Efficient Processing of Top-k Queries in Uncertain Databases

George Kollios
Boston University
http://www.cs.bu.edu/~gkollios/

Abstract:

In this talk I am going to discuss algorithms for processing top-k queries in uncertain databases, under the generally adopted model of x-relations. An x-relation consists of a number of x-tuples, and each x-tuple randomly instantiates into one tuple from one or more alternatives. Previous work identified two important variations of top-k queries in uncertain databases and provided simple algorithms to answer these queries. Unfortunately, the running time of these algorithms is exponential to k and/or the number of tuples that need to be retrieved. In this work, we present new methods that run in near linear or low polynomial time and cover both types of top-k queries in uncertain databases. We provide both a theoretical analysis and an extensive experimental evaluation to demonstrate the superiority of the new approaches over existing solutions.

Joint Work with: Ke Yi, Feifei Li, and Divesh Srivastava