

Figure 8.1
 A network that has learned to detect mirror symmetry in the input vector. The numbers on the arcs are weights and the numbers inside the nodes are biases. The learning required 1,425 sweeps through the set of 64 possible input vectors, with the weights being adjusted on the basis of the accumulated gradient after each sweep. The values of the parameters in equation (9) were $\epsilon = 0.1$ and $\alpha = 0.9$. The initial weights were random and were uniformly distributed between -0.3 and 0.3. The key property of this solution is that for a given hidden unit, weights that symmetric about the middle of the input vector are equal in magnitude and opposite in sign. So if a symmetrical pattern is presented, both hidden units will receive a net input units, and, because the hidden units have a negative bias, both will be off. In this case the output unit, having a positive bias, will be on. Note that the weights on each side of the midpoint are in the ratio 1:2:4. This ensures that each of the eight patterns that can occur above the midpoint sends a unique activation sum to each hidden unit, so the only pattern below the midpoint that can exactly balance this sum is the summetrical one. For all non-symmetrical patterns, both hidden units will receive non-zero activations from the input units. The two hidden units have identical patterns of weights but with opposite signs, so for every non-symmetric pattern one hidden unit will come on and suppress the output unit.




Figure 8.2
 Two isomorphic families of relations: <person 1> <relation 1> <person 2>. The first family is: band, wife, son, daughter. The second family is: 'know' these triples in two terms are encoded by activation levels in a proposition by activation.




Figure 8.3
 Activity levels in a network for representing <person 1> <relation 1> <person 2>. Each of the two groups of white squares inside the first group represent 'aunt'. Each of the two groups of black dots inside the second layer represent 'uncle'. The second layer is the penultimate layer of units, each of which is marked by black dots out spatially with the