

Linear Regression: Analytic Solution*(find and work with a partner)*

In the case of *multiple* linear regression, each example has p features. We still have an associated response (output) variable y for each example. We typically add a “fake 1” to the features of each example, so that $\vec{x} = [1 \ x_1 \ x_2 \ \cdots \ x_p]^T$. Our model is now the dot product between the weight vector $\vec{w}$ and the features:

$$h_{\vec{w}}(\vec{x}) = w_0 + w_1x_1 + \cdots w_px_p = \vec{w} \cdot \vec{x}$$

If we consider $\mathbf{X}$ to be the entire matrix of features (with dimensions $n \times (p+1)$) then we can reframe the linear regression problem in terms of vectors and matrices. In this case the analytic solution for the weight vector is:

$$\hat{w} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \vec{y}$$

1. *Warmup.* Given matrices $\mathbf{A}$ and $\mathbf{B}$ below, compute both $\mathbf{AB}$ and $\mathbf{BA}$. Is matrix multiplication commutative?

$$\mathbf{A} = \begin{bmatrix} 1 & -1 \\ 3 & 2 \end{bmatrix} \quad \text{and} \quad \mathbf{B} = \begin{bmatrix} 0 & 2 \\ -4 & 1 \end{bmatrix}$$

2. Going back to our small example from the last handout, we will use this matrix/vector form to verify our solution for $\vec{w}$ (even though $p = 1$ we can still use the multiple linear regression derivation). We again have the data: $(x_1, y_1) = (1, 0)$ and $(x_2, y_2) = (0, 1)$. First, rewrite $\mathbf{X}$ with the added column of 1's. Then write $\vec{y}$ as a vector.
3. Now use the matrix/vector analytic solution to verify the weight vector $\vec{w}$.

Linear Regression: SGD solution*(find and work with a partner)*

In linear regression, we seek to minimize the sum of squared errors between the actual response and our prediction. We often call this RSS (residual sum of squares) or SSE (sum of squared errors). As an objective function, we often call it J and include a $\frac{1}{2}$ in front to make the derivatives work out nicely.

$$J(\vec{w}) = \frac{1}{2} \sum_{i=1}^n (h_{\vec{w}}(\vec{x}_i) - y_i)^2$$

For linear regression in general, one iteration of stochastic gradient descent includes the following updates (usually with the data points shuffled):

```
for i = 1, 2, ..., n:  
     $\vec{w} \leftarrow \vec{w} - \alpha(\vec{w} \cdot \vec{x}_i - y_i)\vec{x}_i$ 
```

We will begin with our same data from the previous two handouts: $(x_1, y_1) = (0, 1)$ and $(x_2, y_2) = (1, 0)$, except we will reverse the order of the points to make the progress of gradient descent a bit clearer. So in this case our matrix/vector formulation is:

$$\mathbf{X} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}, \quad \vec{y} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

1. Assuming $\alpha = 0.1$ and our initial values are $w_0 = 0$ and $w_1 = 0$, what are w_0 and w_1 after the just the first data point is used to update the gradient?
2. What are w_0 and w_1 after the second data point is used? Since we only have two examples here, your result would be the weight vector after the first iteration of SGD.
3. What is the value of the objective function (cost) after this initial iteration?
4. *(optional)* Sketch the linear model after this initial iteration.