

1.2 DEFINING THE PROBLEM

What would it mean to say that we “understand” the neural code in a particular region of the nervous system? How do we quantify the notion that the spike train of a single cell “conveys information” about the sensory world? In what sense is a particular sequence of spikes the “right answer” to some computational problem faced by the brain? We search for sharper versions of these questions by forcing ourselves to adopt a more precise and more mathematical language. In talking about the nervous system we routinely make colloquial use of terms such as code, information, and reliability. All of these words can be given precise mathematical definitions, and we hope, through the remainder of the text, to convince the reader that these definitions provide a clear guide to the design and analysis of new experiments. In striving for precision we shall see the emergence of some new ideas. We begin, however, by revisiting an old idea, the homunculus.

The homunculus is an often derided concept in discussions of the brain. We recall that this metaphor conjures up a little man—or, in a lovely variant by Michael Land and Simon Laughlin (Fig. 1.7), a little fly—who observes the responses of his own sensory neurons and finally forms the percepts that the organism experiences. The problem with this picture is that it never gets to the essence of what it means to perceive and to experience the world. On the other hand, as explorers of the nervous system we place ourselves, inevitably, in the position of the homunculus—we observe the responses of sensory neurons and try to decide what these responses could mean to the organism. This problem of assigning meaning to the activity of sensory neurons is the central issue in our discussion of the neural code.

It is easy to imagine that the task of the homunculus is trivial—after all, he just watches a projected image of the world as it flashes through the brain. But this projected image is *encoded* in the patterns of action potentials generated by the sensory neurons. It is not at all clear what the homunculus would have

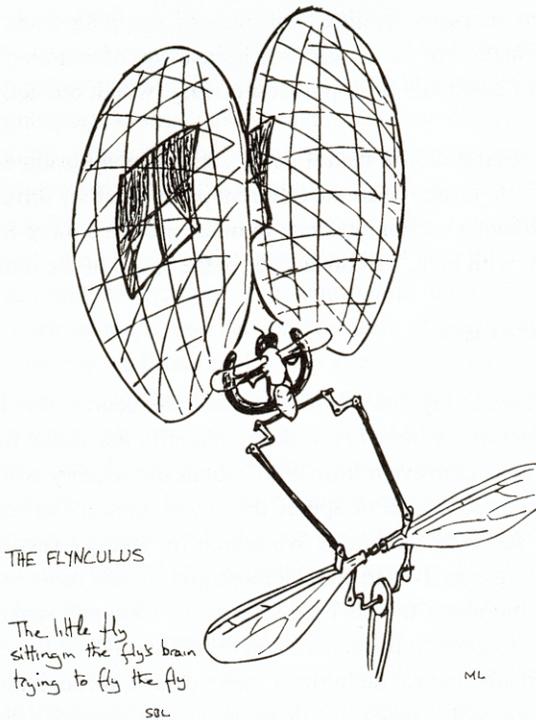


Figure 1.7

The Flyculus. Doodle by M. F. Land, quotation, “The little fly sitting in the fly’s brain trying to fly the fly,” from S. B. Laughlin, with permission.

to do, even in principle, to make sense out of these encoded data. We propose that “understanding the neural code” means that we would know how to make sense out of the bewildering array of spike trains streaming in from the sense organs: If we understand the code, we can function as the homunculus.

When we ask what a spike train means, or what it can tell us about the world, we need to set some boundaries for the question, or, equivalently, a context for the answer. If we live in a world with only two possible sensory stimuli, we can ask how the homunculus could best use the spike train data to make a decision about which stimulus in fact occurred. This decision rule would constitute a complete understanding of the neural code, assuming that the world offers just two possible signals.

In many psychophysical discrimination experiments (Green and Swets 1966), a world of two alternatives is created artificially, and the subject must

solve the problem of choosing between these alternatives. This binary decision problem provides a convenient context for asking questions about the reliability of our perceptions, and we shall see that it is also useful for investigating the reliability of neurons. But it is not enough to build a homunculus that functions in a world of two alternatives; we want to ask our question about the meaning of spike trains in a context that approaches the complexity of the natural world.

Under natural conditions, the stimulus that will appear in the next brief time window is not known to us in advance. Instead the stimulus is chosen from an infinite set of alternatives. On the other hand, these alternatives are not all equally likely. While there are blue spruce trees, green trees do not suddenly turn blue (or red or yellow either). Natural stimuli develop in time, and these dynamics have some underlying regularity or structure. This structure has a deterministic component, as when a leaf falls downward according to Newton's laws. But since we do not know all the forces that shape the dynamics of sensory stimuli, some aspects of these stimuli are unpredictable, as when the falling leaf is deflected by a gust of wind. The result is that natural signals are presented to us at random, but these signals have correlations that reflect their origins in deterministic physical processes.

Rather than inhabiting a world of two alternatives, we thus inhabit a world of random but correlated time dependent signals. The time dependence is crucial, because it means that we cannot wait forever to decide what we are looking at. Not only does biology press for quick decisions—we must catch our prey and not be caught by predators—the physics of our environment is such that any simple averaging for long periods of time will average away the very signals that interest us. The task of the homunculus, then, is not to create a static image of the sensory world from the input spike trains, but rather to give a sort of running commentary or simultaneous translation. We emphasize that this running commentary need not be, and most likely cannot be, a comprehensive reconstruction of the world around us.

To give meaning to the spike trains nonetheless requires that we recreate at least some aspects of the continuous time dependent world that is encoded in discrete sequences of spikes. From our experience in the laboratory we know that when forced to interpret rapidly changing signals we are very susceptible to noise; usually we try to combat noise by averaging in time or by averaging over repeated presentations of the same signal. But the homunculus is not free to set arbitrary averaging times, and he certainly cannot ask for a second, identical copy of the immediate past. On the contrary, the homunculus (and

the animal as well!) has to reach conclusions about the world from just one example of the spike train in each of his sensory neurons.

In generating a running commentary on the meaning of spike trains we shall have to deal with whatever level of noise is present in these data. Ideally, our interpretation of the spike trains should be as reliable as possible given this noise, and the statistically sophisticated homunculus would report confidence levels on his estimates of what is happening in the world. If understanding the neural code means building a homunculus, we can compare two different candidate homunculi—two different candidates for the structure of the neural code—by comparing the accuracy of their inferences about events in the sensory world.

We are closing in, then, on a more precise definition of the problems in understanding the neural code. We place ourselves in the position of the homunculus, monitoring the spike trains of sensory neurons as stimuli vary in time along some unknown trajectory. We must generate a running commentary on the identity of these stimuli, using only the spike train data as input. Our inferences about events in the world will have some limited accuracy, and we shall have to quantify this accuracy. Out of many possible homunculi, there is one that tells us as much as possible about the world given the noise in the spike train data itself. The performance of this best homunculus will reflect a compromise between averaging in time to combat noise and responding quickly to keep up with the dynamics of the world, and we shall have to be precise about these time scales.

The construction of a complete homunculus, or even the complete flynculus of Fig. 1.7, is a daunting task. In the fly, visual signals stream in along thousands of parallel paths reflecting the array of lenses in the compound eye, and in ourselves and our primate cousins the corresponding numbers are three orders of magnitude larger. There are a few special cases, such as the moths discussed in section 4.1.1 (Roeder 1963), for which it might be possible to monitor all of the spike trains that encode one sensory modality, but in general this is hopeless. As we have noted, however, there is a long tradition of trying to make sense out of the responses of single neurons, always recognizing that one cell can tell us about only a small piece of the sensory environment. In this tradition, most of this book is about the problem of an impoverished homunculus who looks at the spike train of just one neuron at a time; we take a brief look at the problem of multiple neurons in section 5.1. We thus have a clear question, amenable to experimental investigation: What can the spike train of this one neuron tell us about events in the world?

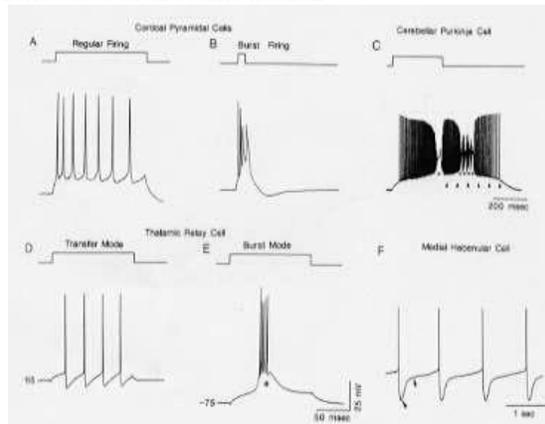
5 Neural Coding

5.1 Receptive fields

- A neuron responds to a specific set of input patterns, known as its receptive field
 - retinal ganglion cells respond to spots of light in particular places of the visual field
 - *simple* cells in area V1 of the visual cortex respond to bars of light of particular orientation in particular places of the visual field
 - *complex* cells in V1 respond to oriented bars moving in particular directions
- What are high-level receptive fields?
 - this is still much debated
 - *grandmother* cell hypothesis suggests that the highest-level receptive fields are quite specific i.e. a neuron that only responds to your grandmother
 - but do we need cells that respond to your grandmother from the front, from the back, sitting, walking etc?
 - there is evidence for neurons that only respond to faces, and even faces with particular expressions
 - we do not have enough neurons to code for every feature of every visual scene we encounter in our lives, so neurons must code for some level of general features
 - individual neurons may die during our lifetimes, so there must be multiple neurons coding for similar features for robustness

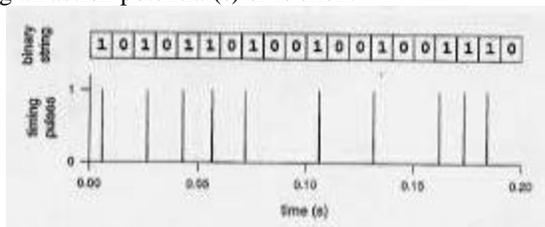
5.2 Single neuron codes

- Output of a neuron is a sequence of action potentials (APs)
 - the temporal pattern of action potential firing can be very different between neuronal types and for different states of the same neuron



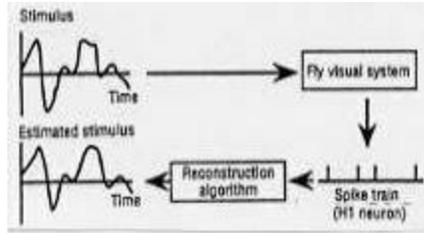
[SH90 Fig. 2.5 pg 47]

- How might information be encoded by APs?
 - binary: either firing an action potential(s) or is silent



[RWRB97 Fig. 3.4 pg 114]

- rate coding: average firing frequency over some time period
 - muscle afferents encode muscle length and rate of change in length by firing rate
 - muscle force is generated as a function of motoneuron firing rate
- interspike interval codes: temporal sequence of exact spike times
 - H1 neuron of fly visual system encodes angular velocity of object in visual field by interspike intervals

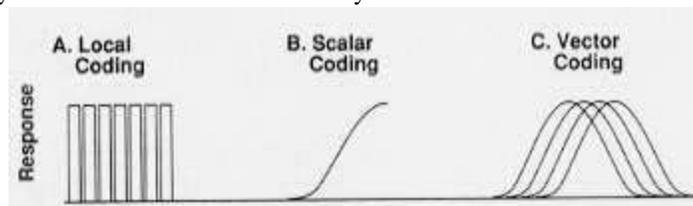


[BRRW91 Fig. 1 pg 1855]

- Dynamic range
 - suppose a neuron can fire from 0 to 200 spikes per second, then how many different neuronal states can be discriminated in 100msecs?
 - if the code is average firing rate, then only approximately 20 different states (rates) can be distinguished by counting spikes in a 100msec period
 - if the code is a binary vector derived by identifying if a spike was fired in each of 20 5msec time bins, then there are 2 to the power of 20 possible states
 - since we have reaction times for visual object recognition of the order of 150msecs, the time available for distinguishing the state of a single neuron is very short (20-30msecs at most)

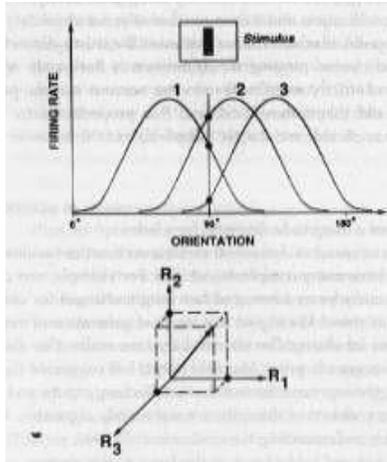
5.3 Population neuronal codes

- Output of a network is a pattern of activity across the population of neurons
- Information may be encoded in three different ways:



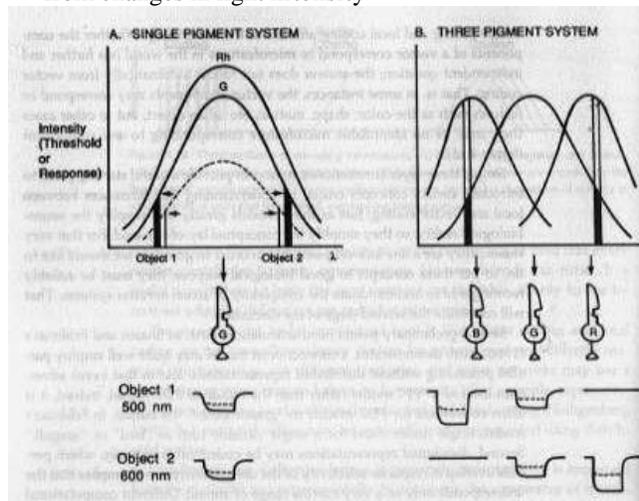
[CS92 Fig. 4.14 pg 164]

- *local coding*: each neuron represents a specific feature that the system distinguishes
- *scalar coding*: firing rate of each neuron encodes a feature
 - redundancy and improved signal-to-noise ratio can be achieved by several neurons coding the same features
- *vector coding*: feature encoded in the firing rates of a subpopulation of neurons that have overlapping tuning curves in the feature space
 - an example is coding for object orientation in the visual system
 - certain neurons respond maximally to bars of light of particular orientation
 - for any given bar, many neurons will respond with different firing rates, depending on how closely the orientation of the bar matches their preferred orientation
 - this *vector* of neuronal firing rates more accurately encodes the exact orientation of the bar than does the firing rate of a single neuron



[CS92 Fig. 4.21 pg 178]

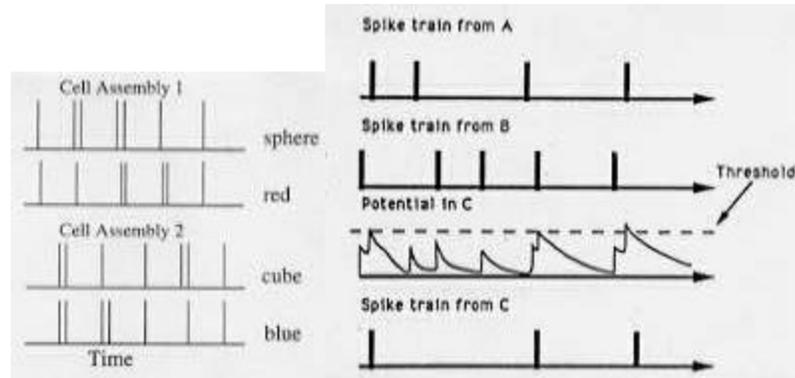
- another example is 3-pigment system for colour coding in the retina
 - 3 types of cones respond preferentially to different frequencies (colour) of light
 - *vector* of cone responses from a population of cones including all three types is needed to encode colour and distinguish changes in colour from changes in light intensity



[CS92 Fig. 4.15 pg 166]

5.4 Temporal binding

- A single object in the real world may be encoded in the synchronous firing of neurons that code for each separate feature of the object
- In this way multiple objects can be represented simultaneously and distinguished by the neurons representing one object firing out of phase with the neurons representing another object
- Coincidence detection by a downstream neuron results in that neuron responding to higher level features e.g. red spheres



[MB99 Fig. 4.4 pg 117; ??]

5.5 Synchronization and oscillations

- There is evidence for an increase in synchrony between cortical neurons responding to a stimulus, without necessarily any change in their average firing rates
 - such synchrony is seen in nearby neurons, between neurons in different cortical areas and even across hemispheres
 - an example is in the auditory system, where in response to a sound stimulus, neurons do not change their average firing rate, but they do fire more in synchrony with each other - this synchrony may be detected by downstream neurons, thus recognising the presence of the sound
- Synchrony may coincide with, but does not depend on, oscillations in neural activity
 - oscillations in the gamma frequency range (20-70Hz) are seen in many cortical and subcortical regions
 - many other oscillation frequencies are seen in different areas and in different circumstances
 - theta frequency (7-12Hz) oscillations in the hippocampus are associated with exploratory behaviour in rats
 - such oscillations may simply be epiphenomena resulting from circuit dynamics
 - however, they can also act as clock signals
 - information about spatial location is contained in the firing rates of cells in the hippocampus (*place cells*)
 - BUT more spatial information is contained in the *phase* of place cell firing relative to the theta rhythm.

5.6 Places to start

- Churchland and Sejnowski (1992) *The Computational Brain*, MIT Press
- Rieke, Warland, de Ruyter van Steveninck and Bialek (1997) *Spikes*, MIT Press
- Maass and Bishop (eds) (1999) *Pulsed Neural Networks*, MIT Press