

- Often the values of properties are not meaningful values but names of individuals.
- It is the properties of these individuals and their relationship to other individuals that needs to be learned.
- Relational learning has been studied under the umbrella of “Inductive Logic Programming” as the representations are often logic programs.

## Example: trading agent

What does Joe like?

Individual	Property	Value
<i>joe</i>	<i>likes</i>	<i>resort_14</i>
<i>joe</i>	<i>dislikes</i>	<i>resort_35</i>
...	...	...
<i>resort_14</i>	<i>type</i>	<i>resort</i>
<i>resort_14</i>	<i>near</i>	<i>beach_18</i>
<i>beach_18</i>	<i>type</i>	<i>beach</i>
<i>beach_18</i>	<i>covered_in</i>	<i>ws</i>
<i>ws</i>	<i>type</i>	<i>sand</i>
<i>ws</i>	<i>color</i>	<i>white</i>
...	...	...

Values of properties may be meaningless names.

## Example: trading agent

Possible theory that could be learned:

$$\begin{aligned} \text{prop}(\text{joe}, \text{likes}, R) \leftarrow \\ \text{prop}(R, \text{type}, \text{resort}) \wedge \\ \text{prop}(R, \text{near}, B) \wedge \\ \text{prop}(B, \text{type}, \text{beach}) \wedge \\ \text{prop}(B, \text{covered\_in}, S) \wedge \\ \text{prop}(S, \text{type}, \text{sand}). \end{aligned}$$

Joe likes resorts that are near sandy beaches.

# Inductive Logic Programming: Inputs and Output

- $A$  is a set of atoms whose definitions the agent is learning.
- $E^+$  is a set of ground atoms observed true: **positive examples**
- $E^-$  is the set of ground atoms observed to be false: **negative examples**
- $B$  is a set of clauses: **background knowledge**
- $H$  is a space of possible hypotheses.  $H$  can be the set of all logic programs defining  $A$ .

The aim is to find a simplest hypothesis  $h \in H$  such that

$$B \wedge h \models E^+ \text{ and}$$

$$B \wedge h \not\models E^-$$

# Inductive Logic Programming: Main Approaches

Single target relation:  $A = \{t(X_1, \dots, X_n)\}$ .

Two main approaches:

- Start with the most general hypothesis and make it more complicated to fit the data.

# Inductive Logic Programming: Main Approaches

Single target relation:  $A = \{t(X_1, \dots, X_n)\}$ .

Two main approaches:

- Start with the most general hypothesis and make it more complicated to fit the data. Most general hypothesis is

$$t(X_1, \dots, X_n).$$

Keep adding conditions, ensuring it always implies the positive examples. Each step, exclude some negative examples.

# Inductive Logic Programming: Main Approaches

Single target relation:  $A = \{t(X_1, \dots, X_n)\}$ .

Two main approaches:

- Start with the most general hypothesis and make it more complicated to fit the data. Most general hypothesis is

$$t(X_1, \dots, X_n).$$

Keep adding conditions, ensuring it always implies the positive examples. Each step, exclude some negative examples.

- Start with a hypothesis that fits the data and keep making it simpler while still fitting the data.

# Inductive Logic Programming: Main Approaches

Single target relation:  $A = \{t(X_1, \dots, X_n)\}$ .

Two main approaches:

- Start with the most general hypothesis and make it more complicated to fit the data. Most general hypothesis is

$$t(X_1, \dots, X_n).$$

Keep adding conditions, ensuring it always implies the positive examples. Each step, exclude some negative examples.

- Start with a hypothesis that fits the data and keep making it simpler while still fitting the data. Initially the logic program can be  $E^+$ . Operators simplify the program, ensuring it fits the training examples.

# Inductive Logic Programming: General to Specific Search

Maintain a logic program  $G$  that entails the positive examples.  
Initially:

$$G = \{t(X_1, \dots, X_n) \leftarrow\}$$

A **specialization operator** takes  $G$  and returns set  $S$  of clauses that specializes  $G$ . Thus  $G \models S$ .

Three primitive specialization operators:

- Split a clause in  $G$  on condition  $c$ . Clause  $a \leftarrow b$  in  $G$  is replaced by two clauses:  $a \leftarrow b \wedge c$  and  $a \leftarrow b \wedge \neg c$ .
- Split clause  $a \leftarrow b$  on variable  $X$ , producing:

$$a \leftarrow b \wedge X = t_1.$$

...

$$a \leftarrow b \wedge X = t_k.$$

where the  $t_i$  are terms.

- Remove any clause not necessary to prove the positive examples.

# Top-down Inductive Logic Program

```
1: procedure TDInductiveLogicProgram( $t, B, E^+, E^-, R$ )
2:      $t$ : an atom whose definition is to be learned
3:      $B$ : background knowledge is a logic program
4:      $E^+$ : positive examples
5:      $E^-$ : negative examples
6:      $R$ : set of specialization operators
7:     Output: logic program that classifies  $E^+$  positively and
    $E^-$  negatively or  $\perp$  if no program can be found
8:          $H \leftarrow \{t(X_1, \dots, X_n) \leftarrow\}$ 
9:     while there is  $e \in E^-$  such that  $B \cup H \models e$  do
10:        if there is  $r \in R$  such that  $B \cup r(H) \models E^+$  then
11:           Choose  $r \in R$  such that  $B \cup r(H) \models E^+$ 
12:            $H \leftarrow r(H)$ 
13:        else
14:           return  $\perp$ 
15:    return  $H$ 
```