The final project for this class is to analyze a longitudinal data set using the methods we have learned, and to write a report about your analysis that describes the steps you took in your analysis. This report should include enough computer output and plots to be clear; not all the output you generate need be included. PDF format is strongly preferred. Please also submit a plain text file with the R commands needed to recreate your analyses.

The report is due by Saturday, May 14.

The guidelines for the analysis will be left intentionally vague, as it will depend somewhat on the data set chosen, and also as deciding what to do is an important part of any analysis. However, all analyses should include these five general parts. Questions are welcomed.

1. Explore the data graphically. Make at least one profile plot and discuss the patterns you see. What kind of models do you think you might try? Is a transformation of the response necessary?
2. Set the fixed effects to be something reasonable and explore various options for the covariance structure, including possibly letting the variance differ by a covariate.
3. Using that covariance structure, explore the fixed effects. Be sure to think about the proper way to handle baseline data, if any, for the data you chose. Confirm that the covariance structure is still reasonable with the fixed effects you ended up with.
4. Compute appropriate residuals and do some model diagnostics. Does your model seem to fit okay? If not, what are your concerns?
5. Using your best model, dig deeper into the analysis to answer the scientific question posed.

Here are three potential data sets. You may also choose one of your own, but you need to have it approved by me first.

- Thall and Vail (1990) give a data set on two-week seizure counts for 59 epileptics. The number of seizures was recorded for a baseline period of 8 weeks, and then patients were randomly assigned to a treatment group or a control group. Counts were then recorded for four successive two-week periods.

  Is there a reduction in seizures due to the treatment? If so, how much? Does this reduction depend on age?

  This data frame has 295 rows and the following 5 columns: 'subject': subject number, 1 to 59. 'age': subject's age, in years. 'trt': treatment, 'placebo' or 'progabide'. 'period': period, 0 to 4. 'count': the count for the period; except for the baseline period where it was divided by four.

  ```
  > read.delim("http://www.stat.umn.edu/~arendahl/Teaching/EPSY8282/final/epilepsy.txt")
  ```

- Five rabbits were studied on two occasions, after treatment with saline (control) and after treatment with the 5-HT_3 antagonist MDL 72222. After each treatment ascending doses of phenylbiguanide were injected intravenously at 10 minute intervals and the responses of mean blood pressure measured.

  Does the cardiogenic chemoreflex elicited by phenylbiguanide depend on the activation of 5-HT_3 receptors? If so, how?

  This data frame contains 60 rows and the following variables: 'BPchange': change in blood pressure relative to the start of the experiment. 'Dose': dose of Phenylbiguanide in micrograms. 'Treatment': placebo or the 5-HT_3 antagonist MDL 72222. 'Animal': label of animal used (1 to 5).

  ```
  > read.delim("http://www.stat.umn.edu/~arendahl/Teaching/EPSY8282/final/rabbitbp.txt")
  ```

- Fifty chicks were measured every approximately every two days from birth (day 0) to three weeks (day 21). These chicks had been randomly assigned to each of four diets.

  Is there an effect of diet on early growth of chicks? If so, what?

  This data frame contains 578 rows and 4 variables: 'weight': the body weight of the chick (gm). 'Time': a numeric vector giving the number of days since birth when the measurement was made. 'Chick': label of the animal (1 to 50). 'Diet': label of the diet the animal was assigned to (1 to 4).

  ```
  > read.delim("http://www.stat.umn.edu/~arendahl/Teaching/EPSY8282/final/chickweight.txt")
  ```