Chapter IV
An Introduction to Neurocomputing

SIMULATING THE BRAIN

Let's build an automated archaeologist!

It is not an easy task. We need a highly complex, nonlinear, and parallel information-processing "cognitive core" able to explain what the robot sees, in terms of causal factors, which not always have an observable nature.

Of course, such a "cognitive core" should not run like a human brain. After all, automated archaeologists do the same tasks as "human archaeologists," but not necessary in the same way. Nevertheless, there is some similitude in the basic mechanism. My suggestion is that an archaeologist, human or "artificial," will perceive archaeological data and, using some basic principles of learning, as those presented in previous chapter, will develop ways of encoding these data to make sense of perceived world. Consequently, we may try to build our artificial archaeologist based on the idea of *learning* and the ability to adapt flexibly epistemic actions to different archaeological problems waiting for a solution.

How much should be programmed in its final form into such a cognitive core and how much will have to be learnt by interacting with some environment, including teachers and other agents? Projects aiming to develop intelligent systems on the basis of powerful and general learning mechanisms start from something close to a "Tabula rasa," however, they risk being defeated by explosive search spaces requiring evolutionary time-scales for success. Biological evolution enables animals to avoid this problem by providing large amounts of "innate" information in the genomes of all species. In the case of humans, this seems to include meta-level information about what kinds of things are good to learn, helping to drive the learning processes as well as specific mechanisms, forms of representation, and architectures to enable them to work. Is it possible to use these ideas for building an "intelligent" machine?

Like its human counterpart, the cognitive core of our automated archaeologist should be made of specialized cells called *neurons* (Figure 4.1). Artificial and biological neurons are relatively similar, and both have the same parts, also called the cell body, axon, synapse, and dendrite (Bechtel & Abrahamson, 1991; Dawson, 2004; Ellis & Humphreys, 1999; O'Reilly & Munakata, 2000; Quinlan, 1991).

Each neuron connects as well as accepts connections from many other neurons, configuring a network of neurons. Those connections are implemented by means of dendrites, while syn-

Figure 4.1. Schematic representation of a neuron

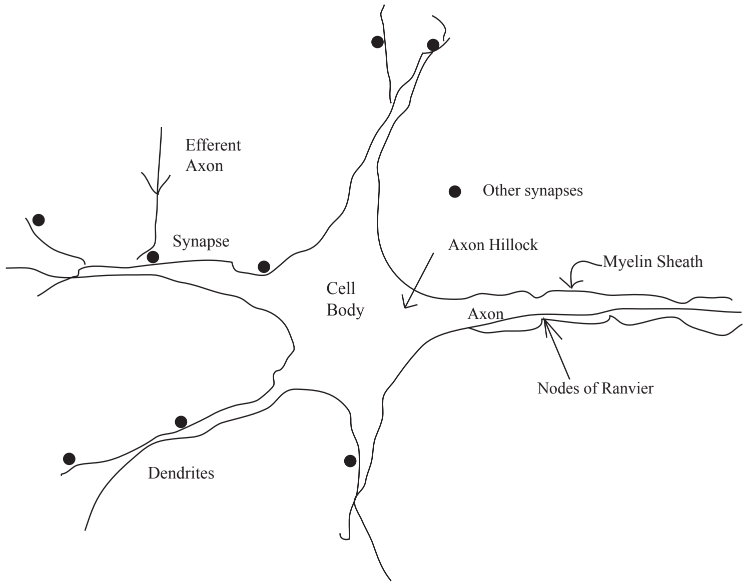

apses are a gateway linked to dendrites coming from other neurons.

We can think about the essential function of each neuron in the network from a computational perspective in terms of a *detector*. First, a detector needs *inputs* that provide the information on which it bases its detection. In human brain, information is expressed in the timing and the frequency neurons communicate among them through electrical pulses. By combining or integrating activation signals or pulses over all the incoming connections (dendrites), each neuron creates some sort of aggregate measure. As a result, the neuron produces a new composite signal, the *output*, transmitted to other neurons, continuing the information-processing cascade through a network of interconnected neurons (Figure 4.2). The chaining of multiple levels of detectors can lead to more powerful and efficient detection capabilities than if everything had to work directly from the raw sensory inputs. However, this chaining implies that the transformation operation is complex because different signals arrive from different sources through different connections, and each connection modifies the information in a particular way. This style of computing—transforming one pattern into another by passing it through a large configuration of synaptic connections—is called *parallel distributed processing*. As the original input pattern distributed across many neurons pass inward from one specialized neural population to the next, and to the next and the next, the original pattern is progressively transformed at each stage by the intervening configuration of synaptic configurations.

On a neural network, the overall pattern of simultaneous activation levels across the assembled neurons of a given population is the primary unit of representation, and the primary vehicle of semantic content. Such patterns are often referred to as "activation vectors" because they can be usefully and uniquely characterized by a sequence of n numbers, where n = the number of neurons in the representing population. Consequently, concepts may be represented as ephemeral patterns of activation across an entire set of units rather than as individuated elements or symbols. These stable patterns then determine further process-