

Practical 2

Jumping Rivers

For this practical we will explore models for the prediction of progression of diabetes for 442 patients. Measurements of their age, gender, body mass index, blood pressure and size blood serum measurements were taken to gether with a numeric measurement of disease progression one year after a baseline.

The data are available in the *jrpyml* package and can be accessed with

```
import jrpyml
diabetes = jrpyml.datasets.diabetes.load_data()
```

The data have already been normalised, so we do not need to worry about this. However we should separate the inputs from the output ready for modelling.

```
X, y = diabetes.drop('y', axis=1), diabetes['y']
```

- It is good practice to have a dedicated test set for final assessment of our chosen models. We can create training and test sets from data using `sklearn.model_selection.train_test_split()`. The following code will partition our data with 10% held out for final testing. The other 90% we will use for training and cross validation of different models.

```
from sklearn.model_selection import train_test_split
```

```
X_train, X_test, y_train, y_test = train_test_split(
    X, y, test_size=0.1,
    random_state=2019, # ensures same random subset
)
```

- Begin by fitting a linear regression to the training set using all available predictor variables.
- Use the `mean_squared_error` function from the *sklearn.metrics* module on the full training set. This will give us the training error.
- Training error gives us a measure of how far from the original data our model is. However it is typically different to test error, which would give us a better idea of how our model generalises to new data. Use 10 fold cross validation to estimate the test error rate for this model.
- How does this compare to the training error

- Fit a lasso regression model to the diabetes data.
- How do the coefficients of the lasso model compare to those of the standard linear regression?
- Try fitting ridge and elastic net models too
- Which model performs best on the test set in terms of mean squared error?