

Induction in Nonmonotonic Causal Theories for a Domestic Service Robot

Jianmin Ji¹ and Xiaoping Chen²

¹ Dept. of CSE, The Hong Kong University of Science and Technology,
jizheng@mail.ustc.edu.cn

² Sch. of CS, University of Science and Technology of China,
xpchen@ustc.edu.cn

Abstract. It is always possible to encounter an expected scenario which has not been covered by a certain theory for an action domain. This paper proposes an approach to treating this problem. We reduce this learning task into the problem of modifying a causal theory such that the interpretations corresponding to new scenarios become a model of the updated theory, while all the original models keep unchanged. We illustrate our approach through a case study based on a domestic service robot, KeJia.

1 Introduction

Nonmonotonic causal theories [5] are devoted to be a nonmonotonic formalism for representing causal knowledge, which can be used to formalize action domains, including indirect effects of actions, implied action preconditions, concurrent actions, nondeterministic actions, ramification and qualification constraints. However, in practice, it is hard to model a complicated action domain completely and there always exist unexpected scenarios which are not covered by a specific theory. An important problem for developing an intelligent agent is how to automatically modify a causal theory for an action domain so that the updated theory can cover a new scenario once it is encountered during the agent's exploration of the real world.

In this paper, we attack the inductive learning problem in nonmonotonic causal theories. We take domestic service robots as the setting, though the results apply to other applications. Our robot is initially equipped with a causal theory as the model of its environment, with which it provides services for its users. In order to get more complete causal theories, we assume the robot get one or more new scenarios through robot learning from demonstration [1], for each of which the robot recognizes and remembers the entire history of its behaviors and relevant fluents of the environment. Formally, this history is an interpretation of the causal theory, but not a model of it. Therefore, we reduce the learning task into the problem of modifying the causal theory such that this interpretation becomes a model of the updated theory, while all the models of original theory are still models of the updated theory. We illustrate our approach through a case study based on a domestic service robot, KeJia [3, 2]. Only the core part of this work will be described in this short paper.

2 Nonmonotonic Causal Theories

The language of nonmonotonic causal theories [5] is based on a propositional language with two zero-place logical connectives \top for tautology and \perp for contradiction. We denote by *Atom* the set of atoms, and *Lit* the set of literals: $Lit = Atom \cup \{\neg a \mid a \in Atom\}$. Given a literal l , the *complement* of l , denoted by \bar{l} , is $\neg a$ if l is a and a if l is $\neg a$, where a is an atom. A set I of literals is called *complete* if for each atom a , exactly one of $\{a, \neg a\}$ is in I . In this paper we identify an interpretation with a complete set of literals.

A *causal theory* is a set of *casual rules* of the form: $\phi \Rightarrow \psi$, where ϕ and ψ are propositional formulas. For a causal rule r of such form, we let $head(r)$ be its head ψ and $body(r)$ its body ϕ . Intuitively, the causal rule reads as “ ψ is caused if ϕ is true”.

Let T be a causal theory and I an interpretation. The *reduction* T^I of T w.r.t. I is defined as $T^I = \{\psi \mid \text{for some } \phi \Rightarrow \psi \in T \text{ and } I \models \phi\}$. T^I is a propositional theory. We say that I is a *model* of T if I is the unique model of T^I .

For example, given $T_1 = \{p \Rightarrow p, q \Rightarrow q, \neg q \Rightarrow \neg q\}$ whose signature is $\{p, q\}$. Let $I_1 = \{p, q\}$, $T_1^{I_1} = \{p, q\}$ and I_1 is the unique model of $T_1^{I_1}$, then I_1 is a model of T_1 . Let $I_2 = \{\neg p, q\}$, $T_1^{I_2} = \{q\}$, both I_1 and I_2 are models of $T_1^{I_2}$, then I_2 is not a model of T_1 . We can see that, T_1 has two models $\{p, q\}$ and $\{\neg p, q\}$.

As a syntax sugar, a causal rule with variables is viewed as shorthand for the set of its ground instances, that is, for the result of substituting corresponding variable-free terms for variables in all possible ways.

3 Induction in Causal Theories

The induction problem considered in this paper is defined as follows. Given a causal theory T and an interpretation I , I is not a model of T , we need to modify T to a new causal theory T' such that I is a model of T' and each model of T is still a model of T' .

Before propose one of such modifications, we consider the notion of *relevance*. Following the intuition behind the work of Galles and Pearl [4], relevance is concerned with statements of the form “Changing X will alter the value of Y , if Z is fixed”. In the setting of nonmonotonic causal theories, the world is specified by a set of models, an atom a is related to another atom b under a set S of literals, if changing the value of a will alter possible evaluations of b given S .

In specific, given a set M of models and a set S of literals, an atom a is *semantically related* to another atom b under S w.r.t. M , if

- there exists two models $I_1, I_2 \in M$ s.t. $S \cup \{l_a, l_b\} \subseteq I_1$ and $S \cup \{\bar{l}_a, \bar{l}_b\} \subseteq I_2$,
- there does not exist a model $I' \in M$ s.t. $S \cup \{l_a, l_b\} \subseteq I'$,

where $l_a \in \{a, \neg a\}$ and $l_b \in \{b, \neg b\}$. Note that, given M and S , the defined relevance relation is reflexive, symmetric and transitive.

Given a causal theory T , an atom a is *semantically related* to an atom b in T , if there exists a set S of literals such that a is semantically related to b under S

w.r.t. the set of models of T . However, it is NP-hard to decide where two atoms are semantically related in a causal theory. Now, based on syntactic properties, we propose a relaxed definition of relevance which can be computed easily.

Given a causal theory T , an atom a is *syntactically related* to an atom b in T if (1) $a = b$, (2) both a and b occurs in a causal rule of T , or (3) both a and b are syntactically related to another atom c . Note that, the syntactical relevance relation is also reflexive, symmetric and transitive.

Proposition 1. *Given a causal theory T and two atoms a and b . If a is semantically related to b in T , then a is syntactically related to b .*

Now, based on syntactical relevance, we propose an approach for the induction problem. Given a casual theory T and an interpretation I which is not a model of T , let $tr(T, I)$ be the casual theory obtained from T by:

1. modifying each causal rule $r \in T$ to $body(r) \wedge head(r) \Rightarrow head(r)$, if $I \models body(r)$ and $I \not\models head(r)$; and
2. adding a causal rule $L \Rightarrow l$, for each literal $l \in I$ such that $T^I \not\models l$, where L is the conjunction of literals which are belonged to I and in which occurred atoms are syntactically related to the atom occurred in l .

The number of causal rules generated by the conversion is polynomial for an interpretation in the number of literals.

Proposition 2. *Given a causal theory T and an interpretation I which is not a model of T . I is a model of $tr(T, I)$ and every model of T is a model of $tr(T, I)$.*

Note that, we assume that the relevance relation in the domain has been revealed by the original causal theory. With the help of such relevance relation, not limited to I , some other interpretations might become the models of $tr(T, I)$. Consider the causal theory T_1 in Section 2, $I_2 = \{\neg p, q\}$ is not a model of T_1 , p is not related to q , then $tr(T_1, I_2) = T_1 \cup \{\neg p \Rightarrow \neg p\}$. In addition to I_2 , $tr(T_1, I_2)$ has another new model $\{\neg p, \neg q\}$.

Every new causal rule $L \Rightarrow l$ added in $tr(T, I)$ can be generalized by substituting variables for some constants occurred in the rule. Note that, every model of T is still a model of the new theory. Generally, the robot can be taught with multiple examples, which leads to some common inductive learning issues investigated in the literature.

4 A Case Study in a Domestic Service Robot's Domain

In this section, we demonstrate the inductive approach by a case study in a domestic service robot's domain.

As shown in Figure 1, there is a board putting on the edge of a table, with one end sticking out. There is also a can on each end of the board, respectively, keeping it balanced. The task is to pick up the can on the inside end. Note that, the outside can may fall if the robot picks up the inside can.




Fig. 1. Setting of the Case Study

First, following the approach in [5], we use causal rules to formalize the action domain. The atoms are expressions of the form a_t and f_t , where a , f , and t are action, fluent, and time names, respectively. Intuitively, a_t is true if and only if the action a occurs at time t , and f_t is true if and only if the fluent f holds at t .

In this setting, we focus on the robot's ability of 'grasp' and the corresponding properties of the environment. The action names and fluent names used in the specification are following, where X and Y are variables ranging over possible objects in the environment:

- $grasp(X)$: the action of gripping the object X and picking it up.
- $holding(X)$: the fluent that the object X is held in the grip of the robot.
- $on(X, Y)$: the fluent that the object X is on the object Y .
- $falling(X)$: the fluent that the object X is falling on the floor.

In addition, σ is a meta-variable ranging over $\{on(X, Y), \neg on(X, Y), holding(X), \neg holding(X), falling(X), \neg falling(X)\}$.

The effect of executing the action $grasp(X)$ is described as follows:

$$grasp(X)_t \Rightarrow holding(X)_{t+1} \quad (1)$$

$$grasp(X)_t \wedge on(X, Y)_t \Rightarrow \neg on(X, Y)_{t+1} \quad (2)$$

The precondition of grasping requires the grip holds nothing:

$$grasp(X)_t \wedge holding(Y)_t \Rightarrow \perp \quad (3)$$

The occurrence of the action is exogenous to the causal theory:

$$grasp(X)_t \Rightarrow grasp(X)_t \quad (4)$$

$$\neg grasp(X)_t \Rightarrow \neg grasp(X)_t \quad (5)$$

The initial state (at time 0) can be arbitrary:

$$\sigma_0 \Rightarrow \sigma_0 \quad (6)$$

There are some restrictions among these fluents:

$$holding(X)_t \wedge falling(X)_t \Rightarrow \perp \quad (7)$$

$$on(X, Y)_t \wedge falling(Y)_t \Rightarrow falling(X)_t \quad (8)$$

Rule (8) specifies that $falling(X)$ is an indirect effect of some action that causes $falling(Y)$ while X is on Y . The frame problem is overcome by the following "inertia" rules:

$$\sigma_t \wedge \sigma_{t+1} \Rightarrow \sigma_{t+1} \quad (9)$$

The causal theory formed by rules (1)–(9) represents the action domain for the robot’s ability of ‘grasp’. Let $0 \leq t \leq m$, the models of such causal theory correspond to the histories of the action domain whose length is m . In particular, an interpretation I is a model if and only if the state s_{i+1} is a successor state of the state s_i after the concurrent execution of actions A_i , where $s_i = \{f_i \in I \mid f \text{ is a fluent name}\}$ and $A_i = \{a_i \in I \mid a \text{ is an action name}\}$.

During the development of our robot, there are some real scenarios that have not been captured by a certain version of the causal theory. Particularly, in one execution the robot recognizes that s_1 is a successor state of an initial state s_0 after the occurrence of an action a_0 , but the interpretation $I = s_0 \cup \{a_0\} \cup \{\neg a'_0 \mid \text{action name } a' \text{ different from } a\} \cup s_1$ is not a model of the causal theory with $0 \leq t \leq 1$. In this case, we can use our inductive approach to modify the causal theory so that I becomes its model and all other models keep unchanged.

In the setting, with the current causal theory, the robot can not predict the end state in that the outside will fall after the action ‘grasp’ the inside can be executed. This new scenario is formally represented by the interpretation I :

$$\begin{aligned} &\{on(a, bd)_0, on(b, bd)_0, \neg holding(a)_0, \neg holding(b)_0, \neg falling(a)_0, \neg falling(b)_0, \\ &\quad \neg falling(bd)_0, grasp(b)_0, \neg grasp(a)_0, falling(bd)_1, \\ &\quad on(a, bd)_1, \neg on(b, bd)_1, \neg holding(a)_1, holding(b)_1, falling(a)_1, \neg falling(b)_1\} \end{aligned}$$

which is not a model of the causal theory T with two objects a, b , time names 0, 1, and the constant bd standing for the board. The interpretation I also expresses the knowledge that a and bd will fall after the action $grasp(b)$ occurs in the setting.

Using our inductive approach, the robot obtains a new causal theory $tr(T, I)$, which contains all causal rules in T and the new rule: $\bigwedge_{l \in I} l \Rightarrow falling(bd)_1$. Note that, every atom occurred in I is syntactically related to $falling(bd)_1$. Let M be the set of models of T , we can see that none of interpretations in M satisfy $falling(bd)_1$ and $falling(bd)_1 \in I$, then every atom is semantically related to $falling(bd)_1$ under some set of literals w.r.t. $M \cup \{I\}$.

The learned rule can be further generalized by changing 0 to t , 1 to $t + 1$, and some constants like a, b , and bd can be replaced by variables.

According to Proposition 2, I is a model of $tr(T, I)$ and every model of T is still a model of $tr(T, I)$. Then the robot uses $tr(T, I)$ as the new model of the action domain and is aware that if it picks up b in the setting then a has a possibility to fall. Thus, to accomplish the task, the robot would compute a more cautious plan in which it removes a from the sticking-out end of the board first before grasping b ³.

³ A similar case study is featured in the video demo “Towards Robot Learning from Comparative Demonstration” at <http://ai.ustc.edu.cn/en/demo/>, where the robot is taught with a positive and a negative example in the same setting.

5 Discussion and Conclusion

We use nonmonotonic causal theories to model the action domains of the robot, mainly because:

- They are convenient to formalize action domains, including the frame problem, indirect effects of actions, concurrent actions, nondeterministic actions, ramification and qualification constraints.
- They can be easily computed by existing sophisticated solvers. The problem of computing models of a causal theory can be equivalently translated to computing answer sets of a logic program with negative as failure [6] and solved by ASP solvers. The causal theory can also be converted to a propositional theory, whose models can be computed by SAT solvers [6, 5].
- More importantly, this effort provides an evidence that the modification of a theory for an action domain is convenient by using nonmonotonic causal theories. Intuitively, an enlarged causal theory most likely covers new interpretations and original models as its models this way: if no abnormal features related to the service task are observed, then the original knowledge still works; otherwise, the newly generated knowledge from the new scenario should be used. Both the original and the new knowledge are integrated conveniently into the enlarged theory as a whole due to the nonmonotonicity of causal theories.

However, there is few work on inductive learning in nonmonotonic causal theories. On the other hand, the problem is closely related to the context of Nonmonotonic Inductive Logic Programming [7]. Sakama [8] proposed an inductive learning approach in nonmonotonic logic programs. Our approach differs from Sakama’s approach in that we may need to modify rules in the original causal theory while Sakama keeps the rules in the original program unchanged.

References

1. Argall, B., Chernova, S., Veloso, M., Browning, B.: A survey of robot learning from demonstration. *Robotics and Autonomous Systems* 57(5), 469–483 (2009)
2. Chen, X., Ji, J., Jiang, J., Jin, G., Wang, F., Xie, J.: Developing high-level cognitive functions for service robots. In: *Proceedings of AAMAS-10*. pp. 989–996 (2010)
3. Chen, X., Jiang, J., Ji, J., Jin, G., Wang, F.: Integrating nlp with reasoning about actions for autonomous agents communicating with humans. In: *Proceedings of IAT-09*. pp. 137–140 (2009)
4. Galles, D., Pearl, J.: Axioms of causal relevance. *Artificial Intelligence* 97(1-2), 9–43 (1997)
5. Giunchiglia, E., Lee, J., Lifschitz, V., McCain, N., Turner, H.: Nonmonotonic causal theories. *Artificial Intelligence* 153(1-2), 49–104 (2004)
6. Lee, J.: Nondefinite vs. definite causal theories. In: *Proceedings of LPNMR-04*. pp. 141–153 (2004)
7. Sakama, C.: Nonmonotonic inductive logic programming. In: *Proceedings of LPNMR-01*. pp. 62–80 (2001)
8. Sakama, C.: Induction from answer sets in nonmonotonic logic programs. *ACM Transactions on Computational Logic (TOCL)* 6(2), 203–231 (2005)