Adaptive Sequential Experimentation Techniques for A/B Testing and Model Tuning

Scott Clark
SigOpt
scott@sigopt.com

Abstract

We introduce Bayesian Global Optimization as an efficient way to optimize a system's parameters, when evaluating parameters is *time-consuming* or *expensive*. The adaptive sequential experimentation techniques described can be used to help tackle a myriad of problems including optimizing a system's click-through or conversion rate via online A/B testing, tuning parameters of a machine learning prediction method or expensive batch job, designing an engineering system or finding the optimal parameters of a real-world physical experiment.

We explore different tools available for performing these tasks, including Yelp's MOE and SigOpt. We will present the motivation, implementation, and background of these tools. Applications and examples from industry and best practices for using the techniques will be provided.

ACM Classification

G.3 [Probability and Statistics]: Experimental Design, G.1.6 [Optimization]: Global Optimization

Keywords

Design of Experiments, A/B testing, Hyperparmeter Tuning, Optimal Learning, Bayesian Global Optimization

Bio

Scott is currently running SigOpt, an optimization as a service startup that leverages techniques from optimal learning to automatically tune A/B tests, machine learning models, and physical experiments. Before that, he worked on the Ad Targeting team at Yelp Inc leading the charge on academic research and outreach with projects like the Yelp Dataset Challenge and open sourcing tools like MOE. Scott holds a PhD in Applied Mathematics and an MS in Computer Science from Cornell University and BS degrees in Mathematics, Physics, and Computational Physics from Oregon State University.

WWW 2015 Companion, May 18–22, 2015, Florence, Italy.
ACM 978-1-4503-3473-0/15/05.
http://dx.doi.org/10.1145/2740908.2743063