## LEARNING FROM ObSERVATIONS

Chapter 18, SEctions 1-3

Outline
$\diamond$ Learning agents
$\diamond$ Inductive learning
$\diamond$ Decision tree learning
$\diamond$ Measuring learning performance

## 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 element

Design of learning element is dictated by
$\diamond$ what type of performance element is used
$\diamond$ which functional component is to be learned
$\diamond$ how that functional compoent is represented
$\diamond$ what kind of feedback is available
Example scenarios:

| Performance element | Component | Representation | Feedback |
| :--- | :--- | :--- | :--- |
| Alpha-beta search | Eval. fn. | Weighted linear function | Win/loss |
| Logical agent | Transition model | Successor-state axioms | Outcome |
| Utility-based agent | Transition model | Dynamic Bayes net | Outcome |
| Simple reflex agent | Percept-action fn | Neural net | Correct action |

Supervised learning: correct answers for each instance Reinforcement learning: occasional rewards

## Inductive learning (a.k.a. Science)

Simplest form: learn a function from examples (tabula rasa)
$f$ is the target function

An example is a pair $x, f(x)$, e.g., $\frac{$| $O$ | $O$ | $X$ |
| :--- | :--- | :--- |
| $X$ | $X$ |  |
| $X$ |  |  |,$+1+1+1}{}$

Problem: find $\mathrm{a}(\mathrm{n})$ 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 a deterministic, observable "environment"
- Assumes examples are given
- Assumes that the agent wants to learn $f$-why?)
Inductive learning method

Construct/adjust $h$ to agree with $f$ on training set ( $h$ is consistent if it agrees with $f$ on all examples)
E.g., curve fitting:


## Inductive learning method

Construct/adjust $h$ to agree with $f$ on training set
( $h$ is consistent if it agrees with $f$ on all examples)
E.g., curve fitting


Inductive learning method
Construct/adjust $h$ to agree with $f$ on training set
( $h$ is consistent if it agrees with $f$ on all examples)
E.g., curve fitting:


Inductive learning method
Construct/adjust $h$ to agree with $f$ on training set
( $h$ is consistent if it agrees with $f$ on all examples)
E.g., curve fitting:


Inductive learning method
Construct/adjust $h$ to agree with $f$ on training set ( $h$ is consistent if it agrees with $f$ on all examples)
E.g., curve fitting:


Inductive learning method
Construct/adjust $h$ to agree with $f$ on training set ( $h$ is consistent if it agrees with $f$ on all examples)
E.g., curve fitting


Ockham's razor: maximize a combination of consistency and simplicity

## Attribute-based representations

## Decision trees

One possible representation for hypotheses
E.g., here is the "true" tree for deciding whether to wait:


Classification of examples is positive ( T ) or negative ( F )

Expressiveness
Decision trees can express any function of the input attributes.
E.g., for Boolean functions, truth table row $\rightarrow$ path to leaf:


Trivially, there is a consistent decision tree for any training set $\mathrm{w} /$ one path to leaf for each example (unless $f$ nondeterministic in $x$ ) but it probably won't generalize to new examples

Prefer to find more compact decision trees

## Hypothesis spaces

How many distinct decision trees with $n$ Boolean attributes??
How many distinct decision trees with $n$ Boolean attributes??
$=$ number of Boolean functions
$=$ number of distinct truth tables with $2^{n}$ rows

## Hypothesis spaces

How many distinct decision trees with $n$ Boolean attributes??
$=$ number of Boolean functions
$=$ number of distinct truth tables with $2^{n}$ rows $=2^{2^{n}}$

## Hypothesis spaces

How many distinct decision trees with $n$ Boolean attributes??
$=$ number of Boolean functions
$=$ number of distinct truth tables with $2^{n}$ rows $=2^{2^{n}}$
E.g., with 6 Boolean attributes, there are $18,446,744,073,709,551,616$ trees

## Hypothesis spaces

How many distinct decision trees with $n$ Boolean attributes??
$=$ number of Boolean functions
$=$ number of distinct truth tables with $2^{n}$ rows $=2^{2^{n}}$
E.g., with 6 Boolean attributes, there are $18,446,744,073,709,551,616$ trees How many purely conjunctive hypotheses (e.g., Hungry $\wedge \neg$ Rain)??

## Hypothesis spaces

How many distinct decision trees with $n$ Boolean attributes??
$=$ number of Boolean functions
$=$ number of distinct truth tables with $2^{n}$ rows $=2^{2^{n}}$
E.g., with 6 Boolean attributes, there are $18,446,744,073,709,551,616$ trees

How many purely conjunctive hypotheses (e.g., Hungry $\wedge \neg$ Rain)??
Each attribute can be in (positive), in (negative), or out
$\Rightarrow 3^{n}$ distinct conjunctive hypotheses
More expressive hypothesis space

- increases chance that target function can be expressed
- increases number of hypotheses consistent $\mathrm{w} /$ training set $\Rightarrow$ may get worse predictions


## Decision tree learning

Aim: find a small tree consistent with the training examples
Idea: (recursively) choose "most significant" attribute as root of (sub)tree

```
function DTL(examples, attributes, default) returns a decision tree
    if examples is empty then return default
    else if all examples have the same classification then return the classification
    else if attributes is empty then return MODE(examples)
    else
        best \leftarrowCHoose-Attribute(attributes, examples)
        tree \leftarrowa new decision tree with root test best
        for each value vi of best do
            examples}\mp@subsup{\mp@code{S}}{i}{}\leftarrow{\mathrm{ elements of examples with best = vi}
            subtree \leftarrow DTL(examples}\mp@subsup{s}{i}{}\mathrm{ , attributes - best, MODE(examples))
            add a branch to tree with label vi and subtree subtree
        return tree
```


## Choosing an attribute

Idea: a good attribute splits the examples into subsets that are (ideally) "all positive" or "all negative"


Patrons? is a better choice-gives information about the classification

## Information

Information answers questions
The more clueless I am about the answer initially, the more information is contained in the answer

Scale: 1 bit = answer to Boolean question with prior $\langle 0.5,0.5\rangle$
Information in an answer when prior is $\left\langle P_{1}, \ldots, P_{n}\right\rangle$ is

$$
H\left(\left\langle P_{1}, \ldots, P_{n}\right\rangle\right)=\sum_{i=1}^{n}-P_{i} \log _{2} P_{i}
$$

(also called entropy of the prior)

## Information contd.

Suppose we have $p$ positive and $n$ negative examples at the root
$\Rightarrow H(\langle p /(p+n), n /(p+n)\rangle)$ bits needed to classify a new example
E.g., for 12 restaurant examples, $p=n=6$ so we need 1 bit

An attribute splits the examples $E$ into subsets $E_{i}$, each of which (we hope) needs less information to complete the classification

Let $E_{i}$ have $p_{i}$ positive and $n_{i}$ negative examples
$\Rightarrow H\left(\left\langle p_{i} /\left(p_{i}+n_{i}\right), n_{i} /\left(p_{i}+n_{i}\right)\right\rangle\right)$ bits needed to classify a new example
$\Rightarrow$ expected number of bits per example over all branches is
$\sum_{i} \frac{p_{i}+n_{i}}{p+n} H\left(\left\langle p_{i} /\left(p_{i}+n_{i}\right), n_{i} /\left(p_{i}+n_{i}\right)\right\rangle\right)$
For Patrons?, this is 0.459 bits, for Type this is (still) 1 bit
$\Rightarrow$ choose the attribute that minimizes the remaining information needed

Example contd.
Decision tree learned from the 12 examples:


Substantially simpler than "true" tree-a more complex hypothesis isn't justified by small amount of data

## Performance measurement

How do we know that $h \approx f$ ? (Hume's Problem of Induction)

1) Use theorems of computational/statistical learning theory
2) Try $h$ on a new test set of examples
(use same distribution over example space as training set)
Learning curve $=\%$ correct on test set as a function of training set size


## Summary

Learning needed for unknown environments, lazy designers
Learning agent $=$ performance element + learning element
Learning method depends on type of performance element, available feedback, type of component to be improved, and its representation

For supervised learning, the aim is to find a simple hypothesis that is approximately consistent with training examples

Decision tree learning using information gain
Learning performance $=$ prediction accuracy measured on test set

