

Association rule mining consists of first finding frequent itemsets (sets of items, such as A and B satisfying a minimum support threshold, or percentage of the task- relevant tuples), from which strong association rules in the form of $A \rightarrow B$ are generated.