Outline

- Concept learning as boolean function approximation
- More-General-Than and Consistency
- Version Spaces
- Choosing useful training examples
- What is an unbiased learner?

Copyright Tom Mitchell, 14 September 1994

A Concept Learning Task

Given:

- Instances: Possible days, each described by $Sky, Air\text{–}Temp, Humidity, Wind, Water\text{–}Temp, Forecast$
- Target function: $Enjoy\text{–}Sport?: Day \rightarrow \{Yes, No\}$
- Hypotheses: Each described by a conjunction of attribute values (e.g., $(Sky=Sunny \wedge Water=Cool)$)
- Positive and negative examples of target function

Determine:

- A hypothesis that best predicts the value of $Enjoy\text{–}Sport?$

Copyright Tom Mitchell, 14 September 1994

Inductive Learning Hypothesis: any hypothesis found to approximate the target function well over a sufficiently large set of training examples will also approximate the target function well over other unobserved examples.

Training Examples

#	Sky	Air	Humidity	Wind	Water	Forecast	Enjoy?
1	Sunny	Warm	Normal	Strong	Warm	Same	Yes
2	Sunny	Warm	High	Strong	Warm	Same	Yes
3	Rainy	Cold	High	Strong	Warm	Change	No
4	Sunny	Warm	High	Strong	Cool	Change	Yes

Copyright Tom Mitchell, 14 September 1994

More General Than

Let h_j and h_k be boolean fns. over I

More-General-Than(h_j,h_k) $\equiv \forall(i \in I)[h_k(i) \rightarrow h_j(i)]$

Copyright Tom Mitchell, 14 September 1994

Instance Space and Hypothesis Space

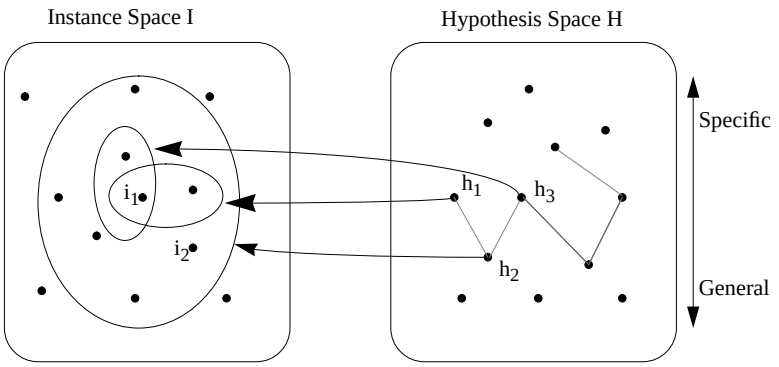

i_1: <Sunny, Warm, High, Strong, Cool, Same>
i_2: <Sunny, Warm, High, Light, Warm, Same>

h_1: <Sunny, ?, ?, Strong, ?, ?>
h_2: <Sunny, ?, ?, ?, ?, ?>
h_3: <Sunny, ?, ?, ?, Cool, ?>

Relationship between the space of instances, I, and the space of hypotheses, H. Each hypothesis covers some set of instances. Lines between hypotheses in H indicate the More-General-Than relation. For example, hypothesis h_2 is More-General-Than hypothesis h_1 because it covers a superset of the instances covered by h_1. Similarly, h_2 is More-General-Than h_3, while neither h_1 nor h_3 is More-General-Than the other. This More-General-Than relation induces a partial ordering over H which allows efficiently organizing the search over the hypothesis space.

Copyright Tom Mitchell, 14 September 1994

Version Spaces

Consistency: Hypothesis h is *consistent* with a set of training examples E, denoted Consistent(h,E), if and only if h correctly classifies each example in E.

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

$$VS_{H,E} = \{h \in H | \text{Consistent}(h,E)\}$$

The *general boundary*, G, of a version space $VS_{H,E}$ is the set of maximally general members of $VS_{H,E}$.

The *specific boundary*, S, of a version space $VS_{H,E}$ is the set of its minimally general (i.e., maximally specific) members of $VS_{H,E}$.

Version space representation theorem: The version space can be represented by S and G.

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

Copyright Tom Mitchell, 14 September 1994

Example Version Space

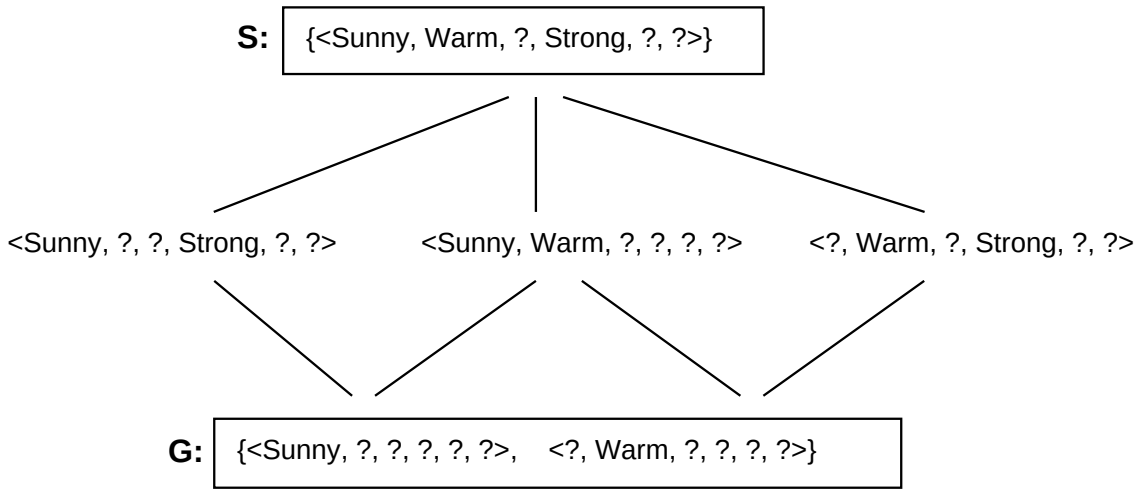

Copyright Tom Mitchell, 14 September 1994

Version Space Candidate Elimination Algorithm

Initialize S and G

For each new positive instance, i_+,

Remove members of G that fail to match i_+

Generalize members of S until they match i_+,
but retain only if they remain more specific than G

Remove any non-maximal members from S

For each new negative instance, i_-,

Remove members of S that match i_-,

Specialize members of G until they
fail to match i_-,
but retain only if they remain more general than S

Remove any non-minimal members from G

Copyright Tom Mitchell, 14 September 1994

What is a Good Next Training Example?

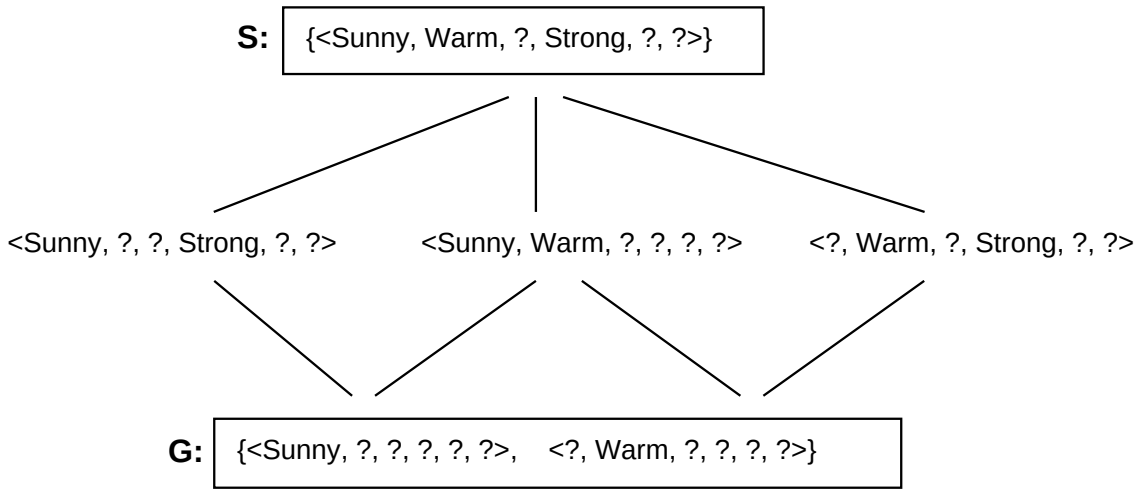

Copyright Tom Mitchell, 14 September 1994

How Should These be Classified?

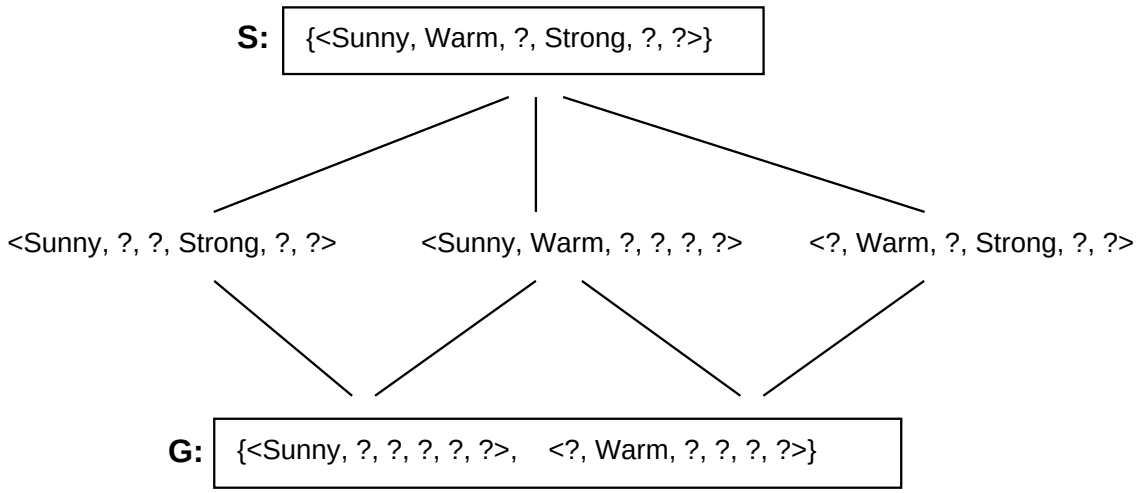

- <Sunny, Warm, Normal, Strong, Cool, Change>
- <Rainy, Cold, Normal, Light, Warm, Same>
- <Sunny, Warm, Normal, Light, Warm, Same>

Copyright Tom Mitchell, 14 September 1994

Version Spaces: Summary

- Guaranteed to converge to correct target concept
    - assuming error-free training data
    - assuming target concept is in H
- Allows choosing useful next example
- Allows confident classification of some new examples using partially learned concept
- Can be used with many different H languages
- For some H languages (e.g., disjunctive) S and G can become very large
- Bias-free algorithm. All bias is in H.

Copyright Tom Mitchell, 14 September 1994

What Justifies This Inductive Leap?

+ <Sunny, Warm, Normal, Strong, Warm, Same>
+ <Rainy, Warm, Normal, Weak, Warm, Same>

_____

S: <?, Warm, Normal, ?, Warm, Same>

Why believe that we can now classify the unseen:

? <Sunny, Warm, Normal, Weak, Warm, Same>

Copyright Tom Mitchell, 14 September 1994

An UNBiased Learner

Idea: Choose an H that expresses every possible subset of I

- H = Power set of I

H' = disjunctions, conjunctions, negations over previous H

Only a biased learner can generalize beyond observed data

But bias is risky...

Copyright Tom Mitchell, 14 September 1994

Inductive Bias

Bias of an inductive learning system: the set of assumptions which, combined with the observed training examples, logically entail subsequent instance classifications made by the learner.

Training-examples + bias $\vdash$ classifications

Copyright Tom Mitchell, 14 September 1994

Bias Characterizes Inductive System by its Equivalent Deductive System

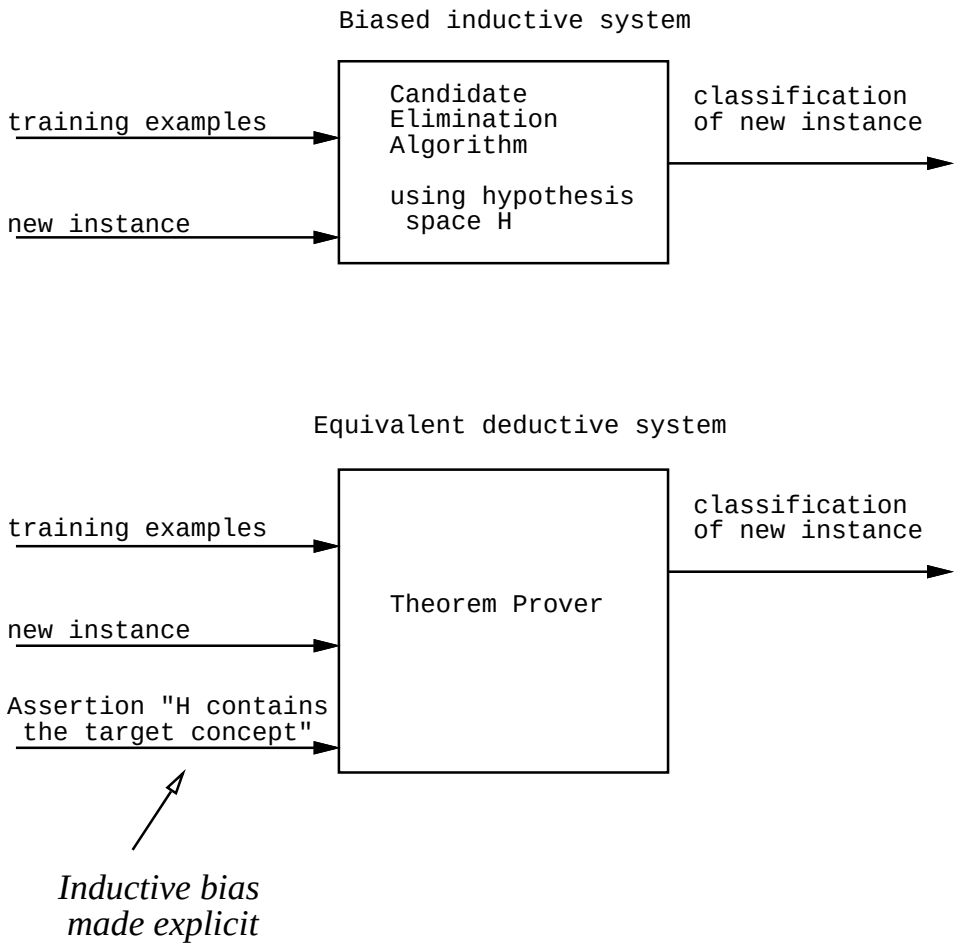

Three Learners with Different Biases

1. **Rote Learner:** Store examples. Classify new iff it matches old.
2. **Version Space Candidate Elimination**. Classify new examples only when VS unanimously classifies it.
3. **Generality Preferrer:** Compute VS. Classify new example positive if it matches any member of G boundary, else negative.

Copyright Tom Mitchell, 14 September 1994

Summary Points

1. Concept learning as search through H
2. General-to-specific ordering over H
3. Version space candidate elimination algorithm
4. S and G boundaries characterize learner's uncertainty
5. Learner can generate useful queries
6. Inductive leaps possible only if learner is biased
7. Inductive learners can be described by equivalent deductive systems

Copyright Tom Mitchell, 14 September 1994