Project proposal
A/B testing with Bayesian Optimisation

Supervisors:
Juergen Branke (juergen.branke@wbs.ac.uk)

Background: A/B testing [1] is a randomised experiment to identify which of two alternatives is better. It is very popular in the domain of website tuning, because tests there are virtually free. For example, if one is not sure whether the "Buy now!" button on a website should be red and flashing, or just be green, simply randomly show users one of the two alternative designs and record over time which design is more effective. In case of more than two alternatives, Ranking & Selection techniques or Bayesian Optimisation [2] can be used to adapt the probability distribution for the different alternatives as more information becomes available. If we split the samples that we collect onto many alternatives, it may take a long time to collect enough information to identify the best, and so samples have to be spread over a substantial period of time. This creates problems if the behaviour of users changes over time, and samples collected are not really comparable.

Mini-project: The goal of the mini-project would be to explore A/B testing strategies for many alternatives over time. There are many possibilities, for example:

- Compare all simultaneously, but drop alternatives from the race if they are detected as inferior.
- Compare A/B on the first day, A/C on the second day, A/D on the third day etc, then some normalisation.
- Compare A/B on the first day, then subsequently the next alternative with the best so far.

At least two such strategies should be implemented and compared empirically under different assumptions on the number of alternatives, the noise, the change of the response behaviour over time etc.

Deliverable: Working code, paper with some empirical results on artificial benchmark problems (as those are quicker to compute)

PhD prospect: A/B testing is widely used in practice, so any improvement will likely have a huge practical impact. There are several possible ways to extend the project to a PhD project. Obviously testing more strategies, and trying to come up with a theoretically optimal (in some specified sense) strategy. Then we can assume that we not only collect responses, but also some information on the environment, or user characteristics. This information can then be used to learn not only which alternative is best overall, but which alternative is best for which user in which environment. If users come back, one may also try to optimise the sequence of alternatives shown to the user.

Student requirements: Programming skills.

References:

[1] R. Kohavi, S. Thomke. 2017. The Surprising Power of Online Experiments. Harvard Business Review
[2] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams and N. de Freitas, "Taking the Human Out of the Loop: A Review of Bayesian Optimization," in *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148-175, 2016