

The 2008 Classic Paper Award: Summary and Significance

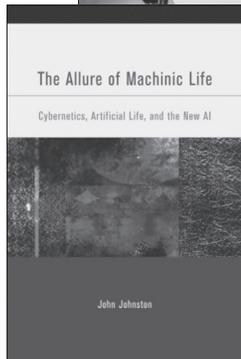
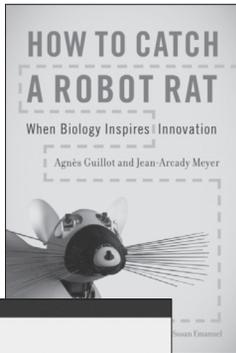
Peter Friedland

■ *We at the NASA laboratory believed that our best work came when we simultaneously advanced AI theory and provided immediately usable solutions for current NASA problems. "Solving Large-Scale Constraint Satisfaction and Scheduling Problems Using a Heuristic Repair Method," by Steve Minton, Mark Johnston, Andy Phillips, and Phil Laird clearly achieved both. It proved that local search and repair was applicable to a wide class of constraint-satisfaction problems and clearly explicated the theory behind that proof.*

As the founder and former director of the Artificial Intelligence Research Branch of NASA Ames Research Center, I was delighted that "Solving Large-Scale Constraint Satisfaction and Scheduling Problems Using a Heuristic Repair Method," by Steve Minton, Mark Johnston, Andy Phillips, and Phil Laird won the prestigious Classic Paper award. The work epitomizes the guiding philosophy of that laboratory: AI research can simultaneously advance the state of the art and provide practical solutions to key problems faced by the Space Agency and its collaborators.

Minton and colleagues developed a heuristic repair method, called "min-conflicts" for solving large-scale constraint-satisfaction problems (CSP), with a particular focus on massive scheduling tasks. Mark Johnston, an astronomer and computer scientist from the Space Telescope Science Institute at Johns Hopkins, served simultaneously as domain expert and codeveloper. He provided full access to one of NASA's most challenging scheduling tasks, the selection and ordering of observations for the Hubble Space Telescope. Between 10,000 and 30,000 observations per year must be scheduled, with complex, interlocking constraints involving the orbit of the telescope, location of astronomical objects, power, stray light sources, scientific and administrative priorities, and so on. The team also applied their method to a classic benchmark constraint satisfaction problem, the n -queens problem (where the goal is to place n queens on an $n \times n$ chessboard so that none attack any other). In both cases the min-conflicts method performed many orders of magnitude better than prior technology.

The inspiration for min-conflicts came from studying a neural-network-based scheduler (called the Guarded Discrete Stochastic or GDS network) that Johnson and colleagues had developed at the Space Telescope Science Institute. It performed better than a traditional backtracking-based scheduler. Minton's team's research discovered that the GDS scheduler operated by updating or "repairing" the neuron with the most conflicts and repeating this process until a complete and correct schedule could be produced. They reproduced this behavior in a simple and efficient hill-climbing algorithm. That system performed even better than the original GDS scheduler/constraint satisfaction system (on the n -queens problem where GDS was able to

 The MIT Press


How to Catch a Robot Rat

When Biology Inspires Innovation

Agnès Guillot and Jean-Arcady Meyer

translated by Susan Emanuel

“Over the last twenty-five years, a subset of computational and robotics researchers around the world have taken to studying biological creatures in order to figure out how to build robots. And at the same time, the constraints they have discovered in building robots have been used to illuminate how the biological systems must work. Guillot and Meyer have been both intellectual and organizational leaders in this field, and in *How to Catch a Robot Rat* they carefully document the history and intellectual currents of the field.” — Rodney Brooks, MIT

232 pp., 103 illus., \$29.95 cloth

Now in Paper

The Allure of Machinic Life

Cybernetics, Artificial Life, and the New AI

John Johnston

“Johnston has done a magnificent job of surveying and digesting the vast literature and producing an extraordinarily clear account of this topic.”

— C. Tappert, *Choice*

480 pp., 51 illus., \$20 paper

To order call 800-405-1619 • <http://mitpress.mit.edu> • Visit our e-books store: <http://mitpress-ebooks.mit.edu>

solve 1 thousand queens problems in 11 minutes, the new min-conflicts systems solved 1 million queens in less than 4 (using comparable computational resources).

As already mentioned, we at the NASA laboratory believed that our best work came when we simultaneously advanced AI theory and provided immediately usable solutions for current NASA problems. This work clearly achieved both. It proved that local search and repair was applicable to a wide class of constraint satisfaction problems and clearly explicated the theory behind that proof. This occurred at a time when the majority of work in CSP focused on various forms of backtracking. Since the min-conflicts approach “blew away” prior techniques on the benchmark n -queens problem by an over four orders of magnitude performance improvement, it clearly indicated a new focus for the field.

This significance over the longer term is shown by the over 350 citations of this paper in the AI literature—an extraordinary large number for a conference proceedings paper. It spawned significant follow-on work in stochastic methods for constraint satisfaction, including the well-known GSAT system of Selman, Levesque, and Mitchell. It also found application in the “just-in-time” sched-

uling and rescheduling work of others in the NASA laboratory and elsewhere in the community. On the practical side, the min-conflicts algorithm replaced the neural net backbone of the Hubble Space Telescope scheduler, becoming the standard tool for selecting observations for that gem of NASA space sciences.

Peter Friedland received AB (chemistry) and BSE (Electrical engineering) degrees in 1974 from Princeton University and a PhD in computer science from Stanford University in 1979, conducting research on the application of artificial intelligence to the design of experiments in molecular biology. He acted as research director of the MOLGEN project at Stanford, 1979–1986, continuing research into planning, knowledge representation, and learning within the biological domain. In 1980 he cofounded two companies, IntelliCorp and Teknowledge, to commercialize the emerging field of knowledge-based or expert systems. In January 1987, Friedland joined NASA Ames Research Center as chief of the Artificial Intelligence Research Branch with responsibilities for creating a fundamental research capability in artificial intelligence within NASA, as well as managing an extensive collaborative program with industry and academe. In addition, Friedland comanaged (with M. Montemerlo) the NASA-wide AI program. He is now a senior scientific advisor to the Air Force Office of Scientific Research. He is also a Fellow of the AAI.