Computing the Mean Is a Linear Least Squares Problem

In this example we demonstrate that computing the mean of student scores on an exam is a linear least squares problem, and in particular it amounts to an orthogonal projection from the high-dimensional space of the set of all sudents onto a one-dimensional space that conveys the overall performance of the class as a whole.

Suppose there are n students with test scores $t_1, t_2, \ldots, t_n$. Then the mean m is given by $m = (\sum_{i=1}^{n} t_i)/n$. Let the test scores be the entries of a vector $\boldsymbol{t} \in \mathbb{R}^n$, and let $\boldsymbol{e}$ be the n-vector having all entries equal to 1. We will now show that computing the mean is equivalent to solving the $n \times 1$ linear least squares problem

$$\boldsymbol{e}\, m \cong \boldsymbol{t}.$$

In other words, the problem has the form $\boldsymbol{Ax} \cong \boldsymbol{b}$, where in this case matrix $\boldsymbol{A}$ is $n \times 1$, $\boldsymbol{x}$ is a scalar, and $\boldsymbol{b} \in \mathbb{R}^n$. To solve this problem we use the normal equations to obtain

$$\boldsymbol{x} = (\boldsymbol{A}^T\boldsymbol{A})^{-1}\boldsymbol{A}^T b = (\boldsymbol{e}^T\boldsymbol{e})^{-1}\boldsymbol{e}^T\boldsymbol{t} = (1/n)\left(\sum_{i=1}^{n} t_i\right) = m.$$

Thus, the solution to the linear least squares problem is simply the mean score, m. To obtain the orthogonal projector $\boldsymbol{P}$, we compute

$$\boldsymbol{P} = \boldsymbol{A}(\boldsymbol{A}^T\boldsymbol{A})^{-1})\boldsymbol{A}^T = \boldsymbol{e}(1/n)\boldsymbol{e}^T = (1/n)\boldsymbol{e}\boldsymbol{e}^T.$$

To verify that $\boldsymbol{P}$ is indeed an orthogonal projector, we observe that it is obviously symmetric, and it is also idempotent, since

$$\boldsymbol{P}^2 = (1/n)\boldsymbol{e}\boldsymbol{e}^T(1/n)\boldsymbol{e}\boldsymbol{e}^T = (1/n^2)\boldsymbol{e}(\boldsymbol{e}^T\boldsymbol{e})\boldsymbol{e}^T = (1/n^2)\, n\, \boldsymbol{e}\boldsymbol{e}^T = (1/n)\boldsymbol{e}\boldsymbol{e}^T = \boldsymbol{P}.$$

To demonstrate that $\boldsymbol{P}$ actually works, we compute

$$\boldsymbol{Pb} = (1/n)\boldsymbol{e}\boldsymbol{e}^T\boldsymbol{t} = (1/n)\boldsymbol{e}\left(\sum_{i=1}^{n} t_i\right) = \boldsymbol{e}\, m = \boldsymbol{Ax} = \boldsymbol{y},$$

the vector in $\text{span}(\boldsymbol{A}) = \text{span}(\boldsymbol{e})$ closest to $\boldsymbol{t}$ in the 2-norm.