TOPICS IN COMPUTATIONAL ADVERTISING

ORGANIZER: DAVID BANKS, DUKE UNIVERSITY, USA

TALK 1: ESTIMATING CONVERSION RATES OF RARE EVENTS THROUGH A MULTIDIMENSIONAL DYNAMIC HIERARCHICAL BAYESIAN FRAMEWORK

SPEAKER: HONGXIA YANG, GOOGLE, USA

We consider the problem of estimating occurrence rates of rare events for extremely sparse data using preexisting hierarchies and selected features to perform inference along multiple dimensions. In particular, we focus on the problem of estimating click rates for {Advertiser, Publisher, User} tuples where both the Advertisers and the Publishers are organized as hierarchies that capture broad contextual in- formation at different levels of granularity. Typically, the click rates are low and the coverage of the hierarchies and dimensions is sparse. To overcome these difficulties, we decompose the joint prior of the three-dimensional Click-Through-Rate (CTR) using tensor decomposition and pro- pose a Multidimensional dynamic Hierarchical Bayesian framework (abbreviated as MadHab). We set up a specific framework of each dimension to model dimension-specific characteristics. More specifically, a hierarchical beta prior for the Advertiser dimension, a hierarchical dynamic Bayesian prior for the Publisher dimension and a feature-dependent mixture prior for the User dimension. Besides the centralized implementation, we propose two distributed algorithms for inference which make the model highly scalable and suited for large scale data mining applications. We demonstrate that on a real world ads campaign platform our framework can effectively discriminate extremely rare events in terms of their click propensity.

TALK 2: CONVEX BICLUSTERING

SPEAKER: GENEVERA ALLEN, RICE UNIVERSITY, USA

In the biclustering problem, we seek to simultaneously group observations and features. Biclustering has applications in a wide array of domains ranging from text mining and collaborative filtering to identifying structure in high dimensional biomedical data. We present a convex formulation of the biclustering problem that possesses a unique global minimizer and an iterative algorithm, COBRA, that is guaranteed to identify it. Our approach generates an entire solution path of possible biclusters as a single tuning parameter is varied. We also show how to reduce the problem of selecting this tuning parameter to solving a trivial modification of the convex biclustering problem. The key contributions of our work are its simplicity, interpretability, and algorithmic guarantees - features that arguably are lacking in the current alternative algorithms. We demonstrate the advantages of our approach, which includes stably and reproducibly identifying biclusterings, on simulated data as well as real applications in genomics and text mining. Joint work with Eric Chi and Richard Baraniuk.

TALK 3: ATTRITION PREDICTION USING CUSTOMER'S ONLINE AND OFFLINE ACTIVITIES

SPEAKER: HONGFEI LI, IBM T. J. WATSON RESEARCH CENTER, USA

Customer attrition prediction provides a market actionable impact to retain customers, improve customer satisfaction and increase loyalty. In order to predict customer attrition rate, it is critical to derive insights of customers' intentions and behaviors that lead to attrition based on the footprints they left through online and offline activities, such as transactions and complain calls. How to transform raw data into features that better represent the customer behaviors to predictive models is challenging and an ongoing research problem. In our work, we describe an approach that applies a variety of transformations to the original customer activity data and obtains behavior derivatives in order to capture customer intentions that indicate attritions. In addition to the feature-learning phase, we develop a time-varying survival model to predict attrition. Conventional survival model is commonly used to understand the factor impact on survival rate. It has limitation on its predictive capability due to unavailability of predicted time-varying covariates. We propose a survival model extension that provides customer attrition prediction for the future based on historical behavior derivatives while not requiring any inputs at future time. We demonstrate our approach via a real example and evaluate the performance using ROC curve. In addition, we illustrate how our solution generates customer attrition alerts that can help the marketing team take actions.