

Lessons learned from the FIMI workshops

Bart Goethals *

May 4, 2019

Abstract of Invited Presentation

In 2003 and 2004, Mohammed Zaki and myself organised the Frequent Itemset Mining Implementations (FIMI) workshops in an attempt to characterize and understand the algorithmic performance space of frequent itemset mining algorithms.

In the years before, a huge number of algorithms had been developed in order to efficiently solve the frequent itemset mining (FIM) problem and several new algorithms were shown by their authors to run faster than previously existing algorithms. Unfortunately, the performance behavior of several of these algorithms was not always as was claimed by its authors when tested on some different datasets. Our goal was to understand precisely why and under what conditions one algorithm would outperform another.

In this talk, I will present some concrete examples of experimental evaluations that motivated the FIMI workshops, discuss the workshops results and whether we reached our goals, and my personal view on the evaluation of data mining methods.

*University of Antwerp