What Went Wrong and Why: Lessons from AI Research and Applications

Mehmet H. Göker[1] and Daniel G. Shapiro[2]

[1]Center for Advanced Research
PricewaterhouseCoopers L.L.P.
Ten Almaden Blvd., Suite 1600
San Jose, CA 95113
mehmet.goker@us.pwc.com

[2]Institue for the Study of Learning and Expertise (ISLE)
2164 Staunton Court
Palo Alto, CA 94306
dgs@stanford.edu

Bugs, glitches, and failures shape research and development by charting the boundaries of technology; they identify errors, reveal assumptions, and expose design flaws. When a system works we focus on its input/output behavior, but when a problem occurs, we examine the mechanisms that generated behavior to account for the flaw and hypothesize corrections. This process produces insight and forces incremental refinement. In a sense, failures are the mother of necessity, and therefore the grandmother of invention.

Unfortunately, bugs, glitches, and failures are *rarely* mentioned in academic discourse. Their role in informing design and development is essentially lost. The first *What Went Wrong and Why* workshop during the 2006 AAAI spring symposium [1,2] started to address this gap by inviting AI researchers and system developers to discuss their most revealing bugs, and relate problems to lessons learned. Revised versions of the articles and the invited talks will be published as a special issue of the AI-Magazine in Summer 2008 [3].

The first workshop clarified that WWWW experiences can be studied at three different levels of abstraction: the Strategic (AI research in general), Tactical (research area) and Execution (project or implementation) levels. An additional category turned out to be the study of how, why and when failures occur in the first place.

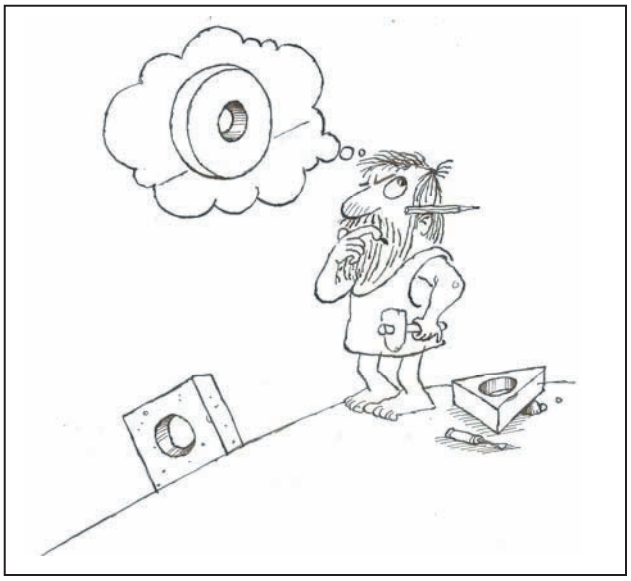

© Yurdagün Göker, 2008

The second workshop continues our analysis of failures in research. In addition to examining the links between failure and insight, we hope to determine if there is a hidden structure behind our tendency to make mistakes that can be utilized to provide guidance in research.

The workshop addresses these themes via papers and three invited talks: Haym Hirsh discusses lessons learned from methodological failures in AI research, Bruce Buchanan extracts lessons

from the expert systems boom, and Steve Chien describes the successes and failures of AI on three deployed spacecraft. The papers connect problems to lessons learned within individual research efforts, although one places the entire history of logic programming in scope. Overall, we encouraged authors to elaborate on what they believe was the source cause of the failure, how the problem helped them arrive at a better solution, and to suggest a broader categorization of failures and how to utilize them. As with the first workshop on What Went Wrong and Why, we expect this workshop to be interesting, informative, and fun.

References

[1] Shapiro, D., Göker, M. (eds.), 'What Went Wrong and Why: Lessons From AI Research and Applications', Papers from the AAAI Spring Symposium, March 27-29, 2006, Stanford, CA. Technical Report SS-06-08, AAAI Press, Menlo Park, 2006.

[2] A. Abdecker, R. Alami, C Baral, T. Bickmore, E. Durfee, T. Fong, M. Göker, N. Green, M. Liberman, C. Lebiere, J. Martin, G. Mentzas, D. Musliner, N. Nicolov, I. Nourbakhsh, F. Salvetti, D. Shapiro, D. Schreckenghost, A. Sheth, L. Stojanovic, V. SunSpiral, R. Wray, "AAAI Spring Symposium Reports" , AI Magazine, VOl 27, Nr. 3, Fall 2006, pp. 107-112, American Association for Artificial Intelligence (AAAI), Menlo Park, 2006

[3] Shapiro, D. Göker, M. (eds.), 'Special Issue on What Went Wrong and Why", AI Magazine, Vol. 29, Number 2, Summer 2008