Symbolic Causal Networks for Reasoning about Actions and Plans

Adnan Darwiche*
Rockwell Science Center
444 High Street
Palo Alto, CA 94301
darwiche@rpal.rockwell.com

Judea Pearl
Computer Science Department
University of California
Los Angeles, CA 90024
pearl@cs.ucla.edu

Abstract

We present an approach for reasoning about actions and plans when domain knowledge is represented by a symbolic causal network, which is a principled, logical representation of a domain that explicates its perceived causal structure. The proposed approach shows that causal structures can play a key role in logical reasoning about actions given their effective role in dealing with some of the problems associated with such reasoning, including the frame, ramification, and concurrency problems.

1 Introduction

A symbolic causal network is a principled representation of a logical database (a set of propositional clauses), which encodes one's perception of causal relationships in a given domain; see Figure 1.

In the same way that a probabilistic causal network represents a probability distribution that is faithful to a given causal structure [14], a symbolic causal network represents a logical database that satisfies similar faithfulness conditions [5].

Causal faithfulness stands for two requirements, one concerns the dynamics of database *revisions* due to new observations, while the second concerns the dynamics of database *updates* due to external actions.

In probabilistic causal networks, revisional faithfulness is encapsulated in conditional independence constraints, the satisfaction of which is guaranteed whenever the distribution is generated by processes configured according to the network's layout [14]. Symbolic causal networks offer similar guarantees with respect to logical databases, but the independence constraints encoded by these networks are logical rather than probabilistic.

Revisional faithfulness and how it can be obtained using symbolic causal networks are treated elsewhere [5]. In this paper, we focus on the action part of faithfulness, ensuring that actions, their effects, their interactions with observations, and their interactions with

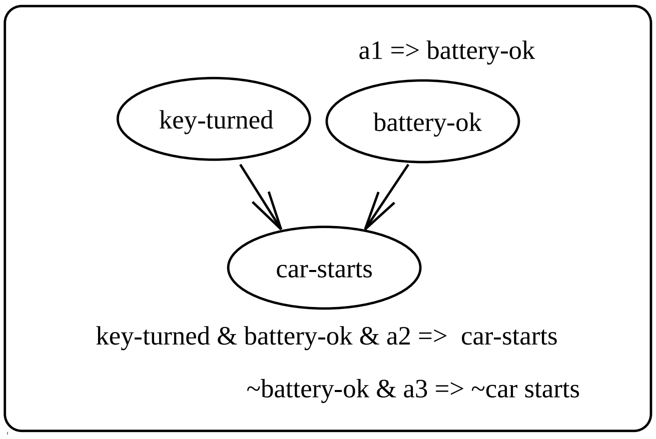

Figure 1: A symbolic causal network.

other actions are consistent with the (perceived) causal structure of the world. Specifically, we will show that symbolic causal networks define a simple update operator on logical databases that meets this faithfulness requirement. Moreover, we will show how this operator can serve as the basis for reasoning about actions and plans when knowledge is encoded using symbolic causal networks.

This paper is structured as follows. Section 2 reviews symbolic causal networks. Section 3 discusses the update operator defined by a symbolic causal network and shows how it can be used to reason about actions. This section also discusses the role of causal structures in dealing with the frame, ramification, and concurrency problems. Section 4 shows how symbolic causal networks can be used to reason about plans (sequences of actions), in addition to illustrating their ability to reason about a mixture of actions and observations and to support abductive as well as predictive reasoning. Finally, Section 5 shows how symbolic causal networks support assumption–based reasoning and discusses the role of this mode of reasoning in dealing with uncertainty.

2 Symbolic causal networks

Consider the networks depicted in Figure 1 and Fig-

*This work has been partially supported by ARPA contract F30602–91–C–0031.

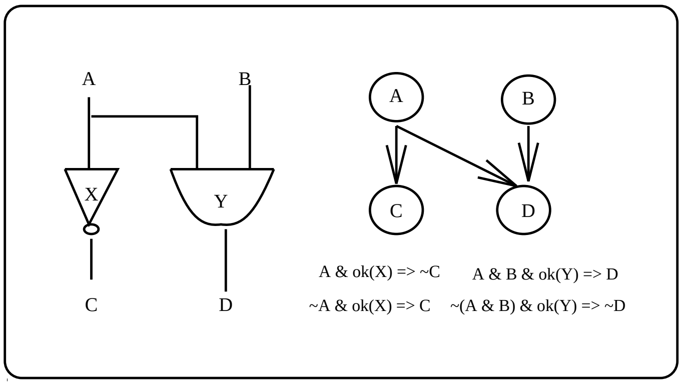

Figure 2: A symbolic causal network describing a circuit.

ure 2. Each of these networks has two components. The first is a *causal structure* that captures perceptions of causal influences. And the second is a set of *micro theories*, each associated with one proposition in the causal structure.

The purpose of each micro theory is to specify the logical relationship between a proposition and its direct causes. For example, the micro theory associated with the Proposition *car_starts* in Figure 1 is

$$\begin{aligned} key\_turned \wedge battery\_ok \wedge a_2 &\supset car\_starts, \\ \neg battery\_ok \wedge a_3 &\supset \neg car\_starts, \end{aligned}$$

which specifies the relationship between this proposition and its direct causes *key_turned* and *battery_ok*. Similarly, the micro theory associated with Proposition D in Figure 2 is

$$\begin{aligned} A \wedge B \wedge OK(Y) &\supset D \\ \neg(A \wedge B) \wedge OK(Y) &\supset \neg D, \end{aligned}$$

which specifies the relationship between Proposition D and its direct causes A and B.

In general, the micro theory associated with a proposition p has two types of material implications:

- *positive causal rules* of the form $\psi \wedge \alpha \supset p$ and
- *negative causal rules* of the form $\phi \wedge \beta \supset \neg p$,

where

1. ψ and ϕ are propositional sentences constructed from the direct causes of p in the causal structure.
2. α and β are propositional sentences constructed from atomic propositions that do not appear in the causal structure (called *assumption symbols*).
3. $\alpha \wedge \beta$ is unsatisfiable whenever $\psi \wedge \phi$ is satisfiable.

For example, in Figure 1, the material implication,

$$\neg battery\_ok \wedge a_3 \supset \neg car\_starts,$$

is a negative causal rule, where $\psi = \neg battery\_ok$ and $\alpha = a_3$. Similarly, in Figure 2, the material implication,

$$A \wedge B \wedge OK(Y) \supset D,$$

is a positive causal rule, where $\phi = A \wedge B$ and $\beta = OK(Y)$ — that is, $OK(Y)$ is an assumption symbol in this case.

The first two conditions above ensure that each micro theory is local to a specific proposition and its direct causes. The last condition is typically self–imposed in causal modeling. In particular, the causal rules $\psi \wedge \alpha \supset p$ and $\phi \wedge \beta \supset \neg p$ entail $\alpha \wedge \beta \supset \neg(\psi \wedge \phi)$. Therefore, if $\alpha \wedge \beta$ and $\psi \wedge \phi$ are both satisfiable, then the micro theory for p — which is intended to specify the relationship between p and its direct causes — is indirectly specifying a relationship between the direct causes of p, which is atypical in causal modeling. For example, one would never specify a relationship between the inputs to a digital gate in the process of specifying the relationship between its inputs and output.

3 Reasoning about action

We observed elsewhere that causal structures impose independence constraints on belief changes that are triggered by observations (belief revisions) [5]. In particular, we characterized the conditional independences imposed by a causal structure on belief revisions:

> Given a state of the assumption symbols, observing the direct causes of a proposition p renders the belief in p independent of observations about its non–effects.

We also showed that the database induced by a symbolic causal network satisfies all the independences encoded by its corresponding causal structure.

But causal structures also impose constraints on belief changes that are triggered by external interventions (belief updates [9]). Therefore, our focus in this paper is on characterizing these constraints and on providing a formal proposal for belief update that respects them.

To motivate the discussion in this section, consider the symbolic causal network in Figure 2. Suppose that gates X and Y are both ok, and that we have no information about the states of wires A, B, C and D. Suppose further that someone intervenes and sets output C to ON by connecting it to a high voltage. How can we predict the effect of this action formally?

Before we answer this question, two points need to be stressed:

1. We cannot account for this action by simply adding C to the database describing the circuit. If we do this, we end up concluding that input A must be OFF and, therefore, that output D must also be OFF. But this is counterintuitive since connecting C to a high voltage should not change our beliefs about input A and output D. This should not be surprising given the recent literature on belief update, which emphasizes the distinction between recording an observation about a static world (observing that C is ON) and recording an observation about a changing world (intervening to make C ON) [9, 8].
2. The propositional database corresponding to the digital circuit does not contain enough information to predict the effect of an action that sets output C to ON. To see why, consider the circuit in Figure 3,

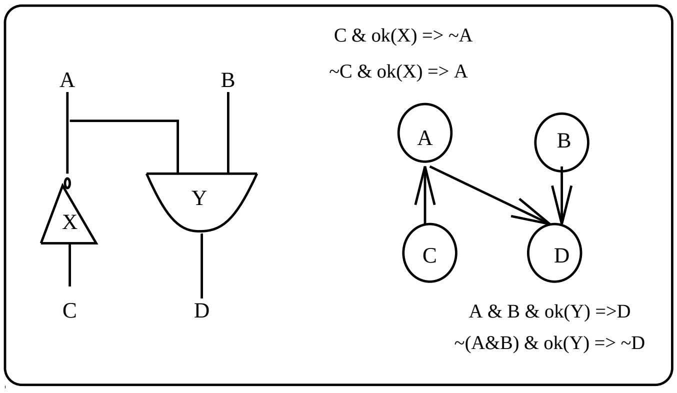

Figure 3: A symbolic causal network describing a circuit.

> which is exactly like the circuit in Figure 2, except that the input and output of inverter X are interchanged. The two circuits have *the same logical description* as can be seen from their corresponding causal networks. However, if we connect C in Figure 3 to a high voltage, then we *would* conclude that A and D are OFF.

What we have here is two circuits with equivalent logical relationships among their wires, but with different reactions to external interventions. The question now is, What extra information about these circuits should we appeal to in order to infer formally their reactions to external interventions?

What we will show next is that the causal structures of these circuits is precisely that extra information. In particular, we will show that although the circuits' logical descriptions do not contain enough information to predict their reactions to setting C to ON, their logical descriptions together with the causal structures do. Moreover, this combined information is exactly what is captured by a symbolic causal network.

We first show in general how the causal structure of a symbolic causal network can be used to update the database specified by the network in response to an action. We then show how the suggested approach works on the previous example.

The update operator is based on the following principle, called *the sufficient cause principle,* which reduces actions to observations:

> Acting to establish the truth of proposition p is equivalent to observing a hypothetical event $Do(p)$, called a sufficient cause of p, which is (1) a direct cause of p, (2) logically entails p and (3) is independent of every non–effect (non–descendant) of p in the causal structure.

According to this principle, predicting the effect of an action that establishes proposition p can be accomplished by performing the following steps. First, we augment the given causal structure with a sufficient cause $Do(p)$ of p, thus changing the (causal) micro theory of p. Second, we simulate the action of establishing p by an observation of the sufficient cause $Do(p)$.

Specifically, after extending the causal structure by adding the sufficient cause $Do(p)$, the micro theory of proposition p,

$$\begin{aligned} \psi \wedge \alpha &\supset p \\ \phi \wedge \beta &\supset \neg p, \end{aligned}$$

is replaced by the following:

$$\begin{aligned} \neg Do(p) \wedge \psi \wedge \alpha &\supset p \\ \neg Do(p) \wedge \phi \wedge \beta &\supset \neg p \\ Do(p) &\supset p. \end{aligned}$$

One must stress the following about the above proposal:

1. The alteration of the micro theory of proposition p is nonmonotonic. For example, the original micro theory of p entails $\psi \wedge \alpha \supset p$, but the new micro theory does not.
2. We are assuming that the action $Do(p)$ has no preconditions.
3. We are assuming that the action $Do(p)$ always succeeds in obtaining its effect of establishing p.[1]
4. The direct effect of the action $Do(p)$ is restricted to its effect on p; all its other effects are logical consequences of its effect on p.[2]

Since the above proposal reduces actions to observations, the results reported in [5] for reasoning about observations become available for reasoning about actions. Among the most important of these results are (1) a characterization of the independences imposed by a causal structure on beliefs, observations and actions; (2) the ability to read these independences and many of their implications directly from the topology of a causal structure using the criterion of d–separation [14]; (3) a proposal for reasoning about actions and observations that is guaranteed to satisfy these independences; and (4) a set of distributed algorithms for computing inferences that are symmetric in their complexity to the algorithms used in probabilistic causal networks.

3.1 An example

Let us see how we can formally predict the effect of setting C to ON in both circuits by appealing to our account of action.

Starting with the circuit in Figure 2, we create a sufficient cause $Do(C)$ for proposition C and update its micro theory, which leads to the causal network in Figure 4 and its corresponding database Γ. We then simulate the action of setting C to ON by observing the sufficient cause $Do(C)$, which gives the intended results:

$$\Gamma \cup \{OK(X), OK(Y), Do(C)\} \models C,$$
$$\Gamma \cup \{OK(X), OK(Y), Do(C)\} \not\models \neg A \vee \neg D.$$

That is, our beliefs about A and D do not change as a result taking an action that sets C to ON.

[1] This is not a limitation, but a simplification. Actions with uncertain effects can be modeled in the same framework, but are outside the scope of this paper.

[2] We note here that actions with similar properties have been treated in probabilistic settings using a probabilistic analogue of the sufficient–cause principle [16, 8].

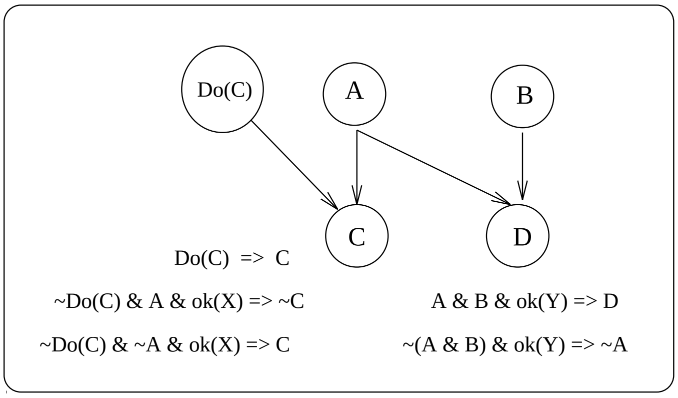

Figure 4: A symbolic causal network extending the one in Figure 2 by including a sufficient cause $Do(C)$ for proposition C. Adding this cause changes the micro theory associated with C only. The rest of the network remains unchanged.

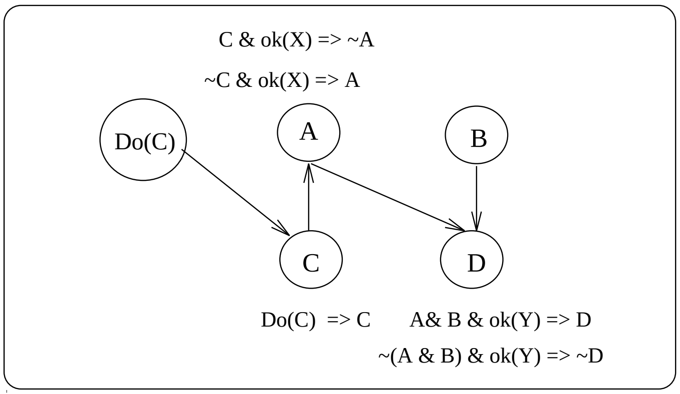

Figure 5: A symbolic causal network extending the one in Figure 3 by including a sufficient cause $Do(C)$ for proposition C. Adding this cause changes the micro theory associated with C only. The rest of the network remains unchanged.

If we perform the same exercise with respect to the circuit in Figure 3, we obtain the symbolic causal network in Figure 5, which specifies a different database Λ. We then simulate the action of setting C to ON by observing the sufficient cause $Do(C)$, which also gives the intended results:

$$\Lambda \cup \{OK(X), OK(Y), Do(C)\} \models C \wedge \neg A \wedge \neg D.$$

That is, our beliefs about A and D change as a result of setting C to ON. Note, however, that adding $Do(C)$ does not perturb the micro theories of either A or D.

The key idea underlying the approach we presented so far is that causal knowledge can efficiently be organized in terms of small mechanisms (described using micro theories), each involving a relatively small number of propositions. Each external intervention overrules just one mechanism, leaving the others intact. The specification of an action then requires only the identification of the mechanism that it overrules. Once this is identified, the effect of an action can then be computed from the constraints imposed by the remaining mechanisms. The simplicity and effectiveness of this approach manifests itself clearly when discussing the way it treats some of the major difficulties in reasoning about actions; that is, the frame, ramification, and concurrency problems.

3.2 Causal structures and the frame problem

A *frame axiom* is a statement identifying an aspect of the world that is not changed by a certain action. For example, "Moving block A on the table does not change its color" is a frame axiom. The *frame problem* is that of succinctly summarizing the frame axioms [12]. A number of proposals for such summarization are discussed in [15].

Summarizing frame axioms is one of the key roles played by the causal structure of a symbolic causal network. In particular, if the effect of an action is predicted according to the sufficient–cause principle, then the following property is guaranteed: An action $Do(p)$ will never change the truth value of a proposition q that is not an effect (descendant) of p in the causal structure. Therefore, by using the sufficient–cause principle to reason about actions, one is implicitly respecting the following frame axioms, which can be read from the topology of a causal structure:

> For each proposition p, and for each proposition q that is not an effect of p in the causal structure, the action $Do(p)$ does not change the truth value of q.

This follows because the non–effects of p in a causal structure are d–separated from $Do(p)$.

For example, the causal structure in Figure 2 represents eighteen frame axioms, and the causal structure in Figure 3 represents sixteen frame axioms. The axiom, "$Do(C)$ does not change the truth value of D" is encoded by the causal structure of Figure 2, but not by the one of Figure 3. On the other hand, the axiom, "$Do(A)$ does not change the truth value of C" is encoded by the causal structure of Figure 3, but not by the one of Figure 2. It is this extra information that a symbolic causal network captures about a system, in addition to the logical description of that system.

3.3 Micro theories and the ramification problem

The ramification problem is the difficulty of describing the indirect effects of actions. This problem is most common in formalisms that solve the frame problem by deriving frame axioms from the completeness assumption of effect axioms [15]. In such a case, prohibiting domain constraints (e.g., the workings of a gate) seems to simplify the derivation of frame axioms because it allows such derivation through only local considerations of effect axioms. The result of this restriction, however, is the inability to deduce indirect effects of actions, which puts the burden on the user to enumerate them in effect axioms.

Domain constraints are represented in symbolic causal networks using the micro theories associated

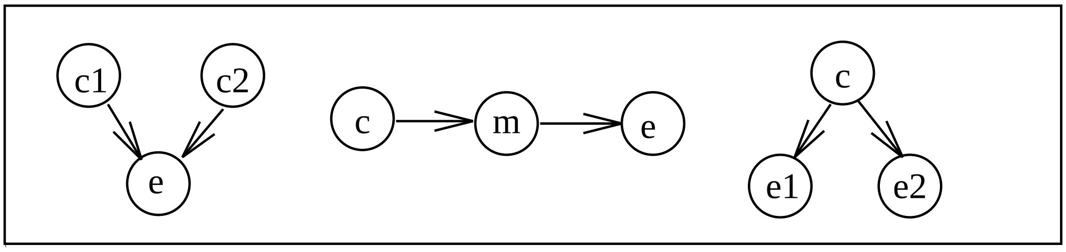

Figure 6: Possible causal interactions.

with network propositions. These constraints are used to infer the indirect effect of actions. For example, in Figure 3, the action $Do(C)$ has the direct effect of setting C to ON, but it also has an indirect effect of making A OFF. But this indirect effect can be inferred using the micro theory of proposition A.

In summary, the ramification problem does not appear in symbolic causal networks because micro theories are domain constraints that are used to infer indirect effects of actions. Moreover, allowing domain constraints in symbolic causal networks does not complicate the frame problem because frame axioms are not derived from effect axioms, but are inferred from the causal structure instead.

3.4 d–separation and concurrent actions

The approach we presented so far does not require a special treatment of concurrent actions [1, 11]. That is, to predict the effect of a set of actions, one needs only to identify the mechanisms that they overrule and then compute the effects of such actions using the constraints imposed by the remaining mechanisms.

More importantly, the causal structure of a symbolic causal network can be used to predict some interactions among concurrent actions without having to specify the involved mechanisms. This should not be surprising, however, given that a causal structure outlines the interactions among these mechanisms.

In particular, the criterion of d–separation tells us that for any propositions p and q in a causal structure, the actions $Do(p)$ and $Do(q)$ are logically independent given a state of assumption symbols. This follows because $Do(p)$ and $Do(q)$ are always d–separated in any causal structure. This property, however, does not necessarily hold in case either p, q, or one of their common descendants is observed. For example, given the state $OK(X) \wedge OK(Y)$ of assumption symbols in Figure 3, observing $\neg D$ makes the actions $Do(\neg C)$ and $Do(B)$ logically dependent; in particular, it makes $Do(\neg C) \wedge Do(B)$ inconsistent. Note, however, that this dependence is only in light of a given state of assumption symbols. If we do not commit to such a state, the combined actions $Do(\neg C) \wedge Do(B)$, together with the observation $\neg D$, logically imply $\neg OK(X) \vee \neg OK(Y)$, thus ruling out the state $OK(X) \wedge OK(Y)$ from being possible.

The criterion of d–separation can also be informative about the interactions between the effects of actions.

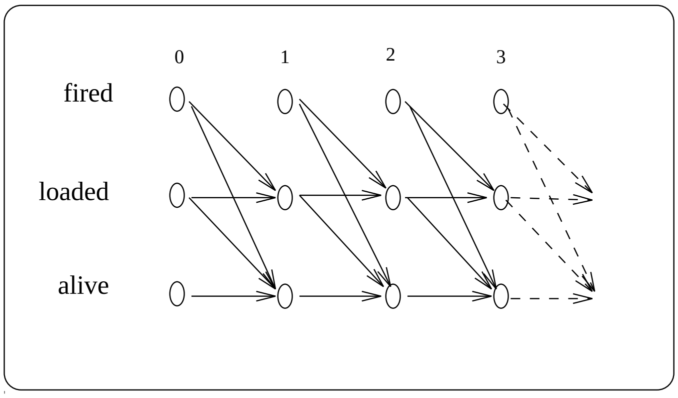

Figure 7: A causal network for reasoning about plans.

In particular, from d–separation we can infer the following principles, which refer to Figure 6. First, if c is an indirect cause of e, and if there is a conflict between the effect of $Do(c)$ and the effect of $Do(e)$, then the effect of $Do(e)$ prevails. Next, if c_1 and c_2 are common causes of e, then the actions $Do(c_1)$ and $Do(c_2)$ can never have a conflicting effect on e. This follows from the definition of a micro theory, which guarantees that c_1 and c_2 are logically independent given a state of assumption symbols. Finally, if e_1 and e_2 are common effects of c, then the cause c is logically independent of the actions $Do(e_1)$ and $Do(e_2)$ given a state of assumption symbols. That is, neither of these actions will have an influence on c.

The reader is referred to [13] for more details on the use of d–separation in making predictions about the effects of actions.

4 Reasoning about plans

Reasoning about plans (sequences of actions) requires one to explicate the temporal order in which actions take place. To that end, symbolic causal networks that support such reasoning can use propositions that are indexed by the time at which they hold.

For example, if one wants to plan in a shooting scenario, then one may construct a symbolic causal network with the causal structure in Figure 7 and the following micro theories:

- The micro theories of $loaded_0$, $alive_0$ and $fired_t$ are empty.

- The micro theory of $loaded_t$, where $t > 0$, is

$$\begin{aligned} fired_{t-1} &\supset \neg loaded_t \\ \neg fired_{t-1} \wedge loaded_{t-1} &\supset loaded_t \\ \neg loaded_{t-1} &\supset \neg loaded_t, \end{aligned}$$

- The micro theory of $alive_t$, where $t > 0$, is

$$\begin{aligned} fired_{t-1} \wedge loaded_{t-1} &\supset \neg alive_t \\ \neg(fired_{t-1} \wedge loaded_{t-1}) \wedge alive_{t-1} &\supset alive_t \\ \neg alive_{t-1} &\supset \neg alive_t. \end{aligned}$$

This completes the definition of the symbolic causal network, which can now be used to answer queries.

We will now consider two examples for reasoning about plans, one illustrating the ability to reason about a plan in light of some observations and the other illustrating the ability to support abductive reasoning.

4.1 Mixing actions and observations

Suppose that the user observes $alive_0 \wedge \neg loaded_0$, decides to load the gun without firing it, $Do(loaded_1) \wedge \neg fired_1$, and then fires it, $Do(fired_2)$. What is the effect of this sequence of actions on $loaded_t$ and $alive_t$?

To answer this query according to the proposal given in Section 3, we change the micro theory of $loaded_1$ from

$$\begin{aligned} fired_0 &\supset \neg loaded_1 \\ \neg fired_0 \wedge loaded_0 &\supset loaded_1 \\ \neg loaded_0 &\supset \neg loaded_1. \end{aligned}$$

to:[3]

$$\begin{aligned} \neg Do(loaded_1) \wedge fired_0 &\supset \neg loaded_1 \\ \neg Do(loaded_1) \wedge \neg fired_0 \wedge loaded_0 &\supset loaded_1 \\ \neg Do(loaded_1) \wedge \neg loaded_0 &\supset \neg loaded_1 \\ Do(loaded_1) &\supset loaded_1. \end{aligned}$$

We also replace the micro theory of $fired_2$, which is empty in this case, with

$$Do(fired_2) \supset fired_2.$$

Now, assuming that the above substitutions lead to database Δ, it is easy to verify that

$$\Delta \cup \{alive_0, \neg loaded_0, Do(loaded_1), \neg fired_1, Do(fired_2)\}$$

implies $loaded_t \wedge alive_t$ for $t = 1, 2$, and that it implies $\neg loaded_t \wedge \neg alive_t$ for $t \geq 3$. That is, the victim is dead and the gun is unloaded after the action is taken and that persists into the future.

4.2 Abductive reasoning

Suppose that the user observes the victim to be alive $alive_0$, decides to fire the gun $Do(fired_1)$, and then finds out that the victim is still alive $alive_2$. How can we explain this observation?

To answer this query according to the proposal given in Section 3, we change the micro theory of $fired_1$, which is empty in this case, with

$$Do(fired_1) \supset fired_1.$$

Now, assuming that the above substitutions lead to database Δ, it is easy to verify that

$$\Delta \cup \{alive_0, Do(fired_1), alive_2\}$$

implies $\neg loaded_1$. That is, the reason why shooting did not kill the victim is explained by the gun being unloaded during the shooting.

5 Assumption–based reasoning

Reasoning about actions and plans is typically done under uncertainty: Actions may not always succeed, and when they do, their effects may not be certain. Even if the direct effects of actions are certain, their indirect effects are often conditional on uncertain propositions.

In probabilistic reasoning about actions, this problem is dealt with by attaching probabilities to propositions. In logical reasoning, uncertainty is typically dealt with using nonmonotonic reasoning. There are many proposals for nonmonotonic reasoning. One of these proposals, called assumption–based reasoning, identifies a set of assumable propositions, assumes their truth values, and then retracts or reverses these assumptions when they prove to be wrong. An ATMS is the basic formalism for implementing this sort of reasoning [6]. In ATMSs, a *label* is attached to each proposition, which characterizes all assumptions under which the proposition holds.

Symbolic causal networks support assumption–based reasoning. In particular, in the same way that probabilistic causal networks compute a probability for each proposition in the network, symbolic causal networks compute an *argument* for each proposition. Arguments are logically equivalent to ATMS labels, but they are not necessarily put in canonical form [3].[4] Consider the network in Figure 3 for an example. Initially, the argument for any proposition is simply *false*, meaning that not any set of assumptions would be enough to entail any proposition. After observing D, however, the arguments for A, B, and $\neg C$ are updated to $OK(Y)$, $OK(Y)$, and $OK(X) \wedge OK(Y)$, respectively. But if we set C to ON, the argument for $\neg A$ is updated to $OK(X)$.

In assumption–based reasoning, one assumes a particular state of assumption symbols, thus leading to some state of belief. But this state is then changed in face of observations that contradict with it. The role that arguments play in this mode of reasoning is two–fold. First, they are needed to decide whether a proposition holds given some assumptions. Second, they are needed to characterize the assumptions that are logically possible after recording some observations. Specifically, proposition p follows from some assumptions precisely when these assumptions entail the argument for

[3] Note that the transformation is nonmonotonic: the original database entails $\neg loaded_0 \supset \neg loaded_1$ but the new database does not entail this.

[4] Therefore, arguments are easier to compute than ATMS labels; the complexity of computing arguments in symbolic causal networks is symmetric to the complexity of computing probabilities in probabilistic causal networks [2, 3].

p. Moreover, the assumptions that are logically possible given the observation O are those that do not entail the argument for $\neg O$ [3].

Another important role of arguments in reasoning under uncertainty is in implementing Dempster–Shafer reasoning. The basic idea here is to assign probabilities to assumption symbols while assuming their probabilistic independence. The probability of the argument for proposition p can then be shown to correspond to the Dempster–Shafer belief in p [10].

We have discussed earlier the role that causal structures play in treating the frame, ramification, and concurrency problems. These structures also play a significant computational role that will be elaborated on in the remainder of this section. Specifically, since a causal structure encodes independences that can be read from its topology, and since the database induced by a symbolic causal network is guaranteed to satisfy these independences, the causal structure of a symbolic causal network can be exploited by distributed algorithms when computing arguments. That is, the topology of a causal structure can be used to guide the decomposition of arguments into smaller arguments that can be computed in parallel [2, 5, 3]. More precisely, whenever propositions X and Y are d–separated by Z in the causal structure [14], we get the following key property:

$$Argument(\psi_{XYZ}) \equiv Argument(\psi_{XZ}) \vee Argument(\psi_{YZ}),$$

where ψ_{XYZ} is a clause over propositions $X \cup Y \cup Z$, while ψ_{XZ} and ψ_{YZ} are the subsets of ψ_{XYZ} over the propositions $X \cup Z$ and $Y \cup Z$, respectively.[5] This decomposition is the basis for distributed algorithms that compute arguments and is analogous to the decomposition used when computing probabilities in probabilistic causal networks [3]. In fact, symbolic causal networks are only an instance of a more general class of networks, called *abstract causal networks* [2], which also include probabilistic causal networks [14] and kappa causal networks [7]. The computational utility of the independences encoded by a causal structure is common to all instances of abstract causal networks as shown in [2], which also presents some formal algorithms and a Lisp implementation, called CNETS [4], for reasoning with abstract causal networks.

Conclusion

The basic contribution of this paper has been a proposal for updating propositional knowledge bases that are represented using symbolic causal networks. The proposal guarantees the faithfulness of belief updates to a given causal structure in a precise sense. It also shows that causal structures can play a key role in logical reasoning about actions given their effective role in dealing with some of the problems associated with such reasoning, including the frame, ramification, and concurrency problems.

[5] This is similar to the guarantee that $Pr(\mathbf{X},\mathbf{Y}|\mathbf{Z}) = Pr(\mathbf{X}|\mathbf{Z})Pr(\mathbf{Y}|\mathbf{Z})$ whenever X and Y are d–separated by Z in a probabilistic causal network.

References

[1] Chitta Baral and Michael Gelfond. Representing concurrent actions in extended logic programs. In *Proceedings of International Joint Conference on Artifical Intelligence (IJCAI)*, pages 866–871, 1993.

[2] Adnan Darwiche. *A Symbolic Generalization of Probability Theory*. PhD thesis, Stanford University, 1992.

[3] Adnan Darwiche. Argument calculus and networks. In *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 420–427, 1993.

[4] Adnan Darwiche. CNETS: A computational environment for generalized causal networks. Technical memorandum, Rockwell International, Palo Alto Laboratory, 1994.

[5] Adnan Darwiche and Judea Pearl. Symbolic causal networks. To appear in the Proceedings of the Twelfth National Conference on Artificial Intelligence, 1994.

[6] Johan de Kleer. An assumption-based TMS. *Artificial Intelligence*, 28:127–162, 1986.

[7] Moisés Goldszmidt and Adnan Darwiche. Action networks: A framework for reasoning about actions and change under uncertainty. In this volume.

[8] Moisés Goldszmidt and Judea Pearl. Rank–based systems: A simple approach to belief revision, belief update and reasoning about evidence and actions. In *Proceedings of the Third Conference on Principles of Knowledge Representation and Reasoning*, pages 661–672. Morgan Kaufmann Publishers, Inc., San Mateo, California, 1992.

[9] H. Katsuno and A. Mendelzon. On the difference between updating a knowledge base and revising it. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference*, pages 387–394, Boston, 1991.

[10] K. B. Laskey and P. E. Lehner. Assumptions, beliefs, and probabilities. *Artificial Intelligence*, 41(1):65–77, 1989.

[11] Fangzhen Lin and Yoav Shoham. Concurrent actions in the situation calculus. In *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI)*, pages 590–595, 1992.

[12] John McCarthy and Patrick Hayes. Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence*, 4:463–502, 1969.

[13] Judea Pearl. From imaging and stochastic control to a calculus of action. In this volume.

[14] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Inc., San Mateo, California, 1988.

[15] Ray Reiter. The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In Vladimir Lifschitz, editor, *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, pages 359–380. Academic Press, San Diego, Ca, 1991.

[16] P. Spirtes, C. Glymour, and R. Schienes. *Causation, Prediction, and Search*. Springer–Verlag, New York, 1993.