




$$V(s) = w_1 X_1(s) + w_2 X_2(s)$$




$$V(s) = \omega_1 X_1(s) + \omega_2 X_2(s)$$


$$\omega_1 = 2 \quad \omega_2 = 0 :$$





$$\omega_1 = 1 \quad \omega_2 = 1 :$$



$$\omega_1 = 0 \quad \omega_2 = 4 :$$





$$V(s) = w_1 X_1(s) + w_2 X_2(s)$$

Tabular equivalent:


$$X_1(s) = \begin{cases} 1 & \text{at } s_1 \\ 0 & \text{elsewhere} \end{cases}$$

$$\vdots$$


$$X_N(s) = \begin{cases} 1 & \text{at } s_N \\ 0 & \text{elsewhere} \end{cases}$$

Thus, $w_i = V(s_i)$


$$w_1 = 2 \quad w_2 = 0 :$$





$$w_1 = 1 \quad w_2 = 1 :$$



$$w_1 = 0 \quad w_2 = 4 :$$



Linear Regression:





$$Y = W_1 X_1 + W_2 X_2$$

$$= W_1 X_1 + W_2 \cdot 1$$

↑
slope

↑
intercept

Linear Regression:



$$Y = W_1 X_1 + W_2 X_2$$

$$= W_1 X_1 + W_2 \cdot 1$$

↑
slope

↑
intercept

Supervised learning:

Given $\langle x, y \rangle$ pairs,
learn $f(x) \rightarrow y$


Discrete $Y \rightarrow$ classification

e.g. "is this image a cat or a dog?"


Continuous $Y \rightarrow$ regression

e.g. "predict height from weight"

$$E(W) = \sum_{i=0}^N (w^T x_i - y_i)^2$$



$$E(w) = \sum_{i=0}^N (w^T x_i - y_i)^2$$





Batch least squares: Solve for $\underset{w}{\operatorname{argmin}} E(w)$ directly

Exact gradient descent: Compute $\frac{dE}{dw}$ exactly from all x_i

SGD: Approximate $\frac{dE}{dw}$ with a small number of samples of X , often only one.

$$E(w) = \sum_{i=0}^N (w^T x_i - y_i)^2$$



Batch least squares: Solve for $\underset{w}{\operatorname{argmin}} E(w)$ directly

Exact gradient descent: Compute $\frac{dE}{dw}$ exactly from all x_i

SGD: Approximate $\frac{dE}{dw}$ with a small number of samples of X , often only one.

RL + Function Approximation
Update Rule:

$$w \leftarrow w + \alpha [G_t - \hat{v}(s_t, w)] \nabla \hat{v}(s_t, w)$$

RL + Function Approximation

Update Rule:

$$w \leftarrow w + \alpha [G_t - \hat{v}(s_t, w)] \nabla \hat{v}(s_t, w)$$

When approximator is linear in features/bases
i.e. $\hat{v}(s_t, w) = w \cdot x(s)$, then:

$$w \leftarrow w + \alpha [G_t - \hat{v}(s_t, w)] x(s_t)$$

since $\frac{d\hat{v}(s)}{dw_i} = x_i(s)$

RL + Function Approximation

Update Rule:

$$w \leftarrow w + \alpha [G_t - \hat{V}(s_t, w)] \nabla \hat{V}(s_t, w)$$

When approximator is linear in features/bases
i.e. $\hat{V}(s_t, w) = w \cdot x(s)$, then:

$$w \leftarrow w + \alpha [G_t - \hat{V}(s_t, w)] x(s_t)$$

since $\frac{d\hat{V}(s)}{dw} = x_i(s)$



Assigns "credit" to
most activated features
for success + failure

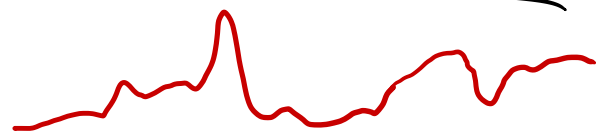
Fourier Basis

w_1 

+ w_2 

+ w_3 


⋮ ⋮

= 

Params :

- order
- interaction terms?


Fourier Basis



Params:

- order
- interaction terms?


Radial Basis Functions



Params:

- Kernel width
 - Tiling density
- (both variable per dimension)


Fourier Basis



Params:

- order
- interaction terms?


Radial Basis Functions



Params:

- Kernel width
 - Tiling density
- (both variable per dimension)

Hand-Designed Features



"Distance to nearest food pellet"

"Number of ghosts within radius of 3"

"1 if exactly this state, 0 otherwise"



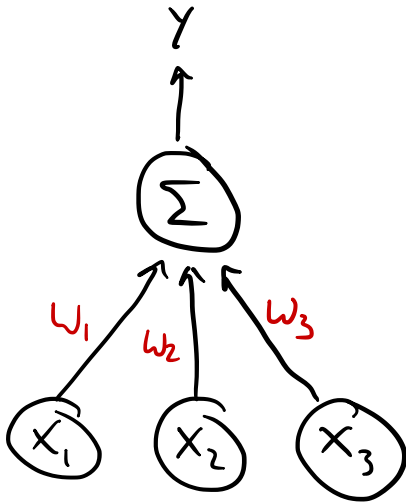
How to choose?

Representative Power?

Locality?



Perceptron




$$Y = x_1 w_1 + x_2 w_2 + x_3 w_3$$

$$\frac{dy}{dw_1} = x_1$$

Y is Linear in X

Perceptron




$$Y = X_1 W_1 + X_2 W_2 + X_3 W_3$$

$$\frac{dy}{dw_1} = X_1$$

Y is Linear in X

Neural Network




$$\frac{dy}{dw_{1,1}} = ?$$

⇒ Differentiation Via Backpropagation

Y is Nonlinear in X!

Nonparametric Methods


K-nearest neighbor



$$\hat{v}(s_q) = \frac{1}{K} \sum_{i=1}^K G(s_i)$$

Nonparametric Methods

K-nearest neighbor



$$\hat{V}(s_q) = \frac{1}{K} \sum_{i=1}^K G(s_i)$$

Kernel Methods

Similarity "kernel" $K(s, s')$


For data set of N states

And query state s_q

$$\hat{V}(s_q) = \sum_{i=1}^N K(s_q, s_i) G(s_i)$$

Nonparametric Methods

K-nearest neighbor



$$\hat{V}(s_q) = \frac{1}{K} \sum_{i=1}^K G(s_i)$$

KNN special case kernel where:
 $K(s, s') = \frac{1}{K}$ if s' is a KNN of s
0 otherwise

Kernel Methods

Similarity "kernel" $K(s, s')$

For data set of N states

And query state s_q

$$\hat{V}(s_q) = \sum_{i=1}^N K(s_q, s_i) G(s_i)$$

"kernel trick" equivalent to
computing \hat{V} in high-dim space
of features


→ Complexity of kernel methods
grows with number of data
points, not features

$$\omega \leftarrow \omega - \frac{1}{2} \alpha \nabla \left[V_{\pi}(s_t) - \hat{v}(s_t, \omega) \right]^2$$

$$\omega \leftarrow \omega + \alpha \left[\mathbf{U}_t - \hat{v}(s_t, \omega) \right] \cdot \nabla \hat{v}(s_t, \omega)$$

$$\omega \leftarrow \omega - \frac{1}{2} \alpha \nabla \left[V_{\pi}(s_t) - \hat{V}(s_t, \omega) \right]^2$$

$$\omega \leftarrow \omega + \alpha \left[U_t - \hat{V}(s_t, \omega) \right] \cdot \nabla \hat{V}(s_t, \omega)$$



MC Estimate:

$$U_t = \sum_{i=t}^T \gamma^{i-t-1} R_i$$


- Unbiased

- fixed when ω changes

converges near global opt

$$w \leftarrow w - \frac{1}{2} \alpha \nabla \left[V_{\pi}(s_t) - \hat{V}(s_t, w) \right]^2$$

$$w \leftarrow w + \alpha \left[U_t - \hat{V}(s_t, w) \right] \cdot \nabla \hat{V}(s_t, w)$$



MC Estimate:

$$U_t = \sum_{i=t}^T \gamma^{i-t-1} R_i$$

- Unbiased
- fixed when w changes
- converges near global opt

Bootstrap Estimate:

$$U_t = R_t + \gamma \hat{V}(s_{t+1}, w)$$

- Biased
- Function of $w \rightarrow$ changes with w !
- Gradient calc treated U_t as a constant!
- converges to TD fixed point (local opt)

For TD fixed point w_{TD} :

$$\overline{VE}(w_{TD}) \leq \frac{1}{1-\gamma} \min_w \left[\overline{VE}(w) \right]$$

