

## Large scale adaptive experiments on the Web

The Internet provides us with an infrastructure to measure the behaviors and preferences of large groups of people over long periods of time. However, this infrastructure cannot merely be used to obtain measurements: we can also use the Web for large-scale experiments. We can often actively manipulate the content that users see to measure the effects of these changes. Online experimentation is used routinely by Web companies, and the possibility to run large randomized (field) experiments using the Web is slowly changing the social and health sciences.

However, while the Internet provides an extremely effective infrastructure for large-scale (field) experiments, setting up such experiments is non-trivial. Often, technical knowledge is required to integrate treatment randomization into existing applications. This process is even more complicated when researchers aim to use not merely treatment randomization, but rather adaptive sequential experimental designs -- as popular in the medical sciences -- or personalized treatment selection schemes. In such cases the obtained measurements directly influence the subsequent treatment assignments. To implement such complex treatment selection schemes researchers often need to resort to custom-build applications that can be difficult to integrate into the existing practice of users.

To enable quick and easy implementation of complex treatment selection schemes in existing applications we present StreamingBandit: an open source (webserver) application that provides a REST API to facilitate complex treatment selections. StreamingBandit is built using Python Tornado and Redis to enable complex treatments selections at very large scales. Furthermore, StreamingBandit allows researchers to easily alter their treatment selection schemes: a random treatment selection scheme can be changed to an adaptive stopping rule without changing the actual field experiment. Using the HTTP REST protocol StreamingBandit integrates easily with most Web applications making it suitable for conducting complex experiments in within existing applications.

During the conference we will introduce the basic building blocks of StreamingBandit, and we will illustrate how to install and deploy the application. Furthermore, we will highlight how StreamingBandit has been used for a number of field experiments. First, we will detail how StreamingBandit was used to examine the effects of personalizing, based on historical responses, the persuasive messages used to motivate purchases in an e-commerce setting. Second, we will demonstrate how StreamingBandit was used to test multiple adaptive pricing schemes to iteratively optimize customer rebate offers. Finally, we will discuss a large-scale experimental study -- conducted on Amazon Mechanical Turk using Streaming Bandit for adaptive treatment assignment -- on the well-known decoy (or asymmetric dominance) effect. These diverse case studies highlight the versatility of the application.

StreamingBandit can be downloaded at:  
<https://github.com/MKaptein/streamingbandit>