

Regression

based on a document by Andrew Ng

September 24, 2008

1 Linear regression

Lets consider a supervised learning example where data points are houses with 2 features (X^1 = living area; X^2 = number of bedrooms) and labels are the prices.

Living area (ft ²)	#bedrooms	price (1000\$s)
2104	3	400
1600	3	330
2400	3	369
1416	2	232
3000	4	540
...

Each datapoint is thus (\mathbf{x}_t, y_t) , where $\mathbf{x}_t = (x_t^1, x_t^2)$ are the living area and the number of bedrooms respectively, and y_t is the price. The main problem addressed by *regression* is: Can we predict the price (or label) y_t from input features \mathbf{x}_t ? Today we study *linear regression*, that is if we can predict the price from the features using a linear regressor

$$h_{\theta}(\mathbf{x}) = \theta^0 + \theta^1 x^1 + \theta^2 x^2$$

where $\theta = (\theta^0, \theta^1, \theta^2)$ are the parameters of the regression function. Within the class of linear functions (regressors) our task shall be to find the best parameters θ . When there is no risk of confusion, we will drop θ from the h notation, and we will assume a dummy feature $x^0 = 1$ for all datapoints such that we can write

$$h(\mathbf{x}) = \sum_{d=0}^D \theta^d x^d$$

where d iterates through input features $1, 2, \dots, D$ (in our example $D = 2$).

What do we mean by the best regression fit? For today, we will measure the error (or cost) of the regression function by the *mean square error*

$$J(\theta) = \sum_t (h_{\theta}(\mathbf{x}_t) - y_t)^2$$

and we will naturally look for the θ that minimizes the error function. The regression obtained using the square error function J is called *least square regression*, or *least square fit*. We present two methods for minimizing J .

2 Least mean squares via Gradient descent

2.0.1 Gradient descent

The gradient descent algorithm finds a local minima of the objective function (J) by guessing an initial set of parameters θ and then "walking" episodically in the direction of the gradient $\partial J/\partial\theta$. Since θ is vector valued (3 components for our example) we need to perform the update for each component separately

$$\theta^j = \theta^j - \lambda \frac{\partial J(\theta)}{\partial \theta^j}$$

where λ is the *learning rate* parameter or the step of the update. To see this in practice, lets consider an example (other than the mean square error): say $J(x) = (x - 2)^2 + 1$ and the initial guess for a minimum is $x_0 = 3$. The differential of J is

$$\frac{\partial J(x)}{\partial x} = 2x - 4$$

Assuming a fixed $\lambda = 0.25$, the first several episodes of gradient descent are:

$$x_1 = x_0 - \lambda \frac{\partial J(x_0)}{\partial x} = 3 - .25(2 * 3 - 4) = 2.5$$

$$x_2 = x_1 - \lambda \frac{\partial J(x_1)}{\partial x} = 2.5 - .25(2 * 2.5 - 4) = 2.25$$

$$x_3 = x_2 - \lambda \frac{\partial J(x_2)}{\partial x} = 2.25 - .25(2 * 2.25 - 4) = 2.125$$

Figure1 illustrates the episodic process.




Figure 1: Gradient descent iterations

2.1 Least mean squares

In order to apply gradient descent to our problem, we need first to differentiate our error (objective) function J .

$$J(\theta) = \frac{1}{2} \sum_t (h_\theta(\mathbf{x}_t) - y_t)^2$$

Lets look at the differential for the case when there is only one datapoint (so there are no sums)

$$\begin{aligned} \frac{\partial J(\theta)}{\partial \theta^j} &= \frac{\partial \frac{1}{2} (h_\theta(\mathbf{x}) - y)^2}{\partial \theta^j} \\ &= (h_\theta(\mathbf{x}) - y) \frac{\partial (h_\theta(\mathbf{x}) - y)}{\partial \theta^j} \\ &= (h_\theta(\mathbf{x}) - y) \frac{\partial (\sum_j \theta^j x^j - y)}{\partial \theta^j} \\ &= (h_\theta(\mathbf{x}) - y) x^j \end{aligned}$$

Thus the gradient descent update rule (for one data point) is

$$\theta^j = \theta^j - \lambda (h_\theta(\mathbf{x}) - y) x^j$$

The update has to be performed for all components $j = 1, 2, \dots, D$ simultaneously (in parallel).

There are two ways to adapt the rule derived above for the case where many datapoints. The first one, *batch gradient descent*, is to use the error function J for all datapoints when differentiate (essentially make the update for all points at once). The sum inside J expression propagates to the differential expression and the update becomes

$$\theta^j = \theta^j - \lambda \sum_t (h_\theta(\mathbf{x}_t) - y_t) x_t^j$$

for all $j=1, 2, \dots, D$. Note that in batch gradient descent, a single update require analysis of all datapoints; this can be very tedious if the data set is large.

The second way to involve all datapoints is to make the update separately for each datapoint (*stochastic gradient descent*). Each update step becomes then a loop including all datapoints:

LOOP for $t=1$ to m

$$\theta^j = \theta^j - \lambda (h_\theta(\mathbf{x}_t) - y_t) x_t^j \text{ for all } j$$

END LOOP

Stochastic (or online) gradient descent advances by updating based on one datapoint at a time. if the regression problem is not too complicated, often few iterations are enough to converge and so in practice this version of gradient decent is often faster. it is possible (although very unlikely) to have a problem where stochastic updates "dance" around the best θ (that realizes minimum error) without actually converging to it.

To conclude regression via gradient descent, we make one final observation. The objective function J is convex because its second differential is positive

$$\frac{\partial^2 J(\theta)}{\partial (\theta^j)^2} = (x^j)^2 \geq 0 \text{ for all } j$$

which means any local minima is in fact a global minima, thus the gradient descent (or any method that finds local minima) finds a global minima.

3 Least mean square via normal equations

3.1 Matrix derivatives

Let $f : \mathbb{R}^{m \times n} \mapsto \mathbb{R}$ a function that takes a matrix as input and outputs a real number. We define the derivative of f with respect to the matrix A as

$$\nabla_A f(A) = \begin{bmatrix} \frac{\partial f}{\partial a_{11}} & \cdots & \frac{\partial f}{\partial a_{1n}} \\ \cdots & \cdots & \cdots \\ \frac{\partial f}{\partial a_{m1}} & \cdots & \frac{\partial f}{\partial a_{mn}} \end{bmatrix}$$

where a_{ij} is the element of A on row i and column j . For example, consider $A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$ and the function $f(A) = \frac{3}{2}a_{11} + 5a_{12}^2 + a_{21}a_{22}$. Then

$$\nabla_A f(A) = \begin{bmatrix} \frac{3}{2} & 10a_{12} \\ a_{22} & a_{21} \end{bmatrix}$$

Today we will look at the *trace* function as f . The trace of a matrix A is the sum of the elements of the main diagonal

$$\text{tr}(A) = \sum_i a_{ii}$$

The following are simple and well known properties of the trace:

$$\begin{aligned} \text{tr}(AB) &= \text{tr}(BA) \\ \text{tr}(A) &= \text{tr}(A^T) \\ \text{tr}(A + B) &= \text{tr}(A) + \text{tr}(B) \\ \text{tr}(xA) &= x\text{tr}(A) \end{aligned}$$

The following properties of the trace matrix derivative are going to be useful for finding an exact regression solution:

$$\nabla_A \text{tr}(AB) = B^T \tag{1}$$

$$\nabla_{A^T} f(A) = (\nabla_A f(A))^T \tag{2}$$

$$\nabla_A \text{tr}(ABA^T C) = CAB + C^T AB^T \tag{3}$$

Combining the second and third statements above we get

$$\nabla_{A^T} \text{tr}(ABA^T C) = B^T A^T C^T + BA^T C \tag{4}$$

3.2 An exact solution for regression using linear algebra

Given the datapoints (\mathbf{x}_t, y_t) for $t = 1, 2, \dots, m$, with D input dimensions (features), we shall look at them in a matrix form

$$X = \begin{bmatrix} x_1^1 & \cdots & x_1^D \\ \cdots & \cdots & \cdots \\ x_m^1 & \cdots & x_m^D \end{bmatrix} \quad Y = \begin{bmatrix} y_1 \\ \cdots \\ y_m \end{bmatrix}$$

Then the error array associated with our regressor $h_\theta(\mathbf{x}) = \sum_d \theta^d x^d$ is

$$E = \begin{bmatrix} h_\theta(\mathbf{x}_1) - y_1 \\ \cdots \\ h_\theta(\mathbf{x}_m) - y_m \end{bmatrix} = \begin{bmatrix} \mathbf{x}_1 \theta \\ \cdots \\ \mathbf{x}_m \theta \end{bmatrix} - \begin{bmatrix} y_1 \\ \cdots \\ y_m \end{bmatrix} = X\theta - Y$$

(we used $\theta = (\theta^0, \theta^1, \dots, \theta^d)^T$ as a vector column). We can now write the mean square error as

$$J(\theta) = \frac{1}{2} \sum_t (h_\theta(\mathbf{x}_t) - y_t)^2 = \frac{1}{2} E^T E = \frac{1}{2} (X\theta - Y)^T (X\theta - Y)$$

Then

$$\begin{aligned} \nabla_\theta J(\theta) &= \nabla_\theta \frac{1}{2} E^T E = \nabla_\theta \frac{1}{2} (X\theta - Y)^T (X\theta - Y) \\ &= \frac{1}{2} \nabla_\theta (\theta^T X^T X\theta - \theta^T X^T Y - Y^T X\theta + Y^T Y) \\ &= \frac{1}{2} \nabla_\theta \text{tr}(\theta^T X^T X\theta - \theta^T X^T Y - Y^T X\theta + Y^T Y) \\ &= \frac{1}{2} \nabla_\theta (\text{tr}(\theta^T X^T X\theta) - 2\text{tr}(Y^T X\theta)) \\ &= \frac{1}{2} (X^T X\theta + X^T X\theta - 2X^T Y) \\ &= X^T X\theta - X^T Y \end{aligned}$$

because: in the third step we have $\text{tr}(x) = x$, in the four step we have $\text{tr}(A) = \text{tr}(A^T)$, and in the fifth step we are using equation 4 with $A^T = \theta$, $B = B^T = X^T X$, $C = I$.

Since we are trying to minimize J , a convex function, a sure way to find θ that minimizes J is to set its derivative to zero. In doing so we obtain

$$X^T X\theta = X^T Y \text{ or } \theta = (X^T X)^{-1} X^T Y$$

This is the exact θ that minimizes the mean square error.

4 Least mean square probabilistic interpretation

Why mean square error? We show now that the objective J used is a direct consequence of a very common assumption over the data. Lets look at the errors

$$\epsilon_t = h(\mathbf{x}_t) - y_t$$

and lets make the assumption that they are IID according to a gaussian (normal) distribution of mean $\mu = 0$ and variance σ^2 . That we write $\epsilon \sim \mathcal{N}(0, \sigma^2)$ or

$$p(\epsilon) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\epsilon^2}{2\sigma^2}\right)$$

which implies

$$p(y|x; \theta) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(\theta\mathbf{x} - y)^2}{2\sigma^2}\right)$$

Note that above θ is a parameter (array) and not a random variable. Given the input X , what is the probability of Y given the parameters θ ? Equivalently, this is the *likelihood* that θ is the correct parameter for the model

$$L(\theta) = L(\theta; X, Y) = p(Y|X; \theta)$$

Using the IID assumption the likelihood becomes

$$\begin{aligned} L(\theta) &= \prod_t p(y_t|\mathbf{x}_t; \theta) \\ &= \prod_t \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(\theta\mathbf{x}_t - y_t)^2}{2\sigma^2}\right) \end{aligned}$$

since we have now a probabilistic model over the data, a common way to determine the best parameters is to use *maximum likelihood*; in other words find θ that realizes the maximum L . Instead of maximizing L we shall maximize the *log likelihood* $\log L(\theta)$ because it simplifies the math (and produces the same "best" θ)

$$\begin{aligned}
 l(\theta) &= \log L(\theta) \\
 &= \log \prod_t \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(\theta\mathbf{x}_t - y_t)^2}{2\sigma^2}\right) \\
 &= \sum_t \log \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(\theta\mathbf{x}_t - y_t)^2}{2\sigma^2}\right) \\
 &= m \log \frac{1}{\sqrt{2\pi}\sigma} - \frac{1}{2\sigma^2} \sum_t (h_\theta(\mathbf{x}_t) - y_t)^2
 \end{aligned}$$

Hence, the maximizing the likelihood L produces the same θ as minimizing the mean square error (since the front term does not depend on θ). That is to say that, if we believe the errors to be IID normally, then the maximum likelihood is obtained for the parameters θ that minimizes the mean square error.

5 Polynomial regression

6 Classification and logistic regression

In classification, the labels y are not numeric values (like prices), but instead *class labels*. For today, lets assume that we have two classes denoted by 0 and 1; we call this *binary classification* and we write $y \in \{0, 1\}$.

6.1 Logistic transformation

We could, in principle try to run the linear regression we just studied, without making use of the fact that $y \in \{0, 1\}$. (Essentially assume y are simply real numbers). There are several problem with this approach: first, the regression assumes the data supports a linear fit, which might not be true anymore for classification problems; second, our regressor $h(\mathbf{x})$ will take lots of undesirable values (like the ones far outside the interval $[0,1]$).

To make an explicit mapping between the real valued regressor h and the set $\{0,1\}$, we would like a function that preserves differentiability and has a easy interpretable meaning. We choose the *logistic function*, also called *sigmoid*

$$g(z) = \frac{1}{1 + e^{-z}}$$

Note that g acts like a indicator for $\{0,1\}$, but it is much more sensitive than a linear function. We can make it even more sensitive by, for example, doubling the input z before applying the function. Lets state the derivative of g , since we are going to use it later on

$$\begin{aligned}
 g'(z) &= \frac{\partial g(z)}{\partial z} \\
 &= \frac{1}{(1 + e^{-z})^2} e^{-z} \\
 &= \frac{1}{1 + e^{-z}} \left(1 - \frac{1}{1 + e^{-z}}\right) \\
 &= g(z)(1 - g(z))
 \end{aligned}$$




Figure 2: Logistic function

6.2 Logistic regression

We apply g to the linear regression function to obtain a *logistic regression*. Our new hypothesis (or predictor, or regressor) becomes

$$h_{\theta}(\mathbf{x}) = g(\theta\mathbf{x}) = \frac{1}{1 + e^{-\theta\mathbf{x}}} = \frac{1}{1 + e^{-\sum_d \theta^d x^d}}$$

7 Generalized linear models