

craton, where the pattern matches that expected from the gas-hydrate model. Further, values of strontium isotopes during the glacial intervals provide no support for greatly diminished weathering during glaciation, or for greatly enhanced weathering in its immediate aftermath⁴, as required by the snowball Earth model. But they are consistent with the gas-hydrate hypothesis.

Both hypotheses involve short-lived changes in the carbon-isotopic composition of the ocean. But the gas-hydrate model avoids some of the obvious difficulties associated with huge changes in atmospheric CO₂ in the snowball Earth model, such as the long time (ten million years) required for such large concentrations to build up. It also provides a better explanation of the temporal variation of carbon isotopes in marine carbonates through the late Neoproterozoic.

But although Kennedy *et al.* have put forward a powerful case, this is, no doubt, far from the end of the story. ■

Stein B. Jacobsen is in the Department of Earth and Planetary Sciences, Harvard University, 20 Oxford Street, Cambridge, Massachusetts 02138, USA. e-mail: jacobsen@neodymium.harvard.edu

1. Harland, W. B. *Geol. Rundsch.* **54**, 45–61 (1964).
2. Kirshvink, J. L. in *The Proterozoic Biosphere: A Multidisciplinary Study* (eds Schopf, J. W. & Klein, C.) 51–52 (Cambridge Univ. Press, New York, 1992).
3. Kaufman, A. J., Knoll, G. H. & Narbonne, G. M. *Proc. Natl Acad. Sci. USA* **94**, 6600–6605 (1997).
4. Jacobsen, S. B. & Kaufman, A. J. *Chem. Geol.* **161**, 37–57 (1999).
5. Hoffman, P. F. *et al. Science* **281**, 1342–1346 (1998).
6. Kennedy, M. J., Christie-Blick, N. & Sohl, L. E. *Geology* **29**, 443–446 (2001).
7. Hyde, W. T., Crowley, T. J., Baum, S. K. & Peltier, W. R. *Nature* **405**, 425–429 (2000).
8. Crowley, T. J., Hyde, W. T. & Peltier, W. R. *Geophys. Res. Lett.* **28**, 283–286 (2001).
9. Walker, J. C. G., Hays, P. B. & Kasting, J. F. *J. Geophys. Res.* **86**, 9776–9782 (1981).

Neuroscience

Dynamic categories

Michael P. Kilgard

Neuroscientists often study brain function by presenting a stimulus many times and averaging the neural response. A trick for finding brain activity in single trials might reveal how animals mentally categorize information.

One of the most enduring problems in biology is that of how the brain assigns meaning to the stimuli it receives. We naturally divide the world into convenient categories, but the brain mechanisms that underlie this ability are poorly understood. On page 733 of this issue, Ohl and colleagues¹ report that category formation in animals is accompanied by fleeting but identifiable patterns of brain activity. The patterns exist for only a fraction of a second, and would be dismissed as noise by more traditional methods of analysis. Ohl *et al.*'s powerful technique may eventually reveal whether and how similar patterns contribute to our own cognitive abilities.

Historically, the responses of human and animal brains to a stimulus, such as a particular sight or sound, have been analysed by presenting the stimulus many times and averaging the neuronal activity that occurs. But because the responses are averaged, this approach detects only activity that occurs reliably at a certain time after every stimulus (such activity is said to be 'time-locked' to the stimulus). Nonetheless, the technique has yielded considerable insight into brain function.

One type of response, called mismatch negativity (MMN), has been widely used to study how the brain categorizes similar stimuli. This type of brain activity occurs when an unexpected sound (a mismatch) interrupts a series of identical auditory stimuli. Categorization is particularly important in

processing spoken language, where a variety of sounds are given one meaning. For example, the gender, age and emotion of the speaker significantly affect the physical qualities of speech, but 'dog' means 'dog' no matter who says it or in what tone. Of course, speech sounds may have varying meanings in different languages. The MMN response has been used to show that native and non-native speakers of a particular language categorize sounds from that language differently^{2,3}. Language-specific representations of speech sounds develop in infants between six and twelve months of age⁴, but can be altered throughout adulthood with sufficient practice⁵.

Although these studies provide a physiological marker of categorization that changes with learning, the MMN response itself does not represent specific categories, and the neural basis for categorization has remained elusive. To tackle the problem, Ohl and colleagues previously trained gerbils to categorize tones as either rising or falling in pitch⁶. After a few days of training, the animals could categorize a new tone instantly, without a need for training with that stimulus. On hearing a rising tone they moved immediately to another compartment of the box in which they were housed; after hearing a falling tone, they stayed where they were. Lesion of various parts of the brain showed that the right auditory cortex is required to discriminate between rising and falling tones⁷. Neurons in this region respond dif-

ferently to tones that are sweeping upwards and downwards. However, these differences seem to reflect the physical differences in the sounds, rather than their behavioural meaning (their category).

Largely on the basis of studies of the olfactory system, Freeman and colleagues suggested that dynamic patterns of brain activity carry information about the meaning of a stimulus⁸. But these patterns are not reliably time-locked to the stimulus, so they cannot be detected in averages of neuronal activity. Ohl, Scheich and Freeman¹ have now developed a strategy for identifying unusual activity patterns in a single trial, and applied it to the auditory system of gerbils. They find that certain activity patterns, which they call marked states, occur in animals that have learned to distinguish between rising and falling tones (Fig. 1, overleaf). The authors suggest that these states might provide a mechanism by which the animals can recognize the abstract qualities that define a category.

Ohl *et al.*¹ detected the unusual brain states by using a grid of 18 electrodes, positioned directly over the auditory cortex, that detect electroencephalographic (EEG) activity. To isolate 'local' activity patterns that occur within the auditory cortex, the authors removed any activity that was common to all electrodes. The patterns of activity that occurred in response to a single rising tone were then compared with the averaged pattern from 30 trials with falling tones, and vice versa. Unusual patterns were considered marked states if they were more than three standard deviations away from the average pattern produced in response to the opposite tone.

As expected⁹, the rising and falling tones evoked different patterns of activity in the auditory cortex in trained and untrained animals¹. Then, as each animal learned to distinguish the sounds, further marked states developed within a few seconds of each tone¹ (Fig. 1c, black boxes). The crucial finding is that these dynamic patterns were stabilized when animals learned that the behavioural meaning of a stimulus ('move' or 'don't move') was based on the abstract quality of tone direction. Although this finding was replicated in only four animals, the discovery of these activity patterns is an important step towards identifying the biological basis of categorization. It appears that the key to the authors' success lay in removing the irrelevant (global) activity to reveal local patterns that were most closely related to the behavioural meaning of the stimuli. Future work will doubtless look at how the patterns vary from trial to trial and from individual to individual, and whether marked states are related in any way to the neuronal patterns evoked by the two classes of stimulus⁹.

The marked states developed during the decision-making period, before each animal was forced to choose whether to move to another compartment. It will be interesting

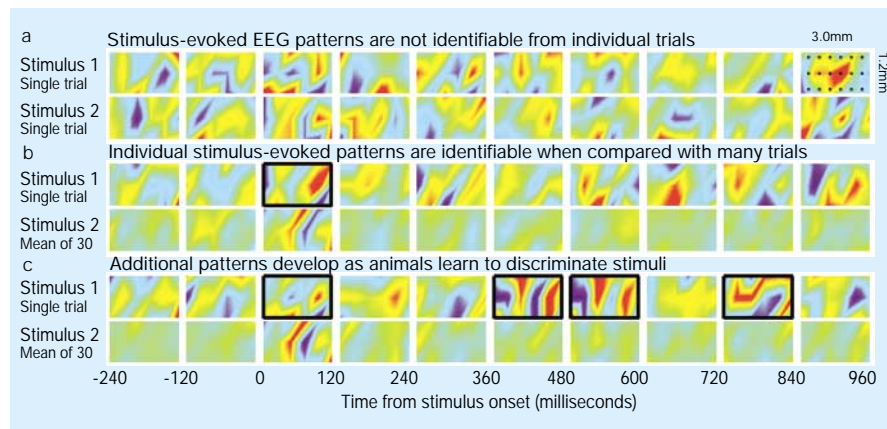


Figure 1 Detecting the brain activity that underlies the formation of categories. Each rectangle represents the pattern of activity in the auditory cortex at an instant in time in response to one of two sounds. Red denotes high activity. The black dots mark the locations of the electrodes recording neuronal activity. **a**, Brain activity can be very noisy, making it hard to identify patterns. **b**, Ohl *et al.*¹ have developed a technique for identifying the unique patterns that arise at some point after each of the sounds is heard. The strategy involves averaging the responses to stimulus 2 (or 1), and comparing them with individual responses to stimulus 1 (or 2). **c**, Additional patterns develop when animals learn to categorize the stimuli. Black boxes mark the patterns of brain activity that can be identified in a single trial using this technique.

to find out how the timing of these states changes when the task contingencies (such as the length of the decision-making period) are altered, and whether similar states develop during other categorizations. Ohl *et al.*¹ have already shown that marked states do not develop in the visual cortex during this auditory task. We now need to know whether they develop in the visual cortex during the categorization of visual stimuli.

Can this new technique be used to study how humans form categories? The number of EEG electrodes used in human studies has been increasing rapidly, yielding higher spatial resolution. So far, most groups have analysed only averaged brain responses. More sophisticated analysis of high-resolution EEG data will probably reveal complex local patterns in humans. The greater challenge will be in making sense of these patterns. Ohl

et al.'s technique may not lead to mind reading, but it does suggest a variety of new tools that may yet reveal the cognitive processes that give meaning to our world. ■

Michael P. Kilgard is at the School of Human Development, University of Texas at Dallas, Richardson, Texas 75083, USA.
e-mail: kilgard@utdallas.edu

1. Ohl, F. W., Scheich, H. & Freeman, W. J. *Nature* **412**, 733–736 (2001).
2. Näätänen, R. *et al.* *Nature* **385**, 432–434 (1997).
3. Sharma, A. & Dorman, M. F. *J. Acoust. Soc. Am.* **107**, 2697–2703 (2000).
4. Cheour, M. *et al.* *Nature Neurosci.* **1**, 351–353 (1998).
5. Tremblay, K., Kraus, N., Carrell, T. D. & McGee, T. J. *J. Acoust. Soc. Am.* **102**, 3762–3773 (1997).
6. Wetzel, W., Wagner, T., Ohl, F. W. & Scheich, H. *Behav. Brain Res.* **91**, 29–39 (1998).
7. Wetzel, W., Ohl, F. W., Wagner, T. & Scheich, H. *Neurosci. Lett.* **252**, 115–118 (1998).
8. Freeman, W. J. *Prog. Brain Res.* **102**, 319–333 (1994).
9. Ohl, F. W., Schulze, H., Scheich, H. & Freeman, W. J. *J. Physiol. (Paris)* **94**, 549–554 (2000).

Planetary science

A new model Moon

Jay Melosh

The most sophisticated simulations yet of the Moon's birth show that it could have been created from an impact of a large body with a fully formed, rather than half-formed, Earth.

The idea that the Moon is a by-product of a stupendous collision between the early Earth and a Mars-size protoplanet has been considered respectable science for more than a decade. It is now the consensus view¹, for no other theory reconciles so many facts about the Earth and Moon (Fig. 1) and provides a relatively seamless explanation of how our large satellite came to be. But there have

been few attempts to directly model the complex physical processes occurring during a planetary impact. Cameron² was one of the first people to suggest the giant impact theory, Hartmann³ being the other principal proponent of the idea in the early days. Cameron and his collaborators have since led the efforts to model the impact numerically, and their view of events has been highly influential.

Now, however, thanks to the advance of computer technology and the development of better impact algorithms, other groups are also able to simulate collisions between planets. On page 708 of this issue⁴, Canup and Asphaug usher in this new era by describing the outcome of the giant impact in terms that in many respects differ from Cameron's latest results⁵.

Computer simulation of a planetary-scale impact is not a task for faint-hearted, point-and-click computer modellers. Not only does such a collision involve all the details of shock physics, melting and vaporization, but the mutual gravitational interactions of all those hot fluids squirting around in space also have to be taken into account. Furthermore, to impart the angular momentum that keeps the Moon in orbit, the impact must have been glancing, not head on, so the problem must be treated in three dimensions. It has taken squadrons of physicists in the United States, Russia and elsewhere nearly 50 years to come up with computers and three-dimensional computer codes that can adequately treat the effects of impacts and explosions under relatively simple conditions in which self-gravity is not important. Adding self-gravity to these codes therefore posed a formidable challenge.

A solution recently arose from an approach⁶ called SPH (smooth particle hydrodynamics), which can, in principle, treat both shock physics and gravity. Benz and Cameron were the first⁷ to apply this method to the origin of the Moon. Other early models^{8,9} relied on conventional 'hydrocodes' that used central gravity, but not self-gravity, and so could accurately model only the early phases of the impact.

The SPH method involves splitting the proto-Earth and the object it collided with into many small computational lumps called 'particles', then computing the interactions between these entities. The results of this procedure become more accurate as more particles are used. Because of computer limitations, the early simulations used only a few thousand particles. It was several years before users recognized that the SPH method requires many more of these particles than originally anticipated. One of the reasons that Canup and Asphaug⁴, who also use an SPH code, achieved results that differ from those of Cameron *et al.* is that they use at least 20,000 particles and so resolve the shock waves and other collision phenomena with much higher accuracy.

Based on his latest simulations⁵, Cameron has held that the collision that created the Moon occurred when the Earth was only about half-formed, about 30 million years after the formation of the first meteorites, as it grew by accretion of material in the early Solar System. This result makes it difficult to understand how the Earth continued to grow without large amounts of