PSYCH-GA.2211/NEURL-GA.2201 – Fall 2014
Mathematical Tools for Cognitive and Neural Science

Homework 5

Due: 19 Dec 2014
(late homeworks penalized 10% per day)

Save the solutions to each numbered problem as sections of a file called `hw5.m` in a folder called `hw5_Lastname`, along with additional files containing any functions you create. Send a zipped copy of this folder as an attachment in an email message with these attributes:

`To: catherio@nyu.edu, asr443@nyu.edu`
`Subject: Math Tools HW5`

Don't wait until the day before the due date... start *now*!

1. **Simulating a 2AFC experiment.** Consider a two-alternative forced choice (2AFC) psychophysical experiment in which a subject sees two stimulus arrays of some intensity on a trial and must say which one contains the target. (One and only one contains the target.) Her probability of being correct on a trial is:

$$p_c(I) = 1/2 + 1/2\Phi(I; \mu, \sigma)$$

where $\Phi(I; \mu, \sigma)$ is the cumulative distribution function of the Gaussian (`normcdf` in matlab) with mean μ and standard deviation σ evaluated at I. The function $p_c(I)$ is known as the *psychometric function*.

   (a) Plot two psychometric functions, for $\{\mu, \sigma\}$ equal to $\{5, 2\}$ and $\{4, 3\}$. (Use $I = [1 : 10]$). Describe the difference between these. If you increase μ, how does the curve change? If you increase σ, how does the curve change? (If you are not sure, make more plots with different parameter values.) What is the range of $p_c(I)$? Explain why this range is appropriate.

   (b) Write a function `C=simpsych(mu,sigma,I,T)` which takes two vectors `(I,T)` of the same length, containing a list of intensities and the number of trials for each intensity, respectively, simulates draws from $p_c(I)$, and returns a vector, C, of the same length as I and T, which contains the number of trials correct out of T, at each intensity I.

   (c) Illustrate the use of `simpsych` with `T=ones(1,7)*100` and `I=1:7` for $\mu = 4$ and $\sigma = 1$. Plot `C ./ T` vs `I` (as points) and plot the psychometric function $p_c(I)$ (as a curve) on the same graph.

   (d) Do the same with `T=ones(1,7)*10` and plot the results (including the psychometric function). What is the difference between this and the plot of the previous question?

2. **Fitting the psychometric function.**

   (a) Write a function `nll = nloglik(mu,sigma,I,T,C)` that returns the negative of the log likelihood of parameters `mu` and `sigma`, for data set `I,T,C`.

   (b) Generate a contour plot (function contour, using 50 lines) of the negative log likelihood of the data set from 1C, for all pairs of `mu` from `muall = [2:0.2:10]` and a `sigma` from `sigmaall = [0.5:0.2:6]`. What is the approximate location of the best fitting pair of parameters from this plot (determined visually)?

(c) Use the function `fminsearch` to find the more precise values `mu, sigma` that minimize the function `nloglik(mu,sigma,)`. Three notes: first, the syntax for calling `nloglik` within `fminsearch` is a bit odd:
`fminsearch(@(x) nloglik(x(1),x(2),I,T,C),[startpoint]).`
Second, `fminsearch` minimizes rather than maximizes, which is why we are using the negative of the log likelihood, so smaller numbers indicate better fit. Third: You need to specify a start point for the search. `[2,2]` is a good choice.

(d) A variant of `fminsearch, fminunc`, also returns the Hessian (matrix of second derivatives) of the negative log likelihood at the optimal `mu` and `sigma`. Using the inverse of the Hessian (which is an estimator for the covariance matrix of the parameter estimates) determine 95% confidence intervals on both parameters. (Hint: The marginal standard errors for each parameter are the square roots of the diagonal entries of the inverse Hessian; a 95% confidence interval is the mean plus or minus 1.96 standard errors.) Do the true parameter values (4 and 1) fall within the confidence intervals? Note: fminunc is a little touchier numerically than fminsearch, and if the optimizer strays too far from the true values, there may be numerical problems due to overflow of the likelihood; in this case, try a different starting point.

(e) Produce a second set of confidence intervals for the parameters using a bootstrap method. For each of the 7 intensities, resample 100 trials (correct or incorrect) from the 100 trials of that intensity in the original data, with replacement. Refit the model to the resampled data using `fminsearch`. Plot the histograms (function hist) of `mu` and `sigma` estimates obtained over 500 such resampled datasets, and define your confidence intervals as the region between the 2.5th and 97.5th percentiles of these distributions. How well do these values agree with those from 2D?

3. **Comparing two psychometric functions.** Suppose we repeat the psychophysical experiment before and after giving the subject an experimental drug. Do the parameters change?

   (a) Simulate the experiment from 1C twice, once using the original parameters and again using the parameters `mu=5, sigma=1`. Fit each dataset using `fminsearch` to recover estimated parameters, and make note of the difference between the two estimates of `mu` and `sigma`.

   (b) Now construct a permutation test of the null hypothesis (i.e., the hypothesis that there has been no change in the parameters). For each intensity, combine the 100 trials from each condition into a total of 200, then randomly partitioning this into two groups of 100. Fit both resampled datasets again, noting the difference between the two `mu`s and the two `sigma`s. Repeat this process 500 times to produce a null distribution of the differences in each parameter. How likely (at what quantile; one-tailed p-value) is the actual difference in `mu` from 3A? What about for `sigma`? Do these results make sense given the true parameter values from which you simulated the datasets?

4. **Comparing two models.** Consider the restricted model given by our previous psychometric model, with `sigma` fixed at 2. This model has only one free parameter (`mu`). Is it a better fit to data?

   (a) Write a new `nloglik` function, `nloglikb`, that returns the likelihood of some dataset for the new model, as a function of `mu` and a dataset. (Note: This function can simply call the previous one, with `sigma=2`.) Using a new invocation of `fminsearch`, fit the model to the data from 1C.

(b) Using the log likelihood of the data at the best fitting parameters, compute the AIC score for the two models, and decide which one is a better fit. Repeat using BIC. Are the results consistent with ground truth?

(c) Simulate a dataset for `mu=5, sigma=2`. Repeat question 4B. Are the results consistent with ground truth?