

Representing and Querying Correlated Tuples in Probabilistic Databases

Prithviraj Sen
Dept. of Computer Science,
University of Maryland,
College Park, MD 20742.
sen@cs.umd.edu

Amol Deshpande
Dept. of Computer Science,
University of Maryland,
College Park, MD 20742.
amol@cs.umd.edu

Abstract

Probabilistic databases have received considerable attention recently due to the need for storing uncertain data produced by many real world applications. The widespread use of probabilistic databases is hampered by two limitations: (1) current probabilistic databases make simplistic assumptions about the data (e.g., complete independence among tuples) that make it difficult to use them in applications that naturally produce correlated data, and (2) most probabilistic databases can only answer a restricted subset of the queries that can be expressed using traditional query languages. We address both these limitations by proposing a framework that can represent not only probabilistic tuples, but also correlations that may be present among them. Our proposed framework naturally lends itself to the possible world semantics thus preserving the precise query semantics extant in current probabilistic databases. We develop an efficient strategy for query evaluation over such probabilistic databases by casting the query processing problem as an *inference* problem in an appropriately constructed *probabilistic graphical model*. We present several optimizations specific to probabilistic databases that enable efficient query evaluation. We validate our approach by presenting an experimental evaluation that illustrates the effectiveness of our techniques at answering various queries using real and synthetic datasets.

1 Introduction

Database research has primarily concentrated on how to store and query *exact* data. This has led to the development of techniques that allow the user to express and efficiently process complex queries in a declarative fashion over large collections of data. Unfortunately, many real-world applications produce large amounts of *uncertain* data. In such cases, databases need to do more than simply store and retrieve, they have to help the user sift through the uncertainty and find the results *most likely to be* the answer.

Numerous approaches have been proposed to handle uncertainty in databases [2, 5, 14, 4, 15, 17, 10, 27]. Among these, *tuple-level uncertainty models* [4, 15, 17, 10, 27], that associate existence probabilities with tuples, are considered more attractive for various reasons: (a) they typically result in relations that are in 1NF, (b) they provide simple and intuitive querying semantics, and (c) they are easier to store and operate on. However, these models often make simplistic and highly restrictive assumptions about the data (e.g., complete independence among base tuples [17, 10]). In particular, they cannot easily model or handle *dependencies/correlations*¹ among tuples. The ability to do so is critical for two reasons:

Natural dependencies in the data: Many application domains naturally produce correlated data. For instance, data integration may result in relations containing duplicate tuples that refer to the same *entity*; such tuples must be modeled as *mutually exclusive* [4, 1]. Real-world datasets such as the Christmas Bird Count [34, 12] naturally contain complex correlations among tuples. Data generated by sensor networks is typically highly correlated, both in time and space [14]. Data produced through use of machine learning techniques (e.g. classification labels) typically exhibits complex correlation patterns.

Dependencies during query evaluation: The problem of handling dependencies among tuples arises naturally during query evaluation *even when one assumes that the base data tuples are independent* (Section 2.1). In other words, the independent tuples assumption is not closed under the relational operators, specifically *join* [17, 10].

Past work on tuple-level uncertainty models has addressed this problem by either restricting the set of queries that can be evaluated against such a database (e.g. *safe plans* [10]), or by restricting the dependencies that can be modeled (e.g. *ProbView* [27]). Neither of these approaches, however, is satisfactory for a large class of real-world applications.

In this paper, we propose a *tuple-level uncertainty model* built on the foundations of statistical modeling techniques that allows us to uniformly handle dependencies in the data, while keeping the basic probabilistic framework simple and intuitive. The salient features of our proposed approach, and our main contributions are as follows:

- We propose a uniform framework for expressing uncertainties and dependencies through use of random variables and joint probability distribution. Unlike prior approaches, our proposed model is closed under relational algebra operations.
- We cast query evaluation on probabilistic databases as an *inference* problem in *probabilistic graphical models*, and develop techniques for efficiently constructing such models during query processing. This allows us to choose from various inference algorithms (exact or approximate) for query evaluation, depending on our requirements of accuracy and speed.

¹From here onwards, we use the terms “dependencies” and “correlations” interchangeably.

D^p			$pwd(D^p)$		$Evaluation$
	S^p		instance	probability	query result
	A	B	$prob$		
s_1	m	1	0.6	$d_1 = \{s_1, s_2, t_1\}$	{p}
s_2	n	1	0.5	$d_2 = \{s_1, s_2\}$	\emptyset
	T^p			$d_3 = \{s_1, t_1\}$	{p}
	C	D	$prob$	$d_4 = \{s_1\}$	\emptyset
t_1	1	p	0.4	$d_5 = \{s_2, t_1\}$	{p}
				$d_6 = \{s_2\}$	\emptyset
				$d_7 = \{t_1\}$	\emptyset
				$d_8 = \emptyset$	\emptyset

(i)
(ii)
(iii)

$Result$	
D	$prob$
p	$prob(d_1) + prob(d_3) + prob(d_5) = 0.32$
\emptyset	$prob(d_2) + prob(d_4) + prob(d_6) + prob(d_7) + prob(d_8) = 0.68$

(iv)

Figure 1: Example reproduced with minor changes from Dalvi and Suciu [10]: (i) A probabilistic database with independent tuples; (ii) corresponding possible worlds; (iii) evaluating $\prod_{\mathbf{D}}(S^p \bowtie_{\mathbf{B}=\mathbf{C}} T^p)$ over $pwd(D^p)$; (iv) computing result probabilities.

- We develop several optimizations specific to probabilistic databases resulting in efficient query execution in spite of the rich dependency modeling that we allow.
- We present experimental results from a prototype implementation over several real and synthetic datasets that demonstrate the need for modeling and reasoning about dependencies and the efficacy of our techniques at evaluating various queries including aggregate operators.

We begin with some background on tuple-level uncertainty models and probabilistic graphical models (Section 2). We then present our proposed model for representing correlated data (Section 3), develop strategies for efficiently executing arbitrary queries over them (Section 4) and present an experimental evaluation over a prototype implementation (Section 5). Finally, we discuss related work in Section 6, and conclude in Section 7.

2 Background

2.1 Independent Tuples Model [17, 10]

One of the most commonly used tuple-level uncertainty models, the *independent tuples model* [17, 10], associates existence probabilities with individual tuples and assumes that the tuples are independent of each other. Figure 1 (i) shows an example of such a database, D^p , with relations S^p (containing tuples s_1

and s_2 with probabilities 0.6 and 0.5 resp) and T^p (containing tuple t_1 with probability 0.4).

Such a probabilistic database can be interpreted as a probability distribution over the set of all possible deterministic database instances, called *possible worlds* (denoted by $pwd(D)$) [22, 17, 10]. Each deterministic instance (world) contains a subset of the tuples present in the probabilistic database, and the probability associated with it can be calculated directly using the independence assumption (by multiplying together the existence probabilities of tuples present in it and non-existence probabilities of tuples not present in it). Figure 1 (ii) shows all the possible worlds for D^p and their associated probabilities. For example, the probability of $d_2 = \{s_1, s_2\}$ is computed as $0.6 \times 0.5 \times (1 - 0.4) = 0.18$.

This possible worlds interpretation lends highly intuitive and precise semantics for query evaluation over probabilistic databases. Let q be a query issued on a probabilistic database D^p . We evaluate such a query against each possible world in $pwd(D^p)$ separately, thus resulting in another set of (*result*) possible worlds (with the same associated probabilities). The final result is obtained by taking a union of all the *result* possible worlds, and by associating a probability with each tuple in them to be the sum of the probabilities of the *result* possible worlds that contain it. For instance, Figure 1 (iii) shows the results of executing $\prod_{\mathbf{D}}(S^p \bowtie_{\mathbf{B}=\mathbf{C}} T^p)$ on each possible world of D^p and Figure 1 (iv) shows the final probability computation.

Evaluating a query via the set of possible worlds is clearly intractable as the number of possible worlds is exponential in the number of tuples contained in the database. Previous literature [17, 10] has suggested two query evaluation strategies instead, called *extensional* and *intensional* semantics. Intensional semantics guarantee results in accordance with possible worlds semantics but are computationally expensive. Extensional semantics, on the other hand, are computationally cheaper but do not guarantee results in accordance with the possible worlds semantics. This is because, even if base tuples are independent of each other, the intermediate tuples that are generated during query evaluation are typically correlated. For instance, consider the example query shown in Figure 1. If the join operation is performed before the projection, the two intermediate tuples that are generated, s_1t_1 and s_2t_1 , are not independent of each other since they share t_1 .

2.2 Tuple Correlations

As we discussed in Section 1, tuple correlations also occur naturally in many application domains, and ignoring such correlations can result in highly inaccurate and unintuitive query results. Consider the four sets of possible worlds shown in Figure 2 (i) derived from the same database shown in Figure 1 (i), but containing different sets of dependencies that we might want to represent:

1. *ind.*: where s_1 , s_2 , and t_1 are independent of each other.
2. *implies*: presence of t_1 implies absence of s_1 and s_2 ($t_1 \Rightarrow \neg s_1 \wedge \neg s_2$).

$pwd(D^p)$ instance	probability			
	<i>ind.</i>	<i>implies</i>	<i>mut. ex.</i>	<i>nxor</i>
$d_1 = \{s_1, s_2, t_1\}$	0.12	0	0	0.2
$d_2 = \{s_1, s_2\}$	0.18	0.5	0.3	0.1
$d_3 = \{s_1, t_1\}$	0.12	0	0	0.2
$d_4 = \{s_1\}$	0.18	0.1	0.3	0.1
$d_5 = \{s_2, t_1\}$	0.08	0	0.2	0
$d_6 = \{s_2\}$	0.12	0	0	0.2
$d_7 = \{t_1\}$	0.08	0.4	0.2	0
$d_8 = \emptyset$	0.12	0	0	0.2

(i)

D	<i>ind.</i>	<i>implies</i>	<i>mut. ex.</i>	<i>nxor</i>
p	0.32	0.00	0.20	0.40

(ii)

Figure 2: (i) $pwd(D^p)$ for various dependencies and the (ii) corresponding query results.

3. *mutual exclusivity (mut. ex.)*: $t_1 \Rightarrow \neg s_1$ and $s_1 \Rightarrow \neg t_1$.
4. *nxor*: high positive correlation between t_1 and s_1 , presence (absence) of one almost certainly implies the presence (absence) of the other.

Figure 2 (ii) shows the result of applying the example query from Figure 1 to these four possible worlds. As we can see, although the tuple probabilities associated with s_1 , s_2 and t_1 are identical, the query results are drastically different across these four databases. Note that, since both the approaches (intensional and extensional semantics) discussed in the previous section assume base tuple independence, neither can be directly used to do query evaluation in such cases.

2.3 Probabilistic Graphical Models and Factored Representations

Probabilistic graphical models form a powerful class of approaches that can compactly represent and reason about complex dependency patterns involving large numbers of correlated random variables [30, 8]. The key idea underlying these approaches is the use of *factored representations* for modeling the correlations.

Let X denote a random variable with a domain $dom(X)$ and let $Pr(X)$ denote a probability distribution over it. Similarly, let $\mathbf{X} = \{X_1, X_2, X_3, \dots, X_n\}$ denote a set of n random variables each with its own associated domain $dom(X_i)$, and $Pr(\mathbf{X})$ denote the join probability distribution over them.

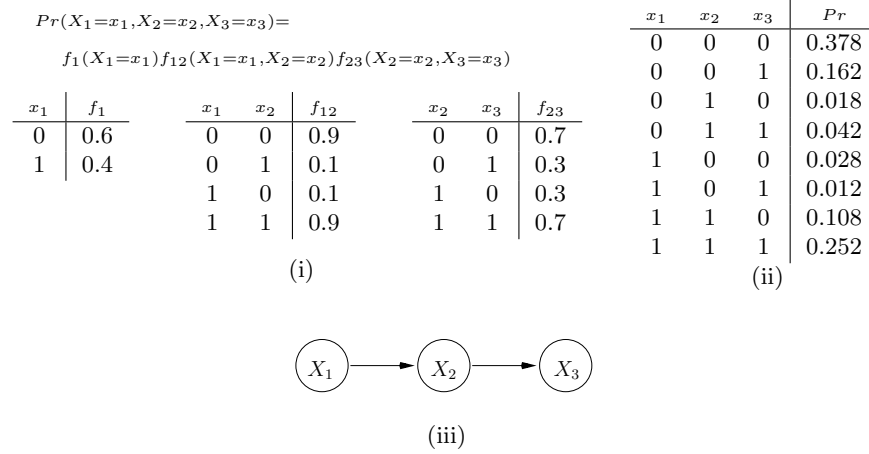


Figure 3: Example involving three dependent random variables each with a binary domain: (i) factored representation (ii) resulting joint probability distribution (iii) graphical model representation.

Definition 2.1. A factor² $f(\mathbf{X})$ is a function of a (small) set of random variables $\mathbf{X} = \{X_1, \dots, X_n\}$ such that $0 \leq f(\mathbf{X} = \mathbf{x}) \leq 1 \forall \mathbf{x} \in \text{dom}(X_1) \times \dots \times \text{dom}(X_n)$.

A factored representation of $Pr(\mathbf{X})$ allows the distribution to be represented compactly as a product of factors:

$$Pr(\mathbf{X} = \mathbf{x}) = \prod_{i=1}^m f_i(\mathbf{X}_i = \mathbf{x}_i) \quad (1)$$

where $\mathbf{X}_i \subseteq \mathbf{X}$ is the set of random variables restricted to factor f_i and \mathbf{x}_i is the corresponding assignment. Figure 3 shows a small example of a factored representation of a joint probability distribution over three random variables.

Computing *marginal probabilities* is a common operation when dealing with such joint probability distributions. It falls under the general class of operations known as *inference*. Given a random variable $X \in \mathbf{X}$ and an assignment $x \in \text{dom}(X)$, the marginal probability computation problem from the joint distribution $Pr(\mathbf{X})$ is:

$$Pr(X = x) = \sum_{\mathbf{x} \sim x} Pr(\mathbf{X} = \mathbf{x}) \quad (2)$$

where $\mathbf{x} \sim x$ denotes an assignment to \mathbf{X} that agrees with $X = x$ and \mathbf{x} is a valid assignment to \mathbf{X} . We will discuss inference techniques in Section 4.1.

²Factors can be seen as a generalization of *Conditional probability tables* in Bayesian networks [30].

$Pr(\mathbf{X}_{DP}) = f_{t_1}^{ind}(X_{t_1}) f_{t_1, s_2}^{implies}(X_{t_1}, X_{s_2})$ $f_{t_1, s_1}^{implies}(X_{t_1}, X_{s_1})$ <table style="margin-left: auto; margin-right: auto; border-collapse: collapse;"> <thead> <tr> <th style="border-right: 1px solid black; padding: 5px;">X_{t_1}</th> <th style="border-right: 1px solid black; padding: 5px;">X_{s_1}</th> <th style="padding: 5px;">$f_{t_1, s_1}^{implies}$</th> </tr> </thead> <tbody> <tr><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="padding: 5px;">0</td></tr> <tr><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="padding: 5px;">1</td></tr> <tr><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="padding: 5px;">1</td></tr> <tr><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="padding: 5px;">0</td></tr> </tbody> </table> <table style="margin-left: auto; margin-right: auto; border-collapse: collapse;"> <thead> <tr> <th style="border-right: 1px solid black; padding: 5px;">X_{t_1}</th> <th style="border-right: 1px solid black; padding: 5px;">X_{s_2}</th> <th style="padding: 5px;">$f_{t_1, s_2}^{implies}$</th> </tr> </thead> <tbody> <tr><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="padding: 5px;">1/6</td></tr> <tr><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="padding: 5px;">5/6</td></tr> <tr><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="padding: 5px;">1</td></tr> <tr><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="padding: 5px;">0</td></tr> </tbody> </table> <p style="text-align: center;">(i)</p>	X_{t_1}	X_{s_1}	$f_{t_1, s_1}^{implies}$	0	0	0	0	1	1	1	0	1	1	1	0	X_{t_1}	X_{s_2}	$f_{t_1, s_2}^{implies}$	0	0	1/6	0	1	5/6	1	0	1	1	1	0	$Pr(\mathbf{X}_{DP}) = f_{s_2}^{ind}(X_{s_2}) f_{t_1, s_1}^{mutex}(X_{t_1}, X_{s_1})$ <table style="margin-left: auto; margin-right: auto; border-collapse: collapse;"> <thead> <tr> <th style="border-right: 1px solid black; padding: 5px;">X_{t_1}</th> <th style="border-right: 1px solid black; padding: 5px;">X_{s_1}</th> <th style="padding: 5px;">f_{t_1, s_1}^{mutex}</th> </tr> </thead> <tbody> <tr><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="padding: 5px;">0</td></tr> <tr><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="padding: 5px;">0.6</td></tr> <tr><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="padding: 5px;">0.4</td></tr> <tr><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="padding: 5px;">0</td></tr> </tbody> </table> <p style="text-align: center;">(ii)</p> $Pr(\mathbf{X}_{DP}) = f_{s_2}^{ind}(X_{s_2}) f_{t_1, s_1}^{nxor}(X_{t_1}, X_{s_1})$ <table style="margin-left: auto; margin-right: auto; border-collapse: collapse;"> <thead> <tr> <th style="border-right: 1px solid black; padding: 5px;">X_{t_1}</th> <th style="border-right: 1px solid black; padding: 5px;">X_{s_1}</th> <th style="padding: 5px;">f_{t_1, s_1}^{nxor}</th> </tr> </thead> <tbody> <tr><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="padding: 5px;">0.4</td></tr> <tr><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="padding: 5px;">0.2</td></tr> <tr><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="border-right: 1px solid black; padding: 5px;">0</td><td style="padding: 5px;">0</td></tr> <tr><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="border-right: 1px solid black; padding: 5px;">1</td><td style="padding: 5px;">0.4</td></tr> </tbody> </table> <p style="text-align: center;">(iii)</p>	X_{t_1}	X_{s_1}	f_{t_1, s_1}^{mutex}	0	0	0	0	1	0.6	1	0	0.4	1	1	0	X_{t_1}	X_{s_1}	f_{t_1, s_1}^{nxor}	0	0	0.4	0	1	0.2	1	0	0	1	1	0.4
X_{t_1}	X_{s_1}	$f_{t_1, s_1}^{implies}$																																																											
0	0	0																																																											
0	1	1																																																											
1	0	1																																																											
1	1	0																																																											
X_{t_1}	X_{s_2}	$f_{t_1, s_2}^{implies}$																																																											
0	0	1/6																																																											
0	1	5/6																																																											
1	0	1																																																											
1	1	0																																																											
X_{t_1}	X_{s_1}	f_{t_1, s_1}^{mutex}																																																											
0	0	0																																																											
0	1	0.6																																																											
1	0	0.4																																																											
1	1	0																																																											
X_{t_1}	X_{s_1}	f_{t_1, s_1}^{nxor}																																																											
0	0	0.4																																																											
0	1	0.2																																																											
1	0	0																																																											
1	1	0.4																																																											

Figure 4: Representing probabilistic databases with dependencies (examples from Figure 2): (i) “implies” dependency; (ii) “mut. ex.” dependency; (iii) “nxor” dependency. $f_{t_1}^{ind}(X_{t_1})$ and $f_{s_2}^{ind}(X_{s_2})$ refer to the independent tuple factors for t_1 and s_2 .

3 Proposed Approach

We now describe how to represent dependencies in probabilistic databases through the use of factors and random variables. We then consider the query evaluation problem over such databases, and begin by presenting an example that grounds the basic ideas underlying our approach. We then follow up with a more detailed description of the overall query evaluation procedure.

3.1 Representing dependencies

Let t denote a tuple belonging to relation R such that t is a mapping from attributes in R to constants. In our framework, every tuple t is associated with a unique boolean valued random variable X_t where 0 represents **false** and 1 represents **true**.

A *probabilistic relation* R consists of a set of tuples with their corresponding random variables and a *probabilistic database* D consists of a collection of *probabilistic relations*. We refer to the collection of all the random variables associated with tuples in the probabilistic database D by the symbol \mathbf{X}_D .

Each instance in $pwd(D)$ can now be expressed as a complete assignment to the set of random variables \mathbf{X}_D denoted by $\mathbf{x}_D \in \{0, 1\}^{|\mathbf{X}_D|}$. For example, d_2 in Figure 2 (i) corresponds to the assignment $X_{s_1} = 1, X_{s_2} = 1, X_{t_1} = 0$.

We can now represent dependencies by defining factors on the tuple-based

random variables in the database. The probability of an instance can be computed by computing the joint probability of the assignment to \mathbf{X}_D which can in turn be obtained by multiplying all factors defined on the tuple-based random variables in the database (Eq. (1)).

Example: Representing Independent Tuples.

We illustrate our approach by expressing the probabilistic database in Figure 1 (i) (with three independent tuples) in our formulation. This can be achieved by defining one factor per independent tuple:

X_{s_1}	$f_{s_1}^{ind}$	X_{s_2}	$f_{s_2}^{ind}$	X_{t_1}	$f_{t_1}^{ind}$
0	0.4	0	0.5	0	0.6
1	0.6	1	0.5	1	0.4

To compute the probability for an instance we multiply these factors. For instance:

$$\begin{aligned}
 Pr(d_2 = \{s_1, s_2\}) &= Pr(X_{s_1} = 1, X_{s_2} = 1, X_{t_1} = 0) \\
 &= f_{s_1}^{ind}(X_{s_1} = 1) f_{s_2}^{ind}(X_{s_2} = 1) f_{t_1}^{ind}(X_{t_1} = 0) \\
 &= 0.6 \times 0.5 \times 0.6 = 0.18
 \end{aligned}$$

Figure 4 shows the factored representations for the other three probabilistic databases in Figure 2 (i).

3.2 Query Evaluation: Example

We begin our discussion on query evaluation by presenting a small example. Consider the database shown in Figure 2 (i) with the “nxor” dependency (Figure 4 (iii)). Figure 5 describes the execution of $\prod_D(S^p \bowtie_{B=C} T^p)$ on this database.

Consider the intermediate tuples introduced during the execution of this query. The tuples i_1 and i_2 , produced by the join (Figure 5), are clearly uncertain tuples since they are not produced in every instance of the database. Similarly, the result tuple r_1 is also a probabilistic tuple. Let us take a closer look at the inter-tuple dependencies:

- i_1 is produced by possible world d iff d contains both s_1 and t_1 (i.e., $s_1 \wedge t_1 \Leftrightarrow i_1$).
- Similarly, $s_2 \wedge t_1 \Leftrightarrow i_2$.
- Finally, r_1 is produced iff either i_1 or i_2 is produced.

Figure 5 shows the factors for these dependencies:

- f_{i_1, s_1, t_1}^{and} and f_{i_2, s_2, t_1}^{and} return 1 when the first argument is the logical-and of the last two arguments (see Figure 6 (i) for the full definition of f_{i_1, s_1, t_1}^{and}).
- f_{r_1, i_1, i_2}^{or} returns 1 when the first argument is the logical-or of the last two arguments (Figure 6 (ii)).

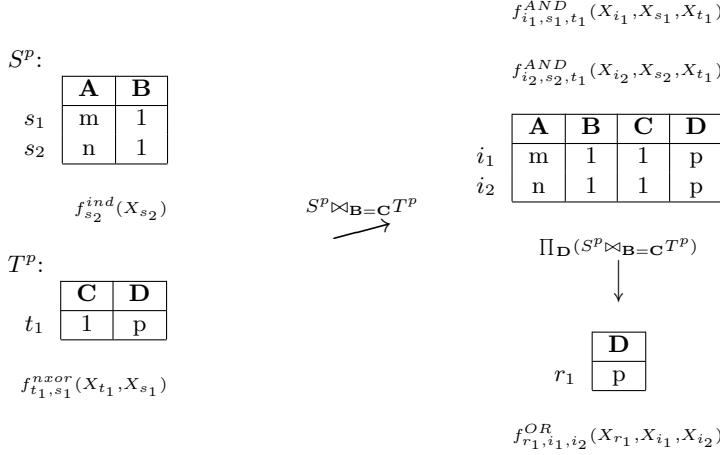


Figure 5: Solving $\prod_D(S^p \bowtie_{B=C} T^p)$ on D^p with “nxor” dependency (Figure 2 (i)).

Consider the factored probability distribution induced by this query, i.e. the product of all the factors introduced including the factors among the base tuples. It turns out that the marginal probability of $Pr(X_{r_1} = 1)$ returns the correct answer 0.40 matching the number in Figure 2 (ii).

3.3 Query Evaluation: Details

3.3.1 Generating factors during query evaluation

The query evaluation procedure presented through the example above requires that we encode dependencies among (intermediate) tuples by introducing factors. We now redefine the three operators σ , \prod , \times to produce factors expressing such dependencies. Let us denote the new operators by σ^p , \prod^p , \times^p where the superscript emphasizes that they are operators for probabilistic databases. We will assume that our query q does not contain two copies of the same relation and consists of only the above three operators.

- select (σ_c^p): Suppose σ_c^p operator with predicate c acts on tuple t and produces new tuple r , both containing the same mapping. There are two cases to consider, if t does not satisfy the predicate c , then r cannot be produced and this is enforced by a **false** factor on X_r that returns 1 if $X_r = 0$ and 0 when $X_r = 1$. The other case is when t satisfies c and in this case, X_r holds true in possible world d if and only if X_t holds true. This is enforced by an **equals** factor that takes arguments X_r and X_t and returns 1 whenever $X_r = X_t$ and 0 o.w.
- join (\times^p): Let tuples t and t' join to return new tuple r . In this case, as we saw in the example, X_r will hold true in possible world d if and only if both X_t and $X_{t'}$ hold true. This can be enforced with the **and** factor that we introduced in the example (Figure 6 (i)).

X_{i_1}	X_{s_1}	X_{t_1}	f_{i_1, s_1, t_1}^{and}
0	0	0	1
0	0	1	0
0	1	0	1
0	1	1	0
1	0	0	1
1	0	1	0
1	1	0	0
1	1	1	1

(i)

X_{r_1}	X_{i_1}	X_{i_2}	f_{r_1, i_1, i_2}^{or}
0	0	0	1
0	0	1	0
0	1	0	0
0	1	1	1
1	0	0	0
1	0	1	1
1	1	0	0
1	1	1	1

(ii)

Figure 6: Definitions of (i) $f_{i_1, s_1, t_1}^{AND}(X_{i_1}, X_{s_1}, X_{t_1})$ and (ii) $f_{r_1, i_1, i_2}^{OR}(X_{r_1}, X_{i_1}, X_{i_2})$ from Figure 5.

- **project ($\prod_{\mathbf{A}}^P$):** Let tuples t_1, \dots, t_n project to form new tuple r . Thus, X_r holds true in possible world d if at least one of X_{t_1}, \dots, X_{t_n} hold true in d . This can be enforced with a straightforward generalization of the **or** factor, introduced in the example (Figure 6 (ii)), $f_{r, t_1, \dots, t_n}^{or}(X_r, X_{t_1}, \dots, X_{t_n})$ that returns 1 if the first argument is the logical-or of the last n arguments.

Figure 7 defines the above operators in functional representation where $\mathcal{F}_q(t)$ denotes the set of factors required to generate tuple t inductively on query plan q . Figure 7 also gives the rules for \cup^p and $-^p$. $q \cup^p q'$ has similar semantics as \prod^p , if either of q or q' produce t then X_t is true and hence the use of f^{or} . t is a result of $q -^p q'$ if t is produced by q and not produced by q' and this is encoded in the factor $f^{true-false}$.

Figure 7 define the operators assuming there are no self-joins. It is easy to get around this assumption by introducing **if ... then ... else ...** conditions in the definitions. For example, consider executing the query $R \bowtie R$. Consider the result of joining tuples $t, t' \in R$ where $t \neq t'$, since t and t' are associated with different random variables we will introduce a factor $f_{r, t, t'}^{and}(X_r, X_t, X_{t'})$ where r denotes the result tuple. On the other hand if t joins with itself, then $r = t \bowtie t$ and the uncertainty associated with r is only associated with the uncertainty of t , hence in this case we only introduce a factor $f_{r, t}^{equals}(X_r, X_t)$.

3.3.2 Formulating the Query Evaluation Problem

We now justify our approach of generating factors for query evaluation. Let $Pr_q(t)$ denote the probability associated with result tuple t produced by query q issued on the database D . Thus, from possible world semantics:

$$Pr_q(t) = \sum_{\mathbf{x} \in \{0,1\}^{|\mathbf{x}_D|}, t \sim q(\mathbf{x})} P(\mathbf{X}_D = \mathbf{x})$$

where $t \sim q(\mathbf{x})$ is true if t results from possible world \mathbf{x} .

Let X_t denote the random variable associated with t . We would like to express the above computation as a marginal probability computation problem

$$\begin{aligned}
\mathcal{F}_{\sigma_c^p(q)}(t) &= \begin{cases} f^{false}(X_t) & \text{if } c(t) \text{ is false} \\ f^{equals}(X_t, X_{t'}) \cup \mathcal{F}_q(t') & \text{if } c(t) \text{ is true} \\ & \text{where } t = t' \end{cases} \\
\mathcal{F}_{\prod_A^p q}(t) &= f^{or}(X_t, \{X_{t'} | t = \prod_A t'\}) \cup \{\mathcal{F}_q(t') | t = \prod_A t'\} \\
\mathcal{F}_{q \times p_{q'}}(t, t') &= f^{and}(X_{t \times t'}, X_t, X_{t'}) \cup \mathcal{F}_q(t) \cup \mathcal{F}_{q'}(t') \\
\mathcal{F}_{q \cup p_{q'}}(t) &= f^{or}(X_t, X_{t'}, X_{t''}) \cup \mathcal{F}_q(t') \cup \mathcal{F}_{q'}(t'') \text{ where } t = t' = t'' \\
\mathcal{F}_{q - p_{q'}}(t) &= f^{true-false}(X_t, X_{t'}, X_{t''}) \cup \mathcal{F}_q(t') \cup \mathcal{F}_{q'}(t'') \text{ where } t = t' = t''
\end{aligned}$$

Figure 7: Definitions for relational algebra operators.

on the distribution $Pr(\mathbf{X}_D, X_t)$ such that:

$$Pr_q(t) = \sum_{\mathbf{x} \in \{0,1\}^{|\mathbf{X}_D|}} Pr(\mathbf{X}_D = \mathbf{x}, X_t = 1)$$

where $Pr(\mathbf{X}_D = \mathbf{x}, X_t = 1)$ is the distribution obtained by the product of factors introduced during query evaluation and present amongst the base tuples with the unnecessary random variables summed out.

For the above reformulation to hold we need to show that $Pr(\mathbf{X}_D = \mathbf{x}, X_t = 1)$ satisfies the following constraints:

$$Pr(\mathbf{X}_D = \mathbf{x}, X_t = 1) = \begin{cases} Pr(\mathbf{X}_D = \mathbf{x}) & \text{if } t \sim q(\mathbf{x}) \\ 0 & \text{o.w.} \end{cases} \quad (3)$$

In this subsection we will show that for any intermediate tuple t introduced while evaluating query q $Pr(X_t = 1, \mathbf{X} = \mathbf{x})$ satisfies Eq. (3).

For ease of exposition, let us rewrite Eq. (3) in terms of conditional probabilities. Given two sets of random variables \mathbf{X}_1 and \mathbf{X}_2 , the conditional probability of $\mathbf{X}_1 = \mathbf{x}_1$ given $\mathbf{X}_2 = \mathbf{x}_2$ is defined as:

$$Pr(\mathbf{X}_1 = \mathbf{x}_1 | \mathbf{X}_2 = \mathbf{x}_2) = \frac{Pr(\mathbf{X}_1 = \mathbf{x}_1, \mathbf{X}_2 = \mathbf{x}_2)}{Pr(\mathbf{X}_2 = \mathbf{x}_2)}$$

Thus, using the chain rule of probability on $Pr(X_t = 1, \mathbf{X}_D = \mathbf{x})$:

$$Pr(\mathbf{X}_D = \mathbf{x}, X_t = 1) = Pr(X_t = 1 | \mathbf{X}_D = \mathbf{x}) Pr(\mathbf{X}_D = \mathbf{x})$$

which means Eq. (3) can be equivalently expressed as:

$$Pr(X_t = 1 | \mathbf{X}_D = \mathbf{x}) = \begin{cases} 1 & \text{if } t \sim q(\mathbf{x}) \\ 0 & \text{o.w.} \end{cases}$$

Since a conditional distribution is also a probability distribution and X_t is a boolean random variable, the full set of constraints on $Pr(X_t = 1 | \mathbf{X}_D = \mathbf{x})$ can be expressed as:

	$Pr(X_t = 0 \mathbf{X}_D = \mathbf{x})$	$Pr(X_t = 1 \mathbf{X}_D = \mathbf{x})$	
if $t \sim q(\mathbf{x})$	0	1	(4)
o. w.	1	0	

Instead of justifying Eq. (3) it suffices to justify Eq. (4) to show that the reformulation is admissible.

Before providing the justification we first require some terminology which will help identify some properties of the query evaluation procedure introduced in the previous subsection.

Definition 3.1. A conditional probability distribution (CPD) is given by a factor $f(\mathbf{X})$ over a set of random variables \mathbf{X} such that \mathbf{X} can be partitioned into \mathbf{X}_π (a set of random variables) and X_c (a single random variable) and $X_c \notin \mathbf{X}_\pi$, $X_c \cup \mathbf{X}_\pi = \mathbf{X}$. Let $Val(\mathbf{X}_\pi)$ and $Val(X_c)$ denote the possible assignments to \mathbf{X}_π and X_c respectively, then $f(\mathbf{X})$ satisfies:

$$\sum_{v \in Val(X_c)} f(X_c = v, \mathbf{X}_\pi = V) = 1, \quad \forall V \in Val(\mathbf{X}_\pi) \quad (5)$$

Notice that in Figure 7, we use only factors representing CPDs (the factors over the base tuple random variables can, however, involve factors that are not CPDs). As an illustration, for $f_{r_1, i_1, i_2}^{or}(X_{r_1}, X_{i_1}, X_{i_2})$ in Figure 6 (ii) setting $\mathbf{X}_\pi = \{X_{i_1}, X_{i_2}\}$ and $X_c = X_{r_1}$ shows how $f_{r_1, i_1, i_2}^{or}(X_{r_1}, X_{i_1}, X_{i_2})$ satisfies the above definition and represents a CPD.

It is traditional to express a CPD using a $|$ to separate X_c from \mathbf{X}_π . Further, in the previous subsection we identified factors by their semantics. For example, $f_{r_1, i_1, i_2}^{or}(X_{r_1}, X_{i_1}, X_{i_2})$ (Figure 6 (ii)) identified a factor with OR semantics. In the subsequent discussion, however, we will identify factors by the relational algebra operators they are associated with and this we will denote by $F^{op}(X_c | \mathbf{X}_\pi)$. Thus, as an illustration, $F^\Pi(X_{r_1} | X_{i_1}, X_{i_2})$ would denote $f_{r_1, i_1, i_2}^{or}(X_{r_1}, X_{i_1}, X_{i_2})$.

Another aspect of the design of our query evaluation algorithm is that for each relational algebra operator op , the corresponding F^{op} contains as X_c the random variable corresponding to the tuple generated at the time of introduction of F^{op} . For example, in the query evaluation example presented in Section 3.2, $f_{r_1, i_1, i_2}^{or}(X_{r_1}, X_{i_1}, X_{i_2})$ was introduced while generating tuple r_1 (see Figure 5) and thus $X_c = X_{r_1}$ for $f_{r_1, i_1, i_2}^{or}(X_{r_1}, X_{i_1}, X_{i_2})$.

Definition 3.2. A deterministic conditional probability distribution (det. CPD) is given by a CPD $f(\mathbf{X})$ over a set of random variables \mathbf{X} such that for every instantiation $v \in Val(\mathbf{X})$:

$$f(\mathbf{X} = v) = \text{either } 1 \text{ or } 0, \quad \forall v \in Val(\mathbf{X})$$

All factors in Figure 7 represent det. CPDs. Given an assignment to \mathbf{X}_π of a det. CPD, there is one and exactly one assignment to the corresponding X_c such that the probability of the joint assignment is non-zero. For example, in

$f_{r_1, i_1, i_2}^{or}(X_{r_1}, X_{i_1}, X_{i_2})$ (Figure 6 (ii)), for $X_{i_1} = 0$, $X_{i_2} = 1$ the only assignment to X_{r_1} that gives a joint assignment with non-zero probability is $X_{r_1} = 1$.

We now present a lemma that vouches for the “correctness” of the factor F^{op} and will be useful while presenting our justification.

We denote the generation of a new tuple t by the application of relational algebra operator op on the collection of tuples S by $S \xrightarrow{op} t$. Note that not all tuples in S need to be true or present to generate t . For example, during projection, any one tuple that projects to the result tuple t need be present in any instantiation of the probabilistic database to produce t . To denote this, we employ the more expressive notation $\mathbf{X}_S = \mathbf{x}_S \xrightarrow{op} X_t = x_t$ where \mathbf{X}_S denotes the random variables corresponding to the collection S , X_t denotes the random variable corresponding to t and \mathbf{x}_S and x_t denote the corresponding assignments.

Lemma 3.1. *Each relational operator op and its corresponding factor f^{op} match in semantics. In other words:*

$$F^{op}(X_t = x_t | \mathbf{X}_S = \mathbf{x}_S) = \begin{cases} 1 & \text{if } \mathbf{X}_S = \mathbf{x}_S \xrightarrow{op} X_t = x_t \\ 0 & \text{if } \mathbf{X}_S = \mathbf{x}_S \xrightarrow{op} X_t = \bar{x}_t \end{cases}$$

Proof. By construction. □

Note that the above lemma was possible partly because $F^{op}()$ is a det. CPD. Otherwise, given any assignment to \mathbf{X}_S we would have obtained a distribution over assignments to X_t .

We are now ready to state the main theorem of this subsection.

Theorem 3.1. *For any tuple t generated by applying query q to probabilistic database D , the probability associated with t under possible world semantics $Pr_q(t)$ is equal to the marginal probability corresponding to $X_t = 1$ from the distribution induced by the factors introduced by query evaluation operations (Figure 7) along with the factors present amongst the base tuples. More precisely,*

$$Pr_q(t) = \sum_{\mathbf{x} \in \{0,1\}^{|\mathbf{X}_D|}} Pr(\mathbf{X}_D = \mathbf{x}, X_t = 1)$$

where $Pr(\mathbf{X}_D = \mathbf{x}, X_t = 1)$ is the distribution obtained by the product of factors introduced during query evaluation and present amongst the base tuples with the unnecessary random variables summed out.

Proof. As discussed earlier, instead of directly proving the theorem we will, instead, show that $Pr(X_t = x_t, \mathbf{X}_D = \mathbf{x})$ satisfies Eq. (4).

The proof is by induction on the operators in query q . We first consider the base case when q consists of only one operator.

Base Case: Let op denote a relational algebra operator that acts on the base tuples in probabilistic database D to produce the tuple t .

Let $\mathbf{X}_i \subseteq \mathbf{X}_D$ denote the random variable corresponding to the tuples that are involved in generating t .

$$\begin{aligned} Pr(X_t = 1 | \mathbf{X}_D = \mathbf{x}) &= \frac{Pr(X_t = 1, \mathbf{X}_D = \mathbf{x})}{Pr(\mathbf{X}_D = \mathbf{x})} \\ &= \frac{F^{op}(X_t = 1 | \mathbf{X}_i = \mathbf{x}_i) Pr(\mathbf{X}_D = \mathbf{x})}{Pr(\mathbf{X}_D = \mathbf{x})} \\ &= F^{op}(X_t = 1 | \mathbf{X}_i = \mathbf{x}_i) \end{aligned}$$

which means $Pr(X_t = 1 | \mathbf{X}_D = \mathbf{x})$ is either zero or one since $F^{op}(X_t = 1 | \mathbf{X}_i = \mathbf{x}_i)$ is a det. CPD. All that needs to be shown now is that $Pr(X_t = 1 | \mathbf{X}_D = \mathbf{x})$ is one when $t \sim q(\mathbf{x})$ and zero otherwise.

$t \sim q(\mathbf{x})$ if $\mathbf{X}_i = \mathbf{x}_i \xrightarrow{op} X_t = 1$ in which case by Lemma (3.1) $Pr(X_t = 1 | \mathbf{X}_D = \mathbf{x}) = F^{op}(X_t = 1 | \mathbf{X}_i = \mathbf{x}_i) = 1$. Also, t is not present in $q(\mathbf{x})$ if $\mathbf{X}_i = \mathbf{x}_i \xrightarrow{op} X_t = 0$ in which case by Lemma (3.1) $Pr(X_t = 1 | \mathbf{X}_D = \mathbf{x}) = F^{op}(X_t = 1 | \mathbf{X}_i = \mathbf{x}_i) = 0$. Thus the base case holds.

Inductive Hypothesis: Let tuple t be the product of some relational algebra operator op on the (intermediate) tuples $\{i_1, i_2, \dots, i_n\}$. Let Eq. (4) hold for each i_j , $j = 1, \dots, n$. Thus:

	$Pr(X_{i_j} = 0 \mathbf{X}_D = \mathbf{x})$	$Pr(X_{i_j} = 1 \mathbf{X}_D = \mathbf{x})$	
if $i_j \sim q(\mathbf{x})$	0	1	
o. w.	1	0	(6)

Let us first take a close look at what Eq. (6) implies about the expression:

$$\begin{aligned} &\sum_{x_i \in \{0,1\}^n} Pr(\{X_{i_j} = x_{i_j} | j = 1, \dots, n\} | \mathbf{X}_D = \mathbf{x}) \\ &= \sum_{x_i \in \{0,1\}^n} \prod_{j=1, \dots, n} Pr(X_{i_j} = x_{i_j} | \{X_{i_k} = x_{i_k} | k = j+1, \dots, n\}, \mathbf{X}_D = \mathbf{x}) \\ &= \sum_{x_i \in \{0,1\}^n} \prod_{j=1, \dots, n} Pr(X_{i_j} = x_{i_j} | \mathbf{X}_D = \mathbf{x}) \end{aligned}$$

where the second line follows from conditional independence of random variables in \mathbf{X}_i .

Construct an assignment to \mathbf{X}_i by choosing x_{i_k} such that $Pr(X_{i_k} = x_{i_k} | \mathbf{X}_D = \mathbf{x}) = 1$. From Eq. (6) we know such an assignment exists for each X_{i_k} . Denote this joint assignment by \mathbf{y} .

The only term in the summation above that is non-zero is the term corresponding to $\mathbf{X}_i = \mathbf{y}$, all other terms result in zeroes. Because of this we say $\mathbf{X}_D = \mathbf{x} \xrightarrow{*} \mathbf{X}_i = \mathbf{y}$. Thus:

$$\sum_{x_i \in \{0,1\}^n} Pr(\{X_{i_j} = x_{i_j} | j = 1, \dots, n\} | \mathbf{X}_D = \mathbf{x}) = Pr(X_i = \mathbf{y} | \mathbf{X}_D = \mathbf{x}) = 1 \quad (7)$$

Thus the inductive hypothesis implies an assignment to \mathbf{X}_i (\mathbf{y}) and if $t \sim q(\mathbf{x})$ then t needs to be produced from this assignment. In other words, $t \sim q(\mathbf{x})$ if $\mathbf{X}_i = \mathbf{y} \xrightarrow{op} X_t = 1$. We now need to show that when $\mathbf{X}_i = \mathbf{y} \xrightarrow{op} X_t = 1$ $Pr(X_t = 1 | \mathbf{X}_D = \mathbf{x}) = 1$.

$$\begin{aligned}
Pr(X_t = 1 | \mathbf{X}_D = \mathbf{x}) &= \sum_{x_i \in \{0,1\}^n} Pr(X_t = 1, \{X_{i_j} = x_{i_j} | j = 1, \dots, n\} | \mathbf{X}_D = \mathbf{x}) \\
&= \sum_{x_i \in \{0,1\}^n} F^{op}(X_t = 1 | \{X_{i_j} = x_{i_j} | j = 1, \dots, n\}) Pr(\{X_{i_j} = x_{i_j} | j = 1, \dots, n\} | \mathbf{X}_D = \mathbf{x}) \\
&= F^{op}(X_t = 1 | X_i = y) Pr(X_i = y | \mathbf{X}_D = \mathbf{x}) \dots \text{(from Eq. (7))} \\
&= F^{op}(X_t = 1 | X_i = y)
\end{aligned}$$

$i \sim q(\mathbf{x})$ when $\mathbf{X}_D = \mathbf{x} \xrightarrow{*} \mathbf{X}_i = \mathbf{y}$ such that $\mathbf{X}_i = \mathbf{y} \xrightarrow{op} X_i = 1$ in which case by Lemma (3.1), $Pr(X_i = 1 | \mathbf{X}_D = \mathbf{x}) = F^{op}(X_i = 1 | \mathbf{X}_i = \mathbf{y}) = 1$. i is not present in $q(\mathbf{x})$ if $\mathbf{X}_D = \mathbf{x} \xrightarrow{*} \mathbf{X}_i = \mathbf{y}$ such that $\mathbf{X}_i = \mathbf{y} \xrightarrow{op} X_i = 0$ in which case, again by Lemma (3.1), $Pr(X_i = 1 | \mathbf{X}_D = \mathbf{x}) = F^{op}(X_i = 1 | \mathbf{X}_i = \mathbf{y}) = 0$. \square

4 Query Execution

In this section, we provide further details as to how we execute queries in probabilistic databases with dependencies. We begin with a description of variable elimination, the inference algorithm used in our current implementation, and discuss various optimizations to perform variable elimination efficiently. After that we discuss ways to store probabilistic databases with correlated tuples and to execute arbitrary Select-Project-Join (SPJ) queries over such data.

4.1 Inference in Graphical Models

Exact probabilistic inference is known to be NP-hard in general [7]. However, many applications provide graphical models with a graph structure that allow efficient probabilistic computation [36]. Variable elimination (VE), also known as *bucket elimination*, [36, 13] is an exact inference algorithm that has the ability to exploit this structure. VE can be used to compute the marginal probabilities of a single random variable from a joint distribution. The main advantages of VE are simplicity and generality.

Computing the marginals of a random variable X requires that we sum out all the other random variables present in the joint distribution (Eq. (2)). Figure 8 (i) shows how VE computes the marginal probability corresponding to $X_3 = x_3$ for the joint distribution described in Figure 3. In Figure 8 (i) we first sum over X_1 producing a new factor $\mu_1(X_2)$ and then sum over X_2 producing a factor $\mu_2(X_3)$ from which we can retrieve the required result.

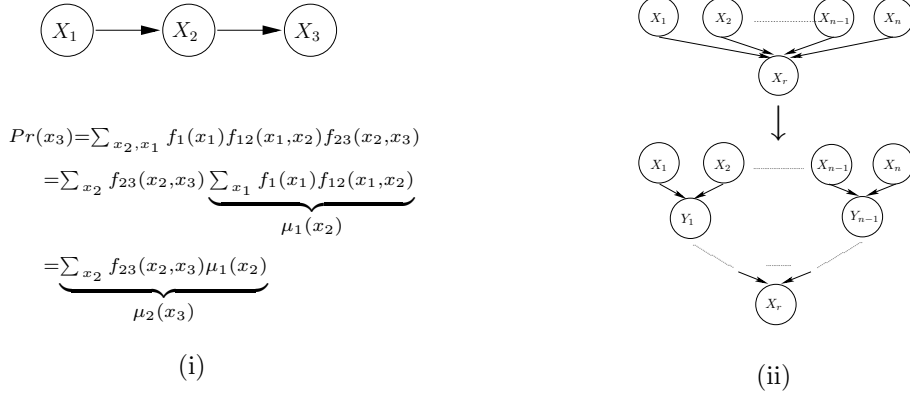


Figure 8: (i) Marginal probability computation for X_3 using variable elimination on the distribution shown in Figure 3. (ii) Transforming a graphical model using decomposability.

The complexity of VE depends on some natural parameters relating to the connectivity of the graph underlying the graphical model corresponding to the joint probability distribution [31]. The inference problem is easy if the graphical model is or closely resembles a tree and the problem becomes progressively harder as the graphical model deviates more from being a tree.

Another possible reason for inference being difficult is due to the presence of factors involving a large number of random variables. Projection and aggregate operations can produce large factors but we can easily reduce the size of these factors by exploiting *decomposability* [37, 31]. This allows us to break any large projection factor into numerous (linear in the number of tuples involved in the projection) constant-sized 3-argument factors. Figure 8 (ii) shows the pictorial representation of this optimization. The top graphical model represents n tuples projecting into one result tuple producing a large factor. The bottom graphical model shows the transformed graphical model that contains $n - 1$ new random variables and consists of only 3-argument factors. We refer the interested reader to Zhang and Poole [37], Rish [31] for details. All aggregate operators (e.g., **sum**, **max** etc.) are also decomposable with the exception of **avg**. To compute **avg**, we first compute **sum** and **count**, both of which are decomposable, and then compute the average.

4.2 Safe plans and Tree-structured Graphical Models

In the context of probabilistic databases with independent base tuples, Dalvi and Suciu [10] identified a class of queries that allow efficient evaluation (queries with *safe plans*). In this subsection, we show that safe plans give rise to inference problems with tree-structured graphical models where running inference is easy. Throughout this subsection we will be concerned with probabilistic databases containing independent base tuples and conjunctive queries (composed of rela-

tional algebra operators σ , \Join and \times) without self-joins.

We first describe how to obtain the probabilistic graphical model for query q issued on probabilistic database D . Once we have run the query and obtained all the factors required for query evaluation construct graph $G_q = (V_q, E_q)$ as follows:

- Construct V_q by introducing a node v for each tuple-based random variable.
- Construct E_q by performing the following steps. For each factor $F^{op}(X_t|\mathbf{X}_S)$ obtained during evaluating operator op on set of tuples S to generate tuple t :
 - Introduce an edge $e \in E_q$ from $X \rightarrow X_c, \forall X \in \mathbf{X}_S$.
 - Annotate the edge e with the operator op .

In our query evaluation procedure, we are usually concerned with performing inference to obtain the probability specific to result tuple t . Thus, we are more concerned with the probabilistic graphical model specific to t which we will denote by $G_q(t) = (V_q(t), E_q(t))$. We can extract this subgraph as follows:

- $V_q(t) = \{v|v \in V_q, v \rightsquigarrow X_t\}$ where X_t denotes the node corresponding to the random variable for t and $v_i \rightsquigarrow v_j$ is true if there is a path from v_i to v_j . Note that $X_t \in V_q(t)$.
- $E_q(t) = \{e = (v_i, v_j)|e \in E_q, v_i, v_j \in V_q(t)\}$

As discussed earlier, $G_q(t)$ provides easy inference if it resembles a tree [31].

We now discuss some simple lemmas related to the graphical model for tuple t produced by (conjunctive) query q (without self-joins) on probabilistic database D .

Lemma 4.1. *For two nodes $v_i, v_j \in V_q(t)$ such that the corresponding probabilistic tuples belong to the same probabilistic relation, there must exist:*

- edge e_i on the path from v_i to X_t such that e_i is annotated with \Join
- edge e_j on the path from v_j to X_t such that e_j is annotated with \Join

Proof. The lemma follows from the behaviour of relational algebra operators. Basically, it says that q should include a projection operator for it to involve two probabilistic tuples from the same relation in $G_q(t)$. \square

Lemma 4.2. *If $G_q(t)$ is not a tree then it must contain nodes $v_i, v_j, v_k \in V_q(t)$ with edges $e_j, e_k \in E_q(t)$ such that $e_j = v_i \rightarrow v_j$ and $e_k = v_i \rightarrow v_k$.*

Proof. Since $G_q(t)$ is not a tree then there must exist a pair of vertices v and w (w could be X_t) such there is more than one path connecting them otherwise $G_q(t)$ is a tree (by definition).


Let t_1 denote the first path from v to w which follows the sequence of nodes $x_0 = v \rightarrow x_1 \rightarrow x_2 \rightarrow \dots \rightarrow x_n = w$ and let t_2 denote the second path $y_0 = v \rightarrow y_1 \rightarrow y_2 \dots \rightarrow y_m = w$. Let x_a and y_b denote the first vertices in which the node sequences t_1 and t_2 differ, respectively. Therefore, set $v_i = x_{a-1} = y_{b-1}, v_j = x_a, v_k = y_b$. \square

Lemma 4.3. *If $G_q(t)$ is not tree-structured then q must involve a projection operator \amalg .*

Proof. By Lemma (4.2), $G_q(t)$ must contain a pair of edges $v_i \xrightarrow{op_1} v_j, v_i \xrightarrow{op_2} v_k \in E_q(t)$ where $v_i, v_j, v_k \in V_q(t)$. We are going to show that if q is restricted to σ and \times operators then no assignment to op_1 and op_2 are possible.

First, it should be clear that neither op_1 nor op_2 can be σ . The operation we defined for σ^p (Figure 7) makes sure that if $v_i \xrightarrow{\sigma} v_j \in E_q(t)$ then there exist no other edge $e \in E_q(t)$ such that $e = v_i \rightarrow v_k$.

The only remaining case to consider is $op_1 = op_2 = \times$. Thus $G_q(t)$ consists of the following subgraph:



The only possibility of this happening is when v_i corresponds to a tuple from one (intermediate) relation that joins with two tuples from another (intermediate) relation corresponding to $v_{j'}$ and $v_{k'}$. But notice that now v_j and v_k belong to the same (intermediate) relation. Thus, by Lemma (4.1) there has to be a projection operator in q for $op_1 = op_2 = \times$ to be true. Thus q has to involve a projection operator if $G_q(t)$ is not tree-structured. \square

We next provide the main theorem connecting safe plans to tree-structured graphical models. Before that we need to first give some background on safe plans. The analysis for safe plans involves tuple-specific event ids. Each independent tuple in the database is assigned a unique event id and each intermediate tuple t is associated with the complex event expression formed by the combination of event ids of the base tuples involved in the generation of t .

Theorem 4.1. *For any safe plan [10] q (representing a conjunctive query without self-joins) and a result tuple t , $G_q(t)$ is tree-structured.*

Proof. (By contradiction). Let $G_q(t)$ not be tree-structured. Thus, by Lemma (4.2), $G_q(t)$ must contain a pair of edges $v_i \xrightarrow{op_1} v_j, v_i \xrightarrow{op_2} v_k \in E_q(t)$ where $v_i, v_j, v_k \in V_q(t)$ and, by Lemma (4.3), q contains atleast one projection operator. We next attempt to locate this projection operator.

First, notice that v_i has a fan-out ≥ 2 and this can happen only when $op_1 = op_2 = \times$, assigning either op_1 or op_2 either of σ or \amalg will not achieve this. Thus $op_1 = op_2 = \times$. Thus, by Lemma (4.1), there exists edges e_j on the

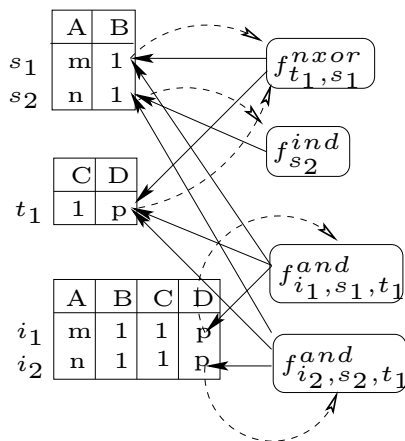


Figure 9: Partitions for the “nxor” dependency shown in Figure 2 (i) with the partitions introduced by the join shown in Figure 5.

path $v_j \rightsquigarrow X_t$ and e_k on $v_k \rightsquigarrow X_t$ both annotated with \square . In other words, if we represent q as a query plan tree then there exists a \square above the \times corresponding to op_1 and op_2 .

Since q is a safe plan, by Theorem 4 in Dalvi and Suciu [11], every operator in q is safe. We now show that the projection operator located above is not safe. Since $v_i \rightarrow v_j$, v_j is associated with a complex event expression that contains v_i ’s complex event expression as a sub-expression. Similarly, v_k also contains v_i ’s complex event expression. Thus the projection operator after this join projects two tuples with complex events that are *not* independent rendering it unsafe (see Dalvi and Suciu [11] Theorem 3 bullet 3). Thus we have a contradiction and hence $G_q(t)$ is tree-structured. \square

4.3 Representing Probabilistic Relations

Earlier approaches represented probabilistic relations by storing uncertainty with each tuple in isolation. This approach is inadequate for our purposes since the same tuple can be involved in multiple dependencies. In our implementation, we store the data and uncertainty parts separately. The tuples are stored as part of relations. To store uncertainty, we introduce the concept of a *partition*. **A partition consists of a factor and a set of references to the tuples whose random variables form the arguments to that factor.** Besides, each tuple t in a relation also contains a list of pointers to the partitions that contain references to the tuple. The relations and partitions together form a doubly-linked data structure so that we can move from one to the other. Figure 9 shows the arrangement of relations and partitions for the database with the “nxor” dependency (Figure 4 (iii)) where dotted lines represent the pointers from tuples to partitions and solid lines show the references.

4.4 Query Execution Steps

The naive way of executing query q would be to execute it and construct the full probabilistic graphical model G_q (for a brief description of how to construct G_q see Section 4.2). One of the concerns with this approach is that the full graphical model (involving all the result tuples) might exceed available memory. In this subsection, we describe some simple strategies to handle this issue.

In most cases, the probabilistic graphical model G_q will be a collection of disjoint subgraphs. In such cases we can improve the situation by extracting one disjoint subgraph at a time and computing the probabilities of the result tuples present in that subgraph. Constructing G_q is not an issue. One can follow the standard database approach of taking each operator from the query and reading in the relations they act upon, constructing new nodes corresponding to intermediate tuples and new edges for G_q and storing these on disk. The main implementation issue here is to find disjoint subgraphs from G_q without actually reading the whole of G_q into memory.

One simple strategy to extract disjoint subgraphs from G_q is to use external memory graph algorithms [33] (e.g., depth-first search to find connected components) for these tasks. In the case when the database consists of independent base tuples, the only operations that can connect components of G_q belonging to different relations are join operations. In this case it is useful to perform all join operations together so one can concentrate on determining disjoint components produced by tuples specific to the joins. Besides these tricks we also employ standard tricks of pushing projections and selections down in the query plan.

Here is a summary of the steps we follow for SPJ queries:

1. *Perform early selections:* We push down and execute as many selections as possible.
2. *Perform early projections:* We also push down as many projections as possible.
3. *Join phase:* In this phase, we perform all the join operations and perform them as part of one multi-join operation, creating partitions and intermediate relations (Figure 9 shows the partitions introduced due to a join operation on the database shown in Figure 4 (iii)).
4. *Perform final projection:* Here we project onto the result attributes specified by the user in the query.
5. *Probability computation:* For each result tuple, we recursively collect all partitions required to perform inference and compute the required marginal probability.

4.5 Further Optimizations

When exact probabilistic inference turns out to be too expensive we have the flexibility of switching to approximate inference techniques. Just like exact

inference, approximate inference is also known to be NP-hard [9]. However there exist a fairly large variety of approximate inference algorithms that perform well in practice in a variety of cases each varying in speed and accuracy e.g., Markov Chain Monte Carlo techniques [20], Variational Methods [26] etc. Depending on the user’s requirements we can easily switch between inference algorithms in our approach.

5 Experimental Study

In this section, we present an experimental evaluation demonstrating the need for modeling tuple correlations, and the effectiveness of our techniques at modeling such correlations and evaluating queries over them. This evaluation was done using a prototype system that we are currently building on top of the Apache Derby Java DBMS [23]. Our system supports the query execution strategies discussed in the previous section.

Following Fuhr and Rolleke [17] and Dalvi and Suciu [10], we generate probabilistic relations by issuing *similarity predicates* over deterministic relations. Briefly, given a predicate of the form $R.a \approx k$, where a is a string attribute, and k is a string constant, the system assigns a probability to each tuple t , based on how *similar* $t.a$ is to k . Similarity is ascertained by computing the distance between $t.a$ and k using a standard distance function (following prior work [10], we used *3-gram distance* [32]). This distance is then converted to a posterior probability by assuming that the distance is normally distributed with mean 0, and variance σ

$$e^{-\text{distance}^2(t.a,k)/\sigma^2}$$

where σ is a parameter to the system. The details can be found in Dalvi and Suciu [10].

5.1 Need for Modeling Dependencies

Consider a *publications* database containing two relations: (1) PUBLS(PID, Title), and (2) AUTHS(PID, Name), where PID is the unique publication id, and consider the task of retrieving all publications with title y written by an author with name x . Assuming that the user is not sure of the spellings x and y , we might use the following query to perform the above task:

$$\prod_{\text{Title}}(\sigma_{\text{Name} \approx x}(\text{AUTHS}) \bowtie \sigma_{\text{Title} \approx y}(\text{PUBLS}))$$

As discussed above, the similarity predicates will cause both the relations to be converted into probabilistic relations, AUTHS^p and PUBLS^p . However, note that AUTHS^p contains natural mutual exclusion dependencies with respect to this query. Since the user is looking for publications by a single author with name x , it is *not possible for x to match two AUTHS^p tuples corresponding to the same publication in the same possible world*. Thus, any two AUTHS^p tuples with the same PID exhibit a mutual exclusion dependency, and a possible world containing both of them should be assigned zero probability.

Title	Title
Reinforcement learning with hidden states (by L. Lin, T. Mitchell)	Reinforcement learning with hidden states (by L. Lin, T. Mitchell)
Feudal Reinforcement Learning (by C. Atkeson, P. Dayan, ...)	Feudal Reinforcement Learning (by C. Atkeson, P. Dayan, ...)
Reasoning (by C. Bereiter, M. Scardamalia)	Reasoning (by C. Bereiter, M. Scardamalia)
...	...

(i) *MUTEX_DB* results at $\sigma = 10, 50, 100$

Title	Title
Feudal Reinforcement Learning (by C. Atkeson, P. Dayan, ...)	Decision making and problem solving (G. Dantzig, R. Hogarth, ...)
Decision making and problem solving (G. Dantzig, R. Hogarth, ...)	HERMES: A heterogeneous reasoning and mediator system (by S. Adali, A. Brink, ...)
Multimodal Learning Interfaces (by U. Bub, R. Houghton, ...)	Induction and reasoning from cases (by K. Althoff, E. Auriol, ...)
...	...

(ii) *IND_DB* results at $\sigma = 10$

Title	Title
Feudal Reinforcement Learning (by C. Atkeson, P. Dayan, ...)	Decision making and problem solving (G. Dantzig, R. Hogarth, ...)
Decision making and problem solving (G. Dantzig, R. Hogarth, ...)	HERMES: A heterogeneous reasoning and mediator system (by S. Adali, A. Brink, ...)
Multimodal Learning Interfaces (by U. Bub, R. Houghton, ...)	Induction and reasoning from cases (by K. Althoff, E. Auriol, ...)
...	...

(iii) *IND_DB* results at $\sigma = 50$

Title	Title
Feudal Reinforcement Learning (by C. Atkeson, P. Dayan, ...)	Decision making and problem solving (G. Dantzig, R. Hogarth, ...)
Decision making and problem solving (G. Dantzig, R. Hogarth, ...)	HERMES: A heterogeneous reasoning and mediator system (by S. Adali, A. Brink, ...)
Multimodal Learning Interfaces (by U. Bub, R. Houghton, ...)	Induction and reasoning from cases (by K. Althoff, E. Auriol, ...)
...	...

(iv) *IND_DB* results at $\sigma = 100$

Figure 10: Top three results for a similarity query: (i) shows results from *MUTEX_DB*; (ii), (iii) and (iv) show results from *IND_DB*.

To illustrate the drawbacks of ignoring these mutual exclusion dependencies, we ran the above query with $x = \text{“T. Michel”}$ and $y = \text{“Reinforcement Learning hidden state”}$ on two probabilistic databases, one assuming complete independence among tuples (*IND_DB*) and another that models the dependencies (*MUTEX_DB*). We ran the query on an extraction of 860 publications from the real-world CiteSeer dataset [19]. We report results across various settings of σ .

Figure 10 shows the top three results obtained from the two databases at three different settings of σ (we also list the author names to aid the reader’s understanding). *MUTEX_DB* returns intuitive and similar results at all three values of σ . *IND_DB* returns reasonable results only at $\sigma = 10$, whereas at $\sigma = 50, 100$ it returns very odd results (“Decision making and problem solving” does not match the string “Reinforcement Learning hidden state” very closely and yet it is assigned the highest rank at $\sigma = 100$). Figure 11 shows the cumulative recall graph for *IND_DB* for various values of σ , where we plot the fraction of the top N results returned by *MUTEX_DB* that were present in the top N results returned by *IND_DB*. As we can see, at $\sigma = 50$ and 100 , *IND_DB* exhibits poor recall.

Figure 10 shows that *IND_DB* favors publications with long author lists. This does not affect the results at low values of σ ($=10$) because, in that case, we use a “peaked” gaussian which assigns negligible probabilities to possible worlds with multiple AUTHS^p from the same publication. At larger settings of σ ,




Figure 11: Cumulative recall graph comparing results of *IND_DB* and *MUTEX_DB* for $\sigma = 10, 50, 100$.

however, these possible worlds are assigned larger probabilities and *IND_DB* returns poor results. *MUTEX_DB* assigns these possible worlds zero probabilities by modeling dependencies on the base tuples.

We would like to note that, although setting the value of σ carefully may have resulted in a good answer for *IND_DB* in this case, choosing σ is not easy in general and depends on various factors such as user preferences, distributions of the attributes in the database etc [10]. Modeling mutual exclusion dependencies explicitly using our approach naturally alleviates this problem.

5.2 Scalability

Next we study the scalability of our proposed query execution strategies using a randomly generated TPC-H dataset of size 10MB. For simplicity, we assume complete independence among the base tuples (though the intermediate tuples may still be correlated).

Figure 12 shows the execution times on TPC-H queries Q2 to Q8 (modified to remove the top-level aggregations). The first bar on each query indicates the time it took for our implementation to run the full query including all the database operations and the probabilistic computations. The second bar on each query indicates the time it took to run only the database operations using our Java implementation. Here are the summary of the results:

- As we can see in Figure 12, for most queries the additional cost of probability computations is comparable to the cost of normal query processing.
- The two exceptions are Q3 and Q4 which severely tested our probabilistic inference engine. By removing the aggregate operations, Q3 resulted in a

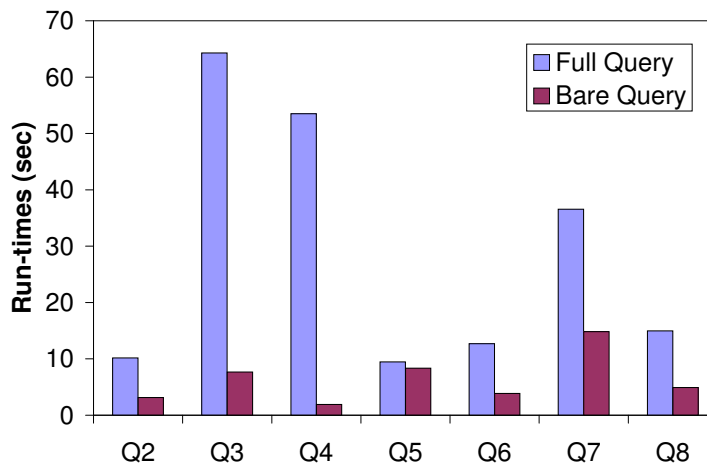


Figure 12: Run-times on TPC-H data.

relation of size in excess of 60,000 result tuples. Although Q4 resulted in a very small relation, each result tuple was associated with a probabilistic graphical model of size exceeding 15,000 random variables. Each of these graphical models are fairly sparse but book-keeping for such large data structures took a significant amount of time. We expect that a more careful and efficient implementation to significantly improve these run times.

- Q7 and Q8 are queries without safe plans [10] yet their run-times are surprisingly fast. By taking a closer look we notice that both these queries gave rise to tree structured graphical models justifying our belief that there are queries that lend to efficient probabilistic inference even though they might not have safe plans.

5.3 Aggregate operations

Our system also naturally supports efficient computation of a variety of aggregate operators over probabilistic relations. As discussed in Section 4, we can perform these operations efficiently due to aggregate operators being *decomposable* [37, 31]. Figure 13 shows the result of running an *average* query over a synthetically generated dataset containing 500 tuples. As we can see, the final result can be a fairly complex probability distribution, which is quite common for aggregate operations. Effective computation of aggregates over large probabilistic databases is an open research problem that we plan to study in future.




Figure 13: AVG aggregate computed over 500 randomly generated tuples with attribute values ranging from 1 to 5.

6 Related Work

There has been much work on managing probabilistic, uncertain, incomplete, and/or fuzzy data in database systems (see e.g. [25, 2, 27, 28, 21, 17, 3, 5, 10, 34]). With a rapid increase in the number of application domains such as data integration, pervasive computing etc., where uncertain data arises naturally, this area has seen renewed interest in recent years (see [18] for a survey of the ongoing research).

We will briefly discuss some of the closely related work in the area of *probabilistic* data management. Previous work in this area can be differentiated based on (a) whether probabilities are associated with *attributes* or with *tuples*, (b) whether the resulting relations are in the first normal form (which is highly desirable given the complexity of managing and querying data which is not in 1NF), and (c) whether the correlations typically present in real-world data can be modeled. Barbara et al. [2] propose an approach that associates probabilities with attributes and can model arbitrary correlations between attributes *within a single tuple*. However the resulting relations are not in 1NF and further the semantics of some of the query operators are messy, both of which seriously limit the applicability of their approach. More recently, Choenni et al. [6] discuss conceptually how to make the query semantics more consistent through use of Dempster-Schafer theory.

Cavallo and Pittarelli [4] and Dey and Sarkar [15] propose and study tuple-level uncertainty models that explicitly capture *mutual exclusivity*. More recently, Andritsos et al. [1] use a similar basic model to capture *dirty* data, and develop query evaluation algorithms for the same. In a series of papers, Fuhr and Rolleke [17] propose using tuple-level probabilities to model uncertain data

in *information retrieval* context, and present the intensional and extensional query evaluation techniques discussed in Section 2. Extending this work, Dalvi and Suciu [10] define *safe query plans* to be those for which extensional and intensional query evaluation produces identical results, and show how to generate a safe query plan for a query if one exists. Tuple independence is assumed for most of the work by both these groups. Lakshmanan et al. [27] attempt to combine these different approaches by associating *probability intervals* with tuples in their *ProbView* system. Their model also supports various conjunction and disjunction *strategies* that allow a limited encoding of tuple interdependences.

Cheng et al. [5], Xia et al. [35] associate (continuous) probability distributions with attributes, and propose several query evaluation and indexing techniques over such data. Trio [34, 12] aims to provide a unified treatment of *data uncertainty*, *data lineage* and *data accuracy* in a single system. They also study the issues of completeness and closure under various alternative models for representing uncertainty, though their notion of uncertainty does not currently include probabilities. In recent years, there has also been renewed interest in probabilistic extensions of object-relational [16] and XML data models [29, 24].

7 Conclusions

There is an increasing need for database solutions for efficiently managing and querying uncertain data exhibiting complex correlation patterns. In this paper, we presented a simple and intuitive framework, based on probabilistic graphical models, for explicitly modeling correlations among tuples in a probabilistic database, and developed strategies for efficiently executing SQL queries over such data. Our experimental evaluation illustrates both the necessity of modeling tuple correlations, and the effectiveness of our techniques at representing and querying correlated datasets. Our research so far has raised several interesting challenges that we plan to pursue in future. Although conceptually our approach allows for capturing arbitrary tuple correlations, exact query evaluation over large datasets exhibiting complex correlations may not always be feasible. We plan to develop approximate query evaluation techniques that can be used in such cases. We are also planning to develop disk-based query evaluation algorithms so that our techniques can scale to very large datasets.

References

- [1] Periklis Andritsos, Ariel Fuxman, and Renee J. Miller. Clean answers over dirty databases. In *International Conference on Data Engineering*, 2006.
- [2] Daniel Barbara, Hector Garcia-Molina, and Daryl Porter. The management of probabilistic data. In *IEEE Transactions of Knowledge Data Engineering*, 1992.

- [3] Bill P. Buckles and Frederick E. Petry. A fuzzy model for relational databases. *International Journal of Fuzzy Sets and Systems*, 1982.
- [4] Roger Cavallo and Michael Pittarelli. The theory of probabilistic databases. In *International Conference on Very Large Data Bases*, 1987.
- [5] Reynold Cheng, Dmitri Kalashnikov, and Sunil Prabhakar. Evaluating probabilistic queries over imprecise data. In *International Conference on Management of Data.*, 2003.
- [6] Sunil Choenni, Henk Ernst Blok, and Erik Leertouwer. Handling uncertainty and ignorance in databases: A rule to combine dependent data. In *Database Systems for Advanced Applications*, 2006.
- [7] Gregory F. Cooper. The computational complexity of probabilistic inference using bayesian belief networks. *Artificial Intelligence*, 1990.
- [8] Robert G. Cowell, Steffen L. Lauritzen, and David J. Spiegelhater. *Probabilistic Networks and Expert Systems*. Springer, New York, 1999.
- [9] Paul Dagum and Michael Luby. Approximate probabilistic reasoning in bayesian belief networks is np-hard. *Artificial Intelligence*, 1993.
- [10] Nilesh Dalvi and Dan Suciu. Efficient query evaluation on probabilistic databases. In *International Conference on Very Large Data Bases*, 2004.
- [11] Nilesh Dalvi and Dan Suciu. Efficient query evaluation on probabilistic databases. *International Journal on Very Large Data Bases*, 2006.
- [12] Anish Das Sarma, Omar Benjelloun, Alon Halevy, and Jennifer Widom. Working models for uncertain data. In *International Conference on Data Engineering*, 2006.
- [13] Rina Dechter. Bucket elimination: A unifying framework for probabilistic inference. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence.*, 1996.
- [14] Amol Deshpande, Carlos Guestrin, Sam Madden, Joseph M. Hellerstein, and Wei Hong. Model-driven data acquisition in sensor networks. In *International Conference on Very Large Data Bases*, 2004.
- [15] Debabrata Dey and Sumit Sarkar. A probabilistic relational model and algebra. *ACM Transactions on Database Systems.*, 1996.
- [16] Thomas Eiter, Thomas Lukasiewicz, and Michael Walter. A data model and algebra for probabilistic complex values. *Annals of Mathematics and Artificial Intelligence*, 2001.
- [17] Norbert Fuhr and Thomas Rolleke. A probabilistic relational algebra for the integration of information retrieval and database systems. *ACM Transactions on Information Systems*, 1997.

- [18] Minos Garofalakis and Dan Suciu, editors. *IEEE Data Engineering Bulletin Special Issue on Probabilistic Data Management*. March 2006.
- [19] C. Lee Giles, Kurt Bollacker, and Steve Lawrence. Citeseer: An automatic citation indexing system. In *Conference on Digital Libraries*, 1998.
- [20] Walter R. Gilks, Sylvia Richardson, and David J. Spiegelhalter. *Markov Chain Monte Carlo in Practice*. Chapman & Hall, 1996.
- [21] Gosta Grahne. Horn tables - an efficient tool for handling incomplete information in databases. In *Symposium on Principles of Database Systems*, 1989.
- [22] Joseph Halpern. An analysis of first-order logics for reasoning about probability. *Artificial Intelligence*, 1990.
- [23] <http://db.apache.org/derby>. The Apache Derby project.
- [24] Edward Hung, Lise Getoor, and V. S. Subrahmanian. PXML: A probabilistic semistructured data model and algebra. In *International Conference on Data Engineering*, 2003.
- [25] Tomasz Imielinski and Jr. Witold Lipski. Incomplete information in relational databases. *Journal of the ACM.*, 1984.
- [26] Michael I. Jordan, Zoubin Ghahramani, Tommi S. Jaakkola, and Lawrence K. Saul. An introduction to variational methods for graphical models. *Machine Learning*, 1999.
- [27] Laks V. S. Lakshmanan, Nicola Leone, Robert Ross, and V. S. Subrahmanian. Probview: a flexible probabilistic database system. *ACM Transactions on Database Systems.*, 1997.
- [28] Suk Kyoong Lee. An extended relational database model for uncertain and imprecise information. In *International Conference on Very Large Data Bases*, 1992.
- [29] Andrew Nierman and H. V. Jagadish. Protodb: Probabilistic data in XML. In *International Conference on Very Large Data Bases*, 2002.
- [30] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, 1988.
- [31] Irina Rish. *Efficient Reasoning in Graphical Models*. PhD thesis, University of California, Irvine, 1999.
- [32] Esko Ukkonen. Approximate string matching with q-grams and maximal matches. In *Theoretical Computer Science*, 1992.
- [33] Jeffrey Scott Vitter. External memory algorithms and data structures: Dealing with massive data. *ACM Computing Surveys*, 2001.

- [34] Jennifer Widom. Trio: A system for integrated management of data, accuracy, and lineage. In *Proceedings of the Biennial Conference on Innovative Data Systems Research*, 2005.
- [35] Yuni Xia, Sunil Prabhakar, Shan Lei, Reynold Cheng, and Rahul Shah. Indexing continuously changing data with mean-variance tree. In *ACM Symposium on Applied Computing*, 2005.
- [36] Nevin Lianwen Zhang and David Poole. A simple approach to bayesian network computations. In *Canadian Conference on Artificial Intelligence*, 1994.
- [37] Nevin Lianwen Zhang and David Poole. Exploiting causal independence in bayesian network inference. *Journal of Artificial Intelligence Research*, 1996.