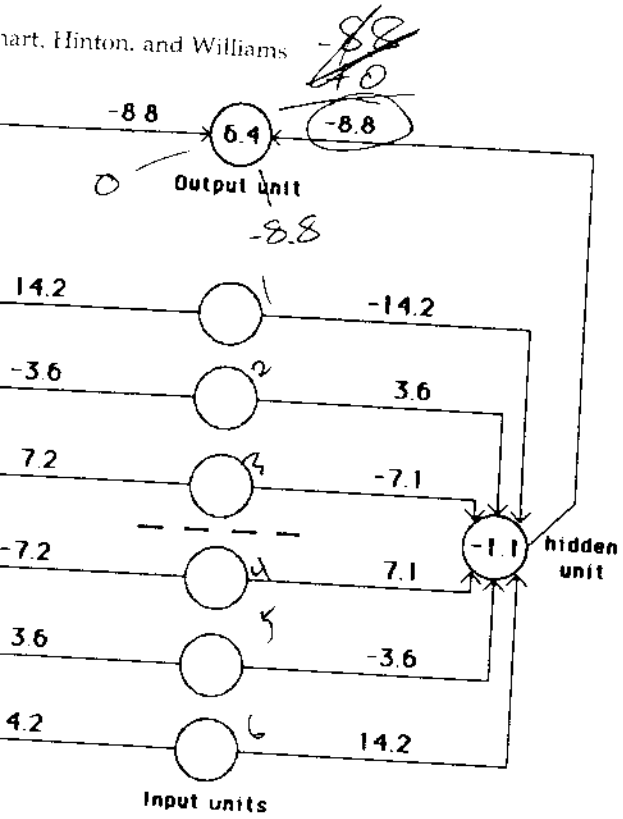

earned to detect mirror symmetry in the input vector. The numbers on the arcs numbers inside the nodes are biases. The learning required 1,425 sweeps possible input vectors, with the weights being adjusted on the basis of the after each sweep. The values of the parameters in equation (9) were $\epsilon = 0.5$ al weights were ramdom and were uniformly distributed between -0.3 and of this solution is that for a given hidden unit, weights that symmetric about ut vector are equal in magnitude and opposite in sign. So if a symmetrical both hidden units will receive a net input units, and, because the hidden units both will be off. In this case the output unit, having a positive bias, will be on on each side of the midpoint are in the ratio 1:2:4. This ensures that each of the occur above the midpoint sends a unique activation sum to each hidden unit ow the midpoint that can exactly balance this sum is the summetrical one. For atterns, both hidden units will receive non-zero activations from the input units have indentical patterns of weights but with opposite signs, so for every one hidden unit will come on and suppress the output unit.

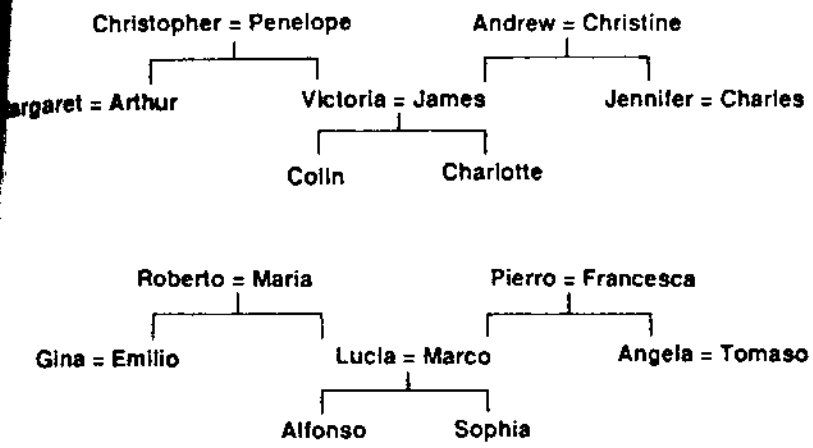

Figure 8.2
Two isomorphic family trees. The information can be expressed as a set of triples of the form ⟨person 1⟩ ⟨relationship⟩ ⟨person 2⟩, where the possible relationships are {father, mother, husband, wife, son, daughter, uncle, aunt, brother, sister, nephew, niece}. A layered net can be said to 'know' these triples if it can produce the third term of each triple when given the first two. The first two terms are encoded by activating two of the input units, and the network must then complete the proposition by activating the output unit that represents the third term.

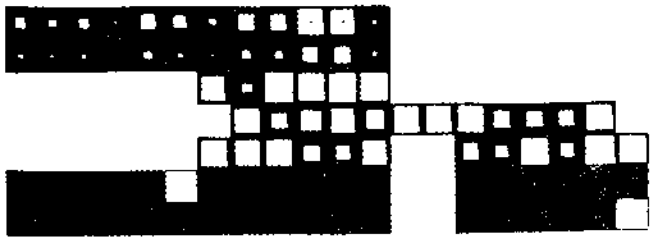

Figure 8.3
Activity levels in a five-layer network after it has learned. The bottom layer 24 input units on the left for representing ⟨person 1⟩ and 12 input units on the right for representing the relationship. The white squares insidethese two groups show the activity levels of the units. There is one active unit in the first group representing Colin and one in the second group representing the relationship 'has-aunt'. Each of the two input groups is totally connected to its own group of 6 units in the second layer. These groups learn to encode people and relationships as distributed patterns and activity. The second layer is totally connected to the central layer of 12 units, and these are connected to the penultimate layer of 6 units. The activity in the penultimate layer must activate the correct output units, each of which stands for a particular ⟨person 2⟩. In this case, there are two correct answers (marked by black dots) because Colin has two aunts. Both the input units and the output units laid out spatially with the English people in one row and the isomorphic Italians immediately below.