## Rule Induction

## Learning Sets of Rules

Rules are very easy to understand; popular in data mining.

- Variable Size. Any boolean function can be represented.
- Deterministic.
- Discrete and Continuous Parameters.

Learning algorithms for rule sets can be described as

- Constructive Search. The rule set is built by adding rules; each rule is constructed by adding conditions.
- Eager.
- Batch.


## Rule Set Hypothesis Space

- Each rule is a conjunction of tests. Each test has the form $x_{j}=v, x_{j} \leq v$, or $x_{j} \geq v$, where $v$ is a value for $x_{j}$ that appears in the training data.

$$
x_{1}=\text { Sunny } \wedge x_{2} \leq 75 \% \Rightarrow y=1
$$

- A rule set is a disjunction of rules. Typically all of the rules are for one class (e.g., $y=1$ ). An example is classified into $y=1$ if any rule is satisfied.

$$
\begin{aligned}
x_{1}=\text { Sunny } \wedge x_{2} \leq 75 \% & \Rightarrow y=1 \\
x_{1}=\text { Overcast } & \Rightarrow y=1 \\
x_{1}=\text { Rain } \wedge x_{3} \leq 20 & \Rightarrow y=1
\end{aligned}
$$

## Relationship to Decision Trees

Any decision tree can be converted into a set of rules. The previous set of rules corresponds to this tree:


## Relationship to Decision Trees

A small set of rules can correspond to a big decision tree, because of the Replication Problem.

$$
x_{1} \wedge x_{2} \Rightarrow y=1 \quad x_{3} \wedge x_{4} \Rightarrow y=1 \quad x_{5} \wedge x_{6} \Rightarrow y=1
$$



## Learning a Single Rule

We grow a rule by starting with an empty rule and adding tests one at a time until the rule "covers" only positive examples.

GrowRule( $S$ )
$R=\{ \}$
repeat
choose best test $x_{j} \Theta v$ to add to $R$, where $\Theta \in\{=, \neq, \leq, \geq\}$
$S:=S$ - all examples that do not satisfy $R \cup\left\{x_{j} \Theta v\right\}$.
until $S$ contains only positive examples.

## Choosing the Best Test

- Current rule $R$ covers $m_{0}$ negative examples and $m_{1}$ positive examples.

Let $p=\frac{m_{1}}{m_{0}+m_{1}}$.

- Proposed rule $R \cup\left\{x_{j} \Theta v\right\}$ covers $m_{0}^{\prime}$ and $m_{1}^{\prime}$ examples.

Let $p^{\prime}=\frac{m_{1}^{\prime}}{m_{0}^{\prime}+m_{1}^{\prime}}$.

- Gain $=m_{1}^{\prime}\left[(-p \lg p)-\left(-p^{\prime} \lg p^{\prime}\right)\right]$

We want to reduce our surprise (to the point where we are certain), but we also want the rule to cover many examples. This formula tries to implement this tradeoff.

## Learning a Set of Rules (Separate-and-Conquer)

GrowRuleSet( $S$ )
$A=\{ \}$
repeat
$R:=\operatorname{GrowRule}(S)$
Add $R$ to $A$
$S:=S$ - all positive examples that satisfy $R$.
until $S$ is empty.
return $A$

## More Thorough Search Procedures

All of our algorithms so far have used greedy algorithms. Finding the smallest set of rules is NP-Hard. But there are some more thorough search procedures that can produce better rule sets.

- Round-Robin Replacement. After growing a complete rule set, we can delete the first rule, compute the set $S$ of training examples not covered by any rule, and one or more new rules, to cover $S$. This can be repeated with each of the original rules. This process allows a later rule to "capture" the positive examples of a rule that was learned earlier.
- Backfitting. After each new rule is added to the rule set, we perform a few iterations of Round-Robin Replacement (it typically converges quickly). We repeat this process of growing a new rule and then performing Round-Robin Replacement until all positive examples are covered.
- Beam Search. Instead of growing one new rule, we grow $B$ new rules. We consider adding each possible test to each rule and keep the best $B$ resulting rules. When no more tests can be added, we choose the best of the $B$ rules and add it to the rule set.


## Probability Estimates From Small Numbers

When $m_{0}$ and $m_{1}$ are very small, we can end up with

$$
p=\frac{m_{1}}{m_{0}+m_{1}}
$$

being very unreliable (or even zero).

## Two possible fixes

- Laplace Estimate. Add $1 / 2$ to the numerator and 1 to the denominator:

$$
p=\frac{m_{1}+0.5}{m_{0}+m_{1}+1}
$$

This is essentially saying that in the absence of any evidence, we expect $p=1 / 2$, but our belief is very weak (equivalent to $1 / 2$ of an example).

- General Prior Estimate. If you have a prior belief that $p=0.25$, you can add any number $k$ to the numerator and $4 k$ to the denominator.

$$
p=\frac{m_{1}+k}{m_{0}+m_{1}+4 k}
$$

The larger $k$ is, the stronger our prior belief becomes.

Many authors have added 1 to both the numerator and denominator in rule learning cases (weak prior belief that $p=1$ ).

## Learning Rules for Multiple Classes

What if rules for more than one class?
Two possibilities:

- Order rules (decision list)
- Weighted vote (e.g., weight $=$ accuracy $\times$ coverage)


## Learning First-Order Rules

Why do that?

- Can learn sets of rules such as

```
Ancestor (x,y)\leftarrowParent(x,y)
Ancestor (x,y)}\leftarrow\operatorname{Parent (x,z)^Ancestor (z,y)
```

- The Prolog programming language: programs are sets of such rules


## First-Order Rule for Classifying Web Pages

[Slattery, 1997]
course (A) $\leftarrow$
has-word(A, instructor),
$\neg$ has-word (A, good),
link-from(A, B),
has-word(B, assign),
$\neg$ link-from (B, C)
Train: $31 / 31$, Test: $31 / 34$

## FOIL (First-Order Inductive Learner)

Same as propositional separate-and-conquer, except:

- Different candidate specializations (literals)
- Different evaluation function


## Specializing Rules in FOIL

Learning rule: $P\left(x_{1}, x_{2}, \ldots, x_{k}\right) \leftarrow L_{1} \ldots L_{n}$
Candidate specializations add new literal of form:

- $Q\left(v_{1}, \ldots, v_{r}\right)$, where at least one of the $v_{i}$ in the created literal must already exist as a variable in the rule.
- Equal $\left(x_{j}, x_{k}\right)$, where $x_{j}$ and $x_{k}$ are variables already presentin the rule
- The negation of either of the above forms of literals


## Information Gain in FOIL

$$
\text { Foil_Gain }(L, R) \equiv t\left(\log _{2} \frac{p_{1}}{p_{1}+n_{1}}-\log _{2} \frac{p_{0}}{p_{0}+n_{0}}\right)
$$

Where

- $L$ is the candidate literal to add to rule $R$
- $p_{0}=$ number of positive bindings of $R$
- $n_{0}=$ number of negative bindings of $R$
- $p_{1}=$ number of positive bindings of $R+L$
- $n_{1}=$ number of negative bindings of $R+L$
- $t=$ no. of positive bindings of $R$ also covered by $R+L$


## FOIL Example



Target function:

- CanReach $(x, y)$ true iff directed path from $x$ to $y$ Instances:
- Pairs of nodes, e.g $\langle 1,5\rangle$, with graph described by literals $\operatorname{LinkedTo}(0,1)$, $\neg \operatorname{LinkedTo}(0,8)$ etc.
Hypothesis space:
- Each $h \in H$ is a set of Horn clauses using predicates LinkedTo (and CanReach)


## Induction as Inverted Deduction

Induction is finding $h$ such that

$$
\left(\forall\left\langle x_{i}, f\left(x_{i}\right)\right\rangle \in D\right) B \wedge h \wedge x_{i} \vdash f\left(x_{i}\right)
$$

where

- $x_{i}$ is $i$ th training instance
- $f\left(x_{i}\right)$ is the target function value for $x_{i}$
- $B$ is other background knowledge

So let's design inductive algorithm by inverting operators for automated deduction.

## Induction as Inverted Deduction

"Pairs of people $\langle u, v\rangle$ such that child of $u$ is $v$ "
$f\left(x_{i}\right):$ Child(Bob,Sharon)
$x_{i}: \quad \operatorname{Male}(B o b), F e m a l e(S h a r o n)$, Father(Sharon, Bob)
$B: \operatorname{Parent}(u, v) \leftarrow F \operatorname{Father}(u, v)$

What satisfies $\left(\forall\left\langle x_{i}, f\left(x_{i}\right)\right\rangle \in D\right) B \wedge h \wedge x_{i} \vdash f\left(x_{i}\right) ?$

$$
\begin{array}{ll}
h_{1}: & \operatorname{Child}(u, v) \leftarrow F \operatorname{ather}(v, u) \\
h_{2}: & \operatorname{Child}(u, v) \leftarrow \operatorname{Parent}(v, u)
\end{array}
$$

## Induction as Inverted Deduction

We have mechanical deductive operators $F(A, B)=C$, where $A \wedge B \vdash C$

Need inductive operators

$$
O(B, D)=h \text { where }\left(\forall\left\langle x_{i}, f\left(x_{i}\right)\right\rangle \in D\right)\left(B \wedge h \wedge x_{i}\right) \vdash f\left(x_{i}\right)
$$

## Induction as Inverted Deduction

## Positives:

- Subsumes earlier idea of finding $h$ that "fits" training data
- Domain theory $B$ helps define meaning of "fit" the data

$$
B \wedge h \wedge x_{i} \vdash f\left(x_{i}\right)
$$

- Suggests algorithms that search $H$ guided by $B$


## Induction as Inverted Deduction

Negatives:

- Doesn't allow for noisy data. Consider

$$
\left(\forall\left\langle x_{i}, f\left(x_{i}\right)\right\rangle \in D\right)\left(B \wedge h \wedge x_{i}\right) \vdash f\left(x_{i}\right)
$$

- First order logic gives a huge hypothesis space $H$
$\rightarrow$ Overfitting
$\rightarrow$ Intractability of calculating all acceptable $h$ 's


## Deduction: Resolution Rule

| $P$ | $\vee$ | $L$ |
| :---: | :---: | :---: |
| $\neg L$ | $\vee$ | $R$ |
| $P$ | $\vee$ | $R$ |

1. Given initial clauses $C_{1}$ and $C_{2}$, find a literal $L$ from clause $C_{1}$ such that $\neg L$ occurs in clause $C_{2}$
2. Form the resolvent $C$ by including all literals from $C_{1}$ and $C_{2}$, except for $L$ and $\neg L$. More precisely, the set of literals occurring in the conclusion $C$ is

$$
C=\left(C_{1}-\{L\}\right) \cup\left(C_{2}-\{\neg L\}\right)
$$

where $\cup$ denotes set union, and "-" is set difference

## Inverting Resolution



## Inverted Resolution (Propositional)

1. Given initial clauses $C_{1}$ and $C$, find a literal $L$ that occurs in clause $C_{1}$, but not in clause $C$.
2. Form the second clause $C_{2}$ by including the following literals

$$
C_{2}=\left(C-\left(C_{1}-\{L\}\right)\right) \cup\{\neg L\}
$$

## First-Order Resolution

1. Find a literal $L_{1}$ from clause $C_{1}$, literal $L_{2}$ from clause $C_{2}$, and substitution $\theta$ such that $L_{1} \theta=\neg L_{2} \theta$
2. Form the resolvent $C$ by including all literals from $C_{1} \theta$ and $C_{2} \theta$, except for $L_{1} \theta$ and $\neg L_{2} \theta$. More precisely, the set of literals occurring in the conclusion $C$ is

$$
C=\left(C_{1}-\left\{L_{1}\right\}\right) \theta \cup\left(C_{2}-\left\{L_{2}\right\}\right) \theta
$$

## Inverting First-Order Resolution

$$
C_{2}=\left(C-\left(C_{1}-\left\{L_{1}\right\}\right) \theta_{1}\right) \theta_{2}^{-1} \cup\left\{\neg L_{1} \theta_{1} \theta_{2}^{-1}\right\}
$$

## Cigol



## Progol

Progol: Reduce comb explosion by generating the most specific acceptable $h$

1. User specifies $H$ by stating predicates, functions, and forms of arguments allowed for each
2. Progol uses sequential covering algorithm. For each $\left\langle x_{i}, f\left(x_{i}\right)\right\rangle$

- Find most specific hypothesis $h_{i}$ s.t.
$B \wedge h_{i} \wedge x_{i} \vdash f\left(x_{i}\right)$
- actually, considers only $k$-step entailment

3. Conduct general-to-specific search bounded by specific hypothesis $h_{i}$, choosing hypothesis with minimum description length

## Rule Induction: Summary

- Rule grown by adding one antecedent at a time
- Rule set grown by adding one rule at a time
- Propositional or first-order
- Alternative: inverse resolution

