designer of the net has in mind a specific definition of "type" required for a reasonable generalization, and counts it a success if the net generalizes to other instances of this type. All the "continue this sequence" questions found on intelligence tests, for example, really have more than one possible answer but most human beings share a sense of what is simple and reasonable and therefore acceptable. But when the net produces an unexpected association can one say it has failed to generalize? One could equally well say that the net has all along been acting on a different definition of "type" and that that difference has just been revealed.

For an amusing and dramatic case of creative but unintelligent generalization, consider the legend of one of connectionism's first applications. In the early days of the perceptron the army decided to train an artificial neural network to recognize tanks partly hidden behind trees in the woods. They took a number of pictures of a woods without tanks, and then pictures of the same woods with tanks clearly sticking out from behind trees. They then trained a net to discriminate the two classes of pictures. The results were impressive, and the army was even more impressed when it turned out that the net could generalize its knowledge to pictures from each set that had not been used in training the net. Just to make sure that the net had indeed learned to recognize partially hidden tanks, however, the researchers took some more pictures in the same woods and showed them to the trained net. They were shocked and depressed to find that with the new pictures the net totally failed to discriminate between pictures of trees with partially concealed tanks behind them and just plain trees. The mystery was finally solved when someone noticed that the training pictures of the woods without tanks were taken on a cloudy day, whereas those with tanks were taken on a sunny day. The net had learned to recognize and generalize the difference between a woods with and without shadows! Obviously, not what stood out for the researchers as the important difference. This example illustrates the general point that a net must share size, architecture, initial connections, configuration and socialization with the human brain if it is to share our sense of appropriate generalization.

There was also a further problem. The purely associationistic pattern recognition model of learning, adopted by the connectionists, could not explain expert consensus. That is, the connectionist model of the acquisition of the ability to behave intelligently failed to account for the important fact that even though each expert has been exposed to different cases of success and failure in different sequences, experts tend to agree in their response to a given situation.

To address the above two problems we need to again ask the question: How *does* an expert cope intelligently with a domain? Only when we understand this will we have a basis for speculating about the possibilities and limitations of artificial experts produced by neural networks.

Clearly, experience improves coping performance. In considerably lower animals it is fairly certain that trial-and-error experience directly produces synaptic and related brain changes causing raw stimulae detected by the sense organs to map into better and better physical coping responses. The changes that occur during learning almost certainly cannot be even approximately described at some higher level of abstraction such as belief, goal, or mental-domain-model modification.