



1. Consider the following two functions.



- (a) Design features for each function, to approximate them as a linear function of these features. Can you design features to make the approximation exact?
  - (b) Can you design one set of features, that allows you to represent both functions?
2. Consider the following neural network  $\hat{v}$  with one-hidden layer, relu activation function  $g$ , with weights  $\mathbf{W}^{[0]} \in \mathbb{R}^{n \times d}$ ,  $\mathbf{W}^{[1]} \in \mathbb{R}^{n \times 1}$ ,  $\mathbf{b}^{[0]} \in \mathbb{R}^{n \times 1}$ ,  $\mathbf{b}^{[1]} \in \mathbb{R}$ ,

$$\begin{aligned} \hat{v}(\mathbf{x}; \mathbf{W}^{[0]}, \mathbf{W}^{[1]}, \mathbf{b}^{[0]}, \mathbf{b}^{[1]}) &= \mathbf{W}^{[1]} g(\mathbf{W}^{[0]} \mathbf{x} + \mathbf{b}^{[0]}) + \mathbf{b}^{[1]} \\ &= \mathbf{W}^{[1]} g(\boldsymbol{\psi}) + \mathbf{b}^{[1]} \\ &= \sum_{i=1}^n \mathbf{W}_i^{[1]} g(\psi_i) + \mathbf{b}^{[1]} \end{aligned}$$

for  $\boldsymbol{\psi} = \mathbf{W}^{[0]} \mathbf{x} + \mathbf{b}^{[0]} \in \mathbb{R}^n$ . Recall that we get the following gradients

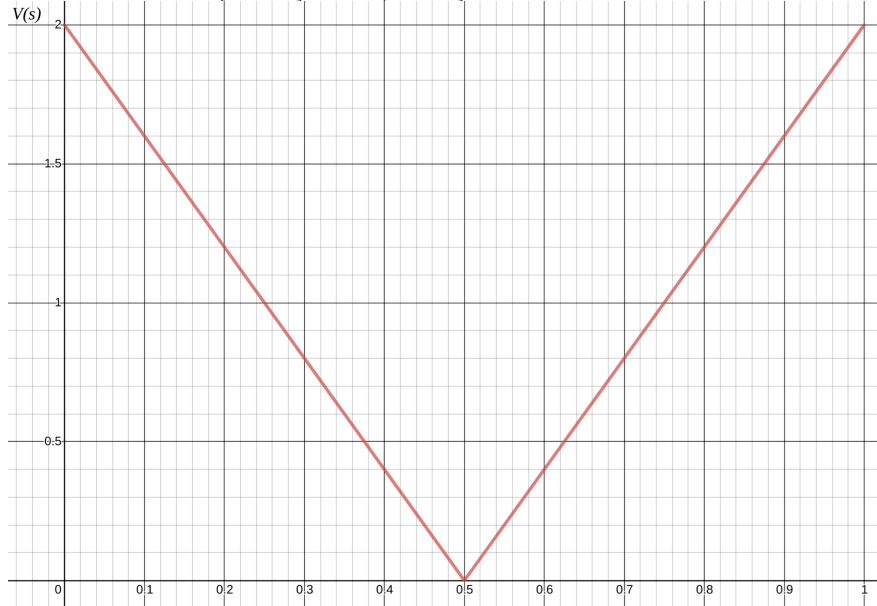
$$\begin{aligned} \frac{\partial \hat{v}}{\partial \mathbf{W}_i^{[1]}} &= g(\psi_i) \\ \frac{\partial \hat{v}}{\partial \mathbf{b}^{[1]}} &= 1 \\ \frac{\partial \hat{v}}{\partial \mathbf{b}_i^{[0]}} &= \mathbf{W}_i^{[1]} \frac{\partial g(\psi_i)}{\partial \mathbf{b}_i^{[0]}} \\ \frac{\partial \hat{v}}{\partial \mathbf{W}_{ij}^{[0]}} &= \sum_k \mathbf{W}_k^{[1]} \frac{\partial g(\psi_k)}{\partial \mathbf{W}_{ij}^{[0]}} \end{aligned}$$

- (a) What are the derivatives specifically for the relu activation  $g$ ?
- (b) We talked about carefully initializing the weights for the NN. For example, each weight can be sampled from a Gaussian distribution. Imagine instead you decided to initialize all the weights to zero. Why would this be a problem? Hint: Consider the derivatives in (a).

3. Consider a problem with the state space,  $\mathcal{S} = \{0, 0.01, 0.02, \dots, 1\}$ . Assume the true value function is

$$v_\pi(s) = 4|s - 0.5|$$

which is visualized below. We decide to create features with state aggregation, and choose to aggregate into two bins:  $[0, 0.5]$  and  $(0.5, 1]$ .



- What are the possible feature vectors for this state aggregation?
  - Imagine you minimize the  $\overline{\text{VE}}(\mathbf{w}) = \sum_{s \in \mathcal{S}} d(s)(v_\pi(s) - \hat{v}(s, \mathbf{w}))^2$  with a uniform weighting  $d(s) = \frac{1}{101}$  for all  $s \in \mathcal{S}$ . What vector  $\mathbf{w}$  is found?
  - Now, if the agent puts all of the weighting on the range  $[0, 0.25]$ , (i.e.  $d(s) = 0$  for all  $s \in (0.25, 1]$ ), then what vector  $\mathbf{w}$  is found by minimizing  $\overline{\text{VE}}$ ?
4. Consider the following general SGD update rule with a general target  $U_t$

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \alpha [U_t - \hat{v}(S_t, \mathbf{w}_t)] \nabla \hat{v}(S_t, \mathbf{w}_t).$$

Assume we are using linear function approximation, *i.e.*  $\hat{v}(S, \mathbf{w}) = x(S)^\top \mathbf{w}$ .

- What happens to the update if we scale the features by a constant and use the new features  $\tilde{x}(S) = 2x(S)$ ? Why might this be a problem?
- In general we want a stepsize that is invariant to the magnitude of the feature vector  $x(S)$ , where the magnitude is measured by the inner product  $x(S)^\top x(S)$ . The book suggests the following stepsize:

$$\alpha = \frac{1}{\tau x(S)^\top x(S)}.$$

What is  $x(S)^\top x(S)$  when using tile coding with 10 tilings? Suppose  $\tau = 1000$ , what is  $\alpha$  if we use tile coding with 10 tilings?