

B(E)3M33UI — Exercise ML02: Linear regression

Petr Pošík, Jiří Spilka

March 5, 2019

The learning goals:

- formulate machine learning task as an optimization problem
- get familiar with the general process of model fitting
- understand the linear regression model
- introduce the numpy package and its features

The code is organized into the following modules:

- `ML02.py` - the main script for the exercise
- `linreg.py` - helper functions for linear regression

After completion, zip these two files and hand in the archive via the BRUTE. **If you do not manage to complete the exercise in the lab, finish it as a homework!**

1 Problem description

In this exercise, we use the `data_ai_course_2015_2018.csv` dataset that contains information about AI course in the past years. It contains attributes such as number of points from first and second semester task and number of points from exam.

For the simple regression exercise, we study the relation of `#hours_exam(x)` and `#exam_points(y)`, i.e. we build the model $\hat{y} = h(x)$.

Task 1: Run the `ML02.py` script. It should end up in an error but it should show scatter plot of data (you can also stop execution using `plt.show()`). Make three hypothesis for linear regression weights – estimate the parameters w_0 and w_1 by hand and draw them into a weight space (x-axis w_1 , y-axis w_0). Further, fill them into `ML02.py`. Before we can use these weights we need to implement several functions.

2 Computing predictions of a linear model given w_0, w_1

From the lecture, you should know that the predictions

$$\hat{\mathbf{y}} = \mathbf{X}\mathbf{w}^T,$$

where \mathbf{X} is the $[|T| \times (D + 1)]$ matrix of inputs in homogeneous coordinates, $\mathbf{w} = (w_0, w_1)$ is the $(D + 1)$ -vector of parameters of the linear model, and $\hat{\mathbf{y}}$ is the $|T|$ -vector of predictions.

Task 2: In `linreg.py`, create function `homogenize(X)` which

- takes a $[|T| \times D]$ matrix of inputs, and

- produces a $[|T| \times (D + 1)]$ matrix of inputs in homogeneous coordinates, i.e. it prepends a column vector of 1s to X .

Hints: Have a look at:

1. The `numpy.ndarray.shape` property. How do you learn the number of rows and columns?
2. The `numpy.ones()` function. How do you create a $[|T| \times 1]$ array of ones?
3. The `numpy.hstack()` function. How do you add the column of ones to the data X ?

Task 3: In `linreg.py`, create function `pred_regr_lin(w,X)` which takes

- a $(D + 1)$ -vector of linear model parameters w and
- a $[|T| \times D]$ matrix of inputs X , and
- produces a $|T|$ -vector of estimates \hat{y} .

Hints:

1. Do not forget to homogenize the inputs.
2. Have a look at the `numpy.ndarray.T` property. How do you get the transposed vector w^T ? Does the transpose have any effect for vectors?
3. Have a look at the `numpy.ndarray.dot()` method or use matrix multiplication symbol `@`. How do you compute the matrix product Xw^T ?

Now, you should see the predictions of the linear model in the figure. What is the error that your hand-crafted model makes on the data?

3 Computing the error of the linear model

The `linreg.py` module contains function `compute_cost_regr_lin(w,X,y)`, i.e. it implements the $J(w, T)$ function which describes how well the linear model with parameters w fits the data T . It first computes the predictions of the model, and then calls function `compute_err_mse(y,yhat)` which should return the mean squared error (MSE). This function, however, is not implemented yet.

Task 4: In `linreg.py`, create function `compute_err_mse(y,yhat)` which

- takes
 1. a $|T|$ -vector of true values of y and
 2. a $|T|$ -vector of the corresponding predictions \hat{y} , and
- computes the mean squared error (MSE), i.e. the score used by the *least squares method*.

Hints:

- you can compute the sum of squares either
 1. by squaring the differences (`diffs**2`) and then summing them (`numpy.sum()`), or
 2. by using the dot product of the vector of differences with itself (`diffs.dot(diffs)`).
- Do not forget to divide the sum of squares by the number of training examples.

Task 5: Now, when you have an objective measure of the model fit, try a few times and find by hand better values of w_0 and w_1 , i.e. values which result in a lower error.

4 Model fitting by minimization of $J(w, T)$

You have just sought for the best values of w_0 and w_1 by hand. We can make the computer to do this operation for us. We shall use the `scipy.optimize` package.

Task 6: In `linreg.py`, create function `fit_regr_lin_by_minimization(X, y)` which

- takes the training data T as input (X, y) and
- returns the (hopefully) optimal values of w as a Numpy array.

Hints:

1. Take a look at the `scipy.optimize.minimize()` function. For us, it will be sufficient to use it in the form

```
result = minimize(f, w_init, method='bfgs')
```

where `f` is the function (of a single argument) to be minimized, `w_init` is the initial guess of the weights, and `method='bfgs'` chooses the BFGS optimization method.

2. The resulting values of w can be found in `result.x`.
3. Function `compute_cost_regr_lin(w, X, y)` is the implementation of $J(w, T)$, which we would like to minimize with respect to w . However, we cannot pass this function to `minimize` directly, since the call

```
v = f(w_init)
```

must pass without any error. Define a local function inside the `fit_regr_lin_by_minimization` like this:

```
def fit_regr_lin_by_minimization(w, X, y):  
    # Possibly some code  
    def f(w): return compute_cost_regr_lin(w, X, y)  
    # Some other code
```

Then you can use function `f(w)` as an argument to `minimize`.

4. As the initial guess `w_init`, you can use your hand-crafted w , or e.g. $(0, 0)$, or $(100, 1)$; it does not matter much in this case.

5 Model fitting using the normal equation

Task 7: In `linreg.py`, create function `fit_regr_lin_by_normal_equation(X, y)` which

- takes the training data T as input (X, y) , and
- produces optimal values of w as a Numpy array by implementing the equation

$$w^* = (X^T X)^{-1} X^T y.$$

Hints:

1. Do not forget to homogenize X .
2. You should know how to do matrix-matrix and matrix-vector multiplication using the `numpy.ndarray.dot()` method or using the `@` symbol.
3. Have a look at `numpy.linalg.inv()` function for inverting a matrix.

Now, you should see the predictions of the optimal linear model in the figure. It is probably not very different from your hand-tuned one.

6 Linear model using `scikit.learn` package

In the above tasks, you learned how to fit a linear regression model manually for educational purposes. Of course, in practice you shall take advantage of a ready-to-use package. In Python, one possibility is to use the `scikit.learn` package. It is a collection of functions and classes for machine learning and modeling. Since we will use it in the forthcoming exercises, let's try the same thing using the `scikit.learn` API.

Task 8: In the respective section of `ML02.py`, fill in the code to train a linear model and get its predictions using `scikit.learn` API.

Hints:

1. Take a look at `sklearn.linear_models.LinearRegression`.
2. The model fitting is done by the `sklearn.linear_models.LinearRegression.fit()` method.
3. The prediction is done by the `sklearn.linear_models.LinearRegression.predict()` method.
4. The parameters of the linear model can be found in `LinearRegression.intercept_` property (w_0), and in `LinearRegression.coef_` property (w_1, \dots, w_D).

7 Multivariate linear regression

So far we have worked with the univariate case, $x = x, D = 1$. Let's try similar thing as above, but for the multivariate case, i.e. try to predict the `#points_exam` on the basis of other variables such as `#hours_exam`, `#hours_per_week`, `#hours_per_week_ai'`, `#lecture_attendance`.

Task 9: In the respective section of `ML02.py`, fill in the code to load the data for multivariate regression, train the multivariate model and compute its error.

Hints: Do you really need any?

The resulting error shall be lower than in the case of simple regression, since the model uses more knowledge about the students, and can thus account for more variability in the data.

8 Summary

In this exercise, we have used three ways of learning the linear regression model: the minimization of cost function J , the normal equation, and the `scikit-learn` implementation.

The approach based on error minimization is very general and can, in principle, be used with any kind of model. In the next exercises, we will mostly use the `scikit-learn` methods.

Complete the exercise as a homework, ask questions on the forum, and upload the solution via Upload system (BRUTE)!