

120-cm-diameter swimming pool. The knock-outs were much slower to find a submerged platform on one side of the pool and less able to remember the platform's position later. Concludes Tonegawa, "These mice were basically incapable of acquiring spatial memory."

Underlying this handicap, the team found, was an altered pattern of neuronal firing. Researchers already knew that spatial learning in rodents involves "neural maps" of firing in the hippocampus. A given cluster of cells, for example, will fire only when a mouse is in a specific spot, say, the southwest corner of a box. That spot is called the cells' "place field." Other cells fire when the mouse is in other locations. Together, overlapping place fields are thought to create a kind of internal map, and researchers had suspected—but never proven—that LTP is what sustains place fields over time.

Now, the MIT team has confirmed this suspicion by inserting electrodes into the brains of living mice and watching pyramidal cells fire as the animals explored variously shaped cages. Place fields in the transgenic mice were less compact and focused than in normal mice, likely accounting for the difficulty the mice had in navigating a new environment.

Meanwhile, working in parallel, the Columbia-SUNY team has come up with a similarly dramatic link between molecular changes in the hippocampus, place fields, and memory. They disrupted hippocampal LTP in a different way, using a combination of promoters that was less tissue specific. And they enhanced, rather than eliminated, expression of the gene encoding CaMKII in certain brain cells, including the pyramidal cells. Constant

activation of CaMKII, they theorized, would disrupt learning. Indeed, their mice also showed spatial memory deficits and less focused place fields. Kandel's group also found that the fields were less stable over time.

These findings, linking molecular, neuronal, and behavioral abnormalities, are moving neurobiologists closer than ever to an understanding of the molecular basis of memory, researchers say. "We're just at the beginning of making broad links between lower, molecular levels of analysis and higher levels of cognition and behavior, but these are certainly important steps," says Daniel Schachter, a cognitive neuroscientist at Harvard Medical School and author of the new book *Searching for Memory*. By offering data at every level, the new studies are likely to prove memorable themselves.

—Wade Roush

COMPUTER SCIENCE

Hedging Bets on Hard Problems

It's a dilemma familiar to every Internet addict. Your Web browser insists that it is connected to a site, but as it happily chugs away, nothing happens. Do you wait, hoping that the data will come in in a few more seconds, or do you give up and try later? It's a gamble either way, because traffic on the Internet fluctuates from second to second. On page 51, computer scientists at the Xerox Palo Alto Research Center have adopted the practices of Wall Street in an approach to solving this and many other problems in computer science that entail the same kind of uncertainty.

Just as investors try to improve their returns by splitting their money among a number of investments, the Xerox team improves a computer's performance on problems ranging from factoring a large number to minimizing a complicated error function by splitting its attention among a number of programs. "Using a 'portfolio' is a really neat way to improve [an algorithm's] performance," says Bernardo Huberman, one of the researchers. Agrees Hal Varian, an economist and the dean of Information Management at the University of California, Berkeley, "It's very clever."


The algorithms that Huberman and his colleagues are trying to speed up rely on a gamble to solve extremely hard problems. These "NP-complete" problems take an enormous amount of effort to work through, the effort rising exponentially with the complexity of the example. But the solutions are crucial in many areas of science. Finding the global minimum of a complicated function—an NP-complete task—is crucial in training a neural net or predicting the shape of a folding protein.

These problems are so difficult that the best a computer can do is to wander around methodically in search of an answer—a so-

called Las Vegas algorithm. These algorithms often begin their calculations from random "seeds." If this starting position is reasonably close to the correct solution, the algorithm breezes through the problem. But a bad seed can send the algorithm up the wrong path, and the computer may grind away fruitlessly for hours.

In hopes of boosting a computer-science version of "return"—the speed with which a computer solves a hard problem—Huberman and his colleagues decided to diversify a computer's algorithm portfolio. They tested the approach on a well-known NP-complete problem: coloring a graph. The computer is given a set of circles connected by lines; given a certain number of colors, the computer is assigned to shade each circle so that no two connected by a line share the same color.

The Xerox team made a portfolio of two copies of a graph-coloring algorithm. The computer then split its attention between the two "investments." In the same way that a gambler might put \$98 on a favored horse and \$2 on a long shot, one copy of the algorithm received most of the computer's attention, while the other, with a different random seed, got only a little bit. Once in a while, the less favored algorithm got lucky and solved the problem quickly. By playing with the computer's stake in each investment, Huberman and his colleagues found a balance where the benefit from the occasional long-shot jackpot more than compensated for the loss of attention to the favored algorithm. Not only was the two-copy



Harder than it looks. Solving a graph-coloring problem takes exponentially longer as the problem grows.

portfolio 22% faster, on average, than a single copy of the algorithm, but the risk of a long wait for an answer dropped by 10%.

The Xerox researchers predict that they can extract a more dramatic speed-up by taking another page from Wall Street. Just as an investor can mix stocks and bonds, the computer portfolio approach can combine several different algorithms for solving the same problem. If those al-

gorithms are complementary—one does very well in exactly the cases where a second fails, and vice-versa—the portfolio would have stunning performance.

Huberman and the Xerox team are now experimenting with changing the amount of time the computer spends on each investment on the fly, much as an investor rebalances the proportion of investments in his portfolio, based on the various time scales over which each algorithm has the best chance of getting an answer. "We get a five- to 10-times speed improvement," claims Huberman.

They are also working on bringing the approach to the aid of frustrated Web surfers. Once the team unravels the causes and the distribution of Internet delays, they hope to develop a portfolio strategy that will lay the groundwork for a multiprocess browser that will speed up access to the Web. Huberman and his colleagues aren't predicting which of these research directions will pay off, but they know the benefits of diversifying. Says Berkeley's Varian, "The research mimics the algorithm."

—Charles Seife

Charles Seife is a science writer in Scarsdale, NY.