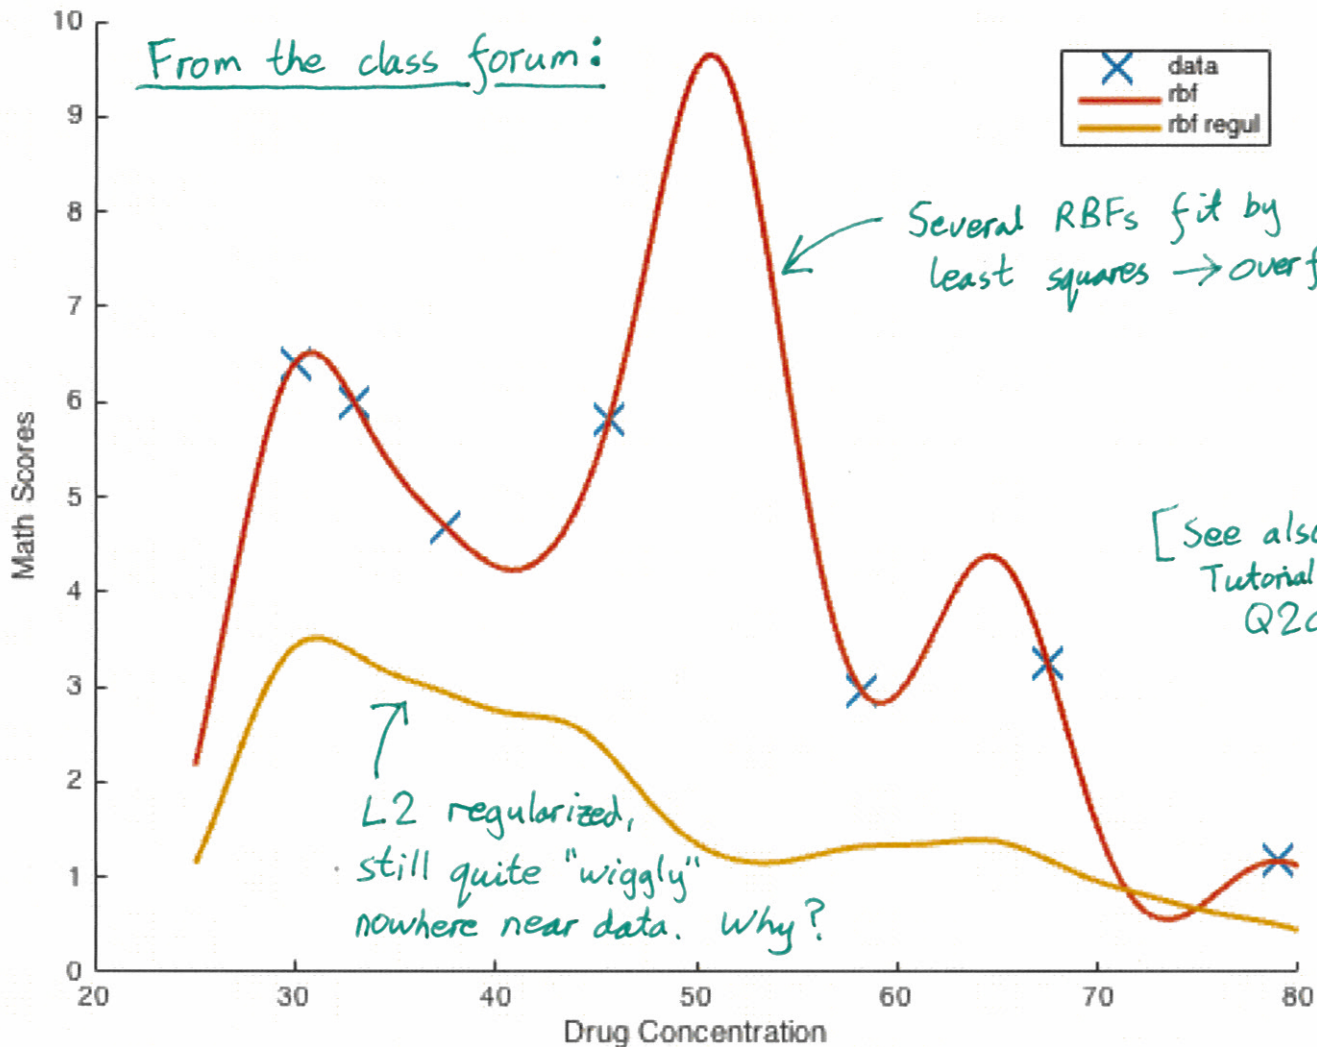



From the class forum:



Neural Networks

First example:





Fit $\{ \{ \underline{v}^{(k)}, b^{(k)} \}, \underline{w}, b \}$ with a gradient-based optimizer. Match f to training targets.

Vocabulary

"Unit" : One of the intermediate computations

"Hidden units" : Also "Latents"

"Visibles" are the inputs, \underline{x} 's



$$f = g(\underline{w}^T \underline{h}^{(3)} + b)$$

$$\underline{h}^{(l)} = g^{(l)}(\underbrace{W^{(l)} \underline{h}^{(l-1)} + \underline{b}^{(l)}}_a)$$

↑ some non-linear function.
applied element-wise


Non-linearities?

These are what we used for basis fns

Sigmoid σ : softly partitions the space.

RBFs: are you near some region

ReLU: Rectified linear unit




Hard ReLU

$$= \max(a, 0), \quad \text{Relu}$$

Soft ReLU

$$= \log(1 + e^a)$$

Easier to fit than:



Practical Fitting Tips

Set initial weight matrices $W^{(l)}$ (and the biases)

Must not set to 0 !

Symmetry means all the hidden units behave the same.

→ Randomly set the weights

Naive idea each weight $\sim N(0, 1)$
randn()

If we have k hidden units connecting to next layer, those are combined

$$h_k^{(l)} = g\left(\underbrace{w_{ki}^{(l)} h_i^{(l-1)}} + b_k^{(l)}\right)$$

Typically very roughly

this value scales with \sqrt{k}

Sigmoid:



Example Initialization:

$$w \sim N\left(0, \left(0.01/\sqrt{k}\right)^2\right)$$