

DM561 – Linear Algebra with Applications

1 Notes on Linear Regression

In *curve fitting* in \mathbb{R}^2 we are given m points (pairs of numbers) $(x_1, y_1), \dots, (x_m, y_m)$ and we want to determine a function $f(x)$ such that

$$f(x_1) \approx y_1, \dots, f(x_m) \approx y_m.$$

The type of function (for example, polynomials, exponential functions, sine and cosine functions) may be suggested by the nature of the problem (the underlying physical law, for instance), and in many cases a polynomial of a certain degree will be appropriate.

Let's assume we have a set of data collected by some measurements. For example, they can be temperature in the atmosphere taken at different days of the year or the cost of houses given their square meters.

For carrying out the computation we will simulate the set of points using the following Python script.

```
#!/usr/bin/python
import numpy as np
import matplotlib.pyplot as plt

np.random.seed(2019)
np.set_printoptions(precision=3)

m=101
x = np.linspace(0, 1, m)
y = x**3-7*x+np.random.exponential(1,m)

plt.plot(x, y, '.')
```

```
#plt.axis.xlabel=('x')
```

```
#plt.axis.Axis.ylabel=('x')
```

```
plt.show()

print(x)
print(y)
```

The script produces two arrays x and y . The i th elements of these arrays give us the i th point, (x_i, y_i) . An example of collected points is shown in Figure 1.

1.1 Linear function

The first case we consider is fitting a straight line $y = a + bx$. Clearly, there is no line that passes through all the points at the same time. However, we can search for the line that minimizes the distance from all the points. More precisely, given the points $(x_1, y_1), \dots, (x_m, y_m)$ we search for the line such that the sum of the squares of the distances of those points from the straight line is minimum, where the distance from (x_i, y_i) is measured in the vertical direction (the y -direction).

Formally, each point i with abscissa x_i has the ordinate $a + bx_i$ in the fitted line. The distance from the actual data (x_i, y_i) is thus $|y_i - a - bx_i|$ and the sum of squares is

$$q = \sum_{i=1}^m (y_i - a - bx_i)^2$$




Figure 1: An example of points generated by the Python script plotted on the plane xy .

Hence, q depends on a and b whose values we are trying to determine, while the values x_i and y_i are given being the coordinates of the points available. From calculus we know that the minimum of a function occurs where the partial derivatives are zero.

Task: Calculate the partial derivatives and find a and b with the help of Python. Report the symbolic derivations, the code you wrote and a plot that shows the points you generated and the line you fitted. The code must be short and use as much as possible matrix calculations by means of numpy or scipy arrays.

[Hint: you can look up in Wikipedia for an explanation of partial derivatives. In short, the partial derivatives of q with respect to the variables a and b , denoted $\frac{\partial q}{\partial a}$ and $\frac{\partial q}{\partial b}$, are, respectively, the derivatives of q with respect to a and b while the other variable is considered like a constant.]

Solution:

A necessary condition for q to be minimum is

$$\frac{\partial q}{\partial a} = -2 \sum_{i=1}^m (y_i - a - bx_i) = 0$$

$$\frac{\partial q}{\partial b} = -2 \sum_{i=1}^m x_i (y_i - a - bx_i) = 0$$

We can rewrite as:

$$\begin{cases} am + b \sum x_i = \sum y_i \\ a \sum x_i + b \sum x_i^2 = \sum x_i y_i \end{cases}$$

which is a system of linear equations in the variables $[a, b]$. Hence, the solution to this systems gives the values of a and b that minimize the square distance.

```
# Task 1
a_11 = len(x)
a_12 = sum(x)
a_22 = sum(x**2)
b_1 = sum(y)
b_2 = sum(x*y)

A = np.array([[a_11, a_12], [a_12, a_22]])
b = np.array([b_1, b_2])

a_11 = len(x)
```




Figure 2:

```

a_12 = sum(x)
a_22 = sum(x**2)
b_1 = sum(y)
b_2 = sum(x*y)

A = np.array([[a_11, a_12], [a_12, a_22]])

coeff = np.linalg.solve(A,b)

# y = coeff[0] + coeff[1]*x

plt.plot(x, y, '.')
```

The resulting plot is reported in Figure 2.

1.2 Polynomial Function

Let us now try to obtain a better fit by using a second order polynomial approximation. We can generalize the polynomial $y = ax + b$ to a polynomial of degree k

$$p(x) = b_0 + b_1x + \dots + b_kx^k$$

where $k \leq m - 1$. Then q takes the form

$$q = \sum_{i=1}^m (y_i - p(x_i))^2$$

and depends on $k + 1$ parameters b_0, \dots, b_k . As before the setting of the parameters that yields the minimum q can be found via partial derivatives.

For the case of a second order polynomial, ie, for $k = 2$, calculate the partial derivatives and find the parameters b_0, \dots, b_2 with the help of Python. Report the symbolic derivations, the code you wrote and a plot that shows the points you generated and the cubic curve you fitted. The code must be short and use as much as possible matrix calculations by means of numpy or scipy arrays.

Discuss which of the two curves fitted, the straight line and the second order polynomial, yields the best model for the points.

Solution:

The necessary condition for q to be minimum gives a $k + 1$ system of linear equations:

$$\begin{aligned}\frac{\partial q}{\partial b_0} &= 0 \\ \frac{\partial q}{\partial b_1} &= 0 \\ &\vdots \\ \frac{\partial q}{\partial b_k} &= 0\end{aligned}$$

For $k = 3$ the system we obtain is (summations are all from 1 to m):

$$\begin{cases} b_0 m + b_1 \sum x_i + b_2 \sum x_i^2 + b_3 \sum x_i^3 = \sum y_i \\ b_0 \sum x_i + b_1 \sum x_i^2 + b_2 \sum x_i^3 + b_3 \sum x_i^4 = \sum x_i y_i \\ b_0 \sum x_i^2 + b_1 \sum x_i^3 + b_2 \sum x_i^4 + b_3 \sum x_i^5 = \sum x_i^2 y_i \\ b_0 \sum x_i^3 + b_1 \sum x_i^4 + b_2 \sum x_i^5 + b_3 \sum x_i^6 = \sum x_i^3 y_i \end{cases}$$

which is a system of linear equations in the variables $[b_0, \dots, b_k]$. Hence, the solution to this systems gives the values of $[b_0, \dots, b_k]$ that minimize the square distance.

```
# Task 2
a_0 = len(x)

a_1 = sum(x)
a_2 = sum(x**2)
a_3 = sum(x**3)

b_1 = sum(y)

a_4 = sum(x**4)
b_2 = sum(x*y)

a_5 = sum(x**5)
b_3 = sum((x**2)*y)

a_6 = sum(x**6)
b_4 = sum((x**3)*y)

A = np.array([[a_0, a_1, a_2, a_3],
              [a_1, a_2, a_3, a_4],
              [a_2, a_3, a_4, a_5],
              [a_3, a_4, a_5, a_6]
             ])

coeff = np.linalg.solve(A, [b_1, b_2, b_3, b_4])

P=np.poly1d(coeff[::-1])
print(P)

plt.plot(x, y, 'r.')
xx = np.linspace(0.0, 1.0, 50)
plt.plot(xx, P(xx), 'b-')
```

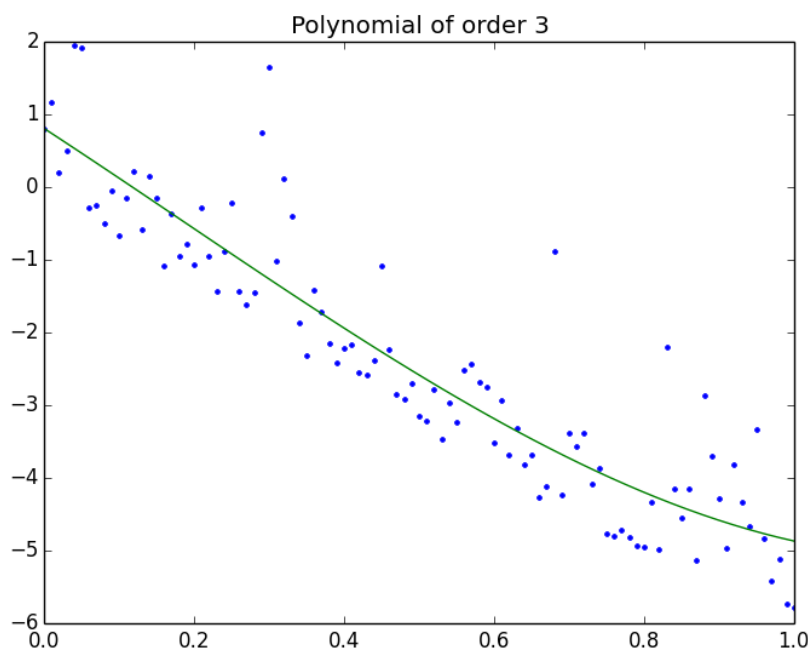


Figure 3:

```
#plt.axhline(y=0)
plt.title('Polynomial of order 3');
plt.show()
```

The resulting plot is reported in Figure 3.

We see that the third degree polynomial fits better the points since the line has more degree of freedom. The decision about which fitting is the best may depend on the application and on the knowledge about the underlying model. If it is known that there is a linear correspondence between the variables then the model to use is the linear. Otherwise, a third degree might be better provided that it does not overfit the data. Cross validation is a technique used in statistical learning to avoid overfitting.

1.3 Linear Algebra Approach

In the previous points we moved forward to a solution using calculus. Now we take a different approach using linear algebra and we will show that we reach the same result. In doing this we generalize the method to the space \mathbb{R}^n . This is useful, for example, in tasks of machine learning where one is interested in learning from data the type of dependency of a variable y on some predictor variables x_1, \dots, x_n . For example, the temperature on the ground may depend, beside on the day of the year, also on the latitude and maybe, in the case of global warming, on the development through years. The price of a house may depend on the square meters but also on the distance from city center and the year of construction.

We assume that the variable y is related to $\mathbf{x} \in \mathbb{R}^n$ linearly, so for some constants b_0 and \mathbf{b} , $y = b_0 + \mathbf{b}^T \mathbf{x}$. Given the set of m points in the ideal case we have that $y_i = b_0 + \mathbf{b}^T \mathbf{x}_i$, for all $i = 1, \dots, m$. In matrix form:

$$\begin{bmatrix} 1 & x_{1,1} & \cdots & x_{1,k} \\ 1 & x_{2,1} & \cdots & x_{2,k} \\ \vdots & \vdots & & \vdots \\ 1 & x_{m,1} & \cdots & x_{m,k} \end{bmatrix} \begin{bmatrix} b_0 \\ b_1 \\ \vdots \\ b_k \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \cdots \\ y_m \end{bmatrix}$$

This can be written as $\mathbf{Az} = \mathbf{y}$.

Considering the same set of data in \mathbb{R}^2 from point 1 of this Task try to find \mathbf{z} in Python. Report what happens and explain why it is so.




Figure 4: $\mathbf{y} \in \mathbb{R}^2$ and A is a 3×2 matrix of rank 2.

Solution:

```
# Task 4

A = np.column_stack([np.ones(101), x])

np.linalg.solve(A, y)
```

Python will complain that the matrix is not square. The system is overdetermined.

In general we cannot expect to find a vector $\mathbf{z} \in \mathbb{R}^{n+1}$ for which $A\mathbf{z}$ equals \mathbf{y} . Instead we can look for a vector \mathbf{z} for which $A\mathbf{z}$ is closest to \mathbf{y} .

For each $\mathbf{z} \in \mathbb{R}^{n+1}$ we can form the *residual*

$$r(\mathbf{z}) = \mathbf{y} - A\mathbf{z}$$

Recalling the definition of the norm or length of a vector, the distance between \mathbf{y} and $A\mathbf{z}$ is given by

$$\|r(\mathbf{z})\| = \|\mathbf{y} - A\mathbf{z}\|$$

Minimizing $\|r(\mathbf{z})\|$ is equivalent to minimize $\|r(\mathbf{z})\|^2$. So what we need is for $A\mathbf{z}$ to be the closest point of the form $A\mathbf{u}$ (for a generic vector $\mathbf{u} \in \mathbb{R}^{n+1}$) to \mathbf{y} . That is, $A\mathbf{z}$ has to be the closest point in $R(A)$ to \mathbf{y} .

Thus a vector $\hat{\mathbf{z}}$ will be the one that minimizes $\|\mathbf{y} - A\mathbf{z}\|^2$ if and only if $\mathbf{p} = A\hat{\mathbf{z}}$ is the vector in $R(A)$ that is closest to \mathbf{y} . The vector \mathbf{p} is said to be the *projection* of \mathbf{y} onto $R(A)$. It follows that

$$r(\hat{\mathbf{z}}) \in R(A)^\perp$$

that is, $r(\hat{\mathbf{z}})$ lays on the subspace of \mathbb{R}^m that is the *orthogonal complement* of $R(A)$. These facts are explained visually in Figure 4 for the case of \mathbb{R}^2 but they can be proved to hold in more general terms for \mathbb{R}^n .

To proceed we need another fact. For an $m \times (n + 1)$ matrix A ,

$$R(A)^\perp = N(A^T),$$

that is, the orthogonal complement of the range of a matrix A is the null space of the transpose of the matrix A .

The proof of this fact is as follows:

A vector \mathbf{w} that belongs to $R(A)$ belongs to the linear span of A , ie, it is a linear combination of the columns of the matrix A . A vector \mathbf{v} that belongs to $R(A)^\perp$, must therefore be orthogonal to each column of the matrix A . Consequently $A^T\mathbf{v} = \mathbf{0}$. Thus, \mathbf{v} must be an element of $N(A^T)$ and hence $N(A^T) = R(A)^\perp$. Thus we have

$$r(\hat{\mathbf{z}}) \in N(A^T)$$

Show that this leads to a $(n + 1) \times (n + 1)$ system of linear equations in the \mathbf{z} variables. The system is actually the same as the one derived earlier via calculus.

Solution:

Since $r(\hat{\mathbf{z}}) = \mathbf{y} - A\mathbf{z}$ we have

$$A^T(\mathbf{y} - A\mathbf{z}) = \mathbf{0}$$

Thus to solve the least square system of $A\mathbf{z} = \mathbf{y}$ we have to solve

$$A^T\mathbf{y} = A^T A\mathbf{z}$$

Remembering that we are solving in the \mathbf{z} variables we have a system of size $(n + 1) \times (n + 1)$ because A is of size $m \times (n + 1)$ and therefore $A^T A$ is of size $(n + 1) \times (n + 1)$.

Task: Show that if A is an $m \times (n + 1)$ matrix of rank $(n + 1)$, then the system at the previous point has a unique solution:

$$\hat{\mathbf{z}} = (A^T A)^{-1} A^T \mathbf{y}$$

(Hint: assume that $A^T A\mathbf{u} = \mathbf{0}$ has only the trivial solution $\mathbf{u} = \mathbf{0}$.)

Solution:

Under the assumption that $A^T A\mathbf{u} = \mathbf{0}$ has only the trivial solution $\mathbf{u} = \mathbf{0}$ then the matrix $A^T A$ is non-singular and hence invertible and the solution unique.

Task Use the formula derived above to find the least square linear regression of point (a). Compare your result with the result you would obtain by using the function from numpy: `numpy.linalg.lstsq`.

Solution:

In Python, the least square solution as derived here to our original data can be computed as:

```
### http://docs.scipy.org/doc/numpy/reference/generated/numpy.linalg.lstsq.html

A = np.vstack([np.ones(len(x)), x]).T

b_0, b_1 = np.linalg.lstsq(A, y)[0]
print(b_0, b_1)

# Plot the data along with the fitted line:

plt.plot(x, y, 'o', label='Original data', markersize=10)
plt.plot(x, b_1*x + b_0, 'r', label='Fitted line')
plt.legend()
plt.show()
```

1.4 Other Applications of Linear Regression

We now turn to another application in which we know that the data sampled come from a circle. For example, in image processing from a medical application we may collect the data regarding points of a bone and we want to determine the circumference of the bone.

The parametric equations for a circle with center $(3, 1)$ and radius 2 are

$$\begin{aligned} x &= 3 + 2 \cos t \\ y &= 1 + 2 \sin t \end{aligned}$$

We simulate a few points in Python by adding some noise to the points from a circle in correspondence of $t \in \{0.0, 0.5, 1.0, \dots, 5.0, 5.5\}$:

```
#!/usr/bin/python
import numpy as np
import matplotlib.pyplot as plt

CPRN = 3 # your CPR number
np.random.seed(CPRN)
np.set_printoptions(precision=3)

t=np.arange(0,6,0.5)
x=3+2*np.cos(t)+0.3*np.random.rand(len(t))
y=1+2*np.sin(t)+0.3*np.random.rand(len(t))

plt.plot(x, y, '.')
plt.show()
```

Write in your report the symbolic derivations that lead to finding the center c and radius r of the circle that gives the best q that fits to the points, ie:

$$\begin{aligned}x &= c_1 + r \cos(t) \\ y &= c_2 + r \sin(t)\end{aligned}$$

Solve the model you found in Python. Report the code and the plot of the points and the fitted circle.

Solution:

The circle can be written as

$$(x - c_1)^2 + (y - c_2)^2 = r^2$$

Solving in y we obtain

$$y = \pm \sqrt{r^2 - (x - c_1)^2} + c_2$$

Then the least squares gives:

$$q = \sum (y_i \pm \sqrt{r^2 - (x_i - c_1)^2} + c_2)$$

which we need to optimize in r^2 , c_1 , c_2 . This leads to a non linear system of equations that is not convenient for us.

Luckily, we are given the points by means of their radial coordinate t rather than just with the (x, y) coordinates. (This is seldom the case since it implies knowing a priori the model that we are actually trying to find.)

Let's then try to minimize the squared distance of the points from the circle in the parameters r^2 , c_1 , c_2 :

$$\begin{aligned}q &= \sum_{i=1}^m [(x_i - c_1 - r \cos(t_i))^2 + (y_i - c_2 - r \sin(t_i))^2] \\ &= \sum_{i=1}^m [(x_i - c_1)^2 + r^2 \cos^2(t_i) - 2(x_i - c_1)(r \cos(t_i)) + (y_i - c_2)^2 + r^2 \sin^2(t_i) - 2(y_i - c_2)(r \sin(t_i))] \\ &= \sum_{i=1}^m [(x_i - c_1)^2 - 2(x_i - c_1)(r \cos(t_i)) + (y_i - c_2)^2 - 2(y_i - c_2)(r \sin(t_i)) + r^2]\end{aligned}$$

$$\frac{\partial q}{\partial r} = \sum_{i=1}^m [-2(x_i - c_1) \cos(t_i) - 2(y_i - c_2) \sin(t_i) + 2r] = 0$$

$$\frac{\partial q}{\partial c_1} = \sum_{i=1}^m [-2(x_i - c_1) + 2(r \cos(t_i))] = 0$$

$$\frac{\partial q}{\partial c_2} = \sum_{i=1}^m [-2(y_i - c_2) + 2(r \sin(t_i))] = 0$$




Figure 5:

Luckily, this gives rise to a linear system:

$$\begin{bmatrix} m & \sum \cos(t_i) & \sum \sin(t_i) \\ \sum \cos(t_i) & m & 0 \\ \sum \sin(t_i) & 0 & m \end{bmatrix} \begin{bmatrix} r \\ c_1 \\ c_2 \end{bmatrix} = \begin{bmatrix} \sum x_i \cos(t_i) + y_i \sin(t_i) \\ \sum x_i \\ \sum y_i \end{bmatrix}$$

Solving it we find r, c_1, c_2 .

```

m=len(x)
A = np.array([[m,np.sum(np.cos(t)),np.sum(np.sin(t))],
              [np.sum(np.cos(t)),m,0],
              [np.sum(np.sin(t)),0,m]]
             )
b = np.array([np.sum(x*np.cos(t))+y*np.sin(t)),
             np.sum(x),
             np.sum(y)])

rcc = np.linalg.solve(A,b)

print(rcc)
t=np.arange(0,2*np.pi,0.1)
xx=rcc[1]+rcc[0]*np.cos(t)
yy=rcc[2]+rcc[0]*np.sin(t)

plt.plot(x, y, '.')
plt.plot(xx,yy,'-')
plt.show()

```

The result is shown in Figure 5.