

Lifted Inference in Probabilistic Databases

Dan Suciu

University of Washington, USA
suciu@cs.washington.edu

Abstract. Probabilistic Databases (PDBs) extend traditional relational databases by annotating each record with a weight, or a probability. Although PDBs define a very simple probability space, by simply adding constraints one can model much richer probability spaces, such as those represented by Markov Logic Networks or other Statistical Relational Models. While in traditional databases query evaluation corresponds to model checking and is always tractable, in probabilistic databases it becomes model counting, a notoriously computationally hard problem. “Lifted inference” is a new class of weighted model counting methods that exploit the First Order sentence in which the query is specified, rather than processing the Boolean formula directly. By analyzing the First Order sentence, it is often possible to classify the data complexity of weighted model counting problem into polynomial time or #P-hard.