

Fairness-Aware Data Mining

Toon Calders*

*Université Libre de Bruxelles (ULB) - CP 165/15
Avenue F.D. Roosevelt 50, B-1050 Bruxelles
toon.calders@ulb.ac.be
<http://cs.ulb.ac.be/members/tcalders>

Biography

Toon Calders obtained his PhD at the University of Antwerp in Belgium in 2003. He recently joined the ULB in 2012 where he holds the Chair in Business Intelligence, after working for 6 years at the Eindhoven University of Technology in The Netherlands as an assistant professor in the department of mathematics and computer science. His main research interests are Business Intelligence, Knowledge Discovery in Databases, machine learning and data mining. Toon Calders is area editor of the Springer Data Mining journal, he was one of the program chair of the ECMLPKDD conference in September 2014, and will be one of the program chairs of the Discovery Science conference in 2016. Toon Calders published over 60 papers in the data mining area including 17 journal papers in high-level data mining venues (DMKD, KDD, ICDM, SDM, ECML/PKDD). His research interests include pattern mining, entity resolution, discrimination and fairness aware data mining, and data stream processing.

Summary

In data mining we often have to learn from biased data, because, for instance, data comes from different batches or there was a gender or racial bias in the collection of social data. In some applications it may be necessary to explicitly control this bias in the models we learn from the data. Recently this topic received considerable interest both in the research community as well as more general, as witnessed by several recent articles in popular news media such as the New York Times. In this talk I will introduce and motivate research in fairness-aware data mining. Different techniques in unsupervised and supervised data mining will be discussed, dividing these techniques into three categories: algorithms of the first category adapt the input data in such a way to remove harmful biases while the second adapts the learning algorithms and the third category modifies the output models in such a way that its predictions become unbiased. Furthermore different ways to quantify unfairness, and indirect and conditional discrimination will be discussed, each with their own pros and cons. With this talk I hope to convincingly argue the validity and necessity of this often contested research area.

