



HAL
open science

On computing probabilistic abductive explanations

Yacine Izza, Xuanxiang Huang, Alexey Ignatiev, Nina Narodytska, Martin Cooper, Joao Marques-Silva

► **To cite this version:**

Yacine Izza, Xuanxiang Huang, Alexey Ignatiev, Nina Narodytska, Martin Cooper, et al.. On computing probabilistic abductive explanations. *International Journal of Approximate Reasoning*, 2023, 159, pp.108939. 10.1016/j.ijar.2023.108939 . hal-04311424

HAL Id: hal-04311424

<https://hal.science/hal-04311424v1>

Submitted on 6 Jan 2025

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.

On Computing Probabilistic Abductive Explanations

Yacine Izza

*University of Toulouse, Toulouse, France
Monash University, Melbourne, Australia*

YACINE.IZZA@UNIV-TOULOUSE.FR

Xuanxiang Huang

University of Toulouse, Toulouse, France

XUANXIANG.HUANG@UNIV-TOULOUSE.FR

Alexey Ignatiev

Monash University, Melbourne, Australia

ALEXEY.IGNATIEV@MONASH.EDU

Nina Narodytska

VMware Research, Palo Alto, CA, U.S.A.

NNARODYTSKA@VMWARE.COM

Martin C. Cooper

IRIT, UPS, Toulouse, France

MARTIN.COOPER@IRIT.FR

Joao Marques-Silva

IRIT, CNRS, Toulouse, France

JOAO.MARQUES-SILVA@IRIT.FR

Abstract

The most widely studied explainable AI (XAI) approaches are unsound. This is the case with well-known model-agnostic explanation approaches, and it is also the case with approaches based on saliency maps. One solution is to consider intrinsic interpretability, which does not exhibit the drawback of unsoundness. Unfortunately, intrinsic interpretability can display unwieldy explanation redundancy. Formal explainability represents the alternative to these non-rigorous approaches, with one example being PI-explanations. Unfortunately, PI-explanations also exhibit important drawbacks, the most visible of which is arguably their size. Recently, it has been observed that the (absolute) rigor of PI-explanations can be traded off for a smaller explanation size, by computing the so-called relevant sets. Given some positive δ , a set \mathcal{S} of features is δ -relevant if, when the features in \mathcal{S} are fixed, the probability of getting the target class exceeds δ . However, even for very simple classifiers, the complexity of computing relevant sets of features is prohibitive, with the decision problem being NP^{PP} -complete for circuit-based classifiers. In contrast with earlier negative results, this paper investigates practical approaches for computing relevant sets for a number of widely used classifiers that include Decision Trees (DTs), Naive Bayes Classifiers (NBCs), and several families of classifiers obtained from propositional languages. Moreover, the paper shows that, in practice, and for these families of classifiers, relevant sets are easy to compute. Furthermore, the experiments confirm that succinct sets of relevant features can be obtained for the families of classifiers considered.

Contents

1	Introduction	3
2	Preliminaries	5
2.1	Logic Foundations	5
2.2	Classification Problems	5
2.3	Families of Classifiers	6
2.4	Running Examples	8
2.5	Formal Explainability	9
2.6	δ -Relevant Sets	13
3	Relevant Sets – Probabilistic Abductive Explanations	14
3.1	Definitions of Probabilistic AXp’s	14
3.2	Computing Locally-Minimal PAXp’s	16
4	Probabilistic Explanations for Decision Trees	16
4.1	Path Probabilities for DTs	17
4.2	Computing Locally-Minimal PAXp’s for DT’s	18
4.3	Computing Minimum-Size PAXp’s for DTs	18
4.4	Deciding Whether a Locally-Minimal PAXp is a Plain PAXp for DTs	21
4.5	Instance-Based vs. Path-Based Explanations	22
5	Probabilistic Explanations for Naive Bayes Classifiers	23
5.1	Explaining NBCs in Polynomial Time	23
5.2	Counting Models of XLCs	25
5.3	Computing Locally-Minimal PAXp’s for NBCs	28
6	Probabilistic Explanations for Other Families of Classifiers	31
6.1	Propositional Classifiers	31
6.2	Graph-Based Classifiers	32
7	Experiments	33
7.1	Case Study 1: Decision Trees	33
7.2	Case Study 2: Naive Bayes Classifiers	36
7.3	Case Study 3: Graph-Based Classifiers	37
8	Related Work	39
9	Conclusions	39
	References	40

1. Introduction

The advances in Machine Learning (ML) in recent years motivated an ever increasing range of practical applications of systems of Artificial Intelligence (AI). Some uses of ML models are deemed *high-risk* given the impact that their operation can have on people [28]. (Other authors refer to *high-stakes* applications [84].) In some domains, with high-risk applications representing one key example, the deployment of AI systems is premised on the availability of mechanisms for explaining the often opaque operation of ML models [28]. The need for explaining the opaque operation of ML models motivated the emergence of eXplainable AI (XAI). Moreover, the rigor of explanations is a cornerstone to delivering trustworthy AI [64]. For example, rigorous explanations can be mathematically proven correct, and so a human decision maker can independently validate that an explanation is (logically) rigorous. The rigor of explanations is even more significant in high-risk AI systems, where human decision makers must be certain that provided explanations are sound.

Motivated by the importance of understanding the operation of blackbox ML models, recent years have witnessed a growing interest in XAI [72, 33, 86, 87, 71, 85]. The best-known XAI approaches can be broadly categorized as either model-agnostic methods, that include for example LIME [81], SHAP [60] and Anchor [82], or intrinsic interpretability [84, 71], for which the explanation is represented by the actual (interpretable) ML model. For specific ML models, e.g. neural networks, there are dedicated explainability approaches, including those based on saliency maps [91, 8]; these exhibit limitations similar to model-agnostic approaches [1, 55, 99, 92]. Moreover, intrinsic interpretability may not represent a viable option in some uses of AI systems. Also, it has been shown, both in theory and in practice, that intrinsic interpretable models (such as decision trees or sets) also require being explained [45, 41, 46]. On the other hand, model-agnostic methods, even if locally accurate, can produce explanations that are unsound [39], in addition to displaying several other drawbacks [16, 93, 57, 25]. Unsound explanations are hopeless whenever rigor is a key requirement; thus, model-agnostic explanations ought not be used in high-risk settings. Indeed, it has been reported [39] that an explanation \mathcal{X} can be consistent with different predicted classes. For example, for a bank loan application, \mathcal{X} might be consistent with an approved loan application, but also with a declined loan application. An explanation that is consistent with both a declined and an approved loan applications offers no insight to why one of the loan applications was declined. There have been recent efforts on rigorous XAI approaches [89, 43, 22, 7, 61, 13, 32], most of which are based on feature selection, namely the computation of so-called abductive explanations (AXp’s). However, these efforts have mostly focused on the scalability of computing rigorous explanations, with more recent work investigating input distributions [32]. Nevertheless, another important limitation of rigorous XAI approaches is the often unwieldy size of explanations. Recent work studied probabilistic explanations, as a mechanism to reduce the size of rigorous explanations [102, 101, 100]. Some approaches for computing probabilistic explanations have extended model-agnostic approaches [102], and so can suffer from unsoundness. Alternatively, more rigorous approaches to computing probabilistic explanations have been shown to be computationally hard, concretely hard for NP^{PP} [101, 100], and so are all but certain to fall beyond the reach of modern automated reasoners.




Figure 1: Summary of results

This paper builds on recent work [101, 100] on computing rigorous probabilistic explanations, and investigates their practical scalability. However, instead of considering classifiers represented as boolean circuits (as in [101]), the paper studies families of classifiers that are shown to be amenable to the practical computation of relevant sets. As the paper shows, such examples include decision trees (DTs), naive Bayes classifiers (NBCs), but also several propositional classifiers, graph-based classifiers, and their multi-valued variants ¹.

The paper revisits the original definition of δ -relevant set, and considers different variants, in addition to the one proposed in earlier work [101, 100], i.e. smallest relevant sets. These include subset-minimal relevant sets, and locally-minimal relevant sets. Throughout the paper, relevant sets will be referred to as *probabilistic abductive explanations* (PAXp’s). Whereas computing a smallest and a subset-minimal PAXp’s is shown to be in NP for decision trees and different propositional classifiers, computing locally-minimal PAXp’s is shown to be in P for decision trees and different propositional classifiers, and in pseudo-polynomial time for naive Bayes classifiers. Furthermore, the experiments confirm that locally-minimal PAXp’s most often match subset-minimal PAXp’s. As a result, locally-minimal relevant sets are shown to represent a practically efficient approach for computing (in polynomial or pseudo-polynomial time) PAXp’s that are most often subset-minimal. The paper’s contributions are summarized in Figure 1. Figure 1 overviews complexity class membership results for several families of classifiers, most of which are established in this paper. For several families of classifiers, the paper proves that computing one locally-minimal PAXp is either in polynomial time or pseudo-polynomial time. Similarly, computing one subset-

1. This paper aggregates and extends recent preprints that compute relevant sets for concrete families of classifiers [48, 50].

minimal (and also one cardinality-minimal) PAXp is shown to be in NP (or pseudo-NP) for several families of classifiers.

The paper is organized as follows. [Section 2](#) introduces the definitions and notation used throughout. [Section 3](#) offers an overview of probabilistic abductive explanations (or relevant sets). [Section 4](#) investigates the computation of relevant sets in the case of decision trees. [Sections 5](#) and [6](#) replicate the same exercise, respectively in the case of naive Bayes classifiers, and also for graph-based classifiers and classifiers based on propositional languages. [Section 7](#) presents experimental results on computing relevant sets for the classifiers studied in the earlier sections. The paper concludes in [Section 9](#).

2. Preliminaries

Complexity classes. The paper addresses a number of well-known classes of decision and function problems, that include P, NP, FP^{NP} , NP^{PP} , among others. The interested reader is referred to a standard reference on the computational complexity [\[5\]](#).

2.1 Logic Foundations

Propositional logic. The paper assumes the notation and definitions that are standard when reasoning about the decision problem for propositional logic, i.e. the Boolean Satisfiability (SAT) problem [\[12\]](#). SAT is well-known to be an NP-complete [\[17\]](#) decision problem. A propositional formula φ is defined over a finite set of Boolean variables $X = \{x_1, x_2, \dots, x_n\}$. Formulas are most often represented in *conjunctive normal form* (CNF). A CNF formula is a conjunction of clauses, a clause is a disjunction of literals, and a literal is a variable (x_i) or its negation ($\neg x_i$). A term is a conjunction of literals. Whenever convenient, a formula is viewed as a set of sets of literals. A Boolean interpretation μ of a formula φ is a total mapping of X to $\{0, 1\}$ (0 corresponds to *False* and 1 corresponds to *True*). Interpretations can be extended to literals, clauses and formulas with the usual semantics; hence we can refer to $\mu(l)$, $\mu(\omega)$, $\mu(\varphi)$, to denote respectively the value of a literal, clause and formula given an interpretation. Given a formula φ , μ is a *model* of φ if it makes φ *True*, i.e. $\mu(\varphi) = 1$. A formula φ is *satisfiable* ($\varphi \not\models \perp$) if it admits a model, otherwise, it is *unsatisfiable* ($\varphi \models \perp$). Given two formulas φ and ψ , we say that φ *entails* ψ (denoted $\varphi \models \psi$) if all models of φ are also models of ψ . φ and ψ are *equivalent* (denoted $\varphi \equiv \psi$) if $\varphi \models \psi$ and $\psi \models \varphi$.

First Order Logic (FOL) and SMT. When necessary, the paper will consider the restriction of FOL to Satisfiability Modulo Theories (SMT). These represent restricted (and often decidable) fragments of FOL [\[10\]](#). All the definitions above apply to SMT. A SMT-solver reports whether a formula is satisfiable, and if so, may provide a model of this satisfaction. (Other possible features include dynamic addition and retraction of constraints, production of proofs, and optimization.)

2.2 Classification Problems

This paper considers classification problems, which are defined on a set of features (or attributes) $\mathcal{F} = \{1, \dots, m\}$ and a set of classes $\mathcal{K} = \{c_1, c_2, \dots, c_K\}$. Each feature $i \in \mathcal{F}$ takes values from a domain \mathbb{D}_i . In general, domains can be categorical or ordinal, with values that can be boolean or integer. (Although real-valued could be considered for some of

the classifiers studied in the paper, we opt not to specifically address real-valued features.) Feature space is defined as $\mathbb{F} = \mathbb{D}_1 \times \mathbb{D}_2 \times \dots \times \mathbb{D}_m$; $|\mathbb{F}|$ represents the total number of points in \mathbb{F} . For boolean domains, $\mathbb{D}_i = \{0, 1\} = \mathbb{B}$, $i = 1, \dots, m$, and $\mathbb{F} = \mathbb{B}^m$. The notation $\mathbf{x} = (x_1, \dots, x_m)$ denotes an arbitrary point in feature space, where each x_i is a variable taking values from \mathbb{D}_i . The set of variables associated with features is $X = \{x_1, \dots, x_m\}$. Moreover, the notation $\mathbf{v} = (v_1, \dots, v_m)$ represents a specific point in feature space, where each v_i is a constant representing one concrete value from \mathbb{D}_i . With respect to the set of classes \mathcal{K} , the size of \mathcal{K} is assumed to be finite; no additional restrictions are imposed on \mathcal{K} . Nevertheless, with the goal of simplicity, the paper considers examples where $|\mathcal{K}| = 2$, concretely $\mathcal{K} = \{\ominus, \oplus\}$, or alternatively $\mathcal{K} = \{0, 1\}$. An ML classifier \mathbb{M} is characterized by a (non-constant) *classification function* κ that maps feature space \mathbb{F} into the set of classes \mathcal{K} , i.e. $\kappa : \mathbb{F} \rightarrow \mathcal{K}$. An *instance* (or observation) denotes a pair (\mathbf{v}, c) , where $\mathbf{v} \in \mathbb{F}$ and $c \in \mathcal{K}$, with $c = \kappa(\mathbf{v})$.

2.3 Families of Classifiers

Throughout the paper, a number of families of classifiers will be studied in detail. These include decision trees [14, 80], naive Bayes classifiers [29], graph-based classifiers [76, 77, 35] and classifiers based on propositional languages [23, 21].

Decision trees. A decision tree $\mathcal{T} = (V, E)$ is a directed acyclic graph, with $V = \{1, \dots, |V|\}$, having at most one path between every pair of nodes. \mathcal{T} has a root node, characterized by having no incoming edges. All other nodes have one incoming edge. We consider univariate decision trees where each non-terminal node is associated with a single feature x_i . Each edge is labeled with a literal, relating a feature (associated with the edge’s starting node) with some values (or range of values) from the feature’s domain. We will consider literals to be of the form $x_i \in \mathbb{E}_i$. x_i is a variable that denotes the value taken by feature i , whereas $\mathbb{E}_i \subseteq \mathbb{D}_i$ is a subset of the domain of feature $i \in \mathcal{F}$. The type of literals used to label the edges of a DT allows the representation of the DTs generated by a wide range of decision tree learners (e.g. [96]). The set of paths of \mathcal{T} is denoted by \mathcal{R} . $\Phi(R_k)$ denotes the set of features associated with path $R_k \in \mathcal{R}$, one per node in the tree, with repetitions allowed. It is assumed that for any $\mathbf{v} \in \mathbb{F}$ there exists *exactly* one path in \mathcal{T} that is consistent with \mathbf{v} . By *consistent* we mean that the literals associated with the path are satisfied (or consistent) with the feature values in \mathbf{v} . Given \mathbf{v} , the set of paths \mathcal{R} is partitioned into \mathcal{P} and \mathcal{Q} , such that each of the paths in \mathcal{P} yields the prediction $c = \kappa(\mathbf{v})$, whereas each of the paths in \mathcal{Q} yields a prediction in $\mathcal{K} \setminus \{c\}$. For the purposes of this paper, the path consistent with \mathbf{v} as $P_t \in \mathcal{P}$, i.e. the target path. (A more in-depth analysis of explaining decision trees is available in [46].)

Naive Bayes Classifiers (NBCs). An NBC [26] is a Bayesian Network model [29] characterized by strong conditional independence assumptions among the features. Given some observation $\mathbf{x} \in \mathbb{F}$, the predicted class is given by:

$$\kappa(\mathbf{x}) = \operatorname{argmax}_{c \in \mathcal{K}} (\Pr(c|\mathbf{x})) \quad (1)$$

Using Bayes’s theorem, $\Pr(c|\mathbf{x})$ can be computed as follows: $\Pr(c|\mathbf{x}) = \Pr(c, \mathbf{x})/\Pr(\mathbf{x})$. In practice, we compute only the numerator of the fraction, since the denominator $\Pr(\mathbf{x})$ is

constant for every $c \in \mathcal{K}$. Moreover, given the conditional mutual independency of the features, we have:

$$\Pr(c, \mathbf{x}) = \Pr(c) \times \prod_i \Pr(x_i|c)$$

Furthermore, it is also common in practice to apply logarithmic transformations on probabilities of $\Pr(c, \mathbf{x})$, thus getting:

$$\log \Pr(c, \mathbf{x}) = \log \Pr(c) + \sum_i \log \Pr(x_i|c)$$

Therefore, (1) can be rewritten as follows:

$$\kappa(\mathbf{x}) = \operatorname{argmax}_{c \in \mathcal{K}} \left(\log \Pr(c) + \sum_i \log \Pr(x_i|c) \right) \quad (2)$$

For simplicity, and following the notation used in earlier work [62], lPr denotes a logarithmic probability. Thus, we get:

$$\kappa(\mathbf{x}) = \operatorname{argmax}_{c \in \mathcal{K}} \left(\text{lPr}(c) + \sum_i \text{lPr}(x_i|c) \right) \quad (3)$$

(Note that also for simplicity, it is common in practice to add a sufficiently large positive constant T to each probability, which allows using only positive values.)

Graph-based classifiers. The paper considers the class of graph-based classifiers (referred to as Decision Graphs) studied in earlier work [35]. A Decision Graph (DG) is a Directed Acyclic Graph (DAG) consisting of two types of nodes, non-terminal nodes and terminal nodes, and has a single root node. Each non-terminal node is labeled with a feature and has at least one child node, each terminal node is labeled with a class and has no child node. Moreover, for an arbitrary non-terminal node p labeled with feature i , we use $\mathbb{C}_i \subseteq \mathbb{D}_i$ to denote the set of values of feature i which are consistent with any path connecting the root to p . Each outgoing edge of node p represents a literal of the form $x_i \in \mathbb{E}_i$ ($\mathbb{E}_i \neq \emptyset$) where $\mathbb{E}_i \subseteq \mathbb{C}_i$. Furthermore, the following constraints are imposed on DGs:

1. The literals associated with the outgoing edges of p represent a partition of \mathbb{C}_i .
2. Any path connecting the root node to a terminal node contains no inconsistent literals.
3. No dead-ends.

A DG is *read-once* if each feature tested at most once on any path. A DG is *ordered*, if features are tested in the same order on all paths. An Ordered Multi-valued Decision Diagrams (OMDD) is a *read-once* and *ordered* DG such that every \mathbb{E}_i is a singleton, which means multiple-edges between two nodes may exist. Furthermore, the OMDDs considered in this paper are *reduced* [53, 94], i.e.

1. No node p such that all child nodes of p are isomorphic; and
2. No isomorphic subgraphs.

When the domains of features are all boolean and the set of classes is binary, then OMDDs corresponds to Ordered Binary Decision Diagrams (OBDDs) [15].

Propositional languages & classifiers. For classification problems for which the feature space \mathbb{F} is restricted to \mathbb{B}^m , then boolean circuits can be used as binary classifiers. Each boolean circuit is a sentence of some propositional language. We briefly review some well-known propositional languages and queries/transformations that these languages support in polynomial time.

The language *negation normal form* (NNF) is the set of all directed acyclic graphs, where each terminal node is labeled with either \top , \perp , x_i or $\neg x_i$, for $x_i \in X$. Each non-terminal node is labeled with either \wedge (or AND) or \vee (or OR). The language *decomposable* NNF (DNNF) [19, 23] is the set of all NNFs, where for every node labeled with \wedge , $\alpha = \alpha_1 \wedge \dots \wedge \alpha_k$, no variables are shared between the conjuncts α_j . A *deterministic* DNNF (d-DNNF) [23, 73] is a DNNF, where for every node labeled with \vee , $\beta = \beta_1 \vee \dots \vee \beta_k$, each pair β_p, β_q , with $p \neq q$, is inconsistent, i.e. $\beta_p \wedge \beta_q \models \perp$. A *Smooth* d-DNNF (sd-DNNF) [20] is a d-DNNF, where for every node labeled with \vee , $\beta = \beta_1 \vee \dots \vee \beta_k$, each pair β_p, β_q is defined on the same set of variables. We focus in this paper on d-DNNF, but for simplicity of algorithms, sd-DNNF is often considered. Furthermore, *sentential decision diagrams* (SDDs) [21, 98] represent a well-known subset of the d-DNNF. (Furthermore, it should be noted that OBDD is a proper subset of SDD.) SDDs are based on a strongly deterministic decomposition [21], which is used to decompose a Boolean function into the form: $(p_1 \wedge s_1) \vee \dots \vee (p_n \wedge s_n)$, where each p_i is called a *prime* and each s_i is called a *sub* (both primes and subs are sub-functions). Furthermore, the process of decomposition is governed by a variable tree (*vtree*) which stipulates the variable order [21].

The languages d-DNNF, sd-DNNF and SDD satisfy the query *polytime model counting* (**CT**), and the transformation *polytime conditioning* (**CD**). Let Δ represent a propositional formula and let ρ denote a consistent term ($\rho \not\models \perp$). The *conditioning* [23] of Δ on ρ , denoted $\Delta|_\rho$ is the formula obtained by replacing each variable x_i by \top (resp. \perp) if x_i (resp. $\neg x_i$) is a positive (resp. negative) literal of ρ . A propositional language \mathbf{L} satisfies **CT** if there exists a polynomial-time algorithm that maps every formula Δ from \mathbf{L} into a non-negative integer denoting the number of models of Δ . \mathbf{L} satisfies **CD** iff there exists a polynomial-time algorithm that maps every formula Δ from \mathbf{L} and every consistent term ρ into a formula in \mathbf{L} that is logically equivalent to $\Delta|_\rho$. There are additional queries and transformations of interest [23], but these are beyond the goals of this paper. It is important to note that OMDD and OBDD also satisfy **CT** and **CD** [23, 75].

2.4 Running Examples

Example Decision Tree. Figure 2 shows the example DT used throughout the paper. This example DT also illustrates the notation used to represent DTs. The set of paths \mathcal{R} is partitioned into two sets \mathcal{P} and \mathcal{Q} , such that the paths in $\mathcal{P} = \{P_1, P_2, P_3\}$ yield a prediction of \oplus , and such that the paths in $\mathcal{Q} = \{Q_1, Q_2\}$ yield a prediction of \ominus . (In general, \mathcal{P} denotes the paths with prediction $c \in \mathcal{K}$, and \mathcal{Q} denotes the paths with prediction other than c , i.e. any class in $\mathcal{K} \setminus \{c\}$.)

Example Naive Bayes Classifier. Consider the NBC depicted graphically in Figure 3². The features are the boolean random variables R_1, R_2, R_3, R_4 and R_5 . Each R_i can take values **t** or **f** denoting, respectively, whether a listener likes or not that radio station. The boolean random variable G corresponds to an age class: the target class \oplus denotes the prediction that the listener is young and \ominus denotes the prediction that the listener is old. Thus, $\mathcal{K} = \{\ominus, \oplus\}$. Let us consider $\mathbf{v} = (R_1, R_2, R_3, R_4, R_5) = (\mathbf{t}, \mathbf{f}, \mathbf{f}, \mathbf{f}, \mathbf{t})$. We associate r_i to each literal ($R_i = \mathbf{t}$) and $\neg r_i$ to literals ($R_i = \mathbf{f}$). Using (3), we get the values shown

2. This example of an NBC is adapted from [62], and it was first studied in [9, Ch.10].

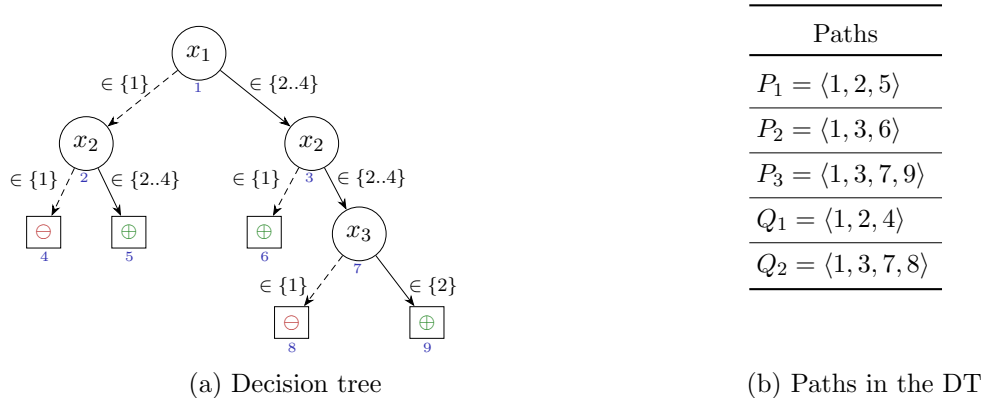


Figure 2: Example DT.




Figure 3: Example NBC.

in Figure 4. (Note that to use positive values, we added $T = +4$ to each $\text{lPr}(\cdot)$.) As can be seen by comparing the values of $\text{lPr}(\oplus|\mathbf{v})$ and $\text{lPr}(\ominus|\mathbf{v})$, the classifier will predict \oplus .

Example graph-based classifiers. Figure 5 shows two examples of graph-based classifiers. Figure 5a shows an OBDD, and Figure 5b an OMDD. Figure 5b represents a function defined on $\mathcal{F} = \{1, 2, 3\}$ and $\mathcal{K} = \{\ominus, \oplus, \otimes\}$, with the domains of features being $\mathbb{D}_1 = \{0, 1\}$, $\mathbb{D}_2 = \mathbb{D}_3 = \{0, 1, 2\}$. If we consider the instance $\mathbf{v} = \{1, 1, 2\}$, the classifier predicts the class \otimes .

2.5 Formal Explainability

In contrast with well-known model-agnostic approaches to XAI [81, 60, 82, 33], formal explanations are model-precise, i.e. their definition reflects the model’s computed function.

Abductive explanations. Prime implicant (PI) explanations [89] denote a minimal set of literals (relating a feature value x_i and a constant $v_i \in \mathbb{D}_i$) that are sufficient for the prediction. PI-explanations are related with abduction, and so are also referred to as

	$\Pr(\oplus)$	$\Pr(r_1 \oplus)$	$\Pr(\neg r_2 \oplus)$	$\Pr(\neg r_3 \oplus)$	$\Pr(\neg r_4 \oplus)$	$\Pr(r_5 \oplus)$	$\text{lPr}(\oplus \mathbf{v})$
$\Pr(\cdot)$	0.10	0.95	0.95	0.98	0.80	0.95	
$\text{lPr}(\cdot)$	1.70	3.95	3.95	3.98	3.78	3.95	21.31

 (a) Computing $\text{lPr}(\oplus|\mathbf{v})$

	$\Pr(\ominus)$	$\Pr(r_1 \ominus)$	$\Pr(\neg r_2 \ominus)$	$\Pr(\neg r_3 \ominus)$	$\Pr(\neg r_4 \ominus)$	$\Pr(r_5 \ominus)$	$\text{lPr}(\ominus \mathbf{v})$
$\Pr(\cdot)$	0.90	0.03	0.05	0.66	0.25	0.03	
$\text{lPr}(\cdot)$	3.89	0.49	1.00	3.58	2.61	0.49	12.06

 (b) Computing $\text{lPr}(\ominus|\mathbf{v})$


 Figure 4: Deciding prediction for $\mathbf{v} = (\mathbf{t}, \mathbf{f}, \mathbf{f}, \mathbf{f}, \mathbf{t})$. (Note that to use positive values, $T = +4$ was added to each $\text{lPr}(\cdot)$.)


Figure 5: Example DD.

abductive explanations (AXp's) [43]³. Formally, given $\mathbf{v} = (v_1, \dots, v_m) \in \mathbb{F}$ with $\kappa(\mathbf{v}) = c$, a set of features $\mathcal{X} \subseteq \mathcal{F}$ is a *weak abductive explanation* [18] (or weak AXp) if the following predicate holds true⁴:

$$\text{WeakAXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c) := \forall (\mathbf{x} \in \mathbb{F}). [\bigwedge_{i \in \mathcal{X}} (x_i = v_i)] \rightarrow (\kappa(\mathbf{x}) = c) \quad (4)$$

3. PI-explanations were first proposed in the context of boolean classifiers based on restricted bayesian networks [89]. Independent work [43] studied PI-explanations in the case of more general classification functions, i.e. not necessarily boolean, and related instead explanations with abduction. This paper follows the formalizations used in more recent work [62, 49, 41, 63, 35, 18, 34, 40, 64].

4. Each predicate associated with a given concept will be noted in sans-serif letterform. When referring to the same concept in the text, the same acronym will be used, but in standard letterform. For example, the predicate name AXp will be used in logic statements, and the acronym AXp will be used throughout the text.

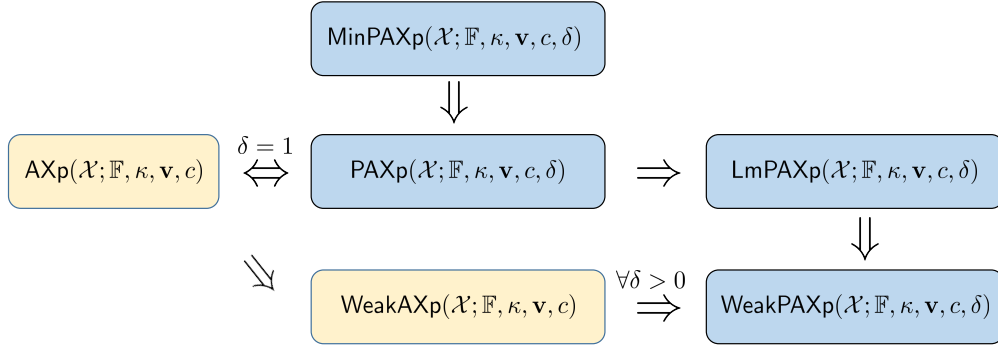


Figure 6: A schematic representation of relationships between different types of explanations. $E_1(\mathcal{X}; \dots) \Rightarrow E_2(\mathcal{X}; \dots)$ means that if \mathcal{X} is an E_1 explanation then \mathcal{X} must be an E_2 explanation.

Moreover, a set of features $\mathcal{X} \subseteq \mathcal{F}$ is an *abductive explanation* (or (plain) AXp) if the following predicate holds true:

$$\begin{aligned} \text{AXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c) &:= \text{WeakAXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c) \wedge \\ &\quad \forall(\mathcal{X}' \subsetneq \mathcal{X}). \neg \text{WeakAXp}(\mathcal{X}'; \mathbb{F}, \kappa, \mathbf{v}, c) \end{aligned} \quad (5)$$

Clearly, an AXp is any weak AXp that is subset-minimal (or irreducible). It is straightforward to observe that the definition of predicate `WeakAXp` is monotone, and so an AXp can instead be defined as follows:

$$\begin{aligned} \text{AXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c) &:= \text{WeakAXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c) \wedge \\ &\quad \forall(j \in \mathcal{X}). \neg \text{WeakAXp}(\mathcal{X} \setminus \{j\}; \mathbb{F}, \kappa, \mathbf{v}, c) \end{aligned} \quad (6)$$

This alternative equivalent definition of abductive explanation is at the core of most algorithms for computing one AXp. (Throughout the paper, we will drop the parameterization associated with each predicate, and so we will write `AXp`(\mathcal{X}) instead of `AXp`($\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c$), when the parameters are clear from the context.)

Example 1. The computation of (weak) AXp’s is illustrated with the DT from Figure 2a. The instance considered throughout is $\mathbf{v} = (v_1, v_2, v_3) = (4, 4, 2)$, with $c = \kappa(\mathbf{v}) = \oplus$. The point \mathbf{v} is consistent with P_3 , and $\Phi(P_3) = \{1, 2, 3\}$. Table 1 (columns 1 to 4) analyzes three sets of features $\{1, 2, 3\}$, $\{1, 3\}$ and $\{3\}$ in terms of being a weak AXp or an AXp. The decision on whether each set is a weak AXp or an AXp can be obtained by analyzing all the 32 points in feature space, or by using an off-the-shelf algorithm. (The analysis of all points in feature space is omitted for brevity.)

It is apparent that (4), (5), and (6) can be viewed as representing a (logic) *rule* of the form:

$$\mathbf{IF} \quad [\wedge_{i \in \mathcal{X}} (x_i = v_i)] \quad \mathbf{THEN} \quad [\kappa(\mathbf{x}) = c] \quad (7)$$

Unless otherwise noted, this interpretation of explanations will be assumed throughout the paper.

S	\mathcal{U}	WeakAXp?	AXp?	$\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c (\mathbf{x}_S = \mathbf{v}_S))$	WeakPAXp?	PAXp?	$\#(S)$	$\#(P_1)$	$\#(P_2)$	$\#(P_3)$	$\#(Q_1)$	$\#(Q_2)$
$\{1, 2, 3\}$	\emptyset	Yes	No	$1 \geq \delta$	Yes	No	1	0	0	1	0	0
$\{1, 3\}$	$\{2\}$	Yes	Yes	$1 \geq \delta$	Yes	No	4	0	1	3	0	0
$\{3\}$	$\{1, 2\}$	No	–	$15/16 = 0.9375 \geq \delta$	Yes	Yes	16	3	3	9	1	0

 Table 1: Examples of sets of fixed features given $\mathbf{v} = (4, 4, 2)$ and $\delta = 0.93$

Abductive explanations can be viewed as answering a ‘Why?’ question, i.e. why is some prediction made given some point in feature space. A different view of explanations is a contrastive explanation [70], which answers a ‘Why Not?’ question, i.e. which features can be changed to change the prediction. The formalization of contrastive explanations revealed a minimal hitting set duality relationship between abductive and contrastive explanations [42]. The paper does not detail further contrastive explanations, as the focus is solely on abductive explanations.

Figure 6 shows relationships between different classes of explanations that we investigate in the paper. $\text{AXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c)$ and $\text{WeakAXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c)$ are deterministic classes of explanations. These classes are shown in yellow boxes. Other classes of explanations, shown in blue boxes, represent probabilistic counterparts of AXp’s. We study them in the following sections.

Progress in formal explainability. The introduction of abductive explanations [89, 43] also revealed practical limitations in the case of bayesian network classifiers [89, 90] and neural networks [43]. However, since then there has been a stream of results, that demonstrate the practical applicability of formal explainability. These results can be broadly organized as follows (a more detailed overview is available in [64]):

- **Tractable explanations.**
Recent work showed that computing one explanation is tractable for naive Bayes classifiers [62], decision trees [45, 35, 46], graph-based classifiers [35], monotonic classifiers [63, 18], and classifiers represented with well-known classes of propositional languages [34]. Additional tractability results were obtained in [18].
- **Efficient explanations.**
For some other families of classifiers, recent work showed that computing one explanation is computationally hard, but it is nevertheless efficient in practice. This is the case with decision lists and sets [41], random forests [49], and tree ensembles in general [43, 39, 40].
- **Explainability queries.**
There has been interest in understanding the complexity of answering different queries related with reasoning about explainability [7, 36, 35, 6]. For example, the feature membership problem is the decision problem of deciding whether some (possibly sensitive) feature occurs in some explanation. Although computationally hard in general [35], it has been shown to be solved efficiently in theory and in practice for specific families of classifiers [35, 37]. Queries related with enumeration of explanations have been extensively studied [62, 42, 41, 63, 35, 34, 40].

- Properties of explanations.

A number of works studied the connections between explanations and robustness [44], and connections between different types of explanations [42].

Despite the observed progress, formal explainability still faces several important challenges. First, for some widely used families of classifiers, e.g. neural networks, formal explainability does not scale in practice [43]. Second, input distributions are not taken into account, since these are not readily available. There is however recent work on accounting for input constraints [18, 32, 104]. Third, the size of explanations may exceed the cognitive limits of human decision makers [69], and computing smallest explanations does not offer a computationally realistic alternative [43]. Recent work studied δ -relevant sets [101, 100], and these are also the focus of this paper.

Finally, we note that there have been different approaches to formal explainability based on the study of the formal logic or the axiomatics of explainers [103, 2, 59]. This paper studies exclusively those approaches for which there is practical supporting evidence of observed progress, as attested above.

2.6 δ -Relevant Sets

δ -relevant sets were proposed in more recent work [101, 100] as a generalized formalization of PI-explanations (or AXp’s). δ -relevant sets can be viewed as *probabilistic* PIs [100], with AXp’s representing a special case of δ -relevant sets where $\delta = 1$, i.e. probabilistic PIs that are actual PIs. We briefly overview the definitions related with relevant sets. The assumptions regarding the probabilities of logical propositions are those made in earlier work [101, 100]. Let $\Pr_{\mathbf{x}}(A(\mathbf{x}))$ denote the probability of some proposition A defined on the vector of variables $\mathbf{x} = (x_1, \dots, x_m)$, i.e.

$$\begin{aligned} \Pr_{\mathbf{x}}(A(\mathbf{x})) &= \frac{|\{\mathbf{x} \in \mathbb{F}: A(\mathbf{x})=1\}|}{|\{\mathbf{x} \in \mathbb{F}\}|} \\ \Pr_{\mathbf{x}}(A(\mathbf{x}) \mid B(\mathbf{x})) &= \frac{|\{\mathbf{x} \in \mathbb{F}: A(\mathbf{x})=1 \wedge B(\mathbf{x})=1\}|}{|\{\mathbf{x} \in \mathbb{F}: B(\mathbf{x})=1\}|} \end{aligned} \tag{8}$$

(Similar to earlier work, it is assumed that the features are independent and uniformly distributed [101]. Moreover, the definitions above can be adapted in case some of the features are real-valued. As noted earlier, the present paper studies only non-continuous features.)

Definition 1 (δ -relevant set [101]). Consider $\kappa : \mathbb{B}^m \rightarrow \mathcal{K} = \mathbb{B}$, $\mathbf{v} \in \mathbb{B}^m$, $\kappa(\mathbf{v}) = c \in \mathbb{B}$, and $\delta \in [0, 1]$. $\mathcal{S} \subseteq \mathcal{F}$ is a δ -relevant set for κ and \mathbf{v} if,

$$\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c \mid \mathbf{x}_{\mathcal{S}} = \mathbf{v}_{\mathcal{S}}) \geq \delta \tag{9}$$

(Where the restriction of \mathbf{x} to the variables with indices in \mathcal{S} is represented by $\mathbf{x}_{\mathcal{S}} = (x_i)_{i \in \mathcal{S}}$. Concretely, the notation $\mathbf{x}_{\mathcal{S}} = \mathbf{v}_{\mathcal{S}}$ represents the constraint $\bigwedge_{i \in \mathcal{S}} x_i = v_i$.)

(Moreover, observe that $\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c \mid \mathbf{x}_{\mathcal{S}} = \mathbf{v}_{\mathcal{S}})$ is often referred to as the *precision* of \mathcal{S} [82, 74].) Thus, a δ -relevant set represents a set of features which, if fixed to some pre-defined value (taken from a reference vector \mathbf{v}), ensures that the probability of the prediction being the same as the one for \mathbf{v} is no less than δ .

Definition 2 (Min- δ -relevant set). Given κ , $\mathbf{v} \in \mathbb{B}^m$, and $\delta \in [0, 1]$, find the smallest k , such that there exists $\mathcal{S} \subseteq \mathcal{F}$, with $|\mathcal{S}| = k$, and \mathcal{S} is a δ -relevant set for κ and \mathbf{v} .

With the goal of proving the computational complexity of finding a minimum-size set of features that is a δ -relevant set, earlier work [101] restricted the definition to the case where κ is represented as a boolean circuit. (Boolean circuits were restricted to propositional formulas defined using the operators \vee , \wedge and \neg , and using a set of variables representing the inputs; this explains the choice of *inputs* over *sets* in earlier work [101].) The main complexity result from earlier work is that the computation of δ -relevant sets is hard for NP^{PP} [101]. Hence, as noted in earlier work [101, 100], it is unlikely that exact computation of δ -relevant sets will be practically feasible.

3. Relevant Sets – Probabilistic Abductive Explanations

In contrast with Min- δ -relevant sets, whose focus are smallest-size explanations, this section investigates alternative definitions of relevant sets (which we will also and indistinguishably refer to as *probabilistic abductive explanations*).

3.1 Definitions of Probabilistic AXp’s

Conceptually, Definition 1 does not need to impose a restriction on the classifier considered (although this is done in earlier work [101]), i.e. the logical representation of κ need not be a boolean circuit. As a result, Definition 1 can also be considered in the case of multi-class classifiers defined on categorical or ordinal (non-continuous) features.

Given the above, a *weak probabilistic AXp* (or weak PAXp) is a pick of fixed features for which the conditional probability of predicting the correct class c exceeds δ , given $c = \kappa(\mathbf{v})$. Thus, $\mathcal{X} \subseteq \mathcal{F}$ is a weak PAXp if the following predicate holds true,

$$\begin{aligned} \text{WeakPAXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c, \delta) & \\ & := \Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c \mid \mathbf{x}_{\mathcal{X}} = \mathbf{v}_{\mathcal{X}}) \geq \delta \\ & := \frac{|\{\mathbf{x} \in \mathbb{F} : \kappa(\mathbf{x}) = c \wedge (\mathbf{x}_{\mathcal{X}} = \mathbf{v}_{\mathcal{X}})\}|}{|\{\mathbf{x} \in \mathbb{F} : (\mathbf{x}_{\mathcal{X}} = \mathbf{v}_{\mathcal{X}})\}|} \geq \delta \end{aligned} \tag{10}$$

which means that the fraction of the number of points predicting the target class and consistent with the fixed features (represented by \mathcal{X}), given the total number of points in feature space consistent with the fixed features, must exceed δ . (Observe that the difference to (1) is solely that features and classes are no longer required to be boolean. Hence, weak PAXp’s can be viewed as generalized δ -relevant sets.) Moreover, a set $\mathcal{X} \subseteq \mathcal{F}$ is a *probabilistic AXp* (or (plain) PAXp) if the following predicate holds true,

$$\begin{aligned} \text{PAXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c, \delta) & := \\ & \text{WeakPAXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c, \delta) \wedge \\ & \forall(\mathcal{X}' \subsetneq \mathcal{X}). \neg \text{WeakPAXp}(\mathcal{X}'; \mathbb{F}, \kappa, \mathbf{v}, c, \delta) \end{aligned} \tag{11}$$

Thus, $\mathcal{X} \subseteq \mathcal{F}$ is a PAXp if it is a weak PAXp that is also subset-minimal,

As can be observed, the definition of weak PAXp (see (10)) does not guarantee monotonicity. In turn, this makes the computation of (subset-minimal) PAXp’s harder. With the purpose of identifying classes of weak PAXp’s that are easier to compute, it will be

convenient to study *locally-minimal* PAXp’s. A set of features $\mathcal{X} \subseteq \mathcal{F}$ is a locally-minimal PAXp if,

$$\begin{aligned} \text{LmPAXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c, \delta) &:= \\ &\text{WeakPAXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c, \delta) \wedge \\ &\forall (j \in \mathcal{X}). \neg \text{WeakPAXp}(\mathcal{X} \setminus \{j\}; \mathbb{F}, \kappa, \mathbf{v}, c, \delta) \end{aligned} \quad (12)$$

As observed earlier in [Section 2.5](#), because the predicate WeakAXp is monotone, subset-minimal AXp’s match locally-minimal AXp’s. An important practical consequence is that most algorithms for computing one subset-minimal AXp, will instead compute a locally-minimal AXp, since these will be the same. Nevertheless, a critical observation is that in the case of probabilistic AXp’s (see (10)), the predicate WeakPAXp is *not* monotone. Thus, there can exist locally-minimal PAXp’s that are not subset-minimal PAXp’s. (As shown in the experiments, computed locally minimal APXp’s are most often PAXp’s. However, exceptions do exist, even though these are rarely observed). Furthermore, the fact that a set of features $\mathcal{X} \subseteq \mathcal{F}$ may satisfy (12) but not (11) imposes that subset-minimal PAXp’s must be computed by using (11); as shown later, this requires more complex algorithms.

Finally, minimum-size PAXp’s (or a smallest PAXp’s) generalize Min- δ -relevant sets in [Definition 2](#). A set of features $\mathcal{X} \subseteq \mathcal{F}$ is a minimum-size AXp if,

$$\begin{aligned} \text{MinPAXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c, \delta) &:= \\ &\text{WeakPAXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c, \delta) \wedge \\ &\forall (\mathcal{X}' \subseteq \mathcal{F}). [(|\mathcal{X}'| < |\mathcal{X}|) \rightarrow \neg \text{WeakPAXp}(\mathcal{X}'; \mathbb{F}, \kappa, \mathbf{v}, c, \delta)] \end{aligned} \quad (13)$$

(As stated earlier, throughout the paper, we will drop the parameterization associated with each predicate, and so we will write $\text{PAXp}(\mathcal{X})$ instead of $\text{PAXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c, \delta)$, when the parameters are clear from the context. Although the parameterization on δ is paramount, we opt instead for simpler notation.)

Example 2. The computation of (probabilistic) AXp’s is illustrated with the DT from [Figure 2](#). The instance considered throughout is $\mathbf{v} = (v_1, v_2, v_3) = (4, 4, 2)$, with $c = \kappa(\mathbf{v}) = \oplus$. Clearly, \mathbf{v} is consistent with P_3 . The goal is to compute a δ -relevant set given $\delta = 0.93$. Let $\#(R_k)$ denote the number of points in feature space that are consistent with path R_k . Moreover, let $\#(\mathcal{X})$ denote the total number of points in feature space that are consistent with the set of *fixed* features $\mathcal{X} \in \mathcal{F}$. [Table 1](#) summarizes the computation of $\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c | \mathcal{X}_{\mathcal{S}} = \mathbf{v}_{\mathcal{S}})$ for different sets \mathcal{S} . The table also includes information on whether each set is a weak AXp, an AXp, a weak PAXp, or a PAXp. The set $\{1, 3\}$ represents an AXp, since for any point consistent with the assignments $x_1 = 4$ and $x_3 = 2$, the prediction is \oplus . However, by setting $\mathcal{S} = \{3\}$, the probability of predicting \oplus given a point consistent with $x_3 = 2$ still exceeds δ , since $15/16 = 93.75\%$. Hence, $\{3\}$ is a PAXp for $\mathbf{v} = (4, 4, 2)$ when $\delta = 0.93$.

Properties of locally-minimal PAXp’s. Let \mathcal{X} denote an AXp. Clearly, \mathcal{X} is also a PAXp. Then, for any locally-minimal $\mathcal{A} \subseteq \mathcal{X}$ (i.e. \mathcal{A} is computed using \mathcal{X} as a seed), we have the following properties, which follow from the definition:

1. $\mathcal{A} \subseteq \mathcal{X}$ (by hypothesis);
2. \mathcal{A} is a weak PAXp (by definition); and
3. There exists at least one PAXp \mathcal{E} such that $\mathcal{E} \subseteq \mathcal{A}$.

Algorithm 1 Computing one locally-minimal PAXp

Input: Features $\{1, \dots, m\}$; feature space \mathbb{F} , classifier κ , instance (\mathbf{v}, c) , threshold δ
Output: Locally-minimal PAXp \mathcal{S}

```

1: procedure findLmPAXp( $\{1, \dots, m\}; \mathbb{F}, \kappa, \mathbf{v}, c, \delta$ )
2:    $\mathcal{S} \leftarrow \{1, \dots, m\}$ 
3:   for  $i \in \{1, \dots, m\}$  do
4:     if WeakPAXp( $\mathcal{S} \setminus \{i\}; \mathbb{F}, \kappa, \mathbf{v}, c, \delta$ ) then
5:        $\mathcal{S} \leftarrow \mathcal{S} \setminus \{i\}$ 
6:   return  $\mathcal{S}$ 

```

Thus, given some AXp \mathcal{X} , we can compute a locally-minimal PAXp \mathcal{A} that is both a subset of \mathcal{X} and a superset of some PAXp, and such that \mathcal{A} exhibits the strong probabilistic properties of relevant sets. Although any locally-minimal PAXp is a subset of a weak AXp, there can exist locally-minimal PAXp’s that are not subsets of some (plain) AXp.

3.2 Computing Locally-Minimal PAXp’s

Algorithm 1 shows one approach for computing a locally-minimal PAXp⁵. As shown, to compute one locally-minimal PAXp, one starts from $\mathcal{F} = \{1, \dots, m\}$ and iterately removes features while it is safe to do so, i.e. while (10) holds for the resulting set. Beside Algorithm 1, one could consider for example variants of the QuickXplain [52] and the Progression [65, 66] algorithms. Both of which also allow computing preferred (locally-minimal) sets subject to anti-lexicographic preferences [52, 67]. Furthermore, we note that the same algorithms (i.e. Deletion, Progression and QuickXplain, among others) can also be used for computing one AXp. Moreover, observe that these algorithms can also be applied to *any classifier* with respect to which we seek to compute one locally-minimal PAXp. Furthermore, another simple observation is that explanations can be enumerated by exploiting hitting set dualization [42], e.g. using a MARCO-like algorithm [58].

Practically efficient computation of relevant sets. Further to the computation of locally-minimal PAXp’s, the next few sections show that the computation of relevant sets (PAXp’s) can be achieved efficiently in practice, for several families of classifiers. Concretely, over the next few sections we analyze decision trees, naive Bayes classifiers, but also several families of propositional and graph-based classifiers, studied in recent work [35, 34].

4. Probabilistic Explanations for Decision Trees

This section shows that the problem of deciding whether a set $\mathcal{X} \subseteq \mathcal{F}$ is a PAXp is in NP when κ is represented by a decision tree⁶. As a result, a minimum-size PAXp can be

5. This simple algorithm is often referred to as the deletion-based algorithm, namely in settings related with solving function problems in propositional logic and constraint programming [66]. However, the same general algorithm can be traced at least to the work of Valiant [97], and some authors [51] argue that it is implicit in works from the 19th century [68].

6. As noted earlier, and for simplicity, the paper considers the case of non-continuous features. However, in the case of DTs, the results generalize to continuous features.

computed with at most a logarithmic number of calls to an NP oracle. (This is a consequence that optimizing a linear cost function, subject to a set of constraints for which deciding satisfiability is in NP, can be achieved with a logarithmic number of calls to a NP oracle.) An SMT formulation of the problem is proposed and the empirical evaluation confirms its practical effectiveness. This section also proposes a polynomial time algorithm to compute one locally-minimal PAXp, thus offering an alternative to computing one PAXp. The results in [Section 7](#) confirm that in practice computed locally-minimal PAXp’s are often subset-minimal, i.e. a locally-minimal PAXp actually represents a (plain) PAXp.

4.1 Path Probabilities for DTs

This section investigates how to compute, in the case of DTs, the conditional probability,

$$\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c \mid \mathbf{x}_{\mathcal{X}} = \mathbf{v}_{\mathcal{X}}) \tag{14}$$

where \mathcal{X} is a set of *fixed* features (whereas the other features are not fixed, being deemed *universal*), and P_t is a path in the DT consistent with the instance (\mathbf{v}, c) . (Also, note that (14) is the left-hand side of (9).) To motivate the proposed approach, let us first analyze how we can compute $\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c)$, where $\mathcal{P} \subseteq \mathcal{R}$ is the set of paths in the DT with prediction c . Let $\Lambda(R_k)$ denote the set of literals (each of the form $x_i \in \mathbb{E}_i$) in some path $R_k \in \mathcal{R}$. If a feature i is tested multiple times along path R_k , then \mathbb{E}_i is the intersection of the sets in each of the literals of R_k on i . The number of values of \mathbb{D}_i consistent with literal $x_i \in \mathbb{E}_i$ is $|\mathbb{E}_i|$. Finally, the features *not* tested along R_k are denoted by $\Psi(R_k)$. For path R_k , the probability that a randomly chosen point in feature space is consistent with R_k (i.e. the *path probability* of R_k) is given by,

$$\Pr(R_k) = \left[\prod_{(x_i \in \mathbb{E}_i) \in \Lambda(R_k)} |\mathbb{E}_i| \times \prod_{i \in \Psi(R_k)} |\mathbb{D}_i| \right] / |\mathbb{F}| \tag{15}$$

As a result, we get that,

$$\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c) = \sum_{R_k \in \mathcal{P}} \Pr(R_k) \tag{16}$$

Given an instance (\mathbf{v}, c) and a set of fixed features \mathcal{X} (and so a set of universal features $\mathcal{F} \setminus \mathcal{X}$), we now detail how to compute (14). Since some features will now be declared universal, multiple paths with possibly different conditions can become consistent. For example, in [Figure 2](#) if feature 1 and 2 are declared universal, then (at least) paths P_1 , P_2 and Q_1 are consistent with some of the possible assignments. Although universal variables might seem to complicate the computation of the conditional probability, this is not the case.

A key observation is that the feature values that make a path consistent are disjoint from the values that make other paths consistent. This observation allows us to compute the models consistent with each path and, as a result, to compute (9). Let $R_k \in \mathcal{R}$ represent some path in the decision tree. (Recall that $P_t \in \mathcal{P}$ is the target path, which is consistent with \mathbf{v} .) Let n_{ik} represent the (integer) number of assignments to feature i that are consistent with path $R_k \in \mathcal{R}$, given $\mathbf{v} \in \mathbb{F}$ and $\mathcal{X} \subseteq \mathcal{F}$. For a feature i , let \mathbb{E}_i denote the set of domain values of feature i that is consistent with path R_k . Hence, for path R_k , we consider a literal $(x_i \in \mathbb{E}_i)$. Given the above, the value of n_{ik} is defined as follows:

1. If i is fixed:

- (a) If i is tested along R_k and the value of x_i is inconsistent with \mathbf{v} , i.e. there exists a literal $(x_i \in \mathbb{E}_i) \in \Lambda(R_k)$ and $\{v_i\} \cap \mathbb{E}_i = \emptyset$, then $n_{ik} = 0$;
 - (b) If i is tested along R_k and the value of x_i is consistent with R_k , i.e. there exists a literal $(x_i \in \mathbb{E}_i) \in \Lambda(R_k)$ and $\{v_i\} \cap \mathbb{E}_i \neq \emptyset$, then $n_{ik} = 1$;
 - (c) If i is not tested along R_k , then $n_{ik} = 1$.
2. Otherwise, i is universal:
- (a) If i is tested along R_k , with some literal $x_i \in \mathbb{E}_i$, then $n_{ik} = |\mathbb{E}_i|$;
 - (b) If i is not tested along R_k , then $n_{ik} = |\mathbb{D}_i|$.

Using the definition of n_{ik} , we can then compute the number of assignments consistent with R_k as follows:

$$\#(R_k; \mathbf{v}, \mathcal{X}) = \prod_{i \in \mathcal{F}} n_{ik} \quad (17)$$

Finally, (14) is given by,

$$\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c \mid \mathbf{x}_{\mathcal{X}} = \mathbf{v}_{\mathcal{X}}) = \sum_{P_k \in \mathcal{P}} \#(P_k; \mathbf{v}, \mathcal{X}) / \sum_{R_k \in \mathcal{R}} \#(R_k; \mathbf{v}, \mathcal{X}) \quad (18)$$

As can be concluded, and in the case of a decision tree, both $\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c \mid \mathbf{x}_{\mathcal{X}} = \mathbf{v}_{\mathcal{X}})$ and $\text{WeakPAXp}(\mathcal{X}; \mathbb{F}, \kappa, \mathbf{v}, c, \delta)$ are computed in polynomial time on the size of the DT.

Example 3. With respect to the DT in Figure 2, and given the instance $((4, 4, 2), \oplus)$, the number of models for each path is shown in Table 1. For example, for set $\{3\}$, we immediately get that $\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c \mid \mathbf{x}_{\mathcal{X}} = \mathbf{v}_{\mathcal{X}}) = 15/(15+1) = 15/16$.

4.2 Computing Locally-Minimal PAXp's for DT's

Recent work showed that, for DTs, one AXp can be computed in polynomial time [45, 35, 46]. A simple polynomial-time algorithm can be summarized as follows. The AXp \mathcal{X} is initialized with all the features in \mathcal{F} . Pick the path consistent with a given instance (\mathbf{v}, c) . The features not in the path are removed from \mathcal{X} . Then, iteratively check whether $\mathcal{X} \setminus \{i\}$ guarantees that all paths to a prediction in $\mathcal{K} \setminus \{c\}$ are still inconsistent. If so, then update \mathcal{X} . As argued in Section 3, we can use a similar (deletion-based) approach for computing one locally-minimal PAXp for DTs. Such an approach builds on Algorithm 1. In the case of DTs, (10) is computed using (18) on some given set $\mathcal{S} \setminus \{i\}$. to decide whether the precision of the approximation $\mathcal{S} \setminus \{i\}$ is no smaller than the threshold δ . As stated earlier, (18) is computed in polynomial time. Hence, Algorithm 1 runs in polynomial time for DTs.

4.3 Computing Minimum-Size PAXp's for DTs

For computing a minimum-size PAXp, we propose two SMT encodings, thus showing that the decision problem is in NP, and that finding a smallest set requires a logarithmic number of calls to an NP-oracle. Regarding the two SMT encodings, one involves the multiplication of integer variables, and so it involves non-linear arithmetic. Given the structure of the problem, we also show that linear arithmetic can be used, by proposing a (polynomially) larger encoding.

A multiplication-based SMT encoding. Taking into account the definition of path probabilities (see Section 4.1), we now devise a model that computes path probabilities based on the same ideas. Let $j \in \mathcal{F}$ denote a given feature. Let n_{jk} denote the number of elements

in \mathbb{D}_j consistent with path R_k (for simplicity, we just use the path index k). Also, u_j is a boolean variable that indicates whether feature j is fixed ($u_j = 0$) or universal ($u_j = 1$). If feature j is not tested along path R_k , then if j is fixed, then $n_{jk} = 1$. If not, then $n_{jk} = |\mathbb{D}_j|$. Otherwise, j is tested along path R_k . n_{jk} is 0 if j is fixed (i.e. $u_j = 0$) and inconsistent with the values of \mathbb{D}_j allowed for path R_k . n_{jk} is 1 if j is fixed and consistent with the values of \mathbb{D}_j allowed for path R_k . If feature j is not fixed (i.e. it is deemed universal and so $u_j = 1$), then n_{jk} denotes the number of domain values of j consistent with path R_k . Let the fixed value of n_{jk} be n_{0jk} and the *universal* value of n_{jk} be n_{1jk} . Thus, n_{jk} is defined as follows,

$$n_{jk} = \text{ite}(u_j, n_{1jk}, n_{0jk}) \quad (19)$$

Moreover, let η_k denote the number of models of path R_k . Then, η_k is defined as follows:

$$\eta_k = \prod_{i \in \Phi(k)} n_{ik} \quad (20)$$

If the domains are boolean, then we can use a purely boolean formulation for the problem. However, if the domains are multi-valued, then we need this formulation.

Recall what we must ensure that (10) holds true. In the case of DTs, since we can count the models associated with each path, depending on which features are fixed or not, then the previous constraint can be translated to:

$$\sum_{R_k \in \mathcal{P}} \eta_k \geq \delta \times \sum_{R_k \in \mathcal{P}} \eta_k + \delta \times \sum_{R_k \in \mathcal{Q}} \eta_k \quad (21)$$

Recall that \mathcal{P} are the paths with the matching prediction, and \mathcal{Q} are the rest of the paths.

Finally, the soft constraints are of the form (u_i) , one for each feature $i \in \mathcal{F} \setminus \Psi(R_k)$, i.e. for the features tested along path R_k . (For each feature i not tested along R_k , i.e. $i \in \Psi(R_k)$, enforce that the feature is universal by adding a hard clause (u_i) .) The solution to the optimization problem will then be a *smallest* weak PAXp, and so also a (plain) PAXp. (The minimum-cost solution is well-known to be computed with a worst-case logarithmic number of calls (on the number of features) to an SMT solver.)

Example 4. For the running example, let us consider $\mathcal{X} = \{3\}$. This means that $u_1 = u_2 = 1$. As a result, given the instance and the proposed encoding, we get Table 2 and Table 3. Finally, by plugging into (21) the values from Table 3, we get: $15 \geq 0.93 \times (15 + 1)$. Thus, \mathcal{X} is a weak PAXp, and we can show that it is both a plain PAXp and a smallest PAXp. Indeed, with $\mathcal{Y} = \emptyset$, we get $\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c \mid \mathbf{x}_Y = \mathbf{v}_Y) = 21/32 = 0.65625 < \delta$. Hence, $\mathcal{X} = \{3\}$ is subset-minimal. Since there can be no PAXp's of smaller size, then \mathcal{X} is also a smallest PAXp.

An alternative addition-based SMT encoding. A possible downside of the SMT encoding described above is the use of multiplication of variables in (20); this causes the SMT problem formulation to involve different theories (which may turn out to be harder to reason about in practice). Given the problem formulation, we can use an encoding that just uses linear arithmetic. This encoding is organized as follows. Let the order of features be $\langle 1, 2, \dots, m \rangle$. Define $\eta_{j,k}$ as the sum of models of path R_k taking into account features 1 up to j , with $\eta_{0,k} = 1$. Given $\eta_{j-1,k}$, $\eta_{j,k}$ is computed as follows:

- Let the domain of feature j be $\mathbb{D}_j = \{v_{j1}, \dots, v_{jr}\}$, and let $s_{j,l,k}$ denote the number of models taking into account features 1 up to $j-1$ and domain values v_{j1} up to v_{jl-1} . Also, let $s_{j,0,k} = 0$.

Feature	Attr.	P_1	P_2	P_3	Q_1	Q_2
1	n_{01k}	0	1	1	0	1
	n_{11k}	1	3	3	1	3
	n_{1k}	$n_{1k} = \text{ite}(u_1, n_{11k}, n_{01k})$				
2	n_{02k}	1	0	1	0	1
	n_{12k}	3	1	3	1	3
	n_{2k}	$n_{2k} = \text{ite}(u_2, n_{12k}, n_{02k})$				
3	n_{03k}	1	1	1	1	0
	n_{13k}	2	2	1	2	1
	n_{3k}	$n_{3k} = \text{ite}(u_3, n_{13k}, n_{03k})$				
Path counts		$\eta_k = n_{1k} \times n_{2k} \times n_{3k}$				

Table 2: SMT encoding for multiplication-based encoding

Path	n_{1k}	n_{2k}	n_{3k}	η_k
R_1	1	3	1	3
R_2	1	3	1	3
R_3	3	3	1	9
R_4	1	1	1	1
R_5	3	3	0	0

 Table 3: Concrete values for the multiplication-based encoding for the case $\mathcal{X} = \{3\}$, i.e. $u_1 = u_2 = 1$ and $u_3 = 0$

- For each value v_{jl} in \mathbb{D}_j , for $l = 1, \dots, r$:
 - If j is tested along path R_k : (i) If v_{jl} is inconsistent with path R_k , then $s_{j,l,k} = s_{j,l-1,k}$; (ii) If v_{jl} is consistent with path R_k and with \mathbf{v} , then $s_{j,l,k} = s_{j,l-1,k} + \eta_{j-1,k}$; (iii) If v_{jl} is consistent with path R_k but not with \mathbf{v} , or if feature j is not tested in path R_k , then $s_{j,l,k} = s_{j,l-1,k} + \text{ite}(u_j, \eta_{j-1,k}, 0)$.
 - If j is not tested along path R_k : (i) If v_{jl} is consistent with \mathbf{v} , then $s_{j,l,k} = s_{j,l-1,k} + \eta_{j-1,k}$; (ii) Otherwise, $s_{j,l,k} = s_{j,l-1,k} + \text{ite}(u_j, \eta_{j-1,k}, 0)$.
- Finally, define $\eta_{j,k} = s_{j,r,k}$.

After considering all the features in order, $\eta_{m,k}$ represents the number of models for path R_k given the assignment to the u_j variables. As a result, we can re-write (21) as follows:

$$\sum_{R_k \in \mathcal{P}} \eta_{m,k} \geq \delta \times \sum_{R_k \in \mathcal{P}} \eta_{m,k} + \delta \times \sum_{R_k \in \mathcal{Q}} \eta_{m,k} \quad (22)$$

As with the multiplication-based encoding, the soft clauses are of the form (u_i) for $i \in \mathcal{F}$.

Example 5. Table 4 summarizes the SMT encoding based on iterated summations for paths with either prediction \oplus or \ominus . The final computed values are then used in the linear

Var.	$R_1 \cong P_1$	$R_2 \cong P_2$	$R_3 \cong P_3$	$R_4 \cong Q_1$	$R_5 \cong Q_2$
$s_{1,0,k}$	$s_{1,0,1} = 0$	$s_{1,0,2} = 0$	$s_{1,0,3} = 0$	$s_{1,0,4} = 0$	$s_{1,0,5} = 0$
$s_{1,1,k}$	$s_{1,0,1} + \text{ite}(u_1, \eta_{0,1}, 0)$	$s_{1,0,2}$	$s_{1,0,3}$	$s_{1,0,4} + \text{ite}(u_1, \eta_{0,4}, 0)$	$s_{1,0,5}$
$s_{1,2,k}$	$s_{1,1,1}$	$s_{1,1,2} + \text{ite}(u_1, \eta_{0,2}, 0)$	$s_{1,1,3} + \text{ite}(u_1, \eta_{0,3}, 0)$	$s_{1,1,4}$	$s_{1,1,5} + \text{ite}(u_1, \eta_{0,5}, 0)$
$s_{1,3,k}$	$s_{1,2,1}$	$s_{1,2,2} + \text{ite}(u_1, \eta_{0,2}, 0)$	$s_{1,2,3} + \text{ite}(u_1, \eta_{0,3}, 0)$	$s_{1,2,4}$	$s_{1,2,5} + \text{ite}(u_1, \eta_{0,5}, 0)$
$s_{1,4,k}$	$s_{1,3,1}$	$s_{1,3,2} + \eta_{0,2}$	$s_{1,3,3} + \eta_{0,3}$	$s_{1,3,4}$	$s_{1,3,5} + \eta_{0,5}$
$\eta_{1,k}$	$s_{1,4,1}$	$s_{1,4,2}$	$s_{1,4,3}$	$s_{1,4,4}$	$s_{1,4,5}$
$s_{2,0,k}$	$s_{2,0,1} = 0$	$s_{2,0,2} = 0$	$s_{2,0,3} = 0$	$s_{2,0,4} = 0$	$s_{2,0,5} = 0$
$s_{2,1,k}$	$s_{2,0,1}$	$s_{2,0,2} + \eta_{1,2}$	$s_{2,0,3}$	$s_{2,0,4} + \text{ite}(u_2, \eta_{1,4}, 0)$	$s_{2,0,5} + \text{ite}(u_2, \eta_{1,5}, 0)$
$s_{2,2,k}$	$s_{2,1,1} + \text{ite}(u_2, \eta_{1,1}, 0)$	$s_{2,1,2}$	$s_{2,1,3} + \text{ite}(u_2, \eta_{1,3}, 0)$	$s_{2,1,4}$	$s_{2,1,5}$
$s_{2,3,k}$	$s_{2,2,1} + \text{ite}(u_2, \eta_{1,1}, 0)$	$s_{2,2,2}$	$s_{2,2,3} + \text{ite}(u_2, \eta_{1,3}, 0)$	$s_{2,2,4}$	$s_{2,2,5}$
$s_{2,4,k}$	$s_{2,3,1} + \eta_{1,1}$	$s_{2,3,2}$	$s_{2,3,3} + \eta_{1,3}, 0$	$s_{2,3,4}$	$s_{2,3,5}$
$\eta_{2,k}$	$s_{2,4,1}$	$s_{2,4,2}$	$s_{2,4,3}$	$s_{2,4,4}$	$s_{2,4,5}$
$s_{3,0,k}$	$s_{3,0,1} = 0$	$s_{3,0,2} = 0$	$s_{3,0,3} = 0$	$s_{3,0,4} = 0$	$s_{3,0,5} = 0$
$s_{3,1,k}$	$s_{3,0,1} + \text{ite}(u_3, \eta_{2,1}, 0)$	$s_{3,0,2} + \text{ite}(u_3, \eta_{2,2}, 0)$	$s_{3,0,3}$	$s_{3,0,4} + \text{ite}(u_3, \eta_{2,4}, 0)$	$s_{3,0,5} + \text{ite}(u_3, \eta_{2,5}, 0)$
$s_{3,2,k}$	$s_{3,1,1} + \eta_{2,1}$	$s_{3,1,2} + \eta_{2,2}$	$s_{3,1,3} + \eta_{2,3}$	$s_{3,1,4} + \eta_{2,4}$	$s_{3,1,5}$
$\eta_{3,k}$	$s_{3,4,1}$	$s_{3,4,2}$	$s_{3,4,3}$	$s_{3,2,4}$	$s_{3,2,5}$

Table 4: Partial addition-based SMT encoding for paths with prediction \oplus , with $(\mathbf{v}, c) = ((4, 4, 2), \oplus)$, and with $\eta_{0,1} = \eta_{0,2} = \eta_{0,3} = 1$

inequality (22), as follows,

$$\eta_{3,1} + \eta_{3,2} + \eta_{3,2} \geq \delta \times (\eta_{3,1} + \eta_{3,2} + \eta_{3,2}) + \delta \times (\eta_{3,4} + \eta_{3,5})$$

The optimization problem also includes $\mathcal{B} = \{(\neg u_1), (\neg u_2), (\neg u_3)\}$ as the soft clauses. For the counting-based encoding, and from Table 4, we get the values shown in Table 5. Moreover, we can then confirm that $15 \geq 0.93 \times 16$, as intended.

Discussion. In this as in the following sections, one might consider the use of a model counter as a possible alternative. However, a model counter would have to be used for each pick of features. Given the complexity of exactly computing the number of models, such approaches are all but assured to be impractical in practice.

4.4 Deciding Whether a Locally-Minimal PAXp is a Plain PAXp for DTs

The problem of deciding whether a set of features \mathcal{X} , representing an LmPAXp, is subset-minimal can be achieved by using one of the models above, keeping the features that are already universal, and checking whether additional universal features can be made to exist. In addition, we need to add constraints forcing universal features to remain universal, and at least one of the currently fixed features to also become universal. Thus, if \mathcal{X} is the set of fixed features, the SMT models proposed in earlier sections is extended with the following constraints:

$$\bigwedge_{j \in \mathcal{F} \setminus \mathcal{X}} (u_j) \wedge \left(\bigvee_{j \in \mathcal{X}} u_j \right) \quad (23)$$

which allow checking whether some set of fixed features can be declared universal while respecting the other constraints.

Var.	$R_1 \cong P_1$	$R_2 \cong P_2$	$R_3 \cong P_3$	$R_4 \cong Q_1$	$R_5 \cong Q_2$
$s_{1,0,k}$	0	0	0	0	0
$s_{1,1,k}$	1	0	0	1	0
$s_{1,2,k}$	1	2	1	1	1
$s_{1,3,k}$	1	2	2	1	2
$s_{1,4,k}$	1	3	3	1	3
$\eta_{1,k}$	1	3	3	1	3
$s_{2,0,k}$	0	0	0	0	0
$s_{2,1,k}$	0	3	0	0	3
$s_{2,2,k}$	1	3	3	1	3
$s_{2,3,k}$	2	3	6	1	3
$s_{2,4,k}$	3	3	9	1	3
$\eta_{2,k}$	3	3	9	1	3
$s_{3,0,k}$	0	0	0	0	0
$s_{3,1,k}$	0	2	0	0	0
$s_{3,2,k}$	3	3	9	1	0
$\eta_{3,k}$	3	3	9	1	0

Table 5: Assignment to variables of addition-based SMT encoding, given $\mathcal{X} = \{3\}$, i.e. $u_1 = u_2 = 1$ and $u_3 = 0$

4.5 Instance-Based vs. Path-Based Explanations

The standard definitions of abductive explanations consider a concrete instance (\mathbf{v}, c) . As argued earlier (see (7)), each (weak) AXp \mathcal{X} can then be viewed as a rule of the form:

$$\mathbf{IF} \quad [\wedge_{i \in \mathcal{X}} (x_i = v_i)] \quad \mathbf{THEN} \quad [\kappa(\mathbf{x}) = c]$$

In the case of DTs, a given \mathbf{v} is consistent with a concrete path P_t . As argued in recent work [46], this enables studying instead generalizations of AXp’s, concretely to so-called *path-based explanations*, each of which can be viewed as representing instead a rule of the form:

$$\mathbf{IF} \quad [\wedge_{i \in \mathcal{X}} (x_i \in \mathbb{E}_i)] \quad \mathbf{THEN} \quad [\kappa(\mathbf{x}) = c] \quad (24)$$

where $\mathbb{E}_i \subseteq \mathbb{D}_i$ and where each literal $x_i \in \mathbb{E}_i$ is one of the literals in the path P_t consistent with the instance (\mathbf{v}, c) .

Clearly, the literals associated with a path R_k offer more information than those associated with a concrete point \mathbf{v} in feature space. As a result, in the case of DTs, we consider a generalization of the definition of relevant set, and seek instead to compute:

$$\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c \mid \mathbf{x}_{\mathcal{X}} \in \mathbb{E}_{\mathcal{X}}) \quad (25)$$

where the notation $\mathbf{x}_{\mathcal{X}} \in \mathbb{E}_{\mathcal{X}}$ represents the constraint $\wedge_{i \in \mathcal{X}} x_i \in \mathbb{E}_i$, and where \mathbb{E}_i denotes the set of values consistent with feature i in path R_k .) Thus, the condition of weak PAXp considers instead the following probability:

$$\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c \mid \mathbf{x}_{\mathcal{X}} \in \mathbb{E}_{\mathcal{X}}) \geq \delta \quad (26)$$

The rest of this section investigates the computation of path probabilities in the case of path-based explanations. For instance-based explanations, the definition of n_{ik} needs to be adapted. Let $P_t \in \mathcal{P}$ be the target path. (For example, P_t can be the path consistent with \mathbf{v} .) Moreover, let $R_k \in \mathcal{R}$ by some path in the decision tree. For a feature i , let E_{ik} denote the set of domain values of feature i that is consistent with path R_k . Hence, for path R_k , we consider a literal $(x_i \in E_{ik})$. Similarly, let E_{it} denote the set of domain values of feature i that is consistent with path P_t . Thus, for path P_t , we consider a literal $(x_i \in E_{it})$. Given the above, the value of n_{ik} is now defined as follows:

1. If i is fixed:
 - (a) If i is tested along R_k and the value of x_i is inconsistent with E_{it} , i.e. there exists a literal $(x_i \in \mathbb{E}_i) \in \Lambda(R_k)$ and $\mathbb{E}_{it} \cap \mathbb{E}_{ik} = \emptyset$, then $n_{ik} = 0$;
 - (b) If i is tested along R_k and the value of x_i is consistent with R_k , i.e. there exists a literal $(x_i \in \mathbb{E}_i) \in \Lambda(R_k)$ and $\mathbb{E}_{it} \cap \mathbb{E}_{ik} \neq \emptyset$, then $n_{ik} = 1$.
 - (c) If i is not tested along R_k , then $n_{ik} = 1$.
2. Otherwise, i is universal:
 - (a) If i is tested along R_k , with some literal $x_i \in \mathbb{E}_{ik}$, then $n_{ik} = |\mathbb{E}_{ik}|$;
 - (b) If i is not tested along R_k , then $n_{ik} = |\mathbb{D}_i|$.

Using the modified definition of n_{ik} , we can now compute (25) as follows:

$$\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c \mid \mathbf{x}_{\mathcal{X}} \in \mathbb{E}_{\mathcal{X}}) = \sum_{P_k \in \mathcal{P}} \#(P_k; \mathcal{F} \setminus \mathcal{X}, \mathbf{v}) / \sum_{R_k \in \mathcal{R}} \#(R_k; \mathcal{F} \setminus \mathcal{X}, \mathbf{v}) \quad (27)$$

The computation of probabilistic explanations proposed in the previous sections can either assume instance-based or path-based explanations. For consistency with the rest of the paper, we opted to investigate instance-based explanations. Changing the proposed algorithms to consider instead path-based explanations would be straightforward, but that is beyond the scope of this paper.

5. Probabilistic Explanations for Naive Bayes Classifiers

This section investigates the computation of relevant sets in the concrete case of NBCs.

5.1 Explaining NBCs in Polynomial Time

This section overviews the approach proposed in [62] for computing AXp's for binary NBCs. The general idea is to reduce the NBC problem into an Extended Linear Classifier (XLC) and then explain the resulting XLC. Our purpose is to devise a new approach that builds on the XLC formulation to compute δ -relevant sets for NBCs. Hence, it is useful to recall first the translation of NBCs into XLCs and the extraction of AXp's from XLCs.

Extended Linear Classifiers. We consider an XLC with categorical features. (Recall that the present paper considers NBCs with binary classes and categorical data.) Each feature $i \in \mathcal{F}$ has $x_i \in \{1, \dots, d_i\}$, (i.e. $\mathbb{D}_i = \{1, \dots, d_i\}$). Let,

$$\nu(\mathbf{x}) \triangleq w_0 + \sum_{i \in \mathcal{F}} \sigma(x_i, v_i^1, v_i^2, \dots, v_i^{d_i}) \quad (28)$$

σ is a selector function that picks the value v_i^r iff x_i takes value r . Moreover, let us define the decision function, $\kappa(\mathbf{x}) = \oplus$ if $\nu(\mathbf{x}) > 0$ and $\kappa(\mathbf{x}) = \ominus$ if $\nu(\mathbf{x}) \leq 0$.

w_0	v_1^1	v_1^2	v_2^1	v_2^2	v_3^1	v_3^2	v_4^1	v_4^2	v_5^1	v_5^2
-2.19	-2.97	3.46	2.95	-2.95	0.4	-2.83	1.17	-1.32	-2.97	3.46

(a) Example reduction of NBC to XLC (Example 6)

Γ	δ_1	δ_5	δ_2	δ_3	δ_4	Φ
9.25	6.43	6.43	5.90	3.23	2.49	15.23

 (b) Computing δ_j 's for the XLC (Example 7)

Figure 7: Values used in the running example (Example 6 and Example 7)

The reduction of a binary NBC, with categorical features, to an XLC is completed by setting: $w_0 \triangleq \text{IPr}(\oplus) - \text{IPr}(\ominus)$, $v_i^1 \triangleq \text{IPr}(x_i = 1|\oplus) - \text{IPr}(x_i = 1|\ominus)$, $v_i^2 \triangleq \text{IPr}(x_i = 2|\oplus) - \text{IPr}(x_i = 2|\ominus)$, \dots , $v_i^{d_i} \triangleq \text{IPr}(x_i = d_i|\oplus) - \text{IPr}(x_i = d_i|\ominus)$. Hence, the argmax in (3) is replaced with inequality to get the following:

$$\text{IPr}(\oplus) - \text{IPr}(\ominus) + \sum_{i=1}^m \sum_{k=1}^{k=d_i} (\text{IPr}(x_i = k|\oplus) - \text{IPr}(x_i = k|\ominus))(x_i = k) > 0 \quad (29)$$

Example 6. Figure 7a shows the resulting XLC formulation for the example in Figure 4. We also let \mathbf{f} be associated with value 1 and \mathbf{t} be associated with value 2, and $d_i = 2$.

Explaining XLCs. We now describe how AXp's can be computed for XLCs. For a given instance $\mathbf{x} = \mathbf{a}$, define a *constant* slack (or gap) value Γ given by,

$$\Gamma \triangleq \nu(\mathbf{a}) = \sum_{i \in \mathcal{F}} \sigma(a_i, v_i^1, v_i^2, \dots, v_i^{d_i}) \quad (30)$$

Computing an AXp corresponds to finding a subset-minimal set of literals $\mathcal{S} \subseteq \mathcal{F}$ such that (5) holds, or alternatively,

$$\forall (\mathbf{x} \in \mathbb{F}). \bigwedge_{i \in \mathcal{S}} (x_i = a_i) \rightarrow (\nu(\mathbf{x}) > 0) \quad (31)$$

under the assumption that $\nu(\mathbf{a}) > 0$. Thus, the purpose is to find the *smallest* slack that can be achieved by allowing the feature not in \mathcal{S} to take any value (i.e. *universal/free* features), given that the literals in \mathcal{S} are fixed by \mathbf{a} (i.e. $\bigwedge_{i \in \mathcal{S}} (x_i = a_i)$).

Let v_i^ω denote the *smallest* (or *worst-case*) value associated with x_i . Then, by letting every x_i take *any* value, the *worst-case* value of $\nu(\mathbf{e})$ is,

$$\Gamma^\omega = w_0 + \sum_{i \in \mathcal{F}} v_i^\omega \quad (32)$$

Moreover, from (30), we have: $\Gamma = w_0 + \sum_{i \in \mathcal{F}} v_i^{a_i}$. The expression above can be rewritten as follows,

$$\begin{aligned} \Gamma^\omega &= w_0 + \sum_{i \in \mathcal{F}} v_i^{a_i} - \sum_{i \in \mathcal{F}} (v_i^{a_i} - v_i^\omega) \\ &= \Gamma - \sum_{i \in \mathcal{F}} \delta_i = -\Phi \end{aligned} \quad (33)$$

where $\delta_i \triangleq v_i^{a_i} - v_i^\omega$, and $\Phi \triangleq \sum_{i \in \mathcal{F}} \delta_i - \Gamma = -\Gamma^\omega$. Recall the goal is to find a subset-minimal set \mathcal{S} such that the prediction is still c (whatever the values of the other features):

$$w_0 + \sum_{i \in \mathcal{S}} v_i^{a_i} + \sum_{i \notin \mathcal{S}} v_i^\omega = -\Phi + \sum_{i \in \mathcal{S}} \delta_i > 0 \quad (34)$$

In turn, (34) can be represented as the following knapsack problem [54]:

$$\begin{aligned} \min \quad & \sum_{i=1}^m p_i \\ \text{such that} \quad & \sum_{i=1}^m \delta_i p_i > \Phi \\ & p_i \in \{0, 1\} \end{aligned} \quad (35)$$

where the variables p_i assigned value 1 denote the indices included in \mathcal{S} . Note that, the fact that the coefficients in the cost function are all equal to 1 makes the problem solvable in log-linear time [62].

Example 7. Figure 7b shows the values used for computing explanations for the example in Figure 4. For this example, the sorted δ_j 's become $\langle \delta_1, \delta_5, \delta_2, \delta_4, \delta_3 \rangle$. By picking δ_1 , δ_2 and δ_5 , we ensure that the prediction is \oplus , independently of the values assigned to features 3 and 4. Thus $\{1, 2, 5\}$ is an AXp for the NBC shown in Figure 3, with the input instance $(v_1, v_2, v_3, v_4, v_5) = (\mathbf{t}, \mathbf{f}, \mathbf{f}, \mathbf{f}, \mathbf{t})$. (It is easy to observe that $\kappa(\mathbf{t}, \mathbf{f}, \mathbf{f}, \mathbf{t}, \mathbf{t}) = \kappa(\mathbf{t}, \mathbf{f}, \mathbf{t}, \mathbf{f}, \mathbf{t}) = \kappa(\mathbf{t}, \mathbf{f}, \mathbf{t}, \mathbf{t}, \mathbf{t}) = \oplus$.)

The next section introduces a pseudo-polynomial time algorithm for computing locally-minimal PAXp's. Although locally-minimal PAXp's are not necessarily subset/cardinality minimal, the experiments (see Section 7) show that the proposed approach computes (in pseudo-polynomial time) succinct [69] and highly precise locally-minimal explanations.

5.2 Counting Models of XLCs

Earlier work [27, 31, 30, 83] proposed the use of dynamic programming (DP) for approximating the number of feasible solutions of the 0-1 knapsack constraint, i.e. the #knapsack problem. Here we propose an extension of the basic formulation, to allow counting feasible solutions of XLCs.

We are interested in the number of solutions of,

$$\sum_{j \in \mathcal{F}} \sigma(x_j, v_j^1, v_j^2, \dots, v_j^{d_j}) > -w_0 \quad (36)$$

where we assume all v_j^i to be integer-valued and non-negative (e.g. this is what our translation from NBCs to XLCs yields after scaling and rounding). Moreover, (36) can be written as follows:

$$\sum_{j \in \mathcal{F}} \sigma(x_j, -v_j^1, -v_j^2, \dots, -v_j^{d_j}) < w_0 \quad (37)$$

which reveals the relationship with the standard knapsack constraint.

For each j , let us sort the $-v_j^i$ in non-decreasing order, collapsing duplicates, and counting the number of duplicates, obtaining two sequences:

$$\begin{aligned} & \langle w_j^1, \dots, w_j^{d'_j} \rangle \\ & \langle n_j^1, \dots, n_j^{d'_j} \rangle \end{aligned}$$

such that $w_j^1 < w_j^2 < \dots < w_j^{d'_j}$ and each $n_j^i \geq 1$ gives the number of repetitions of weight w_j^i .

Counting. Let $C(k, r)$ denote the number of solutions of (37) when the subset of features considered is $\{1, \dots, k\}$ and the sum of picked weights is at most r . To define the solution for the first k features, taking into account the solution for the first $k - 1$ features, we must consider that the solution for r can be obtained due to *any* of the possible values of x_j . As a result, for an XLC the general recursive definition of $C(k, r)$ becomes,

$$C(k, r) = \sum_{i=1}^{d'_k} n_k^i \times C(k-1, r - w_k^i)$$

Moreover, $C(1, r)$ is given by,

$$C(1, r) = \begin{cases} 0 & \text{if } r < w_1^1 \\ n_1^1 & \text{if } w_1^1 \leq r < w_1^2 \\ n_1^1 + n_1^2 & \text{if } w_1^2 \leq r < w_1^3 \\ \dots & \\ \sum_{i=1}^{d'_1} n_1^i & \text{if } w_1^{d'_1} \leq r \end{cases}$$

In addition, if $r < 0$, then $C(k, r) = 0$, for $k = 1, \dots, m$. Finally, the dimensions of the $C(k, r)$ table are as follows:

1. The number of rows is m .
2. The (worst-case) number of columns is given by:

$$W' = \sum_{j \in \mathcal{F}} n_j^{d'_j} \times w_j^{d'_j} \quad (38)$$

W' represents the largest possible value, in theory. However, in practice, it suffices to set the number of columns to $W = w_0 + T$, which is often much smaller than W' .

Example 8. Consider the following problem. There are 4 features, $\mathcal{F} = \{1, 2, 3, 4\}$. Each feature j takes values in $\{1, 2, 3\}$, i.e. $x_j \in \{1, 2, 3\}$. The prediction should be 1 when the sum of the values of the x_j variables is not less than 8. We set $w_0 = -7$, and get the formulation,

$$\sum_{j \in \{1, 2, 3, 4\}} \sigma(x_j, 1, 2, 3) > 7$$

where each x_j picks value in $\{1, 2, 3\}$. We translate to the extended knapsack formulation and obtain:

$$\sum_{j \in \{1, 2, 3, 4\}} \sigma(x_j, -1, -2, -3) < -7$$

We require the weights to be integer and non-negative, and so we add to each w_j^k the complement of the most negative w_j^k plus 1. Therefore, we add +4 to each j and +16 to right-hand side of the inequality. Thus, we get

$$\sum_{j \in \{1, 2, 3, 4\}} \sigma(x_j, 3, 2, 1) < 9$$

k	r												
	0	1	2	3	4	5	6	7	8	9	10	11	12
1	0	1	2	3	3	3	3	3	3	-	-	-	-
2	0	0	1	3	6	8	9	9	9	-	-	-	-
3	0	0	0	1	4	10	17	23	16	-	-	-	-
4	0	0	0	0	1	5	15	31	50	-	-	-	-

Table 6: DP table for Example 8

For this formulation, $x_j = 1$ picks value 3. (For example, we can pick two (but not three) x_j with value 1, which is as expected.)

In this case, the DP table size will be 4×12 , even though we are interested in entry $C(4, 8)$. Table 6 shows the DP table, and the number of solutions for the starting problem, i.e. there are 50 combinations of values whose sum is no less than 8.

By default, the dynamic programming formulation assumes that features can take any value. However, the same formulation can be adapted when features take a given (fixed) value. Observe that this will be instrumental for computing LmPAXp 's.

Consider that feature k is fixed to value l . Then, the formulation for $C(k, r)$ becomes:

$$C(k, r) = n_k^l \times C(k - 1, r - w_k^l) = C(k - 1, r - w_k^l)$$

Given that k is fixed, then we have $n_k^l = 1$.

Example 9. For Example 8, assume that $x_2 = 1$ and $x_4 = 3$. Then, the constraint we want to satisfy is:

$$\sum_{j \in \{1, 3\}} \sigma(x_j, 1, 2, 3) > 3$$

Following a similar transformation into knapsack formulation, we get

$$\sum_{j \in \{1, 3\}} \sigma(x_j, 3, 2, 1) < 5$$

After updating the DP table, with fixing features 2 and 4, we get the DP table shown in Table 7. As a result, we can conclude that the number of solutions is 6.

The table $C(k, r)$ can be filled out in pseudo-polynomial time. The number of rows is m . The number of columns is W (see (38)). Moreover, the computation of each entry uses the values of at most m other entries. Thus, the total running time is: $\Theta(m^2 \times W)$.

From NBCs to positive integer knapsacks. To assess heuristic explainers, we consider NBCs, and use a standard transformation from probabilities to positive real values [78]. Afterwards, we convert the real values to integer values by scaling the numbers. However, to avoid building a very large DP table, we implement the following optimization. The number

k	r												
	0	1	2	3	4	5	6	7	8	9	10	11	12
1	0	1	2	3	3	3	3	3	3	-	-	-	-
2	0	0	0	0	1	2	3	3	3	-	-	-	-
3	0	0	0	0	0	1	3	6	8	-	-	-	-
4	0	0	0	0	0	0	1	3	6	-	-	-	-

Table 7: DP table for Example 9

of decimal places of the probabilities is reduced while there is no decrease in the accuracy of the classifier both on training and on test data. In our experiments, we observed that there is no loss of accuracy if four decimal places are used, and that there is a negligible loss of accuracy with three decimal places.

Assessing explanation precision. Given a Naive Bayes classifier, expressed as an XLC, we can assess explanation accuracy in pseudo-polynomial time. Given an instance \mathbf{v} , a prediction $\kappa(\mathbf{v}) = \oplus$, and an approximate explanation \mathbf{S} , we can use the approach described in this section to count the number of instances consistent with the explanation for which the prediction remains unchanged (i.e. number of points $\mathbf{x} \in \mathbb{F}$ s.t. $(\kappa(\mathbf{x}) = \kappa(\mathbf{v}) \wedge (\mathbf{x}_{\mathcal{S}} = \mathbf{v}_{\mathcal{S}}))$). Let this number be n_{\oplus} (given the assumption that the prediction is \oplus). Let the number of instances with a different prediction be n_{\ominus} ⁷. Hence, the conditional probability (8) can be defined, in the case of NBCs, as follow:

$$\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = \oplus \mid \mathbf{x}_{\mathcal{S}} = \mathbf{v}_{\mathcal{S}}) = \frac{n_{\oplus}}{|\{\mathbf{x} \in \mathbb{F} : (\mathbf{x}_{\mathcal{S}} = \mathbf{v}_{\mathcal{S}})\}|}$$

Observe that the numerator $|\{\mathbf{x} \in \mathbb{F} : \kappa(\mathbf{x}) = \oplus \wedge (\mathbf{x}_{\mathcal{S}} = \mathbf{v}_{\mathcal{S}})\}|$ is expressed by the number of models n_{\oplus} , i.e. the points \mathbf{x} in feature space that are consistent with \mathbf{v} given \mathcal{S} and with prediction \oplus . Further, we have

$$\begin{aligned} \Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = \oplus \mid \mathbf{x}_{\mathcal{S}} = \mathbf{v}_{\mathcal{S}}) &= 1 - \Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = \ominus \mid \mathbf{x}_{\mathcal{S}} = \mathbf{v}_{\mathcal{S}}) \\ &= 1 - \frac{n_{\ominus}}{|\{\mathbf{x} \in \mathbb{F} : (\mathbf{x}_{\mathcal{S}} = \mathbf{v}_{\mathcal{S}})\}|} \end{aligned}$$

where $n_{\ominus} = |\{\mathbf{x} \in \mathbb{F} : \kappa(\mathbf{x}) = \ominus \wedge (\mathbf{x}_{\mathcal{S}} = \mathbf{v}_{\mathcal{S}})\}|$.

5.3 Computing Locally-Minimal PAXp’s for NBCs

Similarly to the case of DTs, we can also use Algorithm 1 for computing locally-minimal PAXp’s in the case of NBCs. The only difference is in the definition of the predicate WeakPAXp. For NBCs, the procedure isWeakPAXp implements the pseudo-polynomial

7. Recall that we are assuming that $\mathcal{K} = \{\ominus, \oplus\}$.

approach, described in the previous section, for model counting. Hence, in the case of NBCs, it is implicit that the DP table is updated at each iteration of the main loop of [Algorithm 1](#). More specifically, when a feature i is just set to universal, its associated cells $C(i, r)$ are recalculated such that $C(k, r) = \sum_{i=1}^{d_k} n_k^i \times C(k-1, r - w_k^i)$; and when i is fixed, i.e. $i \in \mathcal{S}$, then $C(i, r) = C(i-1, r - v_i^j)$ where $v_i^j \triangleq \text{IPr}(v_i = j|c) - \text{IPr}(v_i = j|\neg c)$. Furthermore, we point out that in our experiments, \mathcal{S} is initialized to an AXp \mathcal{X} that we compute initially for all tested instances using the outlined (polynomial) algorithm in [Section 5.1](#). It is easy to observe that features not belonging to \mathcal{X} do not contribute in the decision of $\kappa(\mathbf{v})$ (i.e. their removal does not change the value of n_\ominus that is equal to zero) and thus can be set universal at the initialisation step, which allows us to improve the performance of [Algorithm 1](#).

Moreover, we apply a heuristic order over \mathcal{S} that aims to remove earlier less relevant features and thus to produce shorter approximate explanations. Typically, we order \mathcal{S} following the increasing order of δ_i values, namely the reverse order applied to compute the AXp. Preliminary experiments conducted using a (naive heuristic) lexicographic order over the features produced less succinct explanations.

Finally, notice that [Algorithm 1](#) can be used to compute an AXp, i.e. locally-minimal PAXp when $\delta = 1$. Nevertheless, the polynomial time algorithm for computing AXp's proposed in [\[62\]](#) remains a better choice to use in case of AXp's than [Algorithm 1](#) which runs in pseudo-polynomial time.

Example 10. Let us consider again the NBC of the running example ([Example 6](#)) and $\mathbf{v} = (\mathbf{t}, \mathbf{f}, \mathbf{f}, \mathbf{f}, \mathbf{t})$. The corresponding XLC is shown in [Figure 7b](#) ([Example 7](#)). Also, consider the AXp $\{1, 2, 5\}$ of \mathbf{v} and $\delta = 0.85$. The resulting DP table for $\mathcal{S} = \{1, 2, 5\}$ is shown in [Table 8](#). Note that for illustrating small tables, we set the number of decimal places to zero (greater number of decimal places, i.e. 1, 2, etc, were tested and returned the same results). (Also, note that the DP table reports “—” if the cell is not calculated during the running of [Algorithm 1](#).) Moreover, we convert the probabilities into positive integers, so we add to each w_j^k the complement of the most negative w_j^k plus 1. The resulting weights are shown in [Figure 8](#). Thus, we get $\sum_{i \in \{1, 2, 3, 4, 5\}} \sigma(x_i, w_i^1, w_i^2) < 17$. Observe that the number of models $n_\oplus = C(5, 16)$, and $C(5, 16)$ is calculated using $C(4, 16 - w_5^2) = C(4, 15)$, i.e. $C(4, 15) = C(5, 16)$ (feature 5 is fixed, so it is allowed to take only the value $w_5^2 = 1$). Next, $C(4, 15) = C(3, 15 - w_4^1) + C(3, 15 - w_4^2) = C(3, 12) + C(3, 14)$ (feature 4 is free, so it is allowed to take any value of $\{w_4^1, w_4^2\}$); the recursion ends when $k=1$, namely for $C(1, 5) = C(2, 6) = n_1^2 = 1$, $C(1, 7) = C(2, 7) = n_1^2 = 1$, $C(1, 8) = C(2, 8) = n_1^2 = 1$ and $C(1, 10) = C(2, 11) = n_1^2 = 1$ (feature 1 is fixed and takes value w_1^2). Next, [Table 9](#) (resp. [Table 10](#) and [Table 11](#)) report the resulting DP table for $\mathcal{S} = \{2, 5\}$ (resp. $\mathcal{S} = \{1, 5\}$ and $\mathcal{S} = \{1\}$). It is easy to confirm that after dropping feature 2, the precision of $\mathcal{S} = \{1, 5\}$ becomes 87.5%, i.e. $\frac{7}{8} = 0.875 > \delta$. Furthermore, observe that the resulting \mathcal{S} when dropping feature 1 or 2 and 5, are not weak PAXp's, namely, the precision of $\{2, 5\}$ is $\frac{6}{8} = 0.75 < \delta$ and the precision of $\{1\}$ is $\frac{9}{16} = 0.5625 < \delta$. In summary, [Algorithm 1](#) starts with $\mathcal{S} = \{1, 2, 5\}$, then at iteration #1, feature 1 is tested and since $\{2, 5\}$ is not a weak PAXp then 1 is kept in \mathcal{S} ; at iteration #2, feature 2 is tested and since $\{1, 5\}$ is a weak PAXp, then \mathcal{S} is updated (i.e. $\mathcal{S} = \{1, 5\}$); at iteration #3, feature 5 is tested and since $\{1\}$ is not a weak PAXp, then 5 is saved in \mathcal{S} . As a result, the computed locally-minimal PAXp is $\{1, 5\}$.

k	r																
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	0	—	—	—	—	1	—	1	1	—	1	—	—	—	—	—	—
2	0	—	—	—	—	—	1	—	1	1	—	1	—	—	—	—	—
3	0	—	—	—	—	—	—	—	—	—	—	—	2	—	2	—	—
4	0	—	—	—	—	—	—	—	—	—	—	—	—	—	—	4	—
5	0	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—	4

 Table 8: DP table for $\mathcal{S} = \{1, 2, 5\}$ (Example 10)

k	r																
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	0	—	—	—	—	1	—	1	2	—	2	—	—	—	—	—	—
2	0	—	—	—	—	—	1	—	1	2	—	2	—	—	—	—	—
3	0	—	—	—	—	—	—	—	—	—	—	—	3	—	3	—	—
4	0	—	—	—	—	—	—	—	—	—	—	—	—	—	—	6	—
5	0	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—	6

 Table 9: DP table for $\mathcal{S} = \{2, 5\}$ (Example 10)

We underline that we could initialize \mathcal{S} to \mathcal{F} , in which case the number of models would be 1. However, we opt instead to always start from an AXp. In the example, the AXp is $\{1, 2, 5\}$ which, because it is an AXp, the number of models must be 4 (i.e. 2^2 , since two features are free).

For any proper subset of the AXp, with r free variables, it must be the case that the number of models is strictly less than 2^r . Otherwise, we would have an AXp as a proper subset of another AXp; but this would contradict the definition of AXp. The fact that the number of models is strictly less than 2^r is confirmed by the examples of subsets considered. It must also be the case that if $\mathcal{S}' \subseteq \mathcal{S}$, then the number of models of \mathcal{S}' must not exceed the number of models of \mathcal{S} . So, we can argue that there is monotonicity in the number of models, but not in the precision.

W	w_1^1	w_1^2	w_2^1	w_2^2	w_3^1	w_3^2	w_4^1	w_4^2	w_5^1	w_5^2
16	7	1	1	6	3	6	1	3	7	1

Figure 8: #knapsack problem of Example 10

k	r																
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	0	—	1	1	—	1	—	1	1	—	1	—	—	—	—	—	—
2	0	—	—	—	—	—	1	—	2	2	—	2	—	—	—	—	—
3	0	—	—	—	—	—	—	—	—	—	—	—	3	—	4	—	—
4	0	—	—	—	—	—	—	—	—	—	—	—	—	—	—	7	—
5	0	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—	7

Table 10: DP table for $\mathcal{S} = \{1, 5\}$ (Example 10)

k	r																
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	0	0	1	1	1	1	—	1	1	—	1	—	—	—	—	—	—
2	0	—	0	1	—	1	1	—	2	2	—	2	—	—	—	—	—
3	0	—	—	—	—	—	1	—	1	—	—	—	3	—	4	—	—
4	0	—	—	—	—	—	—	—	—	2	—	—	—	—	—	7	—
5	0	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—	9

Table 11: DP table for $\mathcal{S} = \{1\}$ (Example 10)

6. Probabilistic Explanations for Other Families of Classifiers

This section investigates additional families of classifiers, namely those either based on propositional languages [34] or based on graphs [35]. It should be noted that some families of classifiers can be viewed as both propositional languages and as graphs.

6.1 Propositional Classifiers

This section considers the families of classifiers represented as propositional languages and studied in [34]. Examples include classifiers based on d-DNNFs and SDDs.

Proposition 1. For any language that allows conditioning (**CD**) in polynomial time, a locally-minimal PAXp can be found in polynomial time (for all $\delta \in [0, 1]$) if and only if the language allows *counting* (**CT**) in polynomial time.

Proof. If the language allows **CD** and **CT** in polynomial time, Algorithm 1 finds a locally-minimal PAXp in polynomial time. Conversely, testing whether $\mathcal{X} = \emptyset$ is a locally-minimal PAXp amounts to determining whether the number of models of the classifier κ is at least $\delta 2^m$, and hence a binary search on δ computes the number of models (i.e. **CT**) in polynomial time. \square

Proposition 2. For any language that allows **CT** and **CD** in polynomial time, given an integer k , the problem of deciding the existence of a weak PAXp of size at most k is in NP.

Proof. Let $\mathcal{P} \subseteq \mathcal{F}$ denote a guessed set of picked features of size $k' \leq k$, to be fixed. Since **CT** and **CD** run in polynomial time, the left-hand side of (10), i.e. $\Pr_{\mathbf{x}}(\kappa(\mathbf{x}) = c \mid \mathbf{x}_{\mathcal{X}} = \mathbf{v}_{\mathcal{X}})$ can be computed in polynomial time. Consequently, we can decide in polynomial time whether or not \mathcal{P} is a weak PAXp, and so the decision problem is in NP. \square

Corollary 1. For any language that allows **CT** and **CD** in polynomial time, the problem of finding a minimum-size PAXp belongs to $\text{FP}^{\text{NP}}[\mathcal{O}(\log m)]$.

Proof. It suffices to run a binary search, on the number k of features, using an NP oracle that decides the decision problem introduced in the proof of Proposition 2. This algorithm finds a smallest weak PAXp \mathcal{P} , which is necessarily a PAXp (and hence a smallest PAXp) since no proper subset of \mathcal{P} can be a weak PAXp. \square

Corollary 1 notably applies to the following languages:

- the language EADT of extended affine decision trees [56] which are a strict generalisation of decision trees.
- the language d-DNNF and sublanguages of d-DNNF, such as SDD and OBDD.

6.2 Graph-Based Classifiers

This section considers the DG classifiers introduced in Section 2.3, and in particular two restricted types of DGs: OMDDs and OBDDs. For OMDDs and OBDDs, it is well-known that there are poly-time algorithms for model counting [23, 75]. Hence, we can compute a minimum-size PAXp with a logarithmic number of calls to an NP oracle. We will describe one such algorithm in this section. Moreover, a general method for counting models of DGs will be described as well.

Counting models of OBDD/OMDD. We describe next a dedicated algorithm for counting models of an OMDD, which can also be used with OBDDs.

We associate an indicator n_p to each node p , the indicator of the root node being n_1 , such that the value of n_1 represents the number of models consistent with the class c . And let s_i be the boolean selector of feature i such that $s_i = 1$ if feature i is included in PAXp, otherwise it is free. Besides, we assume that the feature-index order is increasing top-down, and the leftmost outgoing edge of a non-terminal node is consistent with the given instance \mathbf{v} .

1. For a terminal node p , $n_p = 1$ if the label of node p is consistent with the class c ; otherwise $n_p = 0$.
2. For a non-terminal node p label with feature i , that has k child nodes q_1, \dots, q_k , let q be an arbitrary child node of p labeled with feature j , and edge (p, q_1) be the leftmost edge. By assumption, since $i < j$, there may exist some untested features $i < l < j$ along the outgoing edge (p, q) . Let $b_{p,q}$ be the indicator of the edge (p, q) , we have: $b_{p,q} = n_q \times \prod_{i < l < j} \text{ite}(s_l, 1, \mathbb{D}_l) \times |(p, q)|$, so the indicator of node p is $n_p = \text{ite}(s_i, b_{p,q_1}, \sum_{q_1 \leq l \leq q_k} b_{p,l})$. It should be noted that multiple edges may exist; we use $|(p, q)|$ to denote the number of edges between node p and q .

3. The number of models consistent with the class c is n_1 .

In the case of $|\mathcal{K}| = 2$ and $\mathbb{F} = \mathbb{B}^m$, this algorithm is still applicable to OBDDs with some minor modifications. Note that each non-terminal node of an OBDD has exactly two outgoing edges and multiple edges between two nodes are not allowed (otherwise, the OBDD is not *reduced*); this means that $\mathbb{D}_l = \mathbb{B}$, and $|(p, q)| = 1$.

Counting models for unrestricted DGs. For unrestricted DG classifiers, given [Section 2.3](#), any path connecting the root to a terminal node is consistent with some input, and no two paths are consistent with the same input. As a result, in the case of unrestricted DG classifiers, it suffices to enumerate all paths, and count the models associated with the path. Of course, the downside is that the number of paths is worst-case exponential on the size of the graph. To alleviate this, an approach similar to the one outlined above can be considered.

7. Experiments

This section reports the experimental results on computing relevant sets for the classifiers studied in the earlier sections, namely: decision trees, naive Bayes classifiers and graph-based classifiers (in our case OMDDs). For each case study, we describe the prototype implementation and the used benchmarks; we show a table that summarizes the results and then we discuss the results. Furthermore, for the case of DTs, the evaluation includes a comparison with the model-agnostic explainer Anchor [\[82\]](#), aiming at assessing not only the succinctness and precision of computed explanations but also the scalability of our solution. (Observe that for the case of NBCs, earlier work on computing AXp’s [\[62\]](#) have compared their approach with the heuristic methods, e.g. Anchor, SHAP, and show that the latter are slower and do not show a strong correlation between features of their explanations and common features identified from AXp’s. As a result, the comparison with Anchor is restricted to the case of DTs.)

All experiments were conducted on a MacBook Air with a 1.1GHz Quad-Core Intel Core i5 CPU with 16 GByte RAM running macOS Monterey.

7.1 Case Study 1: Decision Trees

Prototype implementation. A prototype implementation⁸ of the proposed algorithms for DTs was developed in Python; whenever necessary, it instruments oracle calls to the well-known SMT solver z3⁹ [\[24\]](#) as described in [Section 4](#). Hence, the prototype implements the LmPAXp procedure outlined in [Algorithm 1](#) and augmented with a heuristic that orders the features in \mathcal{X} . The idea consists in computing the precision loss of the overapproximation of each $\mathcal{X} \setminus \{j\}$ and then sorting the features from the less to the most important one. This strategy often allows us to obtain the closest superset to a PAXp, in contrast to the simple lexicographic order applied over \mathcal{X} . (Recall that \mathcal{X} is initialized to the set of features involved in the decision path.) Algorithm MinPAXp outlined in [Section 4.3](#) implements the two (multiplication- and addition-based) SMT encodings. Nevertheless, preliminary results show that both encodings perform similarly, with some exceptions where the addition-based

8. All sources implemented in these experiments will be publicly available after the paper gets accepted.

9. <https://github.com/Z3Prover/z3/>

Dataset	DT					δ	MinPAXp					LmPAXp					D	Anchor						
	N		Path				Length			Prec	Time	Length			Prec	m_{\subseteq}		Time	Length			Prec	Time	
	N	A	M	m	avg		M	m	avg	avg	avg	M	m	avg	avg	avg		M	m	avg	$F_{\not\subseteq P}$	avg	avg	
adult	1241	89	14	3	10.7	100	11	3	6.8	100	2.34	11	3	6.9	100	100	0.00	d	12	2	7.0	26.8	76.8	0.96
						95	11	3	6.2	98.4	5.36	11	3	6.3	98.6	99.0	0.01	u	12	3	10.0	29.4	93.7	2.20
						90	11	2	5.6	94.6	4.64	11	2	5.8	95.2	96.4	0.01							
dermatology	71	100	13	1	5.1	100	12	1	4.4	100	0.35	12	1	4.4	100	100	0.00	d	31	1	4.8	58.1	32.9	3.10
						95	12	1	4.1	99.7	0.37	12	1	4.1	99.7	99.3	0.00	u	34	1	13.1	43.2	87.2	25.13
						90	11	1	4.0	98.8	0.35	11	1	4.0	98.8	100	0.00							
kr-vs-kp	231	100	14	3	6.6	100	12	2	4.8	100	0.93	12	2	4.9	100	100	0.00	d	36	2	7.9	44.8	69.4	1.94
						95	11	2	3.9	98.1	0.97	11	2	4.0	98.1	100	0.00	u	12	2	3.6	16.6	97.3	1.81
						90	10	2	3.2	95.4	0.92	10	2	3.3	95.4	99.0	0.00							
letter	3261	93	14	4	11.8	100	12	4	8.2	100	16.06	11	4	8.2	100	100	0.00	d	16	3	13.2	43.1	71.3	12.22
						95	12	4	8.0	99.6	18.28	11	4	8.0	99.5	100	0.00	u	16	3	13.7	47.3	66.3	10.15
						90	12	4	7.7	97.7	16.35	10	4	7.8	97.8	100	0.00							
soybean	219	100	16	3	7.3	100	14	3	6.4	100	0.92	14	3	6.5	100	100	0.00	d	35	2	8.6	55.4	33.6	5.43
						95	14	3	6.4	99.8	0.95	14	3	6.4	99.8	100	0.00	u	35	3	19.2	66.0	75.0	38.96
						90	14	3	6.1	98.1	0.94	14	3	6.1	98.2	98.5	0.00							
spambase	141	99	14	3	8.5	0	12	3	7.4	100	1.23	12	3	7.5	100	100	0.01	d	38	2	6.3	65.3	63.3	24.12
						95	9	1	3.7	96.1	2.16	9	1	3.8	96.5	100	0.01	u	57	3	28.0	86.2	65.3	834.70
						90	6	1	2.4	92.4	2.15	8	1	2.4	92.2	100	0.01							
texture	257	100	13	3	6.6	100	12	3	6.2	100	2.01	11	3	6.2	100	100	0.01	d	40	2	16.5	80.6	32.2	532.42
						95	11	3	5.4	99.3	2.19	11	3	5.4	99.4	100	0.01	u	40	5	17.5	84.4	31.6	402.07
						90	11	3	5.4	98.5	2.20	11	3	5.4	99.4	100	0.01							

Table 12: Assessing explanations of MinPAXp, LmPAXp and Anchor for DTs. (For each dataset, we run the explainers on 500 samples randomly picked or all samples if there are less than 500.) In column **DT**, **N** and **A** denote, resp., the number of nodes and the training accuracy of the DT. Column δ reports (in %) the value of the threshold δ . In column **Path**, **avg** (resp. **M** and **m**) denotes the average (resp. max. and min.) depth of paths consistent with the instances. In column **Length**, **avg** (resp. **M** and **m**) denotes the average (resp. max. and min.) length of the explanations; and $F_{\not\subseteq P}$ denotes the avg. % of features in Anchor’s explanations that do not belong to the consistent paths. **Prec** reports (in %) the average precision (defined in (9)) of resulting explanations. m_{\subseteq} shows the number in (%) of LmPAXp’s that are subset-minimal, i.e. PAXp’s. **Time** reports (in seconds) the average runtime to compute an explanation. Finally, **D** indicates which distribution is applied on data given to Anchor: either data distribution (denoted by d) or uniform distribution (denoted by u).

encoding is much larger and so slower. Therefore, the results reported below refer only to the multiplication-based encoding.

Benchmarks. The benchmarks used in the experiments comprise publicly available and widely used datasets obtained from the UCI ML Repository [95]. All the DTs are trained using the learning tool *IAI* (*Interpretable AI*) [11, 38]. The maximum depth parameter in *IAI* is set to 16. As the baseline, we ran Anchor with the default explanation precision of 0.95. Two assessments are performed with Anchor: (i) with the original training data¹⁰ that follows the data distribution; (ii) with using sampled data that follows a uniform distribution. Our setup assumes that all instances of the feature space are equally possible, and so there is no assumed probability distribution over the features. Therefore in order to be fair with Anchor, we further assess Anchor with uniformly sampled data. (Also, we point out that the implementation of Anchor demonstrates that it can generate samples that do not belong to the input distribution. Thus, there is no guarantee that these samples come from the input

10. The same training set used to learn the model.

Dataset	#F	#I	NBC A%	AXp		LmPAXp _{≤9}				LmPAXp _{≤7}				LmPAXp _{≤4}			
				Length	δ	Length	Precision	W%	Time	Length	Precision	W%	Time	Length	Precision	W%	Time
adult	(13	200)	81.37	6.8± 1.2	98	6.8± 1.1	100± 0.0	100	0.003	6.3± 0.9	99.61± 0.6	96	0.023	4.8± 1.3	98.73± 0.5	48	0.059
					95	6.8± 1.1	99.99± 0.2	100	0.074	5.9± 1.0	98.87± 1.8	99	0.058	3.9± 1.0	96.93± 1.1	80	0.071
					93	6.8± 1.1	99.97± 0.4	100	0.104	5.7± 1.3	98.34± 2.6	100	0.086	3.4± 0.9	95.21± 1.6	90	0.093
					90	6.8± 1.1	99.95± 0.6	100	0.164	5.5± 1.4	97.86± 3.4	100	0.100	3.0± 0.8	93.46± 1.5	94	0.103
agaricus	(23	200)	95.41	10.3± 2.5	98	7.7± 2.7	99.12± 0.8	92	0.593	6.4± 3.0	98.75± 0.6	87	0.763	6.0± 3.1	98.67± 0.5	29	0.870
					95	6.9± 3.1	97.62± 2.1	95	0.954	5.3± 3.2	96.59± 1.6	92	1.273	4.8± 3.3	96.24± 1.2	55	1.217
					93	6.5± 3.1	96.65± 2.8	95	1.112	4.8± 3.1	95.38± 1.9	93	1.309	4.3± 3.1	94.92± 1.3	64	1.390
					90	5.9± 3.3	94.95± 4.1	96	1.332	4.0± 3.0	92.60± 2.8	95	1.598	3.6± 2.8	92.08± 1.7	76	1.830
chess	(37	200)	88.34	12.1± 3.7	98	8.1± 4.1	99.27± 0.6	64	0.383	5.9± 4.9	98.70± 0.4	64	0.454	5.7± 5.0	98.65± 0.4	46	0.457
					95	7.7± 3.8	98.51± 1.4	68	0.404	5.5± 4.4	97.90± 0.9	64	0.483	5.3± 4.5	97.85± 0.8	46	0.478
					93	7.3± 3.5	97.56± 2.4	68	0.419	5.0± 4.1	96.26± 2.2	64	0.485	4.8± 4.1	96.21± 2.1	64	0.493
					90	7.3± 3.5	97.29± 2.9	70	0.413	4.9± 4.0	95.99± 2.6	64	0.483	4.8± 4.0	95.93± 2.5	64	0.543
vote	(17	81)	89.66	5.3± 1.4	98	5.3± 1.4	100± 0.0	100	0.000	5.3± 1.3	99.95± 0.2	100	0.007	4.6± 1.1	99.60± 0.4	64	0.014
					95	5.3± 1.4	100± 0.0	100	0.000	5.3± 1.3	99.93± 0.3	100	0.008	4.1± 1.0	98.25± 1.7	64	0.018
					93	5.3± 1.4	100± 0.0	100	0.000	5.2± 1.3	99.78± 1.1	100	0.012	4.1± 0.9	98.10± 1.9	64	0.018
					90	5.3± 1.4	100± 0.0	100	0.000	5.2± 1.3	99.78± 1.1	100	0.012	4.0± 1.2	97.24± 3.1	64	0.022
kr-vs-kp	(37	200)	88.07	12.2± 3.9	98	7.8± 4.2	99.19± 0.5	64	0.387	6.5± 4.7	98.99± 0.4	64	0.427	6.1± 4.9	98.88± 0.3	43	0.457
					95	7.3± 3.9	98.29± 1.4	64	0.416	6.0± 4.3	97.89± 1.1	64	0.453	5.5± 4.5	97.79± 0.9	43	0.462
					93	6.9± 3.5	97.21± 2.5	69	0.422	5.6± 3.8	96.82± 2.2	64	0.448	5.2± 4.0	96.71± 2.1	43	0.468
					90	6.8± 3.5	96.65± 3.1	69	0.418	5.4± 3.8	95.69± 3.0	64	0.468	5.0± 4.0	95.59± 2.8	61	0.487
mushroom	(23	200)	95.51	10.7± 2.3	98	7.5± 2.4	98.99± 0.7	90	0.641	6.5± 2.6	98.74± 0.5	83	0.751	6.3± 2.7	98.70± 0.4	18	0.828
					95	6.5± 2.6	97.35± 1.8	96	1.011	5.1± 2.5	96.52± 1.0	90	1.130	5.0± 2.5	96.39± 0.8	54	1.113
					93	5.8± 2.8	95.77± 2.7	96	1.257	4.4± 2.5	94.67± 1.6	94	1.297	4.2± 2.4	94.48± 1.3	65	1.324
					90	5.3± 3.0	94.01± 3.9	97	1.455	3.8± 2.3	92.36± 2.2	96	1.543	3.6± 2.2	92.07± 1.6	76	1.650
threeO9	(10	103)	83.13	4.2± 0.4	98	4.2± 0.4	100± 0.0	100	0.000	4.2± 0.4	100± 0.0	100	0.000	4.2± 0.4	100± 0.0	78	0.001
					95	4.2± 0.4	100± 0.0	100	0.000	4.2± 0.4	100± 0.0	100	0.000	4.0± 0.2	99.23± 1.4	100	0.002
					93	4.2± 0.4	100± 0.0	100	0.000	4.2± 0.4	100± 0.0	100	0.000	3.9± 0.2	99.20± 1.5	100	0.002
					90	4.2± 0.4	100± 0.0	100	0.000	4.2± 0.4	100± 0.0	100	0.000	3.8± 0.4	98.29± 3.3	100	0.003
xd6	(10	176)	81.36	4.5± 0.9	98	4.5± 0.8	100± 0.0	100	0.000	4.5± 0.8	100± 0.0	100	0.000	4.5± 0.8	100± 0.0	73	0.001
					95	4.5± 0.8	100± 0.0	100	0.000	4.5± 0.8	100± 0.0	100	0.000	4.5± 0.8	100± 0.0	73	0.001
					93	4.5± 0.8	100± 0.0	100	0.000	4.5± 0.8	100± 0.0	100	0.000	4.3± 0.4	98.30± 2.7	73	0.001
					90	4.5± 0.8	100± 0.0	100	0.000	4.5± 0.8	100± 0.0	100	0.000	4.3± 0.4	98.30± 2.7	73	0.002
mamo	(14	53)	80.21	4.9± 0.8	98	4.9± 0.7	100± 0.0	100	0.000	4.9± 0.7	100± 0.0	100	0.000	4.6± 0.6	99.66± 0.5	53	0.007
					95	4.9± 0.7	100± 0.0	100	0.000	4.9± 0.7	100± 0.0	100	0.000	3.9± 0.6	97.80± 1.6	85	0.009
					93	4.9± 0.7	100± 0.0	100	0.000	4.9± 0.7	100± 0.0	100	0.000	3.9± 0.6	97.68± 1.7	85	0.009
					90	4.9± 0.7	100± 0.0	100	0.000	4.9± 0.7	100± 0.0	100	0.000	3.6± 0.8	96.18± 3.2	96	0.011
tumor	(16	104)	83.21	5.3± 0.9	98	5.3± 0.8	100± 0.0	100	0.000	5.2± 0.7	99.96± 0.2	100	0.008	4.1± 0.7	99.41± 0.5	91	0.012
					95	5.3± 0.8	100± 0.0	100	0.000	5.2± 0.6	99.83± 0.7	100	0.012	3.2± 0.6	96.02± 1.5	94	0.016
					93	5.3± 0.8	100± 0.0	100	0.000	5.2± 0.6	99.74± 1.2	100	0.014	3.1± 0.7	95.50± 1.4	95	0.016
					90	5.3± 0.8	100± 0.0	100	0.000	5.1± 0.7	99.67± 1.4	100	0.016	3.0± 0.6	95.30± 1.6	95	0.017

Table 13: Assessing LmPAXp explanations for NBCs. Columns #F and #I show, respectively, number of features and tested instances in the Dataset. Column A% reports (in %) the training accuracy of the classifier. Column δ reports (in %) the value of the parameter δ. LmPAXp_{≤9}, LmPAXp_{≤7} and LmPAXp_{≤4} denote, respectively, LmPAXp’s of (target) length 9, 7 and 4. Columns Length and Precision report, respectively, the average explanation length and the average explanation precision (± denotes the standard deviation). W% shows (in %) the number of success/wins where the explanation size is less than or equal to the target size. Finally, the average runtime to compute an explanation is shown (in seconds) in Time. (Note that the reported average time is the mean of runtimes for instances for which we effectively computed an approximate explanation, namely instances that have AXp’s of length longer than the target length; whereas for the remaining instances the trimming process is skipped and the runtime is 0 sec, thus we exclude them when calculating the average.)

distribution.) Also, the prototype implementation was tested with varying the threshold δ while Anchor runs guided by its own metric.

Results. Table 12 summarizes the results of our experiments for the case study of DTs. One can observe that MinPAXp and LmPAXp compute succinct explanations (i.e. of average size 7 ± 2 [69]), for the majority of tested instances across all datasets, noticeably shorter than consistent-path explanations. More importantly, the computed explanations are trustworthy

and show good quality precision, e.g. *dermatology*, *soybean* and *texture* show average precisions greater than 98% for all values of δ . Additionally, the results clearly demonstrate that our proposed SMT encoding scales for deep DTs with runtimes on average less than 20 sec for the largest encodings while the runtimes of LmPAXp are negligible, never exceeding 0.01 sec. Also, observe from the table (see column \mathbf{m}_{\subseteq}) that the over-approximations computed by LmPAXp are often subset-minimal PAXp’s, and often as short as computed MinPAXp’s. This confirms empirically the advantages of computing LmPAXp’s, i.e. in practice one may rely on the computation of LmPAXp’s, which pays off in terms of (1) performance, (2) sufficiently high probabilistic guarantees of precision, and (3) good quality over-approximation of subset-minimal PAXp’s. In contrast, Anchor is unable to provide precise and succinct explanations in both settings of data and uniform distribution. Moreover, we observe that Anchor’s explanations often include features that are not involved in the consistent path, e.g. for *texture* less than 20% of an explanation is shared with the consistent path. (This trend was also pointed out by [39].) In terms of average runtime, Anchor is overall slower, being outperformed by the computation of LmPAXp by several orders of magnitude.

Overall, the experiments demonstrate that our approach efficiently computes succinct and provably precise explanations for large DTs. The results also substantiate the limitations of model-agnostic explainers, both in terms of explanation quality and computation time.

7.2 Case Study 2: Naive Bayes Classifiers

Prototype implementation. A prototype implementation of the proposed approach for computing relevant sets for NBCs was developed in Python. To compute AXp’s, we use the Perl script implemented by [62]¹¹. The prototype implementation was tested with varying thresholds $\delta \in \{0.90, 0.93, 0.95, 0.98\}$. When converting probabilities from real values to integer values, the selected number of decimal places is 3. (As outlined earlier, we observed that there is a negligible accuracy loss from using three decimal places). In order to produce explanations of size admissible for the cognitive capacity of human decision makers [69], we selected three different target sizes for the explanations to compute: 9, 7 and 4, and we computed a LmPAXp for the input instance when its AXp \mathcal{X} is larger than the target size (recall that \mathcal{S} is initialized to \mathcal{X}); otherwise we consider that the AXp is succinct and the explainer returns \mathcal{X} . For example, assume the target size is 7, an instance \mathbf{v}_1 with an AXp \mathcal{S}_1 of 5 features and a second instance \mathbf{v}_2 with an AXp \mathcal{S}_2 of 8 features, then for \mathbf{v}_1 the output will be \mathcal{S}_1 and for \mathbf{v}_2 the output will be a subset of \mathcal{S}_2 .

Benchmarks. The benchmarks used in this evaluation originate from the UCI ML Repository [95] and Penn ML Benchmarks [79]. The number of training data (resp. features) in the target datasets varies from 336 to 14113 (resp. 10 to 37) and on average is 3999.1 (resp. 20.0). All the NBCs are trained using the learning tool *scikit-learn* [88]. The data split for training and test data is set to 80% and 20%, respectively. Model accuracies are above 80% for the training accuracy and above 75% for the test accuracy. For each dataset, we run the explainer on 200 instances randomly picked from the test data or on all instances if there are less than 200.

11. Publicly available from: <https://github.com/jpmarquessilva/expplc>

Results. Table 13 summarizes the results of our experiments for the case study of NBCs. For all tested values for the parameter threshold δ and target size, the reported results show the sizes and precisions of the computed explanations. As can be observed for all considered settings, the approximate explanations are succinct, in particular the average sizes of the explanations are invariably lower than the target sizes. Moreover, these explanations offer strong guarantees of precision, as their average precisions are strictly greater than δ with significant gaps (e.g. above 97%, in column $\text{LmPAXp}_{\leq 7}$, for datasets *adult*, *vote*, *threeOf9*, *xd6*, *mamo* and *tumor* and above 95% for *chess* and *kr-vs-kp*). An important observation from the results, is the gain of succinctness (explanation size) when comparing AXp’s with LmPAXp’s. Indeed, for some datasets, the AXp’s are too large (e.g. for *chess* and *kr-vs-kp* datasets, the average number of features in the AXp’s is 12), exceeding the cognitive limits of human decision makers [69] (limited to 7 ± 2 features). To illustrate that, one can focus on the dataset *agaricus* or *mushroom* and see that for a target size equal to 7 and $\delta = 0.95$, the average length of the LmPAXp’s (i.e. 5.3 and 5.1, resp.) is 2 times less than the average length of the AXp’s (i.e. 10.3 and 10.7, resp.). Besides, the results show that $\delta = 0.95$ is a good probability threshold to guarantee highly precise and short approximate explanations.

Despite the computational complexity of the proposed approach being pseudo-polynomial, the results demonstrate that in practice the algorithm is effective and scales for large datasets. As can be seen, the runtimes are negligible for all datasets, never exceeding 2 seconds for the largest datasets (i.e. *agaricus* or *mushroom*) and the average is 0.33 seconds for all tested instances across all datasets and all settings. Furthermore, we point out that the implemented prototype was tested with 4 decimal places to assess further the scalability of the algorithm on larger DP tables, and the results show that computing LmPAXp’s is still feasible, e.g. with *agaricus* the average runtime for δ set to 0.95 and target size to 7 is 10.08 seconds and 7.22 seconds for $\delta = 0.98$.

The table also reports the number of explanations being shorter than or of size equal to the target size over the total number of tested instances. We observe that for both settings $\text{LmPAXp}_{\leq 9}$ and $\text{LmPAXp}_{\leq 7}$ and for the majority of datasets and with a few exceptions the fraction is significantly high, e.g. varying for 96% to 100% for *adult* dataset. However, for $\text{LmPAXp}_{\leq 4}$ despite the poor percentage of wins for some datasets, it is the case that the average lengths of computed explanations are close to 4.

Overall, the experiments demonstrate that our approach efficiently computes succinct and provably precise explanations for NBCs. The results also showcase empirically the advantage of the algorithm, i.e. in practice one may rely on the computation of LmPAXp’s, which pays off in terms of (1) performance, (2) succinctness and (3) sufficiently high probabilistic guarantees of precision.

7.3 Case Study 3: Graph-Based Classifiers

Prototype implementation. A prototype implementation of the proposed algorithms from computing models and assessing explanation precision was implemented in Python. A deletion-based procedure was used to compute LmPAX’s and an SMT-based approach, similarly like DTs, was adopted for extracting MinPAXp’s. Hence, the z3 solver was employed to perform SMT oracle calls to the SMT encoding. Moreover, OMDD’s were

Dataset	#I	#F	OMDD		δ	MinPAXp					LmPAXp					
						Length			Prec	Time	Length			Prec	m_{\subseteq}	Time
						#N	A%	M	m	avg	avg	avg	M	m	avg	avg
corral	100	6	15	90.6	100	3	3	3.0	100	0.02	3	3	3.0	100	100	0.00
					95	3	3	3.0	100	0.02	3	3	3.0	100	100	0.00
					90	2	2	2.0	93.8	0.02	2	2	2.0	93.8	100	0.00
lending	100	9	1103	81.7	100	9	6	8.0	100	24.24	9	6	7.9	100	100	1.57
					95	9	5	7.8	99.7	21.48	9	6	7.8	99.8	100	1.49
					90	9	4	7.2	96	24.65	9	5	7.4	97.0	100	1.48
monk2	100	6	70	79.3	100	6	4	5.1	100	0.10	6	4	5.1	100	100	0.03
					95	6	4	5.1	100	0.09	6	4	5.1	100	100	0.03
					90	6	3	4.8	98.1	0.09	6	3	4.8	98.1	100	0.03
postoperative	74	8	109	80	100	8	4	6.1	100	0.26	8	4	6.2	100	100	0.04
					95	8	2	6.0	99.3	0.25	8	2	6.0	99.3	100	0.04
					90	8	2	5.3	95.9	0.23	8	2	5.4	96.6	94.6	0.04
tic_tac_toe	100	9	424	70.3	100	9	5	7.7	100	3.60	9	5	7.8	100	100	0.38
					95	9	5	7.5	99.5	3.24	9	5	7.7	99.6	99.0	0.38
					90	9	3	7.3	98.3	4.06	9	3	7.5	98.6	98.0	0.38
xd6	100	9	76	83.1	100	9	4	4.6	100	0.10	9	4	4.6	100	100	0.03
					95	9	3	3.8	97	0.09	9	3	3.8	97.0	99.0	0.03
					90	9	3	3.3	94.8	0.10	9	3	3.4	94.6	100	0.03

Table 14: Assessing MinPAXp and LmPAXp explanations of OMDDs. Columns **#I**, **#F** denote, resp. the number of tested instances and the number of features. In column **OMDD**, **N** and **A** denote, resp., the number of nodes and the test accuracy of the OMDD. Column δ reports (in %) the value of the threshold δ . In column **Length**, **avg** (resp. **M** and **m**) denotes the average (resp. max. and min.) length of the explanations. **Prec** reports (in %) the average precision (defined in (9)) of resulting explanations. m_{\subseteq} shows the number in (%) of LmPAXp’s that are subset-minimal, i.e. PAXp’s. **Time** reports (in seconds) the average runtime to compute an explanation.

built heuristically using a publicly available package MEDDLY¹², which is implemented in C/C++.

Benchmarks. The benchmarks used in this experimental study all originate from Penn ML Benchmarks [79]. For each dataset, we picked a consistent subset of samples (i.e. no two instances are contradictory) for building OMDDs. The picked subset of the data are randomly split into training (80%) and test set (20%). To assess the explanation precisions and the runtimes of our algorithms, we test for each dataset 100 instances picked randomly or all instances if there are less than 100 rows in the dataset.

Results. Table 14 summarizes the results of our experiments. As can be observed from the column m_{\subseteq} , at least 94.6% of the computed LmPAXp’s are indeed MinPAXp’s. As indicated by the column **Prec**, both LmPAXp and MinPAXp offer good precision. On average, the precision of the computed explanation is greater than or equal to 93.8% for the values of δ we considered. Regarding the succinctness of the computed explanations, compared with the case $\delta = 100$, when $\delta = 95$, the size of computed explanations almost remain unchanged. When $\delta = 90$, for dataset *monk2*, *lending* and *tic_tac_toe*_, the reduction is still negligible.

12. <https://asminer.github.io/meddly/>

But for dataset *corral*, *postoperative* and *xd6*, the reduction of size is at least 13%. Not surprisingly, the time for computing LmPAXp’s is almost negligible, while the computation time for MinPAXp’s cannot be overlooked; for example, the average time for computing one MinPAXp of dataset *lending* is around 20s. Moreover, even though the runtime for computing a MinPAXp is at least an order of magnitude larger than the time for computing a LmPAXp, the average size of the MinPAXp is not significantly smaller than that of the LmPAXp. Overall, the experiments demonstrate that our approach efficiently computes succinct and provably precise explanations for OMDDs.

8. Related Work

Work on probabilistic (abductive) explanations (or probabilistic prime implicants) can be traced to [101, 100]. There has been recent progress on computing probabilistic abductive explanations of decision trees [47, 48, 3, 4]. Moreover, there is also recent work on computing probabilistic abductive explanations of naive Bayes classifiers [50]. The work reported in this paper builds on our own recent work [47, 48, 50] with the focus being the efficient computation of probabilistic explanations, both for decision trees and naive Bayes classifiers. More recent work [3, 4] targets the complexity of computing probabilistic explanations for decision trees. The experimental results (see Section 7) confirm the practicality of the work proposed in this paper.

9. Conclusions

Explanation size is one of the most visible limitations of formal approaches for explaining the predictions of ML classifiers. To address this limitation, recent work established the NP^{PP} -hardness of computing a probabilistic explanation known as a δ -relevant set: a set of features which, when identical to those of the vector to be explained, is sufficient to guarantee the same output with probability at least δ [101, 100]. As acknowledged by earlier work [101, 100], the hardness result makes the problem of exactly computing (minimum-size) δ -relevant sets unrealistic to solve in practice, at least for general classifiers. As a result, instead of considering the general problem of computing (minimum-size) δ -relevant sets, this paper tackles the problem’s complexity by analyzing instead restricted classes of classifiers, and by considering related but different definitions of relevant sets, that include subset-minimal and locally-minimal relevant sets. Furthermore, the paper proves that, for several families of classifiers that include decision trees, graph-based classifiers, and propositional classifiers, subset-minimal relevant sets can be computed in polynomial time using an oracle for NP, and that locally-minimal relevant sets can be computed in polynomial time. In the case of naive Bayes classifiers, we obtain similar results, but in pseudo-polynomial deterministic and non-deterministic time. The experimental results validate the practical interest of computing relevant sets, either locally-minimal or subset-minimal.

Future work will extend the results presented in this paper, e.g. by considering additional families of classifiers, but also by further validating the quality of locally-minimal relevant sets. Additional optimizations to the algorithms proposed can also be envisioned. For example, additional heuristics could be considered for selecting smaller AXp’s, e.g. by picking the smallest of a number of computed AXp’s. Such improvements will not change

the conclusions of the present paper, but can serve to further improve the quality of the results the paper reports.

Acknowledgments. This work was supported by the AI Interdisciplinary Institute ANITI, funded by the French program “Investing for the Future – PIA3” under Grant agreement no. ANR-19-PI3A-0004, and by the H2020-ICT38 project COALA “Cognitive Assisted agile manufacturing for a Labor force supported by trustworthy Artificial intelligence”.

References

- [1] J. Adebayo, J. Gilmer, M. Muehly, I. J. Goodfellow, M. Hardt, and B. Kim. Sanity checks for saliency maps. In *NeurIPS*, pages 9525–9536, 2018.
- [2] L. Amgoud. Non-monotonic explanation functions. In *ECSQARU*, pages 19–31, 2021.
- [3] M. Arenas, P. Barceló, M. Romero, and B. Subercaseaux. On computing probabilistic explanations for decision trees. *CoRR*, abs/2207.12213, 2022.
- [4] M. Arenas, P. Barceló, M. Romero, and B. Subercaseaux. On computing probabilistic explanations for decision trees. In *NeurIPS*, 2022.
- [5] S. Arora and B. Barak. *Computational Complexity - A Modern Approach*. Cambridge University Press, 2009.
- [6] G. Audemard, S. Bellart, L. Bounia, F. Koriche, J. Lagniez, and P. Marquis. On the computational intelligibility of boolean classifiers. In *KR*, pages 74–86, 2021.
- [7] G. Audemard, F. Koriche, and P. Marquis. On tractable XAI queries based on compiled representations. In *KR*, pages 838–849, 2020.
- [8] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one*, 10(7):e0130140, 2015.
- [9] D. Barber. *Bayesian reasoning and machine learning*. Cambridge University Press, 2012.
- [10] C. W. Barrett, R. Sebastiani, S. A. Seshia, and C. Tinelli. Satisfiability modulo theories. In A. Biere, M. Heule, H. van Maaren, and T. Walsh, editors, *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*, pages 825–885. IOS Press, 2009.
- [11] D. Bertsimas and J. Dunn. Optimal classification trees. *Mach. Learn.*, 106(7):1039–1082, 2017.
- [12] A. Biere, M. Heule, H. van Maaren, and T. Walsh, editors. *Handbook of Satisfiability*. IOS Press, 2021.
- [13] R. Boumazouza, F. C. Alili, B. Mazure, and K. Tabia. ASTERYX: A model-agnostic sat-based approach for symbolic and score-based explanations. In *CIKM*, pages 120–129, 2021.

- [14] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification and Regression Trees*. Wadsworth, 1984.
- [15] R. E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Trans. Computers*, 35(8):677–691, 1986.
- [16] O. Camburu, E. Giunchiglia, J. Foerster, T. Lukasiewicz, and P. Blunsom. Can I trust the explainer? Verifying post-hoc explanatory methods. *CoRR*, abs/1910.02065, 2019.
- [17] S. A. Cook. The complexity of theorem-proving procedures. In *STOC*, pages 151–158, 1971.
- [18] M. C. Cooper and J. Marques-Silva. On the tractability of explaining decisions of classifiers. In L. D. Michel, editor, *CP*, pages 21:1–21:18, 2021.
- [19] A. Darwiche. Decomposable negation normal form. *Journal of the ACM (JACM)*, 48(4):608–647, 2001.
- [20] A. Darwiche. On the tractable counting of theory models and its application to truth maintenance and belief revision. *J. Appl. Non Class. Logics*, 11(1-2):11–34, 2001.
- [21] A. Darwiche. SDD: A new canonical representation of propositional knowledge bases. In *IJCAI*, pages 819–826, 2011.
- [22] A. Darwiche and A. Hirth. On the reasons behind decisions. In *ECAI*, pages 712–720, 2020.
- [23] A. Darwiche and P. Marquis. A knowledge compilation map. *J. Artif. Intell. Res.*, 17:229–264, 2002.
- [24] L. M. de Moura and N. S. Bjørner. Z3: an efficient SMT solver. In *TACAS*, pages 337–340, 2008.
- [25] B. Dimanov, U. Bhatt, M. Jamnik, and A. Weller. You shouldn’t trust me: Learning models which conceal unfairness from multiple explanation methods. In *ECAI*, pages 2473–2480, 2020.
- [26] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern classification*. John Wiley & Sons, 1973.
- [27] M. E. Dyer. Approximate counting by dynamic programming. In *STOC*, pages 693–699, 2003.
- [28] EU. Artificial Intelligence Act. <http://tiny.cc/ahcnuz>, 2021.
- [29] N. Friedman, D. Geiger, and M. Goldszmidt. Bayesian network classifiers. *Mach. Learn.*, 29(2-3):131–163, 1997.
- [30] P. Gawrychowski, L. Markin, and O. Weimann. A faster FPTAS for #knapsack. In *ICALP*, pages 64:1–64:13, 2018.

- [31] P. Gopalan, A. R. Klivans, R. Meka, D. Stefankovic, S. S. Vempala, and E. Vigoda. An FPTAS for #knapsack and related counting problems. In *FOCS*, pages 817–826, 2011.
- [32] N. Gorji and S. Rubin. Sufficient reasons for classifier decisions in the presence of domain constraints. In *AAAI*, February 2022.
- [33] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi. A survey of methods for explaining black box models. *ACM Comput. Surv.*, 51(5):93:1–93:42, 2019.
- [34] X. Huang, Y. Izza, A. Ignatiev, M. C. Cooper, N. Asher, and J. Marques-Silva. Tractable explanations for d-DNNF classifiers. In *AAAI*, February 2022.
- [35] X. Huang, Y. Izza, A. Ignatiev, and J. Marques-Silva. On efficiently explaining graph-based classifiers. In *KR*, pages 356–367, 2021.
- [36] X. Huang, Y. Izza, A. Ignatiev, and J. Marques-Silva. On efficiently explaining graph-based classifiers. *CoRR*, abs/2106.01350, 2021.
- [37] X. Huang and J. Marques-Silva. On deciding feature membership in explanations of SDD & related classifiers. *CoRR*, abs/2202.07553, 2022.
- [38] IAI. Interpretable AI. <https://www.interpretable.ai/>, 2020.
- [39] A. Ignatiev. Towards trustable explainable AI. In *IJCAI*, pages 5154–5158, 2020.
- [40] A. Ignatiev, Y. Izza, P. Stuckey, and J. Marques-Silva. Using MaxSAT for efficient explanations of tree ensembles. In *AAAI*, February 2022.
- [41] A. Ignatiev and J. Marques-Silva. SAT-based rigorous explanations for decision lists. In *SAT*, pages 251–269, 2021.
- [42] A. Ignatiev, N. Narodytska, N. Asher, and J. Marques-Silva. From contrastive to abductive explanations and back again. In *AixIA*, pages 335–355, 2020.
- [43] A. Ignatiev, N. Narodytska, and J. Marques-Silva. Abduction-based explanations for machine learning models. In *AAAI*, pages 1511–1519, 2019.
- [44] A. Ignatiev, N. Narodytska, and J. Marques-Silva. On relating explanations and adversarial examples. In *NeurIPS*, pages 15857–15867, 2019.
- [45] Y. Izza, A. Ignatiev, and J. Marques-Silva. On explaining decision trees. *CoRR*, abs/2010.11034, 2020.
- [46] Y. Izza, A. Ignatiev, and J. Marques-Silva. On tackling explanation redundancy in decision trees. *CoRR*, abs/2205.09971, 2022.
- [47] Y. Izza, A. Ignatiev, N. Narodytska, M. C. Cooper, and J. Marques-Silva. Efficient explanations with relevant sets. *CoRR*, abs/2106.00546, 2021.

- [48] Y. Izza, A. Ignatiev, N. Narodytska, M. C. Cooper, and J. Marques-Silva. Provably precise, succinct and efficient explanations for decision trees. *CoRR*, abs/2205.09569, 2022.
- [49] Y. Izza and J. Marques-Silva. On explaining random forests with SAT. In *IJCAI*, pages 2584–2591, 2021.
- [50] Y. Izza and J. Marques-Silva. On computing relevant features for explaining NBCs. *CoRR*, abs/2207.04748, 2022.
- [51] B. Juba. Learning abductive reasoning using random examples. In *AAAI*, pages 999–1007, 2016.
- [52] U. Junker. QUICKXPLAIN: preferred explanations and relaxations for over-constrained problems. In *AAAI*, pages 167–172, 2004.
- [53] T. Kam and R. K. Brayton. Multi-valued decision diagrams. Technical Report UCB/ERL M90/125, EECS Department, University of California, Berkeley, Dec 1990.
- [54] H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack problems*. Springer, 2004.
- [55] P. Kindermans, S. Hooker, J. Adebayo, M. Alber, K. T. Schütt, S. Dähne, D. Erhan, and B. Kim. The (un)reliability of saliency methods. In W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen, and K. Müller, editors, *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, volume 11700 of *Lecture Notes in Computer Science*, pages 267–280. Springer, 2019.
- [56] F. Koriche, J. Lagniez, P. Marquis, and S. Thomas. Knowledge compilation for model counting: Affine decision trees. In *IJCAI*, pages 947–953, 2013.
- [57] H. Lakkaraju and O. Bastani. “How do I fool you?”: Manipulating user trust via misleading black box explanations. In *AIES*, pages 79–85, 2020.
- [58] M. H. Liffiton, A. Previti, A. Malik, and J. Marques-Silva. Fast, flexible MUS enumeration. *Constraints An Int. J.*, 21(2):223–250, 2016.
- [59] X. Liu and E. Lorini. A logic for binary classifiers and their explanation. In *CLAR*, pages 302–321, 2021.
- [60] S. M. Lundberg and S. Lee. A unified approach to interpreting model predictions. In *NeurIPS*, pages 4765–4774, 2017.
- [61] E. L. Malfa, R. Michelmore, A. M. Zbrzezny, N. Paoletti, and M. Kwiatkowska. On guaranteed optimal robust explanations for NLP models. In *IJCAI*, pages 2658–2665, 2021.
- [62] J. Marques-Silva, T. Gerspacher, M. C. Cooper, A. Ignatiev, and N. Narodytska. Explaining naive bayes and other linear classifiers with polynomial time and delay. In *NeurIPS*, 2020.

- [63] J. Marques-Silva, T. Gerspacher, M. C. Cooper, A. Ignatiev, and N. Narodytska. Explanations for monotonic classifiers. In *ICML*, pages 7469–7479, 2021.
- [64] J. Marques-Silva and A. Ignatiev. Delivering trustworthy AI through formal XAI. In *AAAI*, 2022.
- [65] J. Marques-Silva, M. Janota, and A. Belov. Minimal sets over monotone predicates in boolean formulae. In *CAV*, pages 592–607, 2013.
- [66] J. Marques-Silva, M. Janota, and C. Mencia. Minimal sets on propositional formulae. problems and reductions. *Artif. Intell.*, 252:22–50, 2017.
- [67] J. Marques-Silva and A. Previti. On computing preferred MUSes and MCSes. In *SAT*, pages 58–74, 2014.
- [68] J. S. Mill. *A System of Logic, Ratiocinative and Inductive*, volume 1. John W. Parker, 1843.
- [69] G. A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review*, 63(2):81–97, 1956.
- [70] T. Miller. Explanation in artificial intelligence: Insights from the social sciences. *Artif. Intell.*, 267:1–38, 2019.
- [71] C. Molnar. *Interpretable Machine Learning*. Leanpub, 2020. <http://tiny.cc/6c76tz>.
- [72] G. Montavon, W. Samek, and K. Müller. Methods for interpreting and understanding deep neural networks. *Digit. Signal Process.*, 73:1–15, 2018.
- [73] C. J. Muise, S. A. McIlraith, J. C. Beck, and E. I. Hsu. Dsharp: Fast d-DNNF compilation with sharpSAT. In *CAI*, pages 356–361, 2012.
- [74] N. Narodytska, A. A. Shrotri, K. S. Meel, A. Ignatiev, and J. Marques-Silva. Assessing heuristic machine learning explanations with model counting. In *SAT*, pages 267–278, 2019.
- [75] A. Niveau, H. Fargier, and C. Pralet. Representing csps with set-labeled diagrams: A compilation map. In *Graph Structures for Knowledge Representation and Reasoning*, pages 137–171. Springer, 2012.
- [76] J. J. Oliver. Decision graphs – an extension of decision trees. Technical Report 92/173, Monash University, 1992.
- [77] J. J. Oliver, D. L. Dowe, and C. S. Wallace. Inferring decision graphs using the minimum message length principle. In *AAI*, 1992.
- [78] J. D. Park. Using weighted MAX-SAT engines to solve MPE. In *AAAI*, pages 682–687, 2002.
- [79] Penn Machine Learning Benchmarks. <https://github.com/EpistasisLab/penn-ml-benchmarks>.

- [80] J. R. Quinlan. Induction of decision trees. *Mach. Learn.*, 1(1):81–106, 1986.
- [81] M. T. Ribeiro, S. Singh, and C. Guestrin. “Why should I trust you?”: Explaining the predictions of any classifier. In *KDD*, pages 1135–1144, 2016.
- [82] M. T. Ribeiro, S. Singh, and C. Guestrin. Anchors: High-precision model-agnostic explanations. In *AAAI*, pages 1527–1535, 2018.
- [83] R. Rizzi and A. I. Tomescu. Faster FPTASes for counting and random generation of knapsack solutions. *Inf. Comput.*, 267:135–144, 2019.
- [84] C. Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.
- [85] W. Samek, G. Montavon, S. Lapuschkin, C. J. Anders, and K. Müller. Explaining deep neural networks and beyond: A review of methods and applications. *Proc. IEEE*, 109(3):247–278, 2021.
- [86] W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen, and K. Müller, editors. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. Springer, 2019.
- [87] W. Samek and K. Müller. Towards explainable artificial intelligence. In Samek et al. [86], pages 5–22.
- [88] scikit-learn: Machine Learning in Python. <https://scikit-learn.org/>.
- [89] A. Shih, A. Choi, and A. Darwiche. A symbolic approach to explaining bayesian network classifiers. In *IJCAI*, pages 5103–5111, 2018.
- [90] A. Shih, A. Choi, and A. Darwiche. Compiling bayesian network classifiers into decision graphs. In *AAAI*, pages 7966–7974, 2019.
- [91] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *ICLR*, 2014.
- [92] L. Sixt, M. Granz, and T. Landgraf. When explanations lie: Why many modified BP attributions fail. In *ICML*, pages 9046–9057, 2020.
- [93] D. Slack, S. Hilgard, E. Jia, S. Singh, and H. Lakkaraju. Fooling LIME and SHAP: adversarial attacks on post hoc explanation methods. In *AIES*, pages 180–186, 2020.
- [94] A. Srinivasan, T. Ham, S. Malik, and R. K. Brayton. Algorithms for discrete function manipulation. In *1990 IEEE international conference on computer-aided design*, pages 92–93. IEEE Computer Society, 1990.
- [95] UCI Machine Learning Repository. <https://archive.ics.uci.edu/ml>.
- [96] P. E. Utgoff, N. C. Berkman, and J. A. Clouse. Decision tree induction based on efficient tree restructuring. *Mach. Learn.*, 29(1):5–44, 1997.
- [97] L. G. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, 1984.

- [98] G. Van den Broeck and A. Darwiche. On the role of canonicity in knowledge compilation. In *AAAI*, pages 1641–1648, 2015.
- [99] T. J. Viering, Z. Wang, M. Loog, and E. Eisemann. How to manipulate cnns to make them lie: the GradCAM case. *CoRR*, abs/1907.10901, 2019.
- [100] S. Wäldchen. *Towards Explainable Artificial Intelligence – Interpreting Neural Network Classifiers with Probabilistic Prime Implicants*. PhD thesis, Technischen Universität Berlin, March 2022.
- [101] S. Wäldchen, J. MacDonald, S. Hauch, and G. Kutyniok. The computational complexity of understanding binary classifier decisions. *J. Artif. Intell. Res.*, 70:351–387, 2021.
- [102] E. Wang, P. Khosravi, and G. V. den Broeck. [Probabilistic Sufficient Explanations](#). In *IJCAI*, pages 3082–3088, 2021.
- [103] L. Wolf, T. Galanti, and T. Hazan. A formal approach to explainability. In *AIES*, pages 255–261, 2019.
- [104] J. Yu, A. Ignatiev, P. J. Stuckey, N. Narodytska, and J. Marques-Silva. Eliminating the impossible, whatever remains must be true. *CoRR*, abs/2206.09551, 2022.