In Memoriam: Raymond Reiter

June 12, 1939–September 16, 2002

Jack Minker

■ Raymond Reiter, a professor of computer science at the University of Toronto, a fellow of the Royal Society of Canada, and winner of the International Joint Conference on Artificial Intelligence 1993 Outstanding Research Scientist Award, died September 16, 2002, after a year-long struggle with cancer. Reiter, known throughout the world as "Ray," made foundational contributions to artificial intelligence, knowledge representation and databases, and theorem proving.

Raymond Reiter, a professor of computer science at the University of Toronto, a fellow of the Royal Society of Canada, and winner of the International Joint Conference on Artificial Intelligence (IJCAI) 1993 Outstanding Research Scientist Award, died September 16, 2002, after a year-long struggle with cancer. Reiter, known throughout the world as "Ray," made foundational contributions to artificial intelligence, knowledge representation and databases, and theorem proving.

Ray was born in Toronto, Canada, in 1939 to immigrant parents who came from Poland. He received a B.S. in mathematics from the University of Toronto in 1961 and an M.S. degree in mathematics in 1963 from the University of Toronto. He received a Ph.D. in 1967 from the University of Michigan. His dissertation was entitled, "A Study of a Model for Parallel Computations." His dissertation adviser was Harvey Garner, but he was also motivated by Dick Karp who was then on leave from IBM.

Ray wrote many seminal articles, some of which are described later. He also coedited two books (Brachman, Levesque, and Reiter 1989; Brachman, Levesque, and Reiter, 1991) and published a book on dynamic systems (Reiter 2001), discussed later. He served the scientific community by being the program chair or cochair of important conferences and workshops and as an editor, or on the editorial board of journals such as the *Theory and Practice of Logic Programs*. He was a fellow of the Association for Computing Machinery and the American Association for Artificial Intelligence.

Ray Reiter—The Person

Ray was gifted with an analytic mind, an adventurous spirit, and a generous heart. He had a wide range of interests and was accomplished in areas outside computer science. He was a lepidopterist who loved butterflies and moths and published a paper in a journal on the subject. He was interested in literature and music. He counted among his friends several writers. He loved classical music, especially that of Wagner, and sometimes traveled long distances to listen to operas such as those that comprise *Der Ring des Nibelungen*. He and his friend Richard Rosenberg drove to Seattle in 1980 to see the fourth part of the *Ring, Gotterdammerung* (*The Twilight of the Gods*).[1] He was passionate about social justice and the well-being of the planet.

Although he was an interesting conversationalist, Ray was basically a modest person who neither talked much about himself nor boasted about his research. He had different classes of friends and somehow kept them separate. There were his personal friends, the lepidopterists, the writers, and the artificial intelligence group. I believe that Ray understood that his research was of high quality, but he did not seem to value his research as highly as others, such as I, did. When I sent Ray a message to tell him that I was nominating him for the IJCAI award and wanted a copy of his vitae, he demurred and said that I was wasting my time since he would never receive the award. I

Raymond Reiter, June 12, 1939–September 16, 2002

responded that I did not ask for his opinion but for his vitae and that I was going to nominate him with or without his cooperation. He was so confident that he would not win that he offered to take me to the best restaurant in France if he won. True to his word, we celebrated the award at Paul Bocuse's three-star restaurant in Lyons a few days before the presentation of the award.

Ray, like many of us, thought highly of John McCarthy. Ray has been heard to say of John McCarthy that, "no one has the right to be that smart!" As Alan Mackworth has pointed out, the same could be said of Ray, although in "typically Canadian" fashion, Ray tended to undervalue his achievements and talents.

One little-known aspect of Ray is that he was a lepidopterist—he searched for butterflies and moths in the tropical rain forests in such countries as Nepal, New Guinea, Malaysia, Indonesia, and Borneo. He returned from these trips with exotic butterflies and moths that he had caught. He prepared the butterflies and moths himself and filed them in his large collection.

He also went to some of these places to frolic with orangutans and meet primitive cultures. Preceding the 1991 IJCAI conference, he was in Borneo. Heavy rains and sodden landing strips forced the cancellation of a missionary flight that was to pick him up for the trip back down river and then to Australia (where he was program cochair). At great expense, he ended up having to charter a helicopter to bring him back to civilization in time for the opening of IJCAI.

In addition to his adventures in the rain forests, he traveled through Europe on many occasions. On one trip with his boyhood friend Richard Rosenberg and Rich's then-wife Avis, they toured Eastern Europe. They traveled to Rumania, Hungary, Bulgaria, and Czechoslovakia in August 1968. Their trip, although memorable, might not have been pleasant. As noted by Avis Lang, Rich Rosenberg's ex-wife, Ray gave away his copy of an Aleksandr Solzhenitsyn novel to a well-dressed woman somewhere in Rumania, thus presumably contributing to the cache of samizdat that sustained Eastern Europe's intellectuals during the era of Soviet censorship. As they were driving toward Prague on the trip, they frequently encountered military convoys. They aborted their trip and went to Austria. They learned that the Russian army had entered Prague in force shortly after they had departed.

Rich Rosenberg says that Ray was a typical young boy. Rich states,

> Both Ray and I belonged to a religious Zionist youth group for a couple of years when we were about 15. But the primary reasons were neither religion nor Israel—but girls. Against the understood rules of the organization, we held dances and actually did slow dancing. After a while, the religious stuff became too oppressive, and he quit.

Although he was Jewish, and was bar mitzvahed, he was not a religious person. However, Ray's older brother, Jack, said that Ray was outstanding in religious school and read a large section of the Torah portion of the service for his bar mitzvah. His rabbi was so impressed with Ray's talmudic abilities that he offered to pay for his studies if he went to rabbinical school and became a rabbi. Although I am certain Ray would have been a wonderful biblical scholar, and a compassionate rabbi, I am not sure I can envision him in that role. The religious community's loss was the gain of the AI community.

Ray spent his first year in undergraduate school at the University of Toronto in engineering physics before transferring to the arts faculty in mathematics, physics, and chemistry, with a strong interest in applied mathematics. Rich Rosenberg believes that this was an early indication of his resistance to hacks and his strong commitment to formal approaches.

He was given a fellowship to study for a Ph.D. in physics; however, on his way to accepting the fellowship, he changed his mind when he realized that he would prefer to study computer science. He thought he could make a greater mark in computer science than he could in physics.

As a graduate student at the University of Michigan, Ray became fascinated with motorcycles; he and his fellow students John Seely Brown and Abbe Mowshowitz bought large BMW touring bikes. He made trips in Europe and the United States on his BMW.

Ray was also socially conscious of events in the world. While he was a graduate student, the United States was engaged in Vietnam. As a Canadian at a United States university, he was not inhibited from participating on occasions in marches, picketing draft boards, and taking part in sit-ins. He was concerned with social issues all his life and was upset by the Israel-Palestine controversy.

Ray was fundamentally a night person. His graduate students knew not to look for him until some time after 2:00 PM. He prowled around at night and preferred New York and Toronto to Vancouver because he could get a cup of coffee anywhere at 4:00 AM in the big cities. Alan Mackworth speculated that

> Ray was a machine for turning caffeine and nicotine into theorems. To be cut off from this feed stock during the height of his productive hours would be distressing.

Ray was demanding as an adviser. One of his former students, Iluju Kiringa,[2] relates the following:

> I must say right away that Ray was intellectually very challenging. He was not the kind of supervisor that would set up regular meetings where a student would come in, receive instructions on what to work on, walk out, go on to follow these instructions, and then come back to receive further instructions at the next scheduled meetings. To see Ray for the first time for talking about a possible thesis or project topic, you'd better have at least one or two solid topics prepared in a well articulated way, with a convincing example that shows that your ideas are credible. Then you should give him a lecture for roughly half an hour. After that, he would stand up, go to his white board, and show you

all the flaws in your approach. He would almost easily come up with an alternative way of viewing the same concepts, but a way that is more elegant, theoretically sound—without such a soundness, he would not listen to you—and, above all, simple. Ray loved simple theories that you could play with only in your head, without going back to complicated written formulas. He very often said that such theories are those with the most fruitful and intuitive consequences.

Ray Reiter—The Scientist

Ray's research was concerned with the formal foundations of knowledge representation and databases and reasoning in AI. His work was motivated by specific practical problems in AI, which, in turn, drove the theoretical results. He isolated problems and techniques that arose in different application areas in AI, formalized, and, where possible, generalized them, then explored what these theories had to say about the applications that motivated them and about knowledge representation in general. It is clear that his overall objective was to determine unifying reasoning patterns that cut across application domains. From conversations with him, this reflected his belief that a science of AI was possible and that one way to achieve it was by isolating these patterns and studying their formal properties. He was fundamentally looking to break new ground rather than prove a result just because it could be proven. It was this probing of what things do mean, why they do seem to work when we do things in a certain way, and how we can explain a phenomenon in a coherent fashion that distinguished his work. It appears to me that he brought a Talmudic approach to his work that he learned studying for his bar mitzvah. Before he explored an area, he studied the literature of a subject, understood it thoroughly, and tried to explain why things follow from our readings. He was not interested in any problem but in ones that would shed light on a class of problems. Dave Etherington states that Reiter's maxim was, "Just because you can prove it, that doesn't mean that it is interesting."

His research covered a wide range of areas in AI: nonmonotonic reasoning, knowledge representation and databases, logic programming, truth maintenance systems, diagnostic reasoning, computational vision, and representation and reasoning for dynamic worlds.

Nonmonotonic Reasoning

Ray, together with John McCarthy, Drew McDermott, and Jon Doyle, was one of the founders of the field of nonmonotonic reasoning. I believe that this is a major accomplishment that computer scientists have made to the field of mathematical logic. His 1978 papers on closed-world reasoning (Reiter 1978b) and reasoning by default (Reiter 1978c) were among the first to deal with a formal treatment of nonmonotonic reasoning. The work derives from his observations about AI programming languages (PLANNER and PROLOG), databases, the frame problem, and natural language processing that all deal with default situations. In his 1978 paper on default reasoning, he argues that the underlying notion is "...in the absence of evidence to the contrary, assume...." His work on default reasoning was finalized for his article (Reiter 1980) in the seminal 1980 issue of the *AI Journal* that also contained articles by McDermott and Doyle and by McCarthy. This work was supplemented with papers cowritten with Etherington (Etherington and Reiter 1983), which initiated a large body of research on nonmonotonic formalisms for inheritance; with Criscuolo (Reiter and Criscuolo 1983), which described a variety of problematic settings for default reasoning and provided a number of standard benchmark examples for the field; and with Bertossi (Bertossi and Reiter 1992), which provides a circumscriptive characterization of generic objects in geometry. The paper with Bertossi promises a general theory of genericity in mathematics based on circumscription and is one of the first examples of an AI theory with a nontrivial application to mathematics and mathematical logic. The work on nonmonotonic reasoning spawned a whole area of research and biannual conferences on nonmonotonic reasoning and logic programming. It has been shown that a large portion of Reiter's default logic can be translated into logic programs. Ray's "tweety" example used to illustrate the theory has become ubiquitous in nonmonotonic reasoning.

Knowledge Representation and Databases

A significant part of his research in knowledge representation was to generalize databases and develop a theory of deductive databases that includes reasoning capabilities. He was the first to have provided an axiomatization for relational databases and their deductive generalizations (Reiter 1984). This has now become the standard specification for the research community in deductive databases. His most important contributions in AI and databases are as follows: In 1978, he formulated the closed world assumption (CWA) (Reiter 1978b) at the same time Keith Clark developed the Clark completion theory (Clark 1978). These were the first formalizations of the concept of negation in deductive databases. He investigated some of the formal properties of the CWA. This paper was very influential and led to a considerable body of research involving generalizations of the Reiter CWA and its implications for nonmonotonic reasoning and logic programming. Indeed, my work on the generalized closed-world assumption (GCWA) (Minker 1982) was influenced by Reiter's paper. In 1978, he also provided the first proposal for compiling the deductive rules of a database (Reiter 1978a) to obtain an efficient implementation of a deductive system. In 1984, he provided the first formal account of the concept of an integrity constraint (Reiter 1984) and subsequently proposed a more radical approach that appeals to an autoepistemic interpretation of constraints. In 1992, he developed a theory of database updates based on the situation calculus (Reiter 1992). This work borrows from a variety of AI planning ideas in the situation calculus to yield an approach to the formalization of database updates. This is a very nice

example where problems and their solutions from "classical" AI—in this case, the frame problem (Reiter 1991)—have applications to problems outside the field.

Logic Programming

He made several important contributions to logic programming. In 1971, he independently formulated and proved the completeness of the SL (linear resolution with selection function) resolution procedure, developed by Robert Kowalski and Donald Kuehner (Kowalski and Kuehner 1971), which he called the *clause-ordered linear resolution strategy,* that forms the procedural basis for logic programming interpreters (Reiter 1971). His work on the CWA provided an early semantics for PROLOG's negation-as-failure operator. His papers on nonmonotonic reasoning (Reiter 1987a) emphasized both the nonmonotonic character of the negation-as-failure operators of PLANNER and PROLOG and the potential uses of such programming languages for implementing nonmonotonic reasoning systems. This insight provided the motivation for a large body of research on nonmonotonicity and logic programming, including the aforementioned biannual conference on the subject.

Truth Maintenance Systems

Together with Johann de Kleer, he gave the first theoretical foundations for assumption-based truth maintenance systems (de Kleer and Reiter 1987). This revealed their intimate relationship to abductive reasoning, which explains why these systems have enjoyed such widespread applications.

Diagnostic Reasoning

Ray provided the first formalization, in an extremely general setting, of the task of diagnosing faulty systems (Reiter 1987b). This formalism has now become the standard theory of the diagnostic reasoning community.

Computational Vision

Together with Alan Mackworth, he provided the first formal account of high-level image interpretation for computational vision (Reiter and Mackworth 1989). This allowed, for the first time, the integration of background knowledge about the scene domain into the image interpretation process and provided an account of how this knowledge conditions the interpretations computed.

Representation and Reasoning for Dynamic Worlds

Ray's work over the past 10 years focused on problems of representation and reasoning for dynamic worlds. He worked with a number of individuals during this time, some of whom are Hector Levesque, Fangzhen Lin, and Fiora Pirri. He believed that although AI does have various theories of time and change, none of these is sufficiently rich to express everything one wants to say about changing worlds. Together with Hector Levesque, he believed that the situation calculus developed by McCarthy (1963) in 1963 for dynamic worlds had more potential than was commonly believed. They believed that the situation calculus was not only a useful theoretical formalism but could lead to efficient implementations. Together with Levesque and a group at the University of Toronto, he succeeded in achieving many of his objectives. Dynamic systems deal with robotics, databases, software agents, simulation, decision and control theory, computer animation, and disciplines that involve implementing systems that evolve over time. His 2001 book, *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems,* provides a comprehensive exposition of the accomplishments he achieved over the past 10 years. As noted by Reiter in his closing remarks to the book,

> It is remarkable how far we have managed to come with the simple ontology provided by the situation calculus. The only ingredients needed were actions, fluents, and the ability to construct finite action sequences…. With only this minimal ontology, we have developed situation calculus accounts for time, concurrency, procedures, exogenous events, reactivity, sensing and knowledge, probabilistic uncertainty, and decision theory.

Until his untimely death, he was investigating how easily this ontology could be expanded to other aspects of commonsense reasoning. There is no doubt that with the work described in this book, Ray has opened up another important area that will serve as a source of research for many years to come.

Summary

It is clear that Ray made major contributions over a period of approximately 35 years to the field of AI. He has led the way in the formalization of default and nonmonotonic reasoning, knowledge representation and deductive databases, diagnostic reasoning, computational vision, and reasoning in dynamic worlds. He was in the forefront of those who are making AI a science. As Hector Levesque and I noted in talking about Ray after his death, unlike other fields, AI generally has not named important results after the individuals who first recognized them. In this context, Hector and I believe that when writing about the CWA, we should refer to it as Reiter's closed world assumption, and similarly, default logic should be referred to as Reiter's default logic.

His former student David Etherington has captured a great part of the essence of Ray. David stated,

> There seem to be two sides to Ray, one that his peers see and another that he reveals only to his friends.... At first they seem incongruous—like the juxtaposition of the logician and the jungle trekker; reflection shows that his sense of adventure, his desire to explore ahead of the pack and to push on the frontiers, reconcile them.

I had the distinction to nominate Ray for the IJCAI Research Excellence Award. In my remarks introducing Ray's lecture and award presentation I stated, "I believe that Ray deserves to be in the company of the past award winners: John McCarthy, Alan Newell, and Marvin Minsky." Based on Ray's accomplishments before and after he received the award, there is no doubt that my assessment was accurate.

I have had the singular pleasure of

knowing Ray for approximately 30 years. He was an esteemed colleague and a close and dear friend who influenced the direction of my research. I, and all who knew him, will miss his intelligent comments, his adventurous spirit, and his concern for humanity.

Acknowledgments

Much of this memorial is taken from unpublished remarks I made on introducing Ray Reiter on the occasion of his receiving the Research Excellence Award at the IJCAI-1993 conference. To gain a better perspective of Ray as a person, I solicited comments from some of his former students, Dave Etherington and Iluju Kiringa; from his colleagues, Hector Levesque, Fiora Pirri, and Richard Rosenberg; from his friend, Avis Lang; and from his brother, Jack Reiter. A color photograph of Raymond Reiter can be found on the World Wide Web at prism.cs.umd.edu/papers/Min02:reitermemoriam.html.

This memorial was requested by the *Theory and Practice of Logic Programming* and appears in TPLP 3(1): 1–9 (© Cambridge University Press). It is reproduced with their permission.

Notes

1. See Richard Rosenberg's memoir, "Ray Reiter—A Memoir." In Reiter (2001). Rosenberg, a friend of Ray's for 50 years lovingly describes Ray through these 50 years.

2. Personal message from Dr. Kiringa to Jack Minker containing a memorial he wrote on behalf of Ray's students and forwarded to me by e-mail on September 18, 2002.

References

Brachman, R.; Levesque, H.; and Reiter, R. eds. 1991. *Knowledge Representation*. Cambridge, Mass.: MIT Press.

Brachman, R.; Levesque, H.; and Reiter, R., eds. 1989. *First International Conference on Principles of Knowledge Representation and Reasoning*. San Francisco: Morgan Kaufmann.

Bertossi, L., and Reiter, R. 1992. Circumscription and Generic Mathematical Objects. Paper presented at the Fourth International Workshop on Nonmonotonic Reasoning, 28–31 May, Plymouth, Vermont.

Clark, K. L. 1978. Negation as Failure. In *Logic and Data Bases*, eds. H. Gallaire and J. Minker, 293–322. New York: Plenum Press.

de Kleer, J., and Reiter, R. 1987. Foundations for Assumption-Based Truth Maintenance Systems: Preliminary Report. In Proceedings of the Sixth National Conference on Artificial Intelligence, 183–188. Menlo Park, Calif.: American Association for Artificial Intelligence.

Etherington, D. E., and Reiter, R. 1983. On Inheritance Hierarchies with Exceptions. In Proceedings of the Third National Conference on Artificial Intelligence, 104–108. Menlo Park, Calif.: American Association for Artificial Intelligence.

Kowalski, R. A., and Kuehner, D. 1971. Linear Resolution with Selection Function. *Artificial Intelligence* 2(3–4): 227–260.

McCarthy, J. 1963. Situations, Actions, and Causal Laws. In *Semantic Information Processing,* ed. M. Minsky, 410–417. Cambridge, Mass.: MIT Press.

Minker, J. 1982. On Indefinite Databases and the Closed World Assumption. In *Proceedings of the Sixth Conference on Automated Deduction,* 292–308. Lecture Notes in Computer Science 138. Berlin: Springer-Verlag.

Reiter, R. 2001. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamic Systems*. Cambridge, Mass.: The MIT Press.

Reiter, R. 1992. On Formalizing Database Updates: Preliminary Report. Paper presented at the Third International Conference on Extending Database Technology (invited paper), 23–27 March, Vienna, Austria.

Reiter, R. 1991. The Frame Problem in the Situation Calculus: A Simple Solution (Sometimes) and a Completeness Result for Goal Regression. In *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, ed. V. Lifschitz, 359–380. San Diego, Calif.: Academic Press.

Reiter, R 1987a. A Theory of Diagnosis from First Principles. *Artificial Intelligence* 32(1): 57–95.

Reiter, R. 1987b. Nonmonotonic Reasoning. In *Annual Review of Computer Science*, *Volume 2,* 147–186. Palo Alto, Calif.: Annual Reviews.

Reiter, R. 1984. Toward a Logical Reconstruction of Relational Database Theory. In *Data Modeling in Artificial Intelligence, Database, and Programming Languages*, eds. M. L. Brodie, J. Mylopoulos, and J. W. Schmidt, 191–233. Berlin: Springer-Verlag.

Reiter, R., 1980. A Logic for Default Reasoning. *Artificial Intelligence* 13(1–2): 81–132.

Reiter, R. 1978a. Deductive Question Answering on Relational Databases. In *Logic and Databases*, eds. H. Gallaire and J. Minker, 149–177. New York: Plenum Press.

Reiter, R., 1978b. On Closed World Data Bases. In *Logic and Data Bases*, eds. H. Gallaire and J. Minker, 55–76. New York: Plenum Press.

Reiter, R. 1978c. On Reasoning by Default. In Proceedings of Theoretical Issues in Natural Language Processing II, 20. New York: Association for Computing Machinery.

Reiter, R. 1971. Two Results on Ordering for Resolution with Merging and Linear Format. *Journal of the ACM* 18(1): 630–646.

Reiter, R., and Criscuolo, G. 1983. Some Representational Issues in Default Reasoning. *International Journal Computers and Math* (Special Issue on Computational Linguistics) 9(1): 15–27.

Reiter, R., and Mackworth, A. K. 1989. A Logical Framework for Depiction and Image Interpretation. *Artificial Intelligence* 32(2): 125–155.

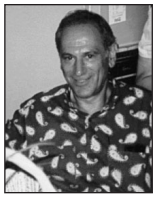

Jack Minker is professor emeritus, Department of Computer Science and the Institute for Advanced Computer Studies, University of Maryland. His research areas are deductive databases, logic programming, AI, and nonmonotonic reasoning. He was the first chairman of the Department of Computer Science at Maryland (1974–1979) and chairman, Advisory Committee on Computing to the NSF (1979–1982). In 1985, Minker received the Association for Computing Machinery's (ACM's) Outstanding Contribution Award for his work in human rights. He is a fellow of the AAAS, a founding fellow of the AAAI, a fellow of the IEEE, and a founding fellow of the ACM. He received the University of Maryland Presidential Medal for 1996 and is a Distinguished Scholar-Teacher for 1997–1998. His e-mail address is minker@cs.umd.edu.