

# Learning

- Learning agents
- Inductive learning
  - Different Learning Scenarios
  - Evaluation
- Neural Networks
  - Perceptrons
  - Multilayer Perceptrons
- Reinforcement Learning
  - Temporal Differences
  - Q-Learning
  - SARSA

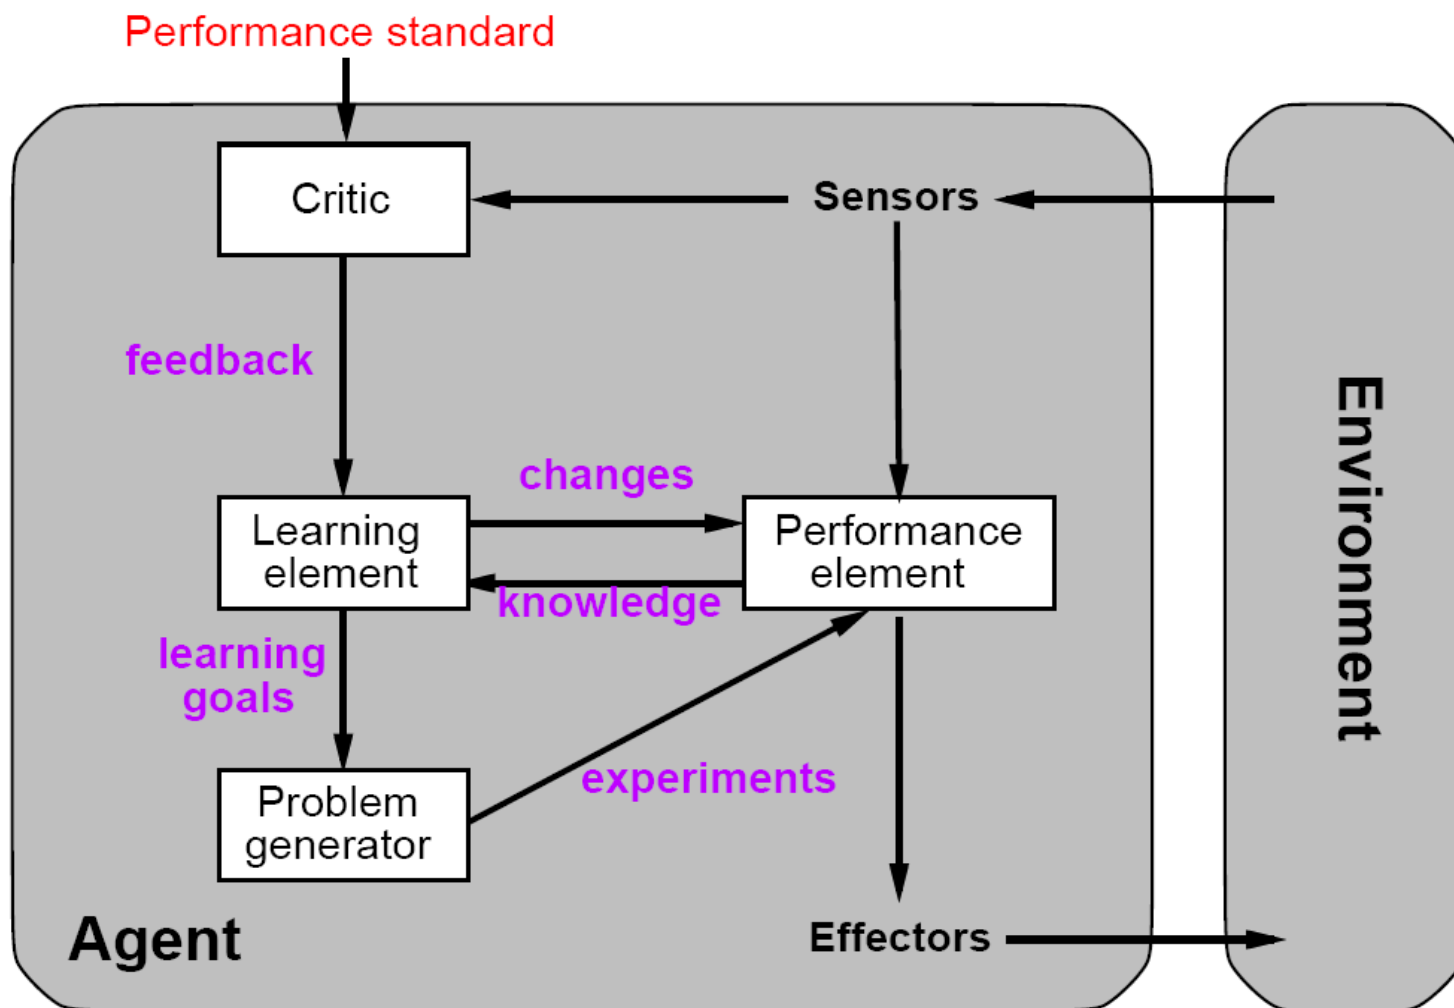
Material from  
Russell & Norvig,  
chapters 18.1,  
18.2, 20.5 and 21

Slides based on Slides  
by Russell/Norvig,  
and Torsten Reil

# Learning

- Learning is essential for **unknown environments**,
  - i.e., when designer lacks omniscience
- Learning is useful as a **system construction method**,
  - i.e., expose the agent to reality rather than trying to write it down
- Learning modifies the agent's decision mechanisms to **improve performance**

# Learning Agents



# Learning Element

- Design of a learning element is affected by
  - Which components of the performance element are to be learned
  - What feedback is available to learn these components
  - What representation is used for the components
- Type of feedback:
  - **Supervised learning**: correct answers for each example
  - **Unsupervised learning**: correct answers not given
  - **Reinforcement learning**: occasional rewards

# Different Learning Scenarios

## Supervised Learning

- A teacher provides the value for the target function for all training examples (labeled examples)
- concept learning, classification, regression

## Reinforcement Learning

- The teacher only provides feedback but not example values

## Semi-supervised Learning

- Only a subset of the training examples are labeled

## Unsupervised Learning

- There is no information except the training examples
- clustering, subgroup discovery, association rule discovery


# Inductive Learning

Simplest form: learn a function from examples

- $f$  is the target function
- An example is a pair  $(x, f(x))$
- Problem: find a hypothesis  $h$  such that  $h \approx f$  given a training set of examples
- This is a highly simplified model of real learning:
  - Ignores prior knowledge
  - Assumes examples are given

# Inductive Learning Method

- Construct/adjust  $h$  to agree with  $f$  on training set ( $h$  is **consistent** if it agrees with  $f$  on all examples)
- Example:
  - curve fitting

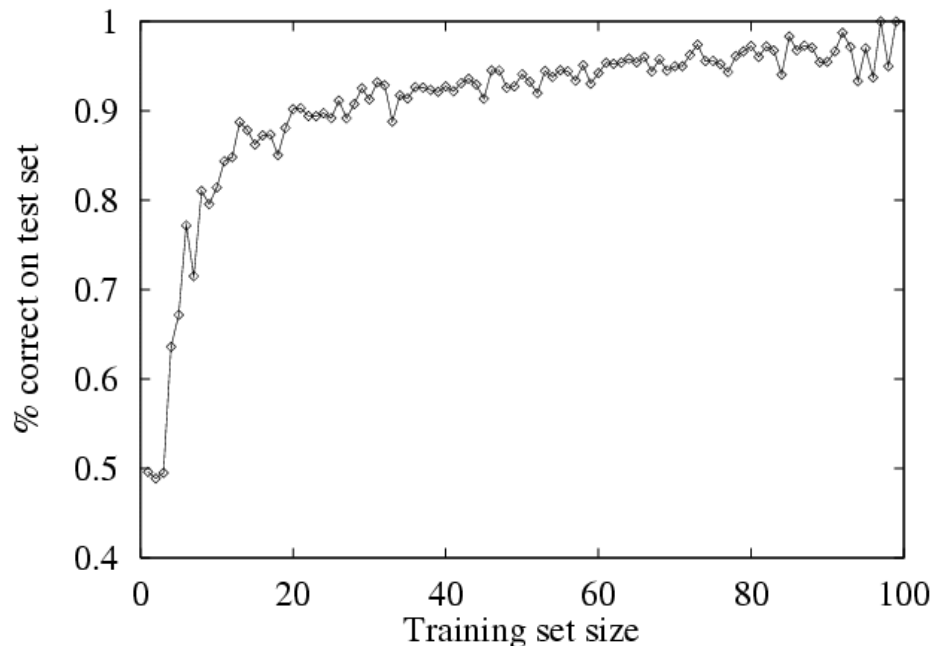


- Ockham's Razor**
  - The best explanation is the simplest explanation that fits the data
- Overfitting Avoidance**
  - maximize a combination of consistency and simplicity

# Performance Measurement

- How do we know that  $h \approx f$ ?
  - Use theorems of computational/statistical learning theory
  - Try  $h$  on a new **test set** of examples where  $f$  is known (use **same distribution** over example space as training set)

**Learning curve** = % correct on test set over training set size







# What are Neural Networks?

- Models of the brain and nervous system
- Highly parallel
  - Process information much more like the brain than a serial computer
- Learning
  
- Very simple principles
- Very complex behaviours
  
- Applications
  - As powerful problem solvers
  - As biological models

# Pigeons as Art Experts

Famous experiment (Watanabe *et al.* 1995, 2001)


- Pigeon in Skinner box
- Present paintings of two different artists (e.g. Chagall / Van Gogh)
- Reward for pecking when presented a particular artist



# Results

- Pigeons were able to discriminate between Van Gogh and Chagall with 95% accuracy (when presented with pictures they had been trained on)
- Discrimination still 85% successful for previously unseen paintings of the artists
- Pigeons do not simply memorise the pictures
- They can extract and recognise patterns (the 'style')
- They generalise from the already seen to make predictions
- This is what neural networks (biological and artificial) are good at (unlike conventional computer)


# A Biological Neuron



- Neurons are connected to each other via synapses
- If a neuron is activated, it spreads its activation to all connected neurons

# An Artificial Neuron

(McCulloch-Pitts, 1943)



- Neurons correspond to nodes or **units**
- A **link** from unit  $i$  to unit  $j$  propagates activation  $a_j$  from  $j$  to  $i$
- The **weight**  $W_{i,j}$  of the link determines the strength and sign of the connection
- The total **input activation** is the sum of the input activations
- The **output activation** is determined by the activation function  $g$

# Perceptron


(Rosenblatt 1957, 1960)

- A single node
  - connecting  $n$  input signals with one output signal
  
- Activation function
  - A simple threshold function:  $a_i = \begin{cases} -1 & \text{if } \sum_{j=0}^n W_{ij} \cdot a_j \leq 0 \\ 1 & \text{if } \sum_{j=0}^n W_{ij} \cdot a_j > 0 \end{cases}$
  
- Thus it implements a **linear separator**
  - i.e., a hyperplane that divides  $n$ -dimensional space into a region with output 0 and a region with output 1

# Perceptrons and Boolean Functions


- a Perceptron can implement all elementary logical functions

(McCulloch & Pitts, 1943)



- more complex functions like XOR cannot be modeled

(Minsky & Papert, 1969)



# Perceptron Learning


- Perceptron Learning Rule for Supervised Learning

$$W_i \leftarrow W_i + \alpha \cdot (f(\mathbf{x}) - h(\mathbf{x})) \cdot x_i$$

learning rate

error

- Example:



## Computation of output signal $h(x)$

$$\text{in}(x) = -1 \cdot 0.2 + 1 \cdot 0.5 + 1 \cdot 0.8 = 1.1$$

$$h(x) = 1 \text{ because } \text{in}(x) > 0$$

## Assume target value $f(x) = -1$ (and $\alpha = 0.5$ )

$$W_0 \leftarrow 0.2 + 0.5 \cdot (-1 - 1) \cdot -1 = 0.2 + 1 = 1.2$$

$$W_1 \leftarrow 0.5 + 0.5 \cdot (-1 - 1) \cdot 1 = 0.5 - 1 = -0.5$$

$$W_2 \leftarrow 0.8 + 0.5 \cdot (-1 - 1) \cdot 1 = 0.8 - 1 = -0.2$$



# Error Minimization

- It is easy to derive a perceptron training algorithm that minimizes the squared error

$$E = \frac{1}{2} Err^2 = \frac{1}{2} (f(\mathbf{x}) - h(\mathbf{x}))^2 = \frac{1}{2} \left( f(\mathbf{x}) - g\left(\sum_{i=0}^n W_j \cdot x_j\right) \right)^2$$



- Change weights into the direction of the steepest descent of the error function

$$\frac{\partial E}{\partial W_j} = Err \cdot \frac{\partial Err}{\partial W_j} = Err \cdot \frac{\partial}{\partial W_j} \left( f(\mathbf{x}) - g\left(\sum_{i=0}^n W_j \cdot x_j\right) \right) = -Err \cdot g'(\text{in}) \cdot x_j$$

- To compute this, we need a continuous and differentiable activation function  $g$ !
- Weight update with learning rate  $\alpha$ :  $W_j = W_j + \alpha \cdot Err \cdot g'(\text{in}) \cdot x_j$ 
  - positive error  $\rightarrow$  increase network output
    - increase weights of nodes with positive input
    - decrease weights of nodes with negative input

# Sigmoid Activation Function

- A commonly used activation function is the sigmoid function
  - similar to the threshold function
  - easy to differentiate
  - non-linear




$$g(x) = \frac{1}{1 + e^{-x}}$$

$$g'(x) = g(x)(1 - g(x))$$


# Multilayer Perceptrons

- Perceptrons may have multiple output nodes
  - may be viewed as multiple parallel perceptrons
- The output nodes may be combined with another perceptron
  - which may also have multiple output nodes
- The size of this **hidden layer** is determined manually



# Multilayer Perceptrons

- Information flow is unidirectional
  - Data is presented to *Input layer*
  - Passed on to *Hidden Layer*
  - Passed on to *Output layer*
- Information is distributed
- Information processing is parallel



# Expressiveness of MLPs

- Every continuous function can be modeled with two layers
- Every function with three layers

# Backpropagation Learning

- The output nodes are trained like a normal perceptron

$$W_{ji} = W_{ji} + \alpha \cdot Err_i \cdot g'(\text{in}_i) \cdot x_j = W_{ji} + \alpha \cdot \Delta_i \cdot x_j$$


- $\Delta_i$  is the error of output node  $i$  times the derivation of its inputs
- the error of the output layers back to the hidden layer

$$\Delta_j = \left( \sum_i W_{ji} \cdot \Delta_i \right) \cdot g'(\text{in}_j) \qquad W_{kj} = W_{kj} + \alpha \cdot \Delta_j \cdot x_k$$

- the training signal of hidden layer node  $j$  is the weighted sum of the errors of the output nodes

# Overfitting

- **Training Set Error** continues to decrease with increasing number of training examples / number of epochs
  - an epoch is a complete pass through all training examples
- **Test Set Error** will start to increase because of overfitting



- Simple training protocol:
  - keep a separate **validation set** to watch the performance
    - validation set is different from training and test sets!
  - stop training if error on validation set gets down