

A Dichotomy

for Queries with **Negation**

in **Probabilistic** Databases



Dan Olteanu


Joint work with Robert Fink

Uncertainty in Computation

Simons Institute for the Theory of Computing

Berkeley Oct 5, 2016

Outline



Probabilistic Databases 101

The Dichotomy

The Interesting but Hard Queries

The Easy Queries

Leftovers

Tuple-Independent Probabilistic Databases

Tuple-independent database of n tuples $(t_i)_{i \in [n]}$:

- Each tuple t_i associated with an independent Boolean random variable x_i .
- $P(x_i = \text{true})$ gives the probability that t_i exists in the database.

Possible-worlds semantics:

- Each possible world defined by an assignment θ of the variables $(x_i)_{i \in [n]}$:
 - ▶ It consists of all tuples t_i for which $\theta(x_i) = \text{true}$.
 - ▶ It has probability $P(\theta) = \prod_{i \in [n]} P(x_i = \theta(x_i))$.
- A tuple-independent database with n tuples has 2^n possible worlds.

Relational Algebra

Popular database query language since Codd times.

- Algebra carrier: set of all finite relations
- Algebra operations: π (**projection**), \times (**Cartesian product**), $-$ (**set difference**), \bowtie (**join**), σ (selection), \cup (set union), δ (renaming)
- As expressive as domain relational calculus (RC)

In this talk: **Relational algebra fragment 1RA⁻**

- Included: Equality joins, selections, projections, difference
- Excluded: Repeating relation symbols, unions

Relational Algebra

Popular database query language since Codd times.

- Algebra carrier: set of all finite relations
- Algebra operations: π (**projection**), \times (**Cartesian product**), $-$ (**set difference**), \bowtie (**join**), σ (selection), \cup (set union), δ (renaming)
- As expressive as domain relational calculus (RC)

In this talk: **Relational algebra fragment 1RA⁻**

- Included: Equality joins, selections, projections, difference
- Excluded: Repeating relation symbols, unions

Examples of (Boolean) 1RA⁻ queries:

- Are there combinations of tuples in (R, T) that are not in (U, V) ?

$$\begin{aligned} \pi_{\emptyset} [(R(A) \times T(B)) - (U(A) \times V(B))] \\ \exists_A \exists_B [(R(A) \wedge T(B)) \wedge \neg (U(A) \wedge V(B))] \quad (\text{in RC}) \end{aligned}$$

- Does relation S “hold hands” with both R and T ?

$$\begin{aligned} \pi_{\emptyset} [R(A) \bowtie S(A, B) \bowtie T(B)] \\ \exists_A \exists_B [R(A) \wedge S(A, B) \wedge T(B)] \quad (\text{in RC}) \end{aligned}$$

The Query Evaluation Problem

For any Boolean 1RA⁻ query Q and tuple-independent database D :

Compute the probability that Q is true in a random world of D .

The case of non-Boolean queries can be reduced to the Boolean case.

We are interested in the *data complexity* of this problem.

- Fix the query Q and take the database D as input.

Outline



Probabilistic Databases 101

The Dichotomy

The Interesting but Hard Queries

The Easy Queries

Leftovers

Data complexity of any $1RA^-$ query Q in tuple-independent databases:

- Polynomial time if Q is hierarchical and $\#P$ -hard otherwise.

Hierarchical 1RA⁻ Queries

(Boolean) 1RA⁻ query Q is *hierarchical* if

- For every pair of distinct query variables A and B in Q ,
- there is no triple of relation symbols R , S , and T in Q such that:
- R has A but not B , S has both A and B , and T has B but not A .

Hierarchical 1RA⁻ Queries

(Boolean) 1RA⁻ query Q is *hierarchical* if

- For every pair of distinct query variables A and B in Q ,
- there is no triple of relation symbols R , S , and T in Q such that:
- R has A but not B , S has both A and B , and T has B but not A .



$R(A)$


$S(A, B)$

$T(B)$

Hierarchical 1RA⁻ Queries

(Boolean) 1RA⁻ query Q is *hierarchical* if



- For every pair of distinct query variables A and B in Q ,
- there is no triple of relation symbols R , S , and T in Q such that:
- R has A but not B , S has both A and B , and T has B but not A .



Hierarchical 1RA⁻ Queries

(Boolean) 1RA⁻ query Q is *hierarchical* if

- For every pair of distinct query variables A and B in Q ,
- there is no triple of relation symbols R , S , and T in Q such that:
- R has A but not B , S has both A and B , and T has B but not A .



Examples

Hierarchical queries:

- $\pi_{\emptyset}[(R(A) \bowtie S(A, B)) - T(A, B)]$
- $\pi_{\emptyset}[(R(A) \times T(B)) - (U(A) \times V(B))]$
- $\pi_{\emptyset}[(M(A) \times N(B)) - [(R(A) \times T(B)) - (U(A) \times V(B))]]$
- $\pi_{\emptyset}[\pi_A[M(A) \times N(B)] - \pi_A[(R(A) \times T(B)) - (U(A) \times V(B))]]$

Examples

Hierarchical queries:

- $\pi_{\emptyset}[(R(A) \bowtie S(A, B)) - T(A, B)]$
- $\pi_{\emptyset}[(R(A) \times T(B)) - (U(A) \times V(B))]$
- $\pi_{\emptyset}[(M(A) \times N(B)) - [(R(A) \times T(B)) - (U(A) \times V(B))]]$
- $\pi_{\emptyset}[\pi_A[M(A) \times N(B)] - \pi_A[(R(A) \times T(B)) - (U(A) \times V(B))]]$

Non-hierarchical queries:

- $\pi_{\emptyset}[R(A) \bowtie S(A, B) \bowtie T(B)]$
- $\pi_{\emptyset}[\pi_B(R(A) \bowtie S(A, B)) - T(B)]$
- $\pi_{\emptyset}[T(B) - \pi_B(R(A) \bowtie S(A, B))]$
- $\pi_{\emptyset}[X(A) \bowtie [R(A) - \pi_A(T(B) \bowtie S(A, B))]]$

Outline



Probabilistic Databases 101

The Dichotomy

The Interesting but Hard Queries

The Easy Queries

Leftovers

Hardness Proof Idea

Reduction from #P-hard model counting problem for positive bipartite DNF:

- Given a non-hierarchical 1RA⁻ query Q and
- Any positive bipartite DNF formula Ψ over disjoint sets \mathbf{X} and \mathbf{Y} of random variables.
- $\#\Psi$ can be computed using linearly many calls to an oracle for $P(Q)$, where Q is evaluated on tuple-independent databases of sizes linear in the size of Ψ .

A Simple Case

Input formula and query:

- $\Psi = x_1y_1 \vee x_1y_2 \vee x_2y_1$ over sets $\mathbf{X} = \{x_1, x_2\}$, $\mathbf{Y} = \{y_1, y_2\}$
- $Q = \pi_{\emptyset} [R(A) \bowtie S(A, B) \bowtie T(B)]$

Construct a database D such that Ψ becomes the grounding of Q wrt D :

- Column Φ holds formulas over random variables.
 - ▶ We use \top for *true* and \perp for *false*.
- Variables also used as constants for A and B .
- $S(x_i, y_j, \top)$: $x_i y_j$ is a clause in Ψ .
- $R(x_i, x_i)$ and $T(y_j, y_j)$: x_i is a variable in \mathbf{X} and y_j is a variable in \mathbf{Y} .

R	T	S	$R \bowtie S \bowtie T$	$\pi_{\emptyset} [R \bowtie S \bowtie T]$
$A \ \Phi$	$B \ \Phi$	$A \ B \ \Phi$	$A \ B \ \Phi$	Φ
$x_1 \ x_1$	$y_1 \ y_1$	$x_1 \ y_1 \ \top$	$x_1 \ y_1 \ x_1y_1$	$(\) \quad \Psi$
$x_2 \ x_2$	$y_2 \ y_2$	$x_1 \ y_2 \ \top$	$x_1 \ y_2 \ x_1y_2$	
		$x_2 \ y_1 \ \top$	$x_2 \ y_1 \ x_2y_1$	

A Simple Case

Input formula and query:

- $\Psi = x_1y_1 \vee x_1y_2 \vee x_2y_1$ over sets $\mathbf{X} = \{x_1, x_2\}$, $\mathbf{Y} = \{y_1, y_2\}$
- $Q = \pi_{\emptyset} [R(A) \bowtie S(A, B) \bowtie T(B)]$

Construct a database D such that Ψ becomes the grounding of Q wrt D :

- Column Φ holds formulas over random variables.
 - ▶ We use \top for *true* and \perp for *false*.
- Variables also used as constants for A and B .
- $S(x_i, y_j, \top)$: $x_i y_j$ is a clause in Ψ .
- $R(x_i, x_i)$ and $T(y_j, y_j)$: x_i is a variable in \mathbf{X} and y_j is a variable in \mathbf{Y} .

R	T	S	$R \bowtie S \bowtie T$	$\pi_{\emptyset} [R \bowtie S \bowtie T]$
$A \ \Phi$	$B \ \Phi$	$A \ B \ \Phi$	$A \ B \ \Phi$	Φ
$x_1 \ x_1$	$y_1 \ y_1$	$x_1 \ y_1 \ \top$	$x_1 \ y_1 \ x_1y_1$	$(\) \quad \Psi$
$x_2 \ x_2$	$y_2 \ y_2$	$x_1 \ y_2 \ \top$	$x_1 \ y_2 \ x_1y_2$	
		$x_2 \ y_1 \ \top$	$x_2 \ y_1 \ x_2y_1$	

This is the only minimal hard pattern for *positive* 1RA⁻ queries!

A Surprising Case

Input formula and query:

- $\Psi = x_1 y_1 \vee x_1 y_2$ over sets $\mathbf{X} = \{x_1\}$, $\mathbf{Y} = \{y_1, y_2\}$
- $Q = \pi_{\emptyset} [R(A) - \pi_A (T(B) \bowtie S(A, B))]$

Construct a database D such that Ψ becomes the grounding of Q wrt D :

- $S(a, b, \top)$: Clause a has variable b in Ψ .
- $R(a, \top)$ and $T(b, \neg b)$: a is a clause and b is a variable in Ψ .

R	T	S	$T \bowtie S$	$\pi_A(T \bowtie S)$	$R - \pi_A(T \bowtie S)$
$A \ \phi$	$B \ \phi$	$A \ B \ \phi$	$A \ B \ \phi$	$A \ \phi$	$A \ \phi$
1 \top	$x_1 \ \neg x_1$	1 $x_1 \ \top$	1 $x_1 \ \neg x_1$	1 $\neg x_1 \vee \neg y_1$	1 $x_1 y_1$
2 \top	$y_1 \ \neg y_1$	1 $y_1 \ \top$	1 $y_1 \ \neg y_1$	2 $\neg x_1 \vee \neg y_2$	2 $x_1 y_2$
	$y_2 \ \neg y_2$	2 $x_1 \ \top$	2 $x_1 \ \neg x_1$		
		2 $y_2 \ \top$	2 $y_2 \ \neg y_2$		

A Surprising Case

Input formula and query:

- $\Psi = x_1 y_1 \vee x_1 y_2$ over sets $\mathbf{X} = \{x_1\}$, $\mathbf{Y} = \{y_1, y_2\}$
- $Q = \pi_{\emptyset} [R(A) - \pi_A (T(B) \bowtie S(A, B))]$

Construct a database D such that Ψ becomes the grounding of Q wrt D :

- $S(a, b, \top)$: Clause a has variable b in Ψ .
- $R(a, \top)$ and $T(b, \neg b)$: a is a clause and b is a variable in Ψ .

R	T	S	$T \bowtie S$	$\pi_A(T \bowtie S)$	$R - \pi_A(T \bowtie S)$
$A \ \phi$	$B \ \phi$	$A \ B \ \phi$	$A \ B \ \phi$	$A \ \phi$	$A \ \phi$
1 \top	$x_1 \ \neg x_1$	1 $x_1 \ \top$	1 $x_1 \ \neg x_1$	1 $\neg x_1 \vee \neg y_1$	1 $x_1 y_1$
2 \top	$y_1 \ \neg y_1$	1 $y_1 \ \top$	1 $y_1 \ \neg y_1$	2 $\neg x_1 \vee \neg y_2$	2 $x_1 y_2$
	$y_2 \ \neg y_2$	2 $x_1 \ \top$	2 $x_1 \ \neg x_1$		
		2 $y_2 \ \top$	2 $y_2 \ \neg y_2$		

This query is already hard when T is the only probabilistic input relation!

A More Involved Case

Input formula and query:

- $\Psi = x_1y_1 \vee x_1y_2 \vee x_2y_1$ over sets $\mathbf{X} = \{x_1, x_2\}$, $\mathbf{Y} = \{y_1, y_2\}$
- $Q = \pi_{\emptyset} [S(A, B) - R(A) \times T(B)]$

We need a different reduction gadget:

- Use additional random variables $\mathbf{Z} = \{z_1, \dots, z_{|E|}\}$, one per clause in $\Psi = \psi_1 \vee \dots \vee \psi_{|E|}$.
- Construct a database D such that the grounding of Q wrt D is $\neg \Upsilon = \neg [\bigvee_{i=1}^{|E|} \neg z_i \neg \psi_i] = \bigwedge_{i=1}^{|E|} (z_i \vee \psi_i)$.

R	T	S	$S - R \times T$	$\pi_{\emptyset} [S - R \times T]$
$A \ \Phi$	$B \ \Phi$	$A \ B \ \Phi$	$A \ B \ \Phi$	Φ
$x_1 \ x_1$	$y_1 \ x_1$	$x_1 \ y_1 \ \neg z_1$	$x_1 \ y_1 \ \neg z_1 \neg (x_1 y_1)$	$() \bigvee_{i=1}^{ E } \neg z_i \neg \psi_i$
$x_2 \ x_2$	$y_2 \ y_2$	$x_1 \ y_2 \ \neg z_2$	$x_1 \ y_2 \ \neg z_2 \neg (x_1 y_2)$	
		$x_2 \ y_1 \ \neg z_3$	$x_2 \ y_1 \ \neg z_3 \neg (x_2 y_1)$	

- Compute $\#\Psi$ using linearly many calls to the oracle for $P_Q = 1 - P(\Upsilon)$.

The Small Print (1/2)

- $\Psi = \bigvee_{(i,j) \in E} x_i y_j = \psi_1 \vee \cdots \vee \psi_{|E|}$ over disjoint variable sets \mathbf{X} and \mathbf{Y}
- Let Θ be the set of assignments of variables $\mathbf{X} \cup \mathbf{Y}$ that satisfy Ψ :

$$\#\Psi = \sum_{\theta \in \Theta: \theta \models \Psi} 1.$$

- Partition Θ into disjoint sets $\Theta_0 \cup \cdots \cup \Theta_{|E|}$, such that $\theta \in \Theta_i$ if and only if θ satisfies exactly i clauses of Ψ :

$$\#\Psi = \sum_{\theta \in \Theta_1: \theta \models \Psi} 1 + \cdots + \sum_{\theta \in \Theta_{|E|}: \theta \models \Psi} 1 = |\Theta_1| + \cdots + |\Theta_{|E|}|.$$

- $|\Theta_1|, \dots, |\Theta_{|E|}|$ can be computed using an oracle for P_{Υ} :

$$\Upsilon = \bigvee_{i=1}^{|E|} \neg z_i \wedge \neg \psi_i \quad \text{or, equivalently} \quad \neg \Upsilon = \bigwedge_{i=1}^{|E|} (z_i \vee \psi_i)$$

The Small Print (2/2)

Express the probability of $\neg\Upsilon = \bigwedge_{i=1}^{|E|} (z_i \vee \psi_i)$ as a function of $|\Theta_1|, \dots, |\Theta_{|E|}|$:

- Fix the probabilities of variables in $\mathbf{X} \cup \mathbf{Y}$ to $1/2$ and of variables in \mathbf{Z} to p_z . Then:







$$\begin{aligned} P_{\neg\Upsilon} &= \sum_{k=0}^{|E|} \underbrace{P\left(\neg\Upsilon \mid \begin{array}{l} \text{exactly } k \text{ clauses} \\ \text{of } \Psi \text{ are satisfied} \end{array}\right)}_{p_z^{|E|-k}} \cdot \underbrace{P\left(\begin{array}{l} \text{exactly } k \text{ clauses} \\ \text{of } \Psi \text{ are satisfied} \end{array}\right)}_{\frac{1}{2}^{|\mathbf{X}|+|\mathbf{Y}|} \cdot |\Theta_k|} \\ &= \frac{1}{2}^{|\mathbf{X}|+|\mathbf{Y}|} \sum_{k=0}^{|E|} p_z^{|E|-k} |\Theta_k| \end{aligned}$$

- This is a polynomial in p_z of degree $|E|$, with coefficients $|\Theta_0|, \dots, |\Theta_{|E|}|$.
- The coefficients can be derived from $|E| + 1$ pairs (p_z, P_{Υ}) using Lagrange's polynomial interpolation formula.
- $|E| + 1$ oracle calls to P_{Υ} suffice to determine $\#\Psi = \sum_{i=0}^{|E|} |\Theta_i|$.

Hard Query Patterns

There are 48 (!) minimal non-hierarchical query patterns.

- Binary trees with leaves A , AB , and B and inner nodes \bowtie or $-$.
 - ▶ Some are symmetric and need not be considered separately:
 A and B can be exchanged, joins are commutative and associative.
 - ▶ Still, many cases left to consider due to the difference operator.



...

...

...







...

- There is a database construction scheme for each pattern.
- Each non-hierarchical query Q matches a pattern $P_{x,y}$.

Hard Query Patterns

There are 48 (!) minimal non-hierarchical query patterns.

- Binary trees with leaves A , AB , and B and inner nodes \bowtie or $-$.
 - ▶ Some are symmetric and need not be considered separately:
 A and B can be exchanged, joins are commutative and associative.
 - ▶ Still, many cases left to consider due to the difference operator.



...

...

...

...


- There is a database construction scheme for each pattern.
- Each non-hierarchical query Q matches a pattern $P_{x,y}$.

In the absence of negation, **P_{1.1}** is the only hard pattern to consider!

Non-hierarchical Queries Match Minimal Hard Patterns


Each non-hierarchical query Q matches a pattern $P_{x,y}$:

- There is a total mapping from $P_{x,y}$ to Q 's parse tree that
 - is identity on inner nodes \bowtie and $-$,
 - preserves ancestor-descendant relationships,
 - maps leaves to relations: A to $R(A)$; AB to $S(A, B)$; and B to $T(B)$.



- The match “preserves” the grounding of the query pattern:
 Q and $P_{x,y}$ have the same grounding for any database using our construction scheme.

Outline



Probabilistic Databases 101

The Dichotomy

The Interesting but Hard Queries

The Easy Queries

Leftovers

Evaluation of Hierarchical 1RA⁻ Queries

Approach based on knowledge compilation

- For any database D , the probability $P_{Q(D)}$ of a 1RA⁻ query Q is the probability P_{Ψ} of Q 's grounding Ψ .
- Compile Ψ into $\text{OBDD}(\Psi)$ in polynomial time.
- Compute probability of $\text{OBDD}(\Psi)$ in time linear in its size.

Evaluation of Hierarchical 1RA⁻ Queries

Approach based on knowledge compilation

- For any database D , the probability $P_{Q(D)}$ of a 1RA⁻ query Q is the probability P_{Ψ} of Q 's grounding Ψ .
- Compile Ψ into $\text{OBDD}(\Psi)$ in polynomial time.
- Compute probability of $\text{OBDD}(\Psi)$ in time linear in its size.

Distinction from existing tractability results [O. & Huang 2008]:

- 1RA⁻ *without* negation: Grounding formulas are read-once.
 - ▶ Read-once formulas admit linear-size OBDDs.
- 1RA⁻: Grounding formulas are not read-once.
 - ▶ They admit OBDDs of sizes linear in the database size but exponential in the query size.

The Inner Workings

From hierarchical $1RA^-$ to RC-hierarchical \exists -consistent RC^\exists :

- Translate query Q into an equivalent disjunction of disjunction-free existential relational calculus queries $Q_1 \vee \dots \vee Q_k$.
 - ▶ k can be very large for queries with projection under difference!
- **RC-hierarchical:**
For each $\exists_X(Q')$, every relation symbol in Q' has variable X .
 - ▶ Each of the disjuncts yields a poly-size OBDD.
- **\exists -consistent:**
The nesting order of the quantifiers is the same in Q_1, \dots, Q_k .
 - ▶ All OBDDs have compatible variable orders and their disjunction is a poly-size OBDD.
- The OBDD width grows exponentially with k , its height stays linear in the size of the database.
 - ▶ Width = maximum number of edges crossing the section between any two consecutive levels.

Query Evaluation Example (1/3)

Consider the following query and tuple-independent database:

$$Q = \pi_{\emptyset} \left[(R(A) \times T(B)) - (U(A) \times V(B)) \right]$$

<u>R</u>	<u>T</u>	<u>U</u>	<u>V</u>	<u>R × T</u>	<u>R × T - U × V</u>
A Φ	B Φ	A Φ	B Φ	A B Φ	A B Φ
1 r ₁	1 t ₁	1 u ₁	1 v ₁	1 1 r ₁ t ₁	1 1 r ₁ t ₁ ¬(u ₁ v ₁)
2 r ₂	2 t ₂	2 u ₂	2 v ₂	1 2 r ₁ t ₂	1 2 r ₁ t ₂ ¬(u ₁ v ₂)
				2 1 r ₂ t ₁	2 1 r ₂ t ₁ ¬(u ₂ v ₁)
				2 2 r ₂ t ₂	2 2 r ₂ t ₂ ¬(u ₂ v ₂)

Query Evaluation Example (1/3)

Consider the following query and tuple-independent database:

$$Q = \pi_{\emptyset} \left[(R(A) \times T(B)) - (U(A) \times V(B)) \right]$$

<i>R</i>	<i>T</i>	<i>U</i>	<i>V</i>	<i>R</i> × <i>T</i>	<i>R</i> × <i>T</i> − <i>U</i> × <i>V</i>
A φ	B φ	A φ	B φ	A B φ	A B φ
1 <i>r</i> ₁	1 <i>t</i> ₁	1 <i>u</i> ₁	1 <i>v</i> ₁	1 1 <i>r</i> ₁ <i>t</i> ₁	1 1 <i>r</i> ₁ <i>t</i> ₁ ¬(<i>u</i> ₁ <i>v</i> ₁)
2 <i>r</i> ₂	2 <i>t</i> ₂	2 <i>u</i> ₂	2 <i>v</i> ₂	1 2 <i>r</i> ₁ <i>t</i> ₂	1 2 <i>r</i> ₁ <i>t</i> ₂ ¬(<i>u</i> ₁ <i>v</i> ₂)
				2 1 <i>r</i> ₂ <i>t</i> ₁	2 1 <i>r</i> ₂ <i>t</i> ₁ ¬(<i>u</i> ₂ <i>v</i> ₁)
				2 2 <i>r</i> ₂ <i>t</i> ₂	2 2 <i>r</i> ₂ <i>t</i> ₂ ¬(<i>u</i> ₂ <i>v</i> ₂)

The grounding of Q is:

$$\Psi = r_1 [t_1(\neg u_1 \vee \neg v_1) \vee t_2(\neg u_1 \vee \neg v_2)] \vee r_2 [t_1(\neg u_2 \vee \neg v_1) \vee t_2(\neg u_2 \vee \neg v_2)].$$

- Variables entangle in Ψ beyond read-once factorization.
- This is the pivotal intricacy introduced by the difference operator.

Query Evaluation Example (2/3)

Translate $Q = \pi_{\emptyset}[(R(A) \times T(B)) - (U(A) \times V(B))]$ into RC^{\exists} :

$$Q_{RC} = \underbrace{\exists_A(R(A) \wedge \neg U(A)) \wedge \exists_B T(B)}_{Q_1} \vee \underbrace{\exists_A R(A) \wedge \exists_B (T(B) \wedge \neg V(B))}_{Q_2}.$$

- Both Q_1 and Q_2 are RC-hierarchical.
- $Q_1 \vee Q_2$ is \exists -consistent: Same order $\exists_A \exists_B$ for Q_1 and Q_2 .

Query grounding:


$$\Psi = \underbrace{(r_1 \neg u_1 \vee r_2 \neg u_2) \wedge (t_1 \vee t_2)}_{\Psi_1} \vee \underbrace{(r_1 \vee r_2) \wedge (t_1 \neg v_1 \vee t_2 \neg v_2)}_{\Psi_2}.$$

- Both Ψ_1 and Ψ_2 admit linear-size OBDDs.
- The two OBDDs have compatible orders and their disjunction is an OBDD whose width is the product of the widths of the two OBDDs.

Query Evaluation Example (3/3)

Compile grounding formula into OBDD:

$$\Psi = \underbrace{(r_1 \neg u_1 \vee r_2 \neg u_2) \wedge (t_1 \vee t_2)}_{\Psi_1} \vee \underbrace{(r_1 \vee r_2) \wedge (t_1 \neg v_1 \vee t_2 \neg v_2)}_{\Psi_2}.$$



Outline



Probabilistic Databases 101

The Dichotomy

The Interesting but Hard Queries

The Easy Queries

Leftovers

Dichotomies Beyond $1RA^-$

Some known dichotomies

- Non-repeating CQ, UCQ [Dalvi & Suciu 2004, 2010]
- Quantified queries, ranking queries [O.& team 2011, 2012]

Non-repeating relational algebra = $1RA^-$ + union.

- Hierarchical property not enough, consistency also needed.
- $\pi_{\emptyset}[(R(A) \bowtie S_1(A, B) \cup T(B) \bowtie S_2(A, B)) - S(A, B)]$ is hard, though it is equivalent to a union of two hierarchical $1RA^-$ queries.

Non-repeating relational calculus

- $S(x, y) \wedge \neg R(x)$ is tractable, $S(x, y) \wedge (R(x) \vee T(y))$ is hard.
 - ▶ Both are non-repeating, yet not expressible in $1RA^-$.
- Possible (though expensive) approach:
 - ▶ Translate to RC^{\exists} and check RC-hierarchical and \exists -consistency.

Full relational algebra (or full relational calculus)

- It is undecidable whether the union of two equivalent relational algebra queries, one hard and one tractable, is tractable.

Thank you!