Specifying Plausibility Levels for Iterated Belief Change in the Situation Calculus

Toryn Q. Klassen and Sheila A. McIlraith and Hector J. Levesque
{**toryn**,sheila,hector}@cs.toronto.edu

Department of Computer Science
University of Toronto

November 1, 2018

Introduction

We will present a framework for

1. iterated belief revision and update
2. modeling of action and change
3. allowing a simple qualitative specification of what the agent considers plausible

Introduction

We will present a framework for

1. iterated belief revision and update
2. modeling of action and change

Shapiro et al. (2011)

3. allowing a simple qualitative specification of what the agent considers plausible

Outline

1. Preliminaries
   - The situation calculus
   - Belief change in the situation calculus (Shapiro et al., 2011)
2. Related work on specifying plausibility levels
   - Only-believing (Schwering and Lakemeyer, 2014)
   - Issues with only-believing
3. Our approach
   - Cardinality-based circumscription
   - Using abnormality fluents to define plausibility
   - Examples
   - Why cardinality-based circumscription?
   - Exogenous actions

The situation calculus (Reiter, 2001)

Key points:

- Situations represent **histories** of actions performed starting from an initial situation.
- Properties that can vary among situations are described using **fluents**, which are predicates (or functions) whose last argument is a situation term, e.g. $P(x, s)$.

The situation calculus (Reiter, 2001)

Key points:

- Situations represent **histories** of actions performed starting from an initial situation.
- Properties that can vary among situations are described using **fluents**, which are predicates (or functions) whose last argument is a situation term, e.g. $P(x, s)$.

Some notation:

- $\mathtt{S}_0$ is the actual initial situation.
- $\mathtt{do}(a, s)$ is the situation that results from performing action a in situation s.
- $\mathtt{do}([a_1, \ldots, a_k], s)$ is the situation resulting from performing actions $a_1, \ldots, a_k$ in order from s.

The situation tree

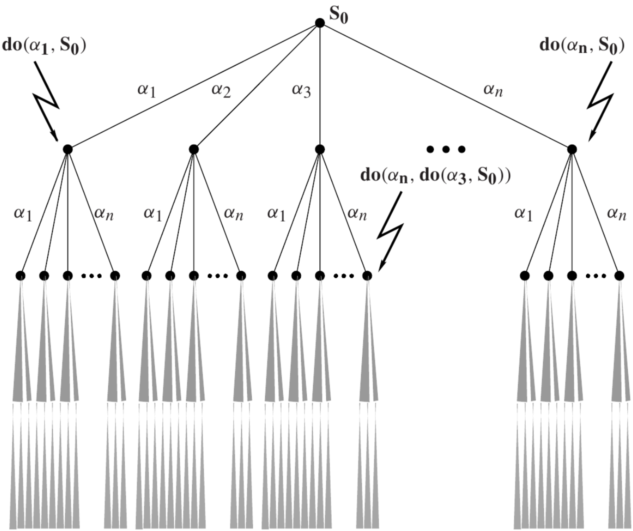

Figure copied from Reiter (2001, Figure 4.1).

Multiple situation trees

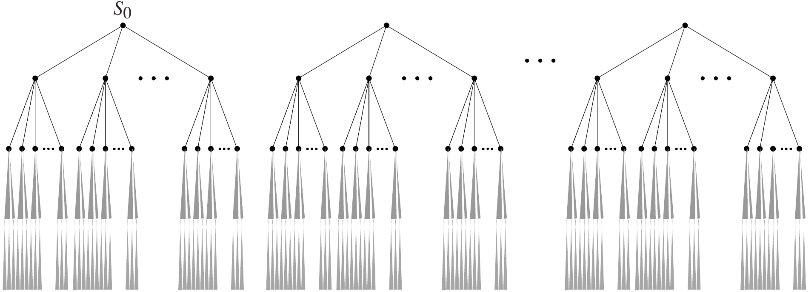

Figure copied from Reiter (2001, Figure 11.7).

Action theories for the situation calculus

The standard way of axiomatizing domains is with some variation of **basic action theories** (Reiter, 2001).

Basic action theories

- **initial state axioms**, which describe the initial situation(s)
- **successor state axioms (SSAs)**, specifying for each fluent how its value in a non-initial situation depends on the previous situation
- (sometimes) **sensing axioms**
- and also some other types (precondition axioms, unique names axioms, foundational axioms)

Iterated belief change in the situation calculus

Shapiro et al. (2011)'s approach has these main points:

- There is an **epistemic accessibility relation** between situations.
- Each initial situation is assigned a numeric **plausibility** level.
- The agent **believes** what is true in all the **most plausible** epistemically accessible situations.
- Sensing actions can make more situations inaccessible (plausibility levels never change).

Outline

1. Preliminaries
    - The situation calculus
    - Belief change in the situation calculus (Shapiro et al., 2011)
2. Related work on specifying plausibility levels
    - Only-believing (Schwering and Lakemeyer, 2014)
    - Issues with only-believing
3. Our approach
    - Cardinality-based circumscription
    - Using abnormality fluents to define plausibility
    - Examples
    - Why cardinality-based circumscription?
    - Exogenous actions

Deriving plausibilities with only-believing

Schwering and Lakemeyer (2014) had an approach for specifying plausibility levels in their modal version of the situation calculus.

Deriving plausibilities with only-believing

Schwering and Lakemeyer (2014) had an approach for specifying plausibility levels in their modal version of the situation calculus.

- $\mathbf{B}(\alpha \Rightarrow \beta)$ holds if β is true in all the most plausible accessible α-worlds.

Deriving plausibilities with only-believing

Schwering and Lakemeyer (2014) had an approach for specifying plausibility levels in their modal version of the situation calculus.

- $\mathbf{B}(\alpha \Rightarrow \beta)$ holds if β is true in all the most plausible accessible α-worlds.
- $\mathbf{O}(\alpha_1 \Rightarrow \beta_1, \ldots, \alpha_k \Rightarrow \beta_k)$ holds only given a particular **unique** assignment of plausibility values.

Deriving plausibilities with only-believing

Schwering and Lakemeyer (2014) had an approach for specifying plausibility levels in their modal version of the situation calculus.

- $\mathbf{B}(\alpha \Rightarrow \beta)$ holds if β is true in all the most plausible accessible α-worlds.
- $\mathbf{O}(\alpha_1 \Rightarrow \beta_1, \ldots, \alpha_k \Rightarrow \beta_k)$ holds only given a particular **unique** assignment of plausibility values.
  - an assignment that entails $\bigwedge_i \mathbf{B}(\alpha_i \Rightarrow \beta_i)$
  - determined like in **System Z** (Pearl, 1990)

Issues with only-believing

1. lack of **independence**:

$$\mathbf{O}(\texttt{True} \Rightarrow P, \texttt{True} \Rightarrow Q) \not\models \mathbf{B}(\neg P \Rightarrow Q)$$

Issues with only-believing

1. lack of **independence**:

$$\mathbf{O}(\mathtt{True} \Rightarrow P, \mathtt{True} \Rightarrow Q) \not\models \mathbf{B}(\neg P \Rightarrow Q)$$

2. can only specify a **finite** number of plausibility levels:

   We can write

$$\mathbf{O}(\mathtt{True} \Rightarrow (\forall x)P(x))$$

   But this is not grammatical:

$$\mathbf{O}((\forall x).\mathtt{True} \Rightarrow P(x))$$

Outline

1. Preliminaries
- The situation calculus
- Belief change in the situation calculus (Shapiro et al., 2011)

2. Related work on specifying plausibility levels
- Only-believing (Schwering and Lakemeyer, 2014)
- Issues with only-believing

3. Our approach
- Cardinality-based circumscription
- Using abnormality fluents to define plausibility
- Examples
- Why cardinality-based circumscription?
- Exogenous actions

Cardinality-based circumscription

Popular idea in non-monotonic reasoning:

> Instead of considering what is true in all models of a sentence, consider what is true in **preferred** models.

Cardinality-based circumscription:

- the preferred models are those where the **cardinalities** of particular predicates are minimized (Liberatore and Schaerf, 1997; Sharma and Colomb, 1997; Moinard, 2000)
- can be described using **second order** logic
- closely related to **lexicographic entailment** (Benferhat et al., 1993; Lehmann, 1995)

Determining the plausibility of situations

How can we apply this to situation calculus?

- Introduce **abnormality fluents**, whose values vary in different initial situations.
- Define the plausibility of a situation by the number of abnormal atoms true there.
  - We can also consider **priorities** – see paper.

How to specify the initial accessibility relation?

- Use **only-knowing** (Lakemeyer and Levesque, 1998).
- $\texttt{OKnows}(\phi, s)$ says that the situations that are epistemically accessible from s are those where ϕ is true.

Example

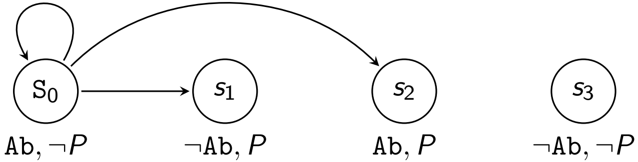

- The accessible situations (from $\mathtt{S}_0$) are those in which $\neg\mathtt{Ab} \supset P$ is true.
- The set of most plausible accessible situations is $\{s_1\}$.
- P is true at all the most plausible accessible situations.
- The agent believes P in $\mathtt{S}_0$.

Immutable abnormality action theories

Differ from Shapiro et al.'s theories in that we

- include an axiom of the form $\mathtt{OKnows}(\phi, \mathtt{S}_0)$ to specify the initial accessibility relation,
- redefine plausibility in terms of abnormality,
- have SSAs for the abnormality fluents (specifying that they never change),
- and include an additional axiom ensuring the existence of enough initial situations among the foundational axioms.

Example 1: independently plausible propositions

Initial state axioms:

$$\neg P(\mathtt{S}_0) \wedge \neg Q(\mathtt{S}_0)$$

$$\mathtt{OKnows}((\neg \mathtt{Ab}_1 \supset P) \wedge (\neg \mathtt{Ab}_2 \supset Q), \mathtt{S}_0)$$

Successor state axioms:

$$P(\mathtt{do}(a, s)) \equiv P(s) \qquad Q(\mathtt{do}(a, s)) \equiv Q(s)$$

Sensing axioms:

$$\mathtt{SF}(\textsc{senseP}, s) \equiv P(s) \qquad \mathtt{SF}(\textsc{senseQ}, s) \equiv Q(s)$$

Example 1: independently plausible propositions

Initially, the accessible situations from $\mathtt{S}_0$ are those initial situations where $(\neg\mathtt{Ab}_1 \supset P) \land (\neg\mathtt{Ab}_2 \supset Q)$ is true.

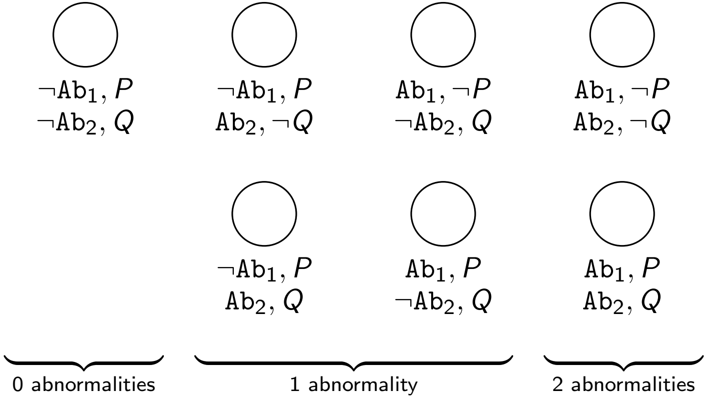

Example 1: independently plausible propositions

After performing SENSEP, the situations where P differs from its true value (false) become **inaccessible**.

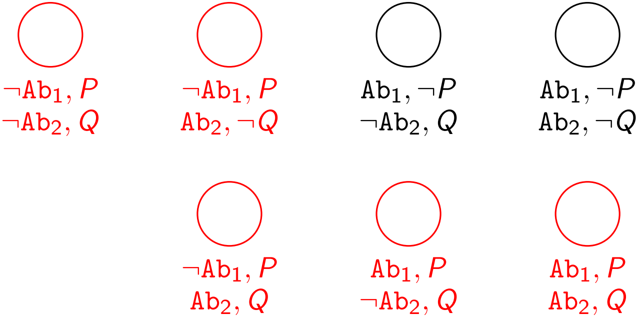

Example 1: independently plausible propositions

After performing SENSEP, the situations where P differs from its true value (false) become **inaccessible**.

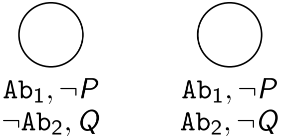

Example 1: independently plausible propositions

After performing SENSEQ, the situations where Q differs from its true value (false) become **inaccessible**.

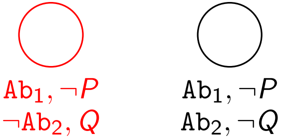

Example 1: independently plausible propositions

After performing SENSEQ, the situations where Q differs from its true value (false) become **inaccessible**.

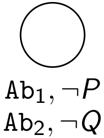

Example 1: independently plausible propositions

$$\neg P(\mathtt{S}_0) \wedge \neg Q(\mathtt{S}_0)$$
$$\mathtt{OKnows}((\neg\mathtt{Ab}_1 \supset P) \wedge (\neg\mathtt{Ab}_2 \supset Q), \mathtt{S}_0)$$
$$\mathtt{SF}(\textsc{senseP}, s) \equiv P(s) \quad \mathtt{SF}(\textsc{senseQ}, s) \equiv Q(s)$$
$$P(\mathtt{do}(a, s)) \equiv P(s) \qquad Q(\mathtt{do}(a, s)) \equiv Q(s)$$

Proposition

Let Σ be the immutable abnormality action theory described above. Then

$$\Sigma \models \mathtt{Bel}(P \wedge Q, \mathtt{S}_0)$$
$$\Sigma \models \mathtt{Bel}(\neg P \wedge Q, \mathtt{do}(\textsc{senseP}, \mathtt{S}_0))$$
$$\Sigma \models \mathtt{Bel}(\neg P \wedge \neg Q, \mathtt{do}([\textsc{senseP}, \textsc{senseQ}], \mathtt{S}_0))$$

Example 2: infinitely many plausibility levels

Initial state axioms:

$$\textsc{Conspirator}(x, \mathtt{S}_0)$$

$$\mathtt{OKnows}\big((\forall x)\neg\mathtt{Ab}(x) \supset \neg\textsc{Conspirator}(x), \mathtt{S}_0\big)$$

Successor state axioms:

$$\textsc{Conspirator}(x, \mathtt{do}(a, s)) \equiv \textsc{Conspirator}(x, s)$$

Sensing axioms:

$$\mathtt{SF}(\textsc{reveal}(x), s) \equiv \textsc{Conspirator}(x, s)$$

Example 2: infinitely many plausibility levels

$$\textsc{Conspirator}(x, \mathrm{S}_0)$$
$$\mathtt{OKnows}\big((\forall x)\neg\mathtt{Ab}(x) \supset \neg\textsc{Conspirator}(x), \mathrm{S}_0\big)$$
$$\textsc{Conspirator}(x, \mathtt{do}(a, s)) \equiv \textsc{Conspirator}(x, s)$$
$$\mathtt{SF}(\textsc{reveal}(x), s) \equiv \textsc{Conspirator}(x, s)$$

Proposition

Let Σ be the immutable abnormality action theory described above, and let $\mathrm{C}_1, \mathrm{C}_2, \mathrm{C}_3, \ldots$ be constant symbols. Then for any k,

$$\Sigma \models \mathtt{Bel}\Big((\forall x)\textsc{Conspirator}(x) \equiv \Big(\bigvee_{i=1}^{k} x = \mathrm{C}_i\Big), \mathtt{do}([\textsc{reveal}(\mathrm{C}_1), \ldots, \textsc{reveal}(\mathrm{C}_k)], s)\Big)$$

Outline

1. Preliminaries
    - The situation calculus
    - Belief change in the situation calculus (Shapiro et al., 2011)
2. Related work on specifying plausibility levels
    - Only-believing (Schwering and Lakemeyer, 2014)
    - Issues with only-believing
3. Our approach
    - Cardinality-based circumscription
    - Using abnormality fluents to define plausibility
    - Examples
    - Why cardinality-based circumscription?
    - Exogenous actions

Why not use regular (subset-based) circumscription?

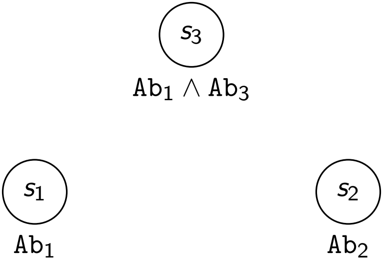

Why not use regular (subset-based) circumscription?

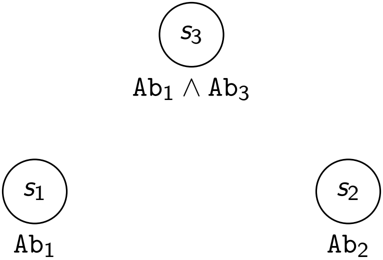

Cardinality-based and regular circumscription **agree** that s_1 and s_2 are the most plausible accessible situations.

Why not use regular (subset-based) circumscription?

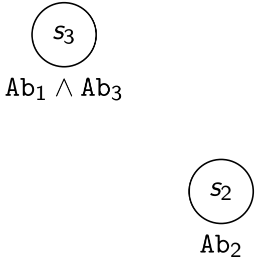

Cardinality-based and regular circumscription **agree** that s_1 and s_2 are the most plausible accessible situations.

Now suppose that s_1 becomes **inaccessible** (e.g. due to sensing).

Why not use regular (subset-based) circumscription?

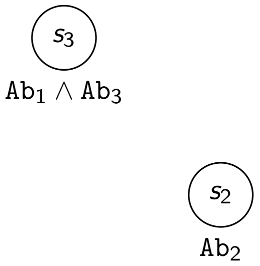

- Cardinality-based circumscription: s_2 is now the most plausible accessible situation

Why not use regular (subset-based) circumscription?

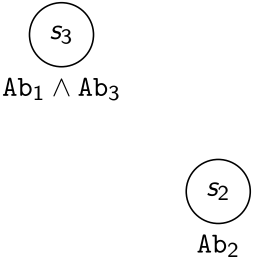

- Cardinality-based circumscription: s_2 is now the most plausible accessible situation
- Regular circumscription: not only s_2 but s_3 is now a most plausible accessible situation

Why not use regular (subset-based) circumscription?

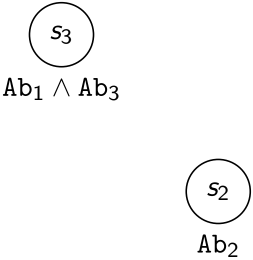

- Cardinality-based circumscription: s_2 is now the most plausible accessible situation
- Regular circumscription: not only s_2 but s_3 is now a most plausible accessible situation
  - leads to violation of **AGM postulates** (Alchourrón et al., 1985)

Exogenous actions

What if we allowed abnormality fluents to change over time?

- **Mutable** abnormality action theories can be used to model **exogenous actions**.
- Exogenous actions were previously considered by Shapiro and Pagnucco (2004), but unlike them we can model that
  - some exogenous actions are **more plausible** than others, and
  - the **non-occurrence** of an exogenous action can be implausible.
- See paper for details.

Example: the fate of abandoned money

- ONSTREET(s): money is on the street
- STEAL: the **exogenous** action of money being stolen

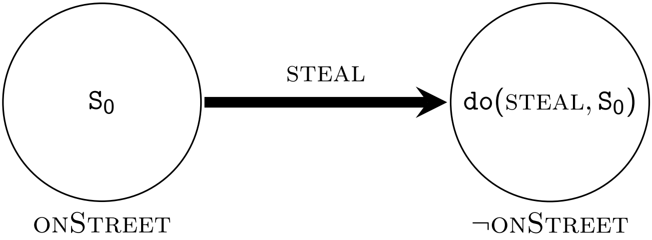

Example: the fate of abandoned money

- $\textsc{onStreet}(s)$: money is on the street
- $\textsc{steal}$: the **exogenous** action of money being stolen

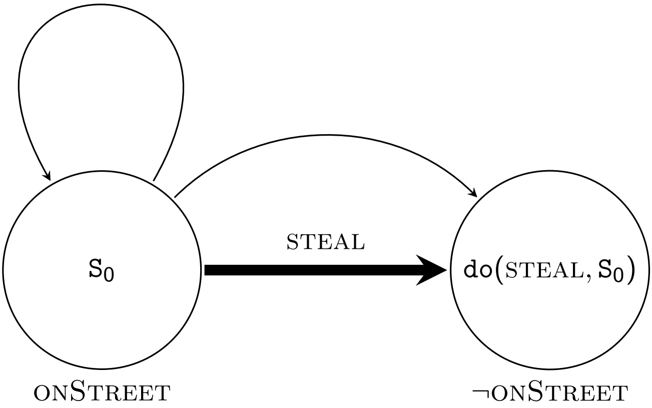

Example: the fate of abandoned money

- ONSTREET(s): money is on the street
- STEAL: the **exogenous** action of money being stolen

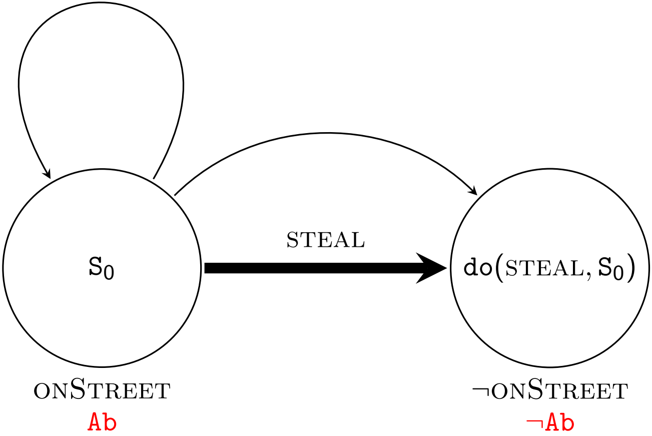

Conclusion

Summary:

We've presented a way of specifying plausibility levels for use in the situation calculus, that avoids some of the issues with Schwering and Lakemeyer's approach.

- We can easily specify propositions as being **independently** plausible.
- We can specify **infinitely** many plausibility levels.

Future work:

- using abnormalities in modelling **non-deterministic actions**
- applications to **story understanding**

References

Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2):510–530, 1985. doi: 10.2307/2274239.

Salem Benferhat, Claudette Cayrol, Didier Dubois, Jerome Lang, and Henri Prade. Inconsistency management and prioritized syntax-based entailment. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence - Volume 1*, IJCAI'93, pages 640–645, 1993.

Gerhard Lakemeyer and Hector J. Levesque. AOL: a logic of acting, sensing, knowing, and only knowing. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 316–327, 1998.

Daniel Lehmann. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence*, 15 (1):61–82, 1995.

Paolo Liberatore and Marco Schaerf. Reducing belief revision to circumscription (and vice versa). *Artificial Intelligence*, 93(1):261–296, 1997.

Yves Moinard. Note about cardinality-based circumscription. *Artificial Intelligence*, 119(1):259 – 273, 2000.

Judea Pearl. System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning About Knowledge*, TARK '90, pages 121–135, 1990.

Raymond Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press, 2001.

Christoph Schwering and Gerhard Lakemeyer. A semantic account of iterated belief revision in the situation calculus. In *ECAI 2014 - 21st European Conference on Artificial Intelligence*, pages 801–806, 2014.

Steven Shapiro and Maurice Pagnucco. Iterated belief change and exogeneous actions in the situation calculus. In *Proceedings of the 16th Eureopean Conference on Artificial Intelligence, ECAI'2004*, pages 878–882, 2004.

Steven Shapiro, Maurice Pagnucco, Yves Lespérance, and Hector J. Levesque. Iterated belief change in the situation calculus. *Artificial Intelligence*, 175(1):165–192, 2011. doi: 10.1016/j.artint.2010.04.003.

Nirad Sharma and Robert Colomb. Towards an integrated characterisation of model-based diagnosis and configuration through circumscription policies. Technical Report 364, Department of Computer Science, University of Queensland, 1997.