CS60020: Foundations of Algorithm Design and Machine Learning

Sourangshu Bhattacharya

Learning a Class from Examples

- Class C of a "family car"
  - Prediction: Is car *x* a family car?
  - Knowledge extraction: What do people expect from a family car?
- Output:

  Positive (+) and negative (–) examples
- Input representation:

  x_1: price, x_2 : engine power

Training set $\mathcal{X}$

$$\mathcal{X} = \{\mathbf{x}^t, r^t\}_{t=1}^N$$

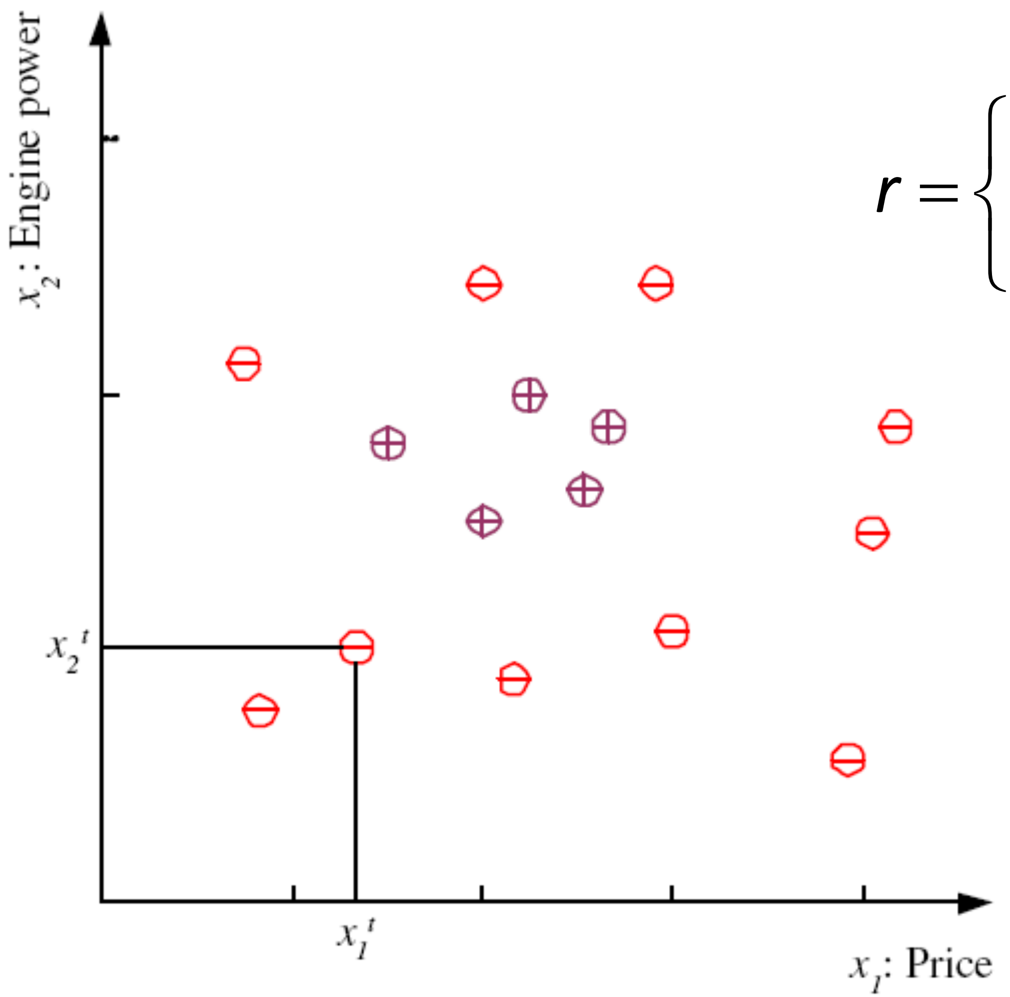

$$r = \begin{cases} 1 \text{ if } \mathbf{x} \text{ is positive} \\ 0 \text{ if } \mathbf{x} \text{ is negative} \end{cases}$$

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Class C

$$(p_1 \le \text{price} \le p_2) \text{ AND } (e_1 \le \text{engine power} \le e_2)$$

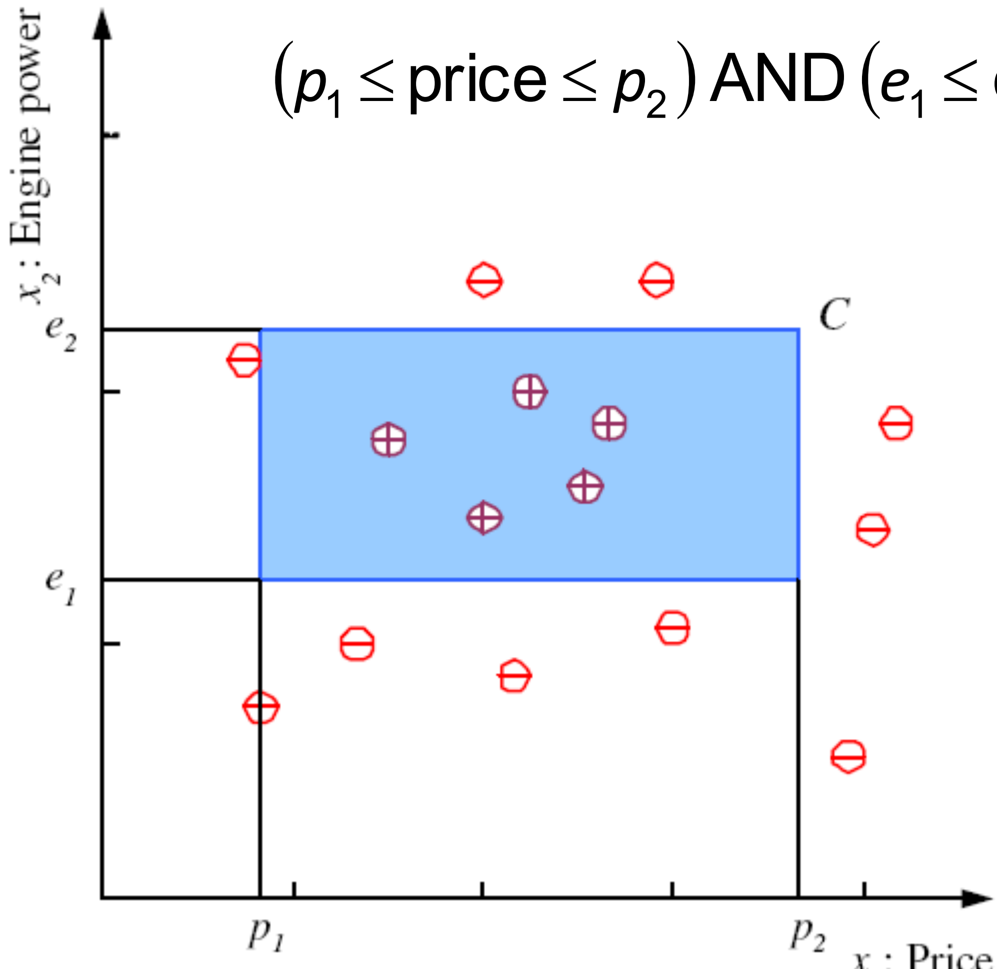

Hypothesis class $\mathcal{H}$

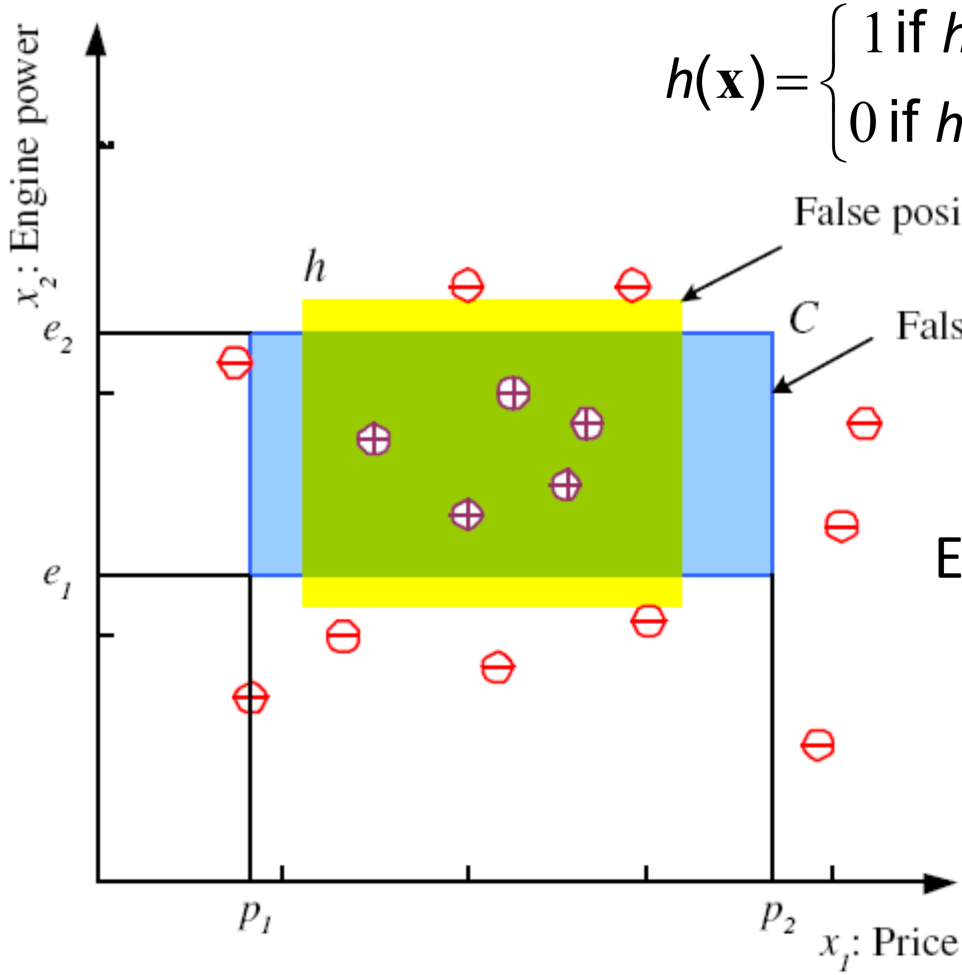

$$h(\mathbf{x}) = \begin{cases} 1 \text{ if } h \text{ says } \mathbf{x} \text{ is positive} \\ 0 \text{ if } h \text{ says } \mathbf{x} \text{ is negative} \end{cases}$$

Error of h on $\mathcal{H}$

$$E(h \mid \mathcal{X}) = \sum_{t=1}^{N} 1\left(h\left(\mathbf{x}^t\right) \neq r^t\right)$$

Concept learning

- Example / Instance: an atomic (real life) situation / object over which we want to learn.
- Instance space: Set of all possible instances.
- Attributes: observable quantities which describe a situation.
- Concept: a Boolean valued function over set of examples.
- Hypothesis space: subset of all Boolean valued functions over instance space.

Concept Learning - example

- Attributes: Sky, Air temp, Humidity, Wind, Weather, Forecast.
- Instance space X. What is the size ?
- Hypothesis space: conjunction of literals (which are conditions over attributes).
- Conditions are of the form: (attr=val) or (attr=?) or (attr=ϕ)
- What is the size of hypothesis space ?

Concept Learning - example

Training Examples for EnjoySport

Sky	Temp	Humid	Wind	Water	Forecst	EnjoySpt
Sunny	Warm	Normal	Strong	Warm	Same	Yes
Sunny	Warm	High	Strong	Warm	Same	Yes
Rainy	Cold	High	Strong	Warm	Change	No
Sunny	Warm	High	Strong	Cool	Change	Yes

What is the general concept?

Inductive learning problem

- Training examples: $D=\{ (x_1,c(x_1)) , \dots , (x_n,(c(x_n)) \}$
- Problem: Given *D*, learn $h \epsilon H$, such that for all $x \epsilon X$, $h(x)=c(x)$.
- Inductive learning assumption:

  Any hypothesis found to approximate target concept well over sufficiently large training set, will also approximate it well over unseen examples.

General to specific ordering

- Example x is said to be positive if $c(x) = 1$, else negative.
- Hypothesis h "satisfies" x, if $h(x) = 1$.
- Hypothesis h_2 is said to be "more general or equal to" h_1 if
  for all x: $h_1(x) = 1$ implies $h_2(x) = 1$

General to specific ordering

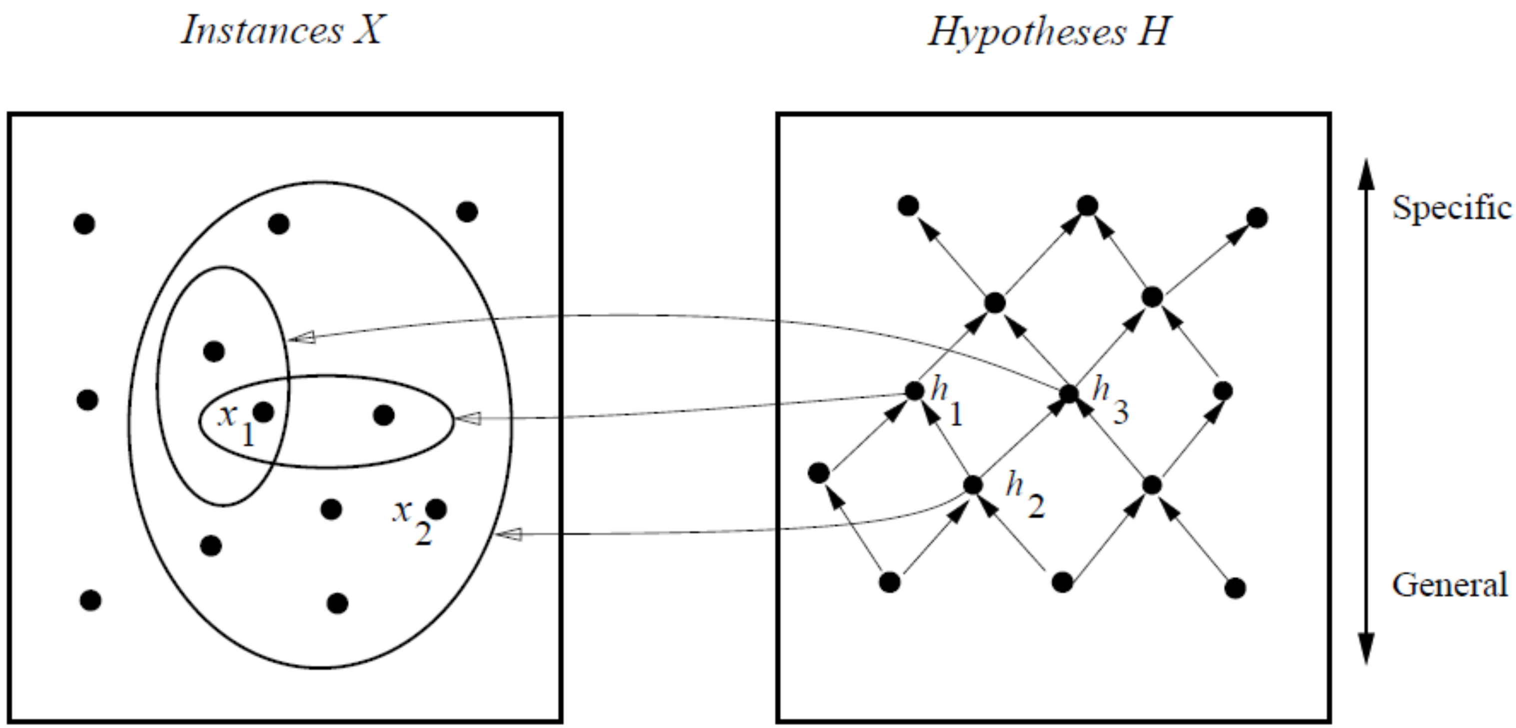

x_1= <*Sunny, Warm, High, Strong, Cool, Same*>
x_2= <*Sunny, Warm, High, Light, Warm, Same*>

h_1= <*Sunny, ?, ?, Strong, ?, ?*>
h_2= <*Sunny, ?, ?, ?, ?, ?*>
h_3= <*Sunny, ?, ?, ?, Cool, ?*>

Find - S

- Finding maximally specific hypothesis

1. Initialize h to the most specific hypothesis in H
2. For each positive training instance x
   - For each attribute constraint a_i in h

     If the constraint a_i in h is satisfied by x

     Then do nothing

     Else replace a_i in h by the next more general constraint that is satisfied by x
3. Output hypothesis h

Find – S Example

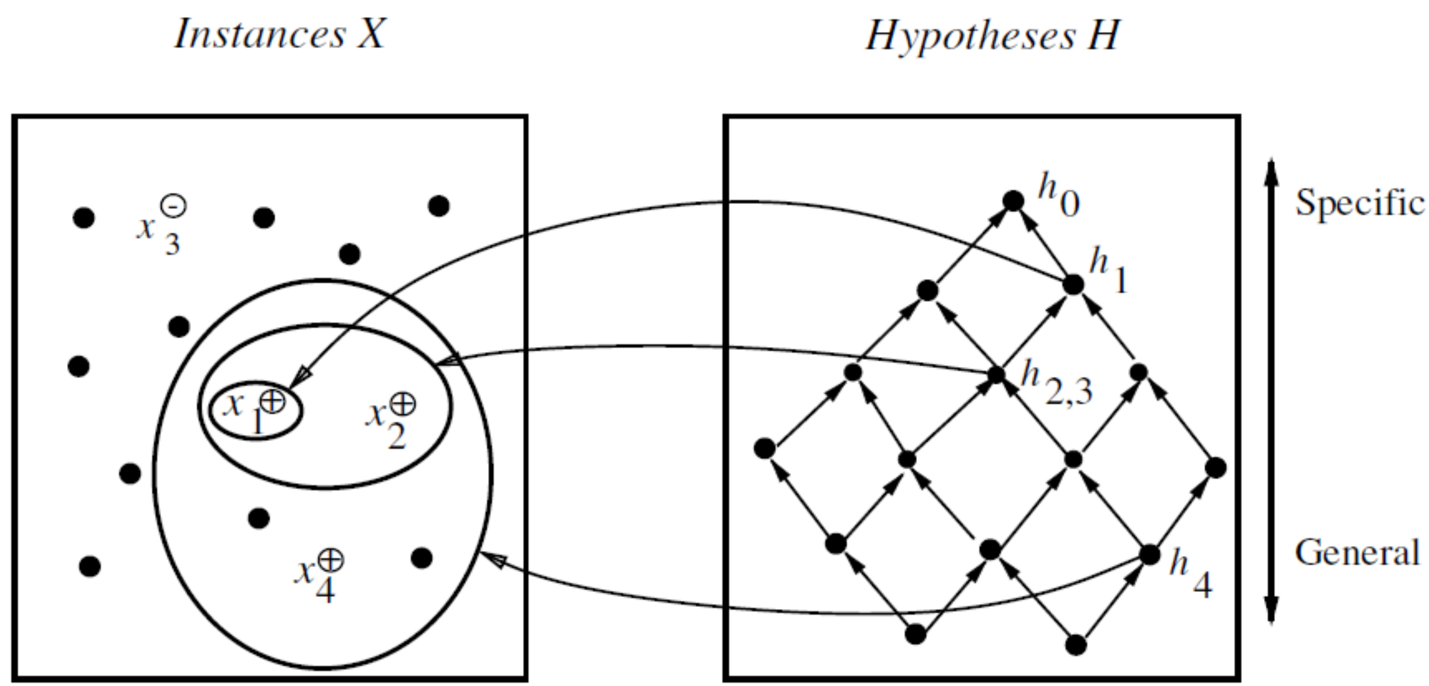

$h_0 = <\emptyset, \emptyset, \emptyset, \ \emptyset, \emptyset, \emptyset>$

x_1 = <*Sunny Warm Normal Strong Warm Same*>, + | h_1 = <*Sunny Warm Normal Strong Warm Sam*

x_2 = <*Sunny Warm High Strong Warm Same*>, + | h_2 = <*Sunny Warm ? Strong Warm Same*>

x_3 = <*Rainy Cold High Strong Warm Change*>, - | h_3 = <*Sunny Warm ? Strong Warm Same*>

x_4 = <*Sunny Warm High Strong Cool Change*>, + | h_4 = <*Sunny Warm ? Strong ? ?* >

Find – S Problems

- Can't tell whether it has learned the concept
- Can't tell whether the data is inconsistent
- Picks maximally specific hypothesis
- There might be several maximally specific hypothesis.

Version Space

A hypothesis h is **consistent** with a set of training examples D of target concept c if and only if $h(x) = c(x)$ for each training example $\langle x, c(x) \rangle$ in D.

$$Consistent(h, D) \equiv (\forall \langle x, c(x) \rangle \in D)\; h(x) = c(x)$$

The **version space**, $VS_{H,D}$, with respect to hypothesis space H and training examples D, is the subset of hypotheses from H consistent with all training examples in D.

$$VS_{H,D} \equiv \{h \in H | Consistent(h, D)\}$$

Version space representation

The **General boundary**, G, of version space $VS_{H,D}$ is the set of its maximally general members

The **Specific boundary**, S, of version space $VS_{H,D}$ is the set of its maximally specific members

Every member of the version space lies between these boundaries

$$VS_{H,D} = \{h \in H | (\exists s \in S)(\exists g \in G)(g \geq h \geq s)\}$$

where $x \geq y$ means x is more general or equal to y

Version space

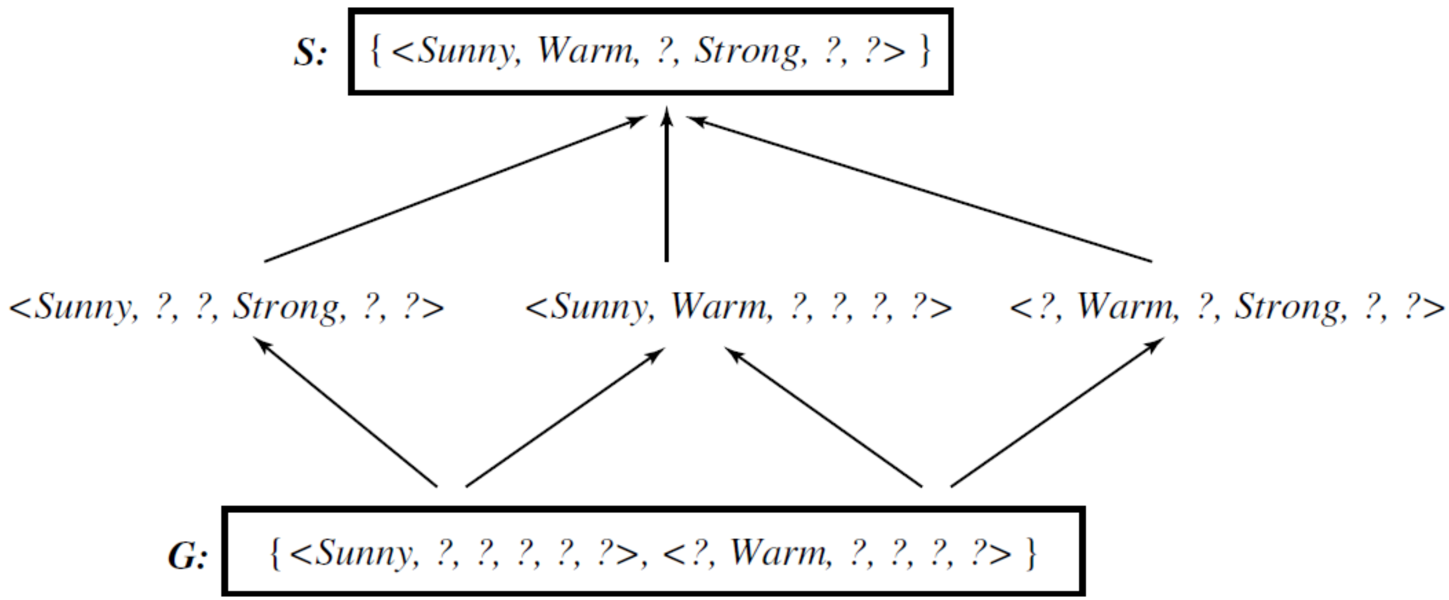

Candidate Elimination

$G \leftarrow$ maximally general hypotheses in H
$S \leftarrow$ maximally specific hypotheses in H
For each training example d, do

- If d is a positive example
  - Remove from G any hypothesis inconsistent with d
  - For each hypothesis s in S that is not consistent with d
    - Remove s from S
    - Add to S all minimal generalizations h of s such that
      1. h is consistent with d, and
      2. some member of G is more general than h
    - Remove from S any hypothesis that is more

Candidate Elimination

- If d is a negative example:
  - Remove from S any hypothesis inconsistent with d
  - For each hypothesis g in G that is not consistent with d
    - * Remove g from G
    - * Add to G all minimal specializations h of g such that
      1. h is consistent with d, and
      2. some member of S is more specific than h
    - * Remove from G any hypothesis that is less general than another hypothesis in G

Example Problem

Training Examples for EnjoySport

Sky	Temp	Humid	Wind	Water	Forecst	EnjoySpt
Sunny	Warm	Normal	Strong	Warm	Same	Yes
Sunny	Warm	High	Strong	Warm	Same	Yes
Rainy	Cold	High	Strong	Warm	Change	No
Sunny	Warm	High	Strong	Cool	Change	Yes

What is the general concept?

Example

- Workout …

Convergence

- Candidate elimination will converge to the target concept if:
  - Training data doesn't have errors.
  - Target concept lies in the hypothesis space.
- Otherwise
  - G and S sets become null.

Partially learned concept

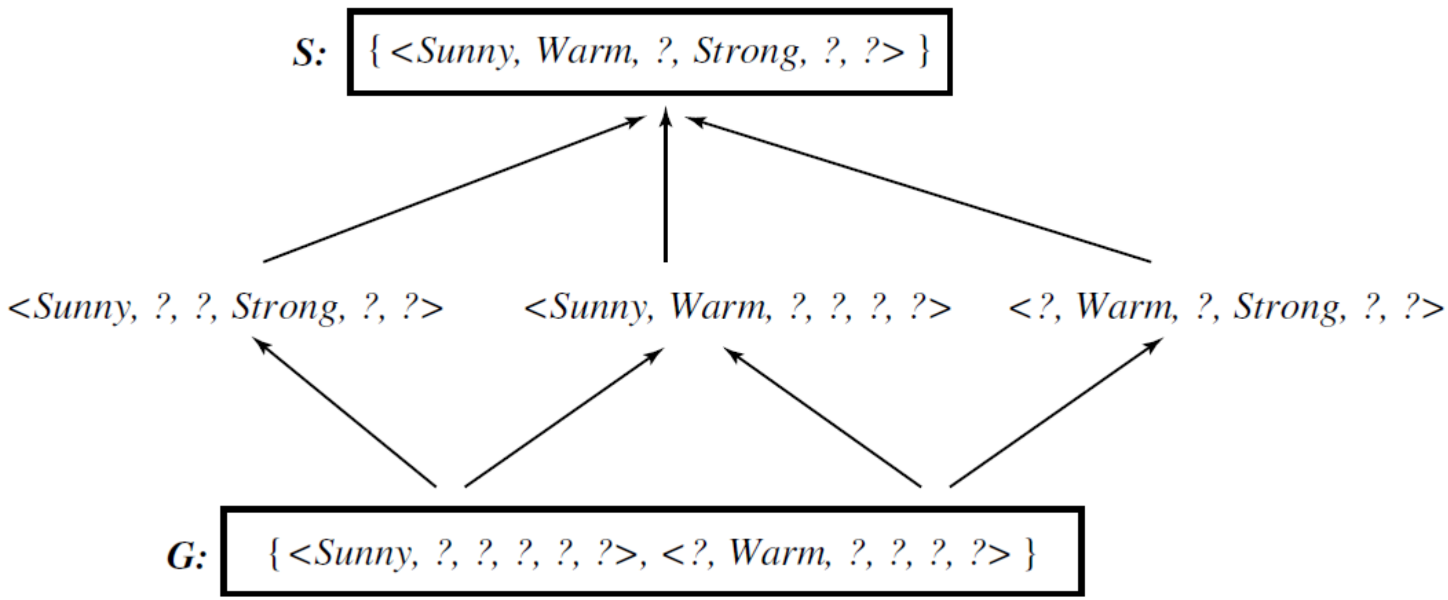

$$\langle Sunny\ Warm\ Normal\ Strong\ Cool\ Change \rangle$$

$$\langle Rainy\ Cool\ Normal\ Light\ Warm\ Same \rangle$$

$$\langle Sunny\ Warm\ Normal\ Light\ Warm\ Same \rangle$$

What next training example ?

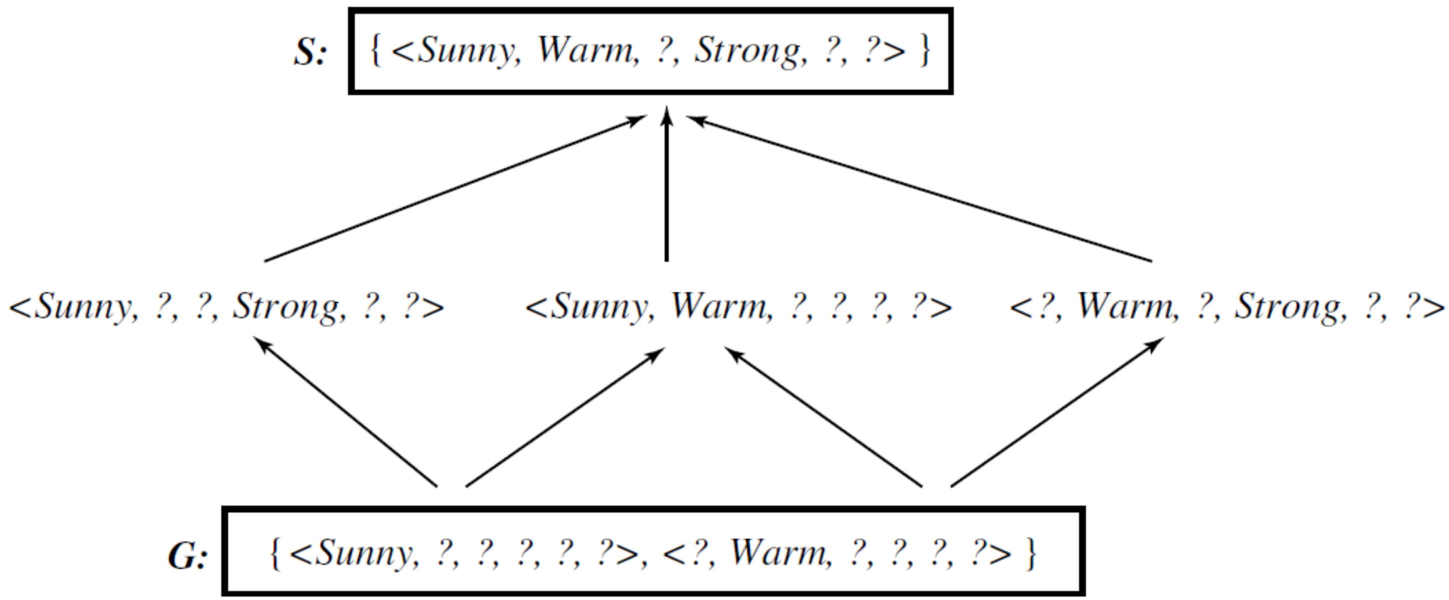

<Sunny, Warm, Normal, Light, Warm, Same>

Observations

- The hypothesis space is biased.
  - Example: XOR concept cannot be expressed.
- Unbiased learner – disjunction of conjunctions.
- Learned Version space:
  - S set: all positive examples
  - G set: compliment of all negative examples
- Can we use the partially learned concept from above ?
  - There is perfect ambiguity for all examples not in training set.

Unbiased learning

- Learning in an unbiased hypothesis space is futile as it cannot generalize to examples other than training examples.

S, G, and the Version Space

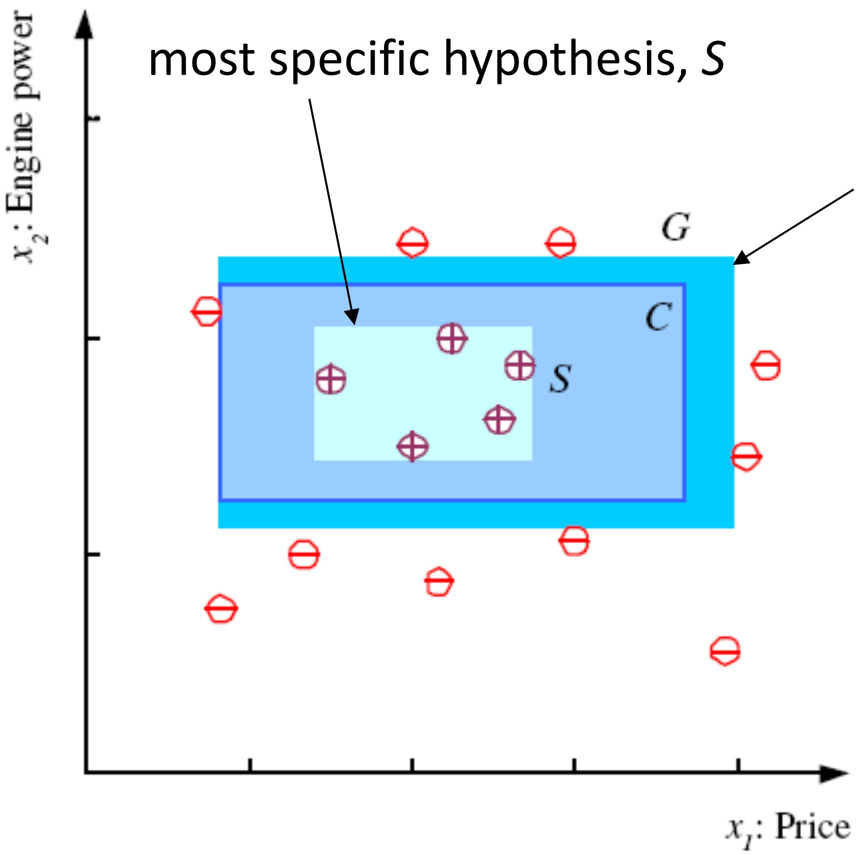

$h \in$ H, between *S* and *G* is consistent
and make up the version space
(Mitchell, 1997)