



HAL
open science

Computing the Shapley Value of Facts in Query Answering

Daniel Deutch, Nave Frost, Benny Kimelfeld, Mikaël Monet

► **To cite this version:**

Daniel Deutch, Nave Frost, Benny Kimelfeld, Mikaël Monet. Computing the Shapley Value of Facts in Query Answering. SIGMOD Conference 2022, Jun 2022, Philadelphia, United States. hal-03514297

HAL Id: hal-03514297

<https://inria.hal.science/hal-03514297v1>

Submitted on 6 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Computing the Shapley Value of Facts in Query Answering

Daniel Deutch
Tel Aviv University
Blavatnik School of
Computer Science, Tel Aviv,
Israel
danielde@post.tau.ac.il

Nave Frost
Tel Aviv University
Blavatnik School of
Computer Science, Tel Aviv,
Israel
navefrost@mail.tau.ac.il

Benny Kimelfeld
Technion - Israel Institute
of Technology
Faculty of Computer
Science, Haifa, Israel
bennyk@cs.technion.ac.il

Mikaël Monet
Univ. Lille, Inria, CNRS,
Centrale Lille, UMR 9189
CRISTAL, F-59000 Lille,
France
mikaël.monet@inria.fr

ABSTRACT

The Shapley value is a game-theoretic notion for wealth distribution that is nowadays extensively used to explain complex data-intensive computation, for instance, in network analysis or machine learning. Recent theoretical works show that query evaluation over relational databases fits well in this explanation paradigm. Yet, these works fall short of providing practical solutions to the computational challenge inherent to the Shapley computation. We present in this paper two practically effective solutions for computing Shapley values in query answering. We start by establishing a tight theoretical connection to the extensively studied problem of query evaluation over probabilistic databases, which allows us to obtain a polynomial-time algorithm for the class of queries for which probability computation is tractable. We then propose a first practical solution for computing Shapley values that adopts tools from probabilistic query evaluation. In particular, we capture the dependence of query answers on input database facts using Boolean expressions (data provenance), and then transform it, via Knowledge Compilation, into a particular circuit form for which we devise an algorithm for computing the Shapley values. Our second practical solution is a faster yet inexact approach that transforms the provenance to a Conjunctive Normal Form and uses a heuristic to compute the Shapley values. Our experiments on TPC-H and IMDB demonstrate the practical effectiveness of our solutions.

1 INTRODUCTION

Explaining query answers has been the objective of extensive research in recent years [12, 15, 20, 29, 30]. A prominent approach is to devise an explanation based on the facts that were used for deriving the answers; these facts are often termed *provenance* or *lineage* of the query answer. For illustration, consider a query asking whether there exists a route from the USA to France with at most one connection over a database of airports and flights. The answer is Boolean, and upon receiving a positive answer one may seek explanations of why it is so. The basis for such explanations would include all details of qualifying routes that are used in the derivation of the answer. Unfortunately, the number of relevant routes might be huge (specifically, quadratic in the database size). Moreover, it is conceivable that different facts differ considerably in their importance to the answer at hand; for instance, some flights may be crucial to enabling the USA-France connection, while others may be easily replaced by alternatives.

To address these issues, there have been several proposals for principled ways of *quantifying* the contribution of input facts to query answers [20, 24, 25, 30]. We focus here on the recent approach of Livshits et al. [20] that applies to this setting the notion of *Shapley*

values [32]—a game-theoretic function for distributing the wealth of a team in a cooperative game. This function has strong theoretical justifications [28], and indeed, it has been applied across various fields such as economics, law, environmental science, and network analysis. It has also been used for explanations in data-centric paradigms such as knowledge representation [16, 39] and machine learning [21, 22]. In the context of relational databases, given a query $q(\bar{x})$, a database D , an input fact $f \in D$ and a tuple \bar{i} of same arity as \bar{x} , the *Shapley value of f in D for query $q(\bar{x})$ and tuple \bar{i}* intuitively represents the contribution of f to the presence (or absence) of \bar{i} in the query result.

Livshits et al. [20, 27] initiated the study of the computational complexity of calculating Shapley values in query answering. They showed mainly lower bounds on the complexity of the problem, with the exception of the sub-class of self-join free SPJ queries called *hierarchical*, where they gave a polynomial-time algorithm. The results are more positive if imprecision is allowed, as they showed that the problem admits a tractable approximation scheme (FPRAS, to be precise) via Monte Carlo sampling. The state of affairs is that the class of known tractable cases (namely the hierarchical conjunctive queries) is highly restricted, and the approximation algorithms with theoretical guarantees are impractical in the sense that they require a large number of executions of the query over database subsets (the samples). Hence, the theoretical analysis of Livshits et al. [20, 27] does not provide sufficient evidence of practical feasibility for adopting the Shapley value as a measure of responsibility in query answering. Moreover, the results of Livshits et al. [20] imply that, for self-join free SPJ queries the class of tractable queries for computing Shapley values coincides with the class of tractable queries in probabilistic tuple-independent databases [7]. Yet, no direct connection has been made between these two problems and, theoretically speaking, it has been left unknown whether algorithms for probabilistic databases can be used for Shapley computation.

Recently, Van den Broeck et al. [36] and Arenas et al. [2, 3] investigated the computational complexity of the *SHAP-score* [22], a notion used in machine learning for explaining the predictions of a model. While both are based on the general notion of Shapley value, the SHAP-score for machine learning and Shapley values for databases are different. In the latter case, the players are the tuples of the database and the game function that is used is simply the value of the query on a subset of the database, while in the former case, the players are the features of the model and the game function is a conditional expectation of the model's output (see Section 6.2 for a more formal definition of the SHAP-score). Remarkably, Van den Broeck et al. [36] have shown that computing the SHAP-score is equivalent (in terms of polynomial-time reductions) to the problem of computing the expected value of the model. One of our

contributions is to show that the techniques developed by [2, 3, 36] can be adapted to the context of Shapley values for databases. For instance, by adapting to our context the proof of Van den Broeck et al. [36] that computing the SHAP-score reduces to computing the expected value of the model, we resolve the aforementioned open question affirmatively: we prove that Shapley computation can be efficiently (polynomial-time) reduced to probabilistic query answering. Importantly, this applies not only to the restricted class of SPJ queries without self-joins, but to *every* database query. Hence, extending theory to practice, one can compute the Shapley values using a query engine for probabilistic databases.

In turn, a common approach that was shown to be practically effective for probabilistic databases is based on *Knowledge Compilation* [11, 18]. In a nutshell, the idea is to first compute the Boolean provenance of a given output tuple in the sense of Imielinski and Lipski [17], and then to “compile” the provenance into a particular circuit form that is more favorable for probability computation. Specifically, the target class of this compilation is that of *deterministic and decomposable circuits* (d-D). In our case, rather than going through probabilistic databases, we devise a more efficient approach that computes the Shapley values directly from the d-D circuit. This is similar to how Arenas et al. [2, 3] directly prove that the SHAP-score can be computed efficiently over such circuits, without using the more general results of [36]. By adapting the proof of [2, 3], we show how, given a d-D circuit representing the provenance of an output tuple, we can efficiently compute the Shapley value of every input fact. While the aforementioned properties of the circuit are not guaranteed in general (beyond the class of hierarchical queries), we empirically show the applicability and usefulness of the approach even for non-hierarchical queries.

Our experimental results (see below) indicate that our exact computation algorithm is fast in most cases, but is too costly in others. For the latter cases, we propose a heuristic approach to retrieve the relative order of the facts by their Shapley values, without actually computing these values. Indeed, determining the most influential facts is in many cases already highly useful, even if their precise contribution remains unknown. The solution that we propose to this end is termed *CNF Proxy*; it is based on a transformation of the provenance to Conjunctive Normal Form (CNF) and using it to compute proxy values intuitively based on (1) the number of clauses in which a variable occur and (2) its alternatives in each clause. These are two aspects that are correlated with Shapley values. The proxy values may be very different from the real Shapley values, and yet, when we order facts according to their proxy values we may intuitively get an ordering that is similar to the order via Shapley. Our experiments validate that this intuition indeed holds for examined benchmarks.

We have experimented with multiple queries from the two standard benchmarks TPC-H and IMDB. Our main findings are as follows. In most cases (98.67% of the IMDB output tuples and 83.83% for TPC-H), our exact computation algorithm terminates in 2.5 seconds or less, given the provenance expression. In the vast majority of remaining cases the execution is very costly, typically running out of memory already in the Knowledge Compilation step. By contrast, our inexact solution *CNF Proxy* is extremely fast even for these hard cases – it typically terminates in a few milliseconds with the worst observed case (an outlier) being 4 seconds. In fact, it is

faster by several orders of magnitude than sampling-based approximation techniques (the Monte Carlo sampling proposed in [20] as well as a popular sampling-based solution for Shapley values in Machine Learning (Kernel SHAP [22])). To measure quality, we use *CNF Proxy* to rank the input tuples, and compared the obtained ranked lists to ranking by actual Shapley values (in cases where exact computation has succeeded), using the standard measures of nDCG and Precision@k. Our solution outperforms the competitors in terms of quality as well.

We then propose a simple hybrid approach: execute the exact algorithm until it either terminates or a timeout elapses. If we have reached the timeout, resort to executing *CNF Proxy* and rank the facts based on the obtained values. We show experiments with different timeout values, justifying our choice of 2.5 seconds.

Hence, our contributions are both of a theoretical and practical nature and can be summarized as follows.

- By adapting the proof technique of [36], we establish a fundamental result about the complexity of computing Shapley values over relational queries: Shapley values can be computed in polynomial time—in data complexity—whenever the query can be evaluated in polynomial time over tuple-independent probabilistic databases (Proposition 3.1). This holds for *every* query.
- By adapting the proof technique of [2, 3], we devise a novel algorithm for computing Shapley values for query evaluation via compilation to a deterministic and decomposable circuit (Proposition 4.4). We show that this algorithm is practical and has the theoretical guarantee of running in polynomial time in the size of the circuit.
- We present a novel heuristic, *CNF Proxy*, that is fast yet inexact, and is practically effective if we are interested in ranking input facts by their contribution rather than computing exact Shapley values (Section 5).
- We describe a thorough experimental study of our algorithms over realistic data and show their efficiency (Section 6).

Related work. Existing models for explaining database query results may roughly be divided in two categories: (1) models that are geared for tracking/presenting provenance of output tuples, e.g., the set of all input facts participating in their computation [6], possibly alongside a description of the ways they were used, in different granularity levels (e.g., [4, 5, 14]); (2) models that quantify contributions of input facts [20, 24, 25, 30], which is the approach that we follow here. Works in the latter context often have connections with the influential line of work on probabilistic databases [33], and we show that this is the case for Shapley computation as well.

As already mentioned, an important point of comparison is the work of Van den Broeck et al. [36] and that of Arenas et al. [2, 3] on the *SHAP-score*. While we show that the proof techniques developed in this area can be adapted to the context of relational databases, we point out that the two sets of results obtained (for SHAP-score and for Shapley values for databases) seem incomparable, as we do not see a way of proving results for Shapley value for query answering using the results on the SHAP-score, or vice-versa. In fact, this adaptation only works up to a certain point. For instance, the *efficiency axiom* of the Shapley value immediately implies that

computing the expected value of a model can be reduced in polynomial time to computing the SHAP-score of its features; in contrast, this axiom does not seem to yield any clear such implication in our context (see our Open Problem 1 and the discussion around it).

Paper organization. We formalize the notion of Shapley values for query answering in Section 2. In Section 3 we present the theoretical connection to probabilistic databases and its implications. Our exact computation algorithm is presented in Section 4 and our heuristic in Section 5. Experimental results are presented in Section 6 and we conclude in Section 7.

2 THE SHAPLEY VALUE OF FACTS

We define here the main notion and illustrate it with an example.

Relational databases and queries. Let $\Sigma = \{R_1, \dots, R_n\}$ be a *signature*, consisting of *relation names* R_i each with its associated *arity* $\text{ar}(R_i) \in \mathbb{N}$, and Const be a set of *constants*. A *fact* over (Σ, Const) is simply a term of the form $R(a_1, \dots, a_{\text{ar}(R)})$, for $R \in \Sigma$ and $a_i \in \text{Const}$. A (Σ, Const) -*database* D , or simply a *database* D , is a finite set of facts over (Σ, Const) . We assume familiarity with the most common classes of query languages and refer the reader to [1] for the basic definitions. In particular, we recall the equivalence between relational algebra and relational calculus [1], and the fact that Select-Project-Join-Union (SPJU) queries are equivalent to unions of conjunctive queries (UCQs). Depending of the context and for consistency with relevant past publications, we will use terminology of either relational calculus or relational algebra. What we call a *Boolean query* is a query q that takes as input a database D and outputs $q(D) \in \{0, 1\}$. If $q(\bar{x})$ is a query with free variables \bar{x} and \bar{i} is a tuple of constants of same length as \bar{x} , we denote by $q[\bar{x}/\bar{i}]$ the Boolean query defined by: $q[\bar{x}/\bar{i}](D) = 1$ if and only if \bar{i} is in the output of $q(\bar{x})$ on D .

Shapley values of facts. Following [20, 27], we use the notion of Shapley values [32] to attribute a contribution to facts of an input database. In this context, the database D is traditionally partitioned into two sets of facts: a set D_x of so-called *exogenous* facts, and a set D_n of *endogenous* facts. The idea is that exogenous facts are considered as given, while endogenous facts are those to which we would like to attribute contributions. Let q be a Boolean query and $f \in D_n$ be an endogenous fact. *The Shapley value of f in D for query q* , denoted $\text{Shapley}(q, D_n, D_x, f)$, is defined as

$$\text{Shapley}(q, D_n, D_x, f) \stackrel{\text{def}}{=} \sum_{E \subseteq D_n \setminus \{f\}} \frac{|E|!(|D_n| - |E| - 1)!}{|D_n|!} (q(D_x \cup E \cup \{f\}) - q(D_x \cup E)). \quad (1)$$

Notice that here, $|E|!(|D_n| - |E| - 1)!$ is the number of permutations of D_n with all endogenous facts in E appearing first, then f , and finally, all the other endogenous facts. Intuitively then, the value $\text{Shapley}(q, D_n, D_x, f)$ represents the contribution of f to the query's output: the higher this value is, the more f helps in satisfying q .

For non-Boolean queries $q(\bar{x})$, we are interested in the Shapley value of the fact f for every individual tuple \bar{i} in the output [20]. The extension to non-Boolean $q(\bar{x})$ is then straightforward: the

Shapley value of the fact f for the answer \bar{i} to $q(\bar{x})$ is the value $\text{Shapley}(q[\bar{x}/\bar{i}], D_n, D_x, f)$. Therefore, the computational challenge reduces to that of the Boolean query $q[\bar{x}/\bar{i}]$. Hence, in the theoretical analysis we focus on Boolean queries, and we go back to considering non-Boolean queries when we study the implementation aspects (starting in Section 4.2).

Example 2.1. Consider the database D and the Boolean query q from Figures 1a and 1c. All facts in table `FLIGHTS` are endogenous, while facts in `AIRPORTS` are exogenous. To alleviate the notation we write, e.g., a_1 for `FLIGHTS(JFK, CDG)`. The query q checks if there are routes from “USA” to “FR” with one or less connecting flights. Let us compute the Shapley value of all endogenous facts. First, we notice that fact a_8 is not part of any valid route, so $\text{Shapley}(q, D_n, D_x, a_8) = 0$ by Equation (1). Next, let us focus on a_1 . Since a_1 is a valid route on its own, adding it to any subset of (endogenous) facts E such that E does not contain a valid route results in $q(D_x \cup E \cup \{a_1\}) - q(D_x \cup E) = 1$ (for all other subsets the difference will be 0). The relevant subsets are the empty set, all singletons $\{a_i\}$ for $2 \leq i \leq 8$ (7 singletons), all the pairs of tuples from a_2, \dots, a_8 excluding the pairs $\{a_2, a_4\}$, $\{a_2, a_5\}$, $\{a_3, a_4\}$, $\{a_3, a_5\}$, and $\{a_6, a_7\}$ (so $\binom{7}{2} - 5 = 16$ pairs), the quadruples $\{a_2, a_3, a_6, a_8\}$, $\{a_2, a_3, a_7, a_8\}$, $\{a_4, a_5, a_6, a_8\}$, and $\{a_4, a_5, a_7, a_8\}$, and overall 14 triplets (left to the reader). Summing it all up results in

$$\begin{aligned} \text{Shapley}(q, D_n, D_x, a_1) &= 1 \cdot \frac{0! \cdot 7!}{8!} + 7 \cdot \frac{1! \cdot 6!}{8!} + 16 \cdot \frac{2!5!}{8!} + \\ &14 \cdot \frac{3!4!}{8!} + 4 \cdot \frac{4!3!}{8!} = \frac{43}{105} \approx 0.4095. \end{aligned}$$

Similarly one can compute the Shapley value of the remaining facts, and find that for $a_i \in \{a_2, a_3, a_4, a_5\}$ it holds that $\text{Shapley}(q, D_n, D_x, a_i) = \frac{23}{210} \approx 0.1095$, and that for $a_i \in \{a_6, a_7\}$ we have $\text{Shapley}(q, D_n, D_x, a_i) = \frac{8}{105} \approx 0.0762$.

3 REDUCTION TO PROBABILISTIC DATABASES

In this section we investigate the complexity of computing Shapley values. As explained in the previous section, the non-Boolean setting of the problem may be reduced to that of Boolean queries, so we will study the following problem for a given Boolean query q .

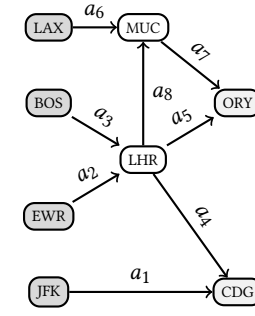
PROBLEM: $\text{Shapley}(q)$
INPUT: A database $D = D_x \cup D_n$ and an endogenous fact $f \in D_n$.
OUTPUT: The value $\text{Shapley}(q, D_n, D_x, f)$.

Note that for the sake of complexity analysis, the query q is assumed to be fixed, so that each query gives rise to a different computational problem; we are then considering what is called the *data complexity* [37]. This assumption is motivated by the fact that in practice, the queries are much smaller than the databases.

The complexity of this problem has been studied in [20, 27], where in particular a dichotomy has been obtained for *self-join-free Boolean conjunctive queries* (sjfbcqs). There, the authors show that, for every sjfbcq q , either q is *hierarchical* (we will not need to define this notion here) and $\text{Shapley}(q)$ can be solved in polynomial time, or q is not hierarchical and then $\text{Shapley}(q)$ is intractable (specifically, $\text{FP}^{\#P}$ -hard). It turns out that the tractability criterion that

FLIGHTS (endo)			AIRPORTS (exo)		
	Src	Dest	Name	Country	
a_1	JFK	CDG	b_1	JFK	USA
a_2	EWR	LHR	b_2	EWR	USA
a_3	BOS	LHR	b_3	BOS	USA
a_4	LHR	CDG	b_4	LAX	USA
a_5	LHR	ORY	b_5	LHR	EN
a_6	LAX	MUC	b_6	MUC	GR
a_7	MUC	ORY	b_7	ORY	FR
a_8	LHR	MUC	b_8	CDG	FR

(a) Database of flights and airports



(b) Flights in graph view. Dark and light gray depict “USA” and “FR” airports respectively

Routes from “USA” to “FR” with one or less connecting flights

$$q_1 = \exists x, y : \text{AIRPORTS}(x, \text{“USA”}) \wedge \text{AIRPORTS}(y, \text{“FR”}) \wedge \text{FLIGHTS}(x, y)$$

$$q_2 = \exists x, y, z : \text{AIRPORTS}(x, \text{“USA”}) \wedge \text{AIRPORTS}(z, \text{“FR”}) \wedge \text{FLIGHTS}(x, y) \wedge \text{FLIGHTS}(y, z)$$

$$q = q_1 \vee q_2$$

(c) q is a Boolean union of conjunctive queries (UCQ)

$$\underbrace{(a_1 \wedge b_1 \wedge b_8)}_{q_1 \text{ provenance}} \vee \underbrace{(a_2 \wedge a_4 \wedge b_2 \wedge b_8) \vee (a_2 \wedge a_5 \wedge b_2 \wedge b_7) \vee (a_3 \wedge a_4 \wedge b_3 \wedge b_8) \vee (a_3 \wedge a_5 \wedge b_3 \wedge b_7) \vee (a_6 \wedge a_7 \wedge b_5 \wedge b_7)}_{q_2 \text{ provenance}}$$

(d) The lineage $\text{Lin}(q, D)$ as a formula in Disjunctive Normal Form (DNF)Figure 1: The database, query q and its lineage used for our running example

is obtained—being hierarchical—is exactly the same as in the context of probabilistic query evaluation (PQE); see, e.g., [7, 8]. In fact, the main result of this section is that this is not a coincidence: we prove that, for every Boolean query q (not just for sjfbcqs), if PQE is tractable for q then so is the problem $\text{Shapley}(q)$. Since PQE has been intensively studied already, our result allows us to vastly extend the tractable cases identified in [20, 27]. We now proceed with the definitions and proof of this result, and explain its consequences.

Probabilistic query evaluation. A tuple-independent (TID) database is a pair consisting of a database D and a function π mapping each fact $f \in D$ to a probability value $\pi(f) \in [0, 1]$. The TID (D, π) defines a probability distribution Pr_π on $D' \subseteq D$, where $\text{Pr}_\pi(D') \stackrel{\text{def}}{=} \prod_{f \in D'} \pi(f) \times \prod_{f \in D \setminus D'} (1 - \pi(f))$. Given a Boolean query q , the probability that q is satisfied by (D, π) is $\text{Pr}(q, (D, \pi)) \stackrel{\text{def}}{=} \sum_{D' \subseteq D \text{ s.t. } q(D')=1} \text{Pr}_\pi(D')$. The probabilistic query evaluation problem for q , $\text{PQE}(q)$ for short, is then defined as follows.

PROBLEM:	$\text{PQE}(q)$
INPUT:	A tuple-independent database (D, π) .
OUTPUT:	The value $\text{Pr}(q, (D, \pi))$.

For two computational problems A and B , we write $A \leq_T^P B$ to assert the existence of a polynomial-time Turing reduction from A to B . We are ready to state the main result of this section.

PROPOSITION 3.1. *For every Boolean query q , we have that $\text{Shapley}(q) \leq_T^P \text{PQE}(q)$.*

This result implies that for any query q for which $\text{PQE}(q)$ is tractable then so is $\text{Shapley}(q)$. Dalvi and Suciu [8] showed a dichotomy for unions of conjunctive queries: for every such query q , either $\text{PQE}(q)$ is solvable in polynomial time, in which case q is

called *safe*¹, or $\text{PQE}(q)$ is $\text{FP}^{\#P}$ -hard (and q is called *unsafe*). Therefore, we obtain as a direct corollary of Proposition 3.1 that $\text{Shapley}(q)$ can be solved in polynomial time for all safe queries.

COROLLARY 3.2. *If q is a safe UCQ then $\text{Shapley}(q)$ can be solved in polynomial time.*

In particular, this corollary generalizes the tractability result obtained in [20], to account for CQs with self-joins and even unions of such queries. We now prove Proposition 3.1.

PROOF OF PROPOSITION 3.1. For a Boolean query q , database $D = D_x \cup D_n$, and integer $k \in \{0, \dots, |D_n|\}$, define

$$\#\text{Slices}(q, D_x, D_n, k) \stackrel{\text{def}}{=} |\{E \subseteq D_n \mid |E| = k \text{ and } q(D_x \cup E) = 1\}|.$$

Then, by grouping by size the terms E from Equation 1 we obtain

$$\text{Shapley}(q, D_n, D_x, f) = \sum_{k=0}^{|D_n|-1} \frac{k!(|D_n| - k - 1)}{|D_n|} \left(\begin{aligned} &\#\text{Slices}(q, D_x \cup \{f\}, D_n \setminus \{f\}, k) \\ &- \#\text{Slices}(q, D_x, D_n \setminus \{f\}, k) \end{aligned} \right). \quad (2)$$

All arithmetical terms (such as $k!$ or $|D_n|!$) can be computed in polynomial time. Therefore, to prove that $\text{Shapley}(q) \leq_T^P \text{PQE}(q)$, it is enough to show that, given an oracle to the problem $\text{PQE}(q)$, we can compute in polynomial time the quantities $\#\text{Slices}(q, D_x, D_n, k)$, for some arbitrary $D = D_x \cup D_n$ and $k \in \{0, \dots, |D_n|\}$. This is what we do next.

The proof is similar to that of [36, Theorem 2] in the context of SHAP-score for machine learning (but, as explained in the Introduction, the two results seem to be incomparable).

We wish to compute $\#\text{Slices}(q, D_x, D_n, k)$, for some database $D = D_x \cup D_n$ and integer $k \in \{0, \dots, |D_n|\}$. Let $n = |D_n|$ be the number of endogenous facts of D . For $z \in \mathbb{Q}$, we define a TID database (D_z, π_z)

¹This notion of safety is distinct from the “usual” notion of query safety [1] that ensures domain independence.

In the last two examples, lineages were represented with Boolean formulas in DNF. Since a lineage is a Boolean function, it can be represented with any formalism that allows to represent Boolean functions. We next review some classes of circuits from the field of knowledge compilation that will be relevant for our work.

Knowledge compilation classes. Let C be a Boolean circuit, featuring \wedge , \vee , \neg , and variable gates, with the usual semantics.² For a gate g of C , we denote by $\text{Vars}(g)$ the set of variables that have a directed path to g . An \wedge -gate g of C is *decomposable* if for every two input gates $g_1 \neq g_2$ of g we have $\text{Vars}(g_1) \cap \text{Vars}(g_2) = \emptyset$. We call C *decomposable* if all \wedge -gates are. An \vee -gate g of C is *deterministic* if the Boolean functions captured by each pair of distinct input gates of g are pairwise disjoint; i.e., no assignment satisfies both. We call C *deterministic* if all \vee -gates in it are. A *deterministic and decomposable (d-D [26]) Boolean circuit* is a Boolean circuit that is both deterministic and decomposable. If C is a Boolean circuit we write $\text{Vars}(C)$ to denote the set of variables that appear in it.

Example 4.3. Recall $\text{ELin}(q, D_n, D_x)$ from Example (4.2) represented as a DNF. Figure 2 depicts a d-D circuit for $\text{ELin}(q, D_n, D_x)$. The output gate, for instance, is a deterministic \vee -gate: indeed, its left child requires a_1 to be 1, whereas its right child requires a_1 to be 0. The right child of the output gate is a decomposable \wedge -gate: indeed, for its left child g_1 we have $\text{Vars}(g_1) = \{a_1\}$, whereas for its right child g_2 we have $\text{Vars}(g_2) = \{a_2, a_3, a_4, a_5, a_6, a_7\}$, and these are indeed disjoint. The reader can easily check that all other \vee -gates are deterministic, and that all other \wedge -gates are decomposable.

4.1 Algorithm

The main result of this section is then the following.

PROPOSITION 4.4. *Given as input a deterministic and decomposable circuit C representing $\text{ELin}(q, D_n, D_x)$ for a database $D = D_n \cup D_x$ and Boolean query q , and an endogenous fact $f \in D_n$, we can compute in polynomial time (in $|C|$) the value $\text{Shapley}(q, D_n, D_x, f)$.*

Next, we prove Proposition 4.4 and present the algorithm, and then explain in Section 4.2 the architecture of the implementation.

Proof of Proposition 4.4. Let C be a deterministic and decomposable circuit representing $\text{ELin}(q, D_n, D_x)$, and let $f \in D_n$. First, we complete the circuit C so that all variables of D_n appear in C . Indeed, it could be the case that $\text{Vars}(C) \subsetneq D_n$: this happens for instance with the deterministic and decomposable circuit in Figure 2, where the endogenous fact a_8 does not appear in the circuit. To do this, we conjunct C with the conjunction $\bigwedge_{f' \in D_n \setminus \text{Vars}(C)} (f' \vee \neg f')$. Note that this does not change the semantics of the circuit (as this conjunction always evaluates to 1) and that the resulting circuit is still deterministic and decomposable. Now, let C_1 (resp., C_2) be the Boolean circuit obtained from C by replacing all variable gates corresponding to the fact f by a constant 1-gate (resp., by a constant 0-gate). Observe then that the variables of C_1 and C_2 are exactly $D_n \setminus \{f\}$, and moreover that C_1 and C_2 are still deterministic and decomposable. By definition of the endogenous lineage, we can rewrite Equation (2) into the following.

²We allow unbounded-fanin \wedge - and \vee -gates, and also allow constant 1-gates and constant 0-gates as, respectively, \wedge -gates with no input and \vee -gates with no inputs.

$$\text{Shapley}(q, D_n, D_x, f) = \sum_{k=0}^{|D_n|-1} \frac{k!(|D_n| - k - 1)}{|D_n|} \left(\# \text{SAT}_k(C_1) - \# \text{SAT}_k(C_2) \right). \quad (3)$$

Proposition 4.4 will thus directly follow from the next lemma.

LEMMA 4.5. *Given as input a deterministic and decomposable Boolean circuit C and an integer $k \in \{0, \dots, |\text{Vars}(C)|\}$, we can compute in polynomial time the quantity $\# \text{SAT}_k(C)$.*

PROOF. Our proof is similar to that of [3, Section 3.2]. Let $X \stackrel{\text{def}}{=} \text{Vars}(C)$ and $n \stackrel{\text{def}}{=} |X|$. First of all, we preprocess C so that the fanin of every \vee - and \wedge -gate is exactly 0 or 2; this can simply be done by rewriting every \wedge -gate of fanin $m > 2$ with $m - 1$ \wedge -gates of fanin 2 (same for \vee -gates), and adding a constant gate of the appropriate type to every \vee - and \wedge -gate of fan-in 1. We then compute, for every gate g of C , the set of variables $\text{Vars}(g)$ upon which the value of g depends. For a gate g of C , let us denote by φ_g the Boolean function over the variables $\text{Vars}(g)$ that is represented by this gate. For a gate g and an integer $\ell \in \{0, \dots, |\text{Vars}(g)|\}$, we define $\alpha_g^\ell \stackrel{\text{def}}{=} \# \text{SAT}_\ell(\varphi_g)$, i.e. the number of assignments of size ℓ to $\text{Vars}(g)$ that satisfy φ_g . We will show how to compute all the values α_g^ℓ for every gate g of C and $\ell \in \{0, \dots, |\text{Vars}(g)|\}$ in polynomial time. This will conclude the proof since, for the output gate g_{output} of C , we have that $\alpha_{g_{\text{output}}}^k = \# \text{SAT}_k(f)$. We will need the following notation: for two disjoint sets of variables X_1, X_2 and two subsets $S_1 \subseteq 2^{X_1}, S_2 \subseteq 2^{X_2}$ of assignments to X_1 and X_2 , we denote by $S_1 \otimes S_2 \subseteq 2^{X_1 \cup X_2}$ the set of assignments of $X_1 \cup X_2$ defined by $S_1 \otimes S_2 \stackrel{\text{def}}{=} \{v_1 \cup v_2 \mid v_1 \in S_1, v_2 \in S_2\}$. We next show how to compute the values α_g^ℓ by bottom-up induction on C .

Variable gate. If g is a variable gate corresponding to some variable y , then $\text{Vars}(g) = \{y\}$. Then, α_g^0 is 0 and α_g^1 is 1.

\neg -gate. If g is a \neg -gate with input gate g' , then $\alpha_g^\ell = \binom{|\text{Vars}(g')|}{\ell} - \alpha_{g'}^\ell$ for every $\ell \in \{0, \dots, |\text{Vars}(g')|\}$.

Deterministic \vee -gate. If g is a deterministic \vee -gate with no input then φ_g is the Boolean function on variables $\text{Vars}(g) = \emptyset$ that is always false, hence $\alpha_g^0 = 0$. Otherwise g has exactly two input gates; let us denote them g_1 and g_2 . Observe that $\text{Vars}(g) = \text{Vars}(g_1) \cup \text{Vars}(g_2)$ by definition. Define $S_1 \stackrel{\text{def}}{=} \text{Vars}(g_2) \setminus \text{Vars}(g_1)$ and similarly $S_2 \stackrel{\text{def}}{=} \text{Vars}(g_1) \setminus \text{Vars}(g_2)$. Since g is deterministic, we have:

$$\text{SAT}(\varphi_g) = (\text{SAT}(\varphi_{g_1}) \otimes 2^{S_1}) \cup (\text{SAT}(\varphi_{g_2}) \otimes 2^{S_2})$$

with the union being disjoint. By intersecting with the assignments of $\text{Vars}(g)$ of size ℓ , we obtain:

$$\text{SAT}_\ell(\varphi_g) = [(\text{SAT}(\varphi_{g_1}) \otimes 2^{S_1}) \cap \{v \subseteq \text{Vars}(g) \mid |v| = \ell\}] \cup [(\text{SAT}(\varphi_{g_2}) \otimes 2^{S_2}) \cap \{v \subseteq \text{Vars}(g) \mid |v| = \ell\}]$$

with again the middle union being disjoint, therefore:

$$\# \text{SAT}_\ell(\varphi_g) = |(\text{SAT}(\varphi_{g_1}) \otimes 2^{S_1}) \cap \{v \subseteq \text{Vars}(g) \mid |v| = \ell\}| + |(\text{SAT}(\varphi_{g_2}) \otimes 2^{S_2}) \cap \{v \subseteq \text{Vars}(g) \mid |v| = \ell\}|$$

We now explain how to compute the first term, that is, $|\text{SAT}(\varphi_{g_1}) \otimes 2^{S_1} \cap \{v \subseteq \text{Vars}(g) \mid |v| = \ell\}|$; the second term is similar. This is equal³ to

$$\sum_{i=\max(0, \ell-|S_1|)}^{\min(\ell, |\text{Vars}(g_1)|)} \alpha_{g_1}^i \times \binom{|S_1|}{\ell-i}.$$

Decomposable \wedge -gate. If g is a decomposable \wedge -gate with no input then φ_g is the Boolean function on variables $\text{Vars}(g) = \emptyset$ that is always true, hence $\alpha_g^0 = 1$. Otherwise, let g_1 and g_2 be the two input gates of g . Since g is decomposable we have $\text{Vars}(g) = \text{Vars}(g_1) \cup \text{Vars}(g_2)$ with the union being disjoint. But then we have:

$$\text{SAT}(\varphi_g) = \text{SAT}(\varphi_{g_1}) \otimes \text{SAT}(\varphi_{g_2})$$

We now intersect with the set of assignments of $\text{Vars}(g)$ of size ℓ to obtain

$$\alpha_g^\ell = \#\text{SAT}_\ell(\varphi_g) = \sum_{i=\max(0, \ell-|\text{Vars}(g_2)|)}^{\min(\ell, |\text{Vars}(g_1)|)} \alpha_{g_1}^i \times \alpha_{g_2}^{\ell-i}$$

This concludes the proof of the lemma, as well as the proof of Proposition 4.4. \square

Algorithm. Algorithm 1 depicts the solution underlying Proposition 4.4. The subroutine `ComputeAll#SATk` takes as input a d-D circuit C and outputs all the values $\#\text{SAT}_0(C), \dots, \#\text{SAT}_{|\text{Vars}(C)|}(C)$. This function computes values α_g^ℓ by bottom-up induction on C just as in the proof of Lemma 4.5, by using the appropriate equations depending on the type of each gate. Then, Lines 1–5 in the algorithm simply follow the part of the proof that starts at the beginning of this section until Lemma 4.5. For instance, the returned value on Line 5 corresponds to Equation (3). A quick inspection of Algorithm 1 reveals that, if one ignores the complexity of performing arithmetic operations (i.e., considering that additions and multiplications take constant time), the running time is $O(|C| \cdot |D_n|^2)$. If one wishes to compute the Shapley value of every endogenous fact (as will be done in the experiments), then the overall complexity is $O(|C| \cdot |D_n|^3)$. Last, we point out that, in the case of non-Boolean queries, this cost is incurred for each potential output tuple that one wants to analyze.

4.2 Implementation Architecture

In this section, we present our architecture for implementing the knowledge compilation approach over realistic datasets. The relevant parts, for now, are the middle and top part of Figure 3, which we next explain. Given a database $D = D_x \cup D_n$, a query $q(\bar{x})$, a tuple \bar{t} of the same arity as \bar{x} , and an endogenous fact $f \in D_n$, we want to compute $\text{Shapley}(q[\bar{x}/\bar{t}], D_x, D_n, f)$. We use two existing tools to help us with this task: `ProvSQL` [31] and the knowledge compiler `c2d` [10]. `ProvSQL` is a tool integrated into `PostgreSQL` that can perform provenance (lineage) computation in various semirings. For our purposes, a knowledge compiler is a tool that takes as input a Boolean function in CNF and outputs an equivalent Boolean function into another formalism. The target formalism that we will use is the so-called “d-DNNF”. A d-DNNF is simply a deterministic

³This comes from the fact that, for disjoint X_1, X_2 and assignments v_1 of X_1 and v_2 of X_2 we have $|v_1 \cup v_2| = |v_1| + |v_2|$, and $\text{Vars}(\varphi_{g_1})$ and S_1 are disjoint.

Algorithm 1: Shapley values from deterministic and decomposable Boolean circuits

Input : Deterministic and decomposable Boolean circuit C with output gate g_{output} representing $\text{ELin}(q, D_x, D_n)$ and an endogenous fact $f \in D_n$.

Output: The value $\text{Shapley}(q, D_x, D_n, f)$.

- 1 Complete C so that $\text{Vars}(g_{\text{output}}) = D_n$;
- 2 Compute $C_1 = C[f \rightarrow 1]$ and $C_2 = C[f \rightarrow 0]$;
// Partial evaluations of C by setting f to 1 and to 0
- 3 $\Gamma = \text{ComputeAll\#SAT}_k(C_1)$; // As an array
- 4 $\Delta = \text{ComputeAll\#SAT}_k(C_2)$; // As an array
- 5 **return** $\sum_{k=0}^{|D_n|-1} \frac{k! (|D_n| - k - 1)!}{|D_n|!} \cdot (\Gamma[k] - \Delta[k])$;
- 6 **Def** `ComputeAll#SATk(C)`:
- 7 | Preprocess C so that each \vee -gate and \wedge -gate has fan-in exactly 0 or 2;
- 8 | Compute the set $\text{Vars}(g)$ for every gate g in C ;
- 9 | Compute values α_g^ℓ for every gate g in C and $\ell \in \{0, \dots, |\text{Vars}(g)|\}$ by bottom-up induction on C using the inductive relations from the proof of Lemma 4.5;
- 10 | **return** $[\alpha_{g_{\text{output}}}^0, \dots, \alpha_{g_{\text{output}}}^{|\text{Vars}(C)|}]$;

and decomposable Boolean circuit such that negation gates are only applied to variables (NNF stands for *negation normal form*).⁴

In our case, we use `ProvSQL` as follows: we feed it the database D , query $q(\bar{x})$ and tuple \bar{t} , and `ProvSQL` computes $\text{Lin}(q[\bar{x}/\bar{t}], D)$ as a Boolean circuit, called C in Figure 3. We note here that for SPJU queries, C can be computed in polynomial-time data complexity. We then set to 1 all the exogenous facts to obtain a Boolean circuit C' for $\text{ELin}(q[\bar{x}/\bar{t}], D_x, D_n)$. Then, ideally, we would like to use the knowledge compiler to transform C' into an equivalent d-DNNF, in order to be able to apply Algorithm 1. Unfortunately, every knowledge compiler that we are aware of takes as input Boolean formulas in conjunctive normal form (CNF), and not arbitrary Boolean circuits. To circumvent it, we use the *Tseytin transformation* [35] to transform the circuit C' into a CNF $\varphi \stackrel{\text{def}}{=} \text{TSEYTIM}(C')$, whose size is linear in that of C' . This CNF φ has the following properties: (1) its variables are the variables of C' plus a set Z of additional variables; (2) for every valuation $v \subseteq \text{Vars}(C')$ that satisfies C' , there exists exactly one valuation $v' \subseteq Z$ such that $\varphi(v \cup v') = 1$; and (3) for every valuation $v \subseteq \text{Vars}(C')$ that does not satisfy C' , there is no valuation $v' \subseteq Z$ such that $\varphi(v \cup v') = 1$. We then

⁴This additional NNF restriction is not important here, but, as far as we know, no knowledge compiler has the more general “deterministic and decomposable circuits” (without NNF) as a target. It is currently unknown whether d-DNNFs and deterministic and decomposable circuits are exponentially separated or not [11, Table 7].

feed φ to the knowledge compiler, which produces a d-DNNF C'' equivalent to φ (the variables of C'' are again $\text{Vars}(C') \cup Z$). We note here that there is no theoretical guarantee that this step is efficient; indeed, the task of transforming a CNF into an equivalent d-D circuit is $\text{FP}^{\#P}$ -hard in general; see Section 6 for an experimental analysis of its tractability in practice. Next, we need to eliminate the additional variables Z in order to be able to apply Algorithm 1. To this end, we use the following Lemma.

LEMMA 4.6. *Given as input a d-DNNF C'' that is equivalent to $\text{TSEYTIM}(C')$ for a Boolean circuit C' , we can compute in time $O(|C''|)$ a d-DNNF C''' that is equivalent to C' (in particular, the variables of C''' are the same as the variables of C' ; in our case, they consist only of endogenous facts).*

PROOF SKETCH. Let Z be the additional variables coming from the Tseytin transformation. First, we remove all the gates of C'' that are not satisfiable, and then we remove all the gates that are not connected to the output gate. Now, let C''' be the circuit that is obtained from this intermediate circuit by replacing every literal z or $\neg z$ for $z \in Z$ by a constant 1-gate. We return C''' . The proof that this algorithm is correct uses the properties (1–3) of the Tseytin transformation, and is omitted due to lack of space. \square

Using this lemma, we obtain a d-DNNF C''' for the endogenous lineage $\text{ELin}(q[\bar{x}/\bar{t}], D_x, D_n)$, to which we can finally apply Algorithm 1 to obtain the value $\text{Shapley}(q[\bar{x}/\bar{t}], D_x, D_n, f)$.

5 INEXACT COMPUTATION

As we will show in the experimental section, the exact computation algorithm that we have proposed performs well in most cases but is too costly in others. In the latter cases, we may wish to resort to methods that do not necessarily compute exact Shapley values, if their results still typically suffice to determine the *order* of facts according to their Shapley contribution. In this section we propose *CNF Proxy*, a heuristic solution that is very efficient, and we will experimentally show that the ranking of facts based on CNF Proxy tends to match the ranking based on the exact Shapley values.

At a high level, CNF Proxy is based on the observation that having a high Shapley score is correlated (albeit in a complex manner) with (1) appearing many times in the provenance and (2) having few “alternatives,” that is, facts that could compensate for the absence of the given fact. The first factor (number of occurrences) may be directly read from the CNF obtained by applying the Tseytin transformation to the provenance circuit (C' in Figure 3). It is also easy to read from the CNF partial information about the second factor (number of alternatives), namely the number of alternatives in each clause (ignoring intricate dependencies between clauses). Next, we present the details of CNF Proxy.

We will start with an auxiliary definition, denoting the Shapley value of a general function $h : 2^X \rightarrow \mathbb{R}$ and a variable $x \in X$ as

$$\text{Shapley}(h, x) \stackrel{\text{def}}{=} \sum_{S \subseteq X \setminus \{x\}} \frac{|S|!(|X| - |S| - 1)!}{|X|!} (h(S \cup \{x\}) - h(S)).$$

Naturally, if $h = \text{ELin}(q, D_n, D_x)$, i.e., the endogenous lineage, and x is a fact in D_n , then $\text{Shapley}(h, x) = \text{Shapley}(q, D_n, D_x, x)$.

Algorithm 2: CNF Proxy

Input : CNF φ and a set of endogenous facts D_n .
Output: The value $\text{Shapley}(\tilde{\varphi}, x)$ for each $x \in D_n$.

```

1  $n \leftarrow |\varphi.\text{clauses}()$ ;
2  $v \leftarrow 0^{|D_n|}$ ; // As an array
3 for  $\psi \in \varphi.\text{clauses}()$  do
4    $\mathcal{L} \leftarrow \psi.\text{literals}()$ ;
5    $m \leftarrow |\mathcal{L}|$ ;
6    $\text{pos} \leftarrow \{\ell \in \mathcal{L} \mid \ell \text{ is positive}\}$ ;
7    $\text{neg} \leftarrow \{\ell \in \mathcal{L} \mid \ell \text{ is negative}\}$ ;
8   for  $\ell \in \text{pos} \cap D_n$  do
9      $v[\ell.\text{var}()] \leftarrow v[\ell.\text{var}()] + \frac{1}{nm \cdot \binom{m-1}{|\text{neg}|}}$ ;
10  end
11  for  $\ell \in \text{neg} \cap D_n$  do
12     $v[\ell.\text{var}()] \leftarrow v[\ell.\text{var}()] - \frac{1}{nm \cdot \binom{m-1}{|\text{pos}|}}$ ;
13  end
14 end
15 return  $v$ 

```

Now, note that for a CNF formula $\varphi = \bigwedge_{i=1}^n \psi_i$ (where each ψ_i is a disjunction of literals) and an assignment v it holds that $\varphi(v) = \prod_{i=1}^n \psi_i(v)$. Instead of calculating the Shapley values of a CNF formula φ (which may be a hard problem), CNF Proxy computes Shapley values with respect to a *proxy function*, denoted $\tilde{\varphi}$. The proxy function of φ is defined as the sum (instead of the product) of the clauses of φ , i.e., $\tilde{\varphi}(v) \stackrel{\text{def}}{=} \sum_{i=1}^n \psi_i(v)$. Intuitively, a fact that appears in many clauses of the CNF φ will occur in many summands of $\tilde{\varphi}$, and when we compute Shapley values with respect to $\tilde{\varphi}$, the number of alternatives in each clause will be reflected in decreased value of the respective summands.

Example 5.1. Consider the CNF formula $\varphi = (x_1 \vee x_2) \wedge (x_1 \vee x_3 \vee x_4)$. The Shapley values of x_1, x_2, x_3, x_4 are $\frac{7}{12}, \frac{3}{12}, \frac{1}{12}, \frac{1}{12}$ respectively. Note that x_1 has the highest influence, which intuitively may be attributed to its appearance in two clauses whereas each other variable appears only in a single clause. The variable x_2 has more influence than x_3 and x_4 , intuitively since it has less alternatives. These comparative features are preserved in $\tilde{\varphi} = (x_1 \vee x_2) + (x_1 \vee x_3 \vee x_4)$, and indeed the Shapley values of x_1, x_2, x_3, x_4 with respect to $\tilde{\varphi}$ are $\frac{5}{6}, \frac{1}{2}, \frac{1}{3}, \frac{1}{3}$ respectively. Observe that although the values assigned to the variables are very different from their actual Shapley values, their order remains intact in this case.

Due to the linearity of Shapley values, their computation with respect to $\tilde{\varphi}$ is much more efficient, as implied by the following lemma (whose proof we omit for space reasons):

LEMMA 5.2. *Let $h = \sum_{i=1}^n \frac{1}{n} \psi_i$, where each ψ_i is a Boolean function representing a disjunction of literals. Without loss of generality let us assume that for each ψ_i there is no variable that appears in more than one literal of ψ_i . Denote by a_i and b_i the number of positive and negative literals in ψ_i respectively. Then, for every variable x , it holds*

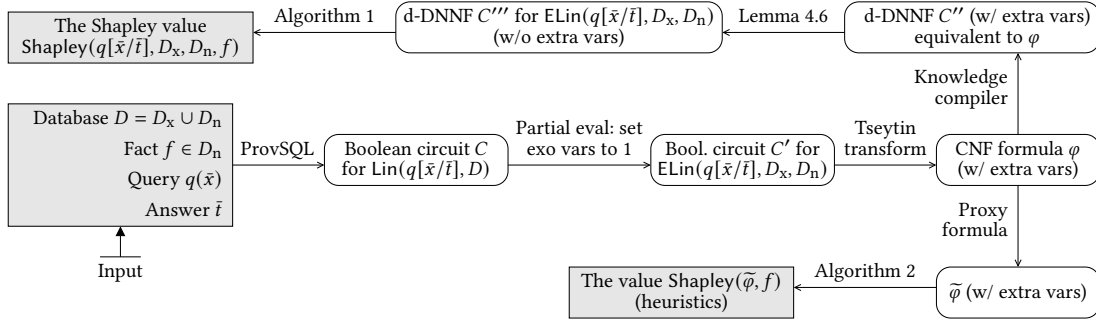


Figure 3: Our implementation architecture.

that $\text{Shapley}(h, x) = \frac{1}{n} \sum_{i=1}^n \Phi(\psi_i, x)$, where

$$\Phi(\psi_i, x) = \begin{cases} \frac{1}{(a_i+b_i) \cdot \binom{a_i+b_i-1}{b_i}} & \text{if } x \text{ appears in } \psi_i \text{ in positive form;} \\ \frac{-1}{(a_i+b_i) \cdot \binom{a_i+b_i-1}{a_i}} & \text{if } x \text{ appears in } \psi_i \text{ in negative form;} \\ 0 & \text{otherwise.} \end{cases}$$

Algorithm 2 then describes the operation of CNF Proxy, computing Shapley values of $\tilde{\phi}$ according to Lemma 5.2. The input to the algorithm is a CNF formula ϕ and the set of endogenous facts D_n . In Line 1, CNF Proxy counts the number n of clauses of ϕ , and in Line 2 it initializes the contribution of every variable of ϕ to zero. Then, it iterates over the clauses (Line 3). For each clause, it counts the number m of literals, and identifies the set of positive and negative literals, pos and neg respectively (Lines 4–7). In Lines 9 and 12 we add (resp., subtract) quantities for each variable in a positive (resp., negative) literal according to Lemma 5.2 (note that m corresponds to $a_i + b_i$ as denoted in the lemma). Finally, in Line 15 the algorithm returns the contribution value of each fact $x \in D_n$ based on $\text{Shapley}(\tilde{\phi}, x)$. Observe that CNF Proxy runs in linear time in the size of ϕ (which itself, being the Tseytin transformation of the provenance circuit C' from Figure 3, is linear in C').

Example 5.3. Recall the queries q_1, q_2, q and their lineages depicted in Figure 1. Using endogenous lineages, we have:

- $\text{ELin}(q_1, D_n, D_x) = a_1$
- $\text{ELin}(q_2, D_n, D_x) = (a_2 \wedge a_4) \vee (a_2 \wedge a_5) \vee (a_3 \wedge a_4) \vee (a_3 \wedge a_5) \vee (a_6 \wedge a_7)$
- $\text{ELin}(q, D_n, D_x) = a_1 \vee (a_2 \wedge a_4) \vee (a_2 \wedge a_5) \vee (a_3 \wedge a_4) \vee (a_3 \wedge a_5) \vee (a_6 \wedge a_7)$

The lineage of q_1 is a CNF with a single variable, thus the contribution of a_1 to q_1 as computed by Algorithm 2 is 1, which is indeed equal to $\text{Shapley}(q_1, D_n, D_x, a_1)$. Applying the Tseytin transformation to the lineage of q_2 introduces 6 new variables ($\{z_i\}_{i=1}^6$) and results in the following equisatisfiable CNF:

$$\begin{aligned} & (z_1) \wedge (z_1 \vee \bar{z}_2) \wedge (z_1 \vee \bar{z}_3) \wedge (z_1 \vee \bar{z}_4) \wedge (z_1 \vee \bar{z}_5) \wedge (z_1 \vee \bar{z}_6) \wedge \\ & (\bar{z}_1 \vee z_2 \vee z_3 \vee z_4 \vee z_5 \vee z_6) \wedge \\ & (\bar{z}_2 \vee a_2) \wedge (\bar{z}_2 \vee a_4) \wedge (z_2 \vee \bar{a}_2 \vee \bar{a}_4) \wedge (\bar{z}_3 \vee a_2) \wedge (\bar{z}_3 \vee a_5) \wedge (z_3 \vee \bar{a}_2 \vee \bar{a}_5) \wedge \\ & (\bar{z}_4 \vee a_3) \wedge (\bar{z}_4 \vee a_4) \wedge (z_4 \vee \bar{a}_3 \vee \bar{a}_4) \wedge (\bar{z}_5 \vee a_3) \wedge (\bar{z}_5 \vee a_5) \wedge (z_5 \vee \bar{a}_3 \vee \bar{a}_5) \wedge \\ & (\bar{z}_6 \vee a_6) \wedge (\bar{z}_6 \vee a_7) \wedge (z_6 \vee \bar{a}_6 \vee \bar{a}_7) \end{aligned}$$

Algorithm 2 iterates over the above clauses, and computes the contribution of the endogenous facts D_n over the proxy function. Note that the facts D_n appear in clauses of two forms. The first form

is $(\bar{z}_j \vee a_i)$; appearance in this type of clause adds $\frac{1}{22 \cdot 2 \cdot \binom{1}{1}} = \frac{1}{44}$ to the contribution of a_i . The second form is $(z_j \vee \bar{a}_i \vee \bar{a}_h)$, which adds $\frac{-1}{22 \cdot 3 \cdot \binom{1}{1}} = \frac{-1}{132}$. Note that each of a_2, a_3, a_4, a_5 has two appearances in clauses of the first form, and one appearance in clauses of the second form. Thus, according to Algorithm 2 the contribution of a_2, a_3, a_4, a_5 is $\frac{5}{132} \approx 0.038$. In contrast, a_6 and a_7 each have a single appearance in a clause of the first form and a single appearance in a clause of the second form, thus their contribution is $\frac{1}{66} \approx 0.015$. We note that the values calculated by Algorithm 2 are very different from the actual Shapley values, as $\text{Shapley}(q_2, D_n, D_x, a_i) = \frac{11}{60} \approx 0.183$ for $a_i \in \{a_2, a_3, a_4, a_5\}$ and $\text{Shapley}(q_2, D_n, D_x, a_i) = \frac{2}{15} \approx 0.133$ for $a_i \in \{a_6, a_7\}$. However, the facts a_2, a_3, a_4, a_5 , are correctly determined to be more influential than a_6 and a_7 .

Our experimental evaluation indicates that in most cases, the ordering of facts according to the values assigned to them by Algorithm 2 agrees with the order obtained by using the actual Shapley values. There is however no theoretical guarantee that this will always be the case, as shown by the following example.

Example 5.4. Applying the Tseytin transformation over the lineage of q will result in a CNF similar to that obtained for q_2 , with a new variable (z_7) and the new clauses $(z_1 \vee \bar{z}_7) \wedge (z_7 \vee \bar{a}_1) \wedge (\bar{z}_7 \vee a_1)$. In addition, the disjunct $\vee z_7$ will be added to the clause $(\bar{z}_1 \vee z_2 \vee z_3 \vee z_4 \vee z_5 \vee z_6)$. Similarly to the case of q_2 , the contributions of a_2, a_3, a_4 , and a_5 are correctly determined to be larger than those of a_6 and a_7 . As for a_1 , its contribution according to Algorithm 2 is 0 while in fact it is the most influential fact.

6 EXPERIMENTS

Our system is implemented in Python 3.6 and using the PostgreSQL 11.10 database engine, and the experiments were performed on a Linux Debian 14.04 machine with 1TB of RAM and an Intel(R) Xeon(R) Gold 6252 CPU @ 2.10GHz processor. ProvenSQL [31] was used to capture the provenance. For knowledge compilation we have used the c2d compiler [9, 10]. The source code of our implementation is available in [13].

Since no standard benchmark for our problem exists, we have created such a benchmark of 40 queries over the TPC-H (1.4GB) and IMDB (1.2GB) databases. The TPC-H queries are based on the ones in [34], where we have only removed nested queries (which ProvenSQL does not handle) and aggregation operations (for which

provenance is not Boolean). The queries for the IMDB database are based on the join queries in [19], where for each query we have added a (last) projection operation over one of the join attributes to make provenance more complex and thus more challenging for our algorithms. The resulting queries are quite complex: in particular, only 4 out of the 40 are hierarchical. See [13] for details of the obtained queries.

6.1 Exact computation

We have evaluated our solution for exact Shapley computation, presented in Section 4.2, on each of the 40 queries. In total, we have obtained 95,803 output tuples along with their provenance expressions (computed with ProVSQL). We have then transformed each provenance expression into a d-DNNF structure using the c2d [9, 10] knowledge compiler. In this experiment, for both the knowledge compilation (KC) and Shapley evaluation steps we have set a timeout of one hour (in fact, as we show below, a much shorter timeout of 2.5 seconds typically suffices). In case the compilation completed successfully within this timeframe, we have computed the Shapley values using Algorithm 1. Table 1 presents the execution times of our solution for 16 representative queries; next we will overview different aspects of the results.

Success rate. We report here the rate of successful executions. The IMDB queries resulted in 95,636 output tuples; the KC step completed successfully for 95,599 out of them, where all 37 failures were the result of insufficient memory. For each of the IMDB output tuples that were successfully compiled into a d-DNNF, we have executed Algorithm 1; only a single execution failed in this step (due to a timeout of one hour). Overall, the exact computation of Shapley values was successful for 95,598 out of the 95,636 IMDB output tuples (i.e., 99.96% success rate). The TPC-H queries resulted in 167 output tuples; the KC step has completed successfully for 141 out of them, again all 26 failures were the result of insufficient memory. For all TPC-H outputs that compiled successfully Algorithm 1 was successful, yielding an overall 84.43% success rate.

Execution time. For each of the 40 queries, we have measured the execution time of each step of the computation. First, we have measured (in the column “Execution time”) the execution time in PostgreSQL, which includes provenance generation for every output tuple using ProVSQL. Then, for each tuple \bar{t} in the output of the query, we have measured the KC execution time and the execution time of Algorithm 1 to compute the contribution of all input facts with respect to the output tuple \bar{t} . For the latter two algorithms, the execution times varied significantly for the different output tuples, and thus we report the execution times for different percentiles (mean, p25, p50, p75 and p99). Observe that the computation is typically efficient; outliers include q11d for which the execution time of Algorithm 1 was over 96 seconds in average.

Figure 4 depicts the running time of the KC step and of the computation of Shapley values from the d-DNNF as a function of different features of the provenance, such as the number of facts appearing in the provenance, the number of clauses in its CNF representation, and the number of gates in its d-DNNF representation.

Scalability. To evaluate the scalability of Algorithm 1 we have further looked at different scales of the TPC-H database. Figure 5

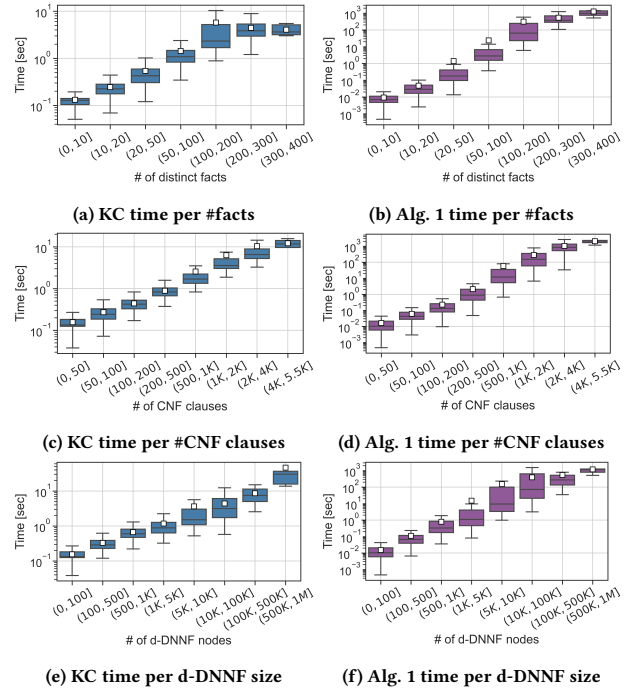


Figure 4: Running times of knowledge compilation and computation of Shapley values from the d-DNNF as function of the number of distinct facts, CNF clauses, and d-DNNF size

depicts the running time of Algorithm 1 over 8 representative query outputs as a function of the number of facts in the `LINEITEM` table. Figure 5a depicts 4 representative query outputs, for which the running time takes a few milliseconds, e.g., for “Q3 23426” (result 23426 of query Q3) and the full TPC-H dataset, the computation of Shapley values for all relevant input facts complete in 4.3ms. In contrast, Figure 5b depicts 4 query outputs for which the exact computation failed to complete over the full TPC-H database. We observe that the algorithm does succeed in these cases if we execute the queries over subsets of the input database, though its execution time may still be high: e.g., if we take a “slice” of the `LINEITEM` table consisting of 480,097 facts, then computation of the contribution of all input facts w.r.t. “Q9 ALGERIA” takes 556sec.

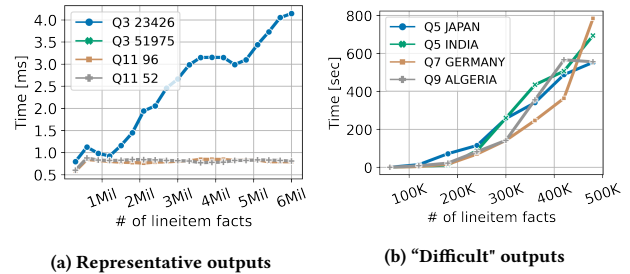


Figure 5: Alg. 1 running time for various TPC-H query outputs as function of table `LINEITEM` size

Table 1: Statistics on the exact computation of Shapley values for 16 representative queries

Dataset	Query	#Joined tables	#Filter conditions	Execution time [sec]	#Output tuples	Success rate	KC execution times [sec]					Alg. 1 execution times [sec]				
							Mean	p25	p50	p75	p99	Mean	p25	p50	p75	p99
TPC-H	3	3	5	20980.71	100	100%	0.06	0.04	0.07	0.08	0.13	0.00	0.00	0.00	0.00	0.01
	5	6	9	48.67	5	0%	-	-	-	-	-	-	-	-	-	-
	7	6	8	30.57	4	0%	-	-	-	-	-	-	-	-	-	-
	10	4	6	4.72	10	100%	0.14	0.13	0.13	0.14	0.19	0.01	0.01	0.01	0.01	0.01
	11	6	7	0.13	10	100%	0.09	0.08	0.08	0.08	0.13	0.01	0.00	0.00	0.01	0.03
	16	3	5	1.25	10	100%	0.26	0.13	0.18	0.33	0.57	0.18	0.01	0.03	0.29	0.88
	18	4	3	37.31	10	100%	0.13	0.08	0.08	0.18	0.23	0.01	0.00	0.01	0.01	0.03
19	2	21	2.04	1	100%	1.20	1.20	1.20	1.20	1.20	156.06	156.06	156.06	156.06	156.06	
IMDB	1a	5	10	0.25	35	100%	0.17	0.08	0.08	0.13	2.10	5.92	0.01	0.01	0.01	206.72
	6b	5	8	2.61	1	100%	0.44	0.44	0.44	0.44	0.44	0.08	0.08	0.08	0.08	0.08
	7c	8	21	77.33	2415	99%	0.82	0.18	0.28	0.63	9.44	24.28	0.02	0.05	0.39	787.12
	8d	7	10	145.10	44517	99.9%	0.26	0.13	0.18	0.29	1.09	0.36	0.01	0.02	0.05	2.28
	11a	8	18	3.20	10	100%	0.75	0.18	0.48	1.14	2.15	23.34	0.01	0.13	1.27	151.99
	11d	8	16	56.99	210	98.1%	0.99	0.19	0.32	0.78	5.97	96.19	0.02	0.06	0.48	1650.72
	13c	9	19	2.44	14	100%	0.22	0.13	0.18	0.28	0.53	0.02	0.01	0.01	0.01	0.06
	15d	9	18	24.25	207	97.6%	1.89	0.24	0.48	1.28	10.01	70.05	0.06	0.30	3.52	1821.13
	16a	8	15	5.56	173	100%	0.18	0.13	0.14	0.20	0.53	0.02	0.01	0.01	0.02	0.12

6.2 Inexact computation

As observed above, computing exact Shapley values using our solution for a given output tuple is typically fast, but may be costly or even fail in some cases. In this section we evaluate inexact computation alternatives.

Algorithms. We compare three algorithms: CNF Proxy (Section 5) and two existing baselines: Monte Carlo, and Kernel SHAP.

Monte Carlo. This is a well-known sampling algorithm [23] for approximating Shapley values in general. To employ the Monte Carlo algorithm in our setting, we feed it a provenance expression h containing n distinct input facts, and a budget of $r \cdot n$ samples, for some $r \in \mathbb{Z}_+$. The Shapley value of each fact f is approximated by sampling r permutations (π_1, \dots, π_r) of the input facts, and then outputting $\frac{1}{r} \cdot \sum_{i=1}^r (h(S_{\pi_i, < f} \cup \{f\}) - h(S_{\pi_i, < f}))$, where $S_{\pi_i, < f}$ is the coalition of all facts preceding f in the permutation π_i .

Kernel SHAP. Lundberg and Lee [22] have defined the notion of SHAP values in the context of ML explainability. Given a function $h : \mathbb{R}^d \rightarrow \mathbb{R}$ (the model whose decisions we want to explain), a probability distribution \mathcal{D} on \mathbb{R}^d (the inputs), and an input vector $\bar{e} \in \mathbb{R}^d$, SHAP values were defined to measure the contribution of \bar{e} 's features to the outcome $h(\bar{e})$. To overcome the issue that h does not operate over subsets of features, the notion of SHAP-score has been defined as follows

$$\text{SHAP}(h, \bar{e}, x) \stackrel{\text{def}}{=} \sum_{S \subseteq X \setminus \{x\}} \frac{|S|!(|X| - |S| - 1)!}{|X|!} (h_{\bar{e}}(S \cup \{x\}) - h_{\bar{e}}(S)),$$

where X is the set of d features, x is a specific feature whose contribution we wish to assess, and $h_{\bar{e}} : 2^X \rightarrow \mathbb{R}$ is defined by $h_{\bar{e}}(S) \stackrel{\text{def}}{=} \mathbb{E}_{\bar{z} \sim \mathcal{D}} [h(\bar{z}) \mid \bar{e}_S = \bar{z}_S]$, where \bar{e}_S and \bar{z}_S denote the vectors \bar{e} and \bar{z} restricted to the features in S .

In [22] the authors have proposed a method for approximating SHAP values, called Kernel SHAP. Kernel SHAP assumes feature independence and estimates the probability by multiplying the marginal distributions $\prod_{i \notin S} \Pr(z_i)$. The marginal probabilities $\Pr(z_i)$ in turn are estimated from a background data T . To approximate SHAP values, Kernel SHAP then samples m coalitions S_1, \dots, S_m of features, and trains a linear model $g : 2^X \rightarrow \mathbb{R}$ by minimizing the

Table 2: Median (resp., mean) performance. Monte Carlo and Kernel SHAP use 50 · #facts samples

	Monte Carlo	Kernel SHAP	CNF Proxy
Execution time	0.079 (1.875)	0.127 (1.978)	7e-4 (0.002)
L1	0.439 (0.448)	0.110 (0.109)	0.317 (0.315)
L2	0.03 (0.034)	0.001 (0.002)	0.010 (0.014)
nDCG	1.0 (0.992)	1.0 (0.998)	1.0 (0.999)
Precision@5	1.0 (0.955)	1.0 (0.961)	1.0 (0.989)
Precision@10	1.0 (0.953)	1.0 (0.961)	1.0 (0.968)

weighted loss $\sum_{i=1}^m w_i \cdot (g(S_i) - \hat{h}_{\bar{e}}(S_i))^2$, where w_i is proportional to the size of S_i and $\hat{h}_{\bar{e}}$ is the estimation of $h_{\bar{e}}$ using the feature independence assumption. The coefficient associated with a feature x in the trained model g is the approximated SHAP value of x .

We adapt Kernel SHAP to our setting and use it to approximate the Shapley values of facts, as follows. The features are the input facts of the database, we set h to be the Boolean function representing the (endogenous) provenance; the vector of interest, \bar{e} , has the value 1 for all facts, and the background data T contains a single example with value 0 in all its entries. Note that for such \bar{e} and T , Kernel SHAP estimates $h_{\bar{e}}(S)$ as the result of applying h to a vector with ones in S features and zeros in the remaining entries.

The input of both Monte Carlo and Kernel SHAP includes a budget of m samples; for provenance expressions containing n distinct facts we have experimented with $m \in \{10n, 20n, 30n, 40n, 50n\}$.

Accuracy metrics To evaluate the performance of the above methods we have used various metrics, specified below. All metrics were computed with respect to the ground truth values obtained by the knowledge compilation approach, and thus these experiments are confined to the cases where the exact computation succeeded.

- **nDCG** is the normalized discounted cumulative gain score [38], used to compare the ordering based on the inexact solution to the ordering based on the ground truth.
- **Precision@k** is the number of facts that appears in the top- k of both the inexact and exact solutions, divided by k . This was evaluated for $k \in \{1, 3, 5, 10\}$.
- **L1** and **L2** are the mean absolute error and squared error, respectively, of the results of an inexact computation method

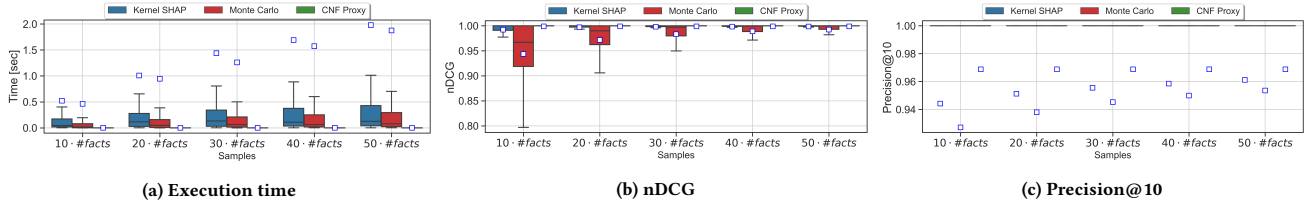


Figure 6: Comparing various metrics for inexact methods as a function of the sampling budget; CNF Proxy does not rely on sampling, thus remains constant for all budgets.

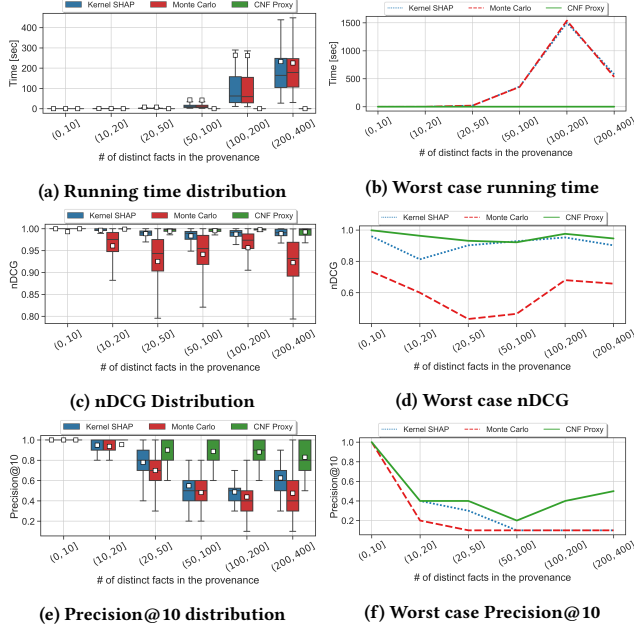


Figure 7: Comparing inexact methods as a function of the number of distinct facts in the provenance (n). Sampling methods (Kernel SHAP and Monte Carlo) presented with sampling budget of $m = 20n$; different budgets led to similar trends.

with respect to the ground truth (i.e. how different the results are from the actual Shapley values).

Execution time. Figure 6a depicts the execution time of the above methods as a function of the sampling budget. The execution times of Monte Carlo and Kernel SHAP are rather similar, while the CNF Proxy method (which does not rely on sampling) is substantially faster, with a median of 0.72 milliseconds and a mean of 2.06 milliseconds for a single query output. Figures 7a and 7b depict the distribution and worst-case running times of the three methods (for Monte Carlo and Kernel SHAP we used a budget of 20 samples per fact) as function of the number of distinct facts in the provenance expressions. CNF Proxy is substantially faster than its competitors: in most cases CNF Proxy completes in few milliseconds and 4 seconds in the worst case, whereas Monte Carlo and Kernel SHAP median execution time for circuits with 101–200 distinct input facts are 59 and 62 seconds respectively and may take up to 1,539 seconds.

Recall our analysis of the execution time of the exact solution in Section 6.1. The computation of Shapley values (KC plus Algorithm 1) takes 27,600 seconds if performed for all output tuples of query q8d in the IMDB dataset; for comparison, the CNF Proxy method has (inexactly, see quality analysis below) computed the Shapley value of the proxy functions for all query outputs within 95 seconds, which is 0.3% of the exact computation time. Surprisingly, the execution time of the exact computation was comparable to Monte Carlo and Kernel SHAP, and even faster for large sampling budgets; For example, Kernel SHAP with $m = 10n$ (where n is the number of distinct facts in the provenance) has completed computation for all output tuples of the query q8d in 16,447 seconds, which is 59% of the exact computation time. For $m = 50n$, Kernel SHAP required 58,161 seconds for completion of this computation, which is 210% of the exact computation time. This means that using such a sampling budget for Kernel SHAP is impractical in this setting, and we will use it to obtain upper bounds on Kernel SHAP’s quality.

Quality analysis. Figures 6b and 6c depict the ranking quality of the above methods as a function of the sampling budget. The rankings were compared for output tuples where the exact computation succeeded, so we have the ground truth. Recall that CNF Proxy does not rely on sampling, and so it remains constant throughout the different budgets. Naturally, as the sampling budget grows, Monte Carlo and Kernel SHAP quality improves. Most notable is Monte Carlo improvement in terms of nDCG (Figure 6b), where its median (resp., mean) nDCG with budget of 10 samples per fact is 0.9669 (resp., 0.9435), while with budget of 50 samples per fact it is 1.0 (resp., 0.9923). Comparison between the two sampling methods (Monte Carlo and Kernel SHAP) reveals that Kernel SHAP is superior w.r.t all metrics, across the entire budgets range. In terms of nDCG, all methods achieve rather high scores, but CNF Proxy performs best. Indeed, the median (resp., mean) nDCG of CNF Proxy is 1.0 (resp., 0.9989), while Kernel SHAP requires 50 samples per fact to get a median (resp., mean) nDCG of 1.0 (resp., 0.9986); as previously noted such budget leads to slower execution than that of the exact computation. It is worth mentioning that also when looking at nDCG@k for $k \in \{1, 3, 5, 10\}$, CNF Proxy performs better than Monte Carlo and Kernel SHAP, even when allowing a budget of 50 samples per fact. In terms of identifying the top influential facts (i.e., Precision@k), CNF Proxy also outperforms the other methods. For example, CNF Proxy’s median (resp., mean) Precision@10 are 1.0 (resp., 0.9688), while Kernel SHAP reaches 1.0 (resp., 0.9611) with 50 samples per fact. Similarly, CNF Proxy outperforms Monte Carlo and Kernel SHAP in terms of Precision@k for $k \in \{1, 3, 5\}$.

Finally, Table 2 zooms in on the results when fixing the budgets of Monte Carlo and Kernel SHAP to 50 samples per fact, which

is the highest budget tested (already for this budget, computation of Monte Carlo and Kernel SHAP is slower than for the exact algorithm). Note that CNF Proxy is much faster than the other methods, while still being superior in terms of ranking (nDCG and Precision@k). As explained above, ranking of facts is indeed the use case we recommend for CNF Proxy. Unsurprisingly, Kernel SHAP achieves better distance from the exact Shapley values (L1 and L2), at the cost of being slower by several orders of magnitude.

Dependency on the provenance size. Figure 7 depicts the performance of Monte Carlo, Kernel SHAP, and CNF Proxy as a function of the number of facts in the provenance expression. The results are aggregated over all output tuples of all queries. Figure 7c presents the quality of facts ranking (nDCG) as a function of the number of distinct provenance facts. CNF Proxy performs the best, and its quality remains steady regardless of the number of facts in the provenance expression. For example, with 1–10 facts the CNF Proxy median (resp., mean) nDCG is 1.0 (resp., 0.9999), and with 201–400 facts it is 0.9977 (resp., 0.9924). Kernel SHAP has a minor deterioration, where it drops from median (resp., mean) nDCG of 1.0 (resp., 0.9998) with 1–10 facts to 0.9906 (resp., 0.9888) with 201–400 facts. Figure 7d zooms in on the aggregated results over the worst case expressions of all output tuples for all queries, and shows that even the worst case CNF Proxy is superior to the alternatives, and that the error in terms of nDCG is small (0.92 in the worst case evaluated). Figure 7e depicts the dependency of Precision@10 on the number of provenance facts. Both Monte Carlo and Kernel SHAP suffer from a massive drop of median (resp., mean) Precision@10, down to 0.4 (resp., 0.476) and 0.6 (resp., 0.6253) respectively with 201–400 facts, while CNF Proxy remains at 0.8 (resp., 0.8293) median and mean Precision@10. A similar trend is observed for Precision@5, whereas for Precision@3 and Precision@1 the drop is less significant. Here again, Figure 7f zooms in on the worst case and again shows the superiority of CNF Proxy.

6.3 Hybrid computation

Recall that in Section 6.1 we have measured the success rate of the exact computation: for IMDB the success rate was 99.96% and for TPC-H the success rate was 84.43%. We saw in Section 6.2 that the inexact method CNF Proxy is very efficient and that the ranking of tuples based on CNF Proxy is typically close to the ranking obtained based on the real Shapley values.

In this section we consider a hybrid approach that works as follows. First, we start by running the exact computation, that is, the knowledge compilation step and Algorithm 1. If the exact computation completes successfully within less than t seconds (where t is configurable) we return its result. Otherwise, we terminate the exact computation and execute the inexact method CNF Proxy, returning only a ranking of input tuples rather than their Shapley values. Figure 8a depicts the success rate of the exact computation given different timeouts. Note that given a timeout of 2.5 seconds, the exact computation succeed for 98.67% of the IMDB output tuples, and 83.83% for TPC-H. Increasing the timeout has a rather minor impact on the success rate: having a 15 seconds timeout increases the success rate for IMDB to 99.52%, while the success rate for TPC-H remains unchanged. Recall that having a timeout of one hour results in success rates of 99.96% and 84.43% for IMDB and

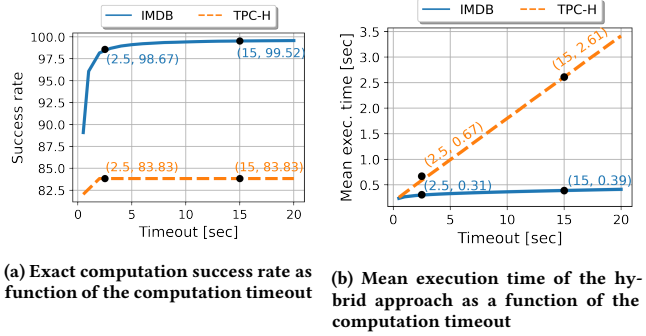


Figure 8: Hybrid approach performance

TPC-H respectively. Figure 8b depicts the mean execution time of the hybrid approach as a function of the chosen timeout t . Observe that given a timeout of 2.5 seconds the mean hybrid execution time is 0.31 seconds and 0.67 seconds for IMDB and TPC-H respectively. The mean execution time of the hybrid approach grows very moderately w.r.t. the timeout for IMDB (since most cases do not reach the timeout); it grows faster for TPC-H, where the difficult cases for which timeout is reached have a more significant effect on the overall mean execution time.

Main conclusions. Our experimental results indicate that for most output tuples (98.67% for IMDB and 83.83% for TPC-H) exact computation of Shapley values terminates within 2.5 seconds. When it does not, we propose an alternative of ranking facts according to CNF Proxy values, which typically only takes several milliseconds (up to 4 seconds for an outlier case). Experimental evidence shows that the ranking obtained via CNF Proxy is both much faster to compute and more accurate (in terms of nDCG and Precision@k, measured in cases where exact computation does succeed, and so we have the ground truth) than the alternative of only using sampling to approximate the actual Shapley values.

7 CONCLUSION

We have proposed in this paper a first practical framework for computing the contributions of database facts in query answering, quantified through Shapley values. The framework includes an exact algorithm that computes the contribution of input facts, and a faster algorithm that is practically effective in ranking contributions of input facts, while producing inexact Shapley values. Our practical implementation is currently designed for SPJU queries (that is, to the class of queries supported by ProVSQL). In addition to these practical contributions, we have also established a theoretical connection between the problem of computing Shapley values and that of probabilistic query evaluation, by showing that, for every query, the former can be reduced in polynomial time to the latter. We leave it open to determine whether there is also a reduction in the other direction (Open Problem 1). Other interesting directions would be to study further constructs such as aggregates and negation, or to extend the framework to bag semantics. Concerning bag semantics we observe that, by differentiating each copy of a same tuple in a bag database (for instance, adding an identifier attribute), our framework can be used as-is. Nevertheless, it would be interesting to see how one could adapt the definitions in order to consider fact multiplicities in a more elaborate way.

ACKNOWLEDGMENTS

This research has been partially funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 804302). The work of Benny Kimelfeld was supported by the Israel Science Foundation (ISF), Grant 768/19, and the German Research Foundation (DFG) Project 412400621 (DIP program).

REFERENCES

- [1] Serge Abiteboul, Richard Hull, and Victor Vianu. 1995. *Foundations of Databases*. Vol. 8. Addison-Wesley Reading. <http://webdam.inria.fr/Alice/>
- [2] Marcelo Arenas, Pablo Barceló, Leopoldo Bertossi, and Mikaël Monet. 2021. On the complexity of SHAP-score-based explanations: Tractability via knowledge compilation and non-approximability results. *arXiv preprint* (2021). <https://arxiv.org/abs/2104.08015>
- [3] Marcelo Arenas, Pablo Barceló, Leopoldo Bertossi, and Mikaël Monet. 2021. The tractability of SHAP-score-based explanations over deterministic and decomposable Boolean circuits. In *Proceedings of AAAI*. <https://arxiv.org/abs/2007.14045>
- [4] Peter Buneman, James Cheney, Wang-Chiew Tan, and Stijn Vansummeren. 2008. Curated databases. In *Proceedings of PODS*. 1–12. <https://homepages.inf.ed.ac.uk/opb/papers/inv.pdf>
- [5] Peter Buneman, Sanjeev Khanna, and Tan Wang-Chiew. 2001. Why and where: A characterization of data provenance. In *ICDT*. Springer, 316–330. https://repository.upenn.edu/cgi/viewcontent.cgi?article=1209&context=cis_papers
- [6] Yingwei Cui, Jennifer Widom, and Janet L. Wiener. 2000. Tracing the lineage of view data in a warehousing environment. *ACM Transactions on Database Systems (TODS)* 25, 2 (2000), 179–227. <http://ilpubs.stanford.edu:8090/252/1/1997-3.pdf>
- [7] Nilesh Dalvi and Dan Suciu. 2007. Efficient query evaluation on probabilistic databases. *Vldb J.* 16, 4 (2007), 523–544. <https://homes.cs.washington.edu/~suciu/vldb-j-probdb.pdf>
- [8] Nilesh Dalvi and Dan Suciu. 2013. The dichotomy of probabilistic inference for unions of conjunctive queries. *Journal of the ACM (JACM)* 59, 6 (2013), 1–87. <https://homes.cs.washington.edu/~suciu/jacm-dichotomy.pdf>
- [9] Adnan Darwiche. 2001. On the tractable counting of theory models and its application to truth maintenance and belief revision. *J. Applied Non-Classical Logics* 11, 1-2 (2001). <https://arxiv.org/abs/cs/0003044>
- [10] Adnan Darwiche. 2004. New advances in compiling CNF to decomposable negation normal form. In *Proceedings of ECAI*. Citeseer, 328–332. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.178.2262>
- [11] Adnan Darwiche and Pierre Marquis. 2002. A knowledge compilation map. *Journal of Artificial Intelligence Research* 17 (2002), 229–264. <https://arxiv.org/abs/1106.1819>
- [12] Daniel Deutch, Nave Frost, and Amir Gilad. 2020. Explaining natural language query results. *The VLDB Journal* 29, 1 (2020), 485–508. <https://arxiv.org/abs/2007.04454>
- [13] Daniel Deutch, Nave Frost, Benny Kimelfeld, and Mikaël Monet. 2021. Shapley for database facts source code. <https://github.com/navefr/ShapleyForDbFacts>.
- [14] Todd J Green, Grigoris Karvounarakis, and Val Tannen. 2007. Provenance semirings. In *Proceedings of PODS*. 31–40. https://repository.upenn.edu/cgi/viewcontent.cgi?article=1022&context=db_research
- [15] Todd J Green and Val Tannen. 2017. The semiring framework for database provenance. In *Proceedings of PODS*. 93–99. <https://dl.acm.org/doi/10.1145/3034786.3056125>
- [16] Anthony Hunter and Sébastien Konieczny. 2010. On the measure of conflicts: Shapley inconsistency values. *Artificial Intelligence* 174, 14 (2010), 1007–1026. <http://www.cril.univ-artois.fr/~konieczny/papers/aij10a.pdf>
- [17] Tomasz Imielinski and Witold Lipski Jr. 1984. Incomplete Information in Relational Databases. *J. ACM* 31, 4 (1984), 761–791. <https://doi.org/10.1145/1634.1886>
- [18] Abhay Jha and Dan Suciu. 2013. Knowledge compilation meets database theory: compiling queries to decision diagrams. *Theory of Computing Systems* 52, 3 (2013), 403–440. <https://link.springer.com/article/10.1007/s00224-012-9392-5>
- [19] Viktor Leis, Andrey Gubichev, Atanas Mirchev, Peter Boncz, Alfons Kemper, and Thomas Neumann. 2015. How good are query optimizers, really? *Proceedings of the VLDB Endowment* 9, 3 (2015), 204–215. <https://www.vldb.org/pvldb/vol9/p204-leis.pdf>
- [20] Ester Livshits, Leopoldo E. Bertossi, Benny Kimelfeld, and Moshe Sebag. 2020. The Shapley value of tuples in query answering. In *ICDT*, Vol. 155. Schloss Dagstuhl, 20:1–20:19. <https://arxiv.org/abs/1904.08679>
- [21] Scott M Lundberg, Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. 2020. From local explanations to global understanding with explainable AI for trees. *Nature machine intelligence* 2, 1 (2020), 2522–5839. <https://arxiv.org/pdf/1905.04610.pdf>
- [22] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Advances in neural information processing systems*. 4765–4774. <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>
- [23] Irwin Mann and LS Shapley. 1960. Values for large games, IV: Evaluating the Electoral College by Monte Carlo Techniques. The Rand Corporation. *Research Memorandum* 2651 (1960). https://www.rand.org/pubs/research_memoranda/RM2651.html
- [24] Alexandra Meliou, Wolfgang Gatterbauer, Katherine F. Moore, and Dan Suciu. 2010. The complexity of causality and responsibility for query answers and non-answers. *PVLDB* 4, 1 (2010), 34–45. <https://www.vldb.org/pvldb/vol4/p34-meliou.pdf>
- [25] Alexandra Meliou, Sudeepa Roy, and Dan Suciu. 2014. Causality and explanations in databases. *Proceedings of the VLDB Endowment (PVLDB)* 7, 13 (2014), 1715–1716. <http://www.vldb.org/pvldb/vol7/p1715-meliou.pdf>
- [26] Mikaël Monet. 2020. Solving a Special Case of the Intensional vs Extensional Conjecture in Probabilistic Databases. In *Proceedings of PODS*. 149–163. <https://arxiv.org/abs/1912.11864>
- [27] Alon Reshef, Benny Kimelfeld, and Ester Livshits. 2020. The impact of negation on the complexity of the Shapley value in conjunctive queries. In *Proceedings of PODS*. 285–297. <https://arxiv.org/abs/1912.12610>
- [28] Alvin E Roth. 1988. *The Shapley Value: Essays in Honor of Lloyd S. Shapley*. Cambridge University Press. <http://www.library.fu.ru/files/Roth2.pdf>
- [29] Sudeepa Roy, Laurel J. Orr, and Dan Suciu. 2015. Explaining query answers with explanation-ready databases. *Proceedings of the VLDB Endowment (PVLDB)* 9, 4 (2015), 348–359. <http://www.vldb.org/pvldb/vol9/p348-roy.pdf>
- [30] Babak Salimi, Leopoldo E. Bertossi, Dan Suciu, and Guy Van den Broeck. 2016. Quantifying causal effects on query answering in databases. In *TaPP*. USENIX Association. <http://web.cs.ucla.edu/~guyvdb/papers/SalimiTaPP16.pdf>
- [31] Pierre Senellart, Louis Jachiet, Silviu Maniu, and Yann Ramusat. 2018. Provsq: Provenance and probability management in postgresql. *Proceedings of the VLDB Endowment (PVLDB)* 11, 12 (2018), 2034–2037. <https://hal.inria.fr/hal-01851538/file/p976-senellart.pdf>
- [32] Lloyd S Shapley. 1953. A value for n-person games. *Contributions to the Theory of Games* 2, 28 (1953), 307–317. <http://www.library.fu.ru/files/Roth2.pdf#page=39>
- [33] Dan Suciu, Dan Olteanu, Christopher Ré, and Christoph Koch. 2011. *Probabilistic Databases*. Morgan & Claypool. <https://www.morganclaypool.com/doi/abs/10.2200/S00362ED1V01Y201105DTM016>
- [34] Transaction Processing Performance Council (TPC). 2017. TPC-H benchmark. http://www.tpc.org/tpc_documents_current_versions/pdf/tpc-h_v2.17.2.pdf
- [35] Grigori S Tseitin. 1983. On the complexity of derivation in propositional calculus. In *Automation of reasoning*. Springer, 466–483. https://link.springer.com/chapter/10.1007/978-3-642-81955-1_28
- [36] Guy Van den Broeck, Anton Lykov, Maximilian Schleich, and Dan Suciu. 2021. On the tractability of shap explanations. In *Proceedings of AAAI*. <https://arxiv.org/abs/2009.08634>
- [37] Moshe Y Vardi. 1982. The complexity of relational query languages. In *STOC*. ACM, 137–146. <http://www.dis.unroma1.it/~degiacom/didattica/semingsoft/SIS05-06/materiale/1-query-congiuntive/riferimenti/vardi-1982.pdf>
- [38] Yining Wang, Liwei Wang, Yuanzhi Li, Di He, Wei Chen, and Tie-Yan Liu. 2013. A theoretical analysis of NDCG ranking measures. In *Proceedings of COLT*, Vol. 8. 6. <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.680.490&rep=rep1&type=pdf>
- [39] Bruno Yun, Srđjan Vesic, Madalina Croitoru, and Pierre Bisquert. 2018. Inconsistency Measures for Repair Semantics in OBDA. In *IJCAI*. ijcai.org, 1977–1983. <https://www.ijcai.org/proceedings/2018/0273.pdf>