

Lab 1: What was the stimulus?

[Home](#)[Up](#)[Practical Information](#)[Calendar](#)[Atlas of Auditory CNS](#)

HST.723J/9.285J - Neural Coding and Perception of Sound

Spring 2009

What was the stimulus?

Bertrand Delgutte

Lab time and location: 2/13/09, 9-11am in [1-115](#)

Report due 2/20/08

Introduction



The purpose of this laboratory exercise is to illustrate how sound stimuli are coded in the temporal discharge patterns of auditory-nerve fibers. The auditory nerve is a good starting point for understanding the neural coding of sounds because responses are stable, well-characterized, and relatively easy to interpret in terms of cochlear mechanisms. The techniques you will be using in this lab are also applicable to neurons in the central auditory system and other sensory systems. One limitation is that only techniques for analyzing response of *single* neurons are considered. A different set of techniques deals with how information is coded in networks of interconnected neurons.

Your task in this lab is very much like detective work. You are given several sets of spike data recorded from auditory-nerve fibers, and your task is to identify the sound stimulus that produced each set of data. This approach is chosen in part because it is fun, but also because this identification task is similar to what the brain must do when decoding the pattern of auditory-nerve activity produced by a sound stimulus. Taking the “brain’s point of view” encourages you to use physiologically-realistic spike analysis techniques.

There are nevertheless two differences between your task and the one the brain faces. One is that the brain simultaneously processes information from all 30,000 auditory-nerve fibers, whereas you are only provided with data from a single fiber. Nevertheless, you will find that a surprising amount of information about the stimulus can be gathered from the temporal discharge patterns of even a single fiber. Another difference is that, whereas the brain can normally identify an auditory object or understand an utterance from a single presentation of a sound stimulus, the same stimulus is presented many times in single-unit experiments in order to obtain statistically-reliable estimates of the temporal discharge patterns. This difference may not be significant because the brain receives essentially the same information from several auditory nerve fibers, at least the ~20 fibers which innervate a single inner hair cell. Moreover, the responses of fibers innervating adjacent hair cells are also very similar because the cochlear filters are relatively broad. Thus, the temporal averaging (over stimulus presentations) that you perform when analyzing single-unit firing statistics may resemble the spatial averaging over multiple, similar auditory-nerve fibers that the brain undoubtedly performs.

Histograms for analyzing neural responses



Because neural discharges (action potential or "spikes") occur at discrete, punctate instants in time, *histograms* are used to analyze and display the temporal discharge patterns. Histograms estimate the distribution of spike times along a temporal dimension which is divided into small time intervals or "*bins*".

The course software supports more types of histograms than you need to complete this lab. The four types you absolutely need to understand are listed in the table and described further below. Additional histograms that may also be useful are also described below.

Histogram	Description	Useful for	Further Processing
Peri-Stimulus Time (PST)	Distribution of spikes over time course of stimulus	All stimuli (transient, periodic, stationary, time-varying)	Spectrum for transient and periodic stimuli Spectrogram for time-varying stimuli
(First-order) Interspike Interval	Distribution of intervals between successive spikes	Stationary and periodic stimuli, spontaneous activity	Spectrum
Period Histogram	Distribution of spikes over stimulus cycle	Periodic stimuli with known period	Spectrum, Synchronization Index
Reverse Correlation ("Revcor")	Average stimulus waveform preceding each spike	Noise with known waveform	Spectrum

Essential Histograms

- **Peristimulus time (PST) histograms** display the distribution of spike times relative to the onset of each stimulus presentation. They are often the most useful way to show temporal discharge patterns. Useful bin widths range from 0.1 ms or less for brief stimuli such as clicks to 10 ms or more for stimuli lasting several seconds.
- **(First-order) interspike interval histograms** display the distribution of time intervals between consecutive spikes. They are most useful for spontaneous activity, and for periodic stimuli such as pure tones, where they reveal the presence of phase locking.
- **Period (a.k.a. cycle) histograms** show the distribution of spike times over the period of a periodic stimulus. They are useful to quantitatively assess the phase locking of spikes to the stimulus. Specifically the *synchronization index* (a.k.a. vector strength), a measure of the strength of phase locking, is the first Fourier coefficient of the period histogram normalized by the mean discharge rate. Unlike interval and autocorrelation histograms, however, period histogram require the stimulus period to be known, which is not the case in this lab. A good strategy is to first identify the unknown period using an interval or autocorrelation histogram, then verify phase locking to this period using a period histogram.
- **Reverse correlation ("Revcor") functions** represent the average stimulus waveform preceding each spike, reversed in time. Mathematically, the revcor is equivalent to the cross-correlation function between the stimulus waveform and the digitized spike train. Revcor functions are most useful with broadband ("white") noise stimuli, where they provide an estimate of the impulse response of the linear filter which does the best job (in a least-squares sense) of predicting the neural response from the stimulus (Wiener theorem). Revcor functions of low-frequency (< 4 kHz) auditory-nerve fibers resemble the impulse response of a bandpass filter centered at the fiber's characteristic frequency.

Additional Histograms

- **PST Raster (a.k.a. dot raster) histograms** show the distribution of spike times relative to stimulus onset along the horizontal axis, and for each stimulus presentation along the vertical axis. A PST histogram is the column-by-column sum of the PST raster. PST rasters give a nice visual display of the raw spike data very much as they are acquired during an experiment, and are useful for assessing the stability of the

neural response over many repetitions of the same stimulus. However, they are rarely used for quantitative analysis of temporal discharge patterns. Appropriate bin widths for PST rasters are the same as for PST histograms.

- **Autocorrelation (a.k.a. all-order interval) histograms** are similar to first-order interval histograms, except they include intervals between non-consecutive spikes as well as intervals between consecutive spikes. For example, if A, B, and C are the occurrence times of 3 consecutive spikes, then a first-order interval histogram would only include the intervals B–A and C–B, whereas an autocorrelation histogram would include the second-order interval C–A as well. As its name indicates, an autocorrelation histogram is mathematically equivalent to the autocorrelation function of the digitized spike train. Because the autocorrelation function reveals all the periodicities in a signal, autocorrelation histograms are useful for detecting periodicities in a spike train, including phase locking to a stimulus with unknown period, as well as a neuron's intrinsic tendency to fire at regular intervals.

Choosing the bin width

The *bin width* of a histogram must be carefully selected depending on the stimulus, the amount of available data (number of spikes), and the type of histogram. Choosing the wrong bin width may mask patterns that would be clearly revealed with another bin width. Choosing a bin width is often a compromise between temporal resolution and statistical reliability: Too large a bin width smears the temporal pattern, while too low a bin width gives noisy histograms where patterns are hard to discern. Some hints on choosing an appropriate bin width are given below.

- **PST Histogram.** A two step procedure is suggested.
 1. **Determine the stimulus duration.** The field **Elapsed Time** in the Histogram GUI gives the total duration of the recording, while the field **Number of Syncs** gives the number of stimulus presentations. Since the stimuli are presented at regular intervals, the stimulus duration (including any silence between stimulus presentations) is the ratio of the elapsed time to the number of syncs.
 2. **Match the histogram duration to the stimulus duration.** The histogram duration is the product of the bin width by the number of bins. Adjust these two numbers (while keeping their product equal to the stimulus duration) to get the best compromise between time resolution and statistical reliability. In some cases, each stimulus presentation is followed by a long silence, in which case the histogram duration can be made shorter to zoom in on the actual stimulus.
- **Interval Histogram.** A bin width of 0.05 or 0.1 ms, together with a histogram duration of 20-50 ms works well in most cases.
- **Period Histogram.** The bin width is the ratio of the period (which you specify) by the number of bins. Make sure to have at least 20 bins per period, more if there is a large number of spikes in the recording.
- **Revcor Histogram:** The bin width is automatically determined by the sampling rate of the stimulus waveform, which is stored in a data file.

Processing histograms



Unlike raw spike trains, histograms are standard discrete-time signals that have a sampling rate equal to the inverse of the bin width. Therefore, they can be processed with the powerful arsenal of digital signal processing techniques, including linear filters and various transforms. Fourier analysis techniques are particularly useful for revealing features of the data that are not always apparent from time-domain histograms. For example, with some care, the discrete Fourier transform of a PST, revcor, or interval histogram can show the frequency components of the stimulus that a neuron phase locks to. For time-varying stimuli, a spectrogram of the PST histogram shows how the frequency components the fiber phase locks to evolve with time.

Specific Instructions



You are given 6 different Matlab data files named after famous detectives (Clarice, Clouseau, Holmes, Marlowe, Poirot, and Scully). Each file represents spike data recorded from an auditory-nerve in response to one of 6 possible stimuli. Not all files represent data from the same fiber, but all fibers had characteristic frequencies below 3000 Hz so as to exhibit clear phase locking to stimuli having energy near their CF. The 6 stimuli are:

1. None (spontaneous activity).
2. A **continuous**, low-frequency **pure tone**.
3. A brief (100- μ sec) **click** presented at a rate of 10/sec.
4. Continuous, **broadband noise** (specifically, periodic noise with a 5-sec period).
5. A sinusoidally frequency-modulated (**FM**) **pure tone**. The tone's instantaneous frequency varies sinusoidally from a minimum to a maximum and back.
6. A **speech utterance** (about 3 sec).

Your task in this lab is to answer the following questions:

1. **Identify the stimuli that produced the neural responses** stored in each of the 6 data files.
2. For the **spontaneous activity**, give the spontaneous discharge rate (in spikes/sec) and estimate the fiber's absolute refractory period.
3. Determine the frequency of the **pure tone**. Also give the fiber's synchronization index to the tone frequency.
4. Estimate the characteristic frequency (CF) of the fiber for which you have the **click** response. Also estimate the cochlear traveling wave delay.
5. Give the CF of the fiber for which you have the response to **broadband noise**. **Hint:** Compute the *reverse correlation* ("revcor") between the noise signal and the neural response. For this purpose, the noise waveform is available in a Matlab file.
6. For the **FM tone**, give the modulation range, i.e. the minimum and maximum instantaneous frequencies and the modulation frequency. Also give a crude estimate of the fiber CF in relation to the minimum and maximum frequencies, i.e. is the CF above the maximum frequency, below the minimum, or between the minimum and the maximum.
7. Determine whether the **speech utterance** was produced by a male or a female speaker and give the number of vowels in the utterance.
8. **Optional.** The file "warshawski" contains the neural response to yet another stimulus. What is it?
9. **Optional - just for fun.** For each famous detective give either the author of the book or the name of the film or TV series in which he/she first appeared as a character.

Lab Report

- Include answers to all the questions in your report. Explain the clues you use for stimulus identification, and how the different clues fit together. Several lines of arguments leading to the same conclusion are better than one. Include all relevant plots and calculations. No credit will be given for answers lacking a

justification.

- Do not repeat the lab instructions and avoid lengthy introductions.
- Your report should not exceed **4 single-spaced pages**, not including figures which can be either interleaved with the text or attached at the end.
- Each figure should be numbered and have a label.
- Conclude your report with a one-sentence summary of what you learned in the lab.
- Please include any suggestions you may have on how to improve the Histogram GUI for both ease of use and power.
- You are expected to share data and figures with your lab partner, and encouraged to discuss all aspects of the lab with each other. However, **you must write the lab report on your own, using your own words.**

Practical Information

Getting Started

- After you log in onto an Athena workstation, type `attach hst.723` into a shell window, then start Matlab from the Athena menu. Run the Matlab desktop.
- In the Matlab command window, enter `addpath(genpath('/mit/hst.723/matlab'))`. This gives you access to the course software and data files. Alternatively, you can set the path with a GUI available from the File menu in the Matlab desktop.
- Enter `HistogramGUI` into the Matlab command window. This starts the graphical user interface you will be using to compute and plot histograms from stored neural data.



The Histogram GUI

The histogram GUI allows you to specify settings for two different histograms called "A" and "B" that are plotted in a separate figure window. Use the tabs to switch between Histograms A and B. Each item in the histogram panel is described from top to bottom. Certain items are only applicable to specific types of histograms, and are only visible when applicable.

Settings

The following items specify what the histogram computes:

- **Data File.** Used to select the neural data file you are analyzing. The files are located in `/mit/hst.723/matlab/LabANF/Data`. Use the **Browse** button to find the file. Once you select the file, two histograms are automatically plotted in a new figure window based on the panel settings below.
- **Type.** Type of histogram. See the Histograms section above for description of the most useful types.
- **Size.** Number of bins in the histogram. For normal (1-D) histograms enter a single integer. For 2-D histograms (such as PST Rasters), enter two integers separated by a space to specify the numbers of rows and columns in the matrix, in that order.
- **Bin Width(s).** Must be a positive number. See the **Histograms** section above for tips on how to set the bin width. This argument is ignored for period histograms. ***The product of the bin width and the number of bins is the total duration of the histogram.*** For interval histograms, a 20-40 ms duration usually works well. See description of **Elapsed Time** below for tips on how to set the duration of a PST histogram.
- **Offset(s).** Usually set to zero. Non-zero offsets shift the time origin and are useful for zooming in on a temporal segment of a long histogram with high resolution.
- **Sync and Spike Channels** should always be 0 and 1, respectively.
- **Period.** Used for period histograms only. Sets the stimulus period in ms. The bin width is automatically set to the period divided by the number of bins.

- **Interval Order.** Used with interval histograms only. Should always be 1.
- **Stimulus.** Used with Revcor functions only. Specifies the stimulus waveform that is averaged before each spike. Use the **Browse** menu to find the stimulus file located in the same folder as the neural data.

The following items are read only. They provide information about the data contained in the analyzed file:

- **Number of Syncs.** Gives the number of stimulus presentations.
- **Number of Spikes.** Gives the number of spikes in the recording.
- **Elapsed Time.** Total duration of the recording in sec. The number of spikes divided by the elapsed time gives the average discharge rate in spikes/sec. *The elapsed time divided by the number of syncs (stimulus presentation) gives the average duration of each stimulus presentation*, which is useful to set the number of bins in a PST histogram.

The following items affect only how the histogram is displayed:

- **Plot Style.** For 1-D histograms, specifies if the histogram is plotted as vertical bars (usually the best) or as continuous lines (good for Revcors). For 2-D histograms, specifies whether the histogram is plotted as an image (usually the best), a set of dots (traditional for PST rasters), or "waterfall" displays (not useful here).
- **Plot threshold.** Only used for plotting 2-D histograms as dots. Specifies the number of spikes per bin that must be exceeded for a bin to be plotted as a dot. A zero threshold usually works well.
- **Plot Scale.** Applies to 1-D histograms only. Specifies if the ordinate represents the number of spikes per bin ("count") or to discharge rate in spikes/sec ("rate").

Buttons

- **Apply.** Not useful in this lab.
- **Compute.** Compute and display a histogram in a figure window based on current settings. If the settings are incomplete or inappropriate, an error message will appear in the Matlab command window.
- **Gate.** Displays the "Histogram Gate" pop-up menu. Gates are used to select a subset of the spikes for including into a histogram computation. Not necessary in this lab.
- **Spectrum.** Compute and display the Fourier magnitude spectrum of the current histogram.
- **Spectrogram.** Compute and display the spectrogram (short-time Fourier transform) of the current histogram. The spectrogram is displayed as an image, with time on the horizontal axis and frequency on the vertical axis. A 20-ms analysis window is used to compute the spectrogram.
- **Synchrony.** Used with period histograms only. Computes the synchronization index to the stimulus period, and displays the result in the adjacent box.



Hints

While the Histogram GUI is easy to use, the numerous options, combined with the unknown nature of the stimulus, may be somewhat baffling at first. Here are some hints on how to proceed.

- Start with a **PST histogram**. See the section "**Choosing the bin width**" to select an appropriate bin width and number of bins.
- For **interval histograms**, start with 200 bins and a 0.1 ms bin width. Any phase locking should be apparent as regularly-spaced modes in the histogram. Decrease the bin width (and increase the number of bins) if the histogram is not too noisy.
- Do not even try a **period histogram** until you are confident that the stimulus is periodic and you have identified the likely period.
- **Interval and autocorrelation histograms** are only useful for "stationary" stimuli, whose spectrum does not change with time. For non-stationary stimuli such as speech, use a PST histogram computed with fine time resolution.
- **Revcor functions** are only useful for the noise stimulus. Start with 200 bins. The bin width is automatically set to the inverse of the noise sampling rate (20 kHz).

For Matlab experts

Histograms are discrete-time signals that can be analyzed using standard digital signal processing (DSP) techniques. If you are familiar with both Matlab and DSP, you can access the data produced by the Histogram GUI and apply further processing beyond that available in the GUI.



- Type `h=HistogramGUI('GetHist');` into the Matlab command window to get the current histogram. The Matlab variable `h` is a special histogram object containing all the histogram settings and data.
- To access the data in the histogram, type `data=double(h)`; This returns a vector or a matrix for 1-D and 2-D histograms, respectively, and containing the number of spikes in each bin.
- To get the bin width, type `bw=get(h, 'BinWidth')` for 1-D histograms, or `bw=get(h, 'BinWidths')` for 2-D histograms.
- You can also access many other histogram properties. Type `help @histogram/get` for further information.

Data files

- The files containing the neural data are, in alphabetical order: **clarice.mat**, **clouseau.mat**, **holmes.mat**, **marlowe.mat**, **poiroi.mat**, and **scully.mat**. Each file contains the response of one fiber to one stimulus as described below.
- Neural responses to each stimulus are stored as a pair of vectors called **t** and **ch** (Figure 1). The vector **t** represents the times of recorded events (either neural spikes or stimulus onsets) expressed in msec since the beginning of recording. The second vector **ch** (for channel) identifies each event as either a stimulus onset (0) or a neural spike (1).
- The file **bbnoise_signal.mat** contains the waveform of the continuous broadband noise stimulus and the noise sampling rate **Fs** in Hz. These are essential for computing the revcor function.
- You may ftp the data files and lab software to your own computer. They are in [/afs/athena.mit.edu/course/other/hst.723/matlab/LabANF](http://afs/athena.mit.edu/course/other/hst.723/matlab/LabANF).

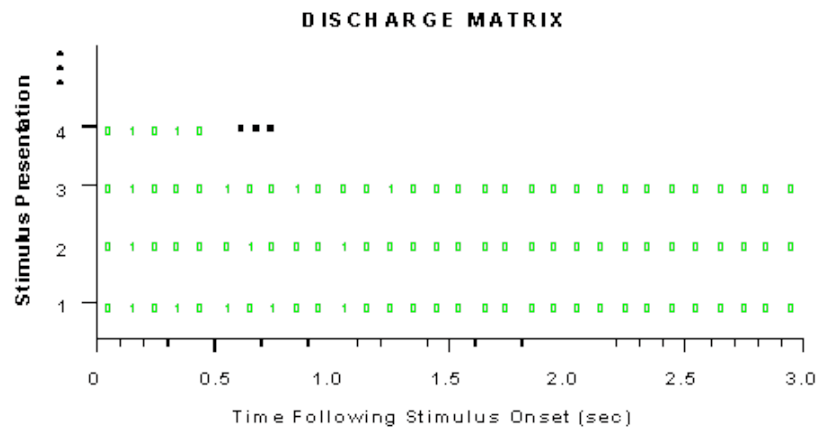



Figure 1. Construction of a PST raster from stored spike times.