## LOGICAL AND BAYESIAN LEARNING

## Outline

- A Logical Formulation of Learning
- Bayesian Learning

Based on the textbook by Stuart Russell \& Peter Norvig:
Artificial Intelligence, A Modern Approach (2nd Edition)
Sections 19.1 and 20.1
© 2006 TKK / Laboratory for Theoretical Computer Science

T-79.5102 / Autumn 2006

## 1. A LOGICAL FORMULATION OF LEARNING

- Inductive learning was previously defined as a process of searching for a hypothesis that agrees with the observed examples.
- For now we concentrate on the case where hypotheses, examples, classifications are represented in terms of logical sentences.
- This form of learning is more general and complex compared to learning decision trees or lists.This approach allows for incremental construction of hypotheses one sentence at a time

The full power of logical inference can be utilised in learning

## Examples and Hypotheses

In the logical representation, attributes become unary predicates.

- The $i^{\text {th }}$ example is generically denoted by $X_{i}$.

Example. The first example in the restaurant domain is described by the following sentence:

$$
\text { Alternate }\left(X_{1}\right) \wedge \neg \operatorname{Bar}\left(X_{1}\right) \wedge \neg F r i / \operatorname{Sat}\left(X_{1}\right) \wedge \operatorname{Hungry}\left(X_{1}\right) \wedge \ldots
$$

The classification of the object is given by $\operatorname{WillWait}\left(X_{1}\right)$.
The generic notations $Q\left(X_{i}\right)$ and $\neg Q\left(X_{i}\right)$ are used for positive and negative examples, respectively

The complete training set corresponds to the conjunction of the respective description and classifications sentences.
© 2006 TKK / Laboratory for Theoretical Computer Science

T-79.5102 / Autumn 2006 Logical and Bayesian Learning

## Candidate Definitions

The aim is to find an equivalent logical expression for the goal predicate $Q$ that can be used to classify examples correctly

- Each hypothesis $H_{i}$ proposes a candidate definition $C_{i}(x)$ for the goal predicate $Q$, i.e. $H_{i}$ takes the form $\forall x\left(Q(x) \leftrightarrow C_{i}(x)\right)$.
> The extension of a hypothesis $H_{i}=\forall x\left(Q(x) \leftrightarrow C_{i}(x)\right)$ is the set of examples $X$ for which $Q(X)$ evaluates to true.

Example. For the decision tree learned in the restaurant example:
$H_{1}=\forall r($ WillWait $(r) \leftrightarrow \operatorname{Patrons}(r$, Some $) \vee$

$$
\begin{aligned}
& (\text { Patrons }(r, \text { Full }) \wedge \neg \operatorname{Hungry}(r) \wedge \text { Type }(r, \text { French })) \vee \\
& (\text { Patrons }(r, \text { Full }) \wedge \neg \operatorname{Hungry}(r) \wedge \text { Type }(r, \text { Thai }) \wedge \text { Fri } / \text { Sat }(r)) \vee \\
& (\text { Patrons }(r, \text { Full }) \wedge \neg \operatorname{Hungry}(r) \wedge \text { Type }(r, \text { Burger })))
\end{aligned}
$$

## Hypothesis Space

- Logically equivalent hypotheses have equal extensions.
- Two hypotheses with different extensions are logically inconsistent with each other, as they differ on at least one example $X_{i}$.
Example. The conjunction of $H_{2}=\forall r($ WillWait $(r) \leftrightarrow$ Hungry $(r))$ and $H_{3}=\forall r($ WillWait $(r) \leftrightarrow \neg \operatorname{Hungry}(r))$ implies a contradiction.
- The hypothesis space $\left\{H_{1}, H_{2}, \ldots, H_{n}\right\}$ is denoted by $\mathbf{H}$.
- It is usually believed that one of the hypotheses in $\mathbf{H}$ is correct,
i.e. the disjunction $H_{1} \vee H_{2} \vee \ldots \vee H_{n}$ evaluates to true.

Example. In decision tree learning, the hypothesis space consists of all decision trees that can be defined in terms of the attributes provided.
© 2006 TKK / Laboratory for Theoretical Computer Science
function CuRRENT-BEST-LEARNING(examples) returns a hypothesis
function CuRRENT-BEST-LEARNING(examples) returns a hypothesis
H\leftarrow
H\leftarrow
for each remaining example in examples do
for each remaining example in examples do
C
C
consistent with examples
consistent with examples
H\leftarrow
H\leftarrow
If no consistent specialization/generalization can be found then fail
If no consistent specialization/generalization can be found then fail
end
end
return }
return }

- Generalisations and specialisations imply logical relationships:
E.g., if $H_{1}=\forall x\left(Q(x) \leftrightarrow C_{1}(x)\right)$ is a generalisation of $H_{2}=\forall x\left(Q(x) \leftrightarrow C_{2}(x)\right)$, then $\forall x\left(C_{2}(x) \rightarrow C_{1}(x)\right)$ holds.

Note that $H_{2}$ is a specialisation of $H_{1}$ in the setting above

## Examples

Example. Recall the training set used in the restaurant domain.

| Example |  |  |  |  |  |  |  |  |  | Goal |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Alt | Bar | Fri | Hun | Pat | Price | Rain | Res | Type | Est | WillWait |
|  | $X_{1}$ | Yes | No | No | Yes | Some | $\$ \$ \$$ | No | Yes | French | $0-10$ |
| $X_{2}$ | Yes | No | No | Yes | Full | $\$$ | Yo | No | Thai | $30-60$ | No |
| $X_{3}$ | No | Yes | No | No | Some | $\$$ | No | No | Burger | $0-10$ | Yes |
| $X_{4}$ | Yes | No | Yes | Yes | Full | $\$$ | No | No | Thai | $10-30$ | Yes |
| $X_{5}$ | Yes | No | Yes | No | Full | $\$ \$ \$$ | No | Yes | French | $>60$ | No |
| $X_{6}$ | No | Yes | No | Yes | Some | $\$ \$$ | Yes | Yes | Italian | $0-10$ | Yes |
| $X_{7}$ | No | Yes | No | No | None | $\$$ | Yes | No | Burger | $0-10$ | No |
| $X_{8}$ | No | No | No | Yes | Some | $\$ \$$ | Yes | Yes | Thai | $0-10$ | Yes |
| $X_{9}$ | No | Yes | Yes | No | Full | $\$$ | Yes | No | Burger | $>60$ | No |
| $X_{10}$ | Yes | Yes | Yes | Yes | Full | $\$ \$ \$$ | No | Yes | Italian | $10-30$ | No |
| $X_{11}$ | No | No | No | No | None | $\$$ | No | No | Thai | $0-10$ | No |
| $X_{12}$ | Yes | Yes | Yes | Yes | Full | $\$$ | No | No | Burger | $30-60$ | Yes |

## Discussion

- The Current-Best-Learning algorithm is non-deterministic: there may be several possible specialisations or generalisations that can be applied at any point.
> The choices made might not lead to the simplest hypothesis.
- If a dead-end (unrecoverable inconsistency) is encountered, the algorithm must backtrack to a previous choice point.
- Checking the consistency of all the previous examples over again for each choice is very expensive.
© 2006 TKK / Laboratory for Theoretical Computer Science

T-79.5102 / Autumn 2006 Logical and Bayesian Learning

## Least-Commitment Search

The original hypothesis space can be seen as a huge disjunction

$$
H_{1} \vee H_{2} \vee \ldots \vee H_{n}
$$

Hypotheses which are consistent with all examples encountered so far form a set of hypotheses called the version space $V$.

Version space is shrunk by the candidate elimination algorithm:

```
function VERSION-SPACE-LEARNING(examples) returns a version spac
    luction VERSION-SPACE-LEARNING(examples) returns a version
    V}\leftarrow\mathrm{ the set of all hypotheses
    for each example e in examples do
        if V is not empty then V\leftarrowVERSIon-SPace-Update(V,e)
    end
function VERSION-SPACE-UPDATE(V,e) returns an updated version space
    V\leftarrow{h\inV:h\mathrm{ is consistent with e}}
```


## Boundary Sets

> The algorithm finds a subset of the version space $V$ that is consistent with all examples in an incremental way.

- Candidate elimination is an example of a least-commitment algorithm, as no arbitrary choices are made among hypotheses.
- Since the hypothesis space $V$ is possibly enormous, it cannot be represented directly as a set of hypotheses or a disjunction.
- The problem can be alleviated by boundary sets $\left\{S_{1}, \ldots, S_{n}\right\}$ (S-set) and $\left\{G_{1}, \ldots, G_{m}\right\}$ (G-set) and a partial ordering among hypotheses induced by specialisation/generalisation.
- Any hypothesis $H$ between a most specific boundary $S_{i}$ and a most general boundary $G_{j}$ is consistent with the examples seen.
© 2006 TKK / Laboratory for Theoretical Computer Science

T-79.5102 / Autumn 2006 Logical and Bayesian Learning

Boundary sets for the version space are illustrated below:


Initially, the S-set contains a single hypothesis $\forall x(Q(x) \leftrightarrow$ False $)$ while the G-set contains $\forall x(Q(x) \leftrightarrow$ True) only.

The remaining problem is how to update S-sets and G-sets for a new example (the job of the Version-Space-Update function).

## Updating Version Space

Upon a false negative/positive example, a most specific boundary $S$ is replaced by all its immediate generalisations / deleted.

Upon a false positive/negative example, a most general boundary $G$ is replaced by all its immediate specialisations / deleted.

These operations on S-sets and G-sets are continued until:

1. There is exactly one hypothesis left in the version space.
2. The version space collapses (i.e., the S-set or G-set becomes empty): there are no consistent hypotheses for the training set.
3. We run out of examples with several hypotheses remaining in the version space: a solution is to take the majority vote.
(C) 2006 TKK / Laboratory for Theoretical Computer Science


- If the domain contains noise or insufficient attributes for exact classification, the version space will always collapse.
- If unlimited disjunction is allowed in the hypothesis space, the S-set will always contain a single most-specific hypothesis (disjunction of positive examples seen to date).

Analogously for the G-set and negative examples.

- A solution is to allow only limited forms of disjunction.
- For certain hypothesis spaces, the number of elements in the S-set and G-set may grow exponentially in the number of attributes.


## 2. BAYESIAN LEARNING

The data, i.e. instantiations of some or all random variables describing the domain, serve as evidence.

- Hypotheses are probabilistic theories of how the domain works.
- The aim is to make a prediction concerning an unknown quantity $X$ given some data and hypotheses.
- In Bayesian learning, the probability of each hypothesis is calculated, given the data, and predictions are made on that basis.
- Predictions are made by using all the hypotheses, weighted by their probabilities, rather than by using a single "best" hypothesis.
© 2006 TKK / Laboratory for Theoretical Computer Science

Example. Our favourite Surprise candy comes in two flavours, cherry and lime, but they are wrapped in an indistinguishable way.
The candy is sold in large (indistinguishable) bags containing various mixtures of the two flavours:

1. $100 \%$ cherry
2. $75 \%$ cherry and $25 \%$ lime
3. $50 \%$ cherry and $50 \%$ lime
4. $25 \%$ cherry and $75 \%$ lime
5. $100 \%$ lime

Given a new bag of candy, the random variable $H$ (for hypothesis) denotes the type of the bag, with possible values $h_{1}$ through $h_{5}$.
[-4 The agent needs to infer a probabilistic model of the world.

## Bayesian learning

- Let $\mathbf{D}$ represent all the data with observed value $\mathbf{d}$.
> The probability of each hypothesis $h_{i}$ is obtained by Bayes' rule:

$$
P\left(h_{i} \mid \mathbf{d}\right)=\alpha P\left(\mathbf{d} \mid h_{i}\right) P\left(h_{i}\right)
$$

Assuming that each $h_{i}$ specifies a complete distribution for an unknown quantity $X$, Bayesian learning is characterised by

$$
\mathbf{P}(X \mid \mathbf{d})=\sum_{i} \mathbf{P}\left(X \mid \mathbf{d}, h_{i}\right) P\left(h_{i} \mid \mathbf{d}\right)=\sum_{i} \mathbf{P}\left(X \mid h_{i}\right) P\left(h_{i} \mid \mathbf{d}\right) .
$$

The key quantities are the hypothesis prior $P\left(h_{i}\right)$ and the likelihood of the data under each hypothesis $P\left(\mathbf{d} \mid h_{i}\right)$.

If the observations are independently and identically distributed (i.i.d. for short), then $P\left(\mathbf{d} \mid h_{i}\right)=\prod_{j} P\left(d_{j} \mid h_{i}\right)$.
© 2006 TKK / Laboratory for Theoretical Computer Science

T-79.5102 / Autumn 2006 Logical and Bayesian Learning

Example. For the candy example, the prior distribution over $h_{1}, \ldots, h_{5}$ is given by $\langle 0.1,0.2,0.4,0.2,0.1\rangle$, as advertised by the manufacturer.
If the bag is really an all-lime bag $\left(h_{5}\right)$ and the first 10 candies are consequently all lime, then $P\left(\mathbf{d} \mid h_{3}\right)=0.5^{10}$.
The posterior probabilities of the five hypotheses change as the sequence of 10 lime candies is observed:



## MAP and ML hypotheses

The true hypothesis eventually dominates Bayesian prediction.

- Unfortunately, the hypothesis space is usually very large or infinite which makes the Bayesian approach intractable.
A common approximation is to use maximum a posteriori (MAP) hypothesis $h_{\mathrm{MAP}}$ - a hypothesis $h_{i}$ that maximises $P\left(h_{i} \mid \mathbf{d}\right)$ :

$$
\mathbf{P}(X \mid \mathbf{d}) \approx \mathbf{P}\left(X \mid h_{\mathrm{MAP}}\right)
$$

To determine $h_{\text {MAP }}$, it is sufficient to maximise $P\left(\mathbf{d} \mid h_{i}\right) P\left(h_{i}\right)$.
In some cases, the prior probabilities $P\left(h_{i}\right)$ can be assumed to be uniformly distributed.

Then maximising $P\left(\mathbf{d} \mid h_{i}\right)$ produces a maximum-likelihood (ML) hypothesis $h_{\mathrm{ML}}-$ a special case of $h_{\mathrm{MAP}}$.
© 2006 TKK / Laboratory for Theoretical Computer Science

## Bayesian Network Learning Problems

The learning problem for Bayesian networks comes in several varieties:

1. Known structure, fully observable: only CPTs are learned and the statistics of the set of examples can be used.
2. Unknown structure, fully observable: this involves heuristic search through the space of structures — guided by the ability of modelling data correctly (MAP or ML probability value).
3. Known structure, hidden variables: analogy to neural networks.
4. Unknown structure, hidden variables: no good/general algorithms are known for learning in this setting.

## SUMMARY

- Learning is essential for dealing with unknown environments.
- Prior knowledge helps learning by eliminating otherwise consistent hypotheses and by "filling in" the explanation of examples, thereby allowing for shorter hypotheses.

Bayesian learning methods formulate learning as a form of probabilistic inference: observations are used to update a prior distribution over hypotheses.

- This approach implements Ockham's razor principle but quickly becomes intractable for complex hypothesis spaces.

Maximum a posteriori (MAP) and maximum likelihood (ML) learning are more tractable approximations of Bayesian learning.

