- Tuesday, April 2, 2002
  Contrastive Backpropagation

  The standard backpropagation learning procedure requires that each training input be accompanied by a specification of the correct output. To apply backpropagation to the unsupervised learning of internal representations of sensory data, it is necessary to find a way of doing without the supervision signal. I shall first describe a way of modelling high-dimensional data distributions, such as ensembles of images, by using a multilayer network in which the activity of each hidden unit represents a"goodness" or "badness" that makes an additive contribution to the log probability that the model assigns to the input vector. Standard maximum likelihood methods cannot be used to train models of this kind because they require intractable computations. There is, however, a fairly efficient new learning procedure developed by Hinton, Teh, Welling and Osindero that changes each parameter in proportion to the difference between two goodness derivatives. One derivative is computed by using backpropagation through the multilayer network when the input is real data. The other derivative is computed when the input has been slightly corrupted to fit better with the network's current model. The learning procedure eliminates the tendency of the model to prefer corruptions of the data to the data itself. I shall demonstrate the effectiveness of the new learning procedure on a number of unsupervised learning tasks and show that when it is presented with patches of natural images it learns topographic maps that represent orientation, location and spatial frequency in a sensible way.

  Geoffrey Hinton
  Comp. Sci., University of Toronto