

frisbee. Take any four frisbees and clamp one to each wall. There will be an oscillation set up as those four frisbees pull on the other ones, and those pull on each other. In time, the oscillations will cease, and the frisbee population will settle into a pattern that reflects an equilibrium among the tensions exerted by the elastic bands. (p. 12)

If one were to change which frisbees are clamped to the wall or the lengths or thicknesses of the bands connecting the frisbees, oscillation would reoccur and a new pattern would settle out.

Because connections among units are assumed to carry different weights of association, learning occurs in the continual adjustment of the weights. Moreover, since processing occurs in parallel, many different adjustments can take place simultaneously, and there can be continuous error adjustments as a function of new information.

Consider how a PDP model might account for the experiences of Wes and Marcy in *The Mechanic* and the *Web Surfer*. In Marcy's case, the units and connections representing her knowledge of car mechanics are likely to be neither extensive nor strong, but some are already stronger than others. It is probably safe to assume, for example, that connections related to steering are stronger than those related to tie rods. Marcy's conversation with Wes serves to activate and strengthen further some of those connections and perhaps introduces new connections (e.g., steering damper may be a new concept to her, although both "steering" and "damper" are familiar). When it comes to recalling the conversation later, then, the stronger connection weights associated with "steering" enable Marcy to remember that as the gist of what was said. Likewise, the very weak connection weights associated with "steering damper" are not enough to prompt its specific recall. A similar analysis could be applied to Wes and what he remembers about the Internet.

PDP models offer a number of advantages over the other models in terms of what they explain about human information processing. First, they seem to account well for the incremental nature of human learning. With constant readjustment of connection weights, they provide a more dynamic picture of human learning than has been suggested heretofore (Estes, 1988). PDP models also offer "for the first time a convenient way of incorporating goals into the dynamics of information processing systems" (Estes, 1988, p. 207). That is, connection weights in most PDP systems are adjusted to reduce disparity between their output and some target output, which may be viewed as a goal.

Finally, there is potential in PDP models to explain cognitive development (McClelland, 1988, 1995). Some knowledge, in terms of prewired connection weights, can be built into the network. Exploring different configurations of initial memory architecture may lead to breakthroughs in determining just how much of human memory is "hard-wired."

Estes (1988) sounded some cautionary notes, however, concerning the conclusions over the long term to which PDP models may lead. He cited the