Received: 22 July 2022 | Accepted: 28 September 2022

DOI: 10.1002/aaai.12072

IN MEMORIAM

In memory of Drew V. McDermott

With deep sorrow, we announce the death of Drew V. McDermott on May 26, 2022 at the age of 72. Drew was a brilliant scientist who made many seminal contributions to AI, an incisive critic of the blind spots and failings of the field, a gifted writer, an inspiring teacher, and a generous and beloved colleague and friend.

Drew first became known for his work, with his advisor Gerry Sussman, on the CONNIVER reasoning *cum* planning system (Sussman and McDermott, 1972). Over the next three decades, he did ground-breaking work across the spectrum of automated reasoning and planning, including seminal papers in reasoning architecture (McDermott, 1976), further work in planning (McDermott, 1978, 1985, 1991); in nonmonotonic logic with Jon Doyle (Doyle and McDermott, 1980, McDermott, 1982); in temporal reasoning (McDermott, 1982); in spatial reasoning with Ernest Davis (McDermott and Davis, 1984); and in semantic web development (Ankolekar et al., 2002). In particular, he served as chair of the committee that developed the Planning Domain Definition Language (PDDL) (Ghallab et al., 1998), which has become a standard for problem specification and domain definition in automated planning research. His discovery, with Steve Hanks, of the "Yale Shooting Problem" (Hanks and McDermott, 1987) hit the world of knowledge-based AI in the same way that Russell's paradox hit Gottlob Frege, exposing fundamental gaps in our understanding of the interaction of plausible inference with temporal reasoning that still have not been resolved.

Drew also co-authored two impactful textbooks. *Introduction to Artificial Intelligence*, with Eugene Charniak (Charniak and McDermott, 1985), presented a view of artificial intelligence as "the study of mental faculties through the use of computational models" through logical analysis. It was, in its time and for years afterward, by far the most comprehensive and deepest analysis of AI as a whole (though, truth to tell, it was almost impossible to teach from). *Artificial Intelligence Programming*, with Charniak and Christopher Riesbeck (Charniak, Riesbeck, and McDermott, 1980) (and with James Meehan in the 2nd edition) gave a generation of AI researchers the tools to build sophisticated AI systems in LISP.

However, Drew's best-known works–and perhaps his most important–were two incisive critiques of the field of AI in its then state. "Artificial Intelligence Meets Natural Stupidity," (McDermott, 1976) rebuked the field for "wishful mnemonics" that mislead the scientist as much as his audience; for describing nonexistent programs; for misunderstanding the nature of language; and for supposing that an AI will magically give the same interpretation to a symbol as its human creator. Many of its criticisms are still cogent 46 years later. "A Critique of Pure Reason" (McDermott, 1987) was more narrowly focused, it exposed the failings and limitations of the logicist program then being pursued by most AI researchers, including Drew himself. Some memorable quotations are given below.

At the center of both his scientific research and his critiques of the field lay a profound intellectual humility before the scope and difficulty of the challenges that we address. He insisted that problems–language, time, reasoning, planning–be addressed in their full complexity, not in toy versions; equally, he insisted that the representations and algorithms employed be well-defined, well-motivated, and clear. All of us who were privileged to work with Drew and learn from him have been deeply influenced by his demand for realism combined with rigor.

In later years, Drew focused his thought on the "big questions": What kinds of computer systems can legitimately be considered intelligent? What is consciousness? How will artificial intelligence specifically and technology generally impact the future of man, for good and ill? His book *Mind and Mechanism* (McDermott, 2001) gave a careful, clear, thought-provoking, philosophical analysis of how a computational model of mind can incorporate free will, qualia, and consciousness.

PERSONAL MEMORIES

Eugene Charniak tells some vivid personal memories of Drew:

At some point Roger Schank came to me as he was looking to hire another AI person at Yale. He had two people as possible candidates, X and Drew. I remember telling him that X was good, but Drew was otherworldly great. A few years later he thanked me saying that Drew was so modest he completely underestimated him.

I also have a memory of Gerry Sussman telling me that there was this guy at the other end of the 8th floor of the building we occupied in Tech Square who understood Micro-planner better than he did.

However, mostly I knew Drew not from MIT but from the time I spent at Yale as a visiting prof. I remember conversations with him on knowledge representation and his insistance on the denotations for the predicates we used. Of course, we all know how well he could cut through nonsense as in "AI and Natural Stupidity." I remember general confusions about the IS-A relation, and some researchers claiming that everything was an IS-A relation. Drew's comment was that if everything is an IS-A relation then nothing is. Rather IS-A just becomes a piece of syntax, like left parenthesis.

SOME QUOTATIONS

Both in writing and in speech, Drew had a rare gift for the telling sentence: clear, penetrating, and witty. It is fitting to close this memorial with some of his own memorable words.

From "Artificial Intelligence Meets Natural Stupidity":

> As a field, artificial intelligence has always been on the border of respectability, and therefore on the border of crackpottery.

> In AI, our programs are mostly problems rather than solutions. If a researcher tries to write an "understanding" problem, it isn't because he has thought of a better way of implementing this well-understood task, but because he thinks he can come closer to writing the *first* implementation. If he calls the main loop of his program "UNDERSTAND" he is (until proven innocent) merely begging the question. He may mislead a lot of people, most prominently himself, and enrage a lot of others. What he should do instead is to refer to this main loop as "G0034" and see if he can convince himself or anyone else that G0034 implements some part of understanding.

> This is an illustration of "contagious wishfulness": because one piece of a system is labeled impressively, the things it interacts with inherit grandiosity. A program called "THINK" is likely inexorably to acquire data structures called "THOUGHTS."

> For instance, although he is aware of complexities, Fahlman proposes that a first cut at representing "Nixon is a Hitler" is

```
HITLER
  ^
  | IS-A
  |
NIXON
```

> It worked for Fido and Dog, didn't it? But we just can't take stuff out of the IS-A concept that we never put in

> It is hard to know where [AI researchers] have gone wronger: in underestimating language or overestimating computer programs.

> AI as a field is starving for a few carefully documented failures.

From an NSF proposal (1981):

> For example, a classic problem of the type studied by problem-solving theorists is the "Tower of Hanoi" … The naıve inquirer is likely to ask, What's so hard about this problem? Are the disks very heavy? Are the pegs far apart? Does someone else want the disks to remain where they are?

From "A Temporal Logic for Reasoning about Processes and Plans"

> A common disclaimer by an AI author is that he has neglected temporal considerations to avoid complication. The implication is nearly made that adding a temporal dimension to the research (on engineering, medical diagnosis, etc.) would be a familiar but tedious exercise that would obscure the new material presented by the author. Actually, of course, no one has ever dealt with time correctly in an AI program, and there is reason to believe that doing it would change everything.

In conversation:

> The end result of Artificial Intelligence will be to show that intelligence is impossible, and that the reported instances of it have been due to experimental error.

> The theory of reactive planning is that you can spend as much time on a task as you want, as long as you don't spend it *planning*.

> The task "Seduce two virgins" is not accomplished by seducing the same virgin twice.

> We should think of human beings as very smart animals, not as very dumb gods.

From "Response to 'The Singularity: A Philosophical Analysis' by David Chalmers" (McDermott, 2012):

> The exponential growth in technology that is the major argument for the Singularity is accompanied by, perhaps made possible by, an exponential growth in the exploitation of finite natural resources (including the atmosphere, viewed as a carbon-dioxide sponge). Our civilization's addiction to a process that simply cannot continue is a sign of insanity, and belief in the Singularity may be one of its most comforting delusions. Even if some of the world's richer citizens get "uploaded," what happens when the power goes off?

From *Mind and Mechanism*:

> We may think that belief in God is a transient stage in the development of civilization, and that, if we survive our own technological achievements, we'll outgrow that belief. I doubt it. We will always be painfully aware of our finiteness, and will always yearn for the Infinite.

> The place where God intervenes in the world is therefore us. If his will is to become effective in the world, it will be because we carry it out.

CONFLICT OF INTEREST

The author declares that there is no conflict.

REFERENCES

Ankolekar, Anupriya, Mark Burstein, Jerry R. Hobbs, Ora Lassila, David Martin, Drew McDermott, Sheila A. McIlraith, et al. 2002. "DAML-S: Web service description for the semantic web." In *Proceedings of the International Semantic Web Conference*, 348–63.

Charniak, Eugene, Christopher Riesbeck, and Drew McDermott. 1980. *Artificial Intelligence Programming*. Hillsdale, NJ: Lawrence Erlbaum Pubs.

Charniak, Eugene, and Drew McDermott. 1985. *Introduction to Artificial Intelligence*. Reading, Mass: Addison Wesley.

Doyle, Jon, and Drew McDermott. 1980. "Nonmonotonic logic I." *Artificial Intelligence* 13(1): 41–72.

Ghallab, Malik, Adele Howe, Craig Knoblock, Drew McDermott, Ashwin Ram, Manuela Veloso, Daniel Weld, and David Wilkins. 1998. "The Planning Domain Definition Language." Technical report CVC TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control.

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

McDermott, Drew. 1976. "Flexibility and Efficiency in a Computer Program for Designing Circuits." PhD thesis, MIT.

McDermott, Drew. 1976. "Artificial intelligence meets natural stupidity." *ACM SIGART Bulletin* 57: 4–9.

McDermott, Drew. 1978. "Planning and acting." *Cognitive Science* 2(2): 71–100.

McDermott, Drew. 1982. "Nonmonotonic logic II: nonmonotonic modal theories." *Journal of the ACM* 29(1): 33–57.

McDermott, Drew. 1982. "A temporal logic for reasoning about processes and plans." *Cognitive Science* 6(2): 101–55.

McDermott, Drew, and Ernest Davis. 1984. "Planning routes through uncertain territory." *Artificial Intelligence* 22: 107–56.

McDermott, Drew. 1985. "Reasoning about plans." In *Formal Theories of the Commonsense World*, edited by J. Hobbs and R. Moore. Norwood: Ablex Pubs.

McDermott, Drew. 1987. "A critique of pure reason." *Computational Intelligence* 3(1): 151–60.

McDermott, Drew. 1991. "A Reactive Plan Language." Research report YALEU/DCS/RR-864, Yale University.

McDermott, Drew. 2001. *Mind and Mechanism*. Cambridge: MIT Press.

McDermott, Drew. 2012. "Response to 'The singularity: a philosophical analysis' by David Chalmers." *Journal of Consciousness* 19: 141–67.

Sussman, Gerald, and Drew McDermott. 1972. "Why Conniving is Better than Planning." MIT AI Lab, Artificial Intelligence Memo 255A.