

# How Inappropriately Heavyweight AI Solutions Dragged Down a Startup

*(and Made Me Realize that Industrial Salaries Are High for a Good Reason)*

*Michael Wooldridge*

The basic idea was to build distributed information systems, with intelligent nodes at each repository, able to reflect upon and reason about repository content in order to route queries appropriately. We came up with a heavyweight agent architecture, using ideas from AI planning and robotics. These sorts of architectures were very much in vogue at the time, and the company wanted its own, proprietary technology. We started thinking about programming languages for the agents and the kinds of knowledge representation and reasoning that would be required. We spent a lot of time and money flying from London to the U.S. West Coast, talking to patent lawyers.

It transpired that the architecture, its decision-making and action models, were completely inappropriate for the problem at hand. By the time we realized we should have been focusing on basic software engineering, quality assurance, and end-user requirements, the company had burned out much of the goodwill—and most of the funds—of our investors. I jumped ship after less than a year, back to academic life (on a much reduced salary). The company folded about six months later, in the spring of 1998 (about the same time that Page and Brin were launching Google).

The experience was traumatic, but I learned a great deal about software engineering, the real priorities of companies, and the kinds of things that do and do not make money.

**Michael Wooldridge** has been a professor of computer science at the University of Liverpool since 1 January 2000. His research is primarily in the theoretical foundations of multiagent systems; he was the recipient of the ACM SIGART Autonomous Agents Research Award in 2006 and is coeditor-in-chief of the journal *Autonomous Agents and Multi-Agent Systems*.