The Frame Problem, Then and Now

Vladimir Lifschitz
University of Texas at Austin

My collaboration with John McCarthy started in 1984. He was interested then in what he called the commonsense law of inertia. That idea is related to actions, such as moving an object to a different location, or, for instance, toggling a light switch. According to this law, whatever we know about the state of affairs before executing an action can be presumed, by default, to hold after the action as well. Formalizing this default would resolve the difficulty known in AI as the frame problem.

Reasoning with defaults is nonmonotonic. John proposed a solution to the frame problem based on the method of nonmonotonic reasoning that he called circumscription [1].

And then something unpleasant happened. Two researchers from Yale University discovered that John's proposed solution was incorrect [2]. Their counterexample involved the actions of loading a gun and shooting, and it became known as the "Yale Shooting Scenario." Twenty years later their paper received the AAAI Classic Paper Award [3].

But that was not all. Automated reasoning is notoriously difficult, and the presence of defaults adds yet another level of complexity. Even assuming that the problem with Yale Shooting is resolved, was there any hope, one could ask, that the commonsense law of inertia would ever become part of usable software?

So John's proposal seemed unsound and non-implementable. It also seemed unnecessary, because other researchers have proposed approaches to the frame problem that did not require nonmonotonic reasoning [4, 5, 6]. There were all indications that his project just wouldn't fly.

But history showed otherwise. It was destined to fly, and, in fact, to fly quite high: in outer space. I'd like to tell you about a program written years later by a group of computer scientists who continued John's research on nonmonotonic reasoning in collaboration with engineers from United Space

Alliance—the company that was responsible for the day-to-day management of the Space Shuttle fleet. The program is called the RCS Advisor [7]. The RCS, or Reaction Control System, was the system aboard the shuttle designed to maneuver it while it was in space. The RCS Advisor was used to verify the possibility of doing that even if several elements of the system malfunction. And that program incorporated a formalization of the commonsense law of inertia. How was this possible in spite of the difficulties that we talked about? First of all, what about the Yale Shooting Problem?

The answer to this question is that simple ways to repair John's original formalization have been found. Some ideas came from experience with the programming language Prolog [8]. Available solutions look so straightforward that it's not easy to explain to students today why the Yale Shooting Scenario attracted so much attention twenty five years ago.

But what about the difficulty of implementing nonmonotonic reasoning?

We have today something that was not available in the 1980s: fast satisfiability solvers for propositional logic [9]. Propositional logic is monotonic, but ideas used in the design of SAT solvers can be applied to nonmonotonic languages also [10, 11]. These languages are closely related to the language of circumscription [12].

But why didn't the creators of the RCS Advisor use simpler, monotonic solutions to the frame problem?

There was a good reason for that. The RCS was a complicated device, and the effects of actions, such as flipping a switch, had to be described in two steps. First, the simple direct effect was stated: when you flip the switch, the state of the switch changes. Then the other effects would logically follow using the rules describing the RCS that were included in the program. Such two-level descriptions of actions become possible when the nonmonotonic approach to the frame problem is adopted [13].

This example shows that John's theory of nonmonotonic reasoning is not only interesting philosophy and beautiful mathematics; it is also computer science with serious applications.

References

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

[2] Steve Hanks and Drew McDermott. Default reasoning, nonmonotonic logics, and the frame problem. In *Proceedings of National Conference on Artificial Intelligence (AAAI)*, 1986.

[3] Tom Mitchell and Hector Levesque. The 2005 AAAI classic paper awards. *AI Magazine*, 26(4), 2006.

[4] Richard Fikes and Nils Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(3–4):189–208, 1971.

[5] Edwin Pednault. ADL and the state-transition model of action. *Journal of Logic and Computation*, 4:467–512, 1994.

[6] Raymond 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, 1991.

[7] Monica Nogueira, Marcello Balduccini, Michael Gelfond, Richard Watson, and Matthew Barry. An A-Prolog decision support system for the Space Shuttle. In *Proceedings of International Symposium on Practical Aspects of Declarative Languages (PADL)*, pages 169–183, 2001.

[8] Chris Evans. Negation-as-failure as an approach to the Hanks and McDermott problem. In *Proceedings of Second International Symposium on Artificial Intelligence*, 1989.

[9] Carla P. Gomes, Henry Kautz, Ashish Sabharwal, and Bart Selman. Satisfiability solvers. In Frank van Harmelen, Vladimir Lifschitz, and Bruce Porter, editors, *Handbook of Knowledge Representation*, pages 89–134. Elsevier, 2008.

[10] Ilkka Niemelä and Patrik Simons. Smodels—an implementation of the stable model and well-founded semanics for normal logic programs. In *Proceedings 4th Int'l Conference on Logic Programming and Nonmonotonic Reasoning (Lecture Notes in Artificial Intelligence 1265)*, pages 420–429. Springer, 1997.

[11] Nicola Leone, Gerald Pfeifer, Wolfgang Faber, Thomas Eiter, Georg Gottlob, Simona Perri, and Francesco Scarcello. The DLV system for knowledge representation and reasoning. *ACM Transactions on Computational Logic*, 7(3):499–562, 2006.

[12] Paolo Ferraris, Joohyung Lee, and Vladimir Lifschitz. Stable models and circumscription. *Artificial Intelligence*, 175:236–263, 2011.

[13] Fangzhen Lin. Embracing causality in specifying the indirect effects of actions. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1985–1991, 1995.