

A Circumscriptive Formalization of the Qualification Problem

G. Neelakantan Kartha

528 Blackfield Drive

Coppell, Tx 75019

gnkartha@yahoo.com

Abstract

The qualification problem refers to the difficulty that arises in formalizing actions, because it is difficult or impossible to specify in advance all the preconditions that should hold before an action can be executed. We study the qualification problem in the setting of the situation calculus and give a simple formalization using nested abnormality theories, a formalism based on circumscription. The formalization that we present allows us to combine a solution to the frame problem with a solution to the qualification problem.

1 Introduction

Two problems that are important in logical formalizations of commonsense reasoning are the frame problem and the qualification problem [McCarthy, 1977; 1980]. While the frame problem has received a lot of attention in the literature in recent years [Lin and Shoham, 1991; Lifschitz, 1991; Kartha and Lifschitz, 1994; Sandewall, 1994; McCain and Turner, 1995; Thielscher, 1997; Giunchiglia *et al.*, 1997; Shanahan, 1997; McIlraith, 2000], the qualification problem has not been studied as carefully (see however the discussion in the section on related work). In this paper, we clearly define what we mean by the qualification problem and introduce a formalization of the problem using nested abnormality theories [Lifschitz, 1995]. This formalization has the advantage that solutions to the frame problem can be incorporated smoothly with the solution to the qualification problem. We also point out problematic aspects of the current approaches to the qualification problem.

The qualification problem [McCarthy, 1980] refers to the difficulty that one has in formalizing an action, because it is difficult or impossible to enumerate all the preconditions that need to be satisfied before the action can be performed. One of the more well known examples that illustrate this problem is the potato in the tailpipe example due to McCarthy. The example is as follows: The action of starting a car results in the engine running. However, there are many implicit preconditions to the action of starting the car, for instance that there is no potato in the tailpipe, the battery is not dead etc. The problem is how to formalize such examples without getting bogged down in the representation of all such preconditions.

Let us examine the example in a little more detail. The action of starting the car (without any qualifications) can be formalized in the situation calculus [McCarthy and Hayes, 1969] as follows:

$$\neg Holds(EngineRunning, s) \supset \\ Holds(EngineRunning, Result(StartCar, s)). \quad (1)$$

To take into account various preconditions of the action *StartCar*, we could modify (1) by adding preconditions to its antecedent as in,

$$\neg Holds(EngineRunning, s) \wedge \neg Holds(PotatoInTailPipe, s) \\ \supset Holds(EngineRunning, Result(StartCar, s)). \quad (2)$$

However, note that we will be able to include in the language of our axiomatization only a subset of *all* the preconditions that can affect the outcome of an action, because we cannot foresee all of them. For instance, a factor that *could* affect the outcome of the action *StartCar* is the presence (or absence) of a mysterious ray. We would include facts about this ray in the language of our axiomatization only when we are given that this is a relevant fact.

Let us then start by fixing the language. Even then, it would be difficult to foresee all the circumstances (in the language) that can affect an action. Thus the qualification problem does not vanish when the language is fixed, but takes on a slightly different flavor. Now the key question becomes how one can axiomatize the domain of interest in such a way that on learning a new precondition for an action, one is able to incorporate this additional knowledge in a modular fashion in the axiomatization. Note that adding preconditions to the antecedent, as we have done in axiom (2) is not very modular.

Thus, it seems that there are really two related, but distinct aspects to the qualification problem.

1. The problem of how to smoothly incorporate various additions to the language of an axiomatization. The addition into the axiomatization of the facts pertaining to the mysterious ray and its effect on starting the car is an instance of this. Note that here the language of the axiomatization changes as a result of incorporating the change. This aspect of the qualification problem is related to McCarthy's notion of elaboration tolerance¹ [McCarthy,

¹According to McCarthy, "a formalism is elaboration tolerant to

1998]. This aspect is also related to the problem of how to move between axiomatizations (in different languages) that axiomatize a domain in different levels of detail.

2. The problem of how to represent the qualifications to an action, when the language of the axiomatization is fixed. The chief issues here are

- (a) Whether we can easily incorporate the fact that another circumstance (in the language) qualifies an action without drastically modifying the original axiomatization. For instance, we would like to achieve this by adding a new sentence into the language. Again, this is because of our desire for an elaboration tolerant solution.
- (b) To see whether the conclusions that are sanctioned formally agree with our intuitions.

Note that the two aspects share a desire for elaboration tolerance, so that a solution for one aspect is likely to contribute to a solution for the other. The current paper deals only with the second of the two aspects of the qualification problem.

The rest of the paper is organized as follows. In the next section, we explain our formalization and give an example of the kind of results that can be proved using our formalization. We then survey related work, summarize the key contributions of this work and point out directions for future work.

2 Formalization

Since our aim is to convey the key ideas, we will explain our formalization via an example, but will indicate at the appropriate places how one would generalize. The example we use is a slightly extended version of the potato in the tailpipe example from [Thielscher, 1996]. The example consists of extending the basic potato in the tailpipe scenario with another action, *PutPotato*. The result of this action is to put the potato in the tailpipe, unless this action is qualified. This action is qualified if the potato is so heavy that it cannot be lifted.

Following McCarthy [McCarthy, 1986], here are some axioms that capture the scenario:

$$\neg \text{Holds}(\text{EngineRunning}, s) \wedge \neg \text{Ab}(\text{StartCar}, s) \supset \text{Holds}(\text{EngineRunning}, \text{Result}(\text{StartCar}, s)), \quad (3)$$

$$\neg \text{Holds}(\text{PotatoInTailPipe}, s) \wedge \neg \text{Ab}(\text{PutPotato}, s) \supset \text{Holds}(\text{PotatoInTailPipe}, \text{Result}(\text{PutPotato}, s)), \quad (4)$$

$$\text{Holds}(\text{PotatoInTailPipe}, s) \supset \text{Ab}(\text{StartCar}, s), \quad (5)$$

$$\text{Holds}(\text{HeavyPotato}, s) \supset \text{Ab}(\text{PutPotato}, s), \quad (6)$$

$$\neg \text{Holds}(\text{EngineRunning}, S_0). \quad (7)$$

the extent it is convenient to modify a set of facts expressed in the formalism to take into account new phenomena or changed circumstances”.

Here *Ab* is an abnormality predicate that will be minimized. The advantage of this representation is that as we think of new qualifications, they can be easily incorporated into the theory by means of axioms such as (5) and (6).

However, it should be noted that we are rarely interested in solving the qualification problem in isolation. We also need to ensure that the truth values of the fluents do not change without a reason, so that we are able to obtain the intuitively expected results. To see this, let *T* be the theory consisting of axioms (3), (4), (5), (6) and (7). Let *T'* be the theory obtained by minimizing *Ab* in *T*. We would like *T'* to entail the formula $\neg \text{Ab}(\text{PutPotato}, S_0)$ and consequently $\neg \text{Holds}(\text{HeavyPotato}, S_0)$. Although, intuitively we would then conclude that the fluent *HeavyPotato* remains false after executing an action in *S*₀ (since there are no actions that change the truth value of that fluent), *T'* is not strong enough to entail formulae such as $\neg \text{Holds}(\text{HeavyPotato}, \text{Result}(\text{StartCar}, S_0))$, $\neg \text{Holds}(\text{HeavyPotato}, \text{Result}(\text{PutPotato}, S_0))$ etc. This shows that we need to combine a solution to the frame problem with a solution to the qualification problem to enable us to draw the intuitively expected conclusions.

The formalism that we use that helps us in tackling the frame problem and the qualification problem in sequence, without the solutions interfering with each other is the formalism of nested abnormality theories (NATs) [Lifschitz, 1995]. The main feature of NATs is that the effect of a circumscription can be restricted to a part of the theory called a block. Formally, a block is of the form $\{C_1, \dots, C_m : \Phi_1, \dots, \Phi_n\}$, where *C*₁, ..., *C*_{*m*} are function and/or predicate constants and $\Phi_1, \dots, \Phi_n$ are axioms, possibly containing an abnormality predicate *Ab*. The block corresponds to the circumscription of *Ab* in the theory consisting of $\Phi_1, \dots, \Phi_n$, with *C*₁, ..., *C*_{*m*} allowed to vary. Note that each of $\Phi_1, \dots, \Phi_n$ itself may be a block.

Formally, the semantics of NATs is defined by a map *g* that translates blocks into sentences of the underlying language. We define *g* as follows:

$$g\{C_1, \dots, C_m : \Phi_1, \dots, \Phi_n\} = \exists ab G(ab),$$

where

$$G(Ab) = \text{CIRC}[g\Phi_1 \wedge \dots \wedge g\Phi_n; Ab; C_1, \dots, C_m].$$

Here CIRC stands for the circumscription of *Ab* in $g\Phi_1 \wedge \dots \wedge g\Phi_n$ with *C*₁, ..., *C*_{*m*} allowed to vary (see [Lifschitz, 1993a] for the formal definitions). For a sentence Φ that does not include *Ab*, $g\Phi$ is equivalent to Φ .

Since we need to consider the frame problem, we will need to axiomatize the commonsense law of inertia, according to which the truth values of fluents persist unless they are abnormal. Using ideas from [Karthi and Lifschitz, 1995], we formalize this as

$$\neg \text{Ab}(f, a, s) \supset \text{Holds}(f, s) \equiv \text{Holds}_R(f, a, s), \quad (8)$$

and denote it by *LI*. Here *Holds*_{*R*} is a new predicate, explicitly defined outside the block solving the frame problem as

$$\text{Holds}_R(f, a, s) \equiv \text{Holds}(f, \text{Result}(a, s)).$$

The role of this predicate for solving the frame problem can be explained roughly as follows—by “factoring away” the result function, we are able to focus attention only on the situation s and the situation that we obtain after action a is performed in s . For more details and intuitions, the reader is referred to [Karthi and Lifschitz, 1995].

In our formalization, we will formalize the effect of performing an action using the $Holds_R$ predicate. Thus, for instance, the axioms (3) and (4) will become

$$\neg Holds(EngineRunning, s) \wedge \neg Ab(StartCar, s) \supset Holds_R(EngineRunning, StartCar, s), \quad (9)$$

$$\neg Holds(PotatoInTailPipe, s) \wedge \neg Ab(PutPotato, s) \supset Holds_R(PotatoInTailPipe, PutPotato, s). \quad (10)$$

Let T_e denote the conjunction of axioms that formalize the effect of performing an action using the predicate $Holds_R$. Then each formula of T_e is of the syntactic form

$$pre_{F,A}^+(s) \wedge \neg Ab(A, s) \supset Holds_R(F, A, s)$$

or

$$pre_{F,A}^-(s) \wedge \neg Ab(A, s) \supset \neg Holds_R(F, A, s).$$

where $pre_{F,A}^\pm(s)$ denotes a finite conjunction of formulas of the form $Holds(F, s)$ or $\neg Holds(F, s)$. In our example, T_e thus stands for the conjunction of (9) and (10).

Similarly, let T_{qua} stand for the conjunction of axioms that formalize the qualifications. Syntactically, these axioms are of the form

$$qual_A(s) \supset Ab(A, s)$$

where $qual_A(s)$ is a finite disjunction of conjunction of formulas of the form $Holds(F, s)$ and $\neg Holds(F, s)$. In our example, T_{qua} is the conjunction of (5) and (6).

Let T_{obs} stand for the conjunction of axioms that formalize the observations. These axioms are a finite conjunction of formulas of the form $Holds(F, S_0)$ and $\neg Holds(F, S_0)$. In our case, T_{obs} is (7).

We also introduce the uniqueness of names and domain closure axioms for fluents and actions. For instance, in our example, the uniqueness of names axiom for fluents consists of the following axiom.

$$\begin{aligned} EngineRunning &\neq HeavyPotato \wedge \\ EngineRunning &\neq PotatoInTailPipe \wedge \\ HeavyPotato &\neq PotatoInTailPipe. \end{aligned}$$

The domain closure axiom for fluents is

$$\begin{aligned} f &= EngineRunning \vee \\ f &= PotatoInTailPipe \vee \\ f &= HeavyPotato. \end{aligned}$$

Let UNA and DCA denote respectively the uniqueness of names axioms and the domain closure axioms for fluents and actions.

We need two technical devices for our axiomatization. Firstly, we need to include an existence of situation axiom ([Baker, 1991]) (denoted ESI) that states that corresponding to every set of fluents, there is a situation in which exactly

these fluents hold. For our example, the existence of situation axiom consists of the conjunction of

$$\begin{aligned} \exists s (\neg Holds(EngineRunning, s) \wedge \\ \neg Holds(PotatoInTailPipe, s) \wedge \\ \neg Holds(HeavyPotato, s)), \end{aligned}$$

and seven other similar formulae asserting the existence of a situation corresponding to other possible truth values of the three fluents.

Secondly, we need to include a formula which states that there is a situation distinct from S_0 where exactly those fluents that hold in S_0 hold. Formally, this formula is

$$\exists s [s \neq S_0 \wedge \forall f [Holds(f, s) \equiv Holds(f, S_0)]].$$

This formula is denoted ESI (for existence of situation for the initial situation).

We are now ready to give our formalization.

$$\begin{aligned} &UNA, \\ &DCA, \\ &\{Holds, Holds_R, Result, S_0 : \\ &\quad ES, \\ &\quad ESI, \\ &\quad Holds_R(f, a, s) \equiv Holds(f, Result(a, s)), \\ &\quad \{Holds_R : \\ &\quad\quad LI, \\ &\quad\quad T_e, \\ &\quad\quad T_{qua}, \\ &\quad\quad T_{obs}, \\ &\quad\quad \} \\ &\quad \}. \end{aligned}$$

The inner block solves the frame problem. Note that the abnormality predicate that is being circumscribed in this block is the ternary one that occurs in LI , not the binary one occurring in T_e . Also, T_{qua} and T_{obs} do not contain any occurrence of $Holds_R$ and $Ab(f, a, s)$, and hence do not play any role in the circumscription of the inner block². The outer block minimizes the binary Ab and solves the qualification problem. The existence of situations axioms ensure that the binary abnormalities are minimized over *all* situations. The strategy of first solving the frame problem and then the qualification problem can also be found in [Ginsberg and Smith, 1988; Lin and Reiter, 1994].

Let N be the NAT that encodes our example. By entailment relation for N , we will mean the circumscriptive entailment relation for NATs (for details, see [Lifschitz, 1995]). As a representative of the kind of results we can prove, we state the following theorem.

Theorem 1 *The following set of formulae are entailed by N .*

1. $\neg Holds(HeavyPotato, S_0) \wedge \neg Holds(PotatoInTailPipe, S_0)$.
2. $\forall as (Holds(HeavyPotato, s) \equiv Holds(HeavyPotato, Result(a, s)))$.
3. $\neg Holds(EngineRunning, Result(PutPotato, S_0))$.
4. $\neg Holds(EngineRunning, Result(StartCar, S_1))$ where $S_1 = Result(PutPotato, S_0)$.

²This can also be achieved by moving T_{qua} and T_{obs} from the inner to the outer block.

The first formula states that the fluents *HeavyPotato* and *PotatoInTailPipe* do not hold in the initial situation. This is intuitively expected, since we are given no information as to whether these fluents hold in the initial situation and we would like to assume that these do not hold if we can do so in order to minimize the abnormalities.

The second formula shows that the value of the fluent *HeavyPotato* does not change as a result of performing an action. Again, this is expected from a solution to the frame problem, as there are no actions that can change the truth value of the fluent *HeavyPotato*. The third formula follows from the solution of the frame problem for the fluent *EngineRunning* and (7).

The fourth formula shows that if you put potato in the tailpipe and then start the car, the engine does not start. This is because as a result of performing the first action *PutPotato*, the second action *StartCar* becomes qualified.

Without the *ES* and *ESI* axioms, we would not be able to obtain the desired conclusions from our formalization. To see this, let us consider a formalization, N_1 , of our example that differs from N only in that N_1 does not include *ES* and *ESI*. Define a structure M with the universe of situations, $|M|_s$ defined as $\{0\}$, the universe of fluents, $|M|_f$ is defined as $\{PotatoInTailPipe, EngineRunning, HeavyPotato\}$, the universe of actions is defined as $\{StartCar, PutPotato\}$, the interpretation of the predicate *Holds* defined as $Holds^M = \{(PotatoInTailPipe, 0)\}$ and the interpretation of *Result*, $Result^M$ defined as $Result^M(A, 0) = 0$, where A ranges over the actions. It is easy to check that M is a model of N_1 . Hence, we see that $N_1 \not\models \neg Holds(PotatoInTailPipe, S_0)$.

We noted that N entails the formula $\neg Holds(HeavyPotato, S_0)$. Clearly, this formula will not be entailed if T_{obs} includes the formula $Holds(HeavyPotato, S_0)$. This shows that adding an observation does not reduce the set of models of N , but changes them. Hence the solution that we present here does not have the restricted monotonicity property [Lifschitz, 1993b] with respect to observation sentences. In general, we do not expect this property to be obeyed by solutions to the qualification problem.

3 Related Work

McCarthy [McCarthy, 1986] has suggested using circumscription to solve the qualification problem. The use of *Ab* in our formalization is based on his abnormality predicate, but several ideas in our formalization such as the use of the $Holds_R$ predicate and the use of the existence of situations axiom for solving the qualification problem are new. Also, we consider how to combine a solution to the frame problem with a solution to the qualification problem.

In his brief paper on the qualification problem, Elkan [Elkan, 1995] suggests using a context mechanism for formalizing the qualification problem, but does not present any such formalization. He criticizes a particular formalization of the qualification problem based on default logic on the ground that all the potential qualifications to an action need to be made explicit in that formalization. While the criticism is on the mark for the particular formalization, it seems that a formal-

ization in default logic that uses an abnormality predicate will not suffer from this criticism.

Lifschitz [Lifschitz, 1987] formalizes the qualification problem by circumscribing a predicate *precond*. His formalism is less expressive than ours, since it does not permit one to express that an action is qualified if two fluents are simultaneously true. Also, his formalism is embedded within a causal language known to have difficulties with domain constraints³ [Baker, 1991].

Thielscher [Thielscher, 1996] claims that global minimization of abnormalities fails to solve the qualification problem. The example that he uses is the one that we have formalized in the previous section as a NAT. The question that Thielscher raises is as to what would happen if we perform the action *PutPotato* followed by the action *StartCar*. He claims that if we globally minimize abnormalities, there would be a model where the action *PutPotato* does not have its intended effect (because that action is qualified) and hence in that model the engine would start after performing the action *StartCar*.

It should be noted that Thielscher gives only an informal argument for the claim, not a formal proof. His claim seems to be based on the fact that he is considering a propositional language and does not consider abnormalities to have situational arguments. As we have shown, the unintended model does not arise in our formalization.

Thielscher represents qualifications in the form

$$a \text{ disqualified after } [a_1, \dots, a_n],$$

which indicates that after performing the sequence of actions $a_1, \dots, a_n$, action a becomes disqualified. To specify such axioms, it is necessary to precompute what the effect of different sequences of actions should be, which is cumbersome. In contrast, our formalization allows us to specify qualifications directly in terms of fluents. However, Thielscher considers the issue of “miraculous qualifications”, an issue that we do not consider.

McIlraith [McIlraith, 2000] addresses the frame, ramification and qualification problem within the context of the situation calculus. In contrast to the approach presented in this paper, McIlraith compiles a set of effect axioms and ramification constraints into a set of successor state axioms. It should be noted that such a compilation is possible only when the effect axioms and ramification constraints follow a syntactic restriction. (Theories with such a syntactic restriction are called solitary stratified theories in [McIlraith, 2000]). Whether the approach presented in this paper coincides with that in [McIlraith, 2000] when restricted to the case of solitary stratified theories remains open.

4 Conclusions and Future Work

We have presented a simple formalization of the qualification problem based on NATs and have shown how to combine a solution to the frame problem with this formalization. We

³By domain constraints, we mean formulae that relate fluents in the same situation. An example of a domain constraint that we may wish to include in our example is $Holds(PotatoInTailPipe, s) \supset \neg Holds(EngineRunning, s)$.

have also argued that this formalization offers several advantages in comparison with existing approaches.

Theorem 1 pertains only to the example presented in this paper. This raises the question as to how we can evaluate the general applicability of the formalization presented here. One way is by proving theorems about what the formalization does and does not entail. For instance, we can prove that the inner block (that solves the frame problem) generates conclusions identical to the ones sanctioned by a *A*-like [Gelfond and Lifschitz, 1993] language. Evaluating that the outer block handles qualifications correctly is harder. One approach is to test that the formalization yields intuitively expected results on a test suite. A more principled approach is to define a language formalizing intuitions about qualifications and then prove soundness and completeness theorems along the lines of [Kartha, 1993]. Both these approaches are currently being pursued.

One direction for future work is to investigate how the ideas presented in this paper can be generalized to handle domain constraints. In the presence of domain constraints, the situation is more complicated for the following reasons.

- The existence of situations axiom needs to be changed, since now, all possible combination of fluents are not consistent.
- As pointed out by [Ginsberg and Smith, 1988] and [Lin and Reiter, 1994], some domain constraints can act as constraints on the performability of an action (these are called qualification constraints) whereas some other syntactically indistinguishable domain constraints give rise to indirect effects (these are called ramification constraints).

The first of the above complications can be handled, for instance, by including a block that formalizes a nonmonotonic version of the existence of situations axiom, as in [Kartha and Lifschitz, 1994].

The second difficulty is more serious. If all the constraints are ramification constraints, then solutions to the frame problem shown here can be easily extended (along the lines of [Giunchiglia *et al.*, 1997]). If we have both ramification and qualification constraints, we need to separate the set of constraints into these two sets and consider them separately.

Acknowledgment

I wish to thank the anonymous referees for several suggestions for improving the clarity of this paper.

References

- [Baker, 1991] Andrew Baker. Nonmonotonic reasoning in the framework of situation calculus. *Artificial Intelligence*, 49:5–23, 1991.
- [Elkan, 1995] Charles Elkan. On solving the qualification problem. In *Working Notes of the Symposium on Extending Theories of Actions*, 1995.
- [Gelfond and Lifschitz, 1993] Michael Gelfond and Vladimir Lifschitz. Representing action and change by logic programs. *Journal of Logic Programming*, 17:301–322, 1993.
- [Ginsberg and Smith, 1988] Matthew L. Ginsberg and David E. Smith. Reasoning about action I: a possible worlds approach. *Artificial Intelligence*, 35(2):165–195, 1988.
- [Giunchiglia *et al.*, 1997] Enrico Giunchiglia, G. Neelakantan Kartha, and Vladimir Lifschitz. Representing action: indeterminacy and ramifications. *Artificial Intelligence*, 95:409–438, 1997.
- [Kartha and Lifschitz, 1994] G. Neelakantan Kartha and Vladimir Lifschitz. Actions with indirect effects. In Jon Doyle, Erik Sandewall, and Piero Torasso, editors, *Proc. of the Fourth Int'l Conf. on Principles of Knowledge Representation and Reasoning*, pages 341–350, 1994.
- [Kartha and Lifschitz, 1995] G. Neelakantan Kartha and Vladimir Lifschitz. A simple formalization of actions using circumscription. In *Proc. of IJCAI-95*, pages 1970–1975, 1995.
- [Kartha, 1993] G. Neelakantan Kartha. Soundness and completeness theorems for three formalizations of action. In *Proc. of IJCAI-93*, pages 724–729, 1993.
- [Lifschitz, 1987] Vladimir Lifschitz. Formal theories of action (preliminary report). In *Proc. of IJCAI-87*, pages 966–972, 1987.
- [Lifschitz, 1991] Vladimir Lifschitz. Towards a metatheory of action. In James Allen, Richard Fikes, and Erik Sandewall, editors, *Proc. of the Second Int'l Conf. on Principles of Knowledge Representation and Reasoning*, pages 376–386, 1991.
- [Lifschitz, 1993a] Vladimir Lifschitz. Circumscription. In D.M. Gabbay, C.J. Hogger, and J.A. Robinson, editors, *The Handbook of Logic in AI and Logic Programming*, volume 3, pages 297–352. Oxford University Press, 1993.
- [Lifschitz, 1993b] Vladimir Lifschitz. Restricted monotonicity. In *Proc. of the Eleventh National Conference of Artificial Intelligence*, pages 432–437, 1993.
- [Lifschitz, 1995] Vladimir Lifschitz. Nested abnormality theories. *Artificial Intelligence*, 74:351–365, 1995.
- [Lin and Reiter, 1994] Fangzhen Lin and Raymond Reiter. State constraints revisited. *Journal of Logic and Computation, Special Issue on Actions and Processes*, 4(5):655–678, 1994.
- [Lin and Shoham, 1991] Fangzhen Lin and Yoav Shoham. Provably correct theories of action. In *Proc. of the Ninth National Conference of Artificial Intelligence*, pages 349–354, 1991.
- [McCain and Turner, 1995] Norman McCain and Hudson Turner. A causal theory of ramifications and qualifications. In *Proc. of IJCAI-95*, pages 1978–1984, 1995.
- [McCarthy and Hayes, 1969] John McCarthy and Patrick Hayes. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, editors, *Machine Intelligence*, volume 4, pages 463–502. Edinburgh University Press, Edinburgh, 1969.

- [McCarthy, 1977] John McCarthy. Epistemological problems of artificial intelligence. In *Proc. of IJCAI-77*, pages 1038–1044, 1977.
- [McCarthy, 1980] John McCarthy. Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence*, 13(1–2):27–39, 1980.
- [McCarthy, 1986] John McCarthy. Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence*, 28(1):89–116, 1986.
- [McCarthy, 1998] John McCarthy. Elaboration tolerance. In *Fourth Symposium on Logical Formalizations of Commonsense Reasoning*, 1998.
- [McIlraith, 2000] Sheila McIlraith. Integrating actions and state constraints: A closed form solution to the ramification problem (sometimes). *Artificial Intelligence*, 116:87–121, 2000.
- [Sandewall, 1994] Erik Sandewall. *Features and Fluents. The representation of Knowledge about Dynamical Systems*. Oxford University Press, 1994.
- [Shanahan, 1997] Murray Shanahan. *Solving the Frame Problem*. MIT Press, Cambridge, MA, 1997.
- [Thielscher, 1996] Michael Thielscher. Causality and the qualification problem. In L. C. Aiello, J. Doyle, and S. C. Shapiro, editors, *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*, Cambridge, MA, November 1996. Morgan Kaufmann.
- [Thielscher, 1997] Michael Thielscher. Ramification and causality. *Artificial Intelligence*, 89(1–2):317–364, 1997.