

Classifiers

Classifier : function ^{positions} attributes ^{action (or meta-actions)} → class

Iris flower data set

MNIST database

<http://yann.lecun.com/exdb/mnist/>

Learning : reinforcement — reward observable

supervised — examples available

unsupervised —
(clustering)



Methods : k-nearest neighbors




2 of 3 nearest neighbors orange, so predict orange

decision trees

nested ifs



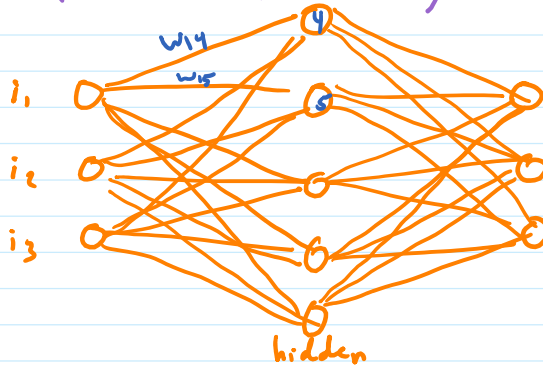
perceptron



$$\text{output} = f\left(\sum_{j=1}^k i_j \cdot w_j\right)$$

$$f(x) = \begin{cases} 1 & \text{if } x \geq b \\ 0 & \text{otherwise} \end{cases}$$

multi-layer perceptron
(artificial neural network)



wright for each connection
activation fn for each node



Deep Q network learning to play Pong

Roadmap

Understanding Alpha Go / Alpha Go Zero

✓ MCTS

✓ reinforcement learning

✓ supervised learning

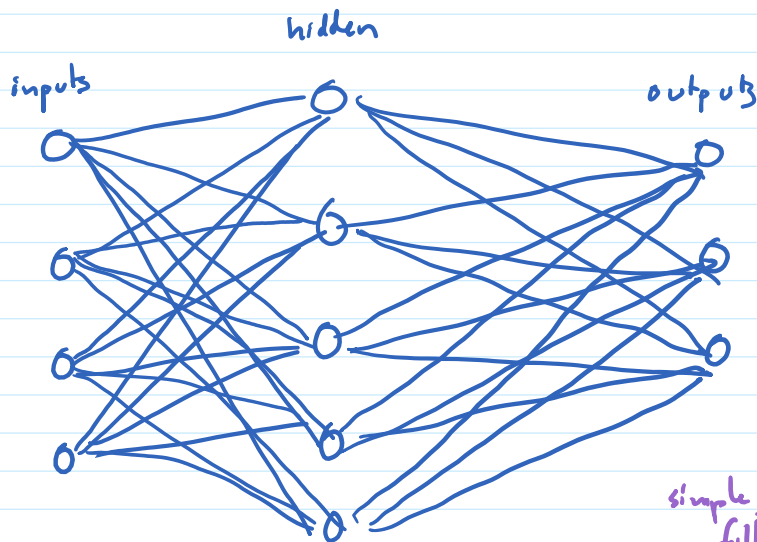

neural networks

convolutional neural networks

1-10	30	90 sec	short pass
1-10	28	87 sec	short pass
1-10	25	85 sec	short pass


want to generalize
from training examples to
new, future input

Artificial Neural Network



simple, common topology: fully connected feedforward

supervised learning: (or reinforcement)
 split examples into
 training
 20% validation set
 test




initialize weights randomly

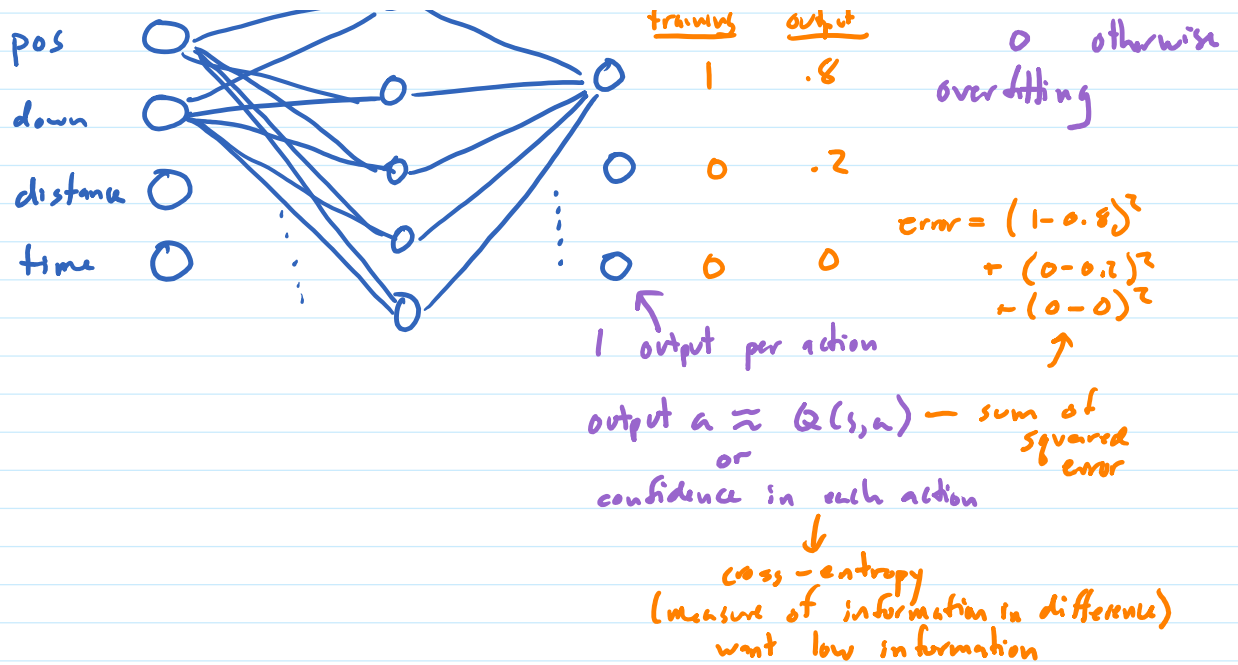
observe rewards following policy defined by NN
 compute output for training examples
 compute distance between network's output and correct output
 adjust weights to decrease that distance
 backpropagation
 gradient descent
 until sufficiently trained
 out of time no improvement on validation set

a differentiable fun to measure error

$$f(x) = \begin{cases} 4 & \text{if } x=2 \\ 5.5 & \text{if } x=4 \\ 7.5 & \text{if } x=5 \\ 9 & \text{if } x=7 \\ 0 & \text{otherwise} \end{cases}$$



overfitting



Input/Output - numeric!

categorical

enumeration?
versicolor 0
virginica 1
setosa 2

75% 0.05? 0% 50%
5% 5% 1.95 58%
0% 95%

one input/output per category
1 = belongs to category
0 = doesn't belong

	versicolor	virginica	setosa		
0	1	0	0	0.0	0.5
0	0	1	0	0.05	0.0
0	0	0	1	0.95	0.5

one-hot

date

seconds since epoch (normalised)

Jan 1 1900 = 0.0
Jan 1 1999 = 0.999

month/day year

Jan 1 1900 = 0.0 0.0
Jan 1 1999 = 0.6 1.0
Dec 31 1950 = 1.0 0.5

fuzzy

Jan

0 0 0 0 0 0 ... Dec 0
Jan 1 .51 .49
Mar 15 0.05 0.9 0.05
Dec 31 .49 .51
Feb 7 .28 .7 .02
Nov 27 .6 .4

Yahtzee State = categories used / unused + upper total + roll + rolls left


12 inputs 0/1 for used/unused

???

Yahtzee (0/50/unused)

upper total
63

one-hot?
rolls left?
2



same roll very different input

no generalization

1 2 4 5
5 4 1 2 1
6 6 6 6 6

(want to ensure that rolls you expect similar results for have similar inputs to facilitate generalization)