SCHOOL OF DATA SCIENCE

香港城市大學
City University of Hong Kong
專業 創新 胸懷全球
Professional·Creative For The World

SEMINAR SERIES

Model-independent algorithms for influence maximization and influence estimation in social networks

Date: 14 January 2019 (Monday)
Time: 2:30pm to 3:30pm
Venue: P7510, 7/F, Yeung Kin Man Academic Building (YEUNG), City University of Hong Kong

Professor Cautis, Bogdan
Department of Computer Science
University of Paris-Sud

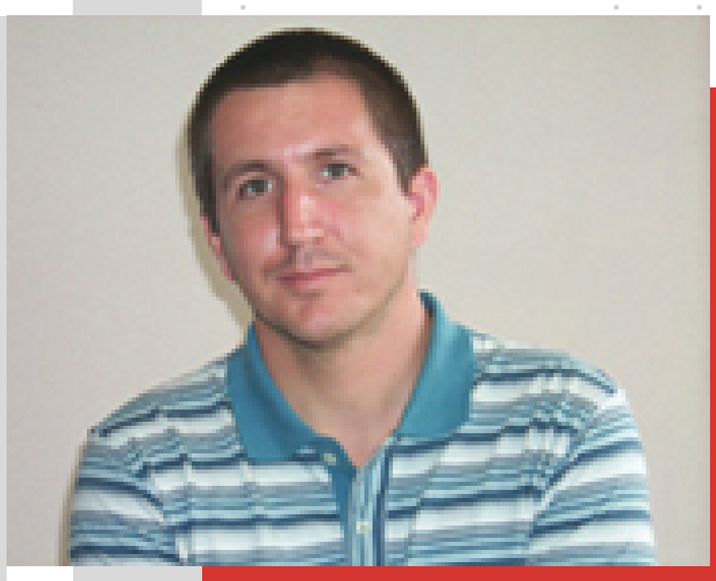

Guest Speaker's profile

Bogdan Cautis is a Professor in the Department of Computer Science at University of Paris-Sud, France, since September 2013. Previously, he was an Associate Professor at Ecole Nationale Superieure des T é l é communications in Paris, between 2007 and 2013. He received Engineering and Master degrees from Ecole Polytechnique, France, and his PhD degree from University of Paris-Sud and Institut National de Recherche en Informatique (INRIA) France, in 2007. His current research interests lie in the broad area of data management, data mining, and information retrieval.

Abstract

Word-of-mouth effects and influence are nowadays key ingredients for successful recommendation campaigns in social networks. I will discuss in this talk some of our recent research on understanding influence patterns and using them when running spread campaigns in social networks. First, we consider influence maximization (IM), the problem of finding influent users (nodes in a graph) so as to maximize the spread or adoption of information. We study a version of IM in which we maximize influence campaigns by adaptively selecting "spread seeds" from a set of candidates, a small subset of the node population. Influencer marketing is one straightforward application of this kind. Importantly, in our study, we make no assumptions on the underlying diffusion model and we work in a setting where neither a diffusion network nor historical adoption traces (so called information cascades) are available. We address this problem by an original approach, based on multi-armed bandit techniques for adaptive learning, and show that it leads to high-quality spreads on both simulated and real-world datasets, while being orders of magnitude faster than state-of-the-art IM methods. Second, we revisit the problem of inferring a diffusion network from information cascades. Once again, we make no assumptions on the underlying diffusion model, in this way obtaining a generic method with broader practical applicability. Our approach exploits the pairwise adoption-time intervals from cascades, with the observation that different kinds of information spread differently. Our experiments on both synthetic data and real-world datasets show that our method significantly outperforms the state-of-the-art algorithms. (This research was presented / will be presented at IEEE ICDM 2017, ACM TKDD 2019, ECML-PKDD 2019, and Springer Data Mining and Knowledge Discovery 2019.)

Enquiries: 3442 7887

All are welcome