

Informatics 1 Cognitive Science – Tutorial 7 Solutions

Matthias Hennig


Week 8

Part 1


In the lectures you have seen that neurons in the visual system primarily respond to changes in visual stimuli. This is true for all sensory systems, neurons usually signal changes rather than producing a constant output that corresponds to the intensity of an ongoing, unchanging stimulus.

Neurons can only produce a limited range of outputs as their firing rates range from zero spikes to a maximum of usually around 30 to 50 spikes per second. This limits the ability of a neuron to encode sensory inputs.

Let's consider a simple situation where a neuron is sensitive to the intensity of a stimulus (say the amplitude of a sound). Now we consider what happens when the overall distribution of amplitudes changes, for instance when we move from a quiet into a noisy environment:



The transfer function of a neuron, which describes the normalised spike rate as a function of the stimulus intensity, may be described by a logistic function. An example looks as follows:




Questions

1. Is the transfer function shown above better suited to encode stimuli in the quiet or noisy environment? Why do you think this is the case?

Solution: The quiet environment as the sensitivity is matched well to the distribution of intensities.

2. Neurons can adapt their transfer function (either through intrinsic mechanisms, or through network effects) to changing input statistics. How should this transfer function change to capture the inputs in the other environment?

Solution: The curve should move to the right to cover higher intensities, and the slope should be reduced to cover the wider range of intensities. In fact the transfer function should match the cumulative distribution of inputs, and the part with the strongest gradient should cover the most frequent stimulus intensities so they can be discriminated best.



3. If this adaptation takes place in the early sensory system, the brain never receives any information about the absolute magnitude of the stimuli. However, we usually can tell if we're in a quiet or loud environment, or if it is bright sunlight or cloudy. How is this possible?

Solution: We make use of other cues to infer this. For instance, traffic noise has certain characteristics that we can recognise, and different lighting creates different shadows etc. Often we're however indeed not aware of the ambient features of our environment. For instance, we quickly adapt to smells from the kitchen that someone just entering the house will immediately notice.

Part 2

The mathematician George Box is reported to have said that “All models are wrong, but some are useful”. In the lectures so far you have seen a wide range of models, ranging from abstract models of language and categories all the way down to models of single neurons in the brain. According to James McClelland, “The essential purpose of cognitive modeling is to allow investigation of the implications of ideas, beyond the limits of human thinking.” (see *The place of modeling in cognitive science*, Topics in Cognitive Science, 2009). So let's discuss how useful models are in cognitive science.

Questions

1. Make a list of different models the course has covered so far.

Solution: There is a large selection, here a partial list: perceptron, multi-layer perceptron and other neural networks, McCulloch-Pitts Neuron, integrate and fire neuron, backpropagation learning rule, words/rules model for past tense, statistical word segmentation model, Bayesian inference, category models, receptive fields of neurons, any statistical model, ... ideally this list has some diversity.

2. Choose a model and come up with an experiment to access one or several quantities the model aims to describe. While all these models are related to the brain in some way, the experiments will likely differ substantially. How?

Solution: The discussion will likely show that some models address microscopic biological quantities such as the activity of neurons (can be measured directly), while other models are more abstract and have no direct biological correlates, but allow, for instance, comparison with behaviour. The models covered also cover several schools of thought, such as connectionism (bio-inspired), Bayesian or symbolic/logic.

3. Now discuss how these models can be useful (given they are all wrong anyway). What is the role of a good model? How can a “wrong” model be useful?

Solution: Following McClelland, a good model will generate testable predictions. Failure to confirm such predictions is useful too as an analysis of the model can show where the error was made. We also heard about post-diction: the model predicts data that was not used to construct the model in the first place, so we avoid circular reasoning.

4. Finally, we can ask if more complex models are more useful, since they are of course more expressive. Is it true that “the best model of a cat is a cat”? (this quote has been attributed to Norbert Wiener)

Solution: Here overfitting can be discussed, and many would prefer simpler models as they enable better understanding (note this assumed that a reductionist approach works in biology - it certainly works very well in other disciplines such as physics). However, current machine learning is diverging from this, building image and language models with billions of parameters. Yet these models are rather hard to understand as a result.