MIT OpenCourseWare
http://ocw.mit.edu

HST.583 Functional Magnetic Resonance Imaging: Data Acquisition and Analysis
Fall 2008

For information about citing these materials or our Terms of Use, visit: http://ocw.mit.edu/terms.

```
Statistical Signal Processing of fMRI
Douglas N. Greve

----------------------------------------------------------------
Experiment with 3 identical stimulus presentations (Nrep=3).
Stimulus schedule - far enough apart for no overlap (FIER)
Results in 3 identical HRFs.
Noise added.
Sampled at TR.
Ntp = number of time points/TRs over experiment.
This is the "Observable" (y).
From this we want to compute the hemodynamic response amplitude
and quantify the uncertainty.

----------------------------------------------------------------
Non-Matrix Formulation:
Construct Post-stimulus delay (PSD) Window (aka Peristimulus)
One bin for each TR in the window
Number of bins = Window/TR (Nbeta)
Each bin corresponds to 3 measurements
For each bin compute: selective sum, count/Nrep, average
Plot the averages (betahat, betas, parameter estimates)

----------------------------------------------------------------
Compute signal estimate (yhat) (resynth)
Compute residual r = y-yhat
SumSquareError = SSE = sum(r^2)
DOF = Ntp - Nbeta
Residual variance = rvar = SSE/DOF = MeanSquareError (MSE)
Residual stddev = sqrt(rvar) = RootMeanSquareError (RMSE)
LMSE = LeastMeanSquareError Estimator
StandardVariance = rvar/Nrep
StandardError = sqrt(rvar/Nrep) = sqrt(StandardVariance)
t = Average/StandardError
Adequate analysis, but not very flexible (eg, how to add more
  conditions, overlay, nuisance variables, etc, ?).

----------------------------------------------------------------
Matrix Formulation: y = X*beta
y = observable (Ntp-by-1)
X = design matrix (Ntp-by-Nbeta)
beta = true parameters/weights (cf betahat vs beta) (Nbeta-by-1)
row = time point
First col of X = 1 where stim are presented, 0 else
Second col = First shifted down one row.
Third col = ...
One col for each delay = Nbeta
cols = post stimulus delay
col vector = regressor
Set of linear equations: Ntp knowns, Nbeta unknowns
Must have: Ntp > Nbeta (overdetermined)
Solve = Fit = Estimate: betahat = inv(X'*X)*X'*y (cf betahat vs beta)
betahat = regression coefficient, parameter estimate, average
Pseudoinverse: inv(X'*X)*X' = X+ (for non-square matrices)
Selective Sum = X'*y
Count = X'*X (strictly only true for FIER)
```

```
1/Count = inv(X'*X)
Average = betahat = SelectiveSum/Count =
    (1/Count)*SelectiveSum = inv(X'*X)*(X'*y)
Signal estimate: yhat = X*betahat
Residual: r = y - yhat = (I - X*inv(X'*X)*X')*y = R*y
R = resisual forming matrix. Sym: R=R'. Idempotent: R*R=R.
DOF = #colsX-#rowsX = Ntp-Nbeta = trace(R)
SSE = r'*r;
Residual Var = rvar = SSE/DOF
StandardVar = BetaVar = rvar * inv(X'*X)
LMSE
Stimulus schedule embedded in X.

----------------------------------------------------------------
Example: TR = 2, Ntp = 10, PSDWin=6sec, Nrep=2 (2s, 12s)
How big are y, X, and beta?
What's the DOF?
How many equations?
How many unknowns?

----------------------------------------------------------------
Statistical Model Summary:
y = X*beta + n (Forward Model)
beta = true beta (unknown)
n = noise, unknown, ~N(0,nvar*Sn),
  nvar=true variance (unknown)
  Sn = true temporal covariance matrix (unknown) (=I for white noise)
HDR Model (FIR):
  = 0 if PSD < 0 or PSD > PSDMax
  anything sampled at TR within the window
Parameters: beta, nvar, Sn, PSDMax
Properties:
  betahat = inv(X'*X)*X'*y
  E(betahat) = beta (unbiased)
  E(rvar) = nvar (unbiased)
  Cov(betahat) = rvar * inv(X'*X) (error bars)
Generally true regardless of X, y, beta, etc.

----------------------------------------------------------------
Overlap in the responses and blocked design.

----------------------------------------------------------------
Assuming a shape to the HDR:
Gamma function: h(psd) = beta * ((psd-D)/tau)^2 * exp((psd-D)/tau)
Sampled within the window (zero outside and less than D)
Three parameters:
  beta = amplitude (linear, unknown)
  D = delay (nonlinear, assumed, eg 2.25 sec)
  tau = dispersion (nonlinear, assumed, eg, 1.25 sec)
Draw plots showing different taus and betas
Assume values for nonlinear parameters, but estimate/fit linear parameters
Draw observable
Draw signal estimate for different betas
Draw residual for different betas (SSE)
Draw SSE as a function of beta

New model: y = Xfir * A * beta
```

```
A = assumed shape with beta=1 as a col vector
Xfir = FIR matrix from before
beta = true amplitude, (Nbeta=1 only one value!)
Can combine X = Xfir*A
y = X*beta
y = observable (Ntp-by-1)
X = design matrix (Ntp-by-1)
beta = (1-by-1)
Ntp equations, 1 unknown
DOF = Ntp - 1 (more than with FIR)
betahat = inv(X'*X)*X'*y = beta at minimum of SSE curve above

Add derivative.

----------------------------------------------------------------
Multiple Conditions:
y = X1*beta1 + X2*beta2
  = [X1 X2] * [beta1; beta2]
  = X*beta
Same number of time points, twice as many unknowns.
betahat = inv(X'*X)*X'*y;
Have to know how design matrix was constructed to interpret betas!

----------------------------------------------------------------
Nuisance Regressors:
y = Xt*betat + Xn*betan
Xt = Task design matrix
Xn = Nuisance regressors
eg, Xn = column of ones for a baseline offset
y = [X Xn]*[beta; betan] = X*beta
betahat = inv(X'*X)*X'*y;
Have to know how design matrix was constructed to interpret betas!

----------------------------------------------------------------
Contrasts
Embodiment of hypothesis (null hypothesis).
Eg, Happy vs Sad faces = HappyHRFAmp - SadHRFAmp
Weighted sum of betas, weights are C = [+1 -1]
gamma = C*beta
Need to know the meaning of each regressor in order to assign weight
Eg, Happy vs Baseline: C = [1 0]
Eg, Sad vs Baseline:  C = [0 1]
Eg, Happy+Sad vs Baseline:  C = [.5 .5]
Does not have to be 0 or +/-1 (overall scale may not be important)
Eg, Aud, Vis, Aud+Vis C = [.5 .5 -1]
ANOVA 2x2: HappyMale, HappyFemale, SadMale, SadFemale
  Interaction between Emotion and Gender:
    (HM-SM) - (HF-SF) = HM - HF - SM +SF : C = [+1 -1 -1 +1]

Multi-variate contrasts (more than one row, OR)
Happy OR Sad vs Baseline:  C = [1 0; 0 1] (multivariate)
ANOVA 2x2, Simple Main Effect of Emotion:
  (HM-SM) OR (HF-SF) : Ca = [+1 0 -1 0], Cb = [0 +1 0 -1]
  C = [+1  0 -1  0]
      [ 0 +1  0 -1]
ANOVA 2x3: HappyMale, HappyFemale, NuetralMale,
           SadMale,   SadFemale,   NetralFemale
```

```
  Interaction bet Emotion and Gender:
    (HM-SM) - (HF-SF) = HM - HF - SM +SF : Ca = [+1 -1  0 -1 +1  0]
    (HM-NM) - (HF-NF) = HM - HF - NM +NF : Cb = [+1  0 -1 -1  0 +1]
      C = [+1 -1  0 -1 +1  0]
          [+1  0 -1 -1  0 +1]
Nuisance - set weights to 0.

Four conditions: C1 C2 C3 C4. Estimate each one separately and test
(C1+C2) - (C3+C4) vs combining C1/C2 and C3/C4 into two regressors (A
and B) and testing A-B.

----------------------------------------------------------------
Hypothesis testing and the quantification of uncertainty and risk
Null hypothesis (H0): nothing is happening
The H0 can be
  1. True (there is truly NOTHING happening)
  2. False (there is truly SOMETHING happening)
We can make the following decisions base on our analysis:
  1. Reject H0 (postive) (there's not nothing happening)
  2. Fail to Reject (negative) (not necessarily accepting H0)
This leads to four possible results (2x2):
         H0True  H0False
  Reject    FP     TP
  Fail      TN     FN
Two types of errors:
  Type  I: False Positive (FP) (alpha, FPR, selectivity; TPR=1-FPR)
  Type II: False Negative (FN) (beta,  FNR, sensitivity; TNR=1-FNR)
Power analysis/ROC is (1-beta) vs alpha = TPR vs FPR

Rate = frequency that we would make a certain decision if we did the
same experiment over and over again without changes to the signal,
design, or noise ("frequentists").

Hard to quantity FNR - need actual TPR.
Can quantify FPR (assuming the model is correct).

Protect yourself against false positives by choosing a threshold

----------------------------------------------------------------
Parametric Statistics and the Linear Model
gamma = C*beta
gammahat = C*betahat

Under H0: beta = 0, y = X*beta+n = n, so
  betahat = inv(X'*X)*X'*n
  gammahat = C*betahat = C*inv(X'*X)*X'*n
Could determine distribution emperically from rest scans

E(gammahat) = gamma
cov(gammahat) = rvar * C*inv(X'*X)*C'  (StandardVar/Err)
gammahat ~N(gamma, cov(gammahat))
t = gammahat/sqrt(cov(gammahat)))
F = gammahat'*inv(cov)*gammahat/J, J = rows in C
  = (C*beta)' * inv(rvar*C*inv(X'*X)*C') * (C*beta)/J


----------------------------------------------------------------
```

Noise modeling:

Noise Sources:

1. Thermal (white) = 50% in CtxGM, 90+% in WM. (3T, unsmoothed)
2. Physio (correlated)
   - Motion - partial voluming, spin history
   - Respiration
   - Heart beat
   - "Resting State Networks"
3. Scanner-related (white)

When Sn != I

Autocorrelation - can predict noise at one tp given previous better than chance (not necessarily deterministic).

Tend to be low freq: respiration, heart beat (aliasing), motion, other physio.

Why should you care:

1. Invalidates t and F (tends to make too liberal)
2. Less efficient than white.

Autocorrelation plot of residual

Time invariant - model as an ACF, generate Sn from ACF

Estimate of ACF from residuals (biased)

W = inv(chol(Sn))

New model (generalized least squares): W*y = W*X*beta

"Fully efficient"

Without pre-whitening: cov(gamhat) = C*inv(X'*X)*X'*Sn*X*inv(X'*X)*C'

With pre-whitening: cov(gamhat) = C*inv(X2'*X2)*C'

Other topics:

Parametric time-varying or weighted (eg, reaction time)

Efficiency and Optimal design

Deconvolution, overlap, simultaneous equations.

In the equation y = X*beta, identify y, X, and beta.

What is the difference between beta and betahat?