KDD 2016 Tutorials

Building Recommender Systems using Photon ML

Xianxing Zhang / LinkedIn

Deepak Agarwal / LinkedIn

Bee-Chung Chen / LinkedIn

Paul Ogilvie / LinkedIn

Tutorials

Recommendation systems have become ubiquitous for web applications. Given significant heterogeneity in user preference, providing personalized recommendations is key to the success of such systems. To achieve this goal at scale, using machine learned models to estimate user preference from user feedback data is essential. Providing an easy-to-use and flexible machine learning library for practitioners to build personalization models is the key to productivity, agility, and developer happiness. In this tutorial, we first give an overview of the components required for building an end-to-end web recommender system and then focus on how to use Photon ML (LinkedIn's open-sourced machine learning library) to train recommendation models and serve the results to users. Participants will get hands-on experience in training models of different levels of granularity to improve model performance and perform the "modeling loop" consisting of training a model, scoring candidate items using the model, seeing recommended items in a web UI, giving feedback to a number of recommended items, and then training a model again using the newly generated feedback.

The outline of the tutorial is as follows:

1. Introduction
    - Introduction to recommender systems
    - Overview of different components needed to build a web-based recommender system end to end
    - Overview of models and algorithms provided by the Photon ML library
2. Hands-on exercise in different recommendation models trained using Photon ML and deployed to a simple web UI.
    - Logistic regression model
    - Generalized linear mixed effect model
    - Interactive modeling with your feedback/ratings given to the recommended items

More details on Photon ML can be found at https://github.com/linkedin/photon-ml

KDD2016

ACM Code of Conduct