Processing communication sounds and learning how to do it

Naftali Tishby and Allison Doupe MCN 2002

Exercises

1. Look at spectrograms of zebra finch songs.

Download the 19 zebra finch songs from the website, http://www.keck.ucsf.edu/∼schenk/MBL/. These are 19 samples of song from one bird (varmus) in .wav format and sampled at 44100 Hz. Also download the song zfaL_11.1.filt.a, you will need this for problem 4.

Part A:

- Load one of these samples into Matlab (using Matlab's *load* or *wavread* functions), plot it and play it. (using Matlab's *soundsc*)[1]
- Use Matlab's *specgram* function to plot the spectrogram of the songs. (Note that the songs have been bandpass filtered with f_{low} = 300Hz and f_{high} = 8000Hz.) Explore a few possible windows (*boxcar, hamming, hanning*, etc.) and several frequency resolutions. For a start, try *specgram(Y,512,44100, hamming(128),100)* and *specgram(Y,512,44100,hamming(400),350)*.
- Make sure you understand the significance of the different parameters.

Part B:

- From the spectrograms of the songs identify and name the individual (sometimes repeating) units, called syllables. As a first step, cut the song into signal (song units) and noise. (You will also need this for problem 2.) As a second step name the signal segments (some will repeat) and draw a diagram of the transitions between them.

2. Clustering using the Information Bottleneck Method

This problem is meant to familiarize you a bit with some of the details of the Information Bottleneck Method (IBN) of clustering. In part (a) you will generate "data" drawn from distributions that are mixtures of differing numbers of Gaussians. You will then compress the identity of these data sets, (each will be a set of samples from some probability distribution). In part (b) you will attempt to use IBN to cluster presegmented zebra finch song into its component syllables. In preparation for this problem download the Matlab routines that implement IBN from http://www.keck.ucsf.edu/∼schenk/MBL/. Put these files in the directory in which you would like to work, or in some directory that is in your Matlab path. (Also download the examples of zebra finch songs that are on this webpage, if you did not do so already for Problem 1. Note that for Part B of this problem you will need to segment one or more of these songs into signal and noise segments.)

---

[1] If you are running linux, *soundsc* will not work. Instead use the command *play -t wav filename* from the command line.

Part A:

- Write a Matlab routine that generates about 20 data sets. Each of these data sets should contain about 1000 samples from distribution that are mixtures of either 2 or 3 gaussians. (You can use the Matlab function RANDN). E.G. generate N data sets drawn from distributions that are mixtures of 2 Gaussians and M data sets that are drawn from distributions that are mixutures of 3 Gaussians. Form the distribution $P(x|y)$ where x is the identity of the data set, and y is the amplitude of the data in each set. This distribution, properly normalized, will be the input for the IBN. (See the example file *ibn-example.m* for help in setting this up.)
- Vary the means and standard deviations of the gaussian mixtures you are drawing your data sets from. How do these manipulations effect the output of the IBN?

Part B:
Now use the IBN to cluster the syllables you cut from a song in Problem 1. To do this using the IBN you must pick a *feature* that you think might be relevant to what you want the clusters to be. As a first pass, use the amplitude distributions of the segments. Form the input matrix for IBN, P(segment indentity|feature) just as you did for Part A above. Does the output of the IBN make sense? Can you think of a feature (or features) that might lead to clustering that more closely follows your by hand labeling of syllables?

3. Song-selective neurons and neural networks that could give rise to them

A sample combination-sensitive HVc neuron and its response to syllables A, B, and pairs of these syllables in particular orders is found on the attached page. Draw a simple ball and stick neural network, using component neurons responsive to A and B, and excitatory and inhibitory mechanisms with typical timescales, that could give rise to such a neuron. Can you imagine other network mechanisms that could do the same thing? How could one test experimentally what mechanisms were really responsible for the song-selective properties of these neurons?

4. Spike trains and song responses

On the website http://phy.ucsf.edu/∼kamal are Matlab readable ascii files of spike data from 10 different sites in the avian forebrain - these are all called "sonenumber.bill". Each set of spike data consists of 10 trials of 8 sec duration; spike arrival times are listed in seconds. For each successive trial, the spike arrival times are incremented by 10 (in order to identify the trial); e.g.l possible spike arrival times for the first trial would be 0.1, 0.5, 1.7, etc, and spike arrival times for the second trial would be 10.2, 10.3, 12.4, for the third 20.4, 20.9, 22.4, etc. All the sensory responses were recorded in response to the song you downloaded in Problem 1, zfaL_11.1.filt.a, played starting at 2 sec into the 8 sec trial.

- Plot the peristimulus time histograms (PSTH) and raster plots for each of the 10 sites. Inspect these and speculate about which feature of the song the neurons might be responding to in each case - are some neurons more complex or difficult to analyze than others?

- Calculate the "spike-triggered average spectrogram" of the 30 ms of the song preceeding each spike. Start with file zfaL_20.4.14.bill, because it has nice high spike rates and good stimulus locking. Hint: To calculate the average spectrogram, take the absolute value of each spectrogram segment (absolute value of *specgram*), square it, (i.e. the power spectrum), and average each value (using *mean*). What difference does it make if you change the size of the stimulus segment that you average?
- Use the IBN algorithm to cluster the (unaveraged) spike-triggered spectrograms for each site and find what are the different stimulus clusters that trigger spikes.