

method work effectively. They do this by making good use of recently developed machine learning algorithms for a special class of neural networks (5, 6).

Hinton and Salakhutdinov's approach uses so-called autoencoder networks—neural networks that learn a compact description of data, as shown in the middle panel of the figure. This is a neural network that attempts to learn to map the three-dimensional data from the spiral down to one dimension, and then back out to three dimensions. The network is trained to reproduce its input on its output—an identity mapping—by the standard backpropagation of error method (7, 8). Although backpropagation is a supervised learning method, by using the input as the teacher, this method becomes unsupervised (or self-supervised). Unfortunately, this network will fail miserably at this task, in much the same way that standard methods such as principal components analysis will fail. This is because even though there is a weighted sum of the inputs (a linear mapping) to a representation of x —the location along the spiral—there is no (semi-)linear function (9) of x that can decode this back to $\sin(x)$ or $\cos(x)$. That is, the network is incapable of even representing the transformation, much less learning it. The best such a network can do is to learn the average of the points, a line down the middle of the spiral. However, if another nonlinear layer is added between the output and the central hidden layer (see the figure, right panel), then the network is powerful enough, and can learn to encode the points as one dimension (easy) but also can learn to decode that one-dimen-

sional representation back out to the three dimensions of the spiral (hard). Finding a set of connection strengths (weights) that will carry out this learning problem by means of backpropagation has proven to be unreliable in practice (10). If one could initialize the weights so that they are near a solution, it is easy to fine-tune them with standard methods, as Hinton and Salakhutdinov show.

The authors use recent advances in training a specific kind of network, called a restricted Boltzmann machine or Harmony network (5, 6), to learn a good initial mapping recursively. First, their system learns an invertible mapping from the data to a layer of binary features. This initial mapping may actually increase the dimensionality of the data, which is necessary for problems like the spiral. Then, it learns a mapping from those features to another layer of features. This is repeated as many times as desired to initialize an extremely deep autoencoder. The resulting deep network is then used as the initialization of a standard neural network, which then tunes the weights to perform much better.

This makes it practical to use much deeper networks than were previously possible, thus allowing more complex nonlinear codes to be learned. Although there is an engineering flavor to much of the paper, this is the first practical method that results in a completely invertible mapping, so that new data may be projected into this very low dimensional space. The hope is that these lower dimensional representations will be useful for important tasks such as pattern recognition, transformation, or

visualization. Hinton and Salakhutdinov have already demonstrated some excellent results in widely varying domains. This is exciting work with many potential applications in domains of current interest such as biology, neuroscience, and the study of the Web.

Recent advances in machine learning have caused some to consider neural networks obsolete, even dead. This work suggests that such announcements are premature.

References and Notes

1. G. E. Hinton, R. R. Salakhutdinov, *Science* **313**, 504 (2006).
2. S. T. Roweis, L. K. Saul, *Science* **290**, 2323 (2000).
3. J. A. Tenenbaum, V. J. de Silva, J. C. Langford, *Science* **290**, 2319 (2000).
4. One can learn a mapping to the manifold (and back), but this is done independently of the original structure-finding method, which does not provide this mapping.
5. G. E. Hinton, *Neural Comput.* **14**, 1771 (2002).
6. P. Smolensky, in *Parallel Distributed Processing*, vol. 1, Foundations, D. E. Rumelhart, J. L. McClelland, PDP Research Group, Eds. (MIT Press, Cambridge, MA, 1986), pp. 194–281.
7. D. E. Rumelhart, G. E. Hinton, R. J. Williams, *Nature* **323**, 533 (1986).
8. G. W. Cottrell, P. W. Munro, D. Zipser, in *Models of Cognition: A Review of Cognitive Science*, N. E. Sharkey, Ed. (Ablex, Norwood, NJ, 1989), vol. 1, pp. 208–240.
9. A so-called semilinear function is one that takes as input a weighted sum of other variables, and applies a monotonic transformation to it. The standard sigmoid function used in neural networks is an example.
10. D. DeMers, G. W. Cottrell, in *Advances in Neural Information Processing Systems*, S. J. Hanson, J. D. Cowan, C. L. Giles, Eds. (Morgan Kaufmann, San Mateo, CA, 1993), vol. 5, pp. 580–587.

10.1126/science.1129813

ATMOSPHERE

What Drives the Ice Age Cycle?

Didier Paillard

The exposure of Earth's surface to the Sun's rays (or insolation) varies on time scales of thousands of years as a result of regular changes in Earth's orbit around the Sun (eccentricity), in the tilt of Earth's axis (obliquity), and in the direction of Earth's axis of rotation (precession). According to the Milankovitch theory, these insolation changes drive the glacial cycles that have dominated Earth's climate for the past 3 million years.

For example, between 3 and 1 million years before present (late Pliocene to early Pleistocene, hereafter LP-EP), the glacial oscillations followed a 41,000-year cycle. These oscillations

correspond to insolation changes driven by obliquity changes. But during this time, precession-driven changes in insolation on a 23,000-year cycle were much stronger than the obliquity-driven changes. Why is the glacial record for the LP-EP dominated by obliquity, rather than by the stronger precessional forcing? How should the Milankovitch theory be adapted to account for this "41,000-year paradox"?

Two different solutions are presented in this issue. The first involves a rethinking of how the insolation forcing should be defined (1), whereas the second suggests that the Antarctic ice sheet may play an important role (2). The two papers question some basic principles that are often accepted without debate.

On page 508, Huybers (1) argues that the summer insolation traditionally used in ice age models may not be the best parameter. Because

Between 3 and 1 million years ago, ice ages followed a 41,000-year cycle. Two studies provide new explanations for this periodicity.

ice mass balance depends on whether the temperature is above or below the freezing point, a physically more relevant parameter should be the insolation integrated over a given threshold that allows for ice melting. This new parameter more closely follows a 41,000-year periodicity, thus providing a possible explanation for the LP-EP record.

On page 492, Raymo *et al.* (2) question another pillar of ice age research by suggesting that the East Antarctic ice sheet could have contributed substantially to sea-level changes during the LP-EP. The East Antarctic ice sheet is land-based and should therefore be sensitive mostly to insolation forcing, whereas the West Antarctic ice sheet is marine-based and thus influenced largely by sea-level changes. Because the obliquity forcing is symmetrical with respect to the hemispheres, whereas the preces-

The author is at the Laboratoire des Sciences du Climat et de l'Environnement, Institut Pierre Simon Laplace, CEA-CNRS-UVSQ, 91191 Gif-sur-Yvette, France. E-mail: didier.paillard@cea.fr