NEUROSCIENCE

New Knockout Mice Point to Molecular Basis of Memory

Shakespeare's Macbeth, seeing his mad, guilt-plagued wife suffer, wondered why her doctor could not "Pluck from the memory a rooted sorrow, / Raze out the written troubles of the brain?" Sorrows can't be surgically excised, but recently biologists have accomplished something close: They've erased the ability to remember from the brains of laboratory mice, by plucking out or adding key proteins in particular clusters of brain cells. These feats at last offer direct confirmation of the reigning theory of how we remember and show how molecular changes affect the patterns of electrical activity of which memories are made.

The work also showcases new methods for probing the workings of the mind. One set of studies, described in three related papers in the 27 December issue of the journal *Cell* by a team at the Massachusetts Institute of Technology (MIT) led by Nobel Prize–winning biologist Susumu Tonegawa and neuroscientist Matthew Wilson, uses an exotic new gene-splicing technique to produce a new kind of "knockout" mice that lack a certain gene. Knockouts are standard in biology these days, but these mice are different: The gene in question, which codes for a receptor for a key neurotransmitter—glutamate—was deleted in only one small group of cells in a brain region called the hippocampus, rather than in every cell in the mouse's body. A fourth paper in *Cell*, and one in the 6 December issue of *Science* (p. 1678) from an independent team, selectively enhanced, rather than deleted, a gene for a particular enzyme (called α-calcium-calmodulin–dependent kinase II or CaMKII). However, this study, led by Eric Kandel and Mark Mayford of Columbia University and Robert Muller of the State University of New York (SUNY) Downstate Medical Center in Brooklyn, had a less precise effect on the brain than did the Tonegawa group's method.

Both manipulations disrupted patterns of neuronal firing in the hippocampus and impaired the animals' ability to learn their way around mazes. Not only do these results lend substantial support to the prevailing theory linking molecular events in the hippocampus to spatial learning; they are also the first to probe memory at all levels in a single set of experiments, from molecular changes through altered patterns of neuronal firing to impaired learning. "It's a dream of neurobiologists to understand some interesting cognitive phenomenon like a memory from the molecular level right up through behavior," says neurobiologist Charles Stevens of the Salk Institute in La Jolla, California. "The articles in *Cell* are a big step in that direction." The selective knockout technique also promises to generate a new wave of progress, as other researchers create their own favorite mouse strains, including those linked to common neurological diseases. Says neuroscientist Michael Stryker of the University of California, San Francisco: "It points the way forward to what the whole field will be doing in the future."

Researchers have long thought that the mammalian brain stores new memories by long-term strengthening of the electrochemical signaling between neurons. And many studies had implicated CaMKII and a particular receptor for the neurotransmitter glutamate, the *N*-methyl-D-aspartate (NMDA) receptor, as crucial players in this process. When glutamate released by a transmitting neuron binds to NMDA receptors on an adjacent, receiving neuron, the receptors open channels in the cell membrane and allow calcium ions to flood in. These ions convert CaMKII into its active form and unleash a biochemical cascade that heightens the receiving neuron's sensitivity to subsequent signals and increases the amount of current it sends on to other neurons.

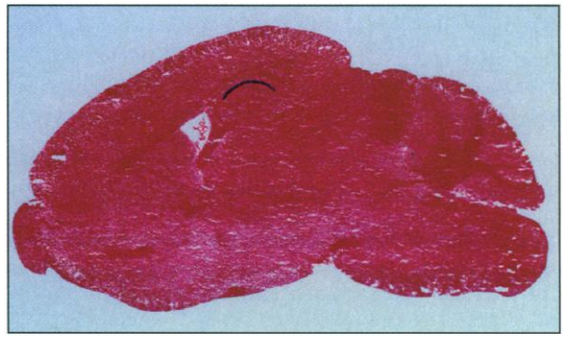

JOE TSIEN/MIT

Pranks for the memories. A clever new method deleted key receptors in only one small part of the brain *(purple)*.

A key site for this memory-building process, known as long-term potentiation (LTP), is thought to be the hippocampus, because injuries there produce severe amnesia. This idea is supported by earlier studies that used drugs to block NMDA receptors, and also by work from the Tonegawa and Kandel labs. The two are friendly rivals, and a few years ago independently created the first knockout mice lacking CaMKII or other key neuronal molecules (*Science*, 10 July 1992, p. 162).

These strategies for blocking the molecular events in LTP created learning impairments. But both were clumsy—the research equivalent of treating a hangnail by amputating one's hand—because they alter the entire brain throughout the life of the animal, including during embryonic development. So researchers couldn't be sure that the memory deficiencies in the altered mice weren't the product of developmental defects or changes outside the hippocampus.

In the new experiments, the MIT team sidestepped that problem with a clever new method, relying on the natural variability in the expression of the CaMKII gene to delete the NMDA receptor gene from discrete parts of the mouse brain. Tonegawa lab member Joe Tsien inserted the gene encoding a DNA-splicing enzyme called Cre in such a way that it was under the control of a "promoter" or on-off switch for the CaMKII gene. This created a collection of mouse strains that expressed Cre in different combinations of tissues, mimicking the pattern of CaMKII expression. "Very fortunately," as Tonegawa puts it, three strains expressed Cre only in certain cells (pyramidal cells) in a particular region of the hippocampus—and only about 3 weeks after birth, after all the normal hippocampal synaptic connections were built.

To affect the NMDA receptors, the group then mated mice from one of these strains to other specially engineered mice that carried twin DNA sequences just before and after the gene encoding the NMDA receptor. In some of the progeny from this match, the Cre enzyme recognized the twin sequences and linked them together, in the process snipping out the NMDA gene in between and creating a mouse in which only the pyramidal cells of the hippocampus lacked the receptors (see photo).

This two-step method is likely to spread like wildfire as neuroscientists adapt it to study the molecular defects underlying such disorders as Alzheimer's and Parkinson's diseases. "This is a major step along the road from global gene knockouts in the early embryo to spatial and temporal control over gene expression," says neurobiologist Steven Hyman, director of the National Institute of Mental Health. And it produced striking results in the NMDA knockouts.

To see how the absence of the NMDA receptors might affect memory, the MIT team analyzed brain slices from their new knockouts, using equipment that artificially shocks neurons and records their responses. After multiple shocks, the pyramidal cells exhibited no increase in the amount of electrical charge they transmitted, proving that intact NMDA receptors are key to LTP. Even more convincing, the team compared the behavior of adult knockout mice to that of their normal littermates in a

120-cm-diameter swimming pool. The knockouts 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.

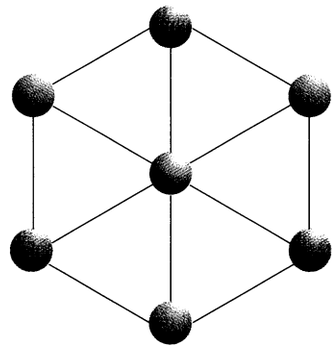

Harder than it looks. Solving a graph-coloring problem takes exponentially longer as the problem grows.

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 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 speedup 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 algorithms 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.