Machine Learning, Spring 2018: Exercise Sheet 7

It's paper and pencil time again.

Problem 1 (A very toy-ish demo of PCA) Assume you have a sample S of four 2-dimensional datapoints from $\mathbb{R}^2$, $S = \{(1,1)', (0,0)', (0,0)', (-1, -1)'\}$. What are the two principal component vectors $\mathbf{u}_1$, $\mathbf{u}_2$ of this dataset?

Problem 2 (understanding an extreme case of linear regression). Consider a slightly weird dataset $(x_i, y_i)_{i=1,\ldots,N} = (2, y_i)_{i=1,\ldots,N}$ where all argument-value pairs have the same argument $x_i = 2$ and $y_i \in \mathbb{R}$. If you carry out a linear regression, what is the regression weight vector w that you get from this dataset? It is rather easy to guess the answer – can you also *prove* it?

Problem 3 (understanding the mathematical nature of a combined PCA-feature-extraction + linear regression learning procedure) In the last paragraph of Section 6 in the lecture notes, I state that $D: \mathbb{R}^n \to \mathbb{R}^k$ is an affine map. That is, D can be written as $D(x) = M\,x + b$, where M is a $k \times n$ sized matrix and b is a k-dimensional vector. Work out the details – write down M and b as constructed from the n-dimensional mean pattern vector μ, the PC matrix U_m, and the regression weight matrix W.