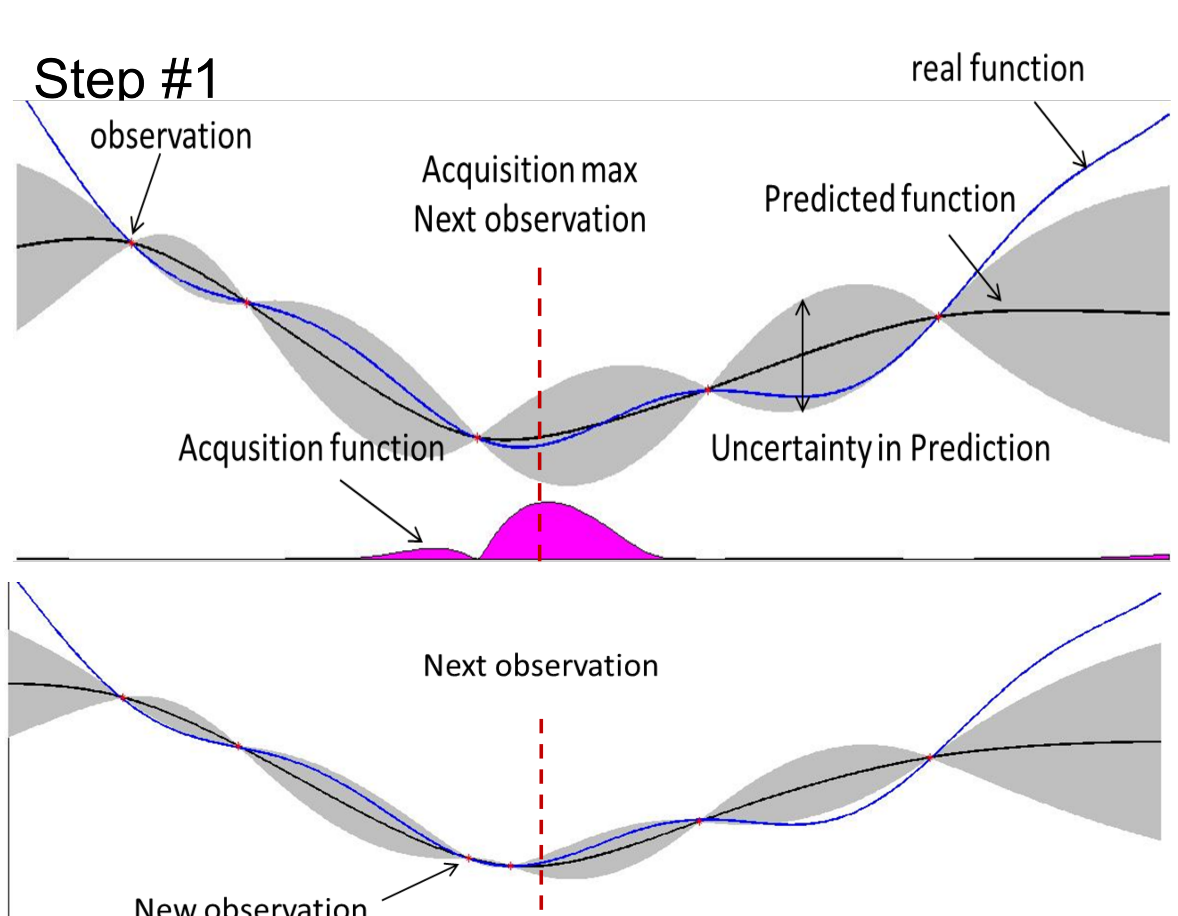


Figure 1: Example of Bayesian Optimisation

## Optimisation under Uncertainty

### Unknown Constraints

Optimising an unknown function with a known set of constraints is a complex problem. Although the constraints limit the range of the function, evaluating the function beyond this range might still hold valuable information. Therefore, the difficulty of optimising with constraints is the decision when it is worthwhile to observe the function beyond its range.

When the constraints are unknown a-priori, the problem becomes more complicated. To resolve the problem, the acquisition function is changed to encode the constraints into the BO's cost/reward. Since the constraints are still unknown, a GP classifier is used to calculate predictions of the constraints based on previous observations [2].

## Application – Robotic Exploration

### How to explore efficiently and safely

Most autonomous exploration algorithms define heuristics that only maximize the information gain. This is a sub-optimal solution, since it does not address the safety of the planned path, or any trade-offs with the associated travel costs.

While a robot travels in an unfamiliar environment, it generates a map of its surrounding using one or more on-board sensors. This map, known as an occupancy map, defines features and obstacles around the robot. There are many techniques to build such a map, however we use Gaussian Process Occupancy Map (GPOM) [3].

Figure 2 shows the simulation result of the first observation done by a robot placed in an unfamiliar environment. The robot uses Laser Range Finder (LRF) to identify obstacles. Using a GP prior, the robot build regression of its surroundings (left) which is then classified into free, occupied and "not sure" zones (right). Note: the robot has no prior knowledge of this structure hence the walls are shown only for reference.

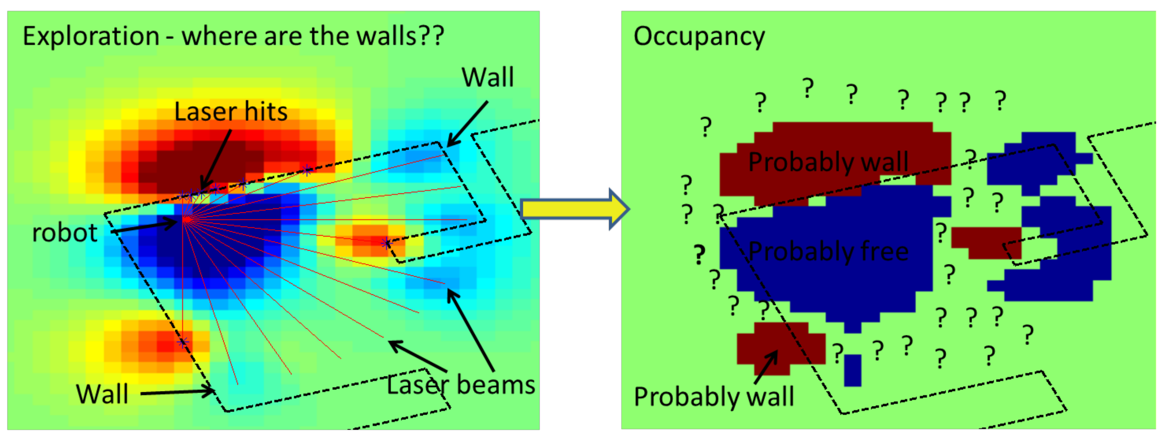


Figure 2: GP Occupancy Maps

To explore, the robot needs to move around the structure. Therefore the robot needs to plan and execute a path. Typically, the goal of each path is to maximize the information gain. However, the robot is limited in its movement by the walls which act as a set of unknown constraints.

Finding the global optimal path is an iterative process as shown in Figure 3. At each step, the Bayesian optimiser suggests a path based on the minimum of the acquisition function. The path validity is checked and if valid, the reward for the suggested path is evaluated using the currently available GP occupancy map.

## Robotic Exploration (cont'd)

The estimated reward/cost of the suggested path are fed back into the Bayesian optimiser to update the underlying GP. The process continues until it converges to the optimal path. The robot then executes the optimal path and new measured laser data update the GP occupancy map and the entire process restarts again.

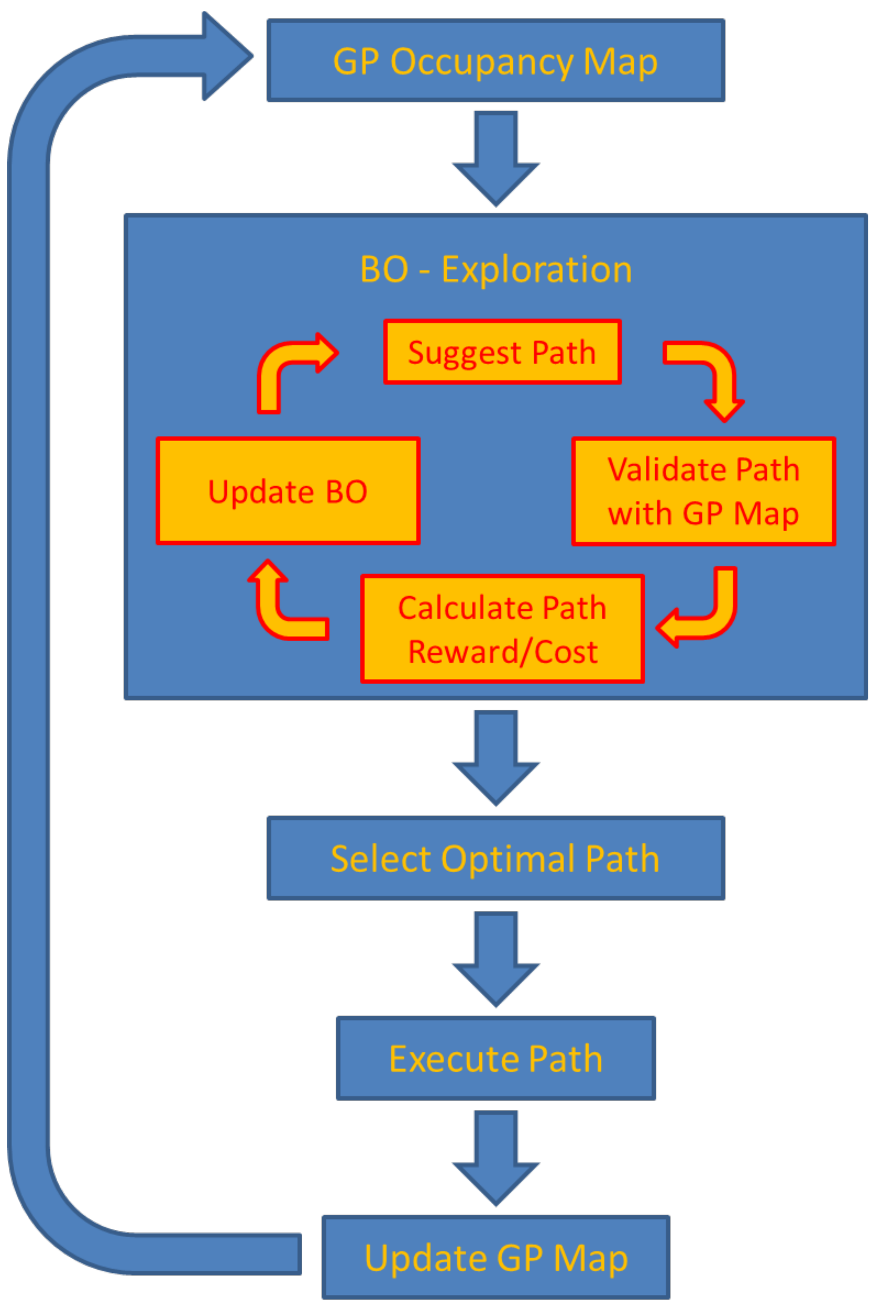


Figure 3: Path planning with BO

Figure 4 shows simulation results for two additional iterations. The automatic balance between exploration and safety is evident by the short and very local path on the first step ("looking around") and the more confident and longer path in step #3.

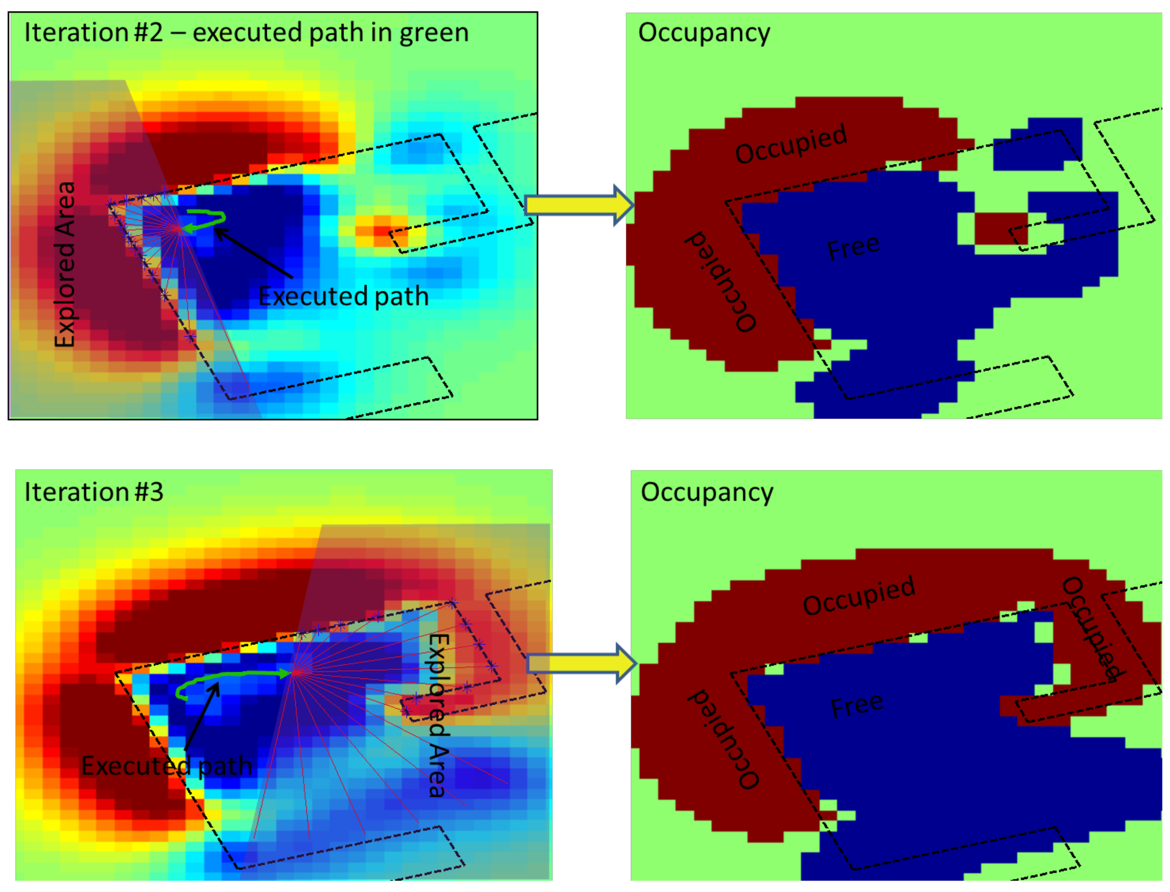


Figure 4: Path planning simulation results

## References

[1] Brochu, E., V.M. Cora, and N. De Freitas, *A tutorial on Bayesian Optimisation of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning*. arXiv preprint arXiv: 1012.2599, 2010.

[2] Gelbart, M.A., J. Snoek, and R.P. Adams, *Bayesian Optimization with Unknown Constraints*. arXiv preprint arXiv: 1403.5607, 2014.

[3] T O'Callaghan, S. and F.T. Ramos, *Gaussian process occupancy maps*. The International Journal of Robotics Research, 2012. 31(1): p. 42-62.