Third miniproject (Bayesian statistics)

Two options:

1. Large sample results in Bayesian statistics
2. Priors for variance parameters in hiearchical models

Large sample methods

For maximum likelihood inference based on observations $Y_1, \ldots, Y_n \sim f(\cdot; \theta)$ we have well-known large sample (large n) results:

$$\sqrt{n}(\hat{\theta} - \theta^*) \to N(0, i(\theta^*)^{-1})$$

for $n \to \infty$ where $i(\theta)$ is the Fisher information based on one observation and θ^* is the true parameter value.

In Bayesian statistics one can obtain similar results:

$$\sqrt{n}(\theta - \hat{\theta})|y_1, \ldots, y_n \to N(0, i(\theta^*)^{-1})$$

i.e. the posterior distribution of θ can be approximated by a normal distribution around the MLE. This also implies ‘consistency’ - i.e. the posterior becomes increasingly concentrated around θ^*.

The idea of this option for the third miniproject is to explore rigorous asymptotic results in Bayesian inference based on page 1-5 in the manuscript by Ghosh, Delampady and Samanta.

The Laplace approximation will appear as an important ingredient in the derivations.

Priors for variances in hierarchical models

Consider the simple mixed model (hierarchical model) (one-way ANOVA)

$$Y_{ij} \sim N(\mu + \alpha_j, \sigma^2)$$
$$\alpha_j \sim N(0, \tau^2)$$

For a Bayesian analysis we need to impose priors for μ, σ^2 and τ^2. It is often difficult to elicit informative priors but using a flat/improper prior for τ^2 can have undesirable consequences.

The second option for the third miniproject is to discuss and try out various choices of prior distributions for variance parameters in hierarchical models based on the paper 'Prior distributions for variance parameters in hierarchical models' by Andrew Gelman.