

Provable Guarantees on the Robustness of Decision Rules to Causal Interventions

Benjie Wang*, Clare Lyle* and Marta Kwiatkowska

University of Oxford
benjie.wang@cs.ox.ac.uk

Abstract

Robustness of decision rules to shifts in the data-generating process is crucial to the successful deployment of decision-making systems. Such shifts can be viewed as interventions on a causal graph, which capture (possibly hypothetical) changes in the data-generating process, whether due to natural reasons or by the action of an adversary. We consider causal Bayesian networks and formally define the *interventional robustness* problem, a novel *model-based* notion of robustness for decision functions that measures worst-case performance with respect to a set of interventions that denote changes to parameters and/or causal influences. By relying on a tractable representation of Bayesian networks as arithmetic circuits, we provide efficient algorithms for computing guaranteed upper and lower bounds on the interventional robustness probabilities. Experimental results demonstrate that the methods yield useful and interpretable bounds for a range of practical networks, paving the way towards provably causally robust decision-making systems.

1 Introduction

As algorithmic decision-making systems become widely deployed, there has been an increasing focus on their safety and robustness, particularly when they are applied to input points outside of the data distribution they were trained on. Much of the work in this area has focused on instance-based robustness properties of classifiers, which guarantee that the prediction does not change in some vicinity of a specific input point [Shih *et al.*, 2018; Narodytska *et al.*, 2018]. However, there are many types of distribution shift that cannot be characterized by robustness against norm-bounded perturbations to individual inputs. Such distribution shifts are often instead characterized by causal interventions on the data-generating process [Quionero-Candela *et al.*, 2009; Zhang *et al.*, 2015; Lipton *et al.*, 2018]. These interventions give rise to a range of different environments (distributions), which can be the

effect of natural shifts (e.g. different country) or actions of other agents (e.g. a hospital changing prescription policy).

To assess the impact of such interventions, we must leverage knowledge about the causal structure of the data-generating distribution. This paper concerns itself with a simple question: given a decision-making system and a posited causal model, is the system robust to a set of plausible interventions to the causal model? Defining and verifying such *model-based* notions of robustness requires a formal representation of the decision-making system. For discrete input features and a discrete output class, regardless of how a classifier is learned, its role in decision-making can be unambiguously represented by its *decision function*, mapping features to an output class. This observation has spurred a recent trend of applying logic for meta-reasoning about classifier properties, such as monotonicity and instance-based robustness, by *compiling* the classifier into a tractable form [Shih *et al.*, 2018; Narodytska *et al.*, 2018; Audemard *et al.*, 2020], for example an ordered decision diagram. We extend this approach to causal modelling by combining logical representations of the decision rule and causal model, and compiling this joint representation into an arithmetic circuit, a tractable representation of probability distributions.

Our main technical contributions are as follows. First, we motivate and formalize the robustness of a decision rule with respect to interventions on a causal model, which we call the *interventional robustness problem*, and characterize its complexity. Second, we develop a *joint compilation* technique which allows us to reason about a causal model and decision function simultaneously. Finally, we develop and evaluate algorithms for computing upper and lower bounds on the interventional robustness problem, enabling the verification of robustness of decision-making systems to causal interventions.

Related Work

The problem of constructing classifiers which are robust to distribution shifts has received much attention from the machine learning perspective [Quionero-Candela *et al.*, 2009; Zhang *et al.*, 2015; Lipton *et al.*, 2018]. Particularly relevant to our work are *proactive* approaches to learning robust classifiers, which aim to produce classifiers that perform well across a range of environments (rather than a specific one) [Rojas-Carulla *et al.*, 2018; Subbaswamy *et al.*, 2019].

A recent line of work analyses the behaviour of machine

*Equal contribution

learning classifiers using symbolic and logical approaches by compiling these classifiers into suitable logical representations [Katz *et al.*, 2017; Narodytska *et al.*, 2018; Shi *et al.*, 2020]. Such representations can be used to answer a range of explanation and verification queries [Audemard *et al.*, 2020] about the classifier tractably, depending on the properties of the underlying propositional language [Darwiche and Marquis, 2002]. Our work uses this premise to tackle defining and verifying robustness to distribution shift, which involves not only the classifier but also a probabilistic causal model such as a causal Bayesian network.

In the Bayesian network literature, *sensitivity analysis* [Chan and Darwiche, 2004] is concerned with examining the effect of (typically small) local changes in parameter values on a target probability. We are concerned with providing worst-case guarantees against a set of possible causal interventions, which can involve changing parameters in multiple CPTs, and even altering the graphical structure of the network. This requires new methods that enable scalability to these large, potentially structural intervention sets. Our causal perspective generalizes and extends the work of Qin [2015], considering a richer class of interventions than previous work and using this perspective to prove robustness properties of a decision function.

2 Background and Notation

In the rest of this paper, we use $\mathbf{V} = (\mathbf{X}, Y, \mathbf{H})$ to denote the set of modelled variables, which includes observable features \mathbf{X} , the prediction target Y , and hidden variables \mathbf{H} . We use lower case (e.g. \mathbf{x}) to denote instantiations of variables.

Decision Functions

Consider the task of predicting Y given features \mathbf{X} . Though many machine learning (ML) techniques exist for this task, once learned, the input-output behaviour of any classifier can be characterized by means of a symbolic decision function F from \mathbf{X} to Y . For many important classes of ML methods, including Bayesian network classifiers, binarized neural networks, and random forests, it is possible to encode the corresponding decision function as a Boolean circuit Σ [Audemard *et al.*, 2020; Narodytska *et al.*, 2018; Shih *et al.*, 2019]. Such logical encodings can then be used to reason about the behaviour of the decision function, for instance providing explanations for decisions and verifying properties.

Causal Bayesian Networks

In this paper, we are interested in robust performance of decision functions under distribution shift caused by changes in the data-generating process (DGP). In order to reason about this, we first need a *causal model* of the DGP which enables such changes to be represented. We first define Bayesian networks, which are a convenient way to specify a joint distribution over the set of variables $\mathbf{V} = \{V_1, V_2, \dots, V_n\}$:

Definition 1 (Bayesian Network). A (discrete) *Bayesian network (BN)* \mathcal{N} over variables \mathbf{V} is a pair (\mathcal{G}, Θ) , where $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ is a directed acyclic graph (DAG) whose nodes correspond to the random variables \mathbf{V} and whose edges indicate conditional dependence, and where Θ denotes the set of conditional probability tables (CPTs)

$\theta_{V_i|\mathbf{U}_i}$ with parameters $\theta_{v_i|\mathbf{u}_i} = P(V_i = v_i | \mathbf{U}_i = \mathbf{u}_i)$ which specify the distribution, where $\mathbf{U}_i = \text{pa}_{\mathcal{G}}(V_i)$ are the parents of V_i in \mathcal{G} . We will denote by $p_{\mathcal{N}}$ the probability distribution defined by the BN \mathcal{N} .

Causal Bayesian networks (CBNs) are defined similarly to Bayesian networks, with the addition of *causal*, or *interventional*, semantics to the joint distribution. Intuitively, an edge (V, V') in a causal Bayesian network indicates that V *causes* V' , and the CPTs correspond to causal mechanisms. An intervention can be defined to be a change to some of these mechanisms, replacing Θ with Θ' . A CBN can thus be characterized as representing a set of distributions, each of which is generated by a different intervention.

We now define a representation of a (causal) Bayesian network, called the *network polynomial*, based on the seminal work of Darwiche [2003].

Definition 2 (Network Polynomial). The *network polynomial of causal BN* \mathcal{N} is defined to be:

$$l_{\mathcal{N}}[\lambda, \Theta] = \sum_{v_1, \dots, v_n} \prod_{i=1}^n \lambda_{v_i} \theta_{v_i|\mathbf{u}_i} \quad (1)$$

where λ_{v_i} denotes an indicator variable for each value v_i in the support of each random variable V_i , and $\theta_{v_i|\mathbf{u}_i}$ denotes each element of a CPT in Θ . Each component of the addition $l_{\mathcal{N}}[\lambda, \Theta] := \prod_{i=1}^n \lambda_{v_i} \theta_{v_i|\mathbf{u}_i}$ is called a *term*, and is associated with an instantiation \mathbf{v} .

Arithmetic circuits

Arithmetic circuits (AC) are computational graphs used to encode probability distributions over a set of discrete variables \mathbf{V} , which can tractably answer a broad range of probabilistic queries, depending on certain structural properties (called decomposability, smoothness and determinism). They were first introduced by Darwiche [2003] as a means of *compiling* Bayesian networks for the purposes of efficient inference. Subsequently they have been considered as objects of study in their own right, with proposals for directly learning ACs from data [Lowd and Domingos, 2008] and extensions relaxing determinism [Poon and Domingos, 2011].

Definition 3 (Arithmetic Circuit). An *arithmetic circuit AC* over variables \mathbf{V} and with parameters Φ is a rooted directed acyclic graph (DAG), whose internal nodes are labelled with $+$ or \times and whose leaf nodes are labelled with indicator variables λ_v , where v is the value of some variable $V \in \mathbf{V}$, or non-negative parameters ϕ .

Crucially, evaluating an arithmetic circuit can be done in time linear in the size (number of edges) of the circuit. When an AC represents a probability distribution, this means that marginals can be computed efficiently.

Like Bayesian networks, arithmetic circuits can be represented as polynomials over indicator and parameter variables, based on subcircuits [Choi and Darwiche, 2017]:

Definition 4 (Complete Subcircuit). A *complete subcircuit* α of an AC is obtained by traversing the circuit top-down, choosing one child of every visited $+$ -node and all children of every visited \times -node. The *term* $\text{term}(\alpha)$ of α is the product of all leaf nodes visited (i.e. all indicator and param-

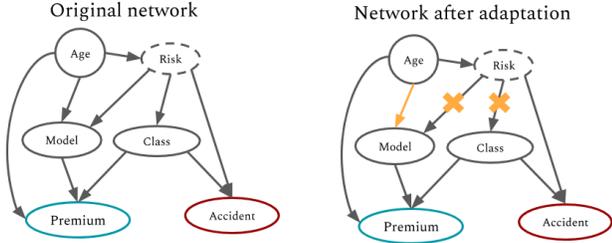


Figure 1: An example causal model describing accident risk for a car insurance problem, and illustrating how strategic adaptation to a classifier can be characterized as a change to a causal model describing how the data was generated.

ter variables). The set of all complete subcircuits is denoted α_{AC} .

Definition 5 (AC Polynomial). The AC polynomial of arithmetic circuit AC is defined to be:

$$l_{AC}[\lambda, \Phi] = \sum_{\alpha \in \alpha_{AC}} \text{term}(\alpha)$$

3 The Intervention Robustness Problem

Many distribution shifts faced by decision-making systems can be characterized by an *intervention* on the data-generating process. For example, if an insurance company gives reduced premiums to drivers who have taken a driving class, the relationship between ‘risk aversion’ and ‘class’ in Figure 1 may change in response as more risk-seeking drivers take driving classes to benefit from reduced premiums. The company therefore seeks to determine whether this policy will be *robust* to changes in the relationship between risk-sensitivity and driving classes before it deploys the policy.

We thus formulate an *intervention robustness* problem which considers the worst-case drop in performance of a classifier in response to changes to a *subset* of the causal mechanisms of the Bayesian network. This is inspired by the principle of *independent causal mechanisms* (ICM) [Peters *et al.*, 2017], which states that causal mechanisms do not inform or influence each other; that is, even as some mechanisms are changed, other mechanisms tend to remain invariant. In the insurance example, this is reflected in that we would not necessarily expect the mechanism for ‘risk aversion’ or ‘accident’ to change, for instance.

While many related notions of robustness exist in the literature, none accurately captures this notion of robustness to causal mechanism changes. Many popular definitions of robustness measure the size of a perturbation necessary to change an input’s classification, without taking into account that such perturbations may change the value which the classifier tries to predict [Shih *et al.*, 2018]. Miller *et al.* [2020] highlight the connection between causal inference and robustness to distribution shifts caused by ‘gaming’ in the *strategic classification* [Hardt *et al.*, 2016] regime. However, Miller *et al.* [2020] does not assume access to a known causal model, and its focus is on identifying classifiers which are robust to gaming, whereas our objective is to verify robustness to a much richer collection of distribution shifts.

3.1 Intervention classes

To reason about the effects of changes to a causal model, we need a formal description of these interventions. We consider interventions as *actions* that modify the mechanisms of a causal Bayesian network $\mathcal{N} = (\mathcal{G}, \Theta)$, thereby changing its joint distribution. In particular, we consider two types of interventions: the first concerns changes to the parameters of the causal model, while the second concerns changes to the existence of cause-effect relationships themselves.

Typically, we might expect that only mechanisms for a subset of variables $\mathbf{W} \subseteq \mathbf{V}$ will change. In what follows, given a subset of variables $\mathbf{W} \subseteq \mathbf{V}$, we will use $\theta_{\mathbf{W}}^{(\mathcal{G})} \subseteq \Theta$ to denote the parameters associated with the CPTs for variables $W \in \mathbf{W}$, where the parents of W are given by graph \mathcal{G} .

Definition 6 (Parametric Interventions). A parametric intervention on variables \mathbf{W} substitutes a subset of parameters $\theta_{\mathbf{W}}^{(\mathcal{G})}$ for new values $\theta_{\mathbf{W}}^{(\mathcal{G})'}$ obtaining a new parameter set Θ' , which yields the BN:

$$\mathcal{N}[\theta_{\mathbf{W}}^{(\mathcal{G})'}] := (\mathcal{G}, \Theta') \quad (2)$$

Parametric interventions encompass the do-interventions discussed by Qin [2015], but allow us to express more complex changes to causal mechanisms than fixing a variable to a set value. We can further consider changes not just to the parameters of the network, but also to its edge structure; such changes to a set of variables \mathbf{W} can be described by a *context function* $C_{\mathbf{W}} : \mathbf{W} \rightarrow \mathcal{P}(\mathbf{V})$, which replaces the parents of $W \in \mathbf{W}$ in \mathcal{G} with $C_{\mathbf{W}}(W)$, producing a new graph \mathcal{G}' . We refer to such interventions as *structural* interventions. In this work we restrict ourselves to context sets which preserve the acyclicity of the DAG.

Definition 7 (Structural interventions). A structural intervention on variables \mathbf{W} modifies the edges \mathbf{E} of the graph $\mathcal{G} = (\mathbf{V}, \mathbf{E})$ according to a context function $C_{\mathbf{W}}$, obtaining a new graph $\mathcal{G}' = (\mathbf{V}, \mathbf{E}')$, and substitutes parameters $\theta_{\mathbf{W}}^{(\mathcal{G})}$ for new values $\theta_{\mathbf{W}}^{(\mathcal{G})'}$, obtaining a new parameter set Θ' , which yields the BN:

$$\mathcal{N}[\theta_{\mathbf{W}}^{(\mathcal{G})'}, C_{\mathbf{W}}] := (\mathcal{G}', \Theta') \quad (3)$$

We will often be interested in considering all of the possible interventions of a given class on some subset $\mathbf{W} \subseteq \mathbf{V}$ of the variables in the causal graph. Letting $\mathbb{P}_{\mathcal{G}}(\mathbf{W})$ denote the set of valid parameter sets for $\mathbf{W} \subseteq \mathbf{V}$ in graph \mathcal{G} , we will write for parametric interventions:

$$\mathcal{I}_{\mathcal{N}}[\mathbf{W}] := \{\mathcal{N}[\theta_{\mathbf{W}}^{(\mathcal{G})'}] \mid \theta_{\mathbf{W}}^{(\mathcal{G})'} \in \mathbb{P}_{\mathcal{G}}(\mathbf{W})\} \quad (4)$$

and for structural interventions,

$$\mathcal{I}_{\mathcal{N}}[\mathbf{W}, C_{\mathbf{W}}] := \{\mathcal{N}[\theta_{\mathbf{W}}^{(\mathcal{G})'}, C_{\mathbf{W}}] \mid \theta_{\mathbf{W}}^{(\mathcal{G})'} \in \mathbb{P}_{\mathcal{G}'}(\mathbf{W})\}. \quad (5)$$

3.2 Problem Definition and Complexity

Accurately assessing the robustness of a decision function F requires an understanding of the causal data-generating process. We propose to model this DGP using a causal Bayesian network \mathcal{N} on all variables \mathbf{V} , thus enabling causal *model-based* analysis of classifiers. In order to reason about the causal structure \mathcal{N} and decision rule F simultaneously, we add an additional node to the CBN \mathcal{N} .

Definition 8 (Augmented BN). For a CBN \mathcal{N} over variables V and a classifier $F : \mathbf{X} \rightarrow Y$, we define the augmented BN \mathcal{N}_F based on \mathcal{N} as follows: $\mathbf{V}_F = \mathbf{V} \cup \{\hat{Y}\}$ with $\text{pa}(\hat{Y}) = \mathbf{X}$ and deterministic CPT $\theta_{\hat{y}|\mathbf{x}} = \mathbb{1}[\hat{y} = F(\mathbf{x})]$.

This produces a well-defined joint distribution over the variables \mathbf{V} and $\hat{Y} = F(\mathbf{X})$, which allows us to specify performance metrics as probabilities of events e . For instance, a classifier’s false positive rate can be expressed as the probability $p_{\mathcal{N}_F}(e)$ of event $e = (\hat{Y} = 1) \wedge (Y = 0)$. More importantly, we can consider how these metrics change as the joint distribution changes, due to hypothetical or observed interventions on the causal model. This provides a basis for model-based notions of robustness of decision rules.

We use intervention sets to represent all interventions which the modeller considers plausible. The *interventional robustness* problem then concerns the worst-case performance of the decision rule over interventions in that set.

Definition 9. Given CBN \mathcal{N} and decision rule F , let $\mathcal{I}_{\mathcal{N}_F}$ be an intervention set for the augmented BN \mathcal{N}_F , e be an assignment of a subset of the variables in \mathbf{V} , and $\epsilon > 0$. The *interventional robustness problem* is that of computing:

$$\text{IntRob}(\mathcal{I}_{\mathcal{N}_F}, e) := \max_{\mathcal{N}' \in \mathcal{I}_{\mathcal{N}_F}} p_{\mathcal{N}'}(e).$$

We also have the corresponding decision problem:

$$\text{IntRob}(\mathcal{I}_{\mathcal{N}_F}, e, \epsilon) := \max_{\mathcal{N}' \in \mathcal{I}_{\mathcal{N}_F}} p_{\mathcal{N}'}(e) > \epsilon.$$

We will be particularly interested in problem instances where $\mathcal{I}_{\mathcal{N}}$ is of the form $\mathcal{I}_{\mathcal{N}}[\mathbf{W}]$, in which case we can view the problem instance as $\text{IntRob}((\mathcal{N}, \mathbf{W}), e)$. We observe that, via a reduction which we defer to the Appendix, the causal semantics of IntRob do not increase the computational hardness of the problem beyond that of MAP inference.

Theorem 1. Let $\mathcal{N} = (\mathcal{G}, \Theta)$ be a causal Bayesian network, with n nodes and maximal in-degree d . Then an instance of MAP can be reduced to an instance of IntRob on a BN \mathcal{N}' of size linear in $|\mathcal{N}|$, and of treewidth $w' \leq w + 2$. An instance of IntRob can be reduced to an instance of MAP on a BN \mathcal{N}' whose CPT Θ' has size polynomial in the size of Θ , and with treewidth $w' \leq 2w$.

4 Verification of Intervention Robustness

In this section, we present our approach to verifying interventional robustness. Due to the difficulty of the problem, we seek to approximate $\text{IntRob}(\mathcal{I}, e)$ by providing guaranteed upper and lower bounds that can be efficiently computed¹.

4.1 Joint Compilation

Our first goal is to *compile* \mathcal{N}_F into an equivalent arithmetic circuit \mathcal{AC} . To do so, we make use of a standard CNF encoding $\Delta_{\mathcal{N}}$ of the causal BN \mathcal{N} , defined over the indicator and parameter variables $\lambda_{\mathbf{V}}, \Theta$ [Chavira and Darwiche, 2005], and additionally an encoding of the decision function F .

A naïve encoding of F is to explicitly enumerate all instantiations of features \mathbf{x} and prediction \hat{y} , and encode these

¹In Appendix F, we show with an example that heuristic approximation using existing methods cannot provide guarantees.

directly as CNF clauses. However, this approach is very inefficient for larger feature sets \mathbf{X} . We instead assume access to an encoding of the classifier as a Boolean circuit Σ over input features \mathbf{X} and prediction \hat{Y} ². Such a circuit can be converted to CNF through the Tseitin transformation, introducing additional intermediate variables \mathbf{T} , obtaining a CNF formula Δ_F over $\lambda_{\mathbf{X}}, \lambda_{\hat{Y}}, \mathbf{T}$. We then combine the encodings of F and \mathcal{N} simply by conjoining the CNF formulae, to produce a new formula $\Delta_{\text{joint}} = \Delta_{\mathcal{N}} \wedge \Delta_F$, over $\lambda_{\mathbf{V}}, \lambda_{\hat{Y}}, \Theta, \mathbf{T}$.

To construct an AC \mathcal{AC} , we now *compile* this CNF encoding into d-DNNF (deterministic decomposable negation normal form), using the C2D compiler [Darwiche, 2004], and then replace \vee -nodes with $+$, \wedge -nodes with \times , and set all negative literals and literals corresponding to \mathbf{T} to 1. This produces an AC with polynomial $l_{\mathcal{AC}}[\lambda, \Theta]$, where $\lambda := \lambda_{\mathbf{V}} \cup \lambda_{\hat{Y}}$. Crucially, this AC is equivalent to the augmented BN, in the following sense:

Proposition 1. $l_{\mathcal{AC}}[\lambda, \Theta]$ is equivalent to $l_{\mathcal{N}_F}[\lambda, \Theta]$. Further, \mathcal{AC} can be used to faithfully evaluate marginal probabilities $p_{\mathcal{N}'}(e)$ under any parametric intervention \mathcal{N}' .

The time and space complexity of this procedure is $O(nw2^w)$, where n is the number of CNF variables and w the treewidth, a measure of the connectivity of the CNF. When jointly compiling a BN and a decision function, we can bound n, w in terms of the individual encodings $\Delta_{\mathcal{N}}, \Delta_F$.

Proposition 2. Suppose $\Delta_{\mathcal{N}}$ has n variables and treewidth w , and Δ_F has n' variables and treewidth w' . Then Δ_{joint} has exactly $n + n' - |\lambda_{\mathbf{X}}|$ variables, and treewidth at most $\max(w, w', \min(w, w') + |\lambda_{\mathbf{X}}|)$.

4.2 Orderings

For the correctness of our upper bounding algorithm, it is necessary to impose some structural constraints on the circuit.

Firstly, any circuit compiled using the procedure described above has the property that every $+$ -node t has two children, and is associated with some CNF variable c , such that one child branch has c true, and the other has c false (information on the identity of this variable for each $+$ -node is provided by the C2D compiler). We need to ensure that the arithmetic circuit only contains $+$ -nodes associated with indicators λ , and not intermediate variables \mathbf{T} . Provided this is the case, the branches of each $+$ -node t will contain contradicting indicators for some unique variable V . We can thus say that t ‘splits’ variable V , as each of its child branches corresponds to different values of V , and we write $\text{split}(t)$ to denote this splitting variable.

Secondly, provided the above holds, we require our circuit to satisfy some constraints of the following form.

Definition 10. An arithmetic circuit \mathcal{AC} satisfies the ordering constraint (V_j, V_i) if:

$$\begin{aligned} \forall t, t', (\text{split}(t) = V_i \wedge \text{split}(t') = V_j) \\ \implies t' \text{ is not a descendant of } t \text{ in } \mathcal{AC} \quad (6) \end{aligned}$$

Intuitively, our algorithm requires that for BN variables in the intervention set \mathbf{W} , the relative position of splitting

²If input features are discrete, they can still be encoded using additional binary variables; see Appendix B.

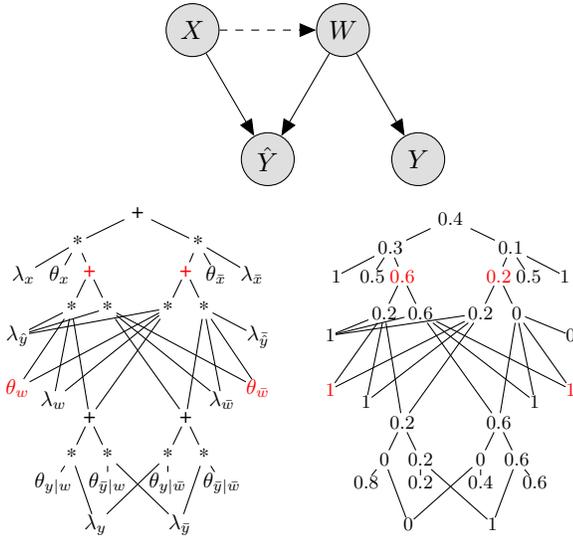


Figure 2: Example augmented BN (top), corresponding AC (bottom left), and execution of Algorithm 1 (bottom right). Nodes which differ from standard AC evaluation are highlighted in red.

$+$ -nodes in the AC agrees with the causal ordering in the BN. More formally, we say that \mathcal{AC} satisfies the ordering constraints associated with intervention set $\mathcal{I}_{\mathcal{N}_F}$, if for all $V_i \in \mathbf{W}$, and all V_j such that $V_j \in \text{pa}_{\mathcal{G}}(V_i)$ (parametric intervention set) or $V_j \in \text{pa}_{\mathcal{G}'}(V_i)$ (structural intervention set), \mathcal{AC} satisfies the ordering constraint (V_j, V_i) . In practice, when computationally feasible, we compile ACs with *topological* and *structural topological* orderings, which satisfy these constraints for all V_i , not just $V_i \in \mathbf{W}$; such orderings have the advantage of being valid for any intervention sets \mathbf{W} .

We enforce these constraints by enforcing corresponding constraints on the *elimination ordering* π over the CNF variables, which is used to construct the *dtree* that is used in the compilation process, and affects the time and space taken by the compilation. Such an elimination ordering is usually chosen using a heuristic such as min-fill. We instead find an elimination ordering by using a *constrained* min-fill heuristic, which ensures that these constraints are satisfied, but may produce an AC which is much larger than can be achieved with an unconstrained heuristic in practice. We provide further details of this in Appendix B.1.

4.3 Upper bounds on Intervention Robustness

In order to compute upper bounds on the interventional robustness quantity, we propose Algorithm 1, which sets parameters in the AC for the CPTs of variables in \mathbf{W} to 1, and applies maximization instead of addition at $+$ -nodes splitting on variables in \mathbf{W} , when evaluating the (appropriately ordered) AC. Algorithm 1 somewhat resembles the well-known MPE algorithm on ACs, introduced by Chan and Darwiche [2006] and used as an upper bound on the MAP problem in Huang *et al.* [2006]. However, our algorithm maximizes over parameters rather than variables and makes use of specific AC structure ensured by our ordering constraints; the reason it produces correct upper bounds is thus also different.

Intuitively, the maximizations represent decision points, where choosing a child branch corresponds to intervening to

Algorithm 1: $UB(\mathcal{AC}, e, \mathbf{W})$ (Upper Bounding)

Input: \mathcal{AC} , the AC; evidence e ; intervenable variables $\mathbf{W} \subseteq \mathbf{V}$;

Result: Output probability p

```

1 for node  $c \in \mathcal{AC}$  (children before parents) do
2   switch type( $c$ ) do
3     case Indicator  $\lambda_v$  do
4       |  $p[c] := 0$  if  $v'$  not consistent with  $e$  else 1
5     case Parameter  $\theta_{v|u}$  do
6       |  $p[c] := 1$  if  $V \in \mathbf{W}$  else  $\theta_{v|u}$ 
7     case  $\times$  do
8       |  $p[c] := \prod_d p[d]$ 
9         where  $d$  are the children of  $c$ 
10    case  $+$  do
11      | if  $c$  splits on some  $W \in \mathbf{W}$  then
12        |  $p[c] := \max_d p[d]$ 
13          where  $d$  are the children of  $c$ 
14      | else
15        |  $p[c] := \sum_d p[d]$ 
16          where  $d$  are the children of  $c$ 
17  Return  $p[c_{root}]$ , where  $c_{root}$  is the root node of  $\mathcal{AC}$ 

```

set a particular parameter $\theta_{w|uw}$ to 1 (and others to 0). For example, consider the augmented Bayesian network in Figure 2, where all variables are binary, $\hat{Y} = X \vee W$, and the context for W is X (represented by the dashed line, which is not part of the original BN). Figure 2 also shows execution of Algorithm 1 for false positive probability (that is, evidence $e = \{Y = 0, \hat{Y} = 1\}$). At the two $+$ -nodes where maximizations occur, the value of X is already “decided”, and the adversary can effectively choose to set $\theta_{w|\bar{x}} = 1$ and $\theta_{w|x} = 1$. In this case, the result 0.4 turns out to be exactly equal to the interventional robustness quantity $\text{IntRob}(\mathcal{I}_{\mathcal{N}_F}, e)$.

We might ask whether this intuition is correct in general. Our next result shows that, while the algorithm cannot always compute $\text{IntRob}(\mathcal{I}_{\mathcal{N}_F}, e)$ exactly, it does produce guaranteed upper bounds (proof in Appendix D):

Theorem 2. *Given a parametric/structural intervention set $\mathcal{I}_{\mathcal{N}_F}$, let \mathcal{AC} be an arithmetic circuit with the same polynomial as \mathcal{N}_F , and satisfying the ordering constraints associated with the intervention set. Then, applying the UB algorithm $UB(\mathcal{AC}, e, \mathbf{W})$ returns a quantity B_U which is an upper bound on the interventional robustness quantity $\text{IntRob}(\mathcal{I}_{\mathcal{N}_F}, e)$.*

This result is quite surprising; it shows that it is possible, through a very simple and inexpensive procedure requiring just a single linear time pass through the AC, to upper bound the worst-case marginal probability over an exponentially sized set of interventions. That this also holds for structural intervention sets, which alter the structure of the Bayesian network which the AC was compiled from, is even more surprising. Further, a compiled AC can be used for *any* intervention set given that it satisfies the appropriate ordering constraints. For instance, an AC compiled using a topological ordering allows us to derive upper bounds for parametric intervention sets involving any subset (of any size) $\mathbf{W} \subseteq \mathbf{V}$, simply by setting the appropriate parameter nodes in the AC to 1 (Line 5). This allows us to amortize the cost of evaluating

robustness against multiple intervention sets.

4.4 Lower bounds via best-response dynamics

In addition to an upper bound on $\text{IntRob}(\mathcal{I}_{N_F}, e)$, we can also straightforwardly lower bound this quantity using any witness. In this section, we will assume the setting of parametric interventions of the form $\mathcal{I}_N[\mathbf{W}]$, or of structural interventions where the search is over a single context function.

Algorithm 2: Lower Bounding

Input: $\mathcal{N} = (\mathcal{G}, \Theta)$, a Bayesian network; evidence e whose probability will be maximized; intervenable variables $\mathbf{W} \subseteq \mathbf{V}$.

Result: Output probability

$$p(e) \leq \max_{\mathcal{N}' \in \mathcal{I}_N[\mathbf{W}]} P_{\mathcal{N}'}(e)$$

```

1 begin
2    $v \leftarrow 0$ ;
3   while  $p_{\mathcal{N}[\Theta_{\mathbf{W}}]}(e) > v$  do
4      $v \leftarrow p_{\mathcal{N}[\Theta_{\mathbf{W}}]}(e)$ ;
5     for CPT  $\theta_{W|u}^{(\mathcal{G})} \in \Theta_{\mathbf{W}}$  do
6        $\theta_{W|u}^{(\mathcal{G})} \leftarrow \arg \max_{\theta_{W|u}'} p_{\mathcal{N}[\theta_{W|u}']}(e)$ ;

```

We obtain such an approach by formalizing the problem of finding an intervention which maximizes $p_{\mathcal{N}'}(e)$ as a multiplayer game, where each instantiation u_W of parents $\text{pa}_{\mathcal{G}'}(W)$ for each $W \in \mathbf{W}$ specifies a player, and where all players share a utility function given by $p_{\mathcal{N}[\Theta]}(e)$. Each player’s strategy set consists of the set of deterministic conditional distributions $\theta_{W|u}$ (we note w.l.o.g. that, by the multilinearity of the network polynomial, the optimal value of $P_{\mathcal{N}[\Theta]}(e)$ is obtained by at least one deterministic interventional distribution). A Nash equilibrium in this game then corresponds to an interventional distribution for which no change in a single parameter can increase $p_{\mathcal{N}'}(e)$. Algorithm 2 follows best-response dynamics in this game. We provide an analysis of the time complexity and convergence of this approach in the proof of the following proposition.

Proposition 3. *Algorithm 2 converges to a locally optimal parametric intervention in finite time. Further, if the algorithm is stopped before termination, the current value v will be a lower bound on $\max_{\mathcal{N}' \in \mathcal{I}[\mathbf{W}]} P_{\mathcal{N}'}(e)$.*

5 Case Study: Insurance

In this case study, we look at an extended version of the car insurance example, using the Bayesian network model shown in Figure 3 [Binder *et al.*, 1997].

Suppose an insurance company wishes to use a classifier to predict MedCost (the medical cost of an insurance claim), and has access to an insurant’s Age , DrivHist , and MakeModel (categorical variables with 3-5 values). MedCost is either *BelowThousand* (0) or *AboveThousand* (1). They fit a Naïve Bayes classifier to historical data, obtaining a decision function F . This is then used as part of their decision-making policy determining what premiums to offer to customers.

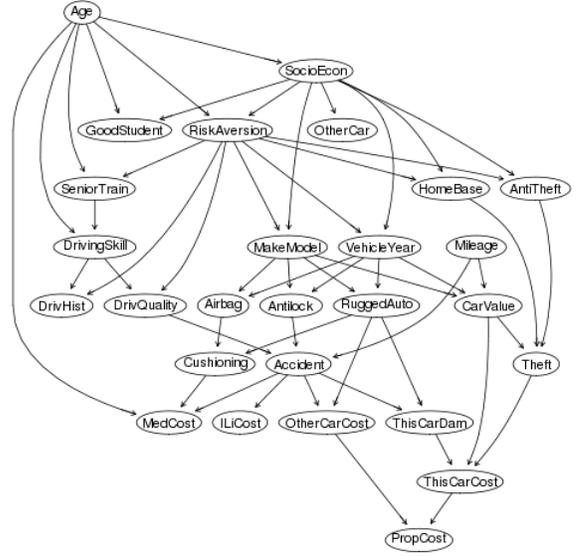


Figure 3: INSURANCE Bayesian network

The company is particularly concerned about false negatives, as this could result in the company losing a lot of money in payouts. Based on the original Bayesian network model (Figure 3) and their new classifier, this should occur 2.5% of the time. However, in reality, insurants may attempt to game the classifier to predict *BelowThousand* (so that they get lower premiums), while actually being likely to have a high medical cost. In our framework, we model this using structural interventions, assuming that insurants can causally intervene on some of DrivHist (perhaps hide some accident history), MakeModel (choose a different type of car than they would normally choose), and Cushioning (upgrade/downgrade the degree of protection inside the car). The company would thus like to understand how *robust* their classifier is to these adaptations.

We will consider a number of structural intervention sets \mathcal{I}_{N_F} , given by intervenable variables \mathbf{W} , which may be any subset of $\{\text{DrivHist}, \text{MakeModel}, \text{Cushioning}\}$. Under each of these intervention sets, we seek to obtain guaranteed upper bounds on these two quantities:

- **FN:** The probability of a false negative $p(F = 0, \text{MedCost} = 1)$, i.e. predicted low medical cost, but high actual medical cost.
- **P:** The probability of a positive $p(\text{MedCost} = 1)$, i.e. high actual medical cost.

The results are shown in Table 1. The insurance company can use these bounds to assess risk, and improve their classifier’s robustness if they deem the false negative rate under intervention unacceptable.

The bounds can also provide further insight. We notice that whenever DrivHist is intervenable, the percentage of false negatives is the same as positives, i.e. the classifier always predicts wrong when MedCost is 1. This turns out to be because the Naïve Bayes classifier always predicts 0 whenever DrivHist is *None*, regardless of the other input variables. Thus, an insurant who can change their DrivHist can al-

Intervenable Variables W	FN	P
Empty Set	2.5%	7.2%
{DrivHist}	7.2%	7.2%
{MakeModel}	5.7%	10.0%
{Cushioning}	6.1%	12.9%
{DrivHist, MakeModel}	10.0%	10.0%
{DrivHist, Cushioning}	12.9%	12.9%
{MakeModel, Cushioning}	13.0%	13.9%
{DrivHist, MakeModel, Cushioning}	13.9%	13.9%

Table 1: Guaranteed upper bounds on FN and P, under different structural intervention sets

ways fool the classifier to predict 0. In addition, the percentage of positives doesn’t increase from the original BN: this can be seen from the causal graph, where `DrivHist` has no causal influence on `MedCost`.

On the other hand, `Cushioning` significantly increases the positive rate. Notice that, in the graph, intervening on `Cushioning` will not have any influence on the inputs to the classifier; thus, the increase in FN to 6.1% is not due to fooling the classifier, but rather making high medical expenses generally more likely, by downgrading the quality of cushioning. In this way, the intervention is “taking advantage” of the classifier not having full information about cushioning.

6 Evaluations

Compilation Performance

In Table 2 we show the performance of our joint compilation approach on a number of benchmark Bayesian networks, where we jointly compile the network and a decision rule (see Appendices G and H for further details). We observe that the sizes of the compiled ACs are significantly smaller than the worst-case bounds would suggest (exponential in treewidth). Further, when we enforce a topological ordering, the size of the compilation increases, but not by more than ~ 100 . Our results provide evidence that our methods can scale to fairly large networks and classifiers, including networks compiled with topological and structural topological orderings.

Lower and Upper Bound tightness

In Table 3 we analyse the quality of our upper and lower bounds on interventional robustness. We compute bounds on false negative probability under different intervention sets. Overall, we find small or nonexistent gaps between the lower and upper bounds across all networks and intervention sets evaluated, suggesting that in many settings of interest it is possible to obtain tight guarantees using our algorithms.

Further, both bounding algorithms are very fast to execute, taking no more than a few seconds for each run. This is remarkable given the sizes of the intervention sets. For instance, for the insurance network, the parametric intervention set P2 covers 6 variables ($|W| = 6$), 248 parameters, and $\sim 10^{36}$ different interventions, making brute-force search clearly infeasible. For worst-case (interventional robustness) analysis, the sensitivity analysis method of Chan and Darwiche [2004] requires $\sim 10^7$ passes through the AC in this case. On the other hand, our upper bounding algorithm requires an ordered AC (which is ~ 5 times larger in this

Net	CSize	Ord	TW	AC size	Time (s)
insurance	3 (41)	N	24	167121	0.5
	3 (41)	T	31	794267	4
	3 (41)	S	33	1270075	8
win95pts	16 (799)	N	51	1210072	3
	16 (799)	T	58	52266950	77
hepar2	12 (946)	N	53	8096874	49
	12 (946)	T	51	123108407	73
	12 (946)	S	51	123164181	75

Table 2: AC sizes and times (s) for the joint compilations used in the UB and LB algorithms. Shown are the number of input features d and the sizes of the Boolean circuits representing the classifier, ordering constraints (none, topological, or structural topological), treewidth of the combined CNF encoding, and AC size and compilation time. Further evaluations provided in Appendix G.1.

Network	IntSet	LBound	UBound	Δ
insurance	P1	0.1181	0.1276	0.0095
	P2	0.3275	0.3433	0.0158
	S1	0.1181	0.1297	0.0116
win95pts	P1	0.2111	0.2111	0.0000
	P2	0.2163	0.2191	0.0028
hepar2	P1	0.09445	0.09445	0.0000
	P2	0.09585	0.09585	0.0000
	S1	0.1029	0.1029	0.0000

Table 3: Analysis of the tightness of bounds (on probability of false negatives) produced by Algorithms 1 and 2. For each network, we have different intervention sets (P/S indicates the intervention set is parametric/structural respectively). Lower and upper bounds, along with the difference, are shown for each intervention set. Further evaluations provided in Appendix G.2.

case), but requires just a *single* pass through the AC, making it $\sim 10^6$ faster. Further, our algorithm is uniquely able to provide guarantees for structural intervention sets.

7 Conclusions

In this work, we have motivated and formalized the interventional robustness problem, developed a compilation technique to produce efficient joint representations for classifiers and DGPs, and provided tractable upper and lower bounding algorithms which we have shown empirically to be tight on a range of networks and intervention sets. The techniques presented here provide ample opportunity for further work, such as extending the upper and lower bounding technique to networks where the modeller has uncertainty over the parameters, and developing learning algorithms for arithmetic circuits which permit reasoning about causal structure.

Acknowledgements This project was funded by the ERC under the European Union’s Horizon 2020 research and innovation programme (FUN2MODEL, grant agreement No. 834115).

References

- Gilles Audemard, Frédéric Koriche, and Pierre Marquis. On Tractable XAI Queries based on Compiled Representations. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning*, pages 838–849, September 2020.
- John Binder, Daphne Koller, Stuart Russell, and Keiji Kanazawa. Adaptive probabilistic networks with hidden variables. *Machine Learning*, 29:213–244, 1997.
- Hei Chan and Adnan Darwiche. Sensitivity analysis in bayesian networks: From single to multiple parameters. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, page 67–75, July 2004.
- Hei Chan and Adnan Darwiche. On the robustness of most probable explanations. In *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence*, page 63–71, July 2006.
- Mark Chavira and Adnan Darwiche. Compiling bayesian networks with local structure. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, page 1306–1312, August 2005.
- Arthur Choi and Adnan Darwiche. On relaxing determinism in arithmetic circuits. In *Proceedings of the 34th International Conference on Machine Learning*, page 825–833, August 2017.
- Adnan Darwiche and Pierre Marquis. A knowledge compilation map. *Journal of Artificial Intelligence Research*, 17:229–264, 2002.
- Adnan Darwiche. A logical approach to factoring belief networks. In *Proceedings of the 8th International Conference on Principles of Knowledge Representation and Reasoning*, page 409–420, April 2002.
- Adnan Darwiche. A differential approach to inference in bayesian networks. *J. ACM*, 50(3):280–305, 2003.
- Adnan Darwiche. New advances in compiling cnf to decomposable negation normal form. In *Proceedings of the 16th European Conference on Artificial Intelligence*, page 318–322, August 2004.
- Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic classification. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, page 111–122, January 2016.
- Jinbo Huang, Mark Chavira, and Adnan Darwiche. Solving map exactly by searching on compiled arithmetic circuits. In *Proceedings of the 21st National Conference on Artificial Intelligence - Volume 2*, page 1143–1148, July 2006.
- Guy Katz, Clark W. Barrett, David L. Dill, Kyle Julian, and Mykel J. Kochenderfer. Reluplex: An efficient SMT solver for verifying deep neural networks. In *Proceedings of the 29th International Conference on Computer Aided Verification - Part I*, pages 97–117, 2017.
- Zachary C. Lipton, Yu-Xiang Wang, and Alexander J. Smola. Detecting and correcting for label shift with black box predictors. In *Proceedings of the 35th International Conference on Machine Learning*, pages 3128–3136, July 2018.
- Daniel Lowd and Pedro Domingos. Learning arithmetic circuits. In *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence*, page 383–392, July 2008.
- John Miller, Smitha Milli, and Moritz Hardt. Strategic classification is causal modeling in disguise. In *Proceedings of 37th International Conference on Machine Learning*, pages 6917–6926, July 2020.
- Nina Narodytska, Shiva Kasiviswanathan, Leonid Ryzhyk, Mooly Sagiv, and Toby Walsh. Verifying properties of binarized deep neural networks. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, pages 6615–6624, February 2018.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press, 2017.
- Hoifung Poon and Pedro Domingos. Sum-product networks: A new deep architecture. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*, page 337–346, July 2011.
- Biao Qin. Differential semantics of intervention in bayesian networks. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, pages 710–716, July 2015.
- Joaquin Quionero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D. Lawrence. *Dataset Shift in Machine Learning*. The MIT Press, 2009.
- Mateo Rojas-Carulla, Bernhard Schölkopf, Richard Turner, and Jonas Peters. Invariant models for causal transfer learning. *Journal of Machine Learning Research*, 19(36):1–34, 2018.
- Tim Roughgarden. Algorithmic game theory. *Communications of the ACM*, 53(7):78–86, 2010.
- Weijia Shi, Andy Shih, Adnan Darwiche, and Arthur Choi. On tractable representations of binary neural networks. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning*, pages 882–892, September 2020.
- Andy Shih, Arthur Choi, and Adnan Darwiche. Formal verification of bayesian network classifiers. In *Proceedings of the Ninth International Conference on Probabilistic Graphical Models*, pages 427–438. PMLR, Sep 2018.
- Andy Shih, Arthur Choi, and Adnan Darwiche. Compiling bayesian network classifiers into decision graphs. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, pages 7966–7974, January 2019.
- Adarsh Subbaswamy, Peter Schulam, and Suchi Saria. Preventing failures due to dataset shift: Learning predictive models that transport. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research, pages 3118–3127, 2019.
- Kun Zhang, Mingming Gong, and Bernhard Schölkopf. Multi-source domain adaptation: A causal view. In *Pro-*

A Proof of Theorem 1

We now prove Theorem 1. To do so, we will formalize the intervention maximization problem, whose decision form is IntRob, as follows.

IntRob is the Intervention Robustness problem. Note that while the definition given in Section 3.2 is agnostic to the form of the interventions, in order to concretely bound the size of the input to the problem, we restrict ourselves to Bayesian networks defined on discrete random variables and intervention sets of the form $\mathcal{I}_{\mathcal{N}}(\mathbf{W})$. Thus, the problem takes as input a Bayesian network \mathcal{N} , a set of intervenable nodes \mathbf{W} , and some evidence e . The goal is to find a parametric intervention on the variables in \mathbf{W} Θ' such that $P_{\mathcal{N}[\Theta']}(e)$ is maximized. The decision version of this problem is intervention robustness.

MAP: is Maximum a Posteriori inference, a well-studied problem which takes as input network \mathcal{N} , variables \mathbf{W} , and evidence e , and whose objective is to find an instantiation $\mathbf{W} = w$ such that $P_{\mathcal{N}}(w, e)$ is maximal.

Theorem 1. *Let $\mathcal{N} = (\mathcal{G}, \Theta)$ be a causal Bayesian network, with n nodes and maximal in-degree d . Then an instance of MAP can be reduced to an instance of IntRob on a BN \mathcal{N}' of size linear in $|\mathcal{N}|$, and of treewidth $w' \leq w+2$. An instance of IntRob can be reduced to an instance of MAP on a BN \mathcal{N}' whose CPT Θ' has size polynomial in the size of Θ , and with treewidth $w' \leq 2w$.*

Without loss of generality, in the proof of this statement we will assume all variables are binary-valued; it is straightforward to then extend the results here to arbitrary discrete random variables supported on a finite set.

Lemma 1. *MAP is reducible to IntRob.*

Proof. Let $\mathcal{N}, \mathbf{W}, e$ be an instantiation of the MAP problem. We can convert this into the IntRob problem by adding the following sets of nodes to \mathcal{N} in order to produce a new network \mathcal{N}' .

1. For each $V \in \mathbf{W}$, add a node V_{θ} with the same support as V , and which has no incoming arrows.
2. Additionally, for each $V \in \mathbf{W}$ add a node A_V with parents V and V_{θ} with support True/False, which is True with probability 1 if $V = V_{\theta}$ and False otherwise. Let $\mathbf{A}_{\mathbf{W}}$ denote the set of all such nodes.

We now show that in this new network \mathcal{N}' , $\text{IntRob}(\mathcal{N}', \mathbf{W}_{\theta}, \{\mathbf{A}_{\mathbf{V}} = \text{True}, e\})$ is equal to $\text{MAP}(\mathcal{N}, \mathbf{W}, e)$.

We first observe that for a single $V \in \mathbf{W}$, we have

$$\begin{aligned} P_{\mathcal{N}}(V = v) &= P_{\mathcal{N}[\theta_{V_{\theta}} = \mathbb{1}\{V_{\theta}=v\}]}(V = V_{\theta}) \\ &= P_{\mathcal{N}[\theta_{V_{\theta}} = \mathbb{1}\{V_{\theta}=v\}]}(A_V = \text{True}). \end{aligned}$$

Because V_{θ} and A_V are independent of the rest of the graph given V , we then straightforwardly obtain that for additional evidence e , the same equality holds for the joint evidence $(V = v, e)$:

$$P_{\mathcal{N}[\theta_{V_{\theta}} = \mathbb{1}\{V_{\theta}=v\}]}(A_V = \text{True}, e) = P_{\mathcal{N}}(V = v, e).$$

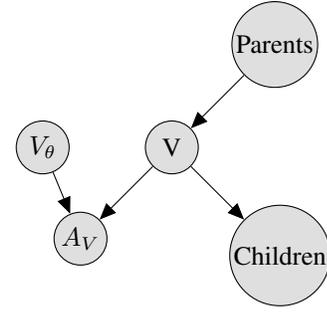


Figure 4: Visualization of the construction of \mathcal{N}' for the proof of Lemma 1.

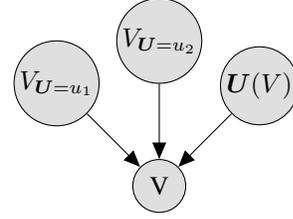


Figure 5: Visualization of the intuition behind the construction of \mathcal{N}' for the proof of Lemma 2.

Finally, this equality can be iterated to incorporate all nodes $V \in \mathcal{I}$, and so for any instantiation $\mathbf{w} = (\mathbf{w}^1, \dots, \mathbf{w}^m)$ of $\mathbf{W} = \{V^1, \dots, V^m\}$ with corresponding parameters $\Theta'_{\mathbf{w}} = \{\mathbb{1}[V_{\theta}^i = \mathbf{w}^i] | V^i \in \mathbf{W}\}$ we obtain

$$P_{\mathcal{N}[\Theta'_{\mathbf{w}}]}(\mathbf{A}_{\mathbf{W}} = \text{True}, e) = P_{\mathcal{N}}(\mathbf{W} = \mathbf{w}, e).$$

So the parametric interventions which maximize $P(\mathbf{A}_{\mathbf{W}} = T, e)$ are equivalent to the variable values \mathbf{w} which maximize $P(\mathbf{W} = \mathbf{w}, e)$.

It is straightforward to show that the size of the resulting BN \mathcal{N}' satisfies the theorem statement. We have increased the number of nodes by $2|\mathbf{W}| \leq 2|\mathbf{V}|$, and added CPTs of size $|\text{supp}(V)|^2$, which for binary variables will be fixed at 4, so $|\Theta'| \leq |\Theta| + 4|\mathbf{V}|$. Finally, we have not increased the treewidth of the network by more than a constant increment of 2 because we have added a fully connected component with 2 additional nodes to each variable $V \in \mathcal{I}$ with no other edges into the graph. \square

Lemma 2. *IntRob is reducible to MAP.*

Proof. For the opposite direction, we show that we can use MAP to solve IntRob(with parametric interventions). Let $\mathcal{N}, \mathbf{W}, e$ be inputs into IntRob. We initially construct \mathcal{N}' as a copy of \mathcal{N} . We will proceed by converting the parameters $\theta_{V|u}$ into variables V_u in \mathcal{N}' , where observing the value $V_u = v$ in \mathcal{N}' is equivalent, up to a constant factor, to setting $\theta_{V|u} := \mathbb{1}[V = v]$ in the original network \mathcal{N} , while avoiding an exponential blowup in the size of this new Bayesian network.

Intuition: For each CPT component $\theta_{V|u}$ in \mathcal{N} such that $V \in \mathbf{W}$, introduce a new “parameter-node” V_u to \mathcal{N}' with support equal to that of V and uniformly distributed (i.p. V_u has no parents), and add an edge into V . Set $P(V =$

$v'|u, V_u) = \mathbb{1}[V_u = v']$. In other words, if the value of $U = u$, then V gets the value of the node V_u deterministically. This construction is visualized in Figure 5.

Naively, this will blow up the CPT for V by a factor exponential in the size of the parent set, and may increase the treewidth of the network by this factor as well. We therefore proceed in the following construction to minimize the impact on the size of the representation of \mathcal{N}' . This construction will still add in the worst case $|\Theta|$ variables to the network, but it will only increase the size of $|\Theta|$ by a linear factor and will only increase the treewidth by at most a factor of 2.

Selector circuits: in a more efficient (and more involved) construction, we add a number of auxiliary variables that act as a filter based on the values of the parents U to pass down the correct value to V . We visualize an example of this construction in Figure 6. We enumerate the parents of V as U_1, \dots, U_n , and add auxiliary variables $\{V_u | u \in \text{supp}(U)\}$ to the graph as described in the intuition above. We then add auxiliary variables S_{ij} , with i ranging from 1 to n and j ranging from 1 to $2^{|U|-i}$ for each level i . To each variable V_u we assign it the binary string $u_n \dots u_1$. For each prefix $u \in \{0, 1\}^{n-1}$ we then draw arrows from U_1, V_{u0} and V_{u1} to $S_{1,u}$ and define the conditional $P(S_{1,u} | V_{u0}, V_{u1}, U_1 = u_1) = \mathbb{1}[S_{1,u} = V_{u,u_1}]$. We inductively define at layer i for a prefix $u \in \{0, 1\}^{i-1}$ the random variable $S_{i,u}$, with parent variables $S_{i-1,u0}, S_{i-1,u1}$, and U_i , and conditional distribution $P(S_{i,u} | S_{i-1,u0}, S_{i-1,u1}, U_i = u_i) = \mathbb{1}[S_{i,u} = V_{u,u_i}]$. The value of $S_{n,\emptyset}$ will therefore be deterministically the value of V_u for $U = u$, and so we can simply set the CPT of V to depend uniquely and deterministically on $S_{n,\emptyset}$.

This procedure will not add more nodes than there are parameters in the CPT to the graph (i.e. the increase in the number of nodes n is bounded by $\exp(w)$) and will not increase the treewidth by more than a linear factor, as the selector circuit has treewidth at most $2|U| + 1$ (this is easily observed by forming a tree decomposition via a depth-first-search procedure). The maximal number of nodes this construction can add to the table is therefore $2|\Theta|$, assuming that all n nodes are set to be intervened on, and we can increase the treewidth by a factor of at most 2, independent of the number of nodes we modify. We therefore obtain that for \mathcal{N}' : $|V| \leq |V| + 2|\Theta|$, $|\Theta| \leq 16|\Theta|$, and $w' \leq 2w$.

The end result of this construction is that we have changed the distribution of the variable V so that it now depends on $2^{|U|}$ additional random variables V_{u_1}, \dots, V_{u_k} , and deterministically takes the value of V_{u_i} whenever its parents satisfy $U = u_i$.

We now claim that in the new network \mathcal{N}' , the conditional distribution on V induced by observing $V_u = v_u$ is equal up to a constant to doing a parametric intervention on $\theta_{v|u}$ in \mathcal{N} which deterministically sets v to v_u conditioned on $U = u$. To see this, let \mathbf{v}_U be an instantiation of the variables \mathbf{V}_U .

$$\begin{aligned} P_{\mathcal{N}'}(V = v | \mathbf{u}, \mathbf{v}_U) &= P_{\mathcal{N}'}[V = v | V_u = v_u, U = \mathbf{u}] \\ &= \mathbb{1}[v = v_u] \\ &= P_{\mathcal{N}[\theta_{v|u}=1]}[V = v | U = \mathbf{u}] \end{aligned}$$

We also observe that for any ancestor U of V , $P_{\mathcal{N}'}(U) =$

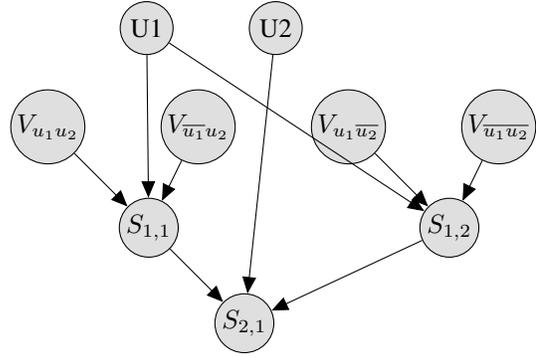


Figure 6: Example ‘selector’ circuit. At each level i , pair up solutions from paths that differ only on variables u_1, \dots, u_i (will be uniquely 2 for binary valued variables), then use the value of u_i to ‘select’ the correct value. I.e. $P(S_{1,1} = V_{u_1 u_2} | U_1) = 1$ if $U_1 = u_1$ else $P(S_{1,1} = V_{\bar{u}_1 u_2}) = 1$ if $U_1 = \bar{u}_1$. After $|U|$ levels, the value of A_n will be the value of V_u , and so will be used to substitute for V .

$P_{\mathcal{N}'}(U)$, since we only changed the conditional distribution of a descendent of U . Further, for any descendent D of V , we obtain $P_B(D | U = \mathbf{u}, V = v) = P_{\mathcal{N}'}(D | V = v, U = \mathbf{u})$. This is again because we did not change any conditional distribution for a descendent of V .

Because V_U is uniformly distributed in \mathcal{N}' , we can therefore straightforwardly derive that for a single variable, INTROB can be reduced to MAP on \mathcal{N}' . We decompose the evidence variables e into e_C and e_A (for descendents, ancestors of V respectively).

$$\max_{\theta_{v|u}} P_{\mathcal{N}[\theta_{v|u}]}(e) = \max_{\theta_{v|u}} P_{\mathcal{N}}(e_A) P_{\mathcal{N}[\theta_{v|u}]}(e_C) \quad (7)$$

$$= P_{\mathcal{N}}(e_A) \max_{v_u} \sum_u P(u | e_A) P(e_C | V = v_u, u) \quad (8)$$

$$= P_{\mathcal{N}'}(e_A) \max_{\mathbf{v}_U} 2^{|U|} P_{\mathcal{N}'}(\mathbf{V}_U = \mathbf{v}_U) \quad (9)$$

$$\sum_u P_{\mathcal{N}'}(U = u | e_A) P(e_C | e_A, u, V = v_u) \quad (10)$$

$$= 2^{|U|} \max_{\mathbf{v}_U} P_{\mathcal{N}'}(e, \mathbf{V}_U = \mathbf{v}_U) \quad (11)$$

We note that because $2^{|U|}$ is a constant, it does not affect the maximization problem and so the two maximization problems will be maximized by equivalent parameter settings. Once the result is established for single variables, this observation is easily extended to the entire set \mathbf{W} by performing an analogous summation over all instantiations of variables which are parents of an element of \mathbf{W} . The only trick is that, when summing over values u of the parents of \mathbf{W} , one must be careful in dealing with elements $v \in \mathbf{W}$ which are also parents of intervenable variables. In such cases, we define $P_{\mathcal{N}'}(\mathbf{w} | \mathbf{u})$ to be zero if the elements at the intersection of \mathbf{W} and U disagree on their assigned values.

$$\max_{\theta_{\mathbf{w}}} P_{\mathcal{N}[\theta_{\mathbf{w}}]}(e) = \max_{\theta_{\mathbf{w}}} P_{\mathcal{N}}(e_A) P_{\mathcal{N}[\theta_{\mathbf{w}}]}(e_C) \quad (12)$$

$$= P_{\mathcal{N}}(e_A) \max_{\mathbf{w}_U} \sum_u P(u | e_A) P(e_C | \mathbf{W} = \mathbf{w}_u, u) \quad (13)$$

$$= P_{\mathcal{N}'}(e_A) \max_{\mathbf{w}_U} 2^{|\mathbf{U}(\mathbf{W})|} P_{\mathcal{N}'}(\mathbf{W}_U = \mathbf{w}_U) \quad (14)$$

$$\sum_u P_{\mathcal{N}'}(\mathbf{U} = \mathbf{u} | e_A) P(e_C | e_A, \mathbf{u}, V = v_u) \quad (15)$$

$$= 2^{|\mathbf{U}|} \max_{\mathbf{v}_{\mathbf{U}}} P_{\mathcal{N}'}(e, \mathbf{V}_{\mathbf{U}} = \mathbf{v}_{\mathbf{U}}) \quad (16)$$

As noted previously, it is sufficient to consider parametric interventions that are deterministic to solve the `IntRob` problem (by the multilinearity of the network polynomial, the optimal value of $P_{\mathcal{N}'}(e)$ is obtained by at least one deterministic interventional distribution). Therefore, searching over the set of distributions induced by deterministic parametric interventions is equivalent to searching over conditional distributions induced by observing $\mathbf{V}_{\mathbf{U}}$, and so this set of conditional distributions is sufficient for solving `IntRob`. \square

B Compilation Details and Proof of Proposition 1

Our goal with compilation is to produce an arithmetic circuit with AC polynomial equivalent to the network polynomial of \mathcal{N}_F . To do this, we first generate logical encodings of the decision function F and original Bayesian network \mathcal{N} .

Any decision function F given as a Boolean circuit Σ can be expressed as a CNF encoding $\Delta_F(\lambda_{\mathbf{X}}, \lambda_{\hat{Y}}, \mathbf{T})$ on indicator variables and intermediate variables. For discrete features \mathbf{X} and discrete prediction \hat{Y} , we assume that the Boolean circuit takes as input Boolean indicator variables $\lambda_{X_i=x_i}, \lambda_{\hat{Y}=\hat{y}}$, which outputs true for an instantiation \mathbf{x}, \hat{y} iff $\hat{y} = F(\mathbf{x})$. For instance, if X_i is binary, we have separate variables $\lambda_{x_i}, \lambda_{\bar{x}_i}$ which take the place of X_i and $\neg X_i$ respectively in the circuit. With the inputs expressed as indicators, we then introduce an additional binary variable for every internal node in the Boolean circuit (forming "intermediate variables" \mathbf{T}), and use the Tseitin transformation to convert to CNF. The resulting CNF encoding $\Delta(\lambda_{\mathbf{X}}, \lambda_{\hat{Y}}, \mathbf{T})$ has the property that, for a given instantiation \mathbf{x} of the features, the CNF encoding has exactly one model (satisfying assignment) with corresponding feature indicators $\lambda_{\mathbf{X}}$, and that model has $\lambda_{\hat{Y}=F(\mathbf{x})}$ true and all other prediction indicators false.

We can encode the Bayesian network \mathcal{N} using the following types of CNF clauses, following [Darwiche, 2002]:

- *Indicator Clauses*: For variable $V \in \mathbf{V}$ with domain v_1, \dots, v_k , we include the following clauses:
 - $\lambda_{v_1} \vee \dots \vee \lambda_{v_k}$
 - $\neg \lambda_{v_i} \vee \neg \lambda_{v_j}$ for $i \neq j$
- *Parameter Clauses*: For each parameter $\theta_{v|\mathbf{u}}$, we include the following clause:
 - $\lambda_v \wedge \lambda_{u_1} \wedge \dots \wedge \lambda_{u_d} \implies \theta_{v|\mathbf{u}}$
 - $\lambda_v \wedge \lambda_{u_1} \wedge \dots \wedge \lambda_{u_d} \iff \theta_{v|\mathbf{u}}$

The models (satisfying assignments) of the resulting encoding $\Delta_{\mathcal{N}}(\lambda_{\mathbf{V}}, \Theta)$ each correspond to an instantiation v of the variables \mathbf{V} , with parameter variables set to true or false depending on whether they apply to that instantiation. Thus, they directly correspond to the terms of the network polynomial $l_{\mathcal{N}}[\lambda_{\mathbf{V}}, \Theta]$.

We now consider the joint encoding $\Delta_{joint} = \Delta_{\mathcal{N}} \wedge \Delta_F$, on variables $\lambda := \lambda_{\mathbf{V}} \cup \lambda_{\hat{Y}}, \Theta, \mathbf{T}$:

Lemma 3. *The models of the joint encoding $\Delta_{joint} = \Delta_{\mathcal{N}} \wedge \Delta_F$ correspond to the terms of the network polynomial $l_{\mathcal{N}_F}[\lambda, \Theta]$ (ignoring intermediate variables \mathbf{T}).*

Proof. Since the mechanism for \hat{Y} in \mathcal{N}_F is deterministic, we can forgo parameters for this mechanism, and write the network polynomial as:

$$l_{\mathcal{N}_F}[\lambda, \Theta] = \sum_{v_1, \dots, v_n} \lambda_{\hat{Y}=F(\mathbf{x})} \prod_{i=1}^n \lambda_{v_i} \theta_{v_i|\mathbf{u}_i}$$

If we conjoin the encodings Δ_F and $\Delta_{\mathcal{N}}$, then since each model of $\Delta_{\mathcal{N}}$ corresponds to an instantiation v (including features \mathbf{x}), the models of the joint encoding are precisely the models of $\Delta_{\mathcal{N}}$, with the prediction indicator $\lambda_{\hat{Y}=F(\mathbf{x})}$ true (and all other prediction indicators false), and the intermediate variables \mathbf{T} taking some values. Thus, each model of the joint encoding directly corresponds to some valid instantiation \mathbf{x}, \hat{y} , i.e. a term in the network polynomial. \square

Given Δ_{joint} , the next step is to compile the CNF into an equivalent sd-DNNF, a rooted DAG with literals (a variable or its negation) as leaves, and conjunctions/disjunctions as internal nodes, satisfying the following properties:

- Decomposability: From every \wedge -node, no two branches can share a variable.
- Smoothness: From every \vee -node, every branch must contain the same variables.
- Determinism: From every \vee -node, every two branches must contradict each other.

We convert the sd-DNNF to an AC by replacing conjunctions with $*$, disjunctions with $+$, and replacing all negative literals and literals corresponding to intermediate variables \mathbf{T} with the value 1.

Proposition 1. *$l_{AC}[\lambda, \Theta]$ is equivalent to $l_{\mathcal{N}_F}[\lambda, \Theta]$. Further, AC can be used to faithfully evaluate marginal probabilities $p_{\mathcal{N}'}(e)$ under any parametric intervention \mathcal{N}' .*

Proof. The logic can be expressed as a disjunction of the complete subcircuits of the d-DNNF (complete subcircuits are generated by traversing the circuit, choosing all children of every \wedge -node and one child of every \vee -node). By decomposability and smoothness, each complete subcircuit is a conjunction which specifies a value (and only one value) for each CNF variable $\lambda, \Theta, \mathbf{T}$. By Proposition 3, each complete subcircuit must thus correspond to a term of the network polynomial. Further, no other complete subcircuit can correspond to that term; otherwise determinism is violated. That is, there is a one-to-one correspondence between the complete subcircuits of the d-DNNF, and the terms of $l_{\mathcal{N}_F}[\lambda, \Theta]$.

Each complete subcircuit thus corresponds to an instantiation v , and takes the form

$$\left(\bigwedge_{i=1}^n \lambda_{v_i} \wedge \theta_{v_i|\mathbf{u}_i} \bigwedge_{v'_i \neq v_i} \neg \lambda_{v'_i} \wedge \neg \theta_{v_i|\mathbf{u}_i} \right)$$

$$\wedge \left(\lambda_{\hat{Y}=F(\mathbf{x})} \bigwedge_{\hat{y}' \neq F(\mathbf{x})} \neg \lambda_{\hat{Y}=\hat{y}'} \right) \\ \wedge \left(\bigwedge_{T \in \mathcal{T}_v} T \bigwedge_{T \in \mathcal{T} \setminus \mathcal{T}_v} \neg T \right)$$

where \mathcal{T}_v are the intermediate variables which are true for instantiation v . If we now exchange conjunctions for \times , disjunctions for $+$, and set negative literals and T -literals to 1, we get that the terms of the complete subcircuits of the AC take the form:

$$\lambda_{\hat{Y}=F(\mathbf{x})} \prod_{i=1}^n \lambda_{v_i} \theta_{V_i | \mathbf{u}_i}$$

Thus $l_{AC}[\lambda, \Theta]$ is equivalent to $l_{NF}[\lambda, \Theta]$. For the second part, for any fixed value of the parameters Θ , the network/AC polynomials will be equivalent multi-linear functions of the indicators, and thus computation of marginals on the respective polynomials will yield the same result. \square

B.1 Elimination Ordering

The C2D compiler makes use of a *dtree* in the compilation process. One method of generating a dtree is using an elimination ordering π : that is, an ordering of the CNF variables λ, Θ, T . We now specify the constraints on this ordering that need to be imposed, in order to compile an \mathcal{AC} to be used in the upper bounding algorithm (Algorithm 1). Each constraint takes the form $c < c'$, which indicates that CNF variable c must come before c' in the ordering.

Firstly, for all $\lambda \in \Lambda, T \in \mathcal{T}$, we impose constraints $T < \lambda$. This ensures that the AC contains only $+$ -nodes associated with indicator variables λ , and not intermediate variables T .

Secondly, for each constraint (V_j, V_i) corresponding to the intervention set, we impose constraints $\lambda_{v_i} < \lambda_{v_j}$ for every value v_i of V_i and every value v_j of V_j . This ensures that the AC satisfies the constraint (V_j, V_i) ; that is, $+$ -nodes splitting on V_j cannot appear after $+$ -nodes splitting on V_i .

Given these constraints, we find an elimination ordering π using the constrained min-fill heuristic. The variables in the ordering are chosen one at a time. At each point, the min-fill heuristic associates a cost with each CNF variable. We pick the lowest-cost variable which would not violate any ordering constraint.

B.2 Encoding Options

Several improvements to the encoding in [Darwiche, 2002] have been proposed, which we address in the context of our joint compilations of causal BNs:

- **Dropping parameter clauses:** [Chavira and Darwiche, 2005] observed that it is possible to drop the second type of parameter clause (reverse implication), which we employ in our experiments. This introduces additional spurious models for the encoding (since the truth values of some parameter variables will be unspecified), which were removed in [Chavira and Darwiche, 2005]

by performing a *minimization* operation on the resulting sd-DNNF. We instead take advantage of the dtree-based compilation approach in the C2D compiler, which does not introduce negative parameter literals into the d-DNNF (since they are not present in the CNF encoding), and forgo smoothing the circuit. This results in a d-DNNF which does not contain negative parameter literals, while satisfying decomposability/smoothness/determinism with respect to the Bayesian network variables.

- **Determinism and Parameter Equality:** Determinism refers to parameters $\theta_{v|u}$ taking the value 0 in the BN, and parameter equality refers to two such parameters in the same CPT taking the same value. Both of these can be encoded using simple modifications to the encoding clauses [Chavira and Darwiche, 2005], potentially resulting in much smaller compiled ACs by removing excess parameter variables. The problem with using such encodings is that there no longer exists a unique parameter variable corresponding to each parameter in the BN. As such, we cannot compute the effect of interventions on the network which affect those parameters.

However, it is in theory possible to encode determinism and parameter equality only for CPTs which we know in advance we will never want to intervene on, while avoiding other CPTs. For simplicity, we avoid encoding determinism or parameter equality entirely in our experiments.

C Proof of Proposition 2

Proposition 2. *Suppose Δ_N has n variables and treewidth w , and Δ_F has n' variables and treewidth w' . Then Δ_{joint} has exactly $n + n' - |\lambda_X|$ variables, and treewidth at most $\max(w, w', \min(w, w') + |\lambda_X|)$.*

Denote the set of n CNF variables in Δ_N as B_N , and the set of n' CNF variables in Δ_F as B_F . The only CNF variables shared between Δ_N and Δ_F are the indicators λ_x for BN variables $X \in \mathbf{X}$. Thus the joint encoding $\Delta_{joint} = \Delta_{bn} \wedge \Delta_F$ has $n + n' - |\lambda_X|$ variables.

The treewidth of a CNF formula is defined to be the treewidth of its interaction graph \mathcal{G}_Δ :

Definition 11. *The interaction graph \mathcal{G}_Δ of a CNF formula Δ is a graph with CNF variables as nodes, and edges between each variables appearing in the same CNF clause.*

Let π be an ordering of the nodes in \mathcal{G}_Δ , called an *elimination order*. The treewidth of π with respect to \mathcal{G}_Δ is defined by the following procedure. Remove the CNF nodes according to π , and connect all pairs of nodes connected to a removed node immediately after removing it. Then the treewidth is the maximum number of edges a variable has immediately before it is removed. The treewidth of \mathcal{G}_Δ is then defined to be the minimum treewidth among all elimination orders π .

Let π_N^* be the optimal such ordering on \mathcal{G}_{Δ_N} , which has treewidth w . Let π_{N-X}^* be that ordering with all CNF variables in λ_X removed. Then consider removing nodes from the interaction graph $\mathcal{G}_{\Delta_{joint}}$ in that order. At the time of removal of any node b in π_{N-X}^* , it will not be connected to any

node in $B_{\mathcal{N}} \setminus \lambda_{\mathcal{X}}$ in $\mathcal{G}_{\Delta_{joint}}$ that it was not connected to at the time of its removal in $\pi_{\mathcal{N}}^*$ in $\mathcal{G}_{\Delta_{\mathcal{N}}}$. This is since the interspersed removals of nodes in $\lambda_{\mathcal{X}}$ in the $\pi_{\mathcal{N}}^*$ ordering can only have added more edges between nodes in $B_{\mathcal{N}} \setminus \lambda_{\mathcal{X}}$. However, in general, b may be connected to any or all of the nodes in $\lambda_{\mathcal{X}}$. Thus the maximal degree of any removed node (at the time of its removal) is at most $w + |\lambda_{\mathcal{X}}|$.

After removing nodes in $B_{\mathcal{N}} \setminus \lambda_{\mathcal{X}}$, only nodes in B_F remain. Let π_F^* be the optimal ordering on \mathcal{G}_{Δ_F} with treewidth w' . Then we can remove those nodes according to that ordering, such that the maximal degree of any removed node (at the time of its removal) is at most w' .

The treewidth of the ordering $\pi = (\pi_{\mathcal{N}-\mathcal{X}}^*, \pi_F^*)$ with respect to $\mathcal{G}_{\Delta_{joint}}$ is at most $\max(w + |\lambda_{\mathcal{X}}|, w')$, and thus the treewidth of Δ_{joint} is at most this quantity. If we consider an ordering which follows a similar procedure to above, except that we remove nodes $B_F \setminus \lambda_{\mathcal{X}}$ first, we get another bound $\max(w, w' + |\lambda_{\mathcal{X}}|)$. Thus the treewidth of Δ_{joint} is at most $\min(\max(w + |\lambda_{\mathcal{X}}|, w'), \max(w, w' + |\lambda_{\mathcal{X}}|)) = \max(w, w', \min(w, w') + |\lambda_{\mathcal{X}}|)$.

D Proof of Theorem 2

Theorem 2. *Given a parametric/structural intervention set $\mathcal{I}_{\mathcal{N}_F}$, let \mathcal{AC} be an arithmetic circuit with the same polynomial as \mathcal{N}_F , and satisfying the ordering constraints associated with the intervention set. Then, applying the UB algorithm $UB(\mathcal{AC}, e, \mathbf{W})$ returns a quantity B_U which is an upper bound on the interventional robustness quantity $IntRob(\mathcal{I}_{\mathcal{N}_F}, e)$.*

In what follows, the intervention set (including intervenable variables \mathbf{W}) and evidence e is fixed throughout, and thus is dropped from notation for brevity.

D.1 Definitions

First, we review the effect of an intervention on the network polynomial of \mathcal{N}_F . For $W_i \in \mathbf{W}$, let U_i denote the parents of W ($pa_G(W)/pa_{G'}(W)$ for parametric/structural intervention sets respectively). Then, any intervention in the set $\mathcal{I}_{\mathcal{N}_F}$ can be specified as a list of functions $\mathbf{f} = (f_1, \dots, f_{|\mathbf{W}|})$, where f_i maps from U_i to W_i , and we write $\mathcal{N}'_{\mathbf{f}}$ to represent the intervened network. This corresponds to setting, for every instantiation \mathbf{u}_i of U_i :

- $\theta_{f(\mathbf{u}_i)|\mathbf{u}_i} := 1$;
- $\theta_{w_i|\mathbf{u}_i} := 0$ for any $w_i \neq f(\mathbf{u}_i)$.

The network polynomial $l_{\mathcal{N}'_{\mathbf{f}}}[\lambda, \Theta]$ of the intervened network \mathcal{N}' is then given by applying these changes to each term of $l_{\mathcal{N}_F}[\lambda, \Theta]$. This can be achieved by setting all parameters $\theta_{f(\mathbf{u}_i)|\mathbf{u}_i}$ for intervenable variables $W_i \in \mathbf{W}$ to 1, then filtering out those which are incompatible with \mathbf{f} . The following definitions of *weight* and *consistency* capture these notions, given evidence e :

Definition 12. *The f-consistency $cs_{\mathbf{f}}(\alpha)$ of a term α of $l_{\mathcal{N}_F}[\lambda, \Theta]$ is defined as:*

$$cs_{\mathbf{f}}(\alpha) = \begin{cases} 1 & \text{each parameter in } \alpha \text{ of the form } \theta_{W|U} \\ & \text{for } W \in \mathbf{W} \text{ is assigned 1 by } \mathbf{f} \\ 0 & \text{otherwise} \end{cases}$$

We say that α is f-consistent if $cs_{\mathbf{f}}(\alpha) = 1$, and f-inconsistent otherwise.

Definition 13. *The weight $w(\alpha)$ of a term α of $l_{\mathcal{N}_F}[\lambda, \Theta]$ is obtained by evaluating α after:*

- Assigning 1 to indicators if they are consistent with e , and 0 otherwise;
- Setting all parameters $\theta_{f(\mathbf{u}_i)|\mathbf{u}_i}$ to 1 (other parameters are unchanged)

Definition 14. *The f-value $val_{\mathbf{f}}(\alpha)$ of a term α of $l_{\mathcal{N}_F}[\lambda, \Theta]$ is defined as:*

$$val_{\mathbf{f}}(\alpha) = w(\alpha) \times cs_{\mathbf{f}}(\alpha)$$

For any set of terms S , we also define:

$$val_{\mathbf{f}}(S) = \sum_{\alpha \in S} val_{\mathbf{f}}(\alpha)$$

Proposition 4. *Let $S()$ be the set of all terms in $l_{\mathcal{N}_F}[\lambda, \Theta]$. Then:*

$$p_{\mathcal{N}'_{\mathbf{f}}}(e) = val_{\mathbf{f}}(S())$$

This proposition simply captures how marginal probabilities for intervened Bayesian networks are calculated.

Now we will define similar concepts for the arithmetic circuit \mathcal{AC} . First, recall the definition of a *complete subcircuit* of an AC:

Definition 4 (Complete Subcircuit). *A complete subcircuit α of an AC is obtained by traversing the circuit top-down, choosing one child of every visited $+$ -node and all children of every visited \times -node. The term $term(\alpha)$ of α is the product of all leaf nodes visited (i.e. all indicator and parameter variables). The set of all complete subcircuits is denoted $\alpha_{\mathcal{AC}}$.*

A *partial subcircuit* is a subset of the nodes (and edges from those nodes) chosen by any complete subcircuit. The *term* of a partial subcircuit α is also the product of all leaf nodes in α . Often it will be useful to partition a complete subcircuit into a *prefix* α_P and *suffix* α_S : that is, two partial subcircuits which partition the nodes of a complete subcircuit, such that all nodes in the prefix are non-descendants of those in the suffix.

Since \mathcal{AC} has the same polynomial as \mathcal{N}_F , each complete subcircuit corresponds to a term of the network polynomial, and we can define very similar notions to those above for the network polynomial, but based on the UB algorithm.

First, the UB algorithm (Algorithm 1) effectively discards some subcircuits by maximizing at some $+$ -nodes:

Definition 15. *The UB-consistency $cs_{UB}(\alpha)$ of a complete or partial subcircuit α is defined as:*

$$cs_{UB}(\alpha) = \begin{cases} 1 & \text{if for every } +\text{-node } t \text{ in } \alpha \text{ splitting on} \\ & W \in \mathbf{W}, \alpha \text{ chooses the same branch} \\ & \text{from } t \text{ as the UB algorithm} \\ 0 & \text{otherwise} \end{cases}$$

We say that α is UB-consistent if $cs_{UB}\mathbf{f}(\alpha) = 1$, and UB-inconsistent otherwise.

Second, note that the UB algorithm sets parameters for intervenable variables $W \in \mathbf{W}$ to 1. Thus, for any complete

subcircuit α (corresponding to a network polynomial term), in a slight abuse of notation, we can define the weight $w(\alpha)$ in the same way as for network polynomial terms. We can also define the weight for partial subcircuits in the obvious way.

Definition 16. *The weight $w(\alpha)$ of a complete or partial subcircuit α is obtained by evaluating $\text{term}(\alpha)$ after:*

- Assigning 1 to indicators if they are consistent with \mathbf{e} , and 0 otherwise;
- Setting all parameters $\theta_{f(\mathbf{u}_i)|\mathbf{u}_i}$ to 1 (other parameters are unchanged)

Definition 17. *The UB-value $\text{val}_{UB}(\alpha)$ of a complete or partial subcircuit α is defined as:*

$$\text{val}_{UB}(\alpha) = w(\alpha) \times \text{cs}_{UB}(\alpha)$$

For any set of complete or partial subcircuits S , we also define:

$$\text{val}_{UB}(S) = \sum_{\alpha \in S} \text{val}_{UB}(\alpha)$$

Note that the only distinction from \mathbf{f} -value is that we use UB-consistency instead of \mathbf{f} -consistency, i.e. we set different complete subcircuits/terms to 0.

Proposition 5. *Let $S()$ be the set of all complete subcircuits in the AC. Then:*

$$UB(\mathcal{AC}, \mathbf{e}, \mathbf{W}) = \text{val}_{UB}(S())$$

This proposition simply states that assigning value 0 to subcircuits unselected by the UB-algorithm and value $w(\alpha)$ to selected subcircuits produces the output of the UB-algorithm.

Recall that our goal is to show that the UB-algorithm produces an upper bound to $\max_{\mathcal{N}' \in \mathcal{I}_{\mathcal{N}'}} p_{\mathcal{N}'}(\mathbf{e}) = \max_{\mathbf{f}} p_{\mathcal{N}'_{\mathbf{f}}}(\mathbf{e})$. By the two Propositions, this is equivalent to showing:

$$\text{val}_{UB}(S()) \geq \max_{\mathbf{f}} \text{val}_{\mathbf{f}}(S()) \quad (17)$$

D.2 Inductive Lemma

In order to prove Equation 17, we will need to prove a Lemma, which relies on some additional definitions.

Definition 18 (Subcircuit sets). *We will use the following notation to denote various sets of complete and partial subcircuits. Here, t denotes a +-node in the AC, and $\mathbf{t} = \{t_1, \dots, t_j\}$ denotes a set of +-nodes such that **no node is a descendant of another**.*

- $S(\mathbf{t})$ denotes all complete subcircuits which cross all +-nodes in \mathbf{t} ;
- $S_P(\mathbf{t})$ denotes all "prefix" partial subcircuits, obtained by selecting all visited nodes which are non-descendants of t from complete subcircuits in $S(\mathbf{t})$;
- $S_S(t)$ denotes all "suffix" partial subcircuits, obtained by traversing the AC top-down starting at t , choosing one child of every +-node and every child of every \times -node.
- $S(\mathbf{t}, \alpha_0)$ denotes a set of complete subcircuits, defined for any $\alpha_0 \in S_P(\mathbf{t}) : S(\mathbf{t}, \alpha_0) = \{(\alpha_0, \alpha_1, \dots, \alpha_j) : \alpha_1 \in S_S(t_1), \dots, \alpha_j \in S_S(t_j)\}$

We note a few important facts regarding these sets. Firstly, $S(\mathbf{t})$ can be empty, if \mathbf{t} includes two nodes which never appear in the same complete subcircuit. Second, $S(\mathbf{t})$ can be represented as the product set $S_P(\mathbf{t}) \times S_S(t_1) \times \dots \times S_S(t_j)$, and $\{S(\mathbf{t}, \alpha_0) : \alpha_0 \in S_P(\mathbf{t})\}$ forms a partition of $S(\mathbf{t})$. Finally, both the UB-consistency and weight of a complete subcircuit $\alpha = (\alpha_0, \alpha_1, \dots, \alpha_j) \in S(\mathbf{t})$ can be decomposed. α is consistent iff its prefix and suffixes are all consistent, that is:

$$\text{cs}_{UB}(\alpha) = \text{cs}_{UB}(\alpha_0) \prod_{i=1}^j \text{cs}_{UB}(\alpha_i)$$

In addition, the weight of α is given by:

$$w(\alpha) = w(\alpha_0) \prod_{i=1}^j w(\alpha_i)$$

We now define a modified version of UB-value for a complete subcircuit α which crosses nodes \mathbf{t} . Informally, this assumes the prefix α_0 to be consistent for the purposes of computing value.

Definition 19. *Given a set of +-nodes \mathbf{t} , and a complete subcircuit α crossing \mathbf{t} which consists of prefix subcircuit $\alpha_0 \in S_P(\mathbf{t})$ and suffix subcircuits $\alpha_i \in S_S(t_i)$, the prefix-consistent BU-value $\text{val}_{\text{pref}, \mathbf{t}}(\alpha)$ is defined as:*

$$\text{val}_{\text{pref}, \mathbf{t}}(\alpha) = w(\alpha) \times \prod_{i=1}^j \text{cs}_{UB}(\alpha_i)$$

For any set of such complete subcircuits S ,

$$\text{val}_{\text{pref}, \mathbf{t}}(S) = \sum_{\alpha \in S} \text{val}_{\text{pref}, \mathbf{t}}(\alpha)$$

Now, we can state our main result:

Lemma 4. *Let \mathbf{t} be a set of +-nodes, such that no $t \in \mathbf{t}$ is a descendant of another. Further, let $\alpha_0 \in S_P(\mathbf{t})$. Then:*

$$\text{val}_{\text{pref}, \mathbf{t}}(S(\mathbf{t}, \alpha_0)) \geq \max_{\mathbf{f}} \text{val}_{\mathbf{f}}(S(\mathbf{t}, \alpha_0)) \quad (18)$$

As previously stated, for many sets \mathbf{t} , $S(\mathbf{t}, \alpha_0)$ may be empty; in such cases, the Lemma is trivially true.

Before proving this Lemma, we demonstrate how this can be used to derive the inequality in Equation 17, and thus, the Theorem. Take \mathbf{t} to be the singleton set $\{r\}$, where r is the root +/MAX-node of the AC. Then the (only) prefix α_0 of \mathbf{t} is empty, so that $\text{val}_{\text{pref}, \mathbf{t}}(S(\mathbf{t}, \alpha_0)) = \text{val}_{UB}(S(\mathbf{t}, \alpha_0))$. Further, $S(\mathbf{t}, \alpha_0) = S()$ consists of all subcircuits in the AC. Thus, in this case, the Lemma reduces to Equation 17.

Proof of Lemma 4. We will prove the Lemma by induction. Let $\pi = \{\pi_1, \dots, \pi_n\}$ be a reverse topological ordering of the n +-nodes in the AC (that is, descendants come before ancestors), and we write $t < t'$ to indicate a node t comes before t' in the ordering. In particular, our inductive hypothesis at step k will be that the lemma holds for all subsets \mathbf{t} containing only nodes in $\pi_{\leq k} = \{\pi_1, \dots, \pi_k\}$.

At step k (for $1 \leq k \leq n$), we need to show the inequality in Equation 18 for any set of +-nodes $\mathbf{t}' = \{t'_1, \dots, t'_{j-1}, t'_j\}$ s.t. $t'_j = \pi_k$, $0 \leq j \leq k$, and $t'_i \in \pi_{\leq k} \forall 1 \leq i \leq j-1$ (so that $t'_j := \pi_k$), and for any $\alpha_0 \in S_P(\mathbf{t}')$. The proof for step

k follows one of 2 templates, depending on the node π_k , in particular, whether it splits on a variable in \mathbf{W} .

Does not split on \mathbf{W}

Suppose that π_k has m branches. For each branch $1 \leq b \leq m$, starting from the prefix α_0 , add π_k , and then traverse down branch b , adding all children of each \times -node, until we reach a leaf node (included) or a $+$ -node (excluded). Denote this extended prefix subcircuit $\alpha_0^{(b)} = (\alpha_0, \alpha^{(b)})$ (where $\alpha^{(b)}$ is defined to be the partial subcircuit added), and the $+$ -nodes reached by $\mathbf{t}^{(b)} = \{t_1^{(b)}, \dots, t_l^{(b)}\}$ (possibly empty). Then define $\mathbf{t}'_{-j} = \{t'_1, \dots, t'_{j-1}\}$, that is, \mathbf{t}' with π_k removed, and define $\mathbf{t}^{(b)} = (\mathbf{t}'_{-j}, \mathbf{t}^{(b)})$. This fulfils the condition that no two nodes are descendants of each other, since $\mathbf{t}^{(b)}$ are immediate descendants of π_k .

Every complete subcircuit $\alpha \in S(\mathbf{t}^{(b)})$ is then of the form $(\alpha_0^{(b)}, \alpha_1, \dots, \alpha_{j-1}, \alpha_1^{(b)}, \dots, \alpha_l^{(b)})$, with prefix $\alpha_0^{(b)} \in S_P(\mathbf{t}^{(b)})$, and suffixes $\alpha_i \in S_S(t'_i)$ ($1 \leq i \leq j-1$) and $\alpha_a^{(b)} \in S_S(t_a^{(b)})$ ($1 \leq a \leq l$).

Recall that $S(\mathbf{t}^{(b)}, \alpha_0^{(b)})$ denotes the set of all complete subcircuits with prefix $\alpha_0^{(b)}$ and any combination of suffix subcircuits. The prefix-consistent BU-value for these subcircuits is greater than the \mathbf{f} -value for any \mathbf{f} :

$$\begin{aligned} & \text{val}_{\text{pref}, \mathbf{t}^{(b)}}(S(\mathbf{t}^{(b)}, \alpha_0^{(b)})) \\ & \geq \max_{\mathbf{f}} \text{val}_{\mathbf{f}}(S(\mathbf{t}^{(b)}, \alpha_0^{(b)})) \end{aligned}$$

This follows directly from the inductive hypothesis for step $k-1$. This is because \mathbf{t}'_{-j} only contains nodes from $\pi_{\leq k-1}$ by definition, and $\mathbf{t}^{(b)}$ are descendants of π_k , so by the fact π is a reverse topological order, they too are taken from $\pi_{\leq k-1}$.

Further, $\{S(\mathbf{t}^{(b)}, \alpha_0^{(b)}) : 1 \leq b \leq m\}$ forms a partition of $S(\mathbf{t}', \alpha_0)$ (each set consists of the set of subcircuits following a particular branch from π_k).

Now consider the UB-consistency of the suffix subcircuit from t'_j , $\alpha_j = (\alpha^{(b)}, \alpha_1^{(b)}, \dots, \alpha_l^{(b)})$. We have $cs_{UB}(\alpha_j) = cs_{UB}(\alpha^{(b)}) \times \prod_{a=1}^l cs_{UB}(\alpha_a^{(b)}) = \prod_{a=1}^l cs_{UB}(\alpha_a^{(b)})$. The final equality follows from the fact that $t'_j = \pi_k$ does not split on $W \in \mathbf{W}$ and is the only $+$ -node in $\alpha^{(b)}$, so $\alpha^{(b)}$ contains no $+$ -nodes splitting on $W \in \mathbf{W}$, and it follows that $cs_{UB}(\alpha^{(b)}) = 1$.

We use this to combine the results for each branch:

$$\begin{aligned} & \text{val}_{\text{pref}, \mathbf{t}'}(S(\mathbf{t}', \alpha_0)) \\ & = \sum_{\alpha \in S(\mathbf{t}', \alpha_0)} w(\alpha) \times \prod_{i=1}^j cs_{UB}(\alpha_i) \\ & = \sum_{b=1}^m \sum_{\alpha \in S(\mathbf{t}^{(b)}, \alpha_0^{(b)})} w(\alpha) \times \prod_{i=1}^j cs_{UB}(\alpha_i) \\ & = \sum_{b=1}^m \sum_{\alpha \in S(\mathbf{t}^{(b)}, \alpha_0^{(b)})} w(\alpha) \times \prod_{i=1}^{j-1} cs_{UB}(\alpha_i) \end{aligned}$$

$$\times \prod_{a=1}^l cs_{UB}(\alpha_a^{(b)})$$

$$\begin{aligned} & = \sum_{b=1}^m \text{val}_{\text{pref}, \mathbf{t}^{(b)}}(S(\mathbf{t}^{(b)}, \alpha_0^{(b)})) \\ & \geq \sum_{b=1}^m \max_{\mathbf{f}} \text{val}_{\mathbf{f}}(S(\mathbf{t}^{(b)}, \alpha_0^{(b)})) \\ & \geq \max_{\mathbf{f}} \sum_{b=1}^m \text{val}_{\mathbf{f}}(S(\mathbf{t}^{(b)}, \alpha_0^{(b)})) \\ & = \max_{\mathbf{f}} \text{val}_{\mathbf{f}}(S(\mathbf{t}', \alpha_0)) \end{aligned}$$

The second equality uses the partition, the third is by the facts about consistency shown above, the fourth equality is by definition, the fifth inequality is by inductive hypothesis, the sixth is a standard sum/max swap, and the final equality is again by the partition.

Does split on \mathbf{W}

In the case where π_k splits on some $W \in \mathbf{W}$, i.e. is a *MAX*-node, the above template doesn't work because the BU algorithm now chooses one of the branches, rather than adding all branches together.

For a given intervention \mathbf{f} , denote the subset of \mathbf{f} -consistent complete subcircuits crossing \mathbf{t}' by $S_{\mathbf{f}}(\mathbf{t}', \alpha_0)$. In this context, we will consider such a complete subcircuit α to be split into a prefix consisting of $\alpha_{-j} = (\alpha_0, \alpha_1, \dots, \alpha_{j-1})$ with $\alpha_0 \in S_P(\mathbf{t}')$, $\alpha_i \in S_S(t'_i)$, and a single suffix $\alpha_j \in S_S(t'_j)$. The following two definitions are useful:

- First, we define the set of " \mathbf{f} -consistent prefixes":

$$S_{-j, \mathbf{f}}(\mathbf{t}') = \{\alpha_{-j} : \exists \alpha_j \text{ s.t. } cs_{\mathbf{f}}(\alpha) = 1\} \quad (19)$$

- Second, for a given prefix α_{-j} , define the set of complete subcircuits $S_{\mathbf{f}}(\pi_k, \alpha_{-j})$ to be the subset of $S(\pi_k, \alpha_{-j})$ which is \mathbf{f} -consistent. This is empty if the prefix α_{-j} is not consistent.

Our strategy will be to show that we can modify \mathbf{f} so that all \mathbf{f} -consistent subcircuits in $S(\mathbf{t}', \alpha_0)$ include the same branch of π_k , without reducing the \mathbf{f} -value. We can then use our inductive hypothesis to bound this value. First, we show that given α_{-j} , any resulting complete subcircuit must go down the same branch:

Lemma 5. *For a given \mathbf{f} and α_{-j} , subcircuits $S_{\mathbf{f}}(\pi_k, \alpha_{-j})$ all choose the same branch at π_k .*

Proof. Here we use our crucial assumption of the AC ordering constraints, namely, that no node that is a descendant of π_k splits on any variable in $\mathbf{U} = pa(W)$. Since subcircuits in $S(\pi_k, \alpha_{-j})$ differ only on their suffix from π_k , they all share the same indicators for variables in $pa(W)$, corresponding to an instantiation \mathbf{u} . Suppose $f_W(\mathbf{u}) = w$. Then, any parameter $\theta_{W=w' | \mathbf{U}=\mathbf{u}}$ for $w' \neq w$ is intervened to 0. Since π_k splits on W , all branches from π_k have different indicators for W . Thus, only the branch with indicator corresponding to $W = w$ can have \mathbf{f} -consistent subcircuits. \square

This does not preclude, however, that there may be \mathbf{f} -consistent subcircuits with different prefixes α_{-j}, β_{-j} which choose different branches. The following Lemma shows that we can find another intervention, \mathbf{f}_{adj} , that does always (effectively) choose the same branch for any prefix, and additionally has greater or equal value.

Lemma 6. *For any intervention \mathbf{f} , there exists another intervention \mathbf{f}_{adj} such that:*

- $val_{\mathbf{f}_{adj}}(S(\mathbf{t}', \alpha_0)) \geq val_{\mathbf{f}}(S(\mathbf{t}', \alpha_0))$
- All subcircuits $\alpha \in S_{\mathbf{f}_{adj}}(\mathbf{t}', \alpha_0)$ with non-zero weight $w(\alpha)$ choose the same branch at π_k (call it b_{adj})

Proof. We will define a non-constructive operation `upgrade`, which takes as input an intervention \mathbf{f} , and outputs an intervention \mathbf{f}_{up} .

Consider all subcircuits $\alpha = (\alpha_0, \alpha_1, \dots, \alpha_{j-1}, \alpha_j) \in S_{\mathbf{f}}(\pi_k, \alpha_{-j})$. Define $S_{j,\mathbf{f}}(\mathbf{t}', \alpha_{-j}) \subseteq S_S(t'_j)$ as follows:

$$S_{j,\mathbf{f}}(\mathbf{t}', \alpha_{-j}) = \{\alpha_j \in S_S(t'_j) : \alpha \in S_{\mathbf{f}}(\pi_k, \alpha_{-j})\}$$

That, is for a given prefix α_{-j} and \mathbf{f} , this is the set of suffix partial subcircuits from π_k which are \mathbf{f} -consistent.

The "optimal set" of suffixes over \mathbf{f} and α_j is defined to be that which attains the greatest combined weight, that is, $\max_{\mathbf{f}} \max_{\alpha_{-j}} w(S_{j,\mathbf{f}}(\mathbf{t}', \alpha_{-j}))$. In general the optimal set will not be unique; we use S_j^* to denote the set of optimal suffix sets. We pick (arbitrarily) some $S_j^* \in S_j^*$, and denote the intervention \mathbf{f} which attains this \mathbf{f}^* , and the prefix which attains this α_{-j}^* .

Now, we will define a set of complete subcircuits $T_{\mathbf{f}}(\mathbf{t}', \alpha_0)$ by taking all \mathbf{f} -consistent prefixes α_{-j} , and taking the product set with the set of suffixes S_j^* . Formally:

$$T_{\mathbf{f}}(\mathbf{t}', \alpha_0) = \{(\alpha_{-j}, \alpha_j) : \alpha_{-j} \in S_{-j,\mathbf{f}}(\mathbf{t}'), \alpha_j \in S_j^*\} \quad (20)$$

Intuitively, we are "increasing the weight" of the \mathbf{f} -consistent subcircuits of $S(\mathbf{t}', \alpha_0)$, while also ensuring consistent subcircuits all choose the same branch (that chosen by S_j^*). However, we cannot be sure that $T_{\mathbf{f}}(\mathbf{t}', \alpha_0)$ actually corresponds to the \mathbf{f}_{up} -consistent subcircuits of any allowed intervention \mathbf{f}_{up} . We will now show that there exists a \mathbf{f}_{up} such that all of these subcircuits are \mathbf{f}_{up} -consistent. That is, $T_{\mathbf{f}}(\mathbf{t}', \alpha_0) \subseteq S_{\mathbf{f}_{up}}(\mathbf{t}', \alpha_0)$.

Suppose for contradiction this was not the case. Then there must exist two subcircuits $\alpha, \beta \in T_{\mathbf{f}}(\mathbf{t}', \alpha_0)$ which are contradictory, that is, for some $W \in \mathbf{W}$ and its parents $U = pa(W)$, α, β must contain indicators for the same value of $U = \mathbf{u}$, but indicators $\lambda_{W=w_\alpha}, \lambda_{W=w_\beta}$ for different values of W respectively (that is, $w_\alpha \neq w_\beta$).

We can write $\alpha = (\alpha_{-j}, \alpha_j)$ and $\beta = (\beta_{-j}, \beta_j)$. Now, we consider three cases according to where the contradictory indicators reside:

- **If** $\lambda_{W=w_\alpha} \in \alpha_{-j}, \lambda_{W=w_\beta} \in \beta_{-j}$

Recall that by definition of $T_{\mathbf{f}}(\mathbf{t}', \alpha_0)$, $\alpha_{-j}, \beta_{-j} \in S_{-j,\mathbf{f}}(\mathbf{t}')$, that is, the set of \mathbf{f} -consistent prefixes. Then there must exist $\alpha'_j, \beta'_j \in S(t'_j)$ such that $\alpha' = (\alpha_{-j}, \alpha'_j)$ and $\beta' = (\beta_{-j}, \beta'_j)$ are both \mathbf{f} -consistent. But this is a contradiction as two complete subcircuits

with conflicting indicators $\lambda_{W=w_\alpha}, \lambda_{W=w_\beta}$ cannot simultaneously be consistent for any \mathbf{f} .

- **If** $\lambda_{W=w_\alpha} \in \alpha_j, \lambda_{W=w_\beta} \in \beta_j$

Recall that by definition of $T_{\mathbf{f}}(\mathbf{t}', \alpha_0)$, $\alpha_j, \beta_j \in S_j^*$, that is, they both belong to $S_{j,\mathbf{f}^*}(\mathbf{t}', \alpha_{-j}^*)$. In other words, $\alpha^* = (\alpha_{-j}^*, \alpha_j)$ and $\beta^* = (\alpha_{-j}^*, \beta_j)$ are both \mathbf{f}^* -consistent. But again this is a contradiction, as they contain conflicting indicators $\lambda_{W=w_\alpha}, \lambda_{W=w_\beta}$.

- **Otherwise**

If neither of the above cases hold, then either $\lambda_{W=w_\alpha} \in \alpha_{-j}, \lambda_{W=w_\beta} \in \beta_j$ or $\lambda_{W=w_\alpha} \in \alpha_j, \lambda_{W=w_\beta} \in \beta_{-j}$. Without loss of generality, let us assume the former. Then consider the subcircuit $\alpha' = (\alpha_{-j}, \beta_j)$. This single complete subcircuit contains both indicators $\lambda_{W=w_\alpha}, \lambda_{W=w_\beta}$, which is a contradiction as no term-complete subcircuit can do so.

Thus, there exists some \mathbf{f}_{up} such that $T_{\mathbf{f}}(\mathbf{t}', \alpha_0) \subseteq S_{\mathbf{f}_{up}}(\mathbf{t}', \alpha_0)$. Our operation `upgrade` outputs any such \mathbf{f}_{up} . \mathbf{f}_{up} has three important properties:

1. **Does not remove consistent prefixes**

$$S_{-j,\mathbf{f}_{up}}(\mathbf{t}') \supseteq S_{-j,\mathbf{f}}(\mathbf{t}')$$

That is, the set of \mathbf{f}_{up} -consistent prefixes subsumes the set of \mathbf{f} -consistent prefixes. This follows from the fact that $T_{\mathbf{f}}(\mathbf{t}', \alpha_0) \subseteq S_{\mathbf{f}_{up}}(\mathbf{t}', \alpha_0)$: $T_{\mathbf{f}}(\mathbf{t}', \alpha_0)$ contains subcircuits with all prefixes in $S_{-j,\mathbf{f}}(\mathbf{t}')$.

2. **All suffixes for previous consistent prefixes are in S_j^* or have zero weight**

$$\forall \alpha_{-j} \in S_{-j,\mathbf{f}}(\mathbf{t}'), \forall \alpha_j \in S_{j,\mathbf{f}_{up}}(\mathbf{t}', \alpha_{-j}), \text{ either } w(\alpha_j) = 0 \text{ or } \alpha_j \in S_j^*.$$

That is, for all \mathbf{f} -consistent prefixes α_{-j} , the corresponding \mathbf{f}_{up} -consistent subcircuits $S_{j,\mathbf{f}_{up}}(\mathbf{t}', \alpha_{-j})$ are either contained in the optimal suffix set S_j^* or have weight 0.

This is shown as follows. From $T_{\mathbf{f}}(\mathbf{t}', \alpha_0) \subseteq S_{\mathbf{f}_{up}}(\mathbf{t}', \alpha_0)$, we know that $S_j^* \subseteq S_{j,\mathbf{f}_{up}}(\mathbf{t}', \alpha_{-j})$ for any $\alpha_{-j} \in S_{-j,\mathbf{f}}(\mathbf{t}')$ (i.e. \mathbf{f} -consistent prefixes). If $S_{j,\mathbf{f}_{up}}(\mathbf{t}', \alpha_{-j})$ contains some $\alpha_j \notin S_j^*$, then it must have weight 0, otherwise $w(S_{j,\mathbf{f}_{up}}(\mathbf{t}', \alpha_{-j})) \geq w(S_j^*)$ which is a contradiction as S_j^* is an optimal suffix set.

3. **Does not decrease value**

$$val_{\mathbf{f}_{up}}(S(\mathbf{t}, \alpha_0)) \geq val_{\mathbf{f}}(S(\mathbf{t}, \alpha_0))$$

This is proven as follows:

$$\begin{aligned} & val_{\mathbf{f}_{up}}(S(\mathbf{t}, \alpha_0)) \\ &= \sum_{\alpha \in S(\mathbf{t}, \alpha_0)} w(\alpha) \times cs_{\mathbf{f}_{up}}(\alpha) \\ &= \sum_{\alpha \in S_{\mathbf{f}_{up}}(\mathbf{t}, \alpha_0)} w(\alpha) \\ &\geq \sum_{\alpha \in T_{\mathbf{f}}(\mathbf{t}, \alpha_0)} w(\alpha) \\ &= \sum_{\alpha_{-j} \in S_{-j,\mathbf{f}}(\mathbf{t}')} \sum_{\alpha_j \in S_j^*} w(\alpha) \end{aligned}$$

$$\begin{aligned}
&= \sum_{\alpha_{-j} \in S_{-j, \mathbf{f}}(\mathbf{t}')} \sum_{\alpha_j \in S_j^*} w(\alpha_{-j}) \times w(\alpha_j) \\
&\geq \sum_{\alpha_{-j} \in S_{-j, \mathbf{f}}(\mathbf{t}')} \sum_{\alpha \in S_{\mathbf{f}}(\pi_k, \alpha_{-j})} w(\alpha_{-j}) \times w(\alpha_j) \\
&= \sum_{\alpha \in S_{\mathbf{f}}(\mathbf{t}', \alpha_0)} w(\alpha_{-j}) \times w(\alpha_j) \\
&= \sum_{\alpha \in S_{\mathbf{f}}(\mathbf{t}', \alpha_0)} w(\alpha) \\
&= \sum_{\alpha \in S(\mathbf{t}', \alpha_0)} w(\alpha) \times cs_{\mathbf{f}}(\alpha) \\
&= val_{\mathbf{f}}(S(\mathbf{t}', \alpha_0))
\end{aligned}$$

Consider applying this operation iteratively, producing a sequence of interventions $\mathbf{f}_{up}^{(1)}, \mathbf{f}_{up}^{(3)}, \dots$. For each intervention $\mathbf{f}_{up}^{(i)}$, we can associate two quantities: $m_1^{(i)}$, the number of \mathbf{f}_{up} -consistent prefixes, and among these prefixes, the number $m_2^{(i)} (< m_1^{(i)})$ of prefixes which have corresponding suffixes all in S_j^* or having zero weight. Property 1 then tells us that $m_1^{(i+1)} \geq m_1^{(i)}$, and Property 2 tells us that $m_2^{(i+1)} \geq m_2^{(i)}$. Since there are a finite number of possible prefixes, we must (in finite time) obtain some \mathbf{f}_{up}^i with $m_1^{(i)} = m_2^{(i)}$, which we will call \mathbf{f}_{adj} . Since all \mathbf{f}_{adj} -consistent prefixes have corresponding suffixes in S_j^* or with weight zero, by Lemma 5 this intervention indeed has the property that all subcircuits take the same branch or have weight zero, and further by Property 3 the \mathbf{f}_{adj} -value of $S(\mathbf{t}, \alpha_0)$ is at least \mathbf{f} -value of $S(\mathbf{t}, \alpha_0)$. \square

The significance of this Lemma is that we can now make use of the inductive hypothesis. As in the case where π_k doesn't split on $W \in \mathbf{W}$, we have $val_{pref, \mathbf{t}'^{(b)}}(S(\mathbf{t}'^{(b)}, \alpha_0^{(b)})) \geq val_{\mathbf{f}_{adj}}(S(\mathbf{t}'^{(b)}, \alpha_0^{(b)}))$ for any branch b . However, since all non-zero weight \mathbf{f}_{adj} -consistent subcircuits follow the same branch, $val_{pref, \mathbf{t}'^{(b_{adj})}}(S(\mathbf{t}'^{(b_{adj})}, \alpha_0^{(b_{adj})}))$ is in fact an upper bound on the \mathbf{f}_{adj} -value summed over subcircuits from all branches: $val_{\mathbf{f}}(S(\mathbf{t}', \alpha_0))$.

Now we use this to prove the inductive hypothesis of Lemma 4 for step k . Consider, for each branch b of π_k , the set of partial subcircuits given by $S_S^{(b)}(\pi_k) := \{(\alpha^{(b)}, \alpha_1^{(b)}, \dots, \alpha_l^{(b)}) : \alpha_a^{(b)} \in S_S(\mathbf{t}_a^{(b)})\}$. The UB algorithm chooses the branch with the greatest UB-value **ignoring consistency of $\alpha^{(b)}$** (i.e. which branch is actually chosen), that is,

$$v_b := w(\alpha^{(b)}) \times \sum_{\alpha_1^{(b)}, \dots, \alpha_l^{(b)}} \prod_{a=1}^l w(\alpha_a^{(b)}) \times cs_{UB}(\alpha_a^{(b)})$$

Suppose that the UB algorithm chooses branch b_{UB} at π_k , so that $cs_{UB}(\alpha^{(b_{UB})}) = 1$ and $cs_{UB}(\alpha^{(b)}) = 0$ for all $b \neq b_{UB}$. Then, for a given \mathbf{f} , we have that:

$$val_{pref, \mathbf{t}'}(S(\mathbf{t}', \alpha_0))$$

$$\begin{aligned}
&= \sum_{\alpha \in S(\mathbf{t}', \alpha_0)} w(\alpha) \times \prod_{i=1}^j cs_{UB}(\alpha_i) \\
&= \sum_{b=1}^m \sum_{\alpha \in S(\mathbf{t}'^{(b)}, \alpha_0^{(b)})} w(\alpha) \times \prod_{i=1}^j cs_{UB}(\alpha_i) \\
&= \sum_{b=1}^m \sum_{\alpha \in S(\mathbf{t}'^{(b)}, \alpha_0^{(b)})} w(\alpha) \times \prod_{i=1}^{j-1} cs_{UB}(\alpha_i) \\
&\quad \times cs_{UB}(\alpha^{(b)}) \times \prod_{a=1}^l cs_{UB}(\alpha_a^{(b)}) \\
&= \sum_{b=1}^m \sum_{\alpha \in S(\mathbf{t}'^{(b)}, \alpha_0^{(b)})} w(\alpha_{-j}) \times \prod_{i=1}^{j-1} cs_{UB}(\alpha_i) \\
&\quad \times cs_{UB}(\alpha^{(b)}) \times w(\alpha^{(b)}) \\
&\quad \times \prod_{a=1}^l w(\alpha_a^{(b)}) \times cs_{UB}(\alpha_a^{(b)}) \\
&= \sum_{\alpha_1, \dots, \alpha_{j-1}} w(\alpha_{-j}) \times \prod_{i=1}^{j-1} cs_{UB}(\alpha_i) \\
&\quad \times \sum_{b=1}^m \left(cs_{UB}(\alpha^{(b)}) \times w(\alpha^{(b)}) \right. \\
&\quad \left. \times \sum_{\alpha_1^{(b)}, \dots, \alpha_l^{(b)}} \prod_{a=1}^l w(\alpha_a^{(b)}) \times cs_{UB}(\alpha_a^{(b)}) \right) \\
&= \sum_{\alpha_1, \dots, \alpha_{j-1}} w(\alpha_{-j}) \times \prod_{i=1}^{j-1} cs_{UB}(\alpha_i) \\
&\quad \times \sum_{b=1}^m v_b \times cs_{UB}(\alpha^{(b)}) \\
&= \sum_{\alpha_1, \dots, \alpha_{j-1}} w(\alpha_{-j}) \times \prod_{i=1}^{j-1} cs_{UB}(\alpha_i) \\
&\quad \times v_{b_{UB}} \\
&\geq \sum_{\alpha_1, \dots, \alpha_{j-1}} w(\alpha_{-j}) \times \prod_{i=1}^{j-1} cs_{UB}(\alpha_i) \\
&\quad \times v_{b_{adj}} \\
&= \sum_{\alpha_1, \dots, \alpha_{j-1}} w(\alpha_{-j}) \times \prod_{i=1}^{j-1} cs_{UB}(\alpha_i) \\
&\quad \times w(\alpha^{(b_{adj})}) \times \\
&\quad \sum_{\alpha_1^{(b_{adj})}, \dots, \alpha_l^{(b_{adj})}} \prod_{a=1}^l w(\alpha_a^{(b_{adj})}) \times cs_{UB}(\alpha_a^{(b_{adj})})
\end{aligned}$$

$$\begin{aligned}
&= \sum_{\alpha \in S(\mathbf{t}'^{(b_{adj})}, \alpha_0^{(b_{adj})})} w(\alpha) \times \prod_{i=1}^{j-1} cs_{UB}(\alpha_i) \\
&\quad \times \prod_{a=1}^l cs_{UB}(\alpha_a^{(b_{adj})}) \\
&= val_{pref, \mathbf{t}'^{(b_{adj})}}(S(\mathbf{t}'^{(b_{adj})}, \alpha_0^{(b_{adj})})) \\
&\geq val_{f_{adj}}(S(\mathbf{t}'^{(b_{adj})}, \alpha_0^{(b_{adj})})) \\
&= \sum_{\alpha \in S(\mathbf{t}'^{(b_{adj})}, \alpha_0^{(b_{adj})})} w(\alpha) \times cs_{f_{adj}}(\alpha) \\
&= \sum_{\alpha \in S(\mathbf{t}', \alpha_0)} w(\alpha) \times cs_{f_{adj}}(\alpha) \\
&= val_{f_{adj}}(S(\mathbf{t}', \alpha_0)) \\
&\geq val_f(S(\mathbf{t}', \alpha_0))
\end{aligned}$$

Thus, in both cases (whether π_k splits on $W \in \mathcal{W}$), we have shown that the inductive hypothesis holds at step k . Thus Lemma 4 is proved. \square

As noted previously, this Lemma immediately implies Theorem 2; thus we are done.

E Proof of Proposition 3

Proposition 3. *Algorithm 2 converges to a locally optimal parametric intervention in finite time. Further, if the algorithm is stopped before termination, the current value v will be a lower bound on $\max_{N' \in \mathcal{I}[\mathcal{W}]} P_{N'}(e)$.*

Proof. Because any intervention $I_P[\Theta](\mathcal{N})$ gives a lower bound on the maximum probability of the evidence e , we know that the algorithm will provide a lower bound on the true value. Assuming the algorithm terminates, it will output a locally maximal intervention; this follows immediately from the termination criterion. It remains to show that the algorithm does indeed terminate. We show this by demonstrating that the algorithm corresponds to best-response dynamics on a potential game, for which it is known that best response dynamics converge to a Nash Equilibrium in potential games [Roughgarden, 2010].

We construct the mapping as follows: for each node $V \in \mathcal{W}$ and each instantiation \mathbf{u} of its parents \mathcal{U} , we construct players $p_{V|\mathbf{u}}$, yielding a finite set \mathcal{P} . The action set of each player $p_{V|\mathbf{u}}$ is the finite set of deterministic distributions over the support of V . Recall that it suffices to consider only *deterministic* interventions to solve the interventional robustness problem, and so this restriction to a finite action set does not prevent the algorithm from obtaining a globally optimal value.

The payoff to all players under strategy set $\Theta'_W = \{\theta'_{V|\mathbf{u}} | p_{V|\mathbf{u}} \in \mathcal{P}\}$ is identical and is equal to $P_{N'[\Theta'_W]}(e)$. Because the players have identical payoffs, we can define a potential function for this game as this joint payoff function $P_{N'[\Theta'_W]}(e)$, and observe that if a player follows best-response dynamics, it will necessarily increase the value of

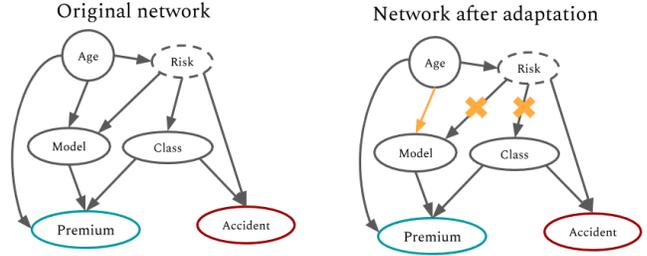


Figure 7: Causal model used by a fictional car insurance company in the example in Section F.

$p(e)$. We therefore obtain that the procedure will eventually converge to a local optimum.

The problem of finding a Nash equilibrium lies in the complexity class PPA [Roughgarden, 2010]. In the worst-case, best-response dynamics can take time proportional to the size of the joint strategy space, which in this case will be exponential in the size of the intervention variable set \mathcal{W} . While anecdotally we do not observe running times anywhere near the worst case, we further note that because the IntRob problem concerns itself with finding a maximum over a set, any element of that set, i.e. any intervention, serves as a witness to give a lower bound on the maximum intervention probability. Thus, the value v of the current intervention is always a lower bound on the IntRob value, and so the algorithm can be prematurely halted after any period of time and still yield a valid lower bound. \square

F Worked example

We will use the simplified car insurance network from Figure 7 to illustrate the interventional robustness problem, and highlight how existing methods cannot provide accurate guarantees for this problem.

Model: we assume the probabilistic model of Figure 7 with CPTs described in Table 4. In particular, we assume that the insurance company decides on premiums based on some combination of the driver’s age, car model, and history of taking driving courses. We will focus on the following decision rule.

```
low if (class = 1) ∧ [(age = <25 ∧ model = budget)
                    ∨ (age = >25 ∧ model = lux)].
```

The target probability in this setting will be $P(e)$ with $e = (\text{Accident} = 1, C1 = \text{low})$, and the intervention set will be $\mathcal{W} = \{\text{model}, \text{class}\}$. The intervention which maximizes $P(e)$ is that in which $\theta'_{\text{model}|\text{age}=>25} = \mathbb{1}[\text{model} = \text{luxury}]$ and $\theta'_{\text{model}|\text{age}=<25} = \mathbb{1}[\text{model} = \text{budget}]$, and $\theta'_{C1|\mathbf{u}} = \mathbb{1}[C = 1]$. We will write these CPTs by $\theta_{w|\mathbf{u}}^*$. Under this intervention, the probability $P_{N'[\Theta^*]}(e) = 0.126$. We set our threshold $\epsilon = 0.1$, and so the classifier C1 is *not* robust to this intervention set.

There are two critical components to answering the interventional robustness problem: the first is accurate estimation of the probability $P_{N'[\Theta^*]}(e)$, and the second is effective search over the intervention space. In what follows, we will show how two approximation methods based on existing methodology fail at each of these components, and therefore fail to

P(accident=1)	P(accident=0)	\mathbf{u}
0.4	0.6	Model=luxury Class=1 risky=1
0.01	0.99	Model=luxury Class=1 risky=0
0.6	0.4	Model=luxury Class=0 risky=1
0.1	0.9	Model=luxury Class=0 risky=0
0.3	0.7	Model=budget Class=1 risky=1
0.05	0.95	Model=budget Class=1 risky=0
0.5	0.5	Model=budget Class=0 risky=1
0.3	0.7	Model=budget Class=0 risky=0

$P(C \mathbf{u})$	$P(C \mathbf{u})$	\mathbf{u}
0.2	0.8	risky=1
0.8	0.2	risky=0

$P(\text{model} = \text{budget} \mathbf{u})$	$P(\text{model} = \text{luxury} \mathbf{u})$	\mathbf{u}
0.3	0.7	age=> 25 risky=1
0.7	0.3	age=> 25 risky=0
0.8	0.2	age=< 25 risky=1
0.2	0.8	age=< 25 risky=0

$P(\text{risky}=0)$	$P(\text{risky}=1)$	\mathbf{u}
0.3	0.7	age=> 25
0.3	0.7	age=< 25

Table 4: CPTs for the insurance network from Figure 1.

effectively answer the intervention robustness problem. This motivates our proposed algorithms.

Sensitivity analysis, when performed exactly, will provide the correct answer to parametric IntRob. However, existing approaches are computationally intractable for large intervention sets. For example, the method of Chan and Darwiche [2004] requires differentiating the network polynomial with respect to every combination of parameters from the intervenable CPTs, which scales very poorly with respect to the size of each CPT and number of intervenable variables $|\mathbf{W}|$. As a workaround, we consider a first-order approximation which consists of summing the partial derivatives of each CPT belonging to an intervenable variable $W \in \mathbf{W}$. This yields the following estimate of $\max_{\Theta'} P_{\mathcal{N}[\Theta']}(e)$.

$$\max_{\Theta'} P_{\mathcal{N}[\Theta']}(e) \approx \sum_{W \in \mathbf{W}} \sum_{\mathbf{u}(W)} \max_w \partial_{\theta_w|\mathbf{u}} P(e)$$

I.e. the approximation estimates the effect of parameter changes by computing the effect of single-CPT changes to $P(e)$, and summing these estimates over all CPTs of interest.

$$\begin{aligned} &= \max_{\text{model}} \partial_{\theta_{\text{model}|\text{age}>25,\cdot}} P(e) + \max_{\text{model}} \partial_{\theta_{\text{model}|\text{age}<25,\cdot}} P(e) \\ &+ \max_{\text{class}} \partial_{\theta_{\text{class}|\cdot}} P(e) \\ &= P_{\mathcal{N}[\theta_{\text{class}|\mathbf{u}}=\mathbb{1}_{\text{class}=1}]}(e) + P_{\mathcal{N}[\theta_{\text{model}|\mathbf{u}}=\theta_{\text{model}|\mathbf{u}}^*]}(e) \\ &= 0.048 + 0.038 \\ &= 0.086 < \epsilon = 0.1 \end{aligned}$$

Thus, under the single-parameter SA approximation, we would incorrectly estimate the classifier to be robust to interventions on model and class.

We could also consider restricting our search over parametric interventions to one over **do-interventions only**. A do-intervention fixes the value of a variable, in contrast to a parametric intervention which changes its conditional distribution. The appeal of do-interventions is that they induce a much smaller search space; whereas the set of all parametric interventions can be of size $2^{2^{|\mathbf{U}|}}$ (in a binary BN), the set of possible do-interventions is only of size $2^{|\mathbf{W}|}$. The differential semantics of do-interventions are discussed by Qin [2015], who show that it is possible to efficiently compute the

effect of a do-intervention on a single variable in time linear in the size of the AC.

In the CBN in Figure 1, $P_{\mathcal{N}}(e)$ is maximized when $\theta_{\text{model}=\text{lux}|\text{age}>25,\cdot} = 1$ and $\theta_{\text{model}=\text{budget}|\text{age}<25,\cdot} = 1$ (i.e. when young people buy budget cars and old people buy luxury cars, independent of their risk appetite), and when all drivers take driving classes. While the latter change in the distribution can be expressed as a do-intervention, the former cannot. The value of model must be set to either 0 or 1 for all drivers. As a result, the worst-case do-intervention sets class to True and model to luxury. Under this intervention Θ' , we see

$$P_{\mathcal{N}[\Theta']}(e) = 0.0635 < 0.1 = \epsilon$$

and so considering only do-interventions is again not sufficient to answer the intervention robustness problem.

It should be noted that while both of these naive approaches are inspired by recent work which performs efficient inference using arithmetic circuits, it is not fair to call either a competing inference approach to our method, as the works cited developed their tools for different objectives. Rather, the purpose of this example is to highlight that many plausible approximation methods to verify the robustness of a classifier can fail on simple examples.

G Additional Circuit Evaluations

G.1 Size of compiled circuits

In Table 5, we show details from joint compilations on five Bayesian networks and decision rules on those networks (further details in Appendix H). First, we notice that the actual size of the compiled ACs is much smaller than that given by the worst case bound (which is exponential in treewidth), due to optimizations in the C2D compiler. Second, notice also that the size of the AC does not increase by 2^d when adding a decision rule to the classification: for instance, in win95pts, the size increases by a factor of ~ 400 , while the decision rule has $2^{16} = 65536$. Interestingly, for the insurance network, the AC size actually *decreases* when compiling with a small decision rule; this is likely due to good fortune with the min-fill heuristic. Finally, when we enforce a topological ordering, the size of the compilation increases, but not

Net	CSize	Ord	TW	AC size	Time (s)
insurance	0 (0)	N	29	362983	1.2
	3 (41)	N	24	167121	0.5
	3 (41)	T	31	794267	4
	3 (41)	S	33	1270075	8
child	0 (0)	N	15	4935	0.007
	8 (326)	N	38	234914	0.7
	8 (326)	T	38	1004786	1
win95pts	0 (0)	N	18	17682	0.04
	16 (799)	N	51	1210072	3
	16 (799)	T	58	52266950	77
hepar2	12 (946)	N	53	8096874	49
	12 (946)	T	51	123108407	73
	12 (946)	S	51	123164181	75
andes	12 (95)	N	41	24787127	272
	12 (95)	PT	43	60865146	778

Table 5: AC sizes and times (s) for the joint compilations used in the UB and LB algorithms. Shown are the number of input features d and the sizes of the Boolean circuits representing the classifier (0 indicates no classifier), ordering constraints (none, partial topological, topological, or structural topological), treewidth of the combined CNF encoding, and size and compilation time. We note that the large increase in network treewidth when adding a classifier is due to the treewidth of the classifier.

by more than ~ 100 . Remarkably, this allows us to upper bound a marginal probability against parametric intervention sets involving any number of intervenable nodes. Our results provide evidence that our methods can scale to fairly large networks and classifiers, including with topological and structural topological orderings.

The ordering types are defined as follows. Recall that, for the correctness of the upper bounding algorithm, we require $V_j < V_i$ for all $V_j \in \text{pa}(V_i)$, and for all $V_i \in \mathbf{W}$. Partial topological orderings impose these constraints $V_j < V_i$. Topological orderings impose the constraints for all $V_i \in \mathbf{V}$, rather than just $V_i \in \mathbf{W}$. This is generally preferred if computationally feasible, as it allows us to compute upper bounds for any parameter intervention set, and further also tends to produce better bounds. However, for the *andes* network, we found that this was too computationally demanding. Structural ordering constraints contain topological constraints, and further for every $V_i \in \mathbf{W}$, we add $V_j < V_i$ for $V_j \in \text{pa}_{G'}(V_i)$, so that we can compute upper bounds on structural intervention sets.

G.2 Tightness of Lower and Upper Bounds

We provide additional results on the tightness of our upper and lower bounds on an expanded set of evaluations, including both false positives and false negatives as evidence, and considering a broader range of networks and intervention sets. Results are detailed in Table 6.

H Experimental Details

In our experiments, we use two types of classifiers F . For the *child*, *win95pts*, *andes* networks, we use Bayesian network classifiers (BNC) trained on the respective networks, in order to predict a root node of the BN. For the *insurance*, *hepar2* network, since the chosen prediction targets are not root nodes of the networks, we used a Naïve Bayes classifier where the conditional probs have been extracted from BN. For *insurance*, *hepar2*, since the prediction targets Y are not root nodes, we can test interventions which change the distribution on Y . Decision functions are obtained from the classifiers by applying a threshold.

To obtain Boolean circuits for the decision functions, we use a BNC-to-ODD compiler [Shih *et al.*, 2019], and then convert the ODD into a Boolean circuit (NNF).

Classifiers

- child*
 - Target: BirthAsphyxia
 - Features: LVHreport, GruntingReport, XrayReport, LowerBodyO2, CardiacMixing, Age, RUQO2, CO2Report
- insurance*
 - Target: MedCost
 - Features: Age, MakeModel, DrivHist
- win95pts*
 - Target: PTROFFLINE
 - Features: Problem3, Problem2, PrtStatMem, PrtStatToner, Problem6, PrtFile, PrtStatOff, PrtIcon, Problem1, REPEAT, HrglssDrtnAftrPrnt, TstpsTxt, PSERRMEM, Problem5, Problem4, PrtStatPaper
- hepar2*
 - Target: Steatosis
 - Features: alt, triglycerides, ggtp, jaundice, alcohol, pain.ruq, cholesterol, ESR, hepatalgia, ast, nausea, fat
- Andes*
 - Target: TRY12
 - Features: TRY15, SNode_14, SNode_19, TRY13, TRY14, GOAL_99, SNode_46, SNode_31, SNode_155, SNode_123, SNode_40, TRY26

Intervention Sets

- child*
 - P1: $\mathbf{W} = \{\text{GruntingReport}\}$
 - P2: $\mathbf{W} = \{\text{ChestXray, Sick, Grunting}\}$
 - P3: $\mathbf{W} = \{\text{LowerBodyO2, RUQO2, CO2Report}\}$
- insurance*

Network	IntSet	False Negatives				False Positives			
		BeforeIntv	LBound	RBound	Δ	BeforeIntv	LBound	UBound	Δ
child	P1	0.06922	0.07098	0.07098	0	0.1629	0.1947	0.1947	0
	P2	0.06922	0.07325	0.07329	0.00004	0.1629	0.2762	0.3069	0.0307
	P3	0.06922	0.06978	0.07127	0.00149	0.1629	0.1717	0.2009	0.0292
insurance	P1	0.02453	0.1181	0.1276	0.0095	0.1981	0.4157	0.4161	0.0004
	P2	0.02453	0.3275	0.3433	0.0158	0.1981	0.9123	0.9130	0.0007
	P3	0.02453	0.02453	0.02453	0	0.1981	0.1981	0.1981	0
	S1	0.02453	0.1181	0.1297	0.0116	0.1981	0.4157	0.4168	0.0011
win95pts	P1	0.2106	0.2111	0.2111	0	0.005170	0.005416	0.005445	0.000029
	P2	0.2106	0.2163	0.2191	0.0028	0.005170	0.007200	0.008665	0.001465
	P3	0.2106	0.2972	0.2985	0.0013	0.005170	0.01430	0.01445	0.00015
	P4	0.2106	0.2109	0.2117	0.0008	0.005170	0.05494	0.05674	0.00180
hepar2	P1	0.03673	0.09445	0.09445	0	0.2360	0.2408	0.2408	0
	P2	0.03673	0.09585	0.09585	0	0.2360	0.9041	0.9041	0
	P3	0.03673	0.1029	0.1029	0	0.2360	0.43758	0.43773	0.00015
	S1	0.03673	0.1029	0.1029	0	0.2360	0.43758	0.43793	0.00035
andes	P1	0.001400	0.001400	0.002540	0.001140	0	0	0	0
	P2	0.001400	0.001400	0.002656	0.001256	0	0	0	0

Table 6: Analysis of the tightness of bounds (on probability of false negatives/false positives) produced by the UB and LB algorithms. For each network, we have different intervention sets. We show the probability of false negatives/false positives in the original Bayesian network (BeforeIntv), along with lower and upper bounds under each intervention set.

- P1: $W = \{\text{MakeModel}, \text{Cushioning}\}$
- P2: $W = \{\text{SocioEcon}, \text{RiskAversion}, \text{Theft}, \text{Mileage}, \text{MakeModel}, \text{Cushioning}\}$
- P3: $W = \{\text{ThisCarDam}, \text{AntiTheft}, \text{OtherCarCost}\}$
- C1: $W = \{\text{MakeModel}, \text{Cushioning}\}$,
 $C(\text{MakeModel}) = \{\text{Age}, \text{AntiTheft}, \text{DrivHist}, \text{DrivingSkill}, \text{GoodStudent}, \text{HomeBase}, \text{Mileage}, \text{OtherCar}, \text{RiskAversion}, \text{SeniorTrain}, \text{SocioEcon}, \text{VehicleYear}, \}$
 $C(\text{MakeModel}) = \{\text{Age}, \text{Airbag}, \text{AntiTheft}, \text{Antilock}, \text{CarValue}, \text{DrivHist}, \text{DrivQuality}, \text{DrivingSkill}, \text{GoodStudent}, \text{HomeBase}, \text{MakeModel}, \text{Mileage}, \text{OtherCar}, \text{RiskAversion}, \text{RuggedAuto}, \text{SeniorTrain}, \text{SocioEcon}, \text{Theft}, \text{VehicleYear}, \}$
- 3. win95pts
 - P1: $W = \{\text{NetOK}, \text{NetPrint}\}$
 - P2: $W = \{\text{AvlblVrtlMmry}, \text{DSApplctn}, \text{DskLocal}, \text{HrglssDrtnAftPrnt}, \text{NtSpd}, \text{DeskPrntSpd}, \text{EPSGrphc}, \text{PSGRAPHIC}, \text{Fl1CrrptdBffr}\}$
 - P3: $W = \{\text{REPEAT}\}$
- P4: $W = \{\text{GDIIN}, \text{PC2PRT}, \text{PSGRAPHIC}, \text{DS_LCLOK}, \text{PSERRMEM}, \text{EMFOK}, \text{DS_NTOK}\}$
- 4. hepar2
 - P1: $W = \{\text{alcoholism}\}$
 - P2: $W = \{\text{alcoholism}\}$
 - P3: $W = \{\text{alcoholism}, \text{hepatomegaly}, \text{alcohol}, \text{itching}, \text{fatigue}, \text{consciousness}, \text{hospital}\}$
 - C1: $W = \{\text{alcoholism}, \text{hepatomegaly}, \text{alcohol}, \text{itching}, \text{fatigue}, \text{consciousness}, \text{hospital}\}$
The context function C specifies additional parents on top of those present in the original BN. All W have age, sex, alcoholism as additional parents. Further, fatigue, consciousness, hospital have anorexia as additional parent.
- 5. andes
 - P1: $W = \{\text{GOAL_49}, \text{GOAL_61}, \text{SNode_26}, \text{SNode_37}\}$
 - P2: $W = \{\text{GOAL_49}, \text{GOAL_61}, \text{SNode_26}, \text{SNode_37}, \text{GOAL_57}, \text{GOAL_149}, \text{GOAL_153}, \text{SNode_74}\}$