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

AAAI Spring Symposium, 2006

The What Went Wrong and Why workshop is dedicated to the proposition that insight often begins with unexpected results, and that clarity often arrives in the response to surprises and apparent problems. For example, fogged film led Curie to discover radium; a failed culture let Fleming find penicillin; microwave noise let Penzias and Wilson verify the Big Bang. Bugs, glitches, and failures also chart the boundaries of technology. They shape research and development by identifying errors, revealing assumptions, and exposing design flaws. Here, a failure is often more informative than a successful demonstration. When a system works we focus on its input/output behavior, but when a problem occurs, we examine the mechanisms that generate behavior to account for the flaw and hypothesize corrections. This produces insight and forces incremental refinement. Thus, trouble is the mother of necessity and therefore the grandmother of invention.

Unfortunately, bugs, glitches, and failures are *rarely* mentioned in academic discourse (unless they are the object of study). Their role in informing design and development is essentially lost. The *What Went Wrong and Why* workshop addresses this gap by providing AI researchers and system developers an opportunity to discuss their most revealing bugs, and share the resulting insights.

The call for papers specifically asked authors to relate problems to lessons learned, such as:

Problem	*Lesson Learned*
• unusual behavior	• clarity on why an algorithm works (or not)
• mistaken assumption	• deep technical insight
• method or algorithm error	• key research problem to solve
• technology/application mismatch	• reframed/reformulated task
• product/project/corporate demise	• new technique or methodology
• fried hardware	• new business opportunity
• physical harm	• changes to a human organization
	• tangible human benefits

The responses published here extract lessons about AI technology from experiences in technology transfer, application development, and commercial product deployment. Two papers go beyond the call and ask where bad theories come from, and how to recognize a WWWW experience in progress so that it can be employed in a constructive sense.

In addition to papers, the workshop features an "open-mike" period where participants can share anecdotes about their most interesting WWWW experiences. Finally, seven prominent contributors to the field of AI give invited talks on the What Went Wrong and Why theme as it applies to technical and commercial realms.