

Advancing AI Research and Applications by Learning from What Went Wrong and Why

Daniel Shapiro and Mehmet H. Göker

■ *This special issue of AI Magazine is dedicated to the proposition that problems populate the path to insight, implying that experiences and lessons learned should be shared.*

I have not failed. I've just found 10,000 ways that won't work.¹

—Thomas Alva Edison

Insight often begins with unexpected results. In this view, problems are signposts, not roadblocks, which guide us towards better solutions. Bugs, surprises, and anomalies also become powerful instructional tools that reveal assumptions, expose design flaws, and chart the boundaries of current technology. This perspective motivated our work on this special issue.

When researchers publish success stories, we commonly leave process out of the narrative. We suppress the bad ideas and conceptual errors (a good thing), but also the interplay between failures and their resolution that rationalizes many design decisions. This interplay can be as valuable as a successful final result, because it alerts other researchers to potential mistakes, and clarifies key insights they can transfer. If you believe this statement, it implies that publications often resemble recipes without the secret sauce—certain key ingredients are unknown. As a corollary, the discussion of what went wrong and why deserves appropriate prominence in professional settings.

Armed with this conviction, we organized a symposium called “What Went Wrong and Why: Lessons from

AI Research and Application”² as part of the 2006 AAAI Spring Symposium Series at Stanford University (Shapiro and Göker 2006, Abdecker et al. 2006). We asked authors to describe the problems they encountered during their work and to couple them with lessons learned in technical, commercial, or organizational realms.

The results were excellent, as well as surprising. The authors extracted lessons not just about the technology of individual projects (as we had expected) but about research methodologies, commercialization models, funding strategies, and even AI in general. The symposium also attracted invited talks from Carl Hewitt, Craig Knoblock, Doug Lenat, John McCarthy, Mike Pazzani, Milind Tambe, and Manuela Veloso. These speakers drew an audience to match, which gave the assembly the aspect of a high level forum on what went wrong and why in AI as a whole. There was a great deal of heated and interesting debate, but we also laughed a lot—enough to disrupt meetings in adjacent rooms.

This special issue is a result of that symposium. As you will see, the articles confirm our belief that technical surprises beget insight that can shape the subsequent discourse. The issue begins with an article by Doug Lenat, who employs David Letterman’s top-10 list-style format to discuss the 12 things that have most gone wrong with AI. It contains his views on what

is happening and what should happen in AI research. In general, it is a call to pursue bold dreams.

Craig Knoblock and his colleagues José Luis Ambite, Mark Carman, Matthew Michelson, Pedro Szekely, and Rattapoom Tuchinda; and Milind Tambe, along with Emma Bowring, Jonathan P. Pearce, Pradeep Varakantham, Paul Scerri, and David V. Pynadath, describe the issues they faced when they unleashed their AI applications on a user community. Both realized new research directions by observing the many things that went wrong. If you enjoy stories about embarrassing mistakes, these papers reveal some of the electric elves' dirty laundry.

James Crawford describes the lessons he learned while developing and deploying large-scale supply-chain planning systems during his time at I2. The message is applicable to any enterprise that needs to tailor or scale up technology for application.

Barbara Hayes-Roth provides a first-person narrative of her experiences in the startup, failure, and resurrection of a business that commercialized agent technology for interactive characters. It illustrates the challenge of crossing the chasm from academic research into application.

Michael L. Anderson, along with his colleagues Scott Fults, Darsana P. Josyula, Tim Oates, Don Perlis, Matthew D. Schmill, Shomir Wilson, and Dean Wright proposes a means of implementing systems that learn from what went wrong and why. They employ a metacognitive loop (make mistakes, learn, and get smarter) as the central control structure within an autonomous agent. This work phrases falling down and recovery as a fundamentally constructive principle.

Jean-Gabriel Ganascia retrospectively analyzes the "what" part of "what went wrong and why." He examines obviously flawed (historical) theories and extracts the underlying assumptions that led a community to believe in them. This research offers insights on how to avoid such pitfalls. It has a wide-ranging, if disquieting relevance to our daily life.

We've added a number of anecdotes to this material to lend a more person-

al flavor to the discussion of problems and lessons learned. These stories came from an open-mike session at the original workshop and in response to a first-ever CFA (call for anecdotes). Several are humorous, others distill lessons from long periods of work, and one documents an "aha!" moment that framed a research career.

We also encountered two of our own "What Went Wrong and Why" experiences in composing this special issue. First, while the workshop attracted a strong response from senior figures of the field, it generated a more limited response from other researchers. Why? This could be the freedom of tenure in action, or risk aversion on the part of less-established contributors. There is a certain endemic resistance to exposing problems, mistakes, and errors, despite the benefits. This resistance is worth examining. The second glitch was that our CFA was not particularly successful. We expected hundreds but only received a handful. Why? Our best guess is that humor and personal stories are too far outside the norm for a technical venue. Here, we hope the community takes the lesson that the anecdotes add insight, spark, and personal interest to the research process, and that the material can be just plain fun.

In summary, our experiences organizing the What Went Wrong and Why symposium and editing this special issue have reaffirmed our belief that every researcher, developer, and entrepreneur has seen bugs give rise to insight, problems reveal new opportunities, and ideas fail while teaching valuable lessons in return. These experiences should be communicated. In addition to saving the community time and effort, they capture, in some sense, the core of the research process. Said simply, if problems are the mother of necessity, they are the grandmother of invention.³

These benefits all hinge on an open discussion of what went wrong and why. We hope, with some fervor, to see similar forums in the future, as this one was more valuable and enjoyable than any other workshop we have attended.

Notes

1. Attributed and most likely paraphrased

(for example, www.quotedb.com/quotes/1351).

2. Yes, we were tempted to insert a typo in the title but we resisted.

3. This theme is popular in the current business literature (such as *Whoever Makes the Most Mistakes Wins*, by Richard Farson and Ralph Keyes, New York: Simon & Schuster, 2002).

References

Abdecker, A.; Alami, R.; Baral, C.; Bickmore, T.; Durfee, E.; Fong, T.; Göker, M.; Green, N.; Liberman, M.; Lebiere, C.; Martin, J.; Mentzas, G.; Musliner, D.; Nicolov, N.; Nourbakhsh, I.; Salvetti, F.; Shapiro, D.; Schreckenghost, D.; Sheth, A.; Stojanovic, L.; SunSpiral, V.; Wray, R. 2006. AAAI Spring Symposium Reports. *AI Magazine* 27(3): 107–112.

Shapiro, D., and Göker, M., eds. 2006. What Went Wrong and Why: Lessons from AI Research and Application, Papers from the AAAI Spring Symposium, March 27–29, 2006, Stanford, CA. Technical Report SS-06-08. Menlo Park, CA: AAAI Press.

Daniel Shapiro (dgs@stanford.edu) is the executive director of the Institute for the Study of Learning and Expertise (ISLE), and president of Applied Reactivity, Inc. (ARi). He conducts research on cognitive architectures and human-level AI at ISLE, and pursues applications of reactive control systems that learn through ARi. He is an affiliate of the Computational Learning Laboratory, and formerly a senior researcher at the Center for the Study of Language and Information (CSLI), both at Stanford University. He received his Ph.D. in management science and engineering from Stanford University in 2001.

Mehmet H. Göker leads the Knowledge Management Research Group at PricewaterhouseCooper's Center for Advanced Research in San Jose, California. He and his group have most recently deployed the Connection Machine, a user-adaptive expertise locator. Previously, Göker ran the professional services of Kaidara Software, a premier case-based reasoning software vendor, and was a senior research scientist at DaimlerBenz (and DaimlerChrysler) Research and Technology. Göker has been developing intelligent, human-centered, industrial-strength knowledge-management systems for 15 years. He has a Dr.-Ing from the Technische Universität Darmstadt.