

Probabilistic Inference for Datalog with Correlated Inputs

JINGBO WANG, Purdue University, USA

SHASHIN HALALINGAIAH, University of Texas at Austin, USA

WEIYI CHEN, Purdue University, USA

CHAO WANG, University of Southern California, USA

IŞIL DILLIG, University of Texas at Austin, USA

Probabilistic extensions of logic programming languages, such as ProbLog, integrate logical reasoning with probabilistic inference to evaluate probabilities of output relations; however, prior work does not account for potential statistical correlations among input facts. This paper introduces PRALINE, a new extension to Datalog designed for precise probabilistic inference in the presence of (partially known) input correlations. We formulate the inference task as a constrained optimization problem, where the solution yields sound and precise probability bounds for output facts. However, due to the complexity of the resulting optimization problem, this approach alone often does not scale to large programs. To address scalability, we propose a more efficient δ -exact inference algorithm that leverages constraint solving, static analysis, and iterative refinement. Our empirical evaluation on challenging real-world benchmarks, including side-channel analysis, demonstrates that our method not only scales effectively but also delivers tight probability bounds.

CCS Concepts: • **Mathematics of computing** → **Probabilistic inference problems**; **Mathematical optimization**; • **Theory of computation** → **Constraint and logic programming**.

Additional Key Words and Phrases: probabilistic logic programming, constrained optimization, type inference

ACM Reference Format:

Jingbo Wang, Shashin Halalingaiah, Weiyi Chen, Chao Wang, and Işıl Dillig. 2025. Probabilistic Inference for Datalog with Correlated Inputs. *Proc. ACM Program. Lang.* 9, OOPSLA2, Article 280 (October 2025), 56 pages. <https://doi.org/10.1145/3763058>

1 Introduction

Logic programming languages are powerful tools for modeling and reasoning about complex systems using well-defined rules, and they are also widely used in graph analysis, bioinformatics, and program analysis tasks [45, 64] such as detecting race conditions [41, 63] and side channels [58, 59]. However, real-world scenarios often involve uncertainty and incomplete information that traditional logic programming cannot handle effectively. Probabilistic extensions to logic programming languages, such as ProbLog [16, 17] and PDDL [5], enhance traditional logic programming by incorporating probabilistic reasoning. For example, in medical diagnosis, symptom-disease associations are inherently probabilistic, and quantifying disease likelihood based on symptoms can enhance diagnostic accuracy. Such probabilistic reasoning is also useful for applications of logic programming in program analysis tasks [44]: in the context of side channel detection, understanding the probability of information leakage can significantly enhance security assessments.

Authors' Contact Information: Jingbo Wang, Purdue University, West Lafayette, USA, wang6203@purdue.edu; Shashin Halalingaiah, University of Texas at Austin, Austin, USA, shashin@cs.utexas.edu; Weiyi Chen, Purdue University, West Lafayette, USA, chen5332@purdue.edu; Chao Wang, University of Southern California, Los Angeles, USA, wang626@usc.edu; Işıl Dillig, University of Texas at Austin, Austin, USA, isil@cs.utexas.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2475-1421/2025/10-ART280

<https://doi.org/10.1145/3763058>

However, these probabilistic extensions to logic programming languages typically assume the independence of input facts to simplify inference [5, 16, 38]. In many real-world scenarios, this assumption does not hold: for example, in medical diagnostics, certain symptoms are often correlated, and treating them as independent can lead to incorrect diagnoses. Similarly, in side-channel analysis, input facts such as memory accesses and branch outcomes are frequently correlated, and ignoring these correlations can produce unsound assessments.

This work aims to develop a novel probabilistic Datalog framework that can accommodate arbitrary statistical correlations between inputs. However, the removal of the independence assumption introduces several challenges: First, it significantly increases the complexity of probabilistic inference, requiring consideration of joint probabilities of inputs rather than treating them in isolation. Second, since the exact conditional dependencies between input facts may not be known a priori, there can be *implicit* dependencies that are not explicitly specified by users. Finally, due to *partial* knowledge about input dependencies, it may not always be possible to compute the *exact* probability of an output fact, necessitating the computation of lower and upper probability bounds. Our proposed framework addresses these challenges, providing an effective approach for accurate and reliable probabilistic reasoning in complex scenarios that involve correlated inputs.

At the core of our approach is a formulation of probabilistic inference as a constrained optimization problem over *joint probability variables*, which capture the joint probability distribution of correlated input facts. Probabilities specified in the Datalog program are translated into a system of constraints, and the probability of each output relation is expressed symbolically in terms of these variables. Our method then solves a constrained optimization problem to compute upper and lower bounds on the probability of the desired output fact. However, a key challenge with this approach is that the generated constrained optimization problems can be very challenging to solve, particularly due to the non-linear nature of the objective function. To address this difficulty, we also present a more scalable δ -exact algorithm [23] that uses the optimization approach in a targeted way to solve simpler sub-problems. Our proposed method first computes *approximate* probability bounds and then iteratively refines them until they are within a user-specified distance δ of the true probability bounds. As illustrated in Figure 1, our technique comprises of three steps. First, it performs a combination of static analysis and constraint solving to infer whether a pair of facts are positively or negatively correlated, or independent. In the second step, it leverages the results of this “correlation type analysis” to compute approximate bounds on the probabilities of each predicate. Finally, it iteratively refines these approximate bounds to a user-specified tightness δ — i.e., upon termination, the computed lower and upper bounds on the probabilities are guaranteed to be within δ of the ground truth.

We have implemented the proposed idea in a new Datalog variant called PRALINE¹ and evaluate it on two domains, namely side channel vulnerability detection and inference for discrete Bayesian networks. Our experiments show that PRALINE can successfully infer accurate probability bounds for output facts, including for large benchmarks with hundreds of thousands of facts. We also conduct ablation studies to evaluate the impact of the three key ideas illustrated in Figure 1 and show that they have a significant impact on the scalability and precision of our approach.

To summarize, this paper makes the following key contributions:

- We introduce the first probabilistic Datalog framework that allows the user to specify arbitrary statistical dependencies between input facts.




Fig. 1. Overview of our approach

¹Stands for PRobAbilistic Logical InfeRence Engine

```

1 % Input facts.
2 % Known input facts about patient
3 1.00 :: chest_pain.
4 % Unknown input facts about patient
5 % Probabilities estimated from average population.
6 0.03 :: abnormal_ecg.
7 0.02 :: troponin_high.
8 ...
9 % Input correlations based on the literature
10 0.90 :: abnormal_ecg | chest_pain, troponin_high.
11 ...
12 % Inference rules for heart disease assessment
13 0.85 :: heart_disease :- abnormal_ecg, troponin_high,
    chest_pain.
14 0.60 :: heart_disease :- abnormal_ecg, high_bnp,
    chest_pain, \+troponin_high.
15 % Query
16 query(heart_disease).

1 % Input Facts about input program from Fig. 3
2 1.0 :: rand(r3). % r3 is a random variable
3 % XOR operation highlighted in Fig. 3
4 1.0 :: xor(n5, n7, n8).
5 ...
6 0.7 :: share(n5, n8). % register sharing
7 0.8 :: dep(n5, r3). % data dependency
8 0.8 :: dep(n8, r3).
9 ...
10 % Input correlations
11 0.9 :: share(n5,n8) | xor(n5,n7,n8), dep(n5,r3).
12 ...
13 % Rules for inferring leaks
14 0.7 :: t(V1,V2) :- dep(V1, R),rand(R),dep(V2, R).
15 0.8 :: leak(V1, V2) :- share(V1, V2), t(V1,V2).
16 ...
17 % Query
18 query(leak(_,_)).

```

Fig. 2. Two simple PRALINE programs. Statistical correlations between inputs (in green) are specified using the syntax $p :: I \mid I_1, I_2$, indicating that probability of I given I_1, I_2 is p , and $\neg I$ indicates the negation of I .

- We propose a constrained optimization approach to compute exact probability bounds for outputs.
- We present a δ -exact algorithm that combines our basic constrained optimization formulation with static analysis, approximation, and iterative refinement to improve scalability.
- We perform an empirical evaluation on 30 real-world probabilistic Datalog programs consisting of large-scale program analysis tasks and Bayesian inference benchmarks. Our results show that PRALINE can produce precise probability bounds, while scaling to large benchmarks.

2 Motivating Examples

In this section, we provide two simple examples to motivate the probabilistic inference capabilities of PRALINE, highlighting scenarios where input facts are correlated but the full joint probability distribution is unknown.

Medical diagnosis. The left side of Figure 2 illustrates how a PRALINE program can perform probabilistic inference for medical diagnosis. Specifically, it models the probability of a patient having heart disease based on observed symptoms (e.g., chest pain) and test results (e.g., ECG, blood biomarkers). Lines 12–14 of the PRALINE program encode how heart disease can be inferred from various diagnostic factors based on well-established medical literature. These factors include an abnormal ECG (abnormal_ecg), elevated cardiac biomarkers such as troponin and B-type Natriuretic Peptide (BNP) (troponin_high, high_bnp), and the presence of chest pain (chest_pain).

For example, the rule at line 13 states that the presence of heart_disease (an output fact) can be deduced based on the input facts abnormal_ecg (test result), troponin_high (blood work result) and chest_pain (symptom). The input fact in line 3 represents the patient’s known symptom (chest pain, with probability 1.0), while lines 6–7 encode test results that are currently unknown because the patient has not yet undergone an ECG or blood work. In clinical settings, when test results are unavailable, probabilities for these factors (e.g., abnormal_ecg) are estimated based on population statistics. However, chest pain is often correlated with other markers of heart disease, such as abnormal ECG findings and elevated troponin levels. This dependence is captured in line 10, where the probability of an abnormal ECG increases to 0.9 given chest pain and high troponin levels. However, the full joint probability distribution of these input facts (e.g., symptoms and test results) is typically unknown, as it depends on numerous latent factors, such as medical history or environmental exposures. The probabilistic reasoning capabilities of PRALINE enable accurate estimation of heart disease risk even under such *partially-known* information. For example,

assuming independence of input facts would lead us to estimate the probability of heart disease given chest pain as 0.07% whereas the true probability is in the range 1.5 – 2.1%.²

Quantitative program analysis. Another motivating scenario for PRALINE is quantitative Datalog-based program analysis, which has emerged as a powerful approach for static reasoning about program behavior [41, 44, 58, 59, 63]. In this example, we consider a (drastically) simplified version of the Datalog-based side channel detection method described in [59]. The PRALINE program shown on the right side of Figure 2 encodes an analysis for detecting information leaks caused by power side channels. In particular, the rules in lines 14–15 define how leakage is inferred based on data dependencies in the input program ($\text{dep}(X, Y)$), randomness ($\text{rand}(X)$), and register sharing between variables ($\text{share}(X, Y)$).

The input facts in lines 1–5 of Figure 2 (right) correspond to properties of the analyzed program (see Figure 3). Some of these facts, such as those in lines 2 and 4, have probability 1.0 because they correspond to known characteristics of the program. For instance, $r3$ is annotated by the user to be a random variable, and the xor operation in line 4 corresponds to a specific highlighted statement in the source program from Figure 3. However, not all input facts are deterministic. Whether two variables share a register, for example, depends on hardware constraints and register allocation policies, which introduce uncertainty. This uncertainty is reflected in line 6, where the probability of register sharing is estimated using empirical data from profiling a code corpus. Similarly, the input facts in lines 7–8 are probabilistic because they result from a pre-analysis [59, 62] that infers semantic data dependencies from syntactic ones. The probability associated with such dependencies is derived from prior empirical studies that measure how often syntactic dependencies lead to actual data dependencies in compiled programs [59].

Since register allocation and data dependencies are influenced by compiler optimizations and architectural constraints, certain input facts are naturally correlated. For example, as shown in line 11, the likelihood of two variables sharing a register is not independent of how frequently the compiler assigns dependent variables (within the same instruction) to the same register across different executions [3]. However, as in the medical diagnosis example, assuming access to the full joint probability distribution is impractical, as it would require exhaustively modeling all interactions between hardware configurations, compiler optimizations, etc. Instead, by combining Datalog inference with partially known probabilities (e.g., derived from empirical measurements and power consumption models), we enable *quantitative* static analysis that estimates the *severity* of an information leak rather than merely providing a binary vulnerability classification.

3 Overview

Before formalizing our technique in detail, we illustrate how PRALINE performs probabilistic inference using the synthetic example in Figure 4 that is crafted to give an overview of our approach without being overly complicated. The first six lines declare *input facts*, specifying graph edges, with associated probabilities (e.g., line 1 states that an edge between nodes 5 and 7 exists with probability 0.7). Lines 10–11 are rules that define the notion of paths in the graph. Line 7 states that $\text{edge}(1, 4)$, $\text{edge}(2, 5)$, and $\text{edge}(2, 6)$ are statistically correlated, while lines 8–9 specify the known conditional dependencies. For instance, line 8 states that, given $\text{edge}(1, 4)$, $\text{edge}(2, 5)$

```

1 bool mChi(bool i1 ... i3,
             bool r1 ... r3) {
2   bool b1, b2, b3... ;
3   b1 = i1 ⊕ r1; b2 = i2 ⊕ r2;
4   b3 = i3 ⊕ r3; n8 = r3 ∧ r2;
5   n7 = r3 ∨ b2;
6   n5 = n7 ⊕ n8;
7   ...}

```

Fig. 3. Masked χ function from MAC-Keccak [18]

²These probabilities are computed based on the simple PRALINE program and do not reflect actual probability of heart disease given chest pain.

```

1 0.7::edge(5,7).
2 0.6::edge(1,2).
3 0.8::edge(6,7).
4 0.6::edge(2,5).
5 0.6::edge(1,4).
6 0.6::edge(2,6).
7 corr(edge(1,4),edge(2,5),edge(2,6)).
8 0.80::edge(2,5) | edge(1,4).
9 0.83::edge(2,6) | edge(1,4).
10 path(X,Y) :- edge(X,Y).
11 path(X,Y) :- path(X,Z), edge(Z,Y).
12 query(path(1,7)).

```

Fig. 4. A PRALINE program




Fig. 5. Derivation graph

exists with probability 0.8. Finally, line 12 queries the probability of $\text{path}(1, 7)$. We refer to the input facts that are correlated (as declared on line 7) as a *correlation class*.

This example exhibits a salient feature we wish to highlight: Given the information specified by the user, the probability of the predicate $\text{path}(1, 7)$ cannot be determined *exactly*; instead, we can *only* derive upper and lower bounds. In particular, to compute this probability exactly, the user would need to give the full conditional probability table [14] between the predicates $\text{edge}(1, 4)$, $\text{edge}(2, 5)$, and $\text{edge}(2, 6)$; but, in the absence of such information, the only sound conclusion that can be reached is that the probability of $\text{path}(1, 7)$ being true is in the range $[0.34, 0.42]$. To the best of our knowledge, there is no existing technique that can perform precise and accurate probabilistic inference for this type of scenario. In the rest of this section, we illustrate how our proposed inference technique addresses this challenge.

3.1 Basic Approach: Inference via Constrained Optimization

Our starting point is a formulation of this probabilistic inference task as a constrained optimization problem. To formulate this optimization problem, our approach introduces a set of *joint probability variables*, each representing the unknown value of an entry in the conditional probability table for a given correlation class. In our example, the predicates $\text{edge}(2, 5)$, $\text{edge}(1, 4)$, and $\text{edge}(2, 6)$ form a correlation class, while $\text{edge}(5, 7)$, $\text{edge}(6, 7)$, and $\text{edge}(1, 2)$ are mutually independent. For the singleton correlation class containing $\text{edge}(5, 7)$, we introduce two joint probability variables, $V_1[1]$ and $V_1[0]$, representing the probabilities of the predicate $\text{edge}(5, 7)$ being true and false, respectively. Since its probability is given as 0.7, the first constraint in Figure 6 sets $V_1[1] = 0.7$ and, by the SumToOne constraint, sets $V_1[0] = 0.3$. The constraints for the correlation class with predicates $\text{edge}(2, 5)$, $\text{edge}(1, 4)$, and $\text{edge}(2, 6)$ are more involved. Since this class contains 3 predicates, it requires 8 joint probability variables, representing all boolean combinations of these predicates. For example, the variable $V_4[100]$ in Figure 6 represents the probability of $\text{edge}(2, 5)$ being true and $\text{edge}(1, 4)$ and $\text{edge}(2, 6)$ being false. Line 4 of Figure 4 enforces the constraint $V_4[100] + V_4[101] + V_4[110] + V_4[111] = 0.6$, ensuring that the sum of joint probability variables where $\text{edge}(2, 5)$ holds is 0.6. Similarly, the conditional probability from line 8 induces the constraint $V_4[110] + V_4[111] = 0.48$, capturing $P(\text{edge}(2, 5) \wedge \text{edge}(1, 4)) = P(\text{edge}(2, 5) | \text{edge}(1, 4)) \times P(\text{edge}(1, 4))$. Figure 6 presents all such constraints for the example in Figure 4.

In addition to generating these constraints, our method also expresses the unknown probability of output facts in terms of the joint probability variables using the logical derivation graph of the input program. In particular, Figure 5 shows the logical derivation of the output predicate $\text{path}(1, 7)$ as

Constraints

$\text{edge}(5,7) = 0.7$
 $\text{edge}(1,2) = 0.6$
 $\text{edge}(6,7) = 0.8$
 $\text{edge}(2,5) = 0.6$
 $\text{edge}(1,4) = 0.6$
 $\text{edge}(2,6) = 0.6$

$0.80 :: \text{edge}(2,5) \mid \text{edge}(1,4).$
 $0.83 :: \text{edge}(2,6) \mid \text{edge}(1,4).$

SumToOne

Input

Objective

$\text{path}(1,7)$

$\mapsto V_1[1] = 0.7$
 $\mapsto V_2[1] = 0.6$
 $\mapsto V_3[1] = 0.8$
 $\mapsto V_4[100] + V_4[101] + V_4[110] + V_4[111] = 0.6$
 $\mapsto V_4[010] + V_4[011] + V_4[110] + V_4[111] = 0.6$
 $\mapsto V_4[001] + V_4[011] + V_4[101] + V_4[111] = 0.6$
 $\mapsto V_4[110] + V_4[111] = 0.8 * 0.6$
 $\mapsto V_4[011] + V_4[111] = 0.83 * 0.6$
 $\mapsto \sum_{b \in \mathbb{B}} V_1[b] = 1 \quad \sum_{b \in \mathbb{B}} V_2[b] = 1 \quad \sum_{b \in \mathbb{B}} V_3[b] = 1 \quad \sum_{b \in \mathbb{B}^3} V_4[b] = 1$
 $\mapsto \forall b \in \mathbb{B}, V_1[b] \in [0, 1], V_2[b] \in [0, 1], V_3[b] \in [0, 1], \forall b \in \mathbb{B}^3, V_4[b] \in [0, 1]$
 $\mapsto V_1[1]V_2[1]V_3[1](V_4[111] + V_4[101]) + V_1[1]V_2[1]V_3[0](V_4[110] + V_4[100]) +$
 $V_1[1]V_2[1]V_3[1](V_4[110] + V_4[100]) + V_1[1]V_2[1]V_3[0](V_4[111] + V_4[101]) +$
 $V_1[0]V_2[1]V_3[1](V_4[011] + V_4[001]) + V_1[0]V_2[1]V_3[1](V_4[101] + V_4[111]) +$
 $V_1[1]V_2[1]V_3[1](V_4[011] + V_4[001])$

Fig. 6. Optimization problem for running example

a graph where nodes correspond to predicates and edges correspond to Datalog rule applications. For instance, according to Figure 5, there are two ways to derive the predicate $\text{path}(1,7)$: one using $\text{edge}(5,7)$ and $\text{path}(1,5)$ and another using $\text{path}(1,6)$ and $\text{edge}(6,7)$. As shown in Figure 6 under Objective, our method uses this information to express the probability of predicate $\text{path}(1,7)$ being true in terms of the joint probability variables. Finally, our method computes the probability bounds for $\text{path}(1,7)$ being true by minimizing and maximizing the objective function subject to the constraints shown in Figure 6. The bold annotations on the derivation graph in Figure 5 show the probability bounds obtained for each relation using our method. It is worth re-iterating that these probabilities are ranges rather than absolute values *not* because of some imprecision in this solution but rather because of unknown conditional dependencies between some of the input facts. The details of this optimization approach are presented in Section 6.

3.2 Scalable δ -exact Inference

The approach introduced above guarantees precise lower and upper bounds, but it can be quite expensive due to the complexity of the resulting constrained optimization problem. Thus, to scale our approach to large Datalog programs with complex statistical correlations, we propose a δ -exact algorithm that first computes (loose) approximate probabilities for each output fact and then iteratively tightens these probabilities until they are within δ of the true lower and upper bounds. Our novel δ -exact algorithm can be reduced to three steps, described next.

Step 1: Correlation type inference. Starting with tighter initial approximations accelerates the overall refinement process, and just knowing the *polarity* of the statistical correlation between predicates can help us compute more precise approximations. For example, if two predicates are positively correlated (i.e., one is more likely true if the other is), we can derive tighter bounds than if the correlation type were unknown. Thus, our method first infers *correlation types* between predicates, classifying them as *positive*, *negative*, *independent*, or *unknown*. To do so, we adapt a simpler version of our constrained optimization method to deduce *how* input facts are correlated. For instance, in our example, while the program specifies that $\text{edge}(2,5)$, $\text{edge}(2,6)$, and $\text{edge}(1,4)$ are correlated, it does not specify *how*. Using our analysis, we first deduce that $\text{edge}(2,5)$, $\text{edge}(2,6)$ are positively




Fig. 7. Iterative strengthening. l^*, u^* are the exact probability bounds; l, u are approximate bounds from Step 2, and l_n, u_n are bounds computed by our method.

correlated, which helps establish that $\text{path}(1, 5)$, $\text{path}(1, 6)$ are *also* positively correlated. We present the details of our correlation analysis method in Section 7.1.

Step 2: Deriving approximate probability bounds. Next, our algorithm uses the correlation analysis results to compute sound but approximate bounds on output probabilities. This involves bottom-up propagation of probabilities from the leaves to the roots of the derivation graph, applying pre-defined approximation rules at each internal node. In our example, to approximate the probability of $\text{path}(1, 7)$, the algorithm first computes that $\text{path}(1, 5)$ and $\text{path}(1, 6)$ each have a probability of 0.36 as they consist of two independent events, each with a probability of 0.6. It then moves to $\text{path}(1, 7)$, which is true if either $\text{path}(1, 5) \wedge \text{edge}(5, 7)$ (event e_1) or $\text{path}(1, 6) \wedge \text{edge}(6, 7)$ (event e_2) is true. To compute the exact probability of $\text{path}(1, 7)$, the conditional probabilities $P(e_1|e_2)$ and $P(e_2|e_1)$ are needed, but are not provided. In the absence of such information, the Fréchet inequality [46] gives the upper bound $P(e_1 \vee e_2) = \min(1, 0.54)$. However, by utilizing the information that $\text{path}(1, 5)$ and $\text{path}(1, 6)$ are positively correlated, we can compute a tighter upper bound of 0.467. Details of this method are in Section 7.2.

Step 3: Iterative refinement. While sound, the probability bounds from the previous step may be far from the true values. Hence, our method iteratively tightens the bounds until they are within a user-specified distance δ of the true bounds. The intuition is to gradually adjust the bounds while checking whether the resultant bound is consistent with the constraint system induced by the program.

As an illustration, consider the output fact $\text{path}(1, 7)$ in our example. Starting from the bounds $[l, u] = [0.288, 0.467]$ computed by the previous step, the iterative refinement step repeatedly attempts to increase l and decrease u by some amount ϵ until the constraint system becomes satisfiable, as shown in Figure 7. When this step terminates, we know that the true lower (resp. upper) bound is between l_n and l_{n+1} (resp. between u_n and u_{n+1}) in Figure 7. In a second tightening step, the algorithm performs binary search between l_n and l_{n+1} (and between u_n and u_{n+1}) until the two values are within distance δ of each other. Thus, upon termination, it can guarantee that the computed lower and upper bounds are always within δ of the ground truth. Assuming a user-specified $\delta = 0.05$ in our running example, this approach refines the initial approximate bounds for $\text{path}(1, 7)$ from $[0.288, 0.467]$ to $[0.338, 0.417]$, ensuring they lie within δ of the true bounds $[0.34, 0.42]$. The details of our iterative refinement procedure are provided in Section 7.3.

4 Preliminaries

A Datalog program consists of a set of rules R , where each rule is a Horn clause of the form $R(\vec{x}) :- \odot R_1(\vec{y}_1), \dots, \odot R_n(\vec{y}_n)$ where \odot denotes an optional negation operator. We refer to $R(\vec{x})$ as the *head* of the rule and the right hand side as its *body*. Given a rule r , we write Body^+ to denote all predicates without a negation symbol in the front and Body^- to denote all predicates that are negated. Finally, we refer to rules without a body as *input facts*. As standard, we can obtain a solution to a Datalog program by first grounding all predicates over the Herbrand universe and then applying the ground rules to a fixed-point. In the rest of this paper, we represent the *solution* to a Datalog program in terms of a *derivation graph*, which shows how each ground output predicate can be derived using ground input predicates.

Definition 1 (Derivation graph). A *derivation graph* for a Datalog program is a hypergraph (V, E) where V is a set of nodes representing ground predicates, and E is a set of directed hyperedges (h, B^+, B^-, r) representing a ground rule r with head h , body $B^+ = \text{Body}^+(r)$ and $B^- = \text{Body}^-(r)$. The source vertex set of this hyper-edge is $\{h\}$ and the target vertices are $B^+ \cup B^-$. Given a node n , we write $\text{Pred}(n)$ to denote the ground predicate represented by n .

Example 1. In the derivation graph of Figure 5, the edge e_3 has source node $\text{path}(1, 5)$ and target nodes $\text{edge}(2, 5)$ and $\text{path}(1, 2)$. The quadruple $(\text{path}(1, 5), \{\text{edge}(2, 5), \text{path}(1, 2)\}, \emptyset, r_2)$ is used to represent e_3 , where r_2 denotes the rule $1 :: \text{path}(1, 5) :- \text{path}(1, 2), \text{edge}(2, 5)$.

Note that different ways of deriving the same ground predicate correspond to multiple hyperedges that share the same source. Given a Datalog program D , we write $\mathcal{G}(D)$ to denote its derivation graph. Intuitively, the derivation graph encodes all possible ways of deriving a ground predicate given the input facts. In the rest of the paper, we assume that the derivation graph is acyclic, as it corresponds to the result of the solver's internal fixed-point computation.³

We also define a *flattened* version of the derivation graph that represents each hyperedge with a set of regular edges. As we will see later, this flattened representation is useful for understanding basic relationships between output and input predicates.

Definition 2 (Flattened derivation graph). Given a derivation graph $G = (V, E)$, $\text{Flatten}(G)$ yields a graph $G' = (V', E')$ where $V' = V$ and there exists an edge $(v, v', \sigma) \in E'$ iff there exists a hyperedge $(v, B^+, B^-, r) \in E$ where $v' \in B^+ \cup B^-$ and $\sigma = 1$ if $v' \in B^+$ and $\sigma = -1$ otherwise.

Example 2. Consider hyperedge e_3 from Figure 5. This corresponds to two separate edges e_3^1 and e_3^2 in the flattened derivation graph, where e_3^1 is from $\text{path}(1, 5)$ to $\text{edge}(2, 5)$ and e_3^2 is from $\text{path}(1, 5)$ to $\text{path}(1, 2)$. Since there are no negated predicates, both edges have a σ value 1.

The flattened derivation graph is useful for determining whether some output facts depend only positively or negatively on an input fact. Given a path $\pi = (v, v_1, \sigma_1), \dots, (v_n, v', \sigma_n)$ from node v to v' in the flattened derivation graph, we define the polarity of the path as $\prod_{i=1}^n \sigma_i$, denoted $\text{Polarity}(\pi)$. We can then classify dependencies between different output and input facts as follows:

Definition 3 (Dependence). Given a Datalog program D with flattened derivation graph G , we say that an output predicate O *depends on* input predicate I , denoted $I \triangleright O$ if there exists a path from O to I in G . We say that O depends positively (resp. negatively) on I iff there exists a path from O to I with a positive (resp. negative) polarity. We write $I \triangleright^+ O$ to denote (logical) positive dependence and $I \triangleright^- O$ for (logical) negative dependence.

The notion of dependence defined above refers to *logical* rather than *statistical* dependence, and it is possible for an output relation to depend both positively *and* negatively on an input fact.

5 PRALINE: Probabilistic Datalog with Correlated Inputs

In this section, we introduce the syntax of PRALINE, our probabilistic Datalog variant that allows correlated input facts. A pair of input facts I_1, I_2 are correlated iff $P(I_1|I_2) \neq P(I_1)$, and a *correlation class* represents a set of input facts that may be correlated with each other.

Definition 4. (Praline program) A PRALINE program is a tuple (C, R, P_R, P_I) where C is a set of correlation classes, R is a set of rules whose heads are output relations, P_R is a mapping from rules in R to their probabilities, and P_I is a set of (conditional) probabilities about input facts.

Syntax. Borrowing notation from prior work [17, 44], we express rule probabilities in PRALINE using the notation $p :: R_i$ meaning that the probability of rule $R_i \in R$ is p , and we assume that rules are statistically independent of each other. However, unlike rules, input facts can be statistically dependent. We express such conditional dependencies using the notation $p :: I | S$, meaning that the probability of input fact I given input fact set S is p . When S is the empty set, such a rule simply declares the probability of an input fact. Finally, we use the notation $\text{Class}(I)$ to denote the

³The derivation graph may contain auxiliary relations that are used for breaking cycles. We refer the interested reader to prior work for details on this semantics-preserving transformation [20].

Algorithm 1 SOLVE(D)**Input:** Datalog program $D = (C, R, P_R, P_I)$ **Output:** Mapping M from output relations to their probability intervals

```

1:  $S_o, G \leftarrow \text{SOLVESTANDARD}(D)$            ▶ SOLVESTANDARD produces output relations  $S_o$  and derivation graph  $G$ 
2:  $V \leftarrow \text{JOINTPROBABILITYVARS}(D)$        ▶ Introduce joint probability variables, as discussed in Section 6.1
3:  $\xi \leftarrow \text{GENEXPRTEMPLATE}(V)$            ▶ Generate probability expression templates as discussed in Section 6.2
4:  $\phi \leftarrow \text{GENCONSTRAINT}(V, D, \xi)$      ▶ Generate constraints  $\phi$  as discussed in Section 6.3
5: for  $R_o \in S_o$  do
6:    $\Phi \leftarrow \text{GENOBJECTIVE}(R_o, V, G, \xi)$  ▶ Generate objective  $\Phi$  as discussed in Section 6.4
7:    $M[R_o] \leftarrow \text{OPTIMIZE}(\Phi, \phi)$      ▶ Use MIP solver to minimize/maximize  $\Phi$  subject to  $\phi$ 
8: return  $M$ 

```

correlation class $C_i \in \mathcal{C}$ that input fact I belongs to, and we assume that each input fact $I \in C_i$ has an index, denoted $\text{Index}(I, C_i) \in [1, |C_i|]$, that can be used to uniquely identify I within C_i .

Semantics. We define the semantics of PRALINE as an extension of the least-fixed-point semantics of Datalog, incorporating probabilities through a *possible-worlds interpretation*. In particular, each Herbrand model of a PRALINE program represents a deterministic instantiation of probabilistic rules and facts, forming a possible world ω . Given a Praline program D , we define an interpretation of D as a pair (ω, μ) where ω is a possible world and μ is a full joint probability distribution over instantiated rules and input facts in D . We say that (ω, μ) is a model of D , denoted $(\omega, \mu) \models D$ if ω is a possible world of D (under the standard Datalog semantics) and μ is consistent with both ω and the probabilities in D . Then, given such a μ , each possible world ω has an exact probability associated with it, denoted as $P_\mu(\omega)$. Given an output fact O , we can now define $P_\mu(O)$ as follows:

$$P_\mu(O) = \sum_{\omega \in \Omega} P_\mu(\omega) \text{ where } \Omega = \{\omega \mid (\omega, \mu) \models D, \omega \models O\}$$

Finally, given an output fact O , the *set* of possible probabilities of O is given by:

$$P(O) = \{P_\mu(O) \mid \exists \omega. (\omega, \mu) \models D\}$$

We refer interested readers to Appendix A.3 in the Supplementary Material for a more formal treatment of the semantics.

Problem definition. We conclude this section by defining the probabilistic inference problem addressed in the remainder of this paper.

Definition 5. (Exact probabilistic inference) Given a PRALINE program $D = (C, R, P_R, P_I)$, the goal of exact probabilistic inference is to produce a mapping from each derived output fact O to a probability interval $[l, u]$ such that $P(O) = \{x \mid l \leq x \leq u\}$.

In Section 6, we provide a method for solving the exact probabilistic inference problem defined above. However, since exact inference is often computationally intractable, we also introduce the notion of δ -exact inference, which we address in Section 7:

Definition 6. (δ -exact probabilistic inference) Given a PRALINE program $D = (C, R, P_R, P_I)$, the goal of δ -exact probabilistic inference is to produce a mapping from each derived output fact O to a probability interval $[l, u]$ such that $l \geq l^* - \delta$ and $u \leq u^* + \delta$ where $P(O) = \{x \mid l^* \leq x \leq u^*\}$.

6 Exact Probabilistic Inference via Constrained Optimization

In this section, we address the *exact* probabilistic inference problem using constrained optimization. As summarized in Algorithm 1, our method consists of five steps. First, for each correlation class C_i , it introduces a set of variables that express the joint probabilities of input facts in that class C_i

(line 2). Second, at line 3, it generates a *probability expression template* ξ over these joint probability variables — the key idea is to express the probabilities of all relations symbolically as instantiations of the same shared template. Then, at line 4, it invokes GENCONSTRAINT to express requirements on the joint probability variables based on the facts in the Datalog program. Finally, for each fact R_o , the algorithm first invokes GENOBJECTIVE (line 6) to express the probability of R_o symbolically as an instantiation Φ of ξ and then optimizes Φ (line 7) subject to the constraints generated earlier.

6.1 Joint Probability Variables

The variables in our encoding represent *joint probabilities* over input facts; hence, we refer to them as *joint probability variables*. For each correlation class C , we introduce a map V_C of joint probability variables, denoted $\text{Rep}(C)$ (for “representative”). Given a correlation class C of size n , V_C contains 2^n variables, one for each possible boolean assignment to input facts in C . We use bitvectors to represent boolean assignments and write $V_C[b]$ to denote the corresponding joint probability variable. For example, for $C = \{I_1, I_2, I_3\}$, $V_C[001]$ denotes the joint probability of I_1, I_2 being false and I_3 being true. We also write \mathbb{B}^n to denote the set of all bitvectors of size n . Given an input predicate I and a joint probability variable v , $v \models I$ (resp. $v \not\models I$) indicates that v represents an event in which I is true (resp. false). In our example, we have $V_C[001] \models I_3$ and $V_C[001] \not\models I_1$.

6.2 Probability Expressions and Templates

As explained earlier, a key idea underlying our algorithm is to express the probability of each output relation as a symbolic expression over the joint probability variables. Because all of these expressions are instantiations of the same template, we first explain what these templates look like.

Definition 7. (Product term) Let $C_1 \dots C_n$ be the set of all correlation classes, i.e., $C_i \in C$. A product term ψ is a product of joint probability variables $v_1 \dots v_n$ where each v_i is a joint probability variable associated with class C_i , i.e., $v_i = \text{Rep}(C_i)[b]$ for some bitvector b .

Intuitively, a product term represents the probability of a particular truth assignment to *all* input relations in the program. For example, if we have two correlation classes $C_1 = \{I_1, I_2\}$ and $C_2 = \{I_3, I_4\}$, then the product term $\psi = V_1[10] \times V_2[11]$ represents the probability of I_2 being false and I_1, I_3, I_4 being true. Extending our previous notation, given a product term ψ and input relation I , we write $\psi \models I$ (resp. $\psi \not\models I$) to denote that ψ represents the probability of an event in which I is true (resp. false). For instance, in our previous example, we have $\psi \models I_1$ and $\psi \not\models I_2$.

Definition 8. (Probability expression template) Let Ψ be the set of all possible product terms. A *probability expression template* ξ is a sum-of-product expression of the form $\sum_{\psi_i \in \Psi} \square_i \times \psi_i$.

For instance, if we have two correlation classes each with a single input fact, the probability expression template would be of the form:

$$\square_1 \times V_1[0]V_2[0] + \square_2 \times V_1[0]V_2[1] + \square_3 \times V_1[1]V_2[0] + \square_4 \times V_1[1]V_2[1]$$

Intuitively, the probability of every relation in the program can be expressed symbolically as an instantiation of a probability expression template, where holes \square are filled by *coefficient terms* λ :

Definition 9. (Coefficient term) A coefficient term λ is an expression of the form $\sum_i \lambda_i^+ \times \lambda_i^-$ where $\lambda_i^+ = r_1 \times \dots \times r_n$ and $\lambda_i^- = (1 - r'_1) \times \dots \times (1 - r'_n)$ and each r_j, r'_j is a *variable* representing the probability of some rules in the Datalog program.

A coefficient term represents the probability of an output fact being derived by applying a specific sequence of rules. However, since rule bodies can include negations, the derivation of some facts may depend on certain rules *not* being applied. As a result, coefficient terms also incorporate factors of the form $(1 - r)$, which represent the probability of a rule *not* being applied.

$$\begin{aligned}
 \llbracket D \rrbracket &= \bigwedge_{C_i \in C} \llbracket C_i \rrbracket \wedge \bigwedge_{I_r \in P_I} \llbracket I_r \rrbracket \wedge \bigwedge_{r \in R} \text{ProbVar}(r) = P_R(r) \\
 \llbracket C \rrbracket &= \forall b \in \mathbb{B}^{|C|}. \text{Rep}(C)[b] \in [0, 1] \wedge \sum_{b \in \mathbb{B}^{|C|}} \text{Rep}(C)[b] = 1 \\
 \llbracket I_r \rrbracket &= \text{Expr}(I) = p && \text{if } I_r = (p :: I | \emptyset) \\
 \llbracket I_r \rrbracket &= \otimes_{i=0}^n \text{Expr}(I_i) = p \times \otimes_{i=1}^n \text{Expr}(I_i) && \text{if } I_r = (p :: I_0 | I_1, \dots, I_n)
 \end{aligned}$$

Fig. 8. Rules for producing constraints from a Datalog program $D = (C, R, P_R, P_I)$. $\text{ProbVar}(r)$ is a fresh variable representing the probability of rule r .

Definition 10. (Template instantiation) An *instantiation* of template ξ either replaces each hole with a $c \in \{0, 1\}$ or with a coefficient term λ . Given a vector of expressions σ , we use the notation $\xi[\sigma]$ to denote a probability expression that is obtained by filling hole \square_i in ξ with expression σ_i .

Intuitively, template instantiations that fill holes with coefficient terms represent probabilities of output facts whereas those that use constants represent probabilities of input relations. Given a relation R , $\text{Expr}(R)$ denotes the symbolic expression representing the probability of R and is obtained through some instantiation σ of the holes in ξ — i.e., $\text{Expr}(R) = \xi[\sigma]$ for some σ .

Definition 11. (Input expression) Let ξ be a template with ψ_i as its i 'th term. The *probability expression* for input fact I , denoted $\text{Expr}(I)$, is given by $\sigma_i = 1$ if $\psi_i \models I$ and $\sigma_i = 0$ otherwise.

In other words, the symbolic expression for an input relation is obtained by choosing 1 as the coefficient of a product term ψ_i if ψ_i represents the probability of an event in which I is true and 0 otherwise. It is easy to see that $\text{Expr}(I)$ symbolically represents the probability of an input relation in terms of the joint probability variables.

6.3 Constraint Generation

We now turn to the `GENCONSTRAINT` procedure used in Algorithm 1, with its implementation summarized in Figure 8. Given a program $D = (C, R, P_R, P_I)$, this procedure generates three types of constraints. First, for each correlation class C_i with representative V_i , it introduces a constraint ensuring that all variables in V_i are valid probabilities in the range $[0, 1]$ and that they sum to 1 (line 2 of Figure 8). Second, the procedure introduces constraints to encode the (conditional) input probabilities in the Datalog program. Specifically, for each rule of the form $p :: I$, a constraint is introduced stating that $\text{Expr}(I) = p$, where $\text{Expr}(I)$ is defined in Definition 11. Additionally, for conditional probability declarations of the form $p :: I_0 | I_1 \dots I_n$ (line 4 of Figure 8), a constraint is introduced stating that $P(I_0 \wedge \dots \wedge I_n) = p \times P(I_1 \wedge \dots \wedge I_n)$. Note that the probability $P(I_0 \wedge \dots \wedge I_n)$ is computed using a special multiplication operator \otimes , which is explained in the next section. Finally, the procedure generates a third type of constraint that relates variables in the coefficient terms to the actual rule probabilities given by P_R (the last conjunct in the first line).

6.4 Generating Optimization Objective

In this final subsection, we focus on the `GENOBJECTIVE` procedure. As a reminder, the purpose of this procedure is to express the probability of each output relation as a symbolic expression over the joint probability variables—i.e., as an instantiation of the probability expression template defined in Definition 10. The key idea behind this procedure is to use the logical derivation graph G to find a coefficient term that can be used to fill each hole in the expression template. In particular, the algorithm performs bottom-up traversal of the derivation graph to construct a symbolic probability expression of each node, utilizing the probability expressions of its children.

Our bottom-up traversal algorithm is presented in Figure 9. The base case is the rule labeled `LEAF`, which corresponds to input relations. Since we already know how to compute the probability expression for an input relation (see Definition 11), the `LEAF` rule simply produces $\text{Expr}(I)$, where

$$\begin{array}{c}
\frac{v \in \text{Leaves}(G) \quad \text{Pred}(v) = I}{G \vdash v \hookrightarrow \text{Expr}(I)} \text{ (LEAF)} \\
\\
\frac{\text{OutgoingEdges}(G, v) = \{e_1, \dots, e_n\} \quad G \vdash e_i \hookrightarrow E_i \text{ for } i \in [1, n] \quad E_1 \oplus \dots \oplus E_n \rightsquigarrow E}{G \vdash v \hookrightarrow E} \text{ (NODE)} \\
\\
\frac{B^+ = \{v_1^+, \dots, v_m^+\} \quad G \vdash v_i^+ \hookrightarrow E_i^+ \text{ for } i \in [1, m] \quad B^- = \{v_1^-, \dots, v_n^-\} \quad G \vdash v_i^- \hookrightarrow E_i^- \text{ for } i \in [1, n] \quad (E_1^+ \otimes \dots \otimes E_m^+) \rightsquigarrow E^+ \quad (\ominus(E_1^-)) \otimes \dots \otimes (\ominus(E_n^-)) \rightsquigarrow E^- \quad E^+ \otimes E^- \rightsquigarrow E}{G \vdash (_, B^+, B^-, r) \hookrightarrow \text{ProbVar}(r) \times E} \text{ (EDGE)}
\end{array}$$

Fig. 9. Rules for inferring symbolic probability expression for output relations.

$$\begin{array}{c}
\frac{E = \sum \lambda_i \psi_i}{\ominus(E) \rightsquigarrow \sum (1 - \lambda_i) \psi_i} \text{ (NEG)} \quad \frac{E_1 = \sum \lambda_{1i} \psi_i \quad E_2 = \sum \lambda_{2i} \psi_i \quad \text{JointProb}(\lambda_{1i}, \lambda_{2i}) \rightsquigarrow \lambda_i}{\vdash E_1 \otimes E_2 \rightsquigarrow \sum \lambda_i \psi_i} \text{ (MUL)} \quad \frac{E_1 = \sum \lambda_{1i} \psi_i \quad E_2 = \sum \lambda_{2i} \psi_i \quad \text{JointProb}(\lambda_{1i}, \lambda_{2i}) \rightsquigarrow \lambda_i}{E_1 \oplus E_2 \rightsquigarrow \sum (\lambda_{1i} + \lambda_{2i} - \lambda_i) \psi_i} \text{ (ADD)} \\
\\
\frac{\lambda_1 = \sum_{i=1}^n e_i \quad \lambda_2 = \sum_{j=1}^m e_j \quad \forall (i, j). e_i \diamond e_j \rightsquigarrow e_{ij}}{\text{JointProb}(\lambda_1, \lambda_2) \rightsquigarrow \sum_{(i,j)=1}^{(n,m)} e_{ij}} \text{ (JOINT)} \quad \frac{e_i = X_i^+ \times X_i^- \quad e_j = Y_j^+ \times Y_j^- \quad X_i^+ \diamond Y_j^+ \rightsquigarrow Z^+ \quad X_i^- \diamond Y_j^- \rightsquigarrow Z^-}{e_i \diamond e_j \rightsquigarrow \text{Disjoint}(X_i, Y_j) ? Z^+ Z^- : 0} (\diamond) \\
\\
\frac{\prod_{i=1}^p \bigcirc(x_i) \diamond \prod_{j=1}^q \bigcirc(y_j) \rightsquigarrow \prod_{z_k \in (\cup_{i=1}^p x_i \cup \cup_{j=1}^q y_j)} \bigcirc(z_k)}{(\diamond)}
\end{array}$$

Fig. 10. Rules defining \ominus , \otimes and \oplus operations. $\text{Disjoint}(X_i, Y_j)$ is true iff $\text{Vars}(X_i^+) \cap \text{Vars}(Y_j^-) = \emptyset \wedge \text{Vars}(X_i^-) \cap \text{Vars}(Y_j^+) = \emptyset$. $\text{Vars}(X)$ denotes the set of rule probability variables present in X . In the \diamond rule, $\bigcirc(x)$ represents a term of the form x or $1 - x$.

I is the input relation represented by the node v . Next, the rule labeled NODE describes how to compute the probability expression for an internal node v representing an output relation. If a node labeled O has n outgoing edges, this means O can be derived in n different ways. Therefore, the rule first computes a probability expression E_i for each edge e_i , representing the probability of a possible derivation of O . It then computes the probability of $P(e_1 \vee \dots \vee e_n)$ using a special \oplus operator, which we explain later. Finally, the last rule, labeled EDGE, computes the probability expression for a derivation. Recall that a hyperedge with source v and target nodes B^+ and B^- represents the application of a rule r , whose body consists of positive facts B^+ and negative facts B^- . This rule first computes the probability expressions E_i^+ for each $v_i^+ \in B^+$ and E_i^- for each $v_i^- \in B^-$, respectively. It then combines these to obtain a probability expression for the edge.

To combine these expressions, we first note that the probability expression for $\neg R$ is given by $\ominus E$, where E denotes the probability expression for R , i.e., $E = \text{Expr}(R)$. Next, we define an operator \otimes (explained later) to compute the probability of $P(v_1 \wedge v_2)$ as $P(v_1) \otimes P(v_2)$. Therefore, the EDGE rule first computes the probability of all positive facts being true as $E^+ = E_1^+ \otimes \dots \otimes E_m^+$, and the probability of all negative facts being true as $E^- = \ominus(E_1^-) \otimes \dots \otimes \ominus(E_n^-)$. The probability of both positive and negative facts being satisfied is then given by $E = E^+ \otimes E^-$. Finally, since the rule itself has a probability, the final probability expression is obtained as $\text{ProbVar}(r) \times E$, where $\text{ProbVar}(r)$ is the variable representing the probability of rule r .

Finally, we turn our attention to the definitions of the \ominus , \otimes , and \oplus operators from Figure 10, which are used to compute probability expressions for conjunctions and disjunctions of events.

$p_1 :: I_1$	$Pr(I_1)$	$= 0 * V[00] + 0 * V[01] + 1 * V[10] + 1 * V[11]$
$p_2 :: I_2$	$Pr(I_2)$	$= 0 * V[00] + 1 * V[01] + 0 * V[10] + 1 * V[11]$
$p_3 :: I_2 I_1$		
$r_1 :: A :- I_1$	$Pr(A)$	$= 0 * V[00] + 0 * V[01] + r_1 * V[10] + r_1 * V[11]$
$r_2 :: B :- A$	$Pr(B)$	$= 0 * V[00] + 0 * V[01] + r_1 r_2 * V[10] + r_1 r_2 V[11]$
$r_3 :: C :- \neg A, I_2$	$Pr(\neg A)$	$= 1 * V[00] + 1 * V[01] + (1 - r_1) V[10] + (1 - r_1) V[11]$
	$Pr(C)$	$= r_3 \times (Pr(\neg A) \otimes Pr(I_2)) = 0 * V[00] + r_3 * V[01] + 0 * V[10] + r_3(1 - r_1) V[11]$
$r_4 :: D :- B, A$	$Pr(D)$	$= r_4 \times (Pr(B) \otimes Pr(A)) = 0 * V[00] + 0 * V[01] + \textcolor{red}{r_1 r_2 r_4} V[10] + \textcolor{red}{r_1 r_2 r_4} V[11]$
$r_5 :: E :- C$	$Pr(E_1)$	$= 0 * V[00] + r_3 r_5 * V[01] + 0 * V[10] + r_3 r_5 (1 - r_1) V[11]$
$r_6 :: E :- D$	$Pr(E_2)$	$= 0 * V[00] + 0 * V[01] + r_1 r_2 r_4 r_6 V[10] + r_1 r_2 r_4 r_6 V[11]$
	$Pr(E)$	$= Pr(E_1) \oplus Pr(E_2)$
		$= 0 * V[00] + r_3 r_5 V[01] + r_1 r_2 r_4 r_6 V[10] + (r_3 r_5 (1 - r_1) + \textcolor{red}{r_1 r_2 r_4 r_6} - 0) V[11]$

Fig. 11. Left: PRALINE program. Right: Probability expressions for both input and output facts.

All of these rules rely on the fact that probability expressions are in a normalized form, consisting of sums of terms of the form $\lambda \times \psi$, where λ is a coefficient term (Def. 9) and ψ is a product term (Def. 7). First, given a symbolic expression E representing the probability of some relation R , the NEG rule computes the probability of $\neg R$ by simply replacing all coefficient terms λ with $1 - \lambda$. Second, the MUL rule computes the probability of a conjunction of events by combining the coefficients of each product term using the JointProb function (explained later). Similarly, the ADD rule computes the probability of a disjunction of events by updating the coefficient of each term ψ_i as $\lambda_{1i} + \lambda_{2i} - \text{JointProb}(\lambda_{1i}, \lambda_{2i})$. Intuitively, this corresponds to an application of the inclusion-exclusion principle.

Next, we focus on the last three rules for computing the joint probability of two coefficient terms λ_1 and λ_2 . The rule labeled JOINT essentially distributes multiplication over addition, as we require each coefficient term to be a sum of products. The second rule, labeled \diamond , considers multiplication expressions of the form $(X_i^+ \times X_i^-)$ and $(Y_i^+ \times Y_i^-)$. In this rule, Z^+ denotes a product of variables v , and Z^- represents a product of terms of the form $1 - v$, where v represents a rule probability. Intuitively, Z^+ indicates the rules that must be applied to derive a fact, while Z^- indicates the rules that must not be applied. Thus, if a variable v appears in X_i^+ (or Y_i^+) and $1 - v$ appears in Y_j^- (or X_i^-), this results in a contradiction, and the probability of the term is zero. Otherwise, if the disjointness condition is satisfied, the probability is computed using the rule labeled \blacklozenge . The intuition behind the \blacklozenge rule is as follows: If a variable v appears in both X_i and Y_i , it indicates that the same fact is derived using the same rule. To avoid overcounting, we do not multiply probabilities repeatedly, as that rule only needs to be applied once. Therefore, the \blacklozenge rule multiplies probabilities after ensuring that variables associated with the same rule are treated uniquely.

Example 3. Consider the PRALINE program on the left side of Figure 11, with the corresponding probability expressions for both input and output facts displayed on the right. In this example, I_1 and I_2 are correlated, so they belong to the same correlation class V — e.g., $V[01]$ represents the joint probability of I_1 being false and I_2 being true. The right side of Figure 11 shows the symbolic probability expressions for each fact. The parts highlighted in red show the need for the disjointness check in the \diamond rule; and the parts highlighted in blue illustrate the need for the \blacklozenge operator.

THEOREM 1. *Let E_1 and E_2 denote the probability expressions of events A and B respectively. Then, we have: (1) If $\ominus E_1 \rightsquigarrow E$, then E represents the probability of $\neg A$; (1) if $E_1 \oplus E_2 \rightsquigarrow E$, then E represents the probability of event $A \vee B$ (3) iff $E_1 \otimes E_2 \rightsquigarrow E$, then E represents the probability of event $A \wedge B$.*

7 δ -Exact Probabilistic Inference

Building on our approach from Section 6, we now present a more practical δ -exact probabilistic inference algorithm, summarized in Algorithm 2. Similar to the previous algorithm, it first uses the

Algorithm 2 $\text{Solve}^{\pm\delta}(D, \delta)$ **Input:** Datalog program D , precision bound δ **Output:** Mapping M that maps output facts to probabilities

- 1: $S_o, G \leftarrow \text{SolveStandard}(D)$ ▷ SolveStandard produces output facts S_o and derivation graph G
- 2: $\phi \leftarrow \text{GENCONSTRAINTS}(D, G)$ ▷ Use technique from Section 6.3 to generate constraints
- 3: $\Sigma \leftarrow \text{INFERCORRELATIONTYPE}(G, \phi)$ ▷ Infer type of statistical correlation between facts
- 4: $M \leftarrow \text{DERIVEAPPROXIMATEBOUNDS}(G, D, \Sigma)$ ▷ Approximate bounds using static analysis and constraint solving
- 5: $M' \leftarrow \text{MAKEDELTAPRECISE}(M, S_o, \phi, \delta)$ ▷ Iterative refinement of bounds
- 6: **return** M'

$$\begin{array}{c}
\frac{p :: (I|\emptyset) \in P_I \quad p < 1}{D, \phi \vdash I \blacktriangleright^+ I} \text{ (ID)} \quad \frac{\text{Class}(I_1) \neq \text{Class}(I_2)}{D, \phi \vdash I_1 \blacktriangleright^\perp I_2} \text{ (INDEP)} \quad \frac{I_1 \blacktriangleright^\star I_2 \quad (\star \in \{+, -, \perp\})}{D, \phi \vdash I_2 \blacktriangleright^\star I_1} \text{ (SYMM)} \\
\\
\frac{E_1 = \text{Expr}(I_1) \quad E_2 = \text{Expr}(I_2) \quad E_1 \otimes E_2 \rightsquigarrow E_\wedge \quad \models \phi \Rightarrow E_\wedge \sqsubseteq E_1 \times E_2 \quad \star = \text{Sign}(\sqsubseteq)}{D, \phi \vdash I_1 \blacktriangleright^\star I_2} \text{ (SEMANTIC)}
\end{array}$$

Fig. 12. Inference of statistical correlation between input variables. Here, $\sqsubseteq \in \{<, >, =\}$, and $\text{Sign}(\sqsubseteq)$ yields $+$ for $>$, $-$ for $<$, and \perp for $=$.

underlying Datalog solver to obtain a derivation graph G . Next, it invokes the $\text{INFERCORRELATIONTYPE}$ function to compute the *type* of the statistical correlation between facts, where a correlation type is either positive, negative, independent, or unknown. In the third step, the algorithm uses this correlation type environment Σ to derive *approximate* probability bounds on output facts. Finally, the call to MAKEDELTAPRECISE at line 5 keeps refining the inferred bounds until the derived bounds are within δ of the ground truth.

7.1 Inference of Correlation Types

In this section, we present the $\text{INFERCORRELATIONTYPE}$ algorithm that can be used to infer whether a pair of relations are positively/negatively correlated or whether they are independent.

Definition 12. (Statistical correlation) Two events X and Y are positively (resp. negatively) correlated if $P(X|Y) > P(X)$ (resp. $P(X|Y) < P(X)$). If $P(X|Y) = P(X)$, then X and Y are independent.

Note that the notion of statistical correlation is symmetric. That is, if X is positively correlated with Y , then Y is also positively correlated with X . A *correlation type* for a pair of Datalog facts is one of Pos ($+$), Neg ($-$), \perp , \top , where \perp , \top indicate independence and unknown correlation respectively. As stated earlier, identifying correlation types allows us to derive tighter approximate bounds than would otherwise be possible. To infer these correlation types, our method proceeds in two phases: First, it infers statistical correlations between input facts. In the second phase, it uses this information to infer statistical correlations between outputs.

Phase 1: Inferring correlation types between input facts. Figure 12 presents our method for inferring correlation types for pairs of input facts using the judgment $D, \phi \vdash I_1 \blacktriangleright^\star I_2$ where $\star \in \{+, -, \perp\}$ indicates positive and negation correlation and statistical independence respectively. Here, D is the Datalog program and ϕ is the set of constraints generated from the Datalog program as described in Section 6.3. The first rule, labeled ID indicates that the input fact I is positively correlated with itself. The next rule, labeled INDEP applies to two input facts that do not belong to the same correlation class. Finally, the last rule labeled SEMANTIC uses the constraint-based technique from Section 6 to check for statistical correlation. The basic idea is to use the algorithm

$\text{Dep}(O)$	$= \{I \mid I \triangleright O\}$	$\text{Dep}^\star(O)$	$= \{I \mid I \triangleright^\star O\}$
$\text{Dep}(E_1 \wedge E_2)$	$= \text{Dep}(E_1) \cup \text{Dep}(E_2)$	$\text{Dep}^\star(E_1 \wedge E_2)$	$= \text{Dep}^\star(E_1) \cup \text{Dep}^\star(E_2)$
$\text{Dep}(E_1 \vee E_2)$	$= \text{Dep}(E_1) \cup \text{Dep}(E_2)$	$\text{Dep}^\star(E_1 \vee E_2)$	$= \text{Dep}^\star(E_1) \cup \text{Dep}^\star(E_2)$
$\text{Dep}(\neg E)$	$= \text{Dep}(E)$	$\text{Dep}^\star(\neg E)$	$= \text{Dep}^\star(E)$

Fig. 13. Auxiliary Dep function. O denotes an output fact and $\star \in \{+, -\}$. $\bar{\star}$ is + if \star is - and vice versa.

$$\begin{array}{c}
\frac{\forall (x, y) \in \text{Dep}(E). x \neq y \rightarrow \text{Class}(x) \neq \text{Class}(y)}{\vdash \chi(E)} \text{ (NONINTERFERE)} \\
\\
\frac{\frac{\exists (x, y) \in \text{Dep}^\star(E_1) \times \text{Dep}^\star(E_2). x \triangleright^{\bar{h}} y}{\vdash E_1 \multimap^{\bar{h}} E_2} \text{ (MAY-1)} \quad \frac{\exists (x, y) \in \text{Dep}^\star(E_1) \times \text{Dep}^{\bar{\star}}(E_2). x \triangleright^{\bar{h}} y}{\vdash E_1 \multimap^{\bar{h}} E_2} \text{ (MAY-2)}}{\vdash \chi(E_1) \quad \vdash \chi(E_2) \quad \vdash E_1 \multimap^+ E_2 \quad \not\vdash E_1 \multimap^- E_2} \text{ (POS)} \quad \frac{\vdash \chi(E_1) \quad \vdash \chi(E_2) \quad \vdash E_1 \multimap^- E_2 \quad \not\vdash E_1 \multimap^+ E_2}{\vdash E_1 \multimap E_2} \text{ (NEG)} \\
\\
\frac{\forall x, y \in \text{Dep}(E_1) \times \text{Dep}(E_2). \text{Class}(x) \neq \text{Class}(y)}{\vdash E_1 \perp E_2} \text{ (INDEP)} \quad \frac{\not\vdash E_1 \multimap E_2 \quad \not\vdash E_1 \multimap E_2 \quad \not\vdash E_1 \perp E_2}{\vdash E_1 \multimap E_2} \text{ (UNKNOWN)}
\end{array}$$

Fig. 14. Inference rules for computing statistical correlation between expressions, where $\bar{h} \in \{+, -\}$. Note that predicates $x \triangleright^\star y$ are derived using Figure 12, and $\bar{\star}$ (resp. \bar{h}) is + if \star (resp. h) is - and vice versa.

of Section 6 to check how the joint probability of $I_1 \wedge I_2$ relates to the product of the individual probabilities of I_1 and I_2 . To do so, it first generates symbolic expressions E_1, E_2, E_\wedge for I_1, I_2 , and $I_1 \wedge I_2$ respectively. Then, it uses a solver to check whether the constraints ϕ (encoding the input probabilities) logically imply whether $E_\wedge \sqcap E_1 \times E_2$, where \sqcap denotes one of $<, >, =$. Note that this semantic approach is not as susceptible to the scalability challenges discussed earlier because we consider *only* input facts and *only* those that belong to the same correlation class.

THEOREM 2. *Suppose that we derive $I_1 \triangleright^\star I_2$ using the rules from Figure 12. Then, $p(I_1|I_2) > p(I_1)$ if $\star = +$, $p(I_1|I_2) < p(I_1)$ if $\star = -$, and $p(I_1|I_2) = p(I_1)$ if $\star = \perp$.*

Phase 2: Inferring correlation types between outputs. The second phase of our inference algorithm uses the results of the first phase to infer correlation types between expressions involving outputs. Note that we could, *in principle*, use the same constraint-based approach presented Figure 12 to infer correlation types between arbitrary expressions; however, such an approach does not scale well. To overcome this scalability bottleneck, we instead utilize lightweight static analysis.

The key idea underlying our method is to utilize the derivation graph, along with the known correlations between input facts, to infer correlation types between arbitrary expressions (i.e., boolean combinations of ground predicates). Given a pair of expressions E_1, E_2 , our method first uses the derivation graph to identify the set S_1, S_2 of input facts, along with their polarity, that E_1 and E_2 *logically* depend on; it then analyzes the *statistical* correlations between elements in S_1, S_2 to decide whether we can determine the correlation type between E_1 and E_2 .

Our analysis is summarized in Figures 13 and 14, where the former defines two auxiliary functions $\text{Dep}, \text{Dep}^\star (\star \in \{+, -\})$ used in Figure 14. As shown in Figure 13, $\text{Dep}(E)$ simply yields the set of all input facts that E is *logically* dependent on according to the derivation graph (recall Def 3). Similarly, $\text{Dep}^+(E)$ (resp. $\text{Dep}^-(E)$) yields the set of input facts that E depends *positively* (resp. *negatively*) on. For example, consider an output fact O that can be derived using $I_1 \wedge \neg I_2$ or using only I_2 . In this case, both $\text{Dep}(O)$ and $\text{Dep}^+(O)$ include I_1, I_2 but $\text{Dep}^-(O)$ only includes I_2 .

Figure 14 uses these auxiliary functions to infer correlation types between arbitrary expressions. The basic idea is to infer whether two expressions E_1 and E_2 *may* be positively or negatively

Algorithm 3 DERIVEAPPROXIMATEBOUNDS(G, D, Σ)

```

1: for node  $n \in \text{InternalNodes}(G)$  do
2:    $E_n \leftarrow \llbracket n \rrbracket(G)$ 
3:    $M[n] \leftarrow \text{APPROXEXPR}(\Sigma, D, E_n)$ 
4: return  $M$ 

```

Fig. 15. Computation of derivation expressions. $\mathbb{P}(e)$ yields the probability of the rule labeling edge e .

Operation	$\star = +$	$\star = -$	$\star = \perp$	$\star = \top$
CL(e_1, e_2, \star)	$e_1 \times e_2$	$\max(e_1 + e_2 - 1, 0)$	$e_1 \times e_2$	$\max(e_1 + e_2 - 1, 0)$
CU(e_1, e_2, \star)	$\min(e_1, e_2)$	$e_1 \times e_2$	$e_1 \times e_2$	$\min(e_1, e_2)$
DL(e_1, e_2, \star)	$\max(e_1, e_2)$	$1 - (1 - e_1)(1 - e_2)$	$1 - (1 - e_1)(1 - e_2)$	$\max(e_1, e_2)$
DU(e_1, e_2, \star)	$1 - (1 - e_1)(1 - e_2)$	$\min(1, e_1 + e_2)$	$1 - (1 - e_1)(1 - e_2)$	$\min(1, e_1 + e_2)$

Table 1. CL/CU/DL/DU computation rules. CL and CU denote the lower and upper bounds of the conjunction operation, respectively, while DL and DU represent the lower and upper bounds of the disjunction operation.

$$\begin{array}{c}
\frac{I \in \text{InputFacts}(D) \quad p :: (I \mid \emptyset) \in \text{InputProbs}(D)}{\Sigma, D \vdash I \rightsquigarrow [p, p]} \text{ (IN)} \quad \frac{\Sigma, D \vdash E \rightsquigarrow [l, u]}{\Sigma, D \vdash \neg E \rightsquigarrow [1 - u, 1 - l]} \text{ (NEG)} \\
\\
\frac{\Sigma, D \vdash E_1 \rightsquigarrow [l_1, u_1] \quad \Sigma, D \vdash E_2 \rightsquigarrow [l_2, u_2] \quad \Sigma(E_1, E_2) = \star}{\Sigma, D \vdash E_1 \wedge E_2 \rightsquigarrow [\text{CL}(l_1, l_2, \star), \text{CU}(u_1, u_2, \star)]} \text{ (CONJUNCT)} \\
\\
\frac{\Sigma, D \vdash E_1 \rightsquigarrow [l_1, u_1] \quad \Sigma, D \vdash E_2 \rightsquigarrow [l_2, u_2] \quad \Sigma(E_1, E_2) = \star}{\Sigma, D \vdash (E_1, p_1) \vee (E_2, p_2) \rightsquigarrow [\text{DL}(l_1 \times p_1, l_2 \times p_2, \star), \text{DU}(u_1 \times p_1, u_2 \times p_2, \star)]} \text{ (DISJUNCT)}
\end{array}$$

Fig. 16. Inference rules for computing approximate probability bounds.

correlated (rules labeled MAX), meaning that a pair of shared input predicates in E_1 and E_2 have the potential to introduce a positive or negative correlation. Then, according to the rules labeled POS and NEG, if we find that E_1 and E_2 may be positively (resp. negatively) correlated and there is nothing that introduces a potential negative (resp. positive) correlation, we can conclude that E_1 and E_2 are definitely positively (resp. negatively) correlated as long as both expressions exhibit a certain non-interference property shown in the NONINTERFERE rule as $\chi(E)$. Intuitively, the non-interference property is necessary because, if an input fact I is positively correlated with I_1 and I_2 individually, it does not necessarily mean that it is positively correlated with $I_1 \wedge I_2$. At the end of the correlation type analysis, the inferred dependencies are stored in Σ and used in Algorithm 2.

THEOREM 3. If $E_1 \xrightarrow{\odot} E_2$ is derivable using the rules in Figure 14, then $p(E_1|E_2) > p(E_1)$. Similarly, $\vdash E_1 \xrightarrow{\ominus} E_2$ implies $p(E_1|E_2) < p(E_1)$ and $\vdash E_1 \xrightarrow{\perp} E_2$ implies $p(E_1|E_2) = p(E_1)$.

7.2 Computing Approximate Probability Bounds

In this section, we present a technique, summarized in Algorithm 3, for deriving *approximate* probability bounds on output relations. For each node in the derivation graph, this algorithm computes a so-called *derivation expression* E that summarizes all ways in which a given relation can be derived (line 2). For example, if there are two rules $p_1 :: R :- A, B$ and $p_2 :: R :- C$, then the derivation expression is of the form $(A \wedge B) \vee C$. However, because we also need to keep track of the rule probabilities, expressions inside a disjunct also have a corresponding probability, represented as $(A \wedge B, p_1) \vee (C, p_2)$ for this example. Figure 15 presents the rules for generating derivation expressions for each node. Then, given the derivation expression E_n for node n , the

Algorithm 4 MAKEDELTAPRECISE(M, S_o, ϕ, δ)**Input:** Mapping M , set of output relations S_o , constraints ϕ , precision bound δ **Output:** Mapping M that maps output facts to probabilities

```

1:  $B \leftarrow \text{BOUNDBOUNDS}(M, \phi, S_o)$ 
2:    $\triangleright$  Compute upper and lower bounds for the bounds
3: for node  $R \in S_o$  do
4:    $(l^-, l^+, u^-, u^+) \leftarrow B[R]$ 
5:    $(l^-, l^+) \leftarrow \text{BINARYSEARCH}(l^-, l^+, \phi, \delta, R, \text{true})$ 
6:    $(u^-, u^+) \leftarrow \text{BINARYSEARCH}(u^-, u^+, \phi, \delta, R, \text{false})$ 
7:    $M[R] \leftarrow (l^-, u^+)$ 
8: return  $M$ 

```

Algorithm 5 BOUNDBOUNDS(M, ϕ, S_o)**Input:** Mapping M , constraints ϕ , set of output relations S_o **Output:** A new mapping B , where $B[n]$ is a quadruple (l^-, l^+, u^-, u^+) where l^-, l^+ (resp. u^-, u^+) are lower and upper bounds for the ground truth lower (resp. upper) bound

```

1: for node  $R \in S_o$  do
2:    $(l, u) \leftarrow M[R]$ 
3:    $(l^-, l^+) \leftarrow \text{MAKESAT}(l, l, \phi, \text{Expr}(R), \text{true})$ 
4:    $(u^-, u^+) \leftarrow \text{MAKESAT}(u, u, \phi, \text{Expr}(R), \text{false})$ 
5:    $B[R] \leftarrow (l^-, l^+, u^-, u^+)$ 
6: return  $B$ 

```

APPROXEXPR procedure (called at line 3 in Algorithm 3) computes the approximate upper and lower bounds for E_n using the rules presented in Figure 16. Given a derivation expression E , these rules derive judgments of the form $\Sigma, D \vdash E \rightsquigarrow [l, u]$ indicating that l, u are lower and upper bounds on the probability of expression E evaluating to true. To compute these lower and upper bounds, we leverage the results of the correlation type analysis (stored in Σ) as well as known statistical inequalities provided in Table 1 for different correlation types [46].

THEOREM 4. *Let D be a PRALINE program with derivation graph G , and suppose Σ is a sound correlation environment for D . Also, let $M' = \text{DERIVEAPPROXBOUNDS}(G, D, \Sigma)$ and let $M = \text{SOLVE}(D)$ (Algorithm 1). For every output relation O of D such that $M(O) = (l^*, u^*)$ and $M'(O) = (l, u)$, we have $l \leq l^* \leq u^* \leq u$.*

7.3 Iterative Refinement of Probability Bounds

In this section, we describe the MAKEDELTAPRECISE algorithm that iteratively tightens the computed probability bounds until it is within some δ of the ground-truth. The key idea is to combine the algorithm from Section 6 with the approximate bounds as illustrated in Figure 7. In this Figure 7, l^* and u^* denote the ground truth (but unknown) probability bounds for relation R , and l and u denote the approximate probability bounds for R , computed as described in Section 7.2. Thus, it is always the case that $l \leq l^*$ and $u^* \leq u$. Our key observation is that the constraint ϕ generated in Section 6.3 partitions this space into three regions:

- **Region 1:** This is the region $\psi_1 = (l \leq \text{Expr}(R) < l^*)$, where $\text{Expr}(R)$ denotes the symbolic expression generated for R , as described in Section 6.4. Since the ground truth is $l^* \leq \text{Expr}(R) \leq u^*$, $\psi_1 \wedge \phi$ must be unsatisfiable.
- **Region 2:** This is the “ground truth” region $\psi_2 = l^* \leq \text{Expr}(R) \leq u^*$; thus, $\phi \wedge \psi_2$ is satisfiable.
- **Region 3:** This is the region $\psi_3 = (u^* < \text{Expr}(R) \leq u)$, so $\psi_3 \wedge \phi$ is again unsatisfiable.

As illustrated in Figure 7, the idea is to repeatedly increase l (resp. u) by ϵ until the formula $l_i \leq \text{Expr}(R) \leq l_i + \epsilon \wedge \phi$ (resp. $u_i - \epsilon \leq \text{Expr}(R) \leq u_i \wedge \phi$) becomes satisfiable. When this procedure terminates, we can obtain lower and upper bounds (l^-, l^+) for l^* as well as bounds (u^-, u^+) for u^* . We can then perform binary search until the distance between the two becomes less than δ .

This discussion is summarized in Algorithm 4. The MAKEDELTAPRECISE procedure first calls BOUNDBOUNDS to compute upper and lower bounds for l^*, u^* , as depicted in Figure 7. As shown in Algorithm 5 (and its auxiliary procedure MAKESAT in Algorithm 6), BOUNDBOUNDS increments (resp. decrements) l (resp. u) until we get into the SAT region in Figure 7. Upon termination of MAKESAT, (l^-, l^+) (resp. (u^-, u^+)) provide lower and upper bounds for l^* (resp. u^*). Then, for each relation R , MAKEDELTAPRECISE calls BINARYSEARCH (Algorithm 7) to find a δ -optimal solution.

Algorithm 6 MAKESAT(l, u, ϕ, e, b)

Input: Lower bound l , upper bound u , constraints ϕ , expression e , boolean flag b . Boolean flag b indicating if we are dealing with a lower or upper bound

Output: Updated lower bound l and upper bound u

```

1: while UNSAT( $\phi \wedge l \leq e \leq u$ ) do
2:    $(l, u) \leftarrow b ? (u, \text{Inc}(u))$ 
3:   :  $(\text{Dec}(l), l)$ 
4: return  $(l, u)$ 

```

Algorithm 7 BINARYSEARCH(l, u, ϕ, δ, R, b)

Input: Lower bound l , upper bound u , constraints ϕ , error bound δ , output relation R , boolean flag b

Output: Updated lower bound l and upper bound u

```

1: while  $(u - l) \geq \delta$  do
2:    $\text{mid} \leftarrow (l + u) / 2$ 
3:    $(l', u') \leftarrow \text{low?}(l, \text{mid}) : (\text{mid}, u)$ 
4:   if UNSAT( $\phi \wedge (l' \leq \text{Expr}(R) \leq u')$ ) then
5:      $(l, u) \leftarrow \text{low?}(\text{mid}, u) : (l, \text{mid})$ 
6:   else  $(l, u) \leftarrow \text{low?}(l, \text{mid}) : (\text{mid}, u)$ 
7: return  $(l, u)$ 

```

When BINARYSEARCH terminates, the returned interval $[l^-, l^+]$ is guaranteed to contain l^* , with $l^+ - l^- \leq \delta$. The same guarantee also applies to the upper bound.

THEOREM 5. *Let (l^*, u^*) be the ground truth probability bounds for relation R , and let (l, u) be the bounds computed by MAKEDELTA-PRECISE. Then, we have $l^* - \delta \leq l \leq l^*$ and $u^* \leq u \leq u^* + \delta$.*

8 Implementation

We have implemented our proposed approach in a tool called PRALINE written in C++. PRALINE instruments the solving procedure of SOUFFLE [30] to generate the derivation graph and utilizes the GUROBI [10] solver for optimization and the CVC5 [6] SMT solver for satisfiability.

Derivation graph generation. Datalog solvers such as SOUFFLE avoid generating the same relation from different rules. For instance, if an output relation O has already been derived using a rule R , the Datalog solver would avoid applying other rules to derive O again. However, to compute the probability of O , we need all possible ways of deriving it; thus, our implementation modifies SOUFFLE to generate the complete derivation graph. It also augments the derivation graph to keep track of rule probabilities.

Inference of correlation classes. While PRALINE allows the user to explicitly specify correlation classes (e.g., via the `corr` declaration), it does not require them to do so. In particular, PRALINE constructs a dependency graph between input facts based on the specified conditional probabilities and assumes that a pair of input facts are in the same correlation class iff they belong to the same connected component. This default behavior can be overridden by users by explicitly specifying correlation classes.

Optimized satisfiability checks. Recall that the iterative refinement technique from Section 7.3 requires repeatedly checking satisfiability until the constraint becomes satisfiable. However, because the overwhelming majority of these calls return unsatisfiable, we simplify the problem by overapproximating the constraints until the overapproximation becomes satisfiable, in which case we switch to the exact encoding. The key idea underlying our over-approximation is as follows. While the exact encoding introduces joint probability variables over the input facts, we can instead introduce joint probability variables over *intermediate relations* that are k steps from the root node and encode the known correlations between them as part of the constraint. We provide an example of this encoding in Appendix A.4. k is not a fixed value; instead, k is selected dynamically. Details on how k is determined are provided in the Appendix.

Handling very large correlation classes. In cases where correlation classes become prohibitively large, even computing joint probability distributions within a *single* correlation class may be infeasible. For instance, some outliers in our experimental evaluation have hundreds of input facts in the same correlation class, making it infeasible to reason precisely about the joint probability for the entire class. In such cases, our implementation retains soundness but may compromise

δ -exactness. In particular, for correlation classes whose size exceeds a predefined threshold, we do not compute correlation types as described in Phase 1 of Section 7.1; however, we still perform the static analysis from Phase 2, assuming that input correlations within the same class are unknown. Second, when performing iterative refinement of the probability bounds, we approximate the satisfiability check as described in the previous paragraph and do *not* switch to the fully precise encoding. However, we emphasize that even this approximate solution for handling excessively large correlation classes utilizes the exact same machinery described in the rest of the paper.

9 Evaluation

In this section, we now describe the results for the evaluation that is designed to answer the following questions:

- **RQ1:** How accurate are the probability bounds inferred by PRALINE?
- **RQ2:** How efficient/scalable is our method in inferring the probability bounds?
- **RQ3:** How impactful are the key technical ingredients underlying PRALINE?

9.1 Application Domains and Benchmarks

We evaluate PRALINE on two different application domains spanning 30 benchmarks, summarized in Table 2. The first category, labeled **SC** (for Side Channel) in Table 2, corresponds to 19 Datalog-based program analyses for detecting power side-channel leaks. The second category, labeled **Bayes**, consists of 11 Bayesian networks sourced from the bnlearn repository. We provide more information about each of these application domains below.

		Min	Max	Avg	Med
SC	# nodes	56	200,164	43,366	9,093
	# edges	3	173,617	33,284	5,919
	CC size	8	731	133	31
Bayes	# nodes	10	223	79	60
	# edges	11	328	128	77
	CC size	2	13	5	4

Table 2. Benchmark statistics, CC denotes “correlation class”.

Side channel benchmarks. Our first application domain is power side channel detection – a problem that has received significant attention in recent work [58, 59]. For this domain, input facts in the Datalog program are derived from the programs under analysis and include implementations of well-known cryptographic protocols such as AES, SHA3, and MAC-Keccak. The Datalog rules describing the side channel analysis are taken from [59]. Hence, each benchmark in the side channel category corresponds to the “cross product” of an existing side channel detector [59] and a real-world implementation of a cryptographic protocol. However, since the original Datalog-based analyzer only outputs a yes/no answer (indicating a potential power side channel vulnerability), we extend the analysis to quantify leakage severity. Our extension retains the existing Datalog rules and input facts as is, but augments them with probabilities as well as conditional dependencies/probabilities between input facts.

In more detail, we obtain the probabilities of the Datalog rules describing the program analysis using a similar methodology to what has been described in prior work on quantitative Datalog-based race detection [44]. This involves instrumenting the program to count rule firings and using these counts to estimate the probability of each rule applying in practice. To obtain the probabilities of input facts, we first note that while some input facts remain deterministic, others are probabilistic due to variations in register allocation algorithms and hardware architectures. For instance, the probability of register sharing is estimated using empirical data from profiling a code corpus. Similarly, certain input facts are probabilistic as they stem from pre-analysis [59, 62] that infers semantic data dependencies from syntactic ones. The probability of such dependencies is derived from prior empirical studies measuring how often syntactic dependencies translate into actual data dependencies in compiled programs [59]. To quantify conditional probabilities between input facts, we leverage empirical co-occurrence statistics from an existing code corpus [58]. We analyze execution traces to measure how often certain conditions –such as register sharing information

and key-related data dependencies – appear together. These conditional probabilities are essential for accurately assessing the severity of detected side channels, as some Datalog rules depend on both register-sharing behavior and secret-dependent data flows between variables. Consequently, the overall probability of data leakage must account for these correlations between (probabilistic) input facts rather than treating them as independent events.

Bayesian network benchmarks. Our second set of benchmarks, labeled **Bayes** in Table 2, consists of 11 discrete Bayesian networks sourced from the bnlearn repository [50], encompassing a range of network sizes, including *small*, *medium*, and *large* examples. We derive these benchmarks by preserving the original network structure, designating nodes without incoming edges as input facts and all other nodes as output facts. To introduce partially known statistical correlations between input facts, we inject dependencies that reflect realistic co-occurrence patterns observed in empirical data. For example, in medical networks, we introduce correlations between demographic and lifestyle factors, such as smoking and being Asian, while in weather models, we correlate atmospheric conditions like humidity and precipitation, which often vary together. These correlations are not deterministically defined, meaning that while input facts are not mutually independent, their exact joint distribution remains unknown. This approach enables us to evaluate inference under conditions where partial dependency information is available.

Benchmark statistics. Table 2 summarizes key statistics for both benchmark categories. **#nodes** and **#edges** represent the number of nodes and hyperedges in the derivation graph, while **CC size** denotes the size of the largest correlation class. We report the minimum, maximum, average, and median values across all benchmarks in each category. As shown in Table 2, the 19 side-channel benchmarks present a greater computational challenge than the 11 Bayesian network benchmarks from [50], highlighting the complexity of probabilistic inference in security applications. These two categories illustrate distinct but complementary use cases for Praline, both of which involve correlated inputs where exact dependencies are not fully known.

9.2 Experimental Methodology and Set-up

Existing probabilistic extensions of logic programming languages do not account for conditional dependencies between inputs. To assess whether PROBLOG[16] could serve as a baseline, we attempted to encode input correlations⁴ using its *evidence* predicate (see Appendix A.1 for details). However, PROBLOG successfully terminated in only 17 of 30 benchmarks and failed to terminate on the remaining 13. Moreover, even when it did terminate, it produced unsound results due to its inability to faithfully represent PRALINE programs. This limitation stems from fundamental expressiveness gaps—accurately encoding a single PRALINE program in PROBLOG would require generating an infinite number of distinct PROBLOG programs (Appendix A.1).

Given these limitations, we evaluate PRALINE against the constrained optimization approach introduced in Section 6, which we use as the baseline to produce the ground truth. Our evaluation compares this baseline with the proposed δ -exact algorithm, which enhances scalability while preserving precise probability bounds. Throughout the remainder of this section, PRALINE refers to the δ -exact method from Section 7, and “CONSTRAINED OPTIMIZATION” refers to the baseline.

All experiments were conducted on macOS Sonoma 14.4.1 with a 3-hour time limit and a memory cap of 16GB. In our evaluation, we set the δ parameter to 0.01 for the δ -exact algorithm, as it offered a practical balance between runtime and accuracy. In general, δ controls a clear trade-off: smaller values produce tighter probability bounds but incur longer runtimes due to additional refinement iterations, whereas larger values yield faster computations at the cost of looser bounds.

⁴For a fair comparison, we also constructed modified versions of our benchmarks where all input facts are treated as independent; in these cases, PRALINE and PROBLOG produced identical results whenever PROBLOG successfully terminated.

9.3 Accuracy Evaluation

In this section, we evaluate the accuracy of the probability bounds inferred by PRALINE in two ways. First, for those benchmarks where exact inference (CONSTRAINED OPTIMIZATION) terminates, we compare the bounds computed by PRALINE against the ground truth. Second, because PRALINE may compromise δ -exactness for very large correlation classes (recall Section 8), we evaluate the percentage of output facts for which PRALINE guarantees δ -exact inference.

9.3.1 Comparison with ground truth. For 12 of the 30 benchmarks used in our evaluation (covering 3,179 queried facts), exact inference terminates within the 3-hour time limit, allowing us to evaluate PRALINE's results against the ground truth values. Table 3 presents the results of this evaluation, where the **LB Error** column reports the average, minimum, and maximum deviations in the lower bound, and the **UB Error** column provides the same for the upper bound. Specifically, **LB Error** corresponds to $l^* - l$, while **UB Error** denotes $u - u^*$, where (l, u) are the probability bounds computed by the δ -exact procedure and (l^*, u^*) are the exact values obtained via constrained optimization. Despite the substantial efficiency gains of the δ -exact method (evaluated more thoroughly in Section 9.5), its computed probability bounds remain highly precise. Across all 12 benchmarks, the average lower and upper bound errors are just **0.00168** and **0.00177**, respectively—well within the specified δ threshold of 0.01 (i.e., 1%). These results demonstrate that our approximate inference method provides an effective alternative to exact inference while maintaining near-optimal accuracy.

	LB Error	UB Error
Average	0.00168	0.00177
Min	0.00000	0.00000
Max	0.00904	0.00935

Table 3. Accuracy Comparison

9.3.2 Evaluation of δ -exactness. As mentioned in Section 8, the implementation of PRALINE gives up on δ -exactness in some cases to scale to programs with very large correlation classes. In this section, we evaluate for what percent of output facts PRALINE can guarantee δ -exactness of inference. In particular, the result of inference for an output fact is guaranteed to be δ -exact if either (1) the length of the inferred interval is $\leq \delta$, or (2) the (final) satisfiability check in Section 7.3 uses the exact encoding over input facts, theoretically guaranteeing δ -exactness. Note that while these constitute sufficient conditions for δ -exact inference, they are not necessary conditions, meaning that the numbers reported here form a *lower bound* on the percentage of δ -exact results. The result of this evaluation is presented in Table 4, where #facts denotes the number of output facts queried by the program. *Overall, at least 73% of the queried output facts are guaranteed to be δ -exact.* As stated earlier, this number is merely a lower bound on the percentage of δ -exact results, owing to the simple reason that we do not have a scalable method of computing ground truths for the remaining facts.

	#facts	$\delta\%$
SC	96,393	73%
Bayes	663	100%
Overall	97,056	73%

Table 4. δ -exactness rate

Result for RQ1: For the benchmarks where ground truth bounds are available, PRALINE produces precise probability bounds, with average lower/upper bound errors of 0.1%. Over all 30 benchmarks comprising 97,056 queried facts, *at least 73%* are guaranteed to be δ -exact.

9.4 Inference Time Evaluation

To answer our second research question, we evaluate PRALINE's inference efficiency and scalability.

9.4.1 Inference time. Table 5 shows the percentage of benchmarks that can be solved within a given time limit. As shown in this table, PRALINE is able to analyze 60% of the benchmarks in under 10 seconds and 76% in under 100 seconds. All 30

Time	Rate
< 1s	50%
1-10s	10%
10-100s	16%
100-1000s	17%
>1000s	7%

Table 5. Runtime

benchmarks terminate within the specified time limit, with the largest benchmark (with 200164 nodes, 173617 edges, 116 as the average correlation class size) taking 1414.23 seconds. This result shows that PRALINE is able to achieve practical inference times even for complex benchmarks.

9.4.2 Scalability Evaluation. To evaluate the scalability of PRALINE with respect to benchmark complexity, Figure 17 plots runtime against two complexity metrics. Figure 17a shows runtime (Y-axis, in seconds) versus the number of nodes (X-axis), while Figure 17b plots runtime against the average correlation class size. In Figure 17a, the orange curve (a *polynomial fit*) aligns closely with the data, achieving a high correlation coefficient of 0.9869, whereas the exponential fit (green dashed line) does not capture the trend well. Similarly, in Figure 17b, the blue curve (an *approximately polynomial fit*) provides a better fit than the exponential alternative.




Fig. 17. Runtime (Y-axis) vs. benchmark complexity (X-axis).

Result for RQ2: PRALINE is able to perform efficient probabilistic inference, with 76% of benchmarks being solved in under 100 seconds. Empirically, PRALINE scales polynomially with both the number of nodes and the average correlation class size.

9.5 Ablation Study

In this section, we describe a series of ablation studies designed to evaluate the impact of key ingredients of our approach. Specifically, we compare PRALINE against the following ablations:

- **CONSTRAINED OPTIMIZATION** This variant implements Algorithm 1.
- **APPROX ONLY (AO)** This is a variant of PRALINE that computes loose approximate bounds using the approximated bound computation technique from Section 7.2. However, it does not utilize correlation types (correlation types are unknown), and it also does not perform refinement.
- **APPROX+CORRELATION (AC)** This variant does not use iterative refinement for tightening the bound. That is, it computes approximate probabilities while leveraging correlation types.
- **APPROX+REFINEMENT (AR)** This ablation does not compute correlation types to assist the approximated bound computation. However, it does perform the refinement method of Section 7.3.

Among these ablations, we note that **AO** and **AC** do not provide any precision guarantees. Next, we evaluate the impact of each key ingredient on both inference time as well as accuracy.

9.5.1 Evaluation of inference time. Figure 18 explores the impact of various design choices on *inference time*. As we can see from Figure 18, the variants of **PRALINE** that do not have precision guarantees can perform inference more efficiently than all of the others. As expected, **CONSTRAINED OPTIMIZATION** is the slowest and times out on the majority of the benchmarks. In contrast, both **PRALINE** and **AR** solve all benchmarks within the three-hour time limit, exhibiting similar performance in inference time. To see why **PRALINE** and **AR** have similar performance, note that computing correlation types adds some overhead, however, it also reduces time to perform iterative refinement, as the refinement algorithm




Fig. 18. Inference time.

starts with tighter bounds. Thus, overall, the computation of correlation types does not end up having a significant impact on inference time.

9.5.2 Accuracy evaluation. We now consider how the different variants of PRALINE perform in terms of accuracy. To quantify the impact of removing a specific feature in terms of accuracy, we consider a *bound tightness ratio (BTR)* metric, defined as follows:

$$\text{BTR}(\text{PRALINE}^\Delta) = \frac{u' - l'}{u - l}$$

		Med	Avg	Max
BTR	AO	22.50	28.20	82.24
	AC	11.33	17.16	52.08
	AR	2.20	2.38	5.53

Here, PRALINE^Δ refers to a specific ablation for PRALINE and $[l, u]$ and $[l', u']$ are the bounds produced by PRALINE and PRALINE^Δ respectively. Intuitively, the closer the BTR is to 1, the more accurate the ablation. The results of this evaluation are shown in Table 6. For each ablation (e.g., **AO**), we compute the BTR value for all output facts, and report the median, average, and maximum values. As is evident from the data, the intervals computed by PRALINE are much tighter compared to the other ablations, with **AO** producing the least precise intervals, followed by **AC**, and then **AR**.

Theoretically, **AR** and **PRALINE** should achieve similar accuracy, as both employ iterative refinement via binary search to tighten probability bounds. In practice, however, a key difference in how satisfiability checking is performed during refinement leads to **PRALINE** achieving at least twice the tightness of **AR**, as shown in Table 6.

This accuracy gap stems from our implementation (Section 8). When correlation classes are large, the exact SAT encoding becomes computationally expensive and often intractable. To mitigate this, we introduce an optimized satisfiability checking strategy that leverages over-approximation. The precision of this optimization depends on the availability of accurate correlation type information. **PRALINE** explicitly infers correlation types in earlier stages (Section 7.1), enabling it to apply tighter over-approximations that closely match the exact encoding. In contrast, **AR** does not perform correlation type inference and conservatively treats all correlations as UNKNOWN, resulting in looser encodings and reduced accuracy.

Result for RQ3: PRALINE strikes an effective balance between precision and inference time, delivering the most precise results across ablations that terminate on all benchmarks. Variants of PRALINE that do not perform refinement result in bounds that are 17 – 28× worse on average.

10 Related Work

Probabilistic logic programming. Probabilistic programming allows programmers to model distributions and perform probabilistic sampling and inference, with systems like Pyro [9], Turing [24], Hakaru [42], SPPL [47], Dice [29] and PPL [52] leading the way. Recently, there has been significant progress in integrating logical reasoning into probabilistic programming to capture richer logical formalisms such as Horn clauses and first-order logic. Notable examples include probabilistic relational models [26], Markov logic networks [43], Bayesian logic programs [31], and probabilistic logic programming languages such as PRISM [49], LPADs [54], Blog [40], CP-logic [53], PDDL [5, 27], Datalogp [21], Scallop [38], and ProbLog [16, 17]. These formalisms extend existing logic programming languages like Prolog and Datalog by associating each rule with probabilities.

Among these languages, ProbLog [16, 17] and Scallop [38] focus on discrete distributions, which are closely related to our work. These techniques reduce probabilistic inference to weighted model counting (WMC) [55] and employ representations like binary decision diagrams (BDD) [12] to support efficient WMC. However, both ProbLog [16, 17] and Scallop [38] largely assume that input facts are independent and do not allow expressing general forms of conditional dependencies

(other than providing the ability to express mutual exclusion between predicates). In contrast, PRALINE offers syntactic support to declare general forms of conditional dependency between input facts, assign numerical probabilities to such dependencies, and assumes independence only in the absence of such declarations. To the best of our knowledge, JudgeD [56] is the only current Datalog extension that allows expressing dependencies between clauses by associating each input fact with a logical sentence. However, it neither supports negations nor does it allow specifying numerical probabilities for conditional dependencies.

Datalog for program analysis. Logic programming languages like Datalog have found numerous applications in program analysis, including for data-race detection [41, 44, 63], thread-modular analysis [35, 36], side-channel detection [58, 59], and points-to analysis [11, 39, 51, 61, 64]. Traditionally, Datalog-based analyses have been qualitative, but recent work [44, 64] has investigated quantitative analysis methods for inferring the likelihood of data races by incorporating probabilistic reasoning. These approaches, however, are constrained by their assumption about independence of input predicates, which our approach aims to address.

Exact probabilistic inference. Our method also relates to exact inference in graphical models like Bayesian networks and Markov networks. Exact inference techniques include weighted model counting [20, 29, 55], symbolic analysis [25], variable elimination [37], conjugacy [28], generating functions [32, 33], and optimization methods [4, 15]. However, applying these techniques directly to our problem domain is challenging for several reasons. Aside from the obvious structural differences between a Bayesian network and a Datalog derivation graph, our work also distinguishes itself from the Bayesian network setting by allowing for the specification of *incomplete* conditional dependencies for which it is not possible to compute a single probability value. Because of these important differences, prior approaches [29, 37] for speeding up probabilistic inference are unlikely to be effective in our setting.

Approximate probabilistic inference. Approximate inference techniques are primarily based on sampling methods [14, 34] such as Importance Sampling (IS), Markov Chain Monte Carlo (MCMC) and variational inference. However, these methods do not provide guarantees for results produced within a finite time frame. Some approaches [1, 8, 13, 19, 48] infer approximate posterior probabilities with guaranteed bounds; however, these methods typically focus on continuous rather than discrete distributions. Additionally, their guarantees rely on a countable set of sampled interval traces, which scale exponentially with the model's dimension [8].

Verifying probabilistic properties. There is large body of work on verifying probabilistic properties of programs, such as differential privacy and demographic fairness. For example, differential privacy can be expressed as relational properties of probabilistic computations involving expected values. Barthe et al. [7] propose a relational refinement type system and use approximate coupling to construct proofs. Albarghouthi and Hsu [2] and Wang et al. [60] simplify approximate coupling proofs to make it more automated. FairSquare [1], on the other hand, uses symbolic solving to verify if a program meets specified demographic fairness properties. While these approaches deal with probabilistic properties, they are largely orthogonal to our approach.

11 Conclusion

In this paper, we introduced a new probabilistic Datalog framework, PRALINE, which allows users to specify arbitrary statistical correlations between input facts, addressing a significant limitation in existing methods. Importantly, PRALINE is designed to handle scenarios where the statistical correlations between inputs are not fully known, allowing accurate probabilistic inference even under partial information. To solve this problem, we first proposed a constrained optimization

approach that can compute exact probability bounds. Then, to address the scalability limitations of exact inference, we used this constrained optimization method in a more lightweight manner as the basis of δ -exact algorithm that can approximate the true bounds, while guaranteeing that they are within distance δ of the ground truth. Our proposed δ -exact approach iteratively strengthens the approximated bounds through a synergistic combination of static analysis, approximation, and iterative refinement. Our empirical evaluation on 30 real-world probabilistic Datalog programs demonstrates that PRALINE can compute precise probability bounds, while scaling to large benchmarks with more than 200,000 relations. In contrast, the ablations of PRALINE that are not δ -exact infer significantly less accurate probability bounds, while exact inference does not scale. These experiments demonstrate that PRALINE strikes an effective balance between precision and inference time.

12 Data-Availability Statement

An artifact supporting the results of this paper is available on Zenodo [57]. Our tool depends on Datalog, SMT, and optimization solvers, which users will need to install separately. One requirement is a free academic license for Gurobi optimization solver, which can be easily obtained using an institutional email address.

Acknowledgments

We thank the anonymous reviewers for their constructive feedback. This work was conducted in a research group supported by NSF awards CCF-1762299, CCF-1918889, CNS-1908304, CCF-1901376, CNS-2120696, CCF- 2210831, and CCF-2319471, CCF-2422130, CCF-2403211 as well as a DARPA award under agreement HR00112590133.

References

- [1] Aws Albarghouthi, Loris D’Antoni, Samuel Drews, and Aditya V Nori. 2017. Fairsquare: probabilistic verification of program fairness. *Proceedings of the ACM on Programming Languages* 1, OOPSLA (2017), 1–30.
- [2] Aws Albarghouthi and Justin Hsu. 2017. Synthesizing coupling proofs of differential privacy. *Proceedings of the ACM on Programming Languages* 2, POPL (2017), 1–30.
- [3] V Aho Alfred, S Lam Monica, and D Ullman Jeffrey. 2007. *Compilers principles, techniques & tools*. pearson Education.
- [4] Stephen H Bach, Matthias Broecheler, Bert Huang, and Lise Getoor. 2017. Hinge-loss markov random fields and probabilistic soft logic. *Journal of Machine Learning Research* 18, 109 (2017), 1–67.
- [5] Vince Bárány, Balder Ten Cate, Benny Kimelfeld, Dan Olteanu, and Zografoula Vagena. 2017. Declarative probabilistic programming with datalog. *ACM Transactions on Database Systems (TODS)* 42, 4 (2017), 1–35.
- [6] Haniel Barbosa, Clark Barrett, Martin Brain, Gereon Kremer, Hanna Lachnitt, Makai Mann, Abdalrhman Mohamed, Mudathir Mohamed, Aina Niemetz, Andres Nötzli, et al. 2022. cvc5: A versatile and industrial-strength SMT solver. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 415–442.
- [7] Gilles Barthe, Marco Gaboardi, Emilio Jesús Gallego Arias, Justin Hsu, Aaron Roth, and Pierre-Yves Strub. 2015. Higher-order approximate relational refinement types for mechanism design and differential privacy. *ACM SIGPLAN Notices* 50, 1 (2015), 55–68.
- [8] Raven Beutner, C-H Luke Ong, and Fabian Zaiser. 2022. Guaranteed bounds for posterior inference in universal probabilistic programming. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. 536–551.
- [9] Eli Bingham, Jonathan P Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D Goodman. 2019. Pyro: Deep universal probabilistic programming. *Journal of machine learning research* 20, 28 (2019), 1–6.
- [10] Bob Bixby. 2007. The gurobi optimizer. *Transp. Re-search Part B* 41, 2 (2007), 159–178.
- [11] Martin Bravenboer and Yannis Smaragdakis. 2009. Strictly declarative specification of sophisticated points-to analyses. In *Proceedings of the 24th ACM SIGPLAN conference on Object oriented programming systems languages and applications*. 243–262.
- [12] Randal E Bryant. 1986. Graph-based algorithms for boolean function manipulation. *Computers, IEEE Transactions on* 100, 8 (1986), 677–691.

- [13] Patrick Cousot and Michael Monerau. 2012. Probabilistic abstract interpretation. In *European Symposium on Programming*. Springer, 169–193.
- [14] Adnan Darwiche. 2009. *Modeling and reasoning with Bayesian networks*. Cambridge university press.
- [15] Erik Daxberger, Anastasia Makarova, Matteo Turchetta, and Andreas Krause. 2021. Mixed-variable Bayesian optimization. In *Proceedings of the Twenty-Ninth International Conference on Artificial Intelligence*. IJCAI-INT JOINT CONF ARTIF INTELL, 2462–2467.
- [16] Luc De Raedt, Angelika Kimmig, and Hannu Toivonen. 2007. ProbLog: A probabilistic Prolog and its application in link discovery. In *IJCAI 2007, Proceedings of the 20th international joint conference on artificial intelligence*. IJCAI-INT JOINT CONF ARTIF INTELL, 2462–2467.
- [17] Anton Dries, Angelika Kimmig, Wannes Meert, Joris Renkens, Guy Van den Broeck, Jonas Vlasselaer, and Luc De Raedt. 2015. ProbLog2: Probabilistic logic programming. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2015, Porto, Portugal, September 7–11, 2015, Proceedings, Part III* 15. Springer, 312–315.
- [18] Hassan Eldib and Chao Wang. 2014. Synthesis of masking countermeasures against side channel attacks. In *Computer Aided Verification: 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18–22, 2014. Proceedings* 26. Springer, 114–130.
- [19] Shenghua Feng, Mingshuai Chen, Han Su, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Naijun Zhan. 2023. Lower bounds for possibly divergent probabilistic programs. *Proceedings of the ACM on Programming Languages* 7, OOPSLA1 (2023), 696–726.
- [20] Daan Fierens, Guy Van den Broeck, Joris Renkens, Dimitar Shterionov, Bernd Gutmann, Ingo Thon, Gerda Janssens, and Luc De Raedt. 2015. Inference and learning in probabilistic logic programs using weighted boolean formulas. *Theory and Practice of Logic Programming* 15, 3 (2015), 358–401.
- [21] Norbert Fuhr. 1995. Probabilistic datalog—a logic for powerful retrieval methods. In *Proceedings of the 18th annual international ACM SIGIR conference on Research and development in information retrieval*. 282–290.
- [22] Norbert Fuhr. 2000. Probabilistic Datalog: Implementing logical information retrieval for advanced applications. *Journal of the American Society for Information Science* 51, 2 (2000), 95–110.
- [23] Sicun Gao, Jeremy Avigad, and Edmund M Clarke. 2012. δ -complete decision procedures for satisfiability over the reals. In *International Joint Conference on Automated Reasoning*. Springer, 286–300.
- [24] Hong Ge, Kai Xu, and Zoubin Ghahramani. 2018. Turing: a language for flexible probabilistic inference. In *International conference on artificial intelligence and statistics*. PMLR, 1682–1690.
- [25] Timon Gehr, Sasa Misailovic, and Martin Vechev. 2016. PSI: Exact symbolic inference for probabilistic programs. In *Computer Aided Verification: 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17–23, 2016, Proceedings, Part I* 28. Springer, 62–83.
- [26] Lise Getoor, Nir Friedman, Daphne Koller, Avi Pfeffer, and Ben Taskar. 2007. Probabilistic relational models. (2007).
- [27] Martin Grohe, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Peter Lindner. 2022. Generative datalog with continuous distributions. *J. ACM* 69, 6 (2022), 1–52.
- [28] Matthew D Hoffman, Matthew J Johnson, and Dustin Tran. 2018. Autoconj: recognizing and exploiting conjugacy without a domain-specific language. *Advances in Neural Information Processing Systems* 31 (2018).
- [29] Steven Holtzen, Guy Van den Broeck, and Todd Millstein. 2020. Scaling exact inference for discrete probabilistic programs. *Proceedings of the ACM on Programming Languages* 4, OOPSLA (2020), 1–31.
- [30] Herbert Jordan, Bernhard Scholz, and Pavle Subotić. 2016. Soufflé: On synthesis of program analyzers. In *Computer Aided Verification: 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17–23, 2016, Proceedings, Part II* 28. Springer, 422–430.
- [31] Kristian Kersting and Luc De Raedt. 2007. Bayesian logic programming: Theory and tool. (2007).
- [32] Lutz Klinkenberg, Christian Blumenthal, Mingshuai Chen, Darion Haase, and Joost-Pieter Katoen. 2024. Exact Bayesian Inference for Loopy Probabilistic Programs using Generating Functions. *Proceedings of the ACM on Programming Languages* 8, OOPSLA1 (2024), 923–953.
- [33] Lutz Klinkenberg, Tobias Winkler, Mingshuai Chen, and Joost-Pieter Katoen. 2023. Exact probabilistic inference using generating functions. *arXiv preprint arXiv:2302.00513* (2023).
- [34] Daphne Koller and Nir Friedman. 2009. *Probabilistic graphical models: principles and techniques*. MIT press.
- [35] Markus Kusano and Chao Wang. 2016. Flow-sensitive composition of thread-modular abstract interpretation. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. 799–809.
- [36] Markus Kusano and Chao Wang. 2017. Thread-modular static analysis for relaxed memory models. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*. 337–348.
- [37] Jianlin Li, Eric Wang, and Yizhou Zhang. 2024. Compiling Probabilistic Programs for Variable Elimination with Information Flow. *Proceedings of the ACM on Programming Languages* 8, PLDI (2024), 1755–1780.
- [38] Ziyang Li, Jiani Huang, and Mayur Naik. 2023. Scallop: A language for neurosymbolic programming. *Proceedings of the ACM on Programming Languages* 7, PLDI (2023), 1463–1487.

- [39] Magnus Madsen, Ming-Ho Yee, and Ondřej Lhoták. 2016. From datalog to flix: A declarative language for fixed points on lattices. *ACM SIGPLAN Notices* 51, 6 (2016), 194–208.
- [40] Brian Milch, Bhaskara Marthi, Stuart Russell, David Sontag, Daniel L Ong, and Andrey Kolobov. 2007. BLOG: Probabilistic models with unknown objects. (2007).
- [41] Mayur Naik, Alex Aiken, and John Whaley. 2006. Effective static race detection for Java. In *Proceedings of the 27th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 308–319.
- [42] Praveen Narayanan, Jacques Carette, Wren Romano, Chung-chieh Shan, and Robert Zinkov. 2016. Probabilistic inference by program transformation in Hakaru (system description). In *Functional and Logic Programming: 13th International Symposium, FLOPS 2016, Kochi, Japan, March 4–6, 2016, Proceedings 13*. Springer, 62–79.
- [43] Feng Niu, Christopher Ré, AnHai Doan, and Jude Shavlik. 2011. Tuffy: Scaling up Statistical Inference in Markov Logic Networks using an RDBMS. *Proceedings of the VLDB Endowment* 4, 6 (2011).
- [44] Mukund Raghothaman, Sulekha Kulkarni, Kihong Heo, and Mayur Naik. 2018. User-guided program reasoning using Bayesian inference. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 722–735.
- [45] Thomas W Reps. 1995. Demand interprocedural program analysis using logic databases. In *Applications of Logic Databases*. Springer, 163–196.
- [46] Ludger Rüschendorf. [n. d.]. Fréchet-bounds and their applications. In *Advances in Probability Distributions with Given Marginals: beyond the copulas*. Springer, 151–187.
- [47] Feras A Saad, Martin C Rinard, and Vikash K Mansinghka. 2021. SPPL: probabilistic programming with fast exact symbolic inference. In *Proceedings of the 42nd acm sigplan international conference on programming language design and implementation*. 804–819.
- [48] Sriram Sankaranarayanan, Aleksandar Chakarov, and Sumit Gulwani. 2013. Static analysis for probabilistic programs: inferring whole program properties from finitely many paths. In *Proceedings of the 34th ACM SIGPLAN conference on Programming language design and implementation*. 447–458.
- [49] Taisuke Sato. 1995. A statistical learning method for logic programs with distribution semantics. (1995).
- [50] Marco Scutari. 2009. Learning Bayesian networks with the bnlearn R package. *arXiv preprint arXiv:0908.3817* (2009).
- [51] Yannis Smaragdakis, George Kastrinis, and George Balatsouras. 2014. Introspective analysis: context-sensitivity, across the board. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 485–495.
- [52] Jan-Willem van de Meent, Brooks Paige, Hongseok Yang, and Frank Wood. 2018. An introduction to probabilistic programming. *arXiv preprint arXiv:1809.10756* (2018).
- [53] Joost Vennekens, Marc Denecker, and Maurice Bruynooghe. 2009. CP-logic: A language of causal probabilistic events and its relation to logic programming. *Theory and practice of logic programming* 9, 3 (2009), 245–308.
- [54] Joost Vennekens, Sofie Verbaeten, and Maurice Bruynooghe. 2004. Logic programs with annotated disjunctions. In *Logic Programming: 20th International Conference, ICLP 2004, Saint-Malo, France, September 6–10, 2004. Proceedings 20*. Springer, 431–445.
- [55] Jonas Vlasselaer, Guy Van den Broeck, Angelika Kimmig, Wannes Meert, and Luc De Raedt. 2016. Tp-compilation for inference in probabilistic logic programs. *International Journal of Approximate Reasoning* 78 (2016), 15–32.
- [56] Brend Wanders, Maurice van Keulen, and Jan Flokstra. 2016. Judged: a probabilistic datalog with dependencies. In *Workshops at the Thirtieth AAAI Conference on Artificial Intelligence*.
- [57] Jingbo Wang, Shashin Halalingaiah, Weiyi Chen, Chao Wang, and Isil Dillig. 2025. Reproduction Package for Article ‘Probabilistic Inference for Datalog with Correlated Inputs’. ACM. <https://doi.org/10.5281/zenodo.15760564>
- [58] Jingbo Wang, Chunggha Sung, Mukund Raghothaman, and Chao Wang. 2021. Data-driven synthesis of provably sound side channel analyses. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*.
- [59] Jingbo Wang, Chunggha Sung, and Chao Wang. 2019. Mitigating power side channels during compilation. In *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 590–601.
- [60] Yuxin Wang, Zeyu Ding, Guan hong Wang, Daniel Kifer, and Danfeng Zhang. 2019. Proving differential privacy with shadow execution. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 655–669.
- [61] John Whaley, Dzintars Avots, Michael Carbin, and Monica S Lam. 2005. Using Datalog with binary decision diagrams for program analysis. In *Asian Symposium on Programming Languages and Systems*. Springer, 97–118.
- [62] Jun Zhang, Pengfei Gao, Fu Song, and Chao Wang. 2018. SCInfer: Refinement-based verification of software counter-measures against side-channel attacks. In *International Conference on Computer Aided Verification*. Springer, 157–177.
- [63] Xin Zhang, Radu Grigore, Xujie Si, and Mayur Naik. 2017. Effective interactive resolution of static analysis alarms. *Proceedings of the ACM on Programming Languages* 1, OOPSLA (2017), 1–30.

- [64] Xin Zhang, Ravi Mangal, Radu Grigore, Mayur Naik, and Hongseok Yang. 2014. On abstraction refinement for program analyses in Datalog. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 239–248.

A Appendix

A.1 Encoding Input Correlations using ProbLog

To the best of our knowledge, ProbLog primarily supports mutual independence as a form of input correlation. To encode more complex types of input correlations, we must creatively use other available syntax structures in ProbLog to implicitly represent them. We have explored two methods to encode input correlations. However, both approaches rely on computations using a complete conditional probability table and can not capture the exact semantics as PRALINE does.

In our evaluation, we use the first method, specifically the *evidence* feature, to encode input correlations in ProbLog.

```

1 0.6::edge(1,4).
2 0.6::edge(2,5).
3 0.6::edge(2,6).
4 0.8::edge(2,5) | edge(1,4).
5 0.83::edge(2,6) | edge(1,4).
6 r :- edge(1,4), edge(2,5), \+edge(2,6).
7 query(r).
```

Fig. 19. An example PRALINE program P^* .

```

1 1/2::edge(2,5).
2 1/2::edge(1,4).
3 1/2::edge(2,6).
4
5 1/2::c :- edge(2,5), edge(1,4), edge(2,6).
6 1/8::c :- edge(2,5), edge(1,4), \+edge(2,6).
7 1/8::c :- edge(2,5), \+edge(1,4), edge(2,6).
8 1/32::c :- edge(2,5), \+edge(1,4), \+edge(2,6).
9 19/128::c :- \+edge(2,5), edge(1,4), edge(2,6).
10 1/128::c :- \+edge(2,5), edge(1,4), \+edge(2,6).
11 1/128::c :- \+edge(2,5), \+edge(1,4), edge(2,6).
12 137/384::c :- \+edge(2,5), \+edge(1,4), \+edge(2,6).
13
14 e_25_14 :- edge(2,5), edge(1,4).
15 e_14_26 :- edge(1,4), edge(2,6).
16 r :- edge(1,4), edge(2,5), \+edge(2,6).
17
18 evidence(c, true).
19
20 query(edge(2,5)). % ProbLog result: 0.6 -> encoding 0.6::edge(2,5).
21 query(edge(1,4)). % ProbLog result: 0.6 -> encoding 0.6::edge(1,4).
22 query(edge(2,6)). % ProbLog result: 0.6 -> encoding 0.6::edge(2,6).
23 query(e_25_14). % ProbLog result: 0.6*0.8 -> encoding 0.8::edge(2,5)|edge(1,4).
24 query(e_14_26). % ProbLog result: 0.6*0.83 -> encoding 0.83::edge(2,6)|edge(1,4).
25 query(r). % ProbLog result: Pr(r) = 0.096
```

Fig. 20. The first ProbLog program encoding the same input fact and conditional probabilities as Figure 19.

```

1 1/2::edge(2,5).
2 1/2::edge(1,4).
3 3/4::edge(2,6).
4
5 5/8::c :- edge(2,5), edge(1,4), edge(2,6).
6 1/2::c :- edge(2,5), edge(1,4), \+edge(2,6).
7 1/32::c :- edge(2,5), \+edge(1,4), edge(2,6).
8 1/2::c :- edge(2,5), \+edge(1,4), \+edge(2,6).
9 377/1920::c :- \+edge(2,5), edge(1,4), edge(2,6).
10 3/640::c :- \+edge(2,5), edge(1,4), \+edge(2,6).
11 263/1920::c :- \+edge(2,5), \+edge(1,4), edge(2,6).
12 1871/1920::c :- \+edge(2,5), \+edge(1,4), \+edge(2,6).
13
14 e_25_14 :- edge(2,5), edge(1,4).
15 e_14_26 :- edge(1,4), edge(2,6).
16 r :- edge(1,4), edge(2,5), \+edge(2,6).
17
18 evidence(c, true).
19
20 query(edge(2,5)). % ProbLog result: 0.6 -> encoding 0.6::edge(2,5).
21 query(edge(1,4)). % ProbLog result: 0.6 -> encoding 0.6::edge(1,4).
22 query(edge(2,6)). % ProbLog result: 0.6 -> encoding 0.6::edge(2,6).
23 query(e_25_14). % ProbLog result: 0.6*0.8 -> encoding 0.8::edge(2,5)|edge(1,4).
24 query(e_14_26). % ProbLog result: 0.6*0.83 -> encoding 0.83::edge(2,6)|edge(1,4).
25 query(r). % ProbLog result: Pr(r) = 0.10105263

```

Fig. 21. The second PROBLOG program encoding the same input fact and conditional probabilities as Figure 19.

1. Using evidence features. In PROBLOG, we can utilize evidence features to encode the conditional probabilities among input facts. The idea is to introduce mutually-independent facts, and constraint them using the derived rules. For instance, given three independent facts x, y, z (i.e., $P(x \wedge y) = P(x)P(y)$), two rules $p_1 :: c :- x, y, z$, $p_2 :: c :- \neg x, y, z$, and the statement *evidence(c)* (i.e., observing that c is true), x, y, z become dependent after observing c , i.e., $P(x \wedge y|c) \neq P(x|c) * P(y|c)$. By reverse engineering the value of p_1 and p_2 , we can encode $P(x|y, c)$ to any value as follows:

$$\begin{aligned}
 P(x|y, c) &= P(x, y|c) / P(y|c) \\
 &= P(x, y, c) / P(y, c) \\
 &= \frac{P(c|x, y)P(x, y)}{P(c|y)P(y)} \\
 &= \frac{[P(c|x, y, z) * P(z) + P(c|x, y, \neg z) * (1 - P(z))] P(x, y)}{[P(c|y, x, z)P(x, z) + P(c|y, \neg x, z)P(\neg x, z) + P(c|y, x, \neg z)P(x, \neg z) + P(c|y, \neg x, \neg z)P(\neg x, \neg z)] P(y)} \\
 &= \frac{[p_1 * P(z) + 0] * P(x) * P(y)}{[p_1 * P(x)P(z) + p_2 * P(\neg x)P(z) + 0 + 0] P(y)}
 \end{aligned}$$

$P(x|y, c)$ specifies the conditional probability between input facts x and y , given c is true.

Example 4. We first provide an example PRALINE program simplified from the motivation example of Figure 4 in Section 3. Given this PRALINE program in Figure 19, we provide two corresponding PROBLOG programs that encode the same conditional probabilities as shown in Figures 20 and 21. For instance, in Figure 20, we denote $\text{edge}(2, 5)$ and $\text{edge}(1, 4)$ as x and y respectively. They are independent initially. We can correlate x and y by using the derived rules (Lines 5-12) and the observation that c is true. Given c is true, x and y are no longer independent. In this PROBLOG encoding, we get $P(x, y|c) = 0.48 = 0.6 * 0.8 \neq P(x|c) * P(y|c) = 0.6 * 0.6$ (line 23), which encodes the input correlation $0.8::\text{edge}(2, 5) :- \text{edge}(1, 4)$.

Given the PRALINE program P^* as shown in Figure 19, we can generate infinitely many PROBLOG programs that encode the same input fact probabilities and conditional probabilities from P^* (Lines 1-6). However, these PROBLOG programs yield different results of $\text{query}(r)$, e.g., **0.096** (Figure 20) and **0.101** (Figure 21).

In short, a single PROBLOG program might not always capture the same semantics as a PRALINE program. PRALINE is equipped to handle scenarios where some conditional probabilities are unknown, allowing it to provide outputs as probability intervals as the *exact* inference result. Conversely, PROBLOG is designed to compute absolute probabilities as its *exact* result, typically requiring complete information to do so.

Figures 20 and 21 are two PROBLOG examples that encode the same input fact and conditional probabilities as the PRALINE program P^* (in Figure 19), by correlating the independent input facts using derived rules, e.g., $\text{edge}(1, 4)$, $\text{edge}(2, 5)$ and $\text{edge}(2, 6)$ are independent input facts but they are dependent given event c is true (e.g., $\text{evidence}(c, \text{true})$). We also want to highlight that to encode the same conditional probabilities using PROBLOG, users would need to reverse engineer all the rule probabilities, e.g., $19/128::c :- \neg \text{edge}(2, 5), \text{edge}(1, 4), \text{edge}(2, 6)$ as shown in Line 9 of Figure 20.

2. Using $p :: I :- S_I$. Another approach involves using the syntax $p :: I :- S_I$ to represent input correlations. In ProbLog, however, this syntax is not intended for articulating conditional probabilities among input facts (e.g., it is not designed to accommodate an input fact I as the rule head). Utilizing this syntax for input correlations might consequently yield unforeseen results.

To demonstrate how such unexpected outcomes occur, we present a new example illustrated in Figure 22. The figure highlights specified input correlations in green, and ProbLog computes the query result for the output fact $\text{path}(1, 2)$ as 0.824, shown in red.

We analyze the derivation graph of $\text{path}(1, 2)$ to trace the source of this discrepancy. As detailed in Figure 22, ProbLog initially selects $\text{edge}(1, 2)$ with a probability of 0.6. If not, it assumes $\text{edge}(1, 2)$ is false with a probability of 0.4 and attempts to derive it using the input correlation rule (colored in green):

$$\begin{aligned} P(\text{edge}(1, 2)) &= P(\text{edge}(1, 2) \wedge \text{edge}(2, 5)) + P(\text{edge}(1, 2) \wedge \neg \text{edge}(2, 5)) \\ &= P(\text{edge}(1, 2) \wedge \text{edge}(2, 5)) + 0 \\ &= 0.8 \times P(\text{edge}(2, 5)) = 0.8 \times 0.7 \end{aligned}$$

This approach leads to an erroneous outcome, as ProbLog is not designed to handle input correlations using such rules and mistakenly processes $\text{edge}(1, 2)$ as both an input and output fact. This results in two conflicting calculations for $P(\text{edge}(1, 2))=0.6$ or $0.4 \times 0.8 \times 0.7$, which are incorrectly aggregated. As a result, $P(\text{path}(1, 2)) = P(\text{edge}(1, 2)) = 0.6 + 0.4 \times 0.8 \times 0.7 = 0.824$.

The unexpected result occurs because ProbLog is not typically designed to use rules such as $p :: R :- S_R$ to directly express input correlations. These rules are primarily intended for specifying the probability of rules where the output fact serves as the rule head.

A.2 Discussion of Syntax Constructs in PRALINE and PROBLOG

In PRALINE, there are two types of rules: $p :: R :- S_R$ and $p :: I | S_I$. In PROBLOG, only the first type is available.

The first type defines how the output fact R is derived. If R can be derived by only one rule, then $P(R) = P(R \wedge S_R) + P(R \wedge \neg S_R) = P(R \wedge S_R) = p \times P(S_R)$, since $P(R \wedge \neg S_R)$ is zero by default.




Fig. 22. Computation procedure of path(1,2) by PROBLOG.

If multiple rules infer R , it is interpreted as follows:

$$\begin{aligned}
 p_1 &:: R \quad :- S_1 \\
 p_2 &:: R \quad :- S_2 \\
 &\dots \\
 p_n &:: R \quad :- S_n
 \end{aligned}$$

$P(R) = P(\underline{R \wedge S_1} \vee \underline{R \wedge S_2} \vee \dots \underline{R \wedge S_n}) + P(R \wedge \neg(S_1 \vee S_2 \dots S_n)) = P(\underline{R \wedge S_1} \vee \underline{R \wedge S_2} \vee \dots \underline{R \wedge S_n})$ where $P(R \wedge \neg(S_1 \vee S_2 \dots S_n))$ is zero. It is because if no rules are applicable to infer R , i.e., $\neg(S_1 \vee S_2 \vee \dots S_n)$, then R is false by default according to probabilistic Datalog semantics.

The second type, $p :: I \mid S_I$, encodes statistical correlations among input facts I , which are considered constraints, rather than the definition of I . It denotes $P(I \mid S_I) = p$. Unlike the first type, $P(I \wedge \neg S_I)$ is unknown and may not be zero. If $P(I \wedge \neg S_I)$ is zero, we may get conflicting probabilities for $P(I)$, such as $P(I)_1 = p_1$ from the rule $p_1 :: I$ and $P(I)_2 = P(I \wedge S_I) + P(I \wedge \neg S_I) = P(I \wedge S_I) = p \times P(S_I)$. To avoid conflicts, $P(I \wedge \neg S_I)$ is not enforced as zero.

When user-provided correlation constraints are conflicting, such as $0.5 :: I_1$, $0.3 :: I_2$, $0.6 :: I_1 \mid I_2$, and $0.7 :: I_1 \mid \neg I_2$, which imply $P(I_1) = 0.5$ is inconsistent with $P(I_1) = P(I_1 \mid I_2)P(I_2) + P(I_1 \mid \neg I_2)(1 - P(I_2)) = 0.6 \times 0.3 + 0.7 \times 0.7$, PRALINE will return "No solution" and produce no output.

The syntax structure $p :: I \mid S_I$ enables users to represent conditional probabilities even when only partially known. For example, consider a scenario where a user is aware of $P(I \mid S_I) = p$ but lacks knowledge about $P(I \mid \neg S_I)$. Under such conditions, determining an exact probability may not be feasible. Our tool, PRALINE, is designed to handle such cases by inferring probabilities and providing them as intervals. This is not a reflection of imprecision in our tool, but rather a result of dealing unknown conditional dependencies. In contrast, traditional probabilistic Datalog tools like PROBLOG calculate absolute probabilities and assume unspecified probabilities as zero, such as $P(I \mid \neg S_I) = 0$, instead of providing outputs as probability intervals.

PRALINE allows expressing conditional dependencies that are not precisely known, which is a common scenario in many real-world applications including program analysis tasks. In such cases, the probability of output facts cannot be determined precisely, but one can compute lower and upper bounds on the probability. In contrast, PROBLOG cannot express such scenarios, as it always produces a single probability value. To capture the same semantics as PRALINE, one would need to write infinitely many PROBLOG programs. Any PROBLOG program inherently corresponds to a

PRALINE program where all joint probabilities are known. Sec A.1 shows a PRALINE program (a simplified version of Figure 4) and why it would take infinitely many PROBLOG programs to encode the same semantics.

A.3 Semantics of PRALINE

In this section, we formally define the semantics of PRALINE. Since PRALINE inherits its syntax from Probabilistic Datalog, we extend the possible-world semantics [22] of probabilistic Datalog to formalize PRALINE.

Probabilistic Datalog

- *Rule*: A sentence of the form:

$$p :: O :- O_1, O_2, \dots, O_n.$$

Here, O is an atom, while O_1, O_2, \dots, O_n are literals. Each atom represents a predicate or relation, and each literal is either an atom or its negation. If the body of a rule is empty, the rule is considered an input fact declaration.

- *Query*: A query is a conjunction of literals.

PRALINE inherits the syntax of probabilistic Datalog but relaxes the mutual independence assumption among rule probabilities. In PRALINE, rules that declare input facts are allowed to be dependent, whereas rules with non-empty bodies are still treated as mutually independent. Additionally, PRALINE introduces new syntax features that allow users to flexibly specify *partially known* or *completely known* conditional dependencies.

- *Correlation classes*: $\text{corr}(I_1, I_2, \dots)$ or $k :: \text{Class}(I)$
- *Conditional dependencies*: $p :: I \mid S$.

The newly added syntax features from PRALINE do not modify the probability expression of any derived fact but only affect its evaluation. We first describe how the probability expression is obtained from a PRALINE program according to the semantics of Probabilistic Datalog and then discuss how it is evaluated based on PRALINE's semantics.

The probability expression is computed based on the probabilities of the ground facts and instantiated rules used in deriving the fact.

Definition 13. A Datalog program D is *modularly stratified* if there exists an assignment of ordinal levels to ground atoms such that: (a) if a ground atom appears negatively in the body of a rule, then the ground atom in the head of that rule is assigned a strictly higher level; (b) if a ground atom appears positively in the body of a rule, then the ground atom in the head is assigned a level that is at least as high.

For a Probabilistic Datalog program D , let S_g denote the set of grounded rules, consisting of S_d and S_p , i.e., $S_g = S_d \cup S_p$. We define S_d as the set of deterministic rules ($p = 1$) and S_p as the set of probabilistic (indeterministic, $0 < p < 1$) rules. The set of all possible deterministic programs of D is given by:

$$S(D) = \{S_d \cup x \mid x \in \text{PowerSet}(S_p)\}$$

LEMMA 1. A probabilistic Datalog program D is modularly stratified if every element in the set of its possible deterministic programs $D' \in S(D)$ is modularly stratified.

LEMMA 2. A modularly stratified Datalog program has a least fixed point.

The least fixed point of a Datalog program corresponds to its least Herbrand model, which captures all facts that can be derived from the program.

Definition 14. The Herbrand universe of a Datalog program D is the set of all ground terms that can be formed using the constants in D . For a deterministic program D' (i.e., $D' \in S(D)$) with a complete, well-founded Herbrand model, let $\omega(D')$ denote this model. Additionally, $\omega(D')$ represents a possible world of D .

For each deterministic program D' , $\omega(D')$ is defined as a set of all ground rules (r) used in D' , i.e., $\omega(D') = \{r\}$. An input fact can also be represented as a rule with a body that is always true.

Definition 15. Given a Probabilistic Datalog program D , the set of possible worlds is defined as

$$W(D) = \{\omega(D') \mid D' \in S(D)\}.$$

We also use the judgment $\omega \models D$ to denote that ω is a possible world consistent with the PRALINE program D .

Definition 16. Let μ be a mapping that assigns each possible world ω ($\omega \in W(D)$) to its full joint probability distribution over the ground terms/rules in ω .

Given that $\omega = \{r\}$, if all input facts and rules are assumed to be mutually independent, μ maps each joint distribution of r to a unique value (*one-to-one mapping*). Otherwise, without the independence assumption, μ maps it to a set of values (*one-to-many mapping*).

Definition 17. For a possible world ω , let $\epsilon : \omega \rightarrow \text{EK}$ be a mapping that assigns each ground rule ($r \in \omega$) to an event key (EK). This mapping satisfies the following constraints:

- $\forall r \in \omega. \quad \mu(r) = 1 \leftrightarrow \epsilon(r) = \text{true}.$
- $\forall r, r' \in \omega. \quad \epsilon(r) = \epsilon(r') \rightarrow (r = r' \vee (\mu(r) = 1 \wedge \mu(r') = 1)).$

If the probability of a rule r is constantly true (i.e., $\mu(r) = 1$), it is not mapped to an EK variable. Instead, it is assigned a constant value. If two rules share the same EK variable, they are either identical (i.e., $r = r'$) or both have a probability of 1.

Given an event key variable v , the inverse mapping $\epsilon^{-1}(v)$ denotes the corresponding ground rule r .

Definition 18. For each output fact O , we define a mapping θ from O to a Boolean expression (\mathbb{B}) over event keys (EK). This mapping is defined as follows:

- For an output fact O ,

$$\theta(O) = \bigvee_{r_o} \theta(r_o),$$

where r_o is a rule whose head matches O and $r_o \in \omega$, $\omega \in W(D)$.

- For a rule r of the form $p :: O :- O_1, O_2, \dots, O_n$,

$$\theta(r) = \epsilon(r) \wedge \theta(O_1) \wedge \theta(O_2) \wedge \dots \wedge \theta(O_n).$$

- For a negated fact $\neg O$,

$$\theta(\neg O) = \neg \theta(O).$$

The subtlety lies in the case of an input fact I . If rule r is used to declare the input fact, i.e., $p :: I$ is a special case of $p :: I :- O_1, O_2, \dots, O_n$ where $O_1 \wedge O_2 \wedge \dots \wedge O_n = \text{true}$, then $\theta(I) = \epsilon(r)$.

Definition 19. We define a mapping \mathcal{W} that assigns each Boolean expression B to the set of possible worlds in which B holds true.

The mapping \mathcal{W} satisfies the following properties:

- $\mathcal{W}(\text{true}) = W(D)$
- $\mathcal{W}(\neg B) = W(D) \setminus \mathcal{W}(B)$
- $\mathcal{W}(B_1 \wedge B_2) = \mathcal{W}(B_1) \cap \mathcal{W}(B_2)$

- $\mathcal{W}(B_1 \vee B_2) = \mathcal{W}(B_1) \cup \mathcal{W}(B_2)$
- $\mathcal{W}(B) = \{\omega \mid r = \epsilon^{-1}(B), r \in \omega, \omega \models D\}$

For the last case, $\mathcal{W}(B)$, the Boolean expression B represents an atomic formula corresponding to an input fact, and r is the rule that declares this input fact.

Specifically, for any output fact O in the PRALINE program D , let $\theta(O)$ denote the Boolean expression over event keys representing O . The set of possible worlds in which O holds is given by $\mathcal{W}(\theta(O))$.

Given a possible world ω , if $\omega \in \mathcal{W}(\theta(O))$, we establish the judgment $\omega \models O$, meaning that O can be inferred from ω .

Definition 20. Given a full joint probability distribution μ , the probability expression of an output fact O is computed as follows:

$$P_\mu(O) = \sum_{\omega \in \Omega} P_\mu(\omega),$$

where $\Omega = \{\omega \mid \omega \models O, (\omega, \mu) \models D\}$.

The judgment $(\omega, \mu) \models D$ indicates that the world ω and the full joint probability distribution μ are consistent with the input PRALINE program D .

For instance, consider a PRALINE program where an input fact is specified as follows: $0.5 :: I1$, denoted as r_1 . If $\mu(r_1) \neq 0.5$, then the judgment $(\omega, \mu) \models D$ is no longer valid.

In a classical probabilistic Datalog program, all rules are assumed to be mutually independent. Consequently, we have:

$$P_\mu(\omega) = P_\mu\left(\bigwedge_{r \in \omega} r\right) = \prod_{r \in \omega} P_\mu(r) = \prod_{r \in \omega} \mu(r).$$

However, in a PRALINE program, this independence assumption is relaxed. Specifically, rules that define input facts (i.e., rules with an empty body) are not required to be independent. As a result, in PRALINE, we may have:

$$P_\mu\left(\bigwedge_{r \in \omega} r\right) \neq \prod_{r \in \omega} P_\mu(r).$$

To determine $P_\mu(\omega)$, users must provide a complete conditional probability distribution μ , which uniquely determines $P_\mu(\bigwedge_{r \in \omega} r)$. However, specifying the full distribution may be impractical in many cases. To address this, PRALINE allows users to provide *partially specified* conditional probability information within the PRALINE program D . Consequently, multiple probability distributions μ may satisfy D , whereas in classical probabilistic Datalog, there is a unique μ that satisfies the input program D .

For example, consider the following input facts and their dependencies in a PRALINE program:

$$0.6 :: I1, \quad 0.7 :: I2, \quad 0.6 :: I3, \quad 0.8 :: I3 \mid I1, I2.$$

In this case, multiple probability distributions μ satisfy these constraints. Consequently, the probability of an output fact $P(O)$ is no longer a singular value but a set of values, defined as follows:

$$P(O) = \{P_\mu(O) \mid \exists \omega. (\omega, \mu) \models D\}.$$

To obtain a valid set of probability distributions μ consistent with the PRALINE program D , this can be formulated as a constraint-solving problem. Any complete model that satisfies the constraints specified by D represents a valid μ . The semantic encoding of these constraints are as

follows:

$$\llbracket p :: I. \rrbracket = P(I) = p \quad (1)$$

$$\llbracket p :: I_0 \mid I_1, I_2, \dots, I_n. \rrbracket = P(I_0 \wedge I_1 \cdots \wedge I_n) = p \times P(I_1 \wedge \cdots \wedge I_n) \quad (2)$$

$$\llbracket k_1 :: \text{Class}(I_1). k_2 :: \text{Class}(I_2). \rrbracket = I_1 \perp\!\!\!\perp I_2 \text{ if } k_1 \neq k_2; \quad (3)$$

$$\text{Otherwise, } I_1 \not\perp\!\!\!\perp I_2 \quad (4)$$

$$\llbracket \text{corr}(I_1, I_2) \rrbracket = I_1 \not\perp\!\!\!\perp I_2 \quad (5)$$

$$\llbracket p :: O :- O_1, O_2, \dots, O_n. \rrbracket = P(r) = p \text{ if } r \text{ denotes the rule } p :: O :- O_1, O_2, \dots, O_n. \quad (6)$$

$$= \forall r, I. \quad r \perp\!\!\!\perp I. \quad (7)$$

$$= \forall r_i, r_j. (i \neq j) \Rightarrow r_i \perp\!\!\!\perp r_j \quad (8)$$

$\perp\!\!\!\perp$ denotes mutual independence. The notation r refers only to rules with non-empty bodies, excluding input declarations. If users do not explicitly declare the correlation class of an input fact I , we assign it a default class ID of -1 , which satisfies the following property: $\forall I_i, I_j. \text{Class}(I_i) = -1 \Rightarrow I_i \perp\!\!\!\perp I_j$. All the above probability computation $P(x)$ can be rewritten as $\mu(x)$. The semantics above naturally extends to multiple correlation class declarations.

$$\llbracket k_1 :: \text{Class}(I_1). \dots k_n :: \text{Class}(I_n). \rrbracket = \begin{cases} I_1 \perp\!\!\!\perp I_2 \perp\!\!\!\perp \dots \perp\!\!\!\perp I_n & \text{if } \forall i \neq j, k_i \neq k_j \\ \bigwedge_{1 \leq i < j \leq n} \begin{cases} I_i \perp\!\!\!\perp I_j & \text{if } k_i \neq k_j \\ I_i \not\perp\!\!\!\perp I_j & \text{if } k_i = k_j \neq -1 \end{cases} & \text{otherwise} \end{cases} \quad (9)$$

$$\llbracket \text{corr}(I_1, I_2, \dots, I_n) \rrbracket = \bigwedge_{1 \leq i < j \leq n} I_i \not\perp\!\!\!\perp I_j \quad (10)$$

A.4 Optimized Satisfiability Check

The iterative refinement from Section 7.3 requires frequent satisfiability checks, most of which return unsatisfiable. To speed up this process, constraints are overapproximated by introducing joint probability variables over intermediate relations that are k steps from the root node.

In the BOUNDBOUNDS procedure, we iteratively check unsatisfiability until it reaches the SAT region. To optimize performance, we propose an overapproximated version ϕ' of the original set of constraints ϕ , such that $\phi \Rightarrow \phi'$. Thus, $\neg(\phi' \wedge l \leq \text{Expr}(R) \leq u) \Rightarrow \neg(\phi \wedge l \leq \text{Expr}(R) \leq u)$. Hence, if we get the UNSAT result from the approximated version ϕ' , it implies the unsatisfiability of the original version ϕ .

To see why the approximated version improves performance and why $\phi \Rightarrow \phi'$, we provide an example to illustrate this.

There is a derivation graph in Figure 23, focusing on computing the probability of the root node (relation A). I_1 till I_8 are input facts. There are two correlation classes $\{I_1, I_2, I_3, I_4\}$ and $\{I_5, I_6, I_7, I_8\}$. The original constraint ϕ is shown on the left of Figure 24. Here, ϕ_{corr} denotes the constraints from the correlation classes, and we omit the details of ϕ_{corr} here.

The overapproximated version is shown on the right of Figure 24. The encoding starts from the *intermediate relations* (B, C, D, E) instead of the input facts (See Figure 23). From the correlation analysis, we know that B, C are dependent and D, E are dependent. We also treat all unknown cases as if they are inside the same correlation class for soundness reasons.

In ϕ' , there are also two correlation classes $V_1 = \{B, C\}$ and $V_2 = \{D, E\}$. For the constraint $\sum_i V_1[1i] \in [l_B, u_B]$, $\sum_i V_1[1i]$ represents the probability of relation B , i.e., $P(B)$, and $[l_B, u_B]$ is




Fig. 23. An example derivation graph.

$ \begin{aligned} \phi : \quad & \sum_{b \in \mathbb{B}^4} V_1[b] = 1 \quad (I_1, I_2, I_3, I_4) \\ & \sum_{b \in \mathbb{B}^4} V_2[b] = 1 \quad (I_5, I_6, I_7, I_8) \\ & \forall b \in \mathbb{B}^4. V_1[b] \in [0, 1], V_2[b] \in [0, 1] \\ & \sum_{i,j,k} V_1[1ijk] = P(I_1) \\ & \sum_{i,j,k} V_1[i1jk] = P(I_2) \\ & \sum_{i,j,k} V_1[ij1k] = P(I_3) \\ & \sum_{i,j,k} V_1[ijk1] = P(I_4) \\ & \sum_{i,j,k} V_2[1ijk] = P(I_5) \\ & \sum_{i,j,k} V_2[i1jk] = P(I_6) \\ & \sum_{i,j,k} V_2[ij1k] = P(I_7) \\ & \sum_{i,j,k} V_2[ijk1] = P(I_8) \\ & \phi_{corr} \end{aligned} $	$ \begin{aligned} \phi' : \quad & \sum_{b \in \mathbb{B}^2} V_1[b] = 1 \quad (B, C) \\ & \sum_{b \in \mathbb{B}^2} V_2[b] = 1 \quad (D, E) \\ & \forall b \in \mathbb{B}^2. V_1[b] \in [0, 1], V_2[b] \in [0, 1] \\ & \sum_i V_1[1i] \in [l_B, u_B] \\ & \sum_i V_1[i1] \in [l_C, u_C] \\ & \sum_i V_2[1i] \in [l_D, u_D] \\ & \sum_i V_2[i1] \in [l_E, u_E] \\ & \phi'_{corr} \end{aligned} $
---	---

 Fig. 24. Original constraints ϕ vs. overapproximated constraints ϕ' .

obtained from the correlation analysis. For example, assuming that l_B^*, u_B^* denote the ground truth lower and upper bounds of $P(B)$ satisfying ϕ . Since our correlation analysis is sound, we know that $l_B \leq l_B^* \leq u_B^* \leq u_B$. As a result, $\phi \implies l_B \leq P(B) \leq u_B$. A similar implication happens to C, D, E . Meanwhile, with the correlation analysis, we can get the correlation type between a pair of relations from $\{B, C, D, E\}$, and encode them as constraints ϕ'_{corr} . Due to the soundness of correlation analysis, according to Theorem 3, we know that $\phi \implies \phi'_{corr}$.

To summarize, $\phi \implies \phi'$.

In terms of performance improvement, it is because the approximated version has fewer joint probability variables compared to the original version. We can see that the number of variables in ϕ' is significantly less than the number of variables in ϕ . For instance, ϕ has 32 variables, while ϕ' only has 8 variables. It would significantly reduce the number of variables when this derivation graph is larger, and we pick a suitable step k to select *intermediate relations*.

It is worth noting that k is not a fixed value. The core idea behind our over-approximation is to partition the derivation graph G into two disjoint sets, S_s and S_t , such that $S_s \cup S_t = \text{Nodes}(G)$ and $S_s \cap S_t = \emptyset$. We omit all edges that connect nodes from S_s to nodes in S_t .

Assume that S_s contains the root node (e.g., the output fact A in Figure 23). Instead of expressing the probability of A in terms of input facts, our over-approximated formulation encodes this

probability based on the leaf nodes within S_s . Here, we define leaf nodes as those without outgoing edges.

For example, in Figure 23, one possible cut yields $S_s = \{A, B, C, D, E\}$ and $S_t = \{I1, I2, I3, I4, I5, I6, I7, I8\}$. Alternatively, a valid cut could be $S_s = \{A, B, C, D, E, I5, I6, I7, I8\}$ and $S_t = \{I1, I2, I3, I4\}$. We encode the cut selection as an ILP (Integer Linear Programming) problem, where the objective is to choose the leaf nodes of S_s along the cut such that the number of incoming edges to these leaf nodes is maximized.

A.5 Inferring Symbolic Probability Expressions

THEOREM 1. Let E_1 and E_2 denote the probability expressions of events A and B respectively. Then, we have: (1) If $\ominus E_1 \rightsquigarrow E$, then E represents the probability of $\neg A$; (1) if $E_1 \oplus E_2 \rightsquigarrow E$, then E represents the probability of event $A \vee B$ (3) iff $E_1 \otimes E_2 \rightsquigarrow E$, then E represents the probability of event $A \wedge B$

In this subsection, we present soundness proofs for \ominus , \otimes and \oplus computations used in deriving symbolic probability expressions. We also present some Lemmas that are used to prove the soundness of both \otimes and \oplus computation.

A.5.1 THEOREM 1 (\ominus). Let E_A denote the probability of event A , then we have:

If $\ominus E_A \rightsquigarrow E$, then E represents the probability of event $\neg A$.

PROOF. Let $Pr(A) = \text{Expr}(A) = \xi[\vec{\sigma}_A]$ where $\vec{\sigma}_A = \{\lambda_0, \lambda_1, \dots, \lambda_i, \dots\}$, as per Definition 10. Hence, $E_A = Pr(A) = \sum_i \lambda_i \psi_i$.

Since $Pr(\neg A) = 1 - Pr(A)$, it follows that:

$$Pr(\neg A) = 1 - E_A = 1 - \sum_i \lambda_i \psi_i.$$

From Definition 8, we know $\sum_i \psi_i = 1$ as it spans the complete joint probability space. This leads to:

$$Pr(\neg A) = 1 - \sum_i \lambda_i \psi_i = \sum_i \psi_i - \sum_i \lambda_i \psi_i = \sum_i (1 - \lambda_i) \psi_i.$$

By the \ominus rule, $\ominus E_A \rightsquigarrow E$ yields $E = \sum_i (1 - \lambda_i) \psi_i$. Therefore, we conclude:

$$Pr(\neg A) = E.$$

□

LEMMA 3. Given two events A and B , if $Pr(A) = \psi_A$, $Pr(B) = \psi_B$ and $\psi_A \neq \psi_B$, then events A and B are mutually exclusive.

PROOF. Based on the definition of product terms (Definition 7), we have $P(A) = \psi_A = \prod_{c \in C} V_c[b]$, $P(B) = \psi_B = \prod_{c \in C} V_c[b']$, and $b \neq b'$. Both b and b' represent bit-vectors. Here, we use V_c to denote $\text{Rep}(c)$ for the ease of presentation.

Given a boolean event A , we can rewrite $A = A_1 \wedge A_2 \wedge A_n$, where $n = |C|$, $P(A_i) = V_i[b]$. Based on the independence definition of distinct correlation classes, we know that $P(A) = P(A_1)P(A_2) \dots P(A_n)$. Similarly, $P(B) = P(B_1)P(B_2) \dots P(B_n)$. For the same correlation class i , $P(A_i) = V_i[b]$ and $P(B_i) = V_i[b']$. For instance, $b = 001$ and $b' = 100$, as for b , it represents the joint probability of the first two events being false and the third event being true, whereas b' represents the joint probability of the first event being true while the last two events being false. Hence, A_i and B_i represent two mutually-exclusive events, i.e., $P(A_i B_i) = 0$. The joint probability of $P(A \wedge B)$ is as follows:

$$P(A \wedge B) = \prod_{1 \leq i \leq |C|} P(A_i B_i) \quad \text{using definitions of correlation class}$$

Hence, $P(A \wedge B) = 0$ as $P(A_i B_i) = 0$. We can also get $P(A|B) = P(A \wedge B)/P(B) = 0$. \square

A.5.2 THEOREM 1 (\otimes). Let E_A and E_B denote the probability of events A and B respectively. Then, we have:

If $\vdash E_A \otimes E_B \rightsquigarrow E$, then E represents the probability of event $A \wedge B$.

PROOF. Given $E_A = P(A)$, $E_B = P(B)$, we will first compute $P(A \wedge B)$ and prove it is equivalent to E .

Assuming that $E_A = \sum \lambda_k^A \psi_k$, $E_B = \sum \lambda_k^B \psi_k$ and $\psi_k = \prod_{c \in C} V_c[b]$. Let e_k be an event, where $P(e_k) = \psi_k$. Let e^- be the event which is a complement of the event $e_1 \vee \dots \vee e_k \dots \vee e_n$.

$$\begin{aligned}
 (1) P(A \wedge B) &= P(A \wedge B \wedge e_1) + \dots P(A \wedge B \wedge e_k) + \dots P(A \wedge B \wedge e_n) + P(A \wedge B \wedge e^-) \\
 (2) &= P(A \wedge B \wedge e_1) + \dots P(A \wedge B \wedge e_k) + \dots P(A \wedge B \wedge e_n) \\
 (3) &= P(A \wedge B|e_1)P(e_1) + \dots P(A \wedge B|e_k)P(e_k) + \dots P(A \wedge B|e_n)P(e_n) \\
 (4) &= P(A_1 \wedge B_1)P(e_1) + \dots P(A_n \wedge B_n)P(e_n) \\
 (5) &= P(\bigvee_{i \in I} A_1^i \wedge \bigvee_{j \in J} B_1^j)P(e_1) + \dots \\
 (6) &= P(\bigvee_{i,j \in I \times J} A_1^i \wedge B_1^j)P(e_1) + \dots \\
 (7) &= (\sum_{i,j \in I \times J} P(A_1^i \wedge B_1^j))P(e_1) + \dots
 \end{aligned}$$

STEP (1): Since each e_i is mutually exclusive (according to Lemma 3), we can get $P(e^-) + \sum_i P(e_i) = 1$. Hence, we can rewrite $P(A \wedge B)$ as $P(A \wedge B \wedge e_1) + \dots P(A \wedge B \wedge e_k) + \dots P(A \wedge B \wedge e_n) + P(A \wedge B \wedge e^-)$.

TRANSFORMATION FROM (1) - (2): Given that $P(A) = \sum \lambda_k^A P(e_k)$ and each e_k is mutually exclusive, we can infer that the boolean event $A = \vee e'_k$, where e'_k is an event with probability $P(e'_k) = \lambda_k^A P(e_k)$. Let's assume an event c with probability $P(c) = \lambda_k^A$. As we know that c is the coefficient which is independent of the term event e_k , we can get $e'_k = e_k \wedge c$. Hence $P(e'_k|e_k) = P(c \wedge e_k|e_k) = P(c)P(e_k)/P(e_k) = \lambda_k^A$, and $P(e_k \wedge e'_k) = P(e'_k)$.

We have the following:

$$\begin{aligned}
 P(e^- \wedge (e'_1 \vee e'_2 \vee \dots e'_n)) &= P(\overline{e^- \wedge e'_1 \vee e^- \wedge e'_2 \vee \dots e^- \wedge e'_n}) \\
 &= P(\overline{e^- \wedge e'_1 \wedge e_1 \vee e^- \wedge e'_2 \wedge e_2 \vee \dots e^- \wedge e'_n \wedge e_n}) \\
 &= P(\overline{e^- \wedge e'_1 \wedge e_1}) + P(\overline{e^- \wedge e'_2 \wedge e_2}) + \dots P(\overline{e^- \wedge e'_n \wedge e_n}) \\
 &= 0
 \end{aligned}$$

From the above, we can get e^- is also mutually exclusive with the union of $e'_1 \vee e'_2 \vee \dots e'_n$. As $A = \vee e'_k$, we can get e^- is mutually exclusive with A . As a result, $P(A \wedge B \wedge e^-) = 0$.

TRANSFORMATION FROM (2) - (3): According to Bayes' theorem, $P(A \wedge B \wedge e_1) = P(A \wedge B|e_1)P(e_1)$. This applies to e_2, \dots, e_n .

TRANSFORMATION FROM (3) - (4): For instance, we introduce events A_1 and B_1 such that $P(A_1) = \lambda_1^A$ and $P(B_1) = \lambda_1^B$. We have $(A_1 \wedge B_1) \perp\!\!\!\perp e_1$ (i.e., the product of coefficients λ are independent of the product terms ψ). Since $A = (A_1 \wedge e_1) \vee \dots (A_n \wedge e_n)$ and $B = (B_1 \wedge e_1) \vee \dots (B_n \wedge e_n)$, we have $A \wedge B = (A_1 \wedge B_1 \wedge e_1) \vee \dots (A_n \wedge B_n \wedge e_n)$. Thus $P(A \wedge B|e_1) = P(A_1 \wedge B_1 \wedge e_1|e_1) = P(A_1 \wedge B_1)$.

TRANSFORMATION FROM (4) - (5): Event A_1 can be represented as a DNF formula: $A_1 = \bigvee_{i \in I} A_1^i$, where each A_1^i is a disjoint conjunctive clause. $A_1^i = \bigwedge_{x^+ \in X^+} x^+ \wedge_{x^- \in X^-} x^-$, where x^+ represents the rule probability variable and x^- represents the negated rule probability variable. We also use P_{x^+} to represent the concrete probability value for variable x^+ , similar to P_{x^-} . The probability of each conjunctive clause $P(A_1^i) = \prod_{x^+ \in X^+} P_{x^+} \prod_{x^- \in X^-} (1 - P_{x^-})$.

TRANSFORMATION FROM (5) - (6): We follow the distributive law of conjunction over disjunction operator.

TRANSFORMATION FROM (6) - (7): Each $A_1^i \wedge B_1^j$ is mutually disjoint from another one $A_1^{i'} \wedge B_1^{j'}$ if $i \neq i' \vee j \neq j'$. Due to this mutual disjoint relation, we can rewrite the probability of a DNF formula over disjoint conjunction as a summation of the conjunction probabilities.

When computing $P(A_1^i \wedge B_1^j)$, we have to consider two cases. Since $P(A_1^i) = \prod_{x^+ \in X^+} P_{x^+} \prod_{x^- \in X^-} (1 - P_{x^-})$ and $P(B_1^j) = \prod_{y^+ \in Y^+} P_{y^+} \prod_{y^- \in Y^-} (1 - P_{y^-})$, if the set $X^+ \cap Y^- \neq \emptyset$ or $X^- \cap Y^+ \neq \emptyset$, we know that one term in A_1^i is a negated version of the term in B_1^j , so that their conjunction is 0. Otherwise, without the contradiction, we can compute the joint probability of A_1^i and B_1^j . If there are no overlapping terms, e.g. $X^+ \cap Y^+ = \emptyset$, due to the mutual independence among rule probabilities, we can have $P(A_1^i \wedge B_1^j) = P(A_1^i)P(B_1^j)$. Otherwise, if there are overlapping terms and denote the conjunction of overlapping terms as S , let's assume that $A_1^i = a_1 \wedge \dots \wedge S \wedge a_n$ and $B_1^j = b_1 \wedge \dots \wedge S \wedge b_n$, then $A_1^i \wedge B_1^j = (a_1 \wedge \dots \wedge S \wedge a_n) \wedge (b_1 \wedge \dots \wedge S \wedge b_n) = a_1 \wedge \dots \wedge S \wedge a_n \wedge b_1 \wedge \dots \wedge b_n$ by removing the overlapping part S .

By following the above computation, we have the formula in $(\sum_{i,j \in I \times J} P(A_1^i \wedge B_1^j))$ shown in Line (7) is evaluated the same as $\text{JointProb}(P(A_1), P(B_1))$. Hence, we can have the following:

$$\begin{aligned} P(A \wedge B) &= (\sum_{i,j \in I \times J} P(A_1^i \wedge B_1^j)) \cdot P(e_1) \dots + \sum_{i,j \in I \times J} P(A_n^i \wedge B_n^j) P(e_n) \\ &= \text{JointProb}(P(A_1), P(B_1)) \cdot P(e_1) \dots + \text{JointProb}(P(A_n), P(B_n)) \cdot P(e_n) \\ &= \sum \text{JointProb}(P(A_i), P(B_i)) \cdot P(e_i) \\ &= P(A) \otimes P(B) \end{aligned}$$

□

A.5.3 THEOREM 1 (\oplus). Let E_A and E_B denote the probability of events A and B respectively. Then, we have:

If $\vdash E_A \oplus E_B \rightsquigarrow E$, then E represents the probability of event $A \vee B$.

PROOF. Here, we do not provide a formal proof for \oplus as it is very similar to the proof of Theorem 1 (\otimes), which mainly utilizes the Lemma 3.

A.6 Inferring Correlation Types

In this section, we give soundness proof for Theorem 2 and Theorem 3.

A.6.1 Soundness of Theorem 2. We prove the soundness of this theorem for each $\star \in \{+, -, \perp\}$:

- **Case $+$:** There are four rules in Figure 12 to infer the positive correlation $I_1 \blacktriangleright^+ I_2$.
 - RULE ID states that $I \blacktriangleright^+ I$, because $P(I|I) = 1 > P(I)$.
 - RULE SYMM states that $I_2 \blacktriangleright^+ I_1$ if $I_1 \blacktriangleright^+ I_2$. Given $I_1 \blacktriangleright^+ I_2$, we can infer that $P(I_1|I_2) > P(I_1)$. Thus, $P(I_2|I_1) = P(I_1|I_2)P(I_2)/P(I_1) > P(I_1)P(I_2)/P(I_1) = P(I_2)$, which implies $I_2 \blacktriangleright^+ I_1$.
 - RULE SEMANTIC states that $I_1 \blacktriangleright^+ I_2$ if $P(I_1 \wedge I_2) > P(I_1)P(I_2)$. With this information, we can infer that $P(I_1|I_2) = P(I_1 \wedge I_2)/P(I_2) > P(I_1)P(I_2)/P(I_2) = P(I_1)$, implying $I_1 \blacktriangleright^+ I_2$.
- **Case $-$:** There are three rules in Figure 12 to infer the negative correlation $I_1 \blacktriangleright^- I_2$.
 - RULE SYMM states that $I_2 \blacktriangleright^- I_1$ if $I_1 \blacktriangleright^- I_2$. Given $I_1 \blacktriangleright^- I_2$, we can infer that $P(I_1|I_2) < P(I_1)$. Thus, $P(I_2|I_1) = P(I_1|I_2)P(I_2)/P(I_1) < P(I_1)P(I_2)/P(I_1) = P(I_2)$, which implies $I_2 \blacktriangleright^- I_1$.
 - RULE SEMANTIC states that $I_1 \blacktriangleright^- I_2$ if $P(I_1 \wedge I_2) < P(I_1)P(I_2)$. With this information, we can infer that $P(I_1|I_2) = P(I_1 \wedge I_2)/P(I_2) < P(I_1)P(I_2)/P(I_2) = P(I_1)$, implying $I_1 \blacktriangleright^- I_2$.
- **Case \perp :** There are three rules in Figure 12 to infer the independence relation $I_1 \blacktriangleright^\perp I_2$.
 - RULE INDEP states that $I_1 \blacktriangleright^\perp I_2$ if $\text{Class}(I_1) \neq \text{Class}(I_2)$. This is correct based on our assumption that elements from different correlation classes are independent.
 - RULE SYMM states that $I_2 \blacktriangleright^\perp I_1$ if $I_1 \blacktriangleright^\perp I_2$. Given $I_1 \blacktriangleright^\perp I_2$, we can infer that $P(I_1|I_2) = P(I_1)$. Thus, $P(I_2|I_1) = P(I_1|I_2)P(I_2)/P(I_1) = P(I_1)P(I_2)/P(I_1) = P(I_2)$, which implies $I_2 \blacktriangleright^\perp I_1$.

- **RULE SEMANTIC** states that $I_1 \blacktriangleright^\perp I_2$ if $P(I_1 \wedge I_2) = P(I_1)P(I_2)$. With this information, we can infer that $P(I_1|I_2) = P(I_1 \wedge I_2)/P(I_2) = P(I_1)P(I_2)/P(I_2) = P(I_1)$, implying $I_1 \blacktriangleright^\perp I_2$.

A.6.2 Soundness of Theorem 3. In the rest of this section, we use the \perp symbol to indicate statistical independence and prove a series of lemmas that are used to prove Theorem 3.

Proof of Theorem 3. To establish the proof, we need to prove the correctness of three rules (Rule INDEP, Rule POS and Rule NEG) in Figure 14.

In the following, we use Lemma 6 to establish the correctness of Rule INDEP, Lemma 10 to establish the correctness of Rule POS and Lemma 14 for Rule NEG.

Lemma 10 establishes the correctness of Rule POS in Figure 14, used for inferring the positive correlation between a pair of expressions E_1 and E_2 . As E_1 and E_2 are arbitrary expressions in Lemma 10, we can use DNF as a general format to represent both E_1 and E_2 . In order to prove Lemma 10, we need to use the following Lemmas 4, 7, 8 and 9 that are also used to prove the soundness of the rule POS but with a restricted format of E_1 and E_2 .

Similarly, to prove Lemma 14 for Rule NEG in Figure 14, we use the following Lemmas: Lemma 4, Lemma 11, Lemma 12, and Lemma 13.

LEMMA 4. If $\frac{\partial f}{\partial \alpha} > 0$ and $\frac{\partial f}{\partial \beta} > 0$, function $f(\alpha, \beta)$ is strictly monotonically increasing, i.e., $f(\alpha_1, \beta_1) > f(\alpha_2, \beta_2)$, where $\alpha_1 > \alpha_2$ and $\beta_1 > \beta_2$.

PROOF of Lemma 4.

We assume the following:

$$f_\alpha(\alpha, \beta) = \frac{\partial f}{\partial \alpha} \quad f_\beta(\alpha, \beta) = \frac{\partial f}{\partial \beta}$$

For a function f of two variables, the Mean Value Theorem states that:

$$f(\alpha_1, \beta_1) - f(\alpha_2, \beta_2) = f_\alpha(\chi, \eta) \cdot (\alpha_1 - \alpha_2) + f_\beta(\chi, \eta) \cdot (\beta_1 - \beta_2)$$

where (χ, η) lies on the line segment joining (α_1, β_1) and (α_2, β_2) .

Given that $\alpha_1 > \alpha_2$ and $\beta_1 > \beta_2$, and knowing that $f_\alpha(\chi, \eta) > 0$ and $f_\beta(\chi, \eta) > 0$, we have:

$$f(\alpha_1, \beta_1) - f(\alpha_2, \beta_2) = f_\alpha(\chi, \eta) \cdot (\alpha_1 - \alpha_2) + f_\beta(\chi, \eta) \cdot (\beta_1 - \beta_2) > 0$$

This is because each term in the sum is positive, given the positivity of the partial derivatives and the non-negativity of the differences $\alpha_1 - \alpha_2$ and $\beta_1 - \beta_2$.

Therefore, $f(\alpha_1, \beta_1) > f(\alpha_2, \beta_2)$ if $\alpha_1 > \alpha_2$ and $\beta_1 > \beta_2$. This completes the proof that $f(\alpha, \beta)$ is strictly increasing in both α and β . \square

LEMMA 5. Given two different correlation classes V_1 and V_2 , if we have two arbitrary boolean expressions $E_1 = f(x)$, $x \in V_1$, and $E_2 = g(y)$, $y \in V_2$, then $E_1 \perp E_2$.

PROOF. We omit the proof here since this is our definition of correlation classes.

LEMMA 6. Given a pair of E_1 and E_2 , if $\forall x, y \in \text{Dep}(E_1) \times \text{Dep}(E_2)$. $\text{Class}(x) \neq \text{Class}(y)$, then $E_1 \perp E_2$.

PROOF of Lemma 6.

Let $S_1 = \text{Dep}(E_1)$ and $S_2 = \text{Dep}(E_2)$. We define the following function:

$$\mathcal{V}(E) = \{V | V = \text{Class}(x), x \in \text{Dep}(E)\} \quad (11)$$

Based on the premises of rule INDEP, we know that $\forall x, y \in S_1 \times S_2$. $\text{Class}(x) \neq \text{Class}(y)$, we can get $\mathcal{V}(E_1) \cap \mathcal{V}(E_2) = \emptyset$. With this, we can assume that all the variables (input facts) appearing in

E_1 belong to a single correlation class $V_1 = \bigcup_{i \in \mathcal{V}(E_1)} V_i$ and all the variables appearing in E_2 belong to a single correlation class $V_2 = \bigcup_{j \in \mathcal{V}(E_2)} V_j$.

We can express E_1 and E_2 as follows:

$$E_1 = f(x_1, x_2, \dots, x_n), x_i \in V_1, E_2 = g(y_1, y_2, \dots, y_n), y_i \in V_2$$

Given that V_1 and V_2 are two different correlation classes, based on Lemma 5, we know that any arbitrary boolean combinations of variables from V_1 are mutually independent from the boolean combinations of variables from V_2 . Thus, $E_1 \perp\!\!\!\perp E_2$.

Note that, in general, independence does not necessarily propagate to conjunctions of conditions, e.g., given $x_1 \perp\!\!\!\perp x$ and $x_2 \perp\!\!\!\perp x$, it does not imply that $x \perp\!\!\!\perp (x_1 \wedge x_2)$. However, this is a valid assumption in our domain. Since all the input facts x, y are atomic, it is natural to assume the mutual independence between arbitrary boolean formulas $f(x)$ and $g(y)$ if x and y are from different correlation classes.

□

LEMMA 7. *Given an input fact $E_1 = I$ and a general DNF formula $E_2 = \bigvee \bigwedge_j I_j$, if $\vdash \chi(E_1), \vdash \chi(E_2), \vdash E_1 \rightarrow^+ E_2$ and $\not\vdash E_1 \rightarrow^- E_2$, then $P(E_1 \wedge E_2) > P(E_1)P(E_2)$, i.e., $E_1 \xrightarrow{\oplus} E_2$.*

PROOF of Lemma 7. We prove this by induction and start with the base case.

Base cases: $E_1 = I, E_2 = \bigwedge_i I_i$.

We can compute the expression $P(E_1 \wedge E_2)$ as follows:

$$\begin{aligned} (1) P(E_1 \wedge E_2) &= P(I \wedge \bigwedge_i I_i) \\ (2) &= P(I \wedge I_k \wedge \bigwedge_{i \neq k} I_i) && \exists I_k \in \text{Dep}(E_2). I_k \blacktriangleright^+ I \text{ (i.e., } \vdash E_1 \rightarrow^+ E_2) \\ (3) &= P(I \wedge I_k) \prod_{i \neq k} P(I_i) && \text{Definition of } \chi(E_2), I_1 \perp\!\!\!\perp I_2 \perp\!\!\!\perp \dots (I_k, I) \perp\!\!\!\perp \dots I_i \\ (4) &= P(I|I_k)P(I_k) \prod_{i \neq k} P(I_i) \\ (5) &= P(I|I_k) \prod_i I_i \end{aligned}$$

The expression of $P(E_1)P(E_2)$ is as follows:

$$\begin{aligned} (1) P(E_1)P(E_2) &= P(I)P(\bigwedge_i I_i) \\ (2) &= P(I) \prod_i I_i && \text{Definition of } \chi(E_2) \end{aligned}$$

Based on the precondition $\vdash E_1 \rightarrow^+ E_2$, we know that $\exists I_k \in \text{Dep}(E_2). I_k \blacktriangleright^+ E_1$, which is $I_k \blacktriangleright^+ I$. It implies that $P(I|I_k) > P(I)$, and thus $P(E_1 \wedge E_2) > P(E_1)P(E_2)$, i.e., $E_1 \xrightarrow{\oplus} E_2$.

Inductive cases: $E_1 = I, E_2 = E'_2 \vee t$, where $E'_2 = \bigvee \bigwedge_i I_i, t = \bigwedge_i I_i$.

Case 1: $E_1 \xrightarrow{\oplus} E'_2$. Given $E_1 = I$, we have the following three subcases.

- Case 1.1: $E_1 \rightarrow^+ t$ violates the precondition $\chi(E_2)$.
- Case 1.2: $E_1 \rightarrow^- t$ violates the precondition $\not\vdash E_1 \rightarrow^- E_2$.
- Case 1.3: $E_1 \perp\!\!\!\perp t$ is valid.

PROOF of Case 1.3

Given $E_1 \xrightarrow{\oplus} E'_2$, and $E_1 = I$, we can get $\mathcal{V}(E_1) \subseteq \mathcal{V}(E'_2)$, where $\mathcal{V}(E_1)$ is defined by Eq 11. Given $\chi(E_2)$, we know that the correlation classes for each $x \in \text{Dep}(E_2)$ are disjoint. With $E_2 = E'_2 \vee t$, we can get $\mathcal{V}(E'_2) \cap \mathcal{V}(t) = \emptyset$. With $\mathcal{V}(E_1) \subseteq \mathcal{V}(E'_2)$, we can get $\mathcal{V}(E_1) \cap \mathcal{V}(t) = \emptyset$, thus $E_1 \perp\!\!\!\perp t$.

The inductive case focuses on proving the following: given $E_1 \xrightarrow{\oplus} E'_2$, and $E_1 \perp\!\!\!\perp t$, it implies that $E_1 \xrightarrow{\oplus} E'_2 \vee t$.

$$\begin{aligned}
 (1) P(E_1 \wedge E_2) &= P(I \wedge E_2) & E_1 = I \\
 (2) &= P((I \wedge E'_2) \vee (I \wedge t)) & E_2 = E'_2 \vee t \\
 (3) &= P(I \wedge E'_2) + P(I \wedge t) - P(I \wedge E'_2 \wedge t) \\
 (4) &= P(I \wedge E'_2) + P(I)P(t) - P(I \wedge E'_2)P(t) & \text{definitions of } \chi(E) : \{I, E'_2\} \perp\!\!\!\perp t. \\
 (1) P(E_1)P(E_2) &= P(I)P(E'_2 \vee t) & E_1 = I \\
 (2) &= P(I)(P(E'_2) + P(t) - P(E'_2)P(t)) & \text{definitions of and } \chi(E) \\
 (3) &= P(I)P(E'_2) + P(I)P(t) - P(I)P(E'_2)P(t)
 \end{aligned}$$

From above expressions, we can compute $P(E_1 \wedge E_2) - P(E_1)P(E_2)$ as Δ :

$$\begin{aligned}
 (1) \Delta &= P(I \wedge E'_2) - P(I \wedge E'_2 \wedge t) - P(I)P(E'_2) + P(I)P(E'_2)P(t) \\
 (2) &= P(I|E'_2)P(E'_2) - P(I|E'_2)P(E'_2)P(t) - P(I)P(E'_2) + P(I)P(E'_2)P(t) \\
 (3) &= (1 - P(t))(P(I|E'_2)P(E'_2) - P(I)P(E'_2))
 \end{aligned}$$

As $E_1 \xrightarrow{\oplus} E'_2$ and $E_1 = I$, we can get $P(I|E'_2)P(E'_2) > P(I)P(E'_2)$, and thus $\Delta > 0$, i.e., $P(E_1 \wedge E_2) > P(E_1)P(E_2)$. Hence, we are able to prove $E_1 \xrightarrow{\oplus} E'_2 \vee t$.

Case 2: $E_1 \perp\!\!\!\perp E'_2$.

- Case 2.1: $E_1 \perp\!\!\!\perp t$ violates the precondition $\vdash E_1 \rightarrow^+ E_2$.
- Case 2.2: $E_1 \rightarrow^- t$ violates the precondition $\nvdash E_1 \rightarrow^- E_2$.
- Case 2.3: $E_1 \rightarrow^+ t$ is valid.

The proof of Case 2.3 is similar to Case 1.3, so we omit it here.

Other cases

Case 3: $E_1 \xrightarrow{\ominus} E'_2$ is invalid as it violates the precondition $\nvdash E_1 \rightarrow^- E_2$. \square

LEMMA 8. *Let E_1, E_2 be conjunctions over input facts, i.e., $E_1 = \bigwedge_i I_i$ and $E_2 = \bigwedge_j I_j$. If $\vdash \chi(E_1), \vdash \chi(E_2), \vdash E_1 \rightarrow^+ E_2$ and $\nvdash E_1 \rightarrow^- E_2$, then $P(E_1 \wedge E_2) > P(E_1)P(E_2)$.*

PROOF of Lemma 8. We prove this by induction and start with the base cases.

Base cases: $E_1 = I_{ai} \wedge I_{bi}$ and $E_2 = I_{aj} \wedge I_{bj}$

To compute $P(E_1 \wedge E_2)$, we first list all the possible dependency information among $\{I_{ai}, I_{bi}, I_{aj}, I_{bj}\}$. Given the precondition $\chi(E_1)$ and $\chi(E_2)$, we know that $I_{ai} \perp\!\!\!\perp I_{bi}$ and $I_{aj} \perp\!\!\!\perp I_{bj}$. Based on the precondition $\vdash E_1 \rightarrow^+ E_2, \nvdash E_1 \rightarrow^- E_2$, we can list all the cases satisfying these conditions:

- (1) $I_{ai} \triangleright^+ I_{aj}, I_{bi} \triangleright^+ I_{bj}$: In this case, these four input facts are grouped to two correlation classes: $\{I_{ai}, I_{aj}\}, \{I_{bi}, I_{bj}\}$, thus $P(E_1)P(E_2) = P(I_{ai} \wedge I_{bi})P(I_{aj} \wedge I_{bj}) = P(I_{ai})P(I_{bi})P(I_{aj})P(I_{bj})$. Meanwhile, we can rewrite the joint probability as follows: $P(E_1 \wedge E_2) = P(I_{ai} \wedge I_{bi} \wedge I_{aj} \wedge I_{bj}) = P(I_{ai} \wedge I_{aj})P(I_{bi} \wedge I_{bj}) = P(I_{ai}|I_{aj})P(I_{aj})P(I_{bi}|I_{bj})P(I_{bj})$. Given $I_{ai} \triangleright^+ I_{aj}, I_{bi} \triangleright^+ I_{bj}$, we know that $P(I_{ai}|I_{aj}) > P(I_{ai})$ and $P(I_{bi}|I_{bj}) > P(I_{bi})$. Thus, $P(E_1 \wedge E_2) = P(I_{ai}|I_{aj})P(I_{aj})P(I_{bi}|I_{bj})P(I_{bj}) > P(I_{ai})P(I_{aj})P(I_{bi})P(I_{bj}) = P(E_1)P(E_2)$.
- (2) $I_{ai} \triangleright^+ I_{bj}, I_{bi} \triangleright^+ I_{aj}$. The proof of this is similar to case (1).
- (3) $I_{ai} \triangleright^+ I_{aj}, I_{bi} \perp\!\!\!\perp I_{bj}$: In this case, these four input facts are grouped to three correlation classes: $\{I_{ai}, I_{aj}\}, \{I_{bi}\}, \{I_{bj}\}$, thus $P(E_1)P(E_2) = P(I_{ai} \wedge I_{bi})P(I_{aj} \wedge I_{bj}) = P(I_{ai})P(I_{bi})P(I_{aj})P(I_{bj})$. Similarly, $P(E_1 \wedge E_2) = P(I_{ai} \wedge I_{bi} \wedge I_{aj} \wedge I_{bj}) = P(I_{ai} \wedge I_{aj})P(I_{bi} \wedge I_{bj}) = P(I_{ai}|I_{aj})P(I_{aj})P(I_{bi}|I_{bj})P(I_{bj})$. Given $I_{ai} \triangleright^+ I_{aj}$ and $I_{bi} \perp\!\!\!\perp I_{bj}$, we know that $P(I_{ai}|I_{aj}) > P(I_{ai})$, and $P(I_{bi}|I_{bj}) = P(I_{bi})$. As a result, $P(E_1 \wedge E_2) > P(E_1)P(E_2)$.
- (4) $I_{ai} \triangleright^+ I_{bj}, I_{bi} \perp\!\!\!\perp I_{aj}$.
- (5) $I_{ai} \perp\!\!\!\perp I_{bj}, I_{bi} \triangleright^+ I_{aj}$.
- (6) $I_{ai} \perp\!\!\!\perp I_{aj}, I_{bi} \triangleright^+ I_{bj}$.

The proof of cases (4)-(6) is similar to case (3), so we omit the proof here.

Inductive cases.

$$\begin{array}{lcl}
 E_1 & = & E'_1 \wedge I_1 \quad \Bigg| \quad E'_1 = \bigwedge_i I_i \\
 E_2 & = & E'_2 \wedge I_2 \quad \Bigg| \quad E'_2 = \bigwedge_j I_j
 \end{array}$$

Case 1: $E'_1 \xrightarrow{\oplus} E'_2$. Based on the precondition $\chi(E_1)$ (resp. $\chi(E_2)$), we know that $E'_1 \perp\!\!\!\perp I_1$ (resp. $E'_2 \perp\!\!\!\perp I_2$).

- Case 1.1 $I_1 \rightarrow^+ I_2$, $I_1 \perp\!\!\!\perp E'_2$, $I_2 \perp\!\!\!\perp E'_1$.

PROOF OF CASE 1.1 $E'_1 \xrightarrow{\oplus} E'_2$, $I_1 \rightarrow^+ I_2$, $I_1 \perp\!\!\!\perp E'_2$, $I_2 \perp\!\!\!\perp E'_1$

$$\begin{aligned}
 P(E_1 \wedge E_2) &= P(E'_1 \wedge I_1 \wedge E'_2 \wedge I_2) \\
 &= P(E'_1 \wedge E'_2)P(I_1 \wedge I_2) && \text{definitions of } \{E'_1, E'_2\} \perp\!\!\!\perp \{I_1, I_2\} \\
 &= P(E'_1|E'_2)P(E'_2)P(I_1|I_2)P(I_2) \\
 P(E_1)P(E_2) &= P(E'_1 \wedge I_1)P(E'_2 \wedge I_2) \\
 &= P(E'_1)P(I_1)P(E'_2)P(I_2) && \text{definitions of } \chi(E_1) \text{ and } \chi(E_2)
 \end{aligned}$$

In case 1.1, given $E'_1 \xrightarrow{\oplus} E'_2$, we can infer that $P(E'_1|E'_2) > P(E'_1)$. Given $I_1 \rightarrow^+ I_2$, as both I_1 and I_2 are input facts, we know that $I_1 \blacktriangleright^+ I_2$, and thus $P(I_1|I_2) > P(I_1)$. As a result, $P(E_1 \wedge E_2) > P(E_1)P(E_2)$.

- Case 1.2 $I_1 \rightarrow^+ I_2$, $I_1 \perp\!\!\!\perp E'_2$, $I_2 \rightarrow^+ E'_1$. It is invalid as it violates the precondition $\chi(E_1)$.
- Case 1.3 $I_1 \rightarrow^+ I_2$, $I_1 \rightarrow^+ E'_2$, $I_2 \perp\!\!\!\perp E'_1$. It is invalid as it violates the preconditions $\chi(E_2)$.
- Case 1.4 $I_1 \rightarrow^+ I_2$, $I_1 \rightarrow^+ E'_2$, $I_2 \rightarrow^+ E'_1$. It is invalid as it violates the precondition $\chi(E_1)$ and $\chi(E_2)$.
- Case 1.5 $I_1 \perp\!\!\!\perp I_2$, $I_1 \perp\!\!\!\perp E'_2$, $I_2 \perp\!\!\!\perp E'_1$

PROOF OF CASE 1.5 $E'_1 \xrightarrow{\oplus} E'_2$, $I_1 \perp\!\!\!\perp I_2$, $I_1 \perp\!\!\!\perp E'_2$, $I_2 \perp\!\!\!\perp E'_1$

$$\begin{aligned}
 P(E_1 \wedge E_2) &= P(E'_1 \wedge I_1 \wedge E'_2 \wedge I_2) \\
 &= P(E'_1 \wedge E'_2)P(I_1)P(I_2) && \text{definitions of } I_1 \perp\!\!\!\perp I_2 \perp\!\!\!\perp \{E'_1, E'_2\} \\
 &= P(E'_1|E'_2)P(E'_2)P(I_1)P(I_2) \\
 P(E_1)P(E_2) &= P(E'_1 \wedge I_1)P(E'_2 \wedge I_2) \\
 &= P(E'_1)P(I_1)P(E'_2)P(I_2) && \text{definitions of } \chi(E_1) \text{ and } \chi(E_2)
 \end{aligned}$$

In case 1.5, given $E'_1 \xrightarrow{\oplus} E'_2$, we can infer that $P(E'_1|E'_2) > P(E'_1)$. Thus, $P(E_1 \wedge E_2) > P(E_1)P(E_2)$.

- Case 1.6 $I_1 \perp\!\!\!\perp I_2$, $I_1 \rightarrow^+ E'_2$, $I_2 \perp\!\!\!\perp E'_1$

PROOF OF CASE 1.6 $E'_1 \xrightarrow{\oplus} E'_2$, $I_1 \perp\!\!\!\perp I_2$, $I_1 \rightarrow^+ E'_2$, $I_2 \perp\!\!\!\perp E'_1$

Given $E'_1 \xrightarrow{\oplus} E'_2$ and $I_1 \rightarrow^+ E'_2$, we split E'_2 into $E'_2(1)$ and $E'_2(2)$, where $E'_1 \xrightarrow{\oplus} E'_2(1)$, $E'_1 \perp\!\!\!\perp E'_2(2)$ and $I_1 \rightarrow^+ E'_2(2)$, $I_1 \perp\!\!\!\perp E'_2(1)$. Meanwhile, given $E'_2 = \bigwedge_j I_j$, E'_2 can be represented as $E'_2 = E'_2(1) \wedge E'_2(2)$.

$$\begin{aligned}
 P(E_1 \wedge E_2) &= P(E'_1 \wedge I_1 \wedge E'_2 \wedge I_2) \\
 &= P(E'_1 \wedge E'_2(1))P(I_1 \wedge E'_2(2))P(I_2) && \text{definitions of } \{E'_1, E'_2(1)\} \perp\!\!\!\perp \{I_1, E'_2(2)\} \perp\!\!\!\perp I_2 \\
 &= P(E'_2(1)|E'_1)P(E'_1)P(E'_2(2)|I_1)P(I_1)P(I_2) \\
 P(E_1)P(E_2) &= P(E'_1 \wedge I_1)P(E'_2 \wedge I_2) \\
 &= P(E'_1)P(I_1)P(E'_2)P(I_2) && \text{definitions of } \chi(E_1) \text{ and } \chi(E_2)
 \end{aligned}$$

In case 1.6, given $E'_1 \xrightarrow{\oplus} E'_2(1)$, we can get $P(E'_2(1)|E'_1) > P(E'_2(1))$. Given the precondition $\not\models E_1 \rightarrow^- E_2$, it implies that $\not\models I_1 \rightarrow^- E'_2$ since $E_1 = E'_1 \wedge I_1$ and $E_2 = E'_2 \wedge I_2$. As $E'_2 = E'_2(1) \wedge E'_2(2)$, it implies that $\not\models I_1 \rightarrow^- E'_2(2)$. Together with $I_1 \rightarrow^+ E'_2(2)$, from Lemma 7, we know that $I_1 \xrightarrow{\oplus} E'_2(2)$. As a result, $P(E'_2(2)|I_1) > P(E'_2(2))$.

Thus, $P(E_1 \wedge E_2) = P(E'_2(1)|E'_1)P(E'_1)P(E'_2(2)|I_1)P(I_1)P(I_2) > P(E'_2(1))P(E'_1)P(E'_2(2))P(I_1)P(I_2) = P(E'_1)P(E'_2)P(I_1)P(I_2) = P(E_1)P(E_2)$.

- Case 1.7 $I_1 \perp\!\!\!\perp I_2$, $I_1 \rightarrow^+ E'_2$, $I_2 \rightarrow^+ E'_1$

PROOF OF CASE 1.7 $E'_1 \xrightarrow{\oplus} E'_2$, $I_1 \perp\!\!\!\perp I_2$, $I_1 \rightarrow^+ E'_2$, $I_2 \rightarrow^+ E'_1$

In this case, we split E'_1 into $E'_1(1)$ and $E'_1(2)$, such that $E'_1(1) \xrightarrow{\oplus} E_2$, $E'_1(2) \rightarrow^+ I_2$, $E'_1(1) \perp\!\!\!\perp I_2$, $E'_1(2) \perp\!\!\!\perp E_2$. Similarly, we split E'_2 into $E'_2(1)$ and $E'_2(2)$, such that $E'_2(1) \xrightarrow{\oplus} E_1$, $E'_2(2) \rightarrow^+ I_1$, $E'_2(1) \perp\!\!\!\perp I_1$, $E'_2(2) \perp\!\!\!\perp E_1$.

$$\begin{aligned}
 P(E_1 \wedge E_2) &= P(E'_1 \wedge I_1 \wedge E'_2 \wedge I_2) \\
 &= P(E'_1(1)E'_2(1))P(I_1 \wedge E'_2(2))P(I_2 \wedge E'_1(2)) \\
 &\quad \text{definitions of } \{E'_1(1), E'_2(1)\} \perp\!\!\!\perp \{I_1, E'_2(2)\} \perp\!\!\!\perp \{I_2, E'_1(2)\} \\
 &= P(E'_2(1)|E'_1(1))P(E'_1(1))P(E'_2(2)|I_1)P(I_1)P(E'_1(2)|I_2)P(I_2) \\
 P(E_1)P(E_2) &= P(E'_1 \wedge I_1)P(E'_2 \wedge I_2) \\
 &= P(E'_1)P(I_1)P(E'_2)P(I_2) \\
 &\quad \text{definitions of } \chi(E_1) \text{ and } \chi(E_2)
 \end{aligned}$$

Given $E'_1(2) \rightarrow^+ I_2$, and $\not\models E'_1(2) \rightarrow^- I_2$, according to Lemma 7, we can get $E'_1(2) \xrightarrow{\oplus} I_2$. Similarly, we can get $E'_2(2) \xrightarrow{\oplus} I_1$. With this, we have $P(E'_1(2)|I_2) > P(E'_1(2))$ and $P(E'_2(2)|I_1) > P(E'_2(2))$. As a result, $P(E_1 \wedge E_2) = P(E'_2(1)|E'_1(1))P(E'_1(1))P(E'_2(2)|I_1)P(I_1)P(E'_1(2)|I_2)P(I_2) > P(E'_2(1))P(E'_1(1))P(E'_2(2))P(E'_1(2))P(I_1)P(I_2) = P(E'_1)P(E'_2)P(I_1)P(I_2) = P(E_1)P(E_2)$. We then prove $E_1 \xrightarrow{\oplus} E_2$.

- Case 1.8 $I_1 \perp\!\!\!\perp I_2$, $I_1 \perp\!\!\!\perp E'_2$, $I_2 \rightarrow^+ E'_1$

This proof of case 1.8 is symmetric to proof of case 1.6, so we omit it here.

- We do not talk about \rightarrow^- relation between a pair of expressions as it violates the precondition $\not\models E_1 \rightarrow^- E_2$.

Case 2: $E'_1 \perp\!\!\!\perp E'_2$. We also know that $E'_1 \perp\!\!\!\perp I_1$ and $E'_2 \perp\!\!\!\perp I_2$.

- Case 2.1 $I_1 \rightarrow^+ I_2$, $I_1 \perp\!\!\!\perp E'_2$, $I_2 \perp\!\!\!\perp E'_1$

This proof of case 2.1 is symmetric to proof of case 1.5, so we omit it here.

- Case 2.2 $I_1 \rightarrow^+ I_2$, $I_1 \perp\!\!\!\perp E'_2$, $I_2 \rightarrow^+ E'_1$. It is invalid as it violates the preconditions $\chi(E_1)$.
- Case 2.3 $I_1 \rightarrow^+ I_2$, $I_1 \rightarrow^+ E'_2$, $I_2 \perp\!\!\!\perp E'_1$. It is invalid as it violates the preconditions $\chi(E_2)$.
- Case 2.4 $I_1 \rightarrow^+ I_2$, $I_1 \rightarrow^+ E'_2$, $I_2 \rightarrow^+ E'_1$. It is invalid as it violates the preconditions $\chi(E_1)$ and $\chi(E_2)$.
- Case 2.5 $I_1 \perp\!\!\!\perp I_2$, $I_1 \perp\!\!\!\perp E'_2$, $I_2 \perp\!\!\!\perp E'_1$. It is invalid as it violates the precondition $\vdash E_1 \rightarrow^+ E_2$.
- Case 2.6 $I_1 \perp\!\!\!\perp I_2$, $I_1 \rightarrow^+ E'_2$, $I_2 \perp\!\!\!\perp E'_1$.

This proof of case 2.6 is symmetric to proof of case 1.5, so we omit it here.

- Case 2.7 $I_1 \perp\!\!\!\perp I_2$, $I_1 \perp\!\!\!\perp E'_2$, $I_2 \rightarrow^+ E'_1$.

This proof of case 2.7 is symmetric to proof of case 1.5, so we omit it here.

- Case 2.8 $I_1 \perp\!\!\!\perp I_2$, $I_1 \rightarrow^+ E'_2$, $I_2 \rightarrow^+ E'_1$.

This proof of case 2.8 is symmetric to proof of case 1.1, so we omit it here.

Case 3: $E'_1 \xrightarrow{\oplus} E'_2$ is invalid as it violates the precondition $\vdash E_1 \rightarrow^- E_2$.

□

LEMMA 9. *Let E_1 be a conjunction over input facts (i.e., $E_1 = \bigwedge_i I_i$) and let $E_2 = \bigvee \bigwedge_j I_j$ be a DNF formula over input facts. If $\vdash \chi(E_1)$, $\vdash \chi(E_2)$, $\vdash E_1 \rightarrow^+ E_2$ and $\not\models E_1 \rightarrow^- E_2$, then $P(E_1 \wedge E_2) > P(E_1)P(E_2)$.*

PROOF of Lemma 9. We prove this by induction and start with the base cases.

Base case: $E_1 = \bigwedge_i I_i$, $E_2 = \bigwedge_j I_j$

Given $\chi(E_1)$, $\chi(E_2)$, $\vdash E_1 \rightarrow^+ E_2$, and $\not\models E_1 \rightarrow^- E_2$, according to Lemma 8, we can get $E_1 \xrightarrow{\oplus} E_2$.

Inductive case:

$$\begin{array}{lcl}
 E_1 & = & E'_1 \\
 E_2 & = & E'_2 \vee C \quad \left| \quad \begin{array}{l} E'_1 = \bigwedge_i I_i \\ E'_2 = \bigvee \bigwedge_j I_j \end{array} \right| \quad C = \bigwedge_{j'} I_{j'}
 \end{array}$$

Case 1 $E'_1 \xrightarrow{\oplus} E'_2$

- Case 1.1 $E'_1 \rightarrow^+ C$

Proof of Case 1.1 $E'_1 \xrightarrow{\oplus} E'_2, E'_1 \rightarrow^+ C$

Given $\neg E_1 \rightarrow^- E_2$, we can infer that $\neg E'_1 \rightarrow^- C$. Given $\chi(E_1)$ and $\chi(E_2)$, we can also infer that $\chi(E'_1)$ and $\chi(C)$ hold.

Given $E'_1 \rightarrow^+ C, \neg E'_1 \rightarrow^- C, \chi(E'_1)$ and $\chi(C)$, we get $E'_1 \xrightarrow{\oplus} C$ by applying Lemma 8.

$$\begin{aligned}
 & \frac{P(E_1)P(E_2)}{P(E'_1)P(E'_2) + P(E'_1)P(C) - P(E'_1)P(E'_2)P(C)} \\
 &= \frac{P(E_1 \wedge E_2)}{P(E'_1 \wedge E'_2 \vee E'_1 \wedge C)} \\
 &= \frac{P(E'_1 \wedge E'_2) + P(E'_1 \wedge C) - P(E'_1 \wedge E'_2 \wedge C)}{P(E'_1|E'_2)P(E'_2) + P(E'_1|C)P(C) - P(E'_1(1) \wedge E'_2)P(E'_1(2) \wedge C)} \quad \text{splitting } E'_1 : P(E'_1) = P(E'_1(1))P(E'_1(2)) \\
 &= \frac{P(E'_1|E'_2)P(E'_2) + P(E'_1|C)P(C) - P(E'_1(1)|E'_2)P(E'_1(2)|C)P(E'_2)P(C)}{P(E'_1|E'_2)P(E'_2) + P(E'_1|C)P(C) - P(E'_1(1)|E'_2)P(E'_1(2)|C)P(E'_2)P(C)}
 \end{aligned}$$

Given $E'_1 \xrightarrow{\oplus} E'_2$, we know that $P(E'_1|E'_2) > P(E'_1)$. Similarly, given $E'_1 \xrightarrow{\oplus} C$, we also have $P(E'_1|C) > P(E'_1)$.

We define a function $f(x, y)$ as follows ($x \in [0, 1], y \in [0, 1]$):

$$f(\alpha, \beta) = \alpha * P(E'_2) + \beta * P(C) - \alpha\beta P(E'_2)P(C)/P(E'_1)$$

We then define the expression of the joint probability as follows: $P(E_1 \wedge E_2) = f(x_1, y_1)$ where $\alpha_1 = P(E'_1|E'_2)$ and $\beta_1 = P(E'_1|C)$. Meanwhile $P(E_1)P(E_2) = f(x_2, y_2)$ where $\alpha_2 = P(E'_1)$, $\beta_2 = P(E'_1)$. To summarize, $\alpha \in \{P(E'_1|E'_2), P(E'_1)\}$, $\alpha \geq P(E'_1)$ and $\beta \geq P(E'_1)$.

The partial derivation $f_\alpha(\alpha, \beta) = \frac{\partial f}{\partial \alpha} = P(E'_2) - \beta P(E'_2)P(C)/P(E'_1) > 0$. Similarly, $f_\beta(\alpha, \beta) = \frac{\partial f}{\partial \beta} = P(C) - \alpha P(E'_2)P(C)/P(E'_1) > 0$. With both $f_\alpha(\alpha, \beta)$ and $f_\beta(\alpha, \beta)$ being positive, according to the Lemma 4, we can get the $f(\alpha_1, \beta_1) > f(\alpha_2, \beta_2)$ if $\alpha_1 > \alpha_2, \beta_1 > \beta_2$. Given $\alpha_1 > \alpha_2$ and $\beta_1 > \beta_2$, we can get $P(E_1 \wedge E_2) = f(\alpha_1, \beta_1) > f(\alpha_2, \beta_2) = P(E_1)P(E_2)$. Therefore, we are able to prove $E_1 \xrightarrow{\oplus} E_2$.

- Case 1.2 $E'_1 \perp\!\!\!\perp C$

Here we omit the proof of case 1.2 since it is similar to the proof of Case 1.1.

Case 2 $E'_1 \perp\!\!\!\perp E'_2$

- Case 2.1 $E'_1 \rightarrow^+ C$

Here we omit the proof of case 2.1 since it is similar to the proof of Case 1.1.

- Case 12.2 $E'_1 \perp\!\!\!\perp C$

Here we omit the proof of case 2.2 since it is similar to the proof of Case 1.1.

Case 3 $E'_1 \xrightarrow{\oplus} E'_2$: This case is invalid as it violates the precondition $\neg E_1 \rightarrow^- E_2$.

□

LEMMA 10 (SOUNDNESS OF RULE POS). *The rule Pos in Figure 14 are sound. In particular, given a pair of DNF formulas $E_1 = \bigvee \bigwedge_i I_i, E_2 = \bigvee \bigwedge_j I_j$, if $\vdash \chi(E_1), \vdash \chi(E_2), \vdash E_1 \rightarrow^+ E_2$ and $\neg E_1 \rightarrow^- E_2$, then $P(E_1 \wedge E_2) > P(E_1)P(E_2)$.*

Proof of Lemma 10.

In Lemma 10, both E_1 and E_2 are arbitrary boolean formulas so we use DNF format to represent both of them, i.e., $E_1 = \bigvee \bigwedge_i I_i$ and $E_2 = \bigvee \bigwedge_j I_j$.

Base case: $E_1 = \bigwedge_i I_i, E_2 = \bigwedge_j I_j$

Given $\chi(E_1), \chi(E_2), \vdash E_1 \rightarrow^+ E_2$, and $\neg E_1 \rightarrow^- E_2$, according to Lemma 8, we can get $E_1 \xrightarrow{\oplus} E_2$.

Inductive case:

$$\begin{array}{lcl} E_1 & = & E'_1 \vee C_1 \mid E'_1 = \bigvee \bigwedge_i I_i \mid C_1 = \bigwedge_{i'} I_{i'} \\ E_2 & = & E'_2 \vee C_2 \mid E'_2 = \bigvee \bigwedge_j I_j \mid C_2 = \bigwedge_{j'} I_{j'} \end{array}$$

Case 1: $E'_1 \xrightarrow{\oplus} E'_2$

Based on the precondition $\chi(E_1)$ (resp. $\chi(E_2)$), we know that $E'_1 \perp\!\!\!\perp C_1$ (resp. $E'_2 \perp\!\!\!\perp C_2$). Following are the valid cases.

- Case 1.1 $C_1 \rightarrow^+ C_2$, $C_1 \perp\!\!\!\perp E'_2$, $C_2 \perp\!\!\!\perp E'_1$

Proof of Case 1.1 $E'_1 \xrightarrow{\oplus} E'_2$, $C_1 \rightarrow^+ C_2$, $C_1 \perp\!\!\!\perp E'_2$, $C_2 \perp\!\!\!\perp E'_1$

$$\begin{aligned} P(E_1 \wedge E_2) &= P((E'_1 \vee C_1) \wedge (E'_2 \vee C_2)) \\ &= P(E'_1 \wedge E'_2 \vee E'_1 \wedge C_2 \vee C_1 \wedge E'_2 \vee C_1 \wedge C_2) \\ &= P(E'_1 \wedge E'_2) + P(E'_1 \wedge C_2) + P(C_1 \wedge E'_2) + P(C_1 \wedge C_2) \\ &\quad - P(E'_1 E'_2 C_2) - P(E'_1 C_1 E'_2) - P(E'_1 C_1 C_2) - P(C_1 E'_2 C_2) + P(E'_1 E'_2 C_1 C_2) \\ &= P(E'_1 | E'_2) \color{red}{P(E'_2)} + \color{blue}{P(E'_1)P(C_2)} + \color{blue}{P(C_1)P(E'_2)} + P(C_1 | C_2) \color{red}{P(C_2)} \\ &\quad - P(E'_1 | E'_2) \color{red}{P(E'_2)P(C_2)} - P(E'_1 | E'_2) \color{red}{P(C_1)P(E'_2)} - P(C_1 | C_2) \color{red}{P(C_2)P(E'_1)} - P(C_1 | C_2) \color{red}{P(E'_2)P(C_2)} \\ &\quad + P(E'_1 | E'_2) P(C_1 | C_2) \color{red}{P(E'_2)P(C_2)} \quad \text{using } C_1 \perp\!\!\!\perp E'_2, C_2 \perp\!\!\!\perp E'_1 \\ P(E_1)P(E_2) &= P(E'_1 \vee C_1)P(E'_2 \vee C_2) \\ &= (P(E'_1) + P(C_1) - P(E'_1 C_1))(P(E'_2) + P(C_2) - P(E'_2 C_2)) \\ &= P(E'_1)P(E'_2) + P(E'_1)P(C_2) + P(C_1)P(E'_2) + P(C_1)P(C_2) \\ &\quad - P(E'_1)P(E'_2 C_2) - P(E'_1 C_1)P(E'_2) - P(E'_1 C_1)P(C_2) - P(C_1)P(E'_2 C_2) + P(E'_1 C_1)P(E'_2 C_2) \\ &= P(E'_1) \color{red}{P(E'_2)} + \color{blue}{P(E'_1)P(C_2)} + \color{blue}{P(C_1)P(E'_2)} + P(C_1) \color{red}{P(C_2)} \\ &\quad - P(E'_1) \color{red}{P(E'_2)P(C_2)} - P(E'_1) \color{red}{P(C_1)P(E'_2)} - P(C_1) \color{red}{P(C_2)P(E'_1)} - P(C_1) \color{red}{P(E'_2)P(C_2)} \\ &\quad + P(E'_1)P(C_1) \color{red}{P(E'_2)P(C_2)} \quad \text{using } C_1 \perp\!\!\!\perp E'_2, C_2 \perp\!\!\!\perp E'_1 \end{aligned}$$

As we can see from above, both $P(E_1 \wedge E_2)$ and $P(E_1)P(E_2)$ share some terms. Here, we use the same color to denote the same terms. The different parts between $P(E_1 \wedge E_2)$ and $P(E_1)P(E_2)$ are still colored in black. Based on this observation, we can define a function $f(\alpha, \beta)$ as follows:

$$\begin{aligned} f(\alpha, \beta) &= \alpha \color{red}{P(E'_2)} + \color{blue}{P(E'_1)P(C_2)} + \color{blue}{P(C_1)P(E'_2)} + \beta \color{red}{P(C_2)} \\ &\quad - \alpha \color{red}{P(E'_2)P(C_2)} - \alpha \color{red}{P(C_1)P(E'_2)} - \beta \color{red}{P(C_2)P(E'_1)} - \beta \color{red}{P(E'_2)P(C_2)} \\ &\quad + \alpha \beta \color{red}{P(E'_2)P(C_2)} \end{aligned}$$

Based on the definition of $f(\alpha, \beta)$, we can infer that $P(E_1 \wedge E_2) = f(\alpha_1, \beta_1)$, $\alpha_1 = P(E'_1 | E'_2)$, $\beta_1 = P(C_1 | C_2)$ whereas $P(E_1)P(E_2) = f(\alpha_2, \beta_2)$, $\alpha_2 = P(E'_1)$, $\beta_2 = P(C_1)$. Hence $\alpha \in \{P(E'_1 | E'_2), P(E_1)\}$, $\beta \in \{P(C_1 | C_2), P(C_1)\}$.

Given $E'_1 \xrightarrow{\oplus} E'_2$, we know that $P(E'_1 | E'_2) > P(E'_1)$, and thus $\alpha \geq P(E'_1)$. Based on the precondition $\not\models E_1 \rightarrow^- E_2$, we can also imply that $\not\models C_1 \rightarrow^- C_2$ since $E_1 = E'_1 \vee C_1$, $E_2 = E'_2 \vee C_2$, and it does not introduce new negations here. In this case 1.1, we know that $C_1 \rightarrow^+ C_2$, and $\not\models C_1 \rightarrow^- C_2$, and both C_1, C_2 is a conjunction of input facts, i.e., $C_1 = \bigwedge_{i'} I_{i'}$. By applying Lemma 8, we can infer that $C_1 \xrightarrow{\oplus} C_2$, and thus $P(C_1 | C_2) > P(C_1)$. As a result, $\beta \geq P(C_1)$.

Compute the partial derivatives of f with respect to α and β :

$$\begin{aligned} f_\alpha(\alpha, \beta) &= \frac{\partial f}{\partial \alpha} = P(E'_2) - P(E'_2)P(C_2) - P(E'_2)P(C_1) + \beta * P(E'_2)P(C_2) \\ f_\beta(\alpha, \beta) &= \frac{\partial f}{\partial \beta} = P(C_2) - P(C_2)P(E'_1) - P(E'_2)P(C_2) + \alpha * P(E'_2)P(C_2) \end{aligned}$$

Prove that the partial derivatives $f_\alpha(\alpha, \beta)$ and $f_\beta(\alpha, \beta)$ are positive for all α and β :

$$\begin{aligned} f_\alpha(\alpha, \beta) &= P(E'_2) - P(E'_2)P(C_2) - P(E'_2)P(C_1) + \beta * P(E'_2)P(C_2) \\ &\geq P(E'_2) - P(E'_2)P(C_2) - P(E'_2)P(C_1) + P(C_1) * P(E'_2)P(C_2) \quad \text{using } \beta \geq P(C_1) \\ &= (1 - P(C_2))(P(E'_2) - P(E'_2)P(C_1)) \quad \text{using } P(C_1), P(C_2) \in [0, 1] \\ &> 0 \end{aligned}$$

$$\begin{aligned}
f_\beta(\alpha, \beta) &= P(C_2) - P(C_2)P(E'_1) - P(E'_2)P(C_2) + \alpha * P(E'_2)P(C_2) \\
&\geq P(C_2) - P(C_2)P(E'_1) - P(E'_2)P(C_2) + P(E'_1) * P(E'_2)P(C_2) \quad \text{using } \alpha \geq P(E'_1) \\
&= (1 - P(E'_1))(P(C_2) - P(C_2)P(E'_2)) \quad \text{using } P(E'_1), P(E'_2) \in [0, 1] \\
&> 0
\end{aligned}$$

As we can get from the above, both $\frac{\partial f}{\partial \alpha}$ and $\frac{\partial f}{\partial \beta}$ are positive. According to Lemma 4, we know that $f(\alpha_1, \alpha_2) > f(\beta_1, \beta_2)$ if $\alpha_1 > \alpha_2$ and $\beta_1 > \beta_2$.

As we know, $P(E_1 \wedge E_2) = f(\alpha_1, \beta_1)$ while $P(E_1)P(E_2) = f(\alpha_2, \beta_2)$. Given $\alpha_1 = P(E'_1|E'_2) > P(E'_1) = \alpha_2$ and $\beta_1 = P(C_1|C_2) > P(C_1) = \beta_2$, we can get $f(\alpha_1, \beta_1) > f(\alpha_2, \beta_2)$, and thus $P(E_1 \wedge E_2) > P(E_1)P(E_2)$.

- Case 1.2 $C_1 \perp\!\!\!\perp C_2$, $C_1 \rightarrow^+ E_2$, $C_2 \rightarrow^+ E_1$

Proof of Case 1.2 $E'_1 \xrightarrow{\oplus} E'_2$, $C_1 \perp\!\!\!\perp C_2$, $C_1 \rightarrow^+ E'_2$, $C_2 \rightarrow^+ E'_1$

Given $E'_1 \xrightarrow{\oplus} E'_2$, $C_1 \rightarrow^+ E'_2$, and $C_2 \rightarrow^+ E'_1$, we split E'_1 (resp. E'_2) into $E'_1(1)$ and $E'_1(2)$ (resp. $E'_2(1)$ and $E'_2(2)$), such that $E'_1(1) \xrightarrow{\oplus} E'_2(1)$, $E'_1(1) \perp\!\!\!\perp C_2$, $E'_1(2) \rightarrow^+ C_2$, $E'_2(1) \perp\!\!\!\perp C_1$, $E'_2(2) \rightarrow^+ C_1$. There are two cases for the splitting, taking E'_1 for example, $E'_1 = E'_1(1) \wedge E'_1(2)$ or $E'_1 = E'_1(1) \vee E'_1(2)$. We take the first form (i.e., $E'_1 = E'_1(1) \wedge E'_1(2)$) to prove, and the proof of the second form is similar to that, so we focus on proving the first form and omit the proof of the second form here.

After splitting, $E'_1(1) \xrightarrow{\oplus} E'_2(1)$, $E'_1(2) \rightarrow^+ C_2$, $E'_2(2) \rightarrow^+ C_1$. We also know that $\nu E'_1(2) \rightarrow^- C_2$, according to Lemma 9, we can get $E'_1(2) \xrightarrow{\oplus} C_2$. Similarly $E'_2(2) \xrightarrow{\oplus} C_1$.

From the proof of Case 1.1, we get the expression of $P(E_1 \wedge E_2)$ and $P(E_1)P(E_2)$ as follows:

$$\begin{aligned}
P(E_1 \wedge E_2) &= P((E'_1 \vee C_1) \wedge (E'_2 \vee C_2)) \\
&= P(E'_1 \wedge E'_2 \vee E'_1 \wedge C_2 \vee C_1 \wedge E'_2 \vee C_1 \wedge C_2) \\
&= P(E'_1 \wedge E'_2) + \overbrace{P(E'_1 \wedge C_2) + P(C_1 \wedge E'_2)}^{\text{blue}} + P(C_1 \wedge C_2) \\
&\quad \text{blue } -P(E'_1 E'_2 C_2) - P(E'_1 C_1 E'_2) - P(E'_1 C_1 C_2) - P(C_1 E'_2 C_2) + P(E'_1 E'_2 C_1 C_2) \\
P(E_1)P(E_2) &= P(E'_1 \vee C_1)P(E'_2 \vee C_2) \\
&= (P(E'_1) + P(C_1) - P(E'_1 C_1))(P(E'_2) + P(C_2) - P(E'_2 C_2)) \\
&= P(E'_1)P(E'_2) + \overbrace{P(E'_1)P(C_2) + P(C_1)P(E'_2)}^{\text{blue}} + P(C_1)P(C_2) \\
&\quad \text{blue } -P(E'_1)P(E'_2)P(C_2) - P(E'_1)P(C_1)P(E'_2) - P(E'_1)P(C_1)P(C_2) - P(C_1)P(E'_2)P(C_2) \\
&\quad \text{blue } +P(E'_1)P(C_1)P(E'_2)P(C_2)
\end{aligned}$$

We now define a new function $f(\alpha, \beta)$ to represent the blue colored formulas in $P(E_1 \wedge E_2)$ and $P(E_1)P(E_2)$.

$$\begin{aligned}
f(\alpha, \beta) &= \alpha P(C_2) + \beta P(C_1) - \alpha \frac{P(E'_1(1)|E'_2)}{P(E'_1(1))} P(E'_2)P(C_2) - \beta \frac{P(E'_2(1)|E'_1)}{P(E'_2(1))} P(E'_1)P(C_1) \\
&\quad - \alpha P(C_1)P(C_2) - \beta P(C_1)P(C_2) + \alpha \beta \frac{P(E'_1(1)E'_2(1))}{P(E'_1(1))P(E'_2(1))} P(I_1)P(I_2)
\end{aligned}$$

Joint probability $P(E_1 \wedge E_2)$ is equivalent to $P(E'_1 \wedge E'_2) + P(C_1 \wedge C_2) + f(\alpha_1, \beta_1)$ while the formula $P(E_1)P(E_2)$ is defined as $P(E'_1)P(E'_2) + P(C_1)P(C_2) + f(\alpha_2, \beta_2)$, where $\alpha_1, \alpha_2, \beta_1, \beta_2$ are defined as follows:

$$\begin{array}{lcl}
\alpha_1 & = & P(E'_1|C_2) \\
\alpha_2 & = & P(E'_1) \\
\beta_1 & = & P(E'_2|C_1) \\
\beta_2 & = & P(E'_2)
\end{array}$$

After splitting, we know that $E'_1(1) \xrightarrow{\oplus} E'_2(1)$, $E'_1(1) \perp\!\!\!\perp C_2$, $E'_1(2) \xrightarrow{\oplus} C_2$, $E'_2(1) \perp\!\!\!\perp C_1$, $E'_2(2) \xrightarrow{\oplus} C_1$. Hence, we can infer that $P(E'_1 \wedge C_2) = P(E'_1(1) \wedge E'_1(2) \wedge C_2) = P(E'_1(1))P(E'_1(2)C_2) > P(E'_1(1))P(E'_1(2))P(C_2) = P(E'_1)P(C_2)$. Thus, $E'_1 \xrightarrow{\oplus} C_2$. Similarly, we can get $E'_2 \xrightarrow{\oplus} C_1$. With this, we can get $\alpha_1 > \alpha_2$ and $\beta_1 > \beta_2$. In addition, $\alpha \geq P(E'_1)$ and $\beta \geq P(E'_2)$.

Similarly, we compute the partial derivative of $f(\alpha, \beta)$ as follows:

$$\begin{aligned}
 \frac{\partial f}{\partial \alpha} &= P(C_2) - \frac{P(E'_1(1)|E'_2)}{P(E'_1(1))} P(E'_2)P(C_2) - P(C_1)P(C_2) + \beta P(C_1)P(C_2) \frac{P(E'_1(1)E'_2(1))}{P(E'_1(1))P(E'_2(1))} \\
 &\geq P(C_2) - \frac{P(E'_1(1)|E'_2)}{P(E'_1(1))} P(E'_2)P(C_2) - P(C_1)P(C_2) + P(E'_2)P(C_1)P(C_2) \frac{P(E'_1(1)E'_2(1))}{P(E'_1(1))P(E'_2(1))} \\
 &= (1 - P(C_1)P(C_2))(1 - \frac{P(E'_2)}{P(E'_1(1))} \frac{P(E'_1(1)E'_2(1))}{P(E'_2(1))}) \\
 &= (1 - P(C_1)P(C_2))(1 - \frac{P(E'_2(2))P(E'_2(1)|E'_1(1))}{P(E'_1(1))}) \\
 &> 0
 \end{aligned}$$

Similarly, we can also prove that the partial derivative $\frac{\partial f}{\partial \beta}$ is positive. Given $\frac{\partial f}{\partial \alpha} > 0$ and $\frac{\partial f}{\partial \beta} > 0$, based on Lemma 4, we can get $f(\alpha_1, \beta_1) > f(\alpha_2, \beta_2)$ if $\alpha_1 > \alpha_2$ and $\beta_1 > \beta_2$.

In addition, we also have $E'_1 \xrightarrow{\oplus} E'_2$, $C_1 \perp\!\!\!\perp C_2$, so $P(E'_1 \wedge E'_2) > P(E'_1)P(E'_2)$ and $P(C_1 \wedge C_2) = P(C_1)P(C_2)$. Thus we are able to prove $P(E_1 \wedge E_2) = P(E'_1 \wedge E'_2) + P(C_1 \wedge C_2) + f(\alpha_1, \beta_1) > P(E'_1)P(E'_2) + P(C_1)P(C_2) + f(\alpha_2, \beta_2) = P(E_1)P(E_2)$. Thus, we can get $E_1 \xrightarrow{\oplus} E_2$.

- Case 1.3 $C_1 \perp\!\!\!\perp C_2$, $C_1 \rightarrow^+ E_2$, $C_2 \perp\!\!\!\perp E_1$
- Case 1.4 $C_1 \perp\!\!\!\perp C_2$, $C_1 \perp\!\!\!\perp E_2$, $C_2 \rightarrow^+ E_1$
- Case 1.5 $C_1 \perp\!\!\!\perp C_2$, $C_1 \perp\!\!\!\perp E_2$, $C_2 \perp\!\!\!\perp E_1$

Case 2: $E'_1 \perp\!\!\!\perp E'_2$

- Case 2.1 $C_1 \rightarrow^+ C_2$, $E'_1 \perp\!\!\!\perp C_2$, $E'_2 \perp\!\!\!\perp C_1$
- Case 2.2 $C_1 \perp\!\!\!\perp C_2$, $E'_1 \rightarrow^+ C_2$, $E'_2 \perp\!\!\!\perp C_1$
- Case 2.3 $C_1 \perp\!\!\!\perp C_2$, $E'_1 \perp\!\!\!\perp C_2$, $E'_2 \rightarrow^+ C_1$
- Case 2.4 $C_1 \perp\!\!\!\perp C_2$, $E'_1 \rightarrow^+ C_2$, $E'_2 \rightarrow^+ C_1$

Case 3: $E'_1 \xrightarrow{\ominus} E'_2$ This case is invalid as it violates the precondition $\not\models E_1 \rightarrow^- E_2$.

Proof of other cases

Here we omit the proof of other valid cases (1.3, 1.4, 1.5, 2.1, 2.2, 2.3, and 2.4) since it is similar to the proof of Case 1.1 and Case 1.2. \square

Assumptions

For the above proof of Lemmas 7, 8, 9 and 10, we made two assumptions. First, we assume that input I is a term, instead of the literal or input fact. For instance, $I \in \{a, \neg a\}$ where $a \in \text{InputFacts}(D)$. Second, given $I_a \rightarrow^+ I_b$, we imply $P(I_a I_b) > P(I_a)P(I_b)$.

Now we give a proof about the second assumption: if $I_a \rightarrow^+ I_b$, $I_a \in \{a, \neg a\}$, $I_b \in \{b, \neg b\}$, $a \in \text{InputFacts}(D)$, and $b \in \text{InputFacts}(D)$, then $P(I_a I_b) > P(I_a)P(I_b)$.

Proof.

Based on the rules MAY-1 and MAY-2, there are four cases for I_a and I_b that can infer $I_a \rightarrow^+ I_b$.

- (1) Case 1: $I_a = a$, $I_b = b$, $a \blacktriangleright^+ b$.
 $P(I_a \wedge I_b) = P(a \wedge b) > P(a)P(b)$ according to $a \blacktriangleright^+ b$.
- (2) Case 2: $I_a = a$, $I_b = \neg b$, $a \blacktriangleright^- b$.
 $P(I_a \wedge I_b) = P(a \wedge \neg b) = P(\neg b|a)P(a) = (1 - P(b|a))P(a)$. Given $a \blacktriangleright^- b$, we have $P(b|a) < P(b)$, and thus $(1 - P(b|a)) > (1 - P(b))$. Hence $P(I_a \wedge I_b) = (1 - P(b|a))P(a) > (1 - P(b))P(a) = P(I_b)P(I_a)$.
- (3) Case 3: $I_a = \neg a$, $I_b = b$, $a \blacktriangleright^- b$.
 $P(I_a \wedge I_b) = P(\neg a \wedge b) = P(\neg a|b)P(b) = (1 - P(a|b))P(b)$. Given $a \blacktriangleright^- b$, we have $P(a|b) < P(a)$, and thus $(1 - P(a|b)) > (1 - P(a))$. Hence $P(I_a \wedge I_b) = (1 - P(a|b))P(b) > (1 - P(a))P(b) = P(I_a)P(I_b)$.

(4) Case 4: $I_a = \neg a$, $I_b = \neg b$, $a \blacktriangleright^+ b$.

$$\begin{aligned}
 P(I_a \wedge I_b) &= P(\neg a \wedge \neg b) = 1 - P(a \vee b) &= 1 - P(a) - P(b) + P(a \wedge b) \\
 &> 1 - P(a) - P(b) + P(a)P(b) && \text{using } a \blacktriangleright^+ b \\
 &= (1 - P(a))(1 - P(b)) = P(I_a)P(I_b)
 \end{aligned}$$

□

LEMMA 11. Given an input fact $E_1 = I$ and a general DNF formula $E_2 = \bigvee \bigwedge_j I_j$, if $\vdash \chi(E_1)$, $\vdash \chi(E_2)$, $\not\vdash E_1 \rightarrow^+ E_2$ and $\vdash E_1 \rightarrow^- E_2$, then $P(E_1 \wedge E_2) < P(E_1)P(E_2)$, i.e., $E_1 \xrightarrow{\ominus} E_2$.

LEMMA 12. Let E_1, E_2 be conjunctions over input facts, i.e., $E_1 = \bigwedge_i I_i$ and $E_2 = \bigwedge_j I_j$. If $\vdash \chi(E_1)$, $\vdash \chi(E_2)$, $\not\vdash E_1 \rightarrow^+ E_2$ and $\vdash E_1 \rightarrow^- E_2$, then $P(E_1 \wedge E_2) < P(E_1)P(E_2)$.

LEMMA 13. Let E_1 be a conjunction over input facts (i.e., $E_1 = \bigwedge_i I_i$) and let $E_2 = \bigvee \bigwedge_j I_j$ be a DNF formula over input facts. If $\vdash \chi(E_1)$, $\vdash \chi(E_2)$, $\not\vdash E_1 \rightarrow^+ E_2$ and $\vdash E_1 \rightarrow^- E_2$, then $P(E_1 \wedge E_2) < P(E_1)P(E_2)$.

LEMMA 14 (**SOUNDNESS OF RULE NEG**). The rule *NEGATIVE* in Figure 14 are sound. In particular, given a pair of relations $E_1 = \bigvee \bigwedge_i I_i$, $E_2 = \bigvee \bigwedge_j I_j$, if $\vdash \chi(E_1)$, $\vdash \chi(E_2)$, $\not\vdash E_1 \rightarrow^+ E_2$ and $\vdash E_1 \rightarrow^- E_2$, then $P(E_1 \wedge E_2) < P(E_1)P(E_2)$.

Proof of Lemma 14.

To prove negative correlation $E_1 \xrightarrow{\ominus} E_2$, we also make two assumptions as we did for Lemma 10. First, we assume that the input I is a term, instead of a literal or input fact. For instance, $I \in \{a, \neg a\}$ where $a \in \text{InputFacts}(D)$. Second, given $I_a \rightarrow^- I_b$, we imply $P(I_a I_b) < P(I_a)P(I_b)$.

Now we give a proof about the second assumption: if $I_a \rightarrow^- I_b$, $I_a \in \{a, \neg a\}$, $I_b \in \{b, \neg b\}$, $a \in \text{InputFacts}(D)$, and $b \in \text{InputFacts}(D)$, then $P(I_a I_b) < P(I_a)P(I_b)$.

Proof.

Based on the rules MAY-1 and MAY-2, there are four cases for I_a and I_b that can infer $I_a \rightarrow^- I_b$.

(1) Case 1: $I_a = a$, $I_b = b$, $a \blacktriangleright^- b$.

$$P(I_a \wedge I_b) = P(a \wedge b) < P(a)P(b) \text{ according to } a \blacktriangleright^- b.$$

(2) Case 2: $I_a = a$, $I_b = \neg b$, $a \blacktriangleright^+ b$.

$$\begin{aligned}
 P(I_a \wedge I_b) &= P(a \wedge \neg b) = P(\neg b|a)P(a) = (1 - P(b|a))P(a). \text{ Given } a \blacktriangleright^+ b, \text{ we have } P(b|a) > P(b), \\
 &\text{and thus } (1 - P(b|a)) < (1 - P(b)). \text{ Hence } P(I_a \wedge I_b) = (1 - P(b|a))P(a) < (1 - P(b))P(a) = P(I_b)P(I_a)
 \end{aligned}$$

(3) Case 3: $I_a = \neg a$, $I_b = b$, $a \blacktriangleright^+ b$.

$$\begin{aligned}
 P(I_a \wedge I_b) &= P(\neg a \wedge b) = P(a|b)P(b) = (1 - P(a|b))P(b). \text{ Given } a \blacktriangleright^+ b, \text{ we have } P(a|b) > P(a), \\
 &\text{and thus } (1 - P(a|b)) < (1 - P(a)). \text{ Hence } P(I_a \wedge I_b) = (1 - P(a|b))P(b) < (1 - P(a))P(b) = P(I_a)P(I_b)
 \end{aligned}$$

(4) Case 4: $I_a = \neg a$, $I_b = \neg b$, $a \blacktriangleright^- b$.

$$\begin{aligned}
 P(I_a \wedge I_b) &= P(\neg a \wedge \neg b) = 1 - P(a \vee b) &= 1 - P(a) - P(b) + P(a \wedge b) \\
 &< 1 - P(a) - P(b) + P(a)P(b) && \text{using } a \blacktriangleright^- b \\
 &= (1 - P(a))(1 - P(b)) = P(I_a)P(I_b)
 \end{aligned}$$

To prove Lemma 14 for Rule *NEG* in Figure 14, we use the following Lemmas: Lemma 4, Lemma 11, Lemma 12 and Lemma 13. The proof of Lemma 14 ($E_1 \xrightarrow{\ominus} E_2$) is symmetric to the proof of Lemma 10 ($E_1 \xrightarrow{\oplus} E_2$). Similarly, the proofs of Lemma 11, Lemma 12, Lemma 13 are symmetric to the the proofs of Lemma 7, Lemma 8, Lemma 9. Hence, we omit the detailed proof here.

□

A.7 Computing approximate probability bounds

In this section, we give proof for Theorem 4. To prove Theorem 4, we first introduce the follow Lemmas which establish the soundness of the rules in Figure 16 and Table 1.

LEMMA 15 (SOUNDNESS OF RULES IN TABLE 1). *Given Boolean expressions E_1 and E_2 with probabilities $P(E_1) \in [l_1, u_1]$ and $P(E_2) \in [l_2, u_2]$, and their correlation type \star , applying the formulas in Table 1 yields:*

$$CL(l_1, l_2, \star) \leq P(E_1 \wedge E_2) \leq CU(u_1, u_2, \star) \text{ and } DL(l_1, l_2, \star) \leq P(E_1 \vee E_2) \leq DU(u_1, u_2, \star).$$

PROOF. We will prove the four cases (CL, CU, DL, DU) respectively.

• **CL CU** This case is used for computing the lower and upper bounds of conjunction $E_1 \wedge E_2$.

- (1) $\star = \perp$: In this case, E_1 and E_2 are independent, and thus $P(E_1 \wedge E_2) = P(E_1)P(E_2)$. Given $u_1 \geq P(E_1) \geq l_1$ and $u_2 \geq P(E_2) \geq l_2$, $u_1 u_2 \geq P(E_1)P(E_2) \geq l_1 l_2$. Thus, $P(E_1 \wedge E_2) \in [l_1 l_2, u_1 u_2]$.
- (2) $\star = \top$: In this case, the correlation between E_1 and E_2 is unknown. Given the conjunction $P(E_1 \wedge E_2) = P(E_1) \cdot P(E_2|E_1) = P(E_2) \cdot P(E_1|E_2)$, $P(E_1|E_2) \in [0, 1]$ and $P(E_2|E_1) \in [0, 1]$, in order to maximize $P(E_1 \wedge E_2)$, either $P(E_1|E_2)$ or $P(E_2|E_1)$ must be one. There are following three cases to compute the upper bound of $P(E_1 \wedge E_2)$:
 - $P(E_1) = P(E_2)$: In this case, $P(E_1|E_2) = P(E_2|E_1) = 1$, the upper bound of $P(E_1 \wedge E_2)$ is $P(E_1)$.
 - $P(E_1) < P(E_2)$: In this case, only $P(E_1|E_2)$ can be one, and $P(E_2|E_1) < 1$, the upper bound of $P(E_1 \wedge E_2)$ is $P(E_1)$.
 - $P(E_1) > P(E_2)$: In this case, only $P(E_2|E_1)$ can be one, and $P(E_1|E_2) < 1$, the upper bound of $P(E_1 \wedge E_2)$ is $P(E_2)$.

To summarize the upper bound of $P(E_1 \wedge E_2)$ is $\min(P(E_1), P(E_2))$. As we focus on the upper bound, $P(E_1 \wedge E_2) < \min(u_1, u_2)$.

As we know, $P(E_1 \vee E_2) = P(E_1) + P(E_2) - P(E_1 \wedge E_2) < 1$, we can get $P(E_1 \wedge E_2) > P(E_1) + P(E_2) - 1$.

As we focus on computing the lower bound, $P(E_1 \wedge E_2) > l_1 + l_2 - 1$. Meanwhile, naturally, $P(E_1 \wedge E_2) \in [0, 1]$, so the lower bound can be rewritten as $\max(0, l_1 + l_2 - 1)$.

As a result, $P(E_1 \wedge E_2) \in [\max(0, l_1 + l_2 - 1), \min(u_1, u_2)]$.

- (3) $\star = +$: In this case, $E_1 \xrightarrow{\oplus} E_2$. According to the definition of $\xrightarrow{\oplus}$, we know that $P(E_1 \wedge E_2) > P(E_1)P(E_2)$. Thus the lower bound of $P(E_1 \wedge E_2)$ is $P(E_1)P(E_2)$, which is $l_1 l_2$. The upper bound in this case is same as the upper bound of the unknown (\top) case, which is $\min(u_1, u_2)$. Hence, $P(E_1 \wedge E_2) \in [l_1 l_2, \min(u_1, u_2)]$.
- (4) $\star = -$: In this case, $E_1 \xrightarrow{\ominus} E_2$. According to the definition of $\xrightarrow{\ominus}$, we know that $P(E_1 \wedge E_2) < P(E_1)P(E_2)$. Thus the upper bound of $P(E_1 \wedge E_2)$ is $P(E_1)P(E_2)$, which is $u_1 u_2$. The lower bound in this case is same as the upper bound of the unknown (\top) case, which is $\max(0, l_1 + l_2 - 1)$. Hence, $P(E_1 \wedge E_2) \in [\max(0, l_1 + l_2 - 1), u_1 u_2]$.

• **DL DU** This case is used for computing the lower and upper bounds of disjunction $E_1 \vee E_2$.

- (1) $\star = \perp$: In this case, E_1 and E_2 are independent, and thus $P(E_1 \vee E_2) = P(E_1) + P(E_2) - P(E_1)P(E_2)$. Given $u_1 \geq P(E_1) \geq l_1$ and $u_2 \geq P(E_2) \geq l_2$, $u_1 u_2 \geq P(E_1)P(E_2) \geq l_1 l_2$. Thus, $P(E_1 \vee E_2) \in [l_1 + l_2 - l_1 l_2, u_1 + u_2 - u_1 u_2]$.
- (2) $\star = \top$: In this case, the correlation between E_1 and E_2 is unknown. Given the disjunction $P(E_1 \vee E_2) = P(E_1) + P(E_2) - P(E_1 \wedge E_2)$, in order to maximize $P(E_1 \vee E_2)$, $P(E_1 \wedge E_2)$ should be zero, i.e., E_1 and E_2 are mutually exclusive. Meanwhile, as $P(E_1 \vee E_2) \in [0, 1]$. We refine the upper bound as $\min(1, P(E_1) + P(E_2))$, which is $\min(1, u_1 + u_2)$.

To compute the lower bound of $P(E_1 \vee E_2)$, we want to maximize $P(E_1 \wedge E_2)$, which is $P(E_1|E_2)P(E_2)$ or $P(E_2|E_1)P(E_1)$. There are two cases as follows:

- $P(E_1) = P(E_2)$: in this case, the maximal value of $P(E_1|E_2)$ is 1 and $P(E_1 \wedge E_2) = P(E_1)$. Thus, $P(E_1 \vee E_2) = P(E_1) + P(E_2) - P(E_1) = P(E_2) = (E_1)$.

– $P(E_1) > P(E_2)$: in this case, the maximal value of $P(E_1|E_2)$ is one and $P(E_1 \wedge E_2) = P(E_2)$. Thus, $P(E_1 \vee E_2) = P(E_1) + P(E_2) - P(E_2) = P(E_1)$.

– $P(E_1) < P(E_2)$: in this case, the maximal value of $P(E_2|E_1)$ is one and $P(E_1 \wedge E_2) = P(E_1)$. Thus, $P(E_1 \vee E_2) = P(E_1) + P(E_2) - P(E_1) = P(E_2)$.

To summarize the lower bound of $P(E_1 \vee E_2)$ is $\max(P(E_1), P(E_2))$. As we focus on the lower bound, $P(E_1 \vee E_2) > \max(l_1, l_2)$.

As a result, $P(E_1 \vee E_2) \in [\max(l_1, l_2), \min(1, u_1 + u_2)]$.

(3) $\star = +$: In this case, $E_1 \xrightarrow{\oplus} E_2$. According to the definition of $\xrightarrow{\oplus}$, we know that $P(E_1 \wedge E_2) > P(E_1)P(E_2)$. Thus $P(E_1 \vee E_2) = P(E_1) + P(E_2) - P(E_1 \wedge E_2) < P(E_1) + P(E_2) - P(E_1)P(E_2)$. The upper bound of $P(E_1 \vee E_2)$ is $u_1 + u_2 - u_1u_2$. The lower bound in this case is same as the lower bound of the unknown (\top) case, which is $\max(l_1, l_2)$. Hence, $P(E_1 \vee E_2) \in [\max(l_1, l_2), u_1 + u_2 - u_1u_2]$.

(4) $\star = -$: In this case, $E_1 \xrightarrow{\ominus} E_2$. According to the definition of $\xrightarrow{\ominus}$, we know that $P(E_1 \wedge E_2) < P(E_1)P(E_2)$. Thus $P(E_1 \vee E_2) = P(E_1) + P(E_2) - P(E_1 \wedge E_2) > P(E_1) + P(E_2) - P(E_1)P(E_2)$. Given a function $f(x, y) = x + y - xy$, it is evident that function f is strictly increasing. Hence, $P(E_1 \vee E_2) > f(P(E_1), P(E_2)) \geq f(l_1, l_2)$. Thus, the lower bound of $f(x, y)$ is $l_1 + l_2 - l_1l_2$.

The upper bound in this case is same as the upper bound of the unknown (\top) case, which is $\min(1, u_1 + u_2)$. Hence, $P(E_1 \vee E_2) \in [l_1 + l_2 - l_1l_2, \min(1, u_1 + u_2)]$.

□

LEMMA 16 (SOUNDNESS OF RULES IN FIGURE 16). *Given a Boolean expression E , an environment mapping Γ , a Datalog program D , and a correlation environment Σ , the rules in Figure 16 compute the lower and upper bounds l and u of $P(E)$ such that $l \leq P(E) \leq u$.*

PROOF. We prove that each rule is sound.

- **IN:** Given $I \in \text{InputFacts}(D)$ and $p :: (I \mid \emptyset) \in \text{InputProbs}(D)$, we can get $P(I) = p$ and thus $P(I) \in [p, p]$.
- **OUT:** Γ is an environment mapping output relations O to a pair of variables representing their upper and lower bounds, i.e., $\Gamma[O] = [l, u]$. Hence, $P(O) \in [l, u]$.
- **NEG:** Given $P(E) = [l, u]$, $P(\neg E) = 1 - P(E) \in [1 - u, 1 - l]$.
- **CONJUNCT:** Given $P(E_1) \in [l_1, u_1]$, $P(E_2) \in [l_2, u_2]$, and $\Sigma(E_1, E_2) = \star$, according to Lemma 15, the probability of $E_1 \wedge E_2$ is $P(E_1 \wedge E_2) \in [\text{CL}(l_1, l_2, \star), \text{CU}(u_1, u_2, \star)]$.
- **DISJUNCT:** Given (E_1, p_1) and (E_2, p_2) , $E_1 = \bigwedge_i I_i$, $E_2 = \bigvee_j I_j$, this rule computes the probability of $(E_1, p_1) \vee (E_2, p_2)$. We prove this rule by listing two possible cases.

(1) **Case 1:** $E_1 = \bigwedge_i I_i$, $E_2 = \bigwedge_j I_j$.

Given $P(E_1) \in [l_1, u_2]$, we use e_1 to denote the event (E_1, p_1) , and $P(e_1) = [p_1l_1, p_1u_1]$. Similarly, we use e_2 to denote the event (E_2, p_2) and $P(e_2) \in [p_2l_2, p_2u_2]$. Given $\Sigma(E_1, E_2) = \star$, it's obvious that $\Sigma(e_1, e_2) = \star$. According to Lemma 15, the probability of $P(e_1 \vee e_2) \in [\text{DL}(p_1l_1, p_2l_2, \star), \text{DU}(p_1u_1, p_2u_2, \star)]$.

(2) **Case 2:** $E_1 = \bigwedge_i I_i$, $E_2 = \bigvee_j I_j$

In the derivation graph, edge E_1 leads to one derivation D_1 with path probability $P(D_1|E_1) = p$. The probability of the union of two derivations is $P(D_i \vee D_j) = P((E_i, p_i) \vee (E_j, p_j))$, thus $P(D_i \vee D_j|(E_i, p_i) \vee (E_j, p_j)) = 1$. Therefore, the path probability of the union of multiple derivations is 1, making the path probability of E_2 equal to 1.

Let e_1 denote (E_1, p_1) and e_2 represent $(E_2, 1)$. Given $P(E_1) \in [l_1, u_1]$ and $P(E_2) \in [l_2, u_2]$, we have $P(e_1) = [p_1l_1, p_1u_1]$ and $P(e_2) \in [l_2, u_2]$. Since $\Sigma(E_1, E_2) = \star$, it follows that $\Sigma(e_1, e_2) = \star$. According to Lemma 15, the probability $P(e_1 \vee e_2)$ falls within $[\text{DL}(p_1l_1, l_2, \star), \text{DU}(p_1u_1, u_2, \star)]$.

□

LEMMA 17. *Given a derivation graph G , a Datalog program D , and a correlation mapping Σ , Algorithm 3 outputs a mapping M that associates each output fact with its probability interval. For an output relation R , the mapping $M[R] = (l, u)$ ensures that the probability $P(R)$ lies within the interval $[l, u]$.*

PROOF.

The probability interval (l, u) is directly computed by the APPROXEXPR function. Since the APPROXEXPR function applies the rules in Figure 16, according to Lemma 16, we know that the computed interval (l, u) is sound, i.e., $l \leq P(R) \leq u$, given the output relation R . Here, R is also an internal node n in G . □

PROOF OF THEOREM 4

Based on Lemma 17, we know that Algorithm 3 outputs a mapping M that associates each output fact with its probability interval. For an output relation R , the mapping $M[R] = (l, u)$ ensures that the probability $P(R)$ lies within the interval $[l, u]$. That is, $l \leq P(R) \leq u$.

For the constrained optimization shown in Algorithm 1, it also returns a mapping M . Let's denote $(l^*, u^*) = M(R)$. Since it is based on optimization, we know that $l^* = \min(P(R))$ and $u^* = \max(P(R))$.

Hence, we have $l \leq P(R)$, which implies $l \leq \min(P(R)) = l^*$, and $u \geq P(R)$, which implies $u \geq \max(P(R)) = u^*$. Thus, we have $l \leq l^* \leq u^* \leq u$. □

A.8 Iterative Refinement of Probability Bounds

In this section, we provide the formal proof of Theorem 5. To establish this proof, we need the following Lemmas: Lemma 18, Lemma 19, and Lemma 20.

Given an output relation R , l^* and u^* denote the ground truth probability bounds for R , i.e., $l^* = \min(P(R))$ and $u^* = \max(P(R))$.

LEMMA 18. *Given the lower bound l , upper bound u , constraints ϕ , expression e and the boolean flag `low`, MAKESAT procedure returns an interval (l, u) such that $l \leq l^* \leq u$ if `low` and $l \leq u^* \leq u$ otherwise.*

PROOF OF LEMMA 18

We start the proof by listing all possible cases as follows:

- `low=true`: in this case, it focuses on the lower bound. It starts from $[l, l]$. As we know that $\min(e) = l^*$, and our lower bound is sound: $l \leq l^*$. Let's further classify this case:
 - (1) $l = l^*$: In this case, $\phi \wedge l \leq e \leq l$ is SAT, it does enter the loop in *lines 1-2* and MAKESAT directly return (l, l) . It is consistent with $l \leq l^* \leq l$.
 - (2) $l < l^*$: In this case, $\phi \wedge l \leq e \leq l$ is UNSAT due to the contradiction between $l < l^* = \min(e)$ and $e \leq l$. Hence, it enters the loop in *lines 1-2* to iteratively increasing l until at t -iteration, $\phi \wedge l_t \leq e \leq \text{Incr}(l_t)$ turns SAT. In this case, MAKESAT returns $(l_t, \text{Incr}(l_t))$. It also implies that at $t - 1$ iteration, $\phi \wedge l_{t-1} \leq e \leq l_t$ is UNSAT. From the UNSAT case in $t - 1$ -th iteration, we know that either $l_t < \min(e)$ or $l_{t-1} > \max(e)$. It's obvious that $l_{t-1} > \max(e)$ is an invalid case. Hence, we can infer $l_t < \min(e)$. From the SAT case, we know that $[\min(e), \max(e)]$ overlaps with $[l_t, \text{Incr}(l_t)]$, thus we can get $l_t \leq \max(e) \wedge \text{Incr}(l_t) \geq \min(e)$. Hence, we have $l_t < \min(e) \leq \text{Incr}(l_t)$, which also implies $l_t \leq \min(e) \leq \text{Incr}(l_t)$. As $l^* = \min(e)$, we have $l_t \leq l^* \leq \text{Incr}(l_t)$.
- `low=false`: in this case, it focuses on the upper bound. It starts from $[u, u]$. As we know that $\max(e) = u^*$, and our lower bound is sound: $u \geq u^*$. Let's further classify this case:

- (1) $u = u^*$: In this case, $\phi \wedge u \leq e \leq u$ is SAT, it does enter the loop in *lines 1-2* and **MAKE**SAT directly return (u, u) . It is consistent with $u \leq u^* \leq u$.
- (2) $u > u^*$: In this case, $\phi \wedge u \leq e \leq u$ is UNSAT due to the contradiction between $u > u^* = \max(e)$ and $u \leq e$. Hence, it enters the loop in *lines 1-2* to iteratively decreasing u until at t -iteration, $\phi \wedge \text{Decr}(u_t) \leq e \leq u_t$ turns SAT. In this case, **MAKE**SAT returns $(\text{Decr}(u_t), u_t)$.
It also implies that at $t - 1$ iteration, $\phi \wedge u_t \leq e \leq u_{t-1}$ is UNSAT. From the UNSAT case in $t - 1$ -th iteration, we know that either $u_{t-1} < \min(e)$ or $u_t > \max(e)$. It's obvious that $u_{t-1} < \min(e)$ is an invalid case. Hence, we can infer $u_t > \max(e)$. From the SAT case, we know that $[\min(e), \max(e)]$ overlaps with $[\text{Decr}(u_t), u_t]$, thus we can get $\text{Decr}(u_t) \leq \max(e) \wedge u_t \geq \min(e)$. Hence, we have $\text{Decr}(u_t) < \max(e) < u_t$, which also implies $\text{Decr}(u_t) < \max(e) \leq u_t$. As $u^* = \max(e)$, we have $\text{Decr}(u_t) \leq u^* \leq u_t$.

□

LEMMA 19. *BOUND*BOUNDS Procedure in Algorithm 5 outputs a mapping B that maps R to a quadruple (l^-, l^+, u^-, u^+) , such that $l^- \leq l^* \leq l^+$ and $u^- \leq u^* \leq u^+$.

PROOF OF LEMMA 19

We start the proof by listing all possible cases as follows:

- *lower bound*: When the signal low is set to true, **MAKE**SAT function in *Line 3* is invoked in Algorithm 5 and returns (l^-, l^+) . According to Lemma 18, we know that $l^- \leq l^* \leq l^+$.
- *upper bound*: When the signal low is set to false, **MAKE**SAT function in *Line 4* is invoked in Algorithm 5 and returns (u^-, u^+) . According to Lemma 18, we know that $u^- \leq u^* \leq u^+$.

□

LEMMA 20. *BINARY*SEARCH procedure in Algorithm 7 takes an output relation R , lower bound l , upper bound u , the boolean flag *low* and error bound δ as inputs, and outputs an interval (L, U) . If *low* is set to true, $L \leq l^* \leq U$ and $|U - L| < \delta$. Otherwise, if *low* is set to false, $L \leq u^* \leq U$ and $U - L < \delta$.

PROOF OF LEMMA 20

We start the proof by listing all possible cases as follows:

- *low = true*:
 - (1) $(u - l) < \delta$: **BINARY**SEARCH procedure directly returns (l, u) . Here, the output $[L, U] = [l, u]$.
 - (2) $(u - l) \geq \delta$: it enters the loop in *lines 1-6*. Given $\text{mid} = (l + u)/2$. If $\phi \wedge l \leq \text{Expr}(R) \leq \text{mid}$ is UNSAT, we know that intervals $[l, \text{mid}]$ and $[l^*, u^*]$ are disjoint, we have $l > u^* \vee \text{mid} < l^*$. It's obvious that $l > u^*$ is false. Given $\text{mid} < l^*$, l^* falls into the interval $[\text{mid}, l]$. Next iteration focuses on splitting $[\text{mid}, l]$ until its distance is less than δ .
Otherwise, if $\phi \wedge l \leq \text{Expr}(R) \leq \text{mid}$ is SAT, it means that two intervals $[l, \text{mid}]$ and $[l^*, u^*]$ overlap. Hence, we can get $l \leq u^* \wedge \text{mid} \geq l^*$. Thus l^* falls into the interval $[l, \text{mid}]$. Next iteration focuses on splitting $[l, \text{mid}]$ until its distance is less than δ .

Hence, after each iteration, the interval $[l, u]$ we get is guaranteed to contain l^* . When it terminates, the loop condition is falsified, and it returns the output interval $[L, U]$ where $U - L < \delta$.
- *low = false*:
 - (1) $(u - l) < \delta$: **BINARY**SEARCH procedure directly returns (l, u) . Here, the output $[L, U] = [l, u]$.
 - (2) $(u - l) \geq \delta$: it enters the loop in *lines 1-6*. Given $\text{mid} = (l + u)/2$. If $\phi \wedge \text{mid} \leq \text{Expr}(R) \leq u$ is UNSAT, we know that intervals $[\text{mid}, u]$ and $[l^*, u^*]$ are disjoint, we have $\text{mid} > u^* \vee u < l^*$. It's obvious that $u < l^*$ is false. Given $\text{mid} > u^*$, u^* falls into the interval $[\text{mid}, l]$. Next iteration focuses on splitting $[\text{mid}, l]$ until its distance is less than δ .
Otherwise, if $\phi \wedge \text{mid} \leq \text{Expr}(R) \leq u$ is SAT, it means that two intervals $[\text{mid}, u]$ and $[l^*, u^*]$ overlap. Thus, we can get $l \leq u^* \wedge \text{mid} \geq l^*$. Thus u^* falls into the interval $[\text{mid}, u]$. Next iteration focuses on splitting $[\text{mid}, u]$ until its distance is less than δ .

Hence, after each iteration, the interval $[l, u]$ we get is guaranteed to contain u^* . When it terminates, the loop condition is falsified, and it returns the output interval $[L, U]$ where $U - L < \delta$. \square

PROOF OF THEOREM 5

In MAKEDELTA-PRECISE procedure, the last step invokes BINARYPROCEDURE. For lower bound, it returns (l^-, l^+) . For upper bound, it returns (u^-, u^+) . According to Lemma 20, we know that $l^- \leq l^* \leq l^+$, $l^+ - l^- < \delta$ and $u^- \leq u^* \leq u^+$, $u^+ - u^- < \delta$. MAKEDELTA-PRECISE procedure returns (l, u) as its probability bounds, where $l = l^-$, $u = u^+$.

Given $l^* < l^+$, we have $l^* - l = l^* - l^- \leq l^+ - l^- < \delta$. Hence, $l > l^* - \delta$. Given $l = l^- \leq l^*$, we can get $l^* - \delta \leq l \leq l^*$.

Given $u^* > u^-$, we have $u - u^* = u^+ - u^* \leq u^+ - u^- < \delta$. Given $u = u^+ \geq u^*$, we can get $u^* \leq u \leq u^* + \delta$. \square

A.9 Correlation type of constantly true and false cases

For cases that are constantly true or false, the correlation type—whether $+$, $-$, \perp , \top —does not affect the probability computation in Table 1. Regardless of the assigned correlation type, the computation remains sound and correct.

- $E_1 = \text{true}, E_2 = \text{true}$. $e_1 = Pr(E_1) = 1, e_2 = Pr(E_2) = 1$.
 - $CL(e_1, e_2, +) = CL(e_1, e_2, -) = CL(e_1, e_2, \perp) = CL(e_1, e_2, \top) = 1$
 - $CU(e_1, e_2, +) = CU(e_1, e_2, -) = CU(e_1, e_2, \perp) = CU(e_1, e_2, \top) = 1$
 - $DL(e_1, e_2, +) = DL(e_1, e_2, -) = DL(e_1, e_2, \perp) = DL(e_1, e_2, \top) = 1$
 - $DU(e_1, e_2, +) = DU(e_1, e_2, -) = DU(e_1, e_2, \perp) = DU(e_1, e_2, \top) = 1$
 As a result, $Pr(E_1 \wedge E_2) \in [1, 1] = 1$ and $Pr(E_1 \vee E_2) \in [1, 1] = 1$ for any correlation type.
- $E_1 = \text{true}, E_2 = \text{false}$. $e_1 = Pr(E_1) = 1, e_2 = Pr(E_2) = 0$.
 - $CL(e_1, e_2, +) = CL(e_1, e_2, -) = CL(e_1, e_2, \perp) = CL(e_1, e_2, \top) = 0$
 - $CU(e_1, e_2, +) = CU(e_1, e_2, -) = CU(e_1, e_2, \perp) = CU(e_1, e_2, \top) = 0$
 - $DL(e_1, e_2, +) = DL(e_1, e_2, -) = DL(e_1, e_2, \perp) = DL(e_1, e_2, \top) = 1$
 - $DU(e_1, e_2, +) = DU(e_1, e_2, -) = DU(e_1, e_2, \perp) = DU(e_1, e_2, \top) = 1$
 As a result, $Pr(E_1 \wedge E_2) \in [0, 0] = 0$ and $Pr(E_1 \vee E_2) \in [1, 1] = 1$ for any correlation type.
- $E_1 = \text{false}, E_2 = \text{true}$. $e_1 = Pr(E_1) = 0, e_2 = Pr(E_2) = 1$.
 - Symmetric to Case $(E_1 = \text{true}, E_2 = \text{false})$, so we will not provide a proof here.
- $E_1 = \text{false}, E_2 = \text{false}$. $e_1 = Pr(E_1) = 0, e_2 = Pr(E_2) = 0$.
 - $CL(e_1, e_2, +) = CL(e_1, e_2, -) = CL(e_1, e_2, \perp) = CL(e_1, e_2, \top) = 0$
 - $CU(e_1, e_2, +) = CU(e_1, e_2, -) = CU(e_1, e_2, \perp) = CU(e_1, e_2, \top) = 0$
 - $DL(e_1, e_2, +) = DL(e_1, e_2, -) = DL(e_1, e_2, \perp) = DL(e_1, e_2, \top) = 0$
 - $DU(e_1, e_2, +) = DU(e_1, e_2, -) = DU(e_1, e_2, \perp) = DU(e_1, e_2, \top) = 0$
 As a result, $Pr(E_1 \wedge E_2) \in [0, 0] = 0$ and $Pr(E_1 \vee E_2) \in [0, 0] = 0$ for any correlation type.

In summary, the correlation type does not affect the computation of always-true and always-false cases. When assigning a positive correlation ($+$) to a pair of constantly-true and constantly-false, the computation remains correct.

A.10 Implied Conditional Dependence from Overview

In this subsection, we provide a proof that input facts $\text{edge}(2, 5)$ and $\text{edge}(2, 6)$ from Figure 4 must be dependent. Let A, B, C denote $\text{edge}(2, 5)$, $\text{edge}(2, 6)$, and $\text{edge}(1, 4)$ respectively. Then, we have $P(A \wedge C) = P(A|C) \times P(C) = 0.8 \times 0.6 = 0.48$. Since $P(A \wedge C) = P(A \wedge C \wedge B) + P(A \wedge C \wedge \neg B)$, we have $P(A \wedge B \wedge C) = 0.48 - x$ where $x = P(A \wedge C \wedge \neg B)$. Next, consider $P(B \wedge C) = P(B|C) \times P(C) = 0.83 \times 0.6 = 0.498$. Next, observe that $P(C) = P(C \wedge B) + P(C \wedge \neg B)$;

thus, $P(C \wedge \neg B) = 0.6 - 0.498 = 0.102$. Now, because $x = P(A \wedge C \wedge \neg B) \leq P(C \wedge \neg B) = 0.102$, we obtain $x \leq 0.102$. Finally, note that $P(A \wedge B) \geq P(A \wedge B \wedge C) = 0.48 - x \geq (0.48 - 0.102) = 0.378$. Thus, clearly, $P(A \wedge B) > P(A) \times P(B) = 0.36$, so $\text{edge}(2, 5)$ and $\text{edge}(2, 6)$ are positively correlated.

Received 2025-03-25; accepted 2025-08-12