Computational statistics, HW 1: Smooth optimization

Liam Paninski

February 5, 2013

Generate some simulated data according to the Poisson regression model

$$n_t \sim Poiss(\lambda_t) \quad (1)$$
$$\lambda_t = \exp(X_t \theta), \quad (2)$$

where the basis functions $\{x_i(t)\}_{1 \leq i \leq p}$ are prespecified. (Choose whatever basis functions you'd like. The parameter θ should be chosen so that λ_t is not too large and not too small; keep $\log \lambda_t$ between -2 and 2 or so.)

In addition, introduce a Gaussian prior distribution on θ,

$$\theta \sim \mathcal{N}(0, C), \quad (3)$$

for some covariance C.

Write two codes to compute the MLE and MAP estimates of θ given the observed data. The first code should use the Newton-Raphson method (with backstepping linesearches, if necessary), and the second should use the conjugate gradient method.

Plot the estimated λ_t for both the MAP and the MLE, along with the true λ_t.

Can you think of a good way to compute approximate confidence (or credible) intervals around your estimates? If so, plot these too.

How does the quality of the solution depend on p, C and the specific basis $\{x_i(t)\}_{1 \leq i \leq p}$?

Plot the cpu time as a function of p; for what bases $\{x_i(t)\}$ and prior covariance C do you observe linear scaling of cpu time as a function of p (while still obtaining good estimates of λ_t)? Are there bases/prior covariances for which Newton can be made faster than CG, and vice versa?