Maximum likelihood estimation and resampling techniques

Maximum likelihood estimation in linear regression

1. Write up the simple linear regression model with one explanatory variable.
2. Write the log likelihood for `n` observations.
3. Show that the maximum likelihood estimator (MLE) is indeed the estimator minimising the mean squared error (MSE).
4. Demonstrate the two estimators numerically on the `trees` dataset with `Volume` as response variable and `Girth` as explanatory variable. Use `optim()` for both MLE and MSE, and compare to the output of `summary(lm(...))`.

Overfitting and cross validation

Use the `trees` dataset (`Volume` explained by `Girth`) in these exercises.

1. Use both non-parametric and parametric bootstrap to estimate standard errors of the parameter estimates in the simple linear regression model. Compare to those obtained from `summary()`.
2. Use cross validation to investigate polynomial and spline regression for the `trees` dataset (`Volume` explained by `Girth`).