

Linear regression for function approximation: a crash course

We want to find an approximator $\hat{g}(x)$ for an unknown function $g(x)$, from a set of input-output samples:

$$\{(x_i, y_i = g(x_i) + w_i) \mid i = 1, \dots, N\}$$

where $g(x_i)$ is corrupted by noise w_i .

We choose a linearly parameterized form for $\hat{g}(x)$:

$$\hat{g}(x) = \varphi_1(x)\Theta_1 + \varphi_2(x)\Theta_2 + \dots + \varphi_n(x)\Theta_n = [\varphi_1(x) \dots \varphi_n(x)] \begin{bmatrix} \Theta_1 \\ \vdots \\ \Theta_n \end{bmatrix}$$

Notations: $\varphi^T(x) \in \mathbb{R}^n$ $\Theta \in \mathbb{R}^n$

where $\varphi_i(x)$ are the regressors (we choose them manually) and Θ_i the parameters (the algorithm will find them automatically.)

We write our objective: \hat{g} should match y_i for the samples we have.

$$\hat{g}(x_1) = \varphi_1(x_1)\Theta_1 + \dots + \varphi_n(x_1)\Theta_n = y_1$$

$$\hat{g}(x_2) = \varphi_1(x_2)\Theta_1 + \dots + \varphi_n(x_2)\Theta_n = y_2$$

⋮

$$\hat{g}(x_N) = \varphi_1(x_N)\Theta_1 + \dots + \varphi_n(x_N)\Theta_n = y_N$$

or, in matrix form:

$$\underbrace{\begin{bmatrix} \varphi_1(x_1) & \dots & \varphi_n(x_1) \\ \varphi_1(x_2) & \dots & \varphi_n(x_2) \\ \vdots & & \vdots \\ \varphi_1(x_N) & \dots & \varphi_n(x_N) \end{bmatrix}}_{\Phi \in \mathbb{R}^{N \times n}} \cdot \underbrace{\begin{bmatrix} \Theta_1 \\ \Theta_2 \\ \vdots \\ \Theta_n \end{bmatrix}}_{\Theta \in \mathbb{R}^n} = \underbrace{\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix}}_{Y \in \mathbb{R}^N}$$

Notations: $\Phi \in \mathbb{R}^{N \times n}$

The system is overdetermined, $N > n$, so we cannot solve it with equality. We will minimize the mean of squared errors between the two sides:

$$MSE = \frac{1}{N} \sum_{i=1}^N [\hat{g}(x_i) - y_i]^2$$

In Matlab, it is easy: the parameter vector Θ minimizing the MSE is found by writing $\Theta = \Phi \backslash Y$ (backslash).

Once Θ is found like this, then \hat{g} can be applied to approximate g at any point x_q not just those in the dataset. Just compute:

$$y_q = \hat{g}(x_q) = \sum_{i=1}^n \varphi_i(x_q)\Theta_i = \varphi^T(x_q)\Theta \quad (\uparrow \text{stands for "query"})$$

To be sure that the solution is good, we'll generally compute the MSE on a different, validation dataset:

$$\{(x_i^v, y_i^v) \mid i = 1, \dots, N^v\}$$

and plot g versus \hat{g} visually!