Book Review:
M. Shanahan, *Solving the Frame Problem*

Vladimir Lifschitz
Department of Computer Sciences
University of Texas at Austin
Austin, TX 78712, USA

Murray Shanahan has been actively involved in research on the frame problem since the late eighties. When a scientist looks back at the past of his field, including his own early work, he can see how poorly the subject was understood years ago, and then two attitudes are possible. One is to think: It's amazing how naive and confused we used to be; much of this work should have never appeared in print, and the rest should have been expressed very differently! Or he can adopt a wiser view of the history of his field as a reasonable—maybe necessary—path towards the better understanding that is available today.

This respect for the past, with all its embarassing mistakes, is an attractive quality of Shanahan's book. He believes that the history of attempts to solve the frame problem

> isn't an arbitrary record of failures which culminates in the discovery of a truth which, in hindsight, seems obvious. Rather, it seems to me to reflect a natural order of enquiry. The dead ends that people have investigated are much more than mistakes. They're obvious possibilities, which any thinker engaged in studying the frame problem is wont to investigate. A deep understanding of the knowledge representation issues surrounding the frame problem can only be acquired by getting a feel for the space of mathematical possibilities and where they lead. [p. *xv*]

The book is about the history of research on the frame problem as much as about our current state of knowledge. This may be the only way a good book can be written today on a subject as volatile as the theory of commonsense knowledge and reasoning.

The book is not a comprehensive survey of the history of the problem: it is mostly about the work done within the "Stanford/Imperial axis," that is to say, about the approaches that use the situation calculus and circumscription (John McCarthy) and the event calculus and logic programming (Bob Kowalski). There are also good pages on several other directions of work, including default logic, the explanation closure method, and action languages.

In the introductory chapters, after discussing the role of knowledge representation in AI and the role of logic in knowledge representation, the frame problem is described as a problem that arises when we attempt to represent properties of actions using logic. In addition to specifying the changes caused by performing a particular kind of action,

> we also have to describe what does *not* change. Otherwise, we find that we cannot use the description to draw any useful conclusions... The frame problem is the problem of constructing a formal framework that enables us to do just this. [p. 1]

The first half of the book is mostly about developing such a framework using circumscription. The second half stresses the ideas that come from logic programming.

Given Shanahan's attention to the history of the subject, it is not surprising that *Yale shooting scenario* is the longest entry in his index. Yale shooting is not mentioned very often in today's research on actions. After all, the example is quite elementary, it has no features that are of interest to us today: no indirect effects, no implicit preconditions, no nondeterminism, no concurrency, no sensing, no complex actions, no continuous time, no counterfactuals. But it remains the most memorable of the past embarassments of the commonsense reasoning community.

This counterexample demonstrated that the first serious attempt to apply circumscription to the frame problem [11] was unsatisfactory. The discussion of possible ways to overcome this difficulty has led the authors of the example to drastic conclusions:

> This problem is shown to be independent of the logic used; nor does it depend on any particular temporal representation. Upon analyzing the failure we find that the nonmonotonic logics we considered are inherently incapable of representing this kind of default reasoning. [7, p. 379]

This seemed particularly convincing because the senior author of the paper, Drew McDermott, was known as one of the creators of formal nonmonotonic reasoning, an enthusiastic supporter of the logic approach to AI, and a co-author of an influential textbook written from the "logicist" position [2]. After the discovery of the shooting example, he changed his views. His other paper written around the same time "draws mostly negative conclusions about... the progress we can expect" in logic-based AI [13].

The reaction of many other "logicists" was opposite: they saw the shooting example as an inspiring challenge. The main lesson they learned from that example is that in logic-based AI, as in any other area of science, mathematics should be taken seriously. The claim in [11] which Hanks and McDermott found to be incorrect ("The following set of situation calculus axioms solves the frame problem") is not proved there. This is understandable: describing the effect of the circumscription proposed in [11] is a difficult mathematical problem. But the question is not even mentioned there as a topic for future research. Interestingly, the discussion of the counterexample in [7] does not include a satisfactory mathematical proof either: the authors restrict attention to the finitely many situations that are explicitly mentioned in the scenario and disregard the existence of infinitely many other situations definable in the language. (This can be easily fixed.)

When mathematicians are lucky, the concepts that they define behave exactly as intuitively expected. Usually they are not so lucky. The Taylor series for a function f may converge to a function that is different from f. Mathematicians have learned to live with facts like this, and they even find them useful. Because of the differences between what is expected and what is true, they make sure to verify all their conjectures, no matter how plausible, by reasoning. It is crucial for further progress that students of commonsense knowledge be prepared to do difficult mathematics.

There is a lot of mathematics in Shanahan's book, most of it about the semantics and properties of nonmonotonic formalisms. The invention of these formalisms was motivated to a large degree by the need to solve the frame problem, and the book can serve as an introduction to formal nonmonotonic reasoning. All terminology and notation are carefully defined, beginning with those related to the language of predicate calculus. Some particularly long proofs are relegated to appendices.

The second lesson of Yale shooting is that not all nonmonotonic formalisms are created equal. Researchers at Imperial were not impressed by the shooting example as much as those at Stanford. Indeed, the logic programming solution to the frame problem described by Shanahan in Section 12.1, which handles the Yale shooting scenario correctly, is not a special trick invented in response to Hanks and McDermott; it is based on what the author remembered from the undergraduate lectures Kowalski gave back in 1981, long before the invention of the counterexample. For Kowalski and his students, Yale shooting was not even a challenge.

Default logic and autoepistemic logic are similar in this sense to logic programming. Each of these systems can be viewed as an extension of the language of Prolog with negation as failure [1, 4], so that the logic programming formalization mentioned above can be easily restated in terms of either formalism. An autoepistemic equivalent of Kowalski's logic program was actually rediscovered in [5]. Moreover, as we learned from [15], the "frame default" from Reiter's 1980 paper in which default logic was originally introduced [14] works correctly—the claims to the contrary by Hanks and McDermott notwithstanding—provided that the rest of the default theory is set up the right way.

Reasoning involved in the Yale shooting story is an example of "temporal projection"—determining the effect of a given sequence of actions executed in a given initial situation. A large part of the book under review is about more complex forms of temporal reasoning, including those that involve incomplete narratives and concurrency.

It is unfortunate that two important lines of research are not mentioned in the book at all: nonmonotonic causal logics and satisfiability planning. The seminal papers by Geffner [3] and by Kautz and Selman [8] are not even included in the bibliography. Perhaps the significance of these papers was not as clear when Shanahan's manuscript was being finalized as it is today.

Geffner's work has led to several interesting solutions to the frame problem. Some of them use circumscription [9, 6], and this circumscription is of a very simple kind. The formulation given in [10] is "rule-based"; it is similar (and closely related) to both logic programming and default logic. The solutions to the frame problem proposed in these three papers would not be out of place in a book that emphasises circumscription and logic programming. The first of the three is actually mentioned by Shanahan, but not discussed in any detail.

The ideas of Kautz and Selman have led to the creation of powerful systems for reasoning about actions and planning, such as BLACKBOX[1] and CCALC[2]. The former is a very efficient (by today's standards) planner that

[1] `http://www.research.att.com/~kautz/blackbox` .

[2] `http://www.cs.utexas.edu/users/mccain/cc` .

accepts a description of actions in the language of STRIPS, designed by "logicists" Fikes and Nilsson, and uses propositional solvers for search. The latter uses a search mechanism of the same kind (although somewhat less sophisticated) and can perform, besides planning, many forms of temporal reasoning discussed in Shanahan's book. Its input language is the language from [10], considerably more expressive than STRIPS. The reader should have a look at these systems if he believes that "logicists" only prove theorems but never write programs. These days, they do both.

In spite of these omissions, *Solving the Frame Problem* is an important contribution to the literature on logical AI. It is clearly written, intellectually stimulating and enjoyable to read.

References

[1] Nicole Bidoit and Christine Froidevaux. Minimalism subsumes default logic and circumscription. In *Proc. LICS-87*, pages 89–97, 1987.

[2] Eugene Charniak and Drew McDermott. *Introduction to Artificial Intelligence*. Addison-Wesley, 1985.

[3] Hector Geffner. Causal theories for nonmonotonic reasoning. In *Proc. AAAI-90*, pages 524–530, 1990.

[4] Michael Gelfond. On stratified autoepistemic theories. In *Proc. AAAI-87*, pages 207–211, 1987.

[5] Michael Gelfond. Autoepistemic logic and formalization of commonsense reasoning. In Michael Reinfrank, Johan de Kleer, Matthew Ginsberg, and Erik Sandewall, editors, *Non-Monotonic Reasoning: 2nd International Workshop (Lecture Notes in Artificial Intelligence 346)*, pages 176–186. Springer-Verlag, 1989.

[6] Joakim Gustaffson and Patrick Doherty. Embracing occlusion in specifying the indirect effects of actions. In *Principles of Knowledge Representation and Reasoning: Proc. Fifth Int'l Conference*, pages 87–98, 1996.

[7] Steve Hanks and Drew McDermott. Nonmonotonic logic and temporal projection. *Artificial Intelligence*, 33(3):379–412, 1987.

[8] Henry Kautz and Bart Selman. Planning as satisfiability. In *Proc. ECAI-92*, pages 359–363, 1992.

[9] Fangzhen Lin. Embracing causality in specifying the indirect effects of actions. In *Proc. IJCAI-95*, pages 1985–1991, 1995.

[10] Norman McCain and Hudson Turner. Causal theories of action and change. In *Proc. AAAI-97*, pages 460–465, 1997.

[11] John McCarthy. Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence*, 26(3):89–116, 1986. Reproduced in [12].

[12] John McCarthy. *Formalizing Common Sense: Papers by John McCarthy.* Ablex, Norwood, NJ, 1990.

[13] Drew McDermott. A critique of pure reason. *Computational Intelligence*, 3:151–160, 1987.

[14] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[15] Hudson Turner. Representing actions in logic programs and default theories: a situation calculus approach. *Journal of Logic Programming*, 31:245–298, 1997.