

# Association Rules : Past, Present, Future

Ramakrishnan Srikant

IBM Almaden Research Center.  
srikant@almaden.ibm.com

**Keywords:** Data Mining, KDD

The problem of mining association rules has received considerable research and industry attention, with widespread applications ranging from market basket analysis to detecting redundant medical tests. The original problem formulation has been extended in many directions, including the incorporation of taxonomies, quantitative associations, long associations, sequential patterns, and finding only those associations that satisfy user-specified constraints. There has also been much work on fast algorithms for both the original formulation and the above variations. In this talk, I will describe the problem and the most popular algorithmic approach to solving it, followed by an overview of the various extensions and current work on the topic, and conclude by presenting some open research topics.