CS 681 Fall 2008
Designing Expert Systems

Learning-based Knowledge Representation

Prof. Gheorghe Tecuci
tecuci@gmu.edu
http://lac.gmu.edu/

Learning Agents Center
and Computer Science Department
George Mason University

Overview

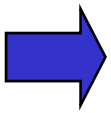

- Concept Representation
- Generalization and Specialization Rules
- Types of Generalizations and Specializations
- Partially Learned Knowledge
- Readings

Object Expressions

Object concept representation

One can define more complex concepts as logical expressions involving the basic concepts from the object ontology.

The following concept represents the set of all pairs (?O1 ?O2) where ?O1 is a PhD student interested in ?O2 which is a PhD research area:

?O1	*instance of*	*PhD student*
	is interested in	*?O2*
?O2	*instance of*	*PhD research area*

Which are some instances of this concept?

Exercise

What does the following concept represents?

?O1	*instance of*	*course*
	has as reading	*?O2*
?O2	*instance of*	*publication*
	has as author	*?O3*
?O3	*instance of*	*professor*

Which is an instance?

Exercise

What does the following concept represents?

$?O_1$	instance of	PhD student
	is interested in	$?O_2$
$?O_2$	instance of	PhD research area
Except When		
$?O_2$	instance of	PhD research area
	requires	"programming"

Overview

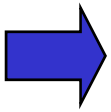

Concept Representation

Generalization and Specialization Rules

Types of Generalizations and Specializations

Partially Learned Knowledge

Readings

Generalization and Specialization Rules

A *generalization rule* is a rule that transforms an expression (or concept) into a more general expression.

A *specialization rule* is a rule that transforms an expression (or concept) into a less general expression.

The reverse of any generalization rule is a specialization rule.

Discussion

Indicate several generalizations of the following sentence:

Students who have majored in Computer Science at George Mason University between 2003 and 2004.

Provide another example of a concept and indicate some of its generalizations.

Discussion

Indicate several specializations of the following sentence:

Students who have majored in Computer Science at George Mason University between 2003 and 2004.

Generalization (and Specialization) Rules

- Turning constants into variables
- Turning occurrences of a variable into variables
- Climbing the generalization hierarchy
- Dropping condition
- Extending intervals
- Extending ordered sets of intervals
- Extending discrete sets
- Using feature definitions
- Using inference rules

Turning Constants into Variables

Generalizes an expression by replacing a constant with a variable.

The set of professors with 55 publications.

?O1 instance of professor
number of publications 55

generalization
55 => ?N1

specialization
?N1=> 55

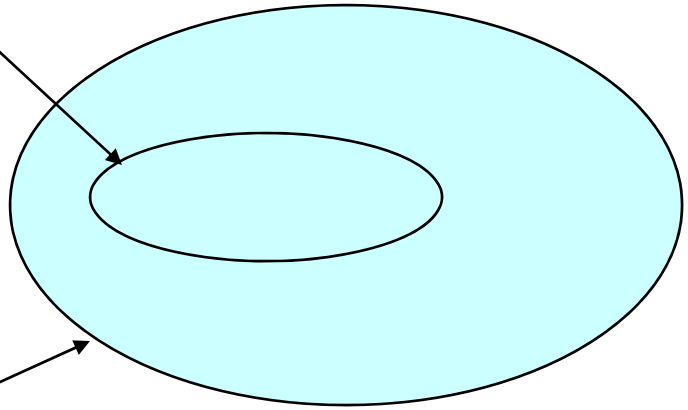

?O1 instance of professor
number of publications ?N1

The set of professors with any number of publications.

Explanation

The top expression represents the following concept: the set of professors with 55 publications.

By replacing 55 with a variable ?N1 that can take any value, we generalize this concept to the following one: the set of professors with any number of publications. In particular ?N1 could be 55. Therefore the second concept includes the first one.

Conversely, by replacing ?N1 with 55, we specialize the bottom concept to the top one.

The important thing to notice here is that by a simple syntactic operation (transforming a number into a variable) we can generalize a concept. This is one way in which an agent generalizes concepts.

Climbing the Generalization Hierarchies

Generalizes an expression by replacing a concept with a more general one.

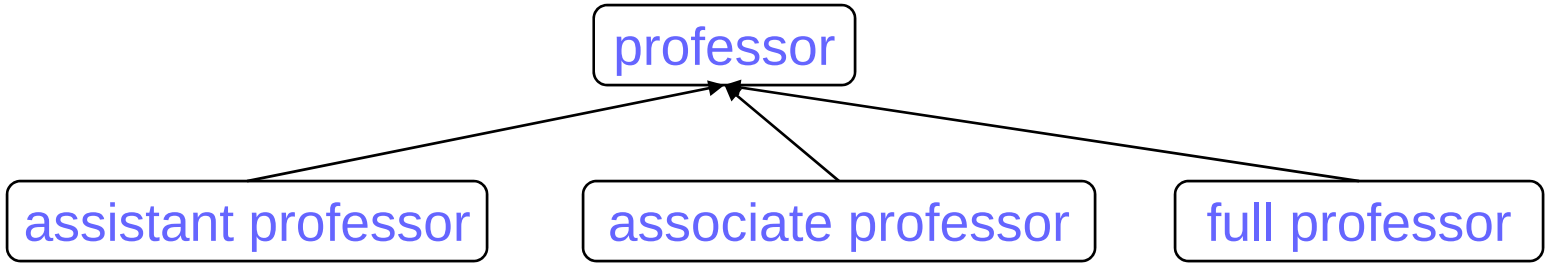

?O1	*instance of*	*assistant professor*
	has as employer	*?O2*
?O2	*instance of*	*state university*

The set of assistant professors employed by state universities

generalization

assistant professor => professor

specialization

professor => assistant professor

?O1	*instance of*	*professor*
	has as employer	*?O2*
?O2	*instance of*	*state university*

The set of professors employed by state universities

Explanation

One can also generalize an expression by replacing a concept from its description with a more general concept, according to some generalization hierarchy.

The reverse operation, of replacing a concept with a less general one, leads to the specialization of an expression.

The agent can also generalize a concept by dropping a condition. That is, by dropping a constraint that its instances must satisfy.
This rule is illustrated in the next slide.

Dropping Conditions

Generalizes an expression by removing a constraint from its description.

The set of assistant professors employed by state universities.

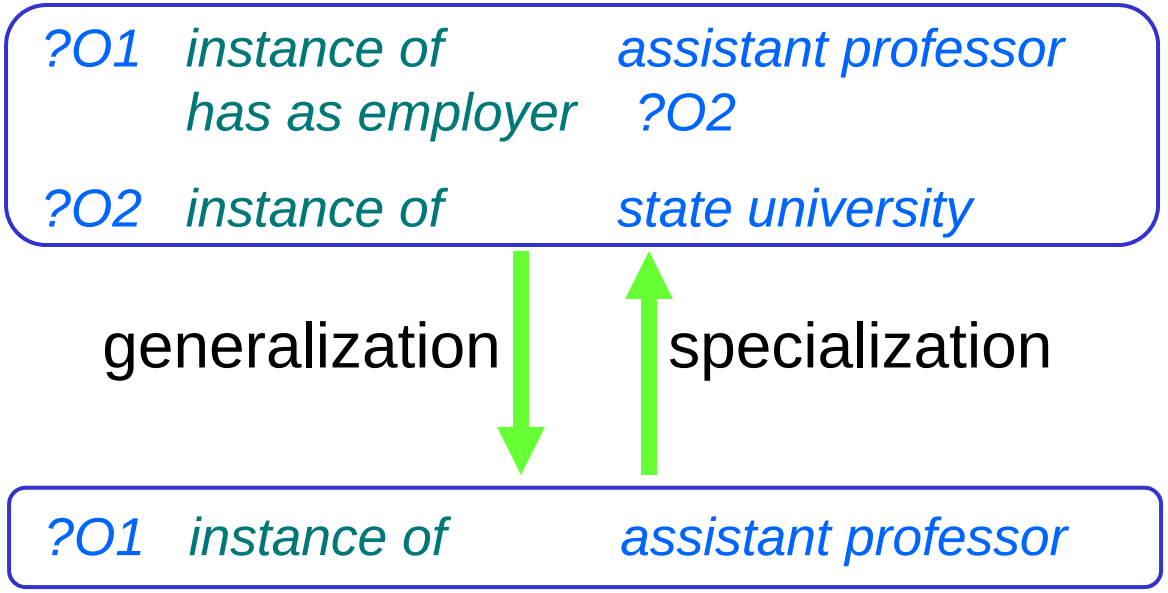

The set of assistant professors

Extending Intervals

Generalizes an expression by replacing a number with an interval, or by replacing an interval with a larger interval.

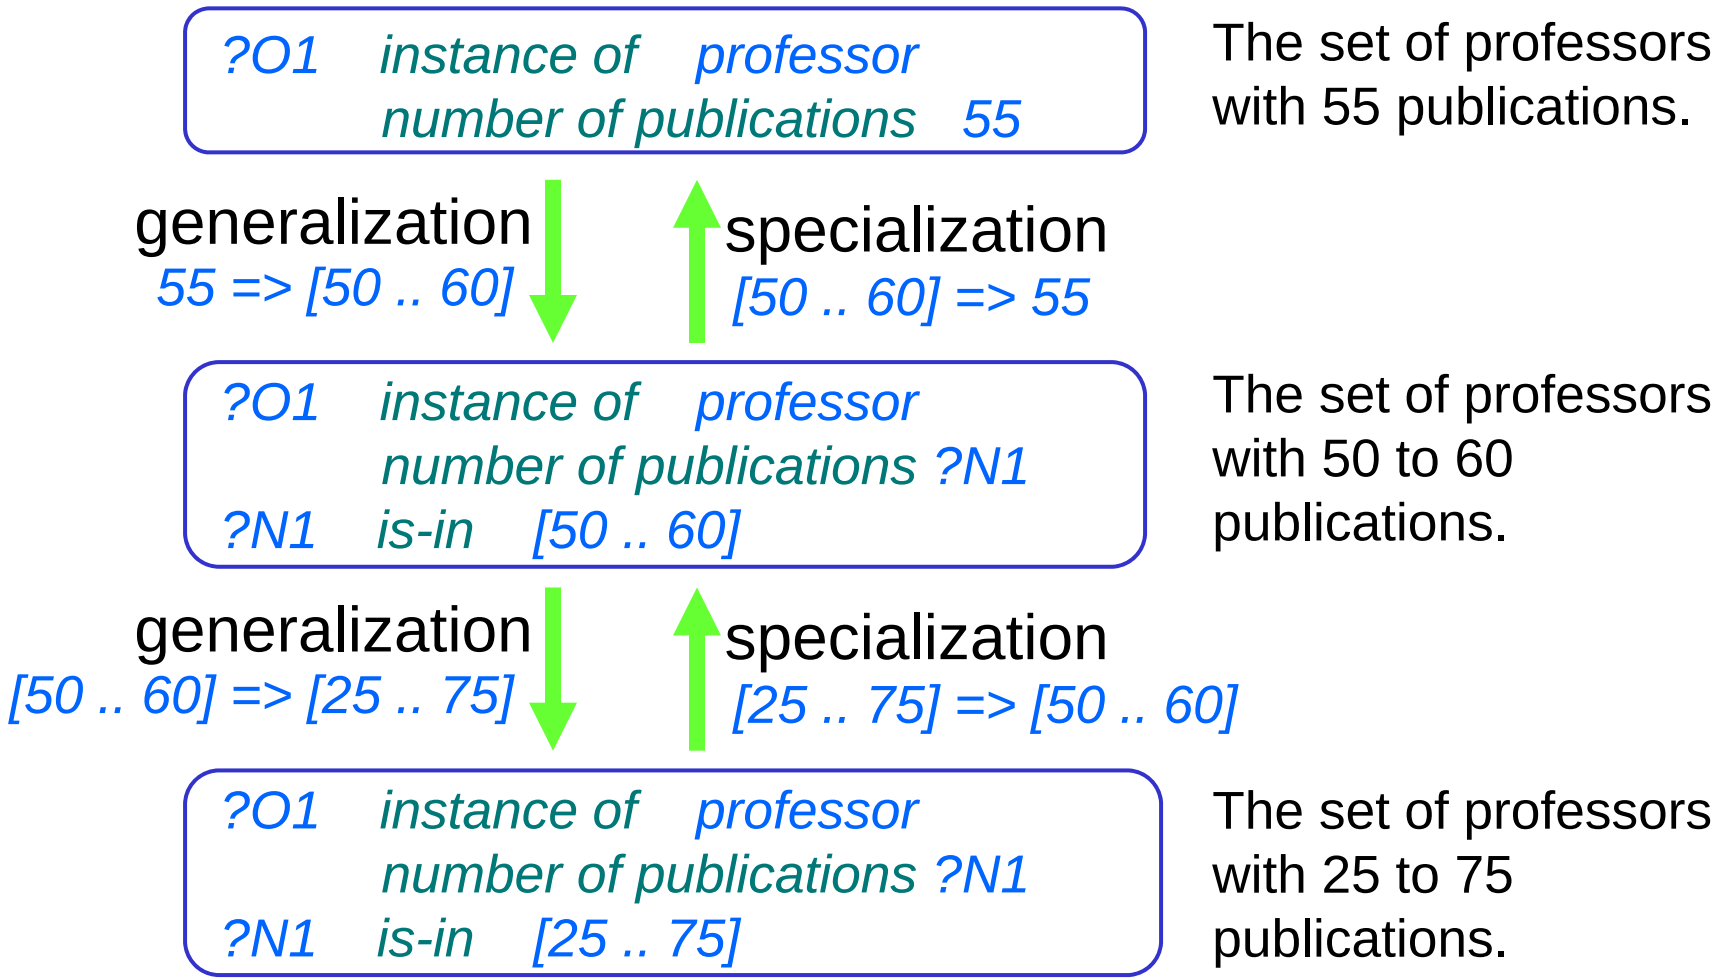

Explanation

A concept may also be generalized by replacing a number with an interval containing it, or by replacing an interval with a larger interval. The reverse operations specialize the concept.

Yet another generalization rule, which is illustrated in the next slide, is to add alternatives.

Extending Ordered Sets of Intervals

Generalizes an expression by replacing a symbolic interval with the larger interval

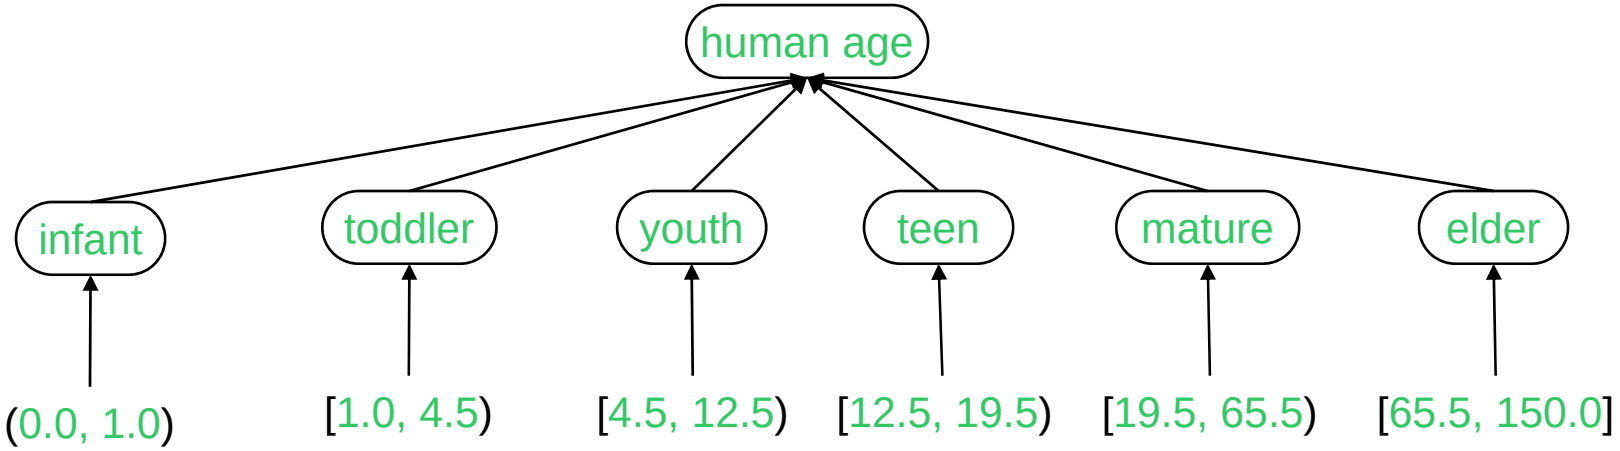

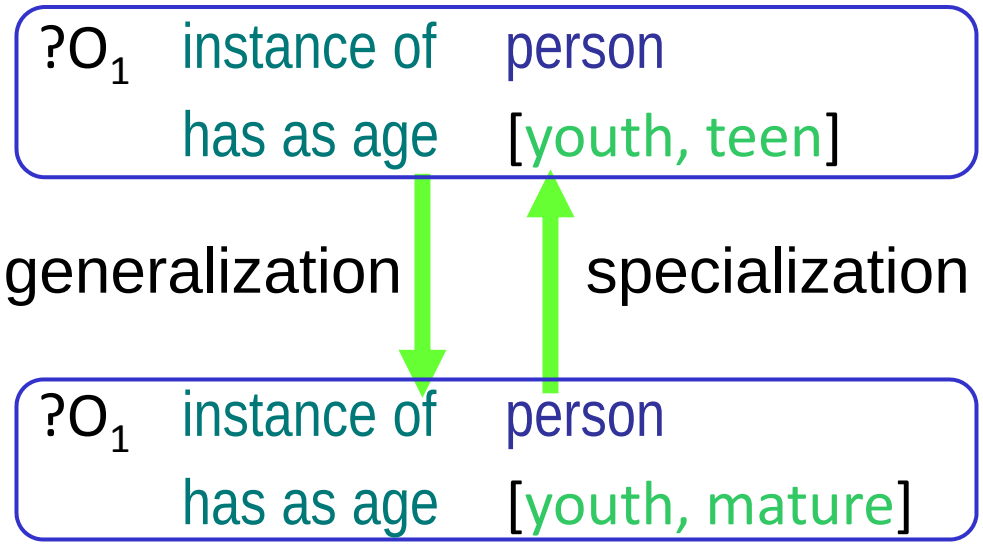

Extending Discrete Sets

Generalizes an expression by replacing a discrete set with a larger set

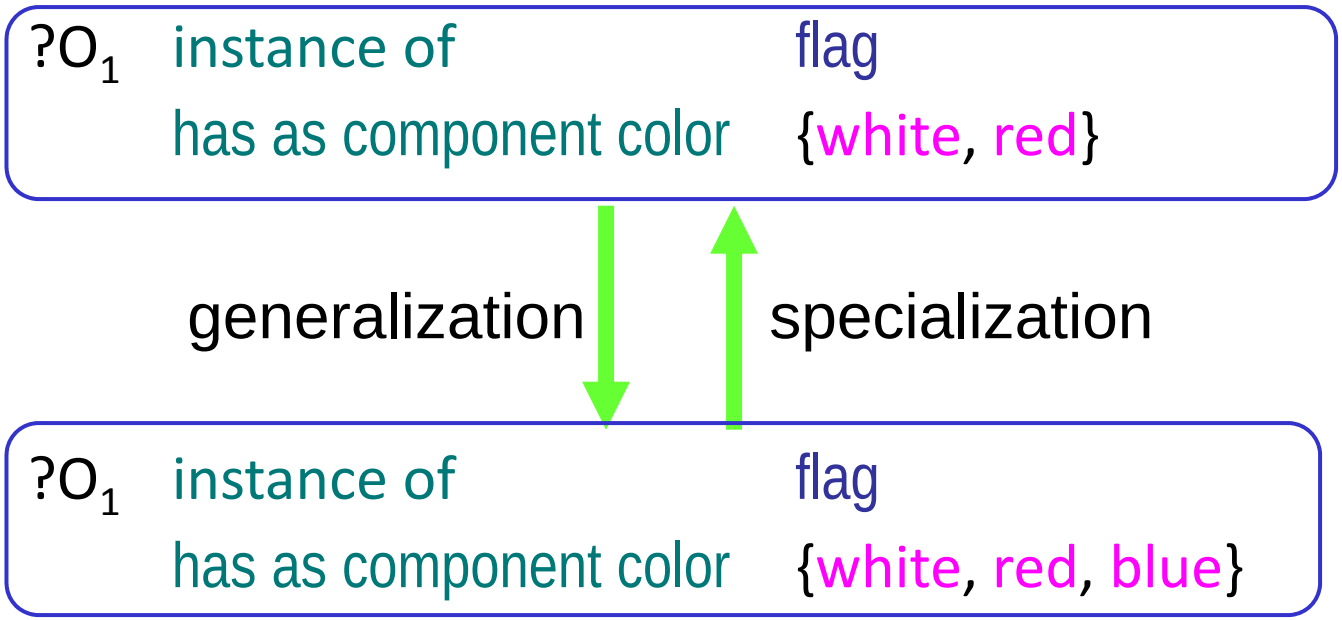

Using Feature Definitions

Generalizes an expression containing “A feature B” by replacing A and B with feature’s domain and range, respectively.

?O1 instance of professor
in expert in ?O2

?O2 instance of Computer Science

generalization

specialization

?O1 instance of person
is expert in ?O2

?O2 instance of area of expertise

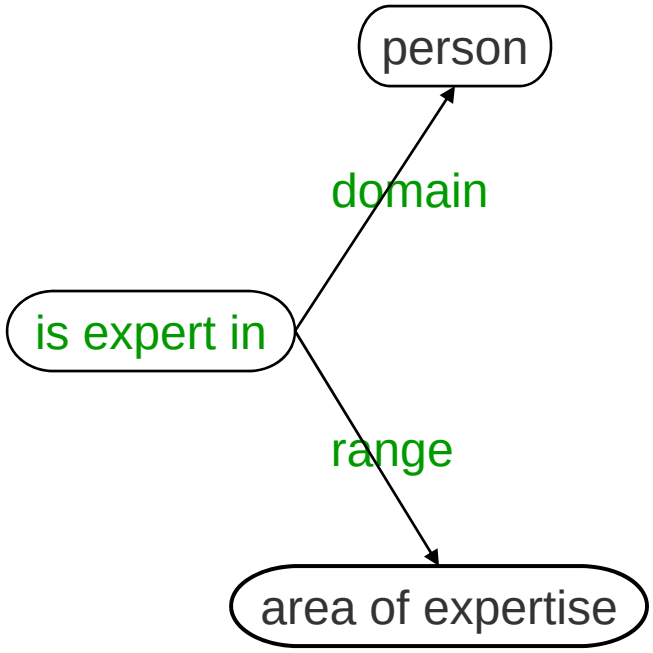

Using Inference Rules

Given an inference rule of the form "A → B" generalizes an expression by replacing A with B.

∀x, ∀y, ((x has as PhD advisor y) → (x knows y))

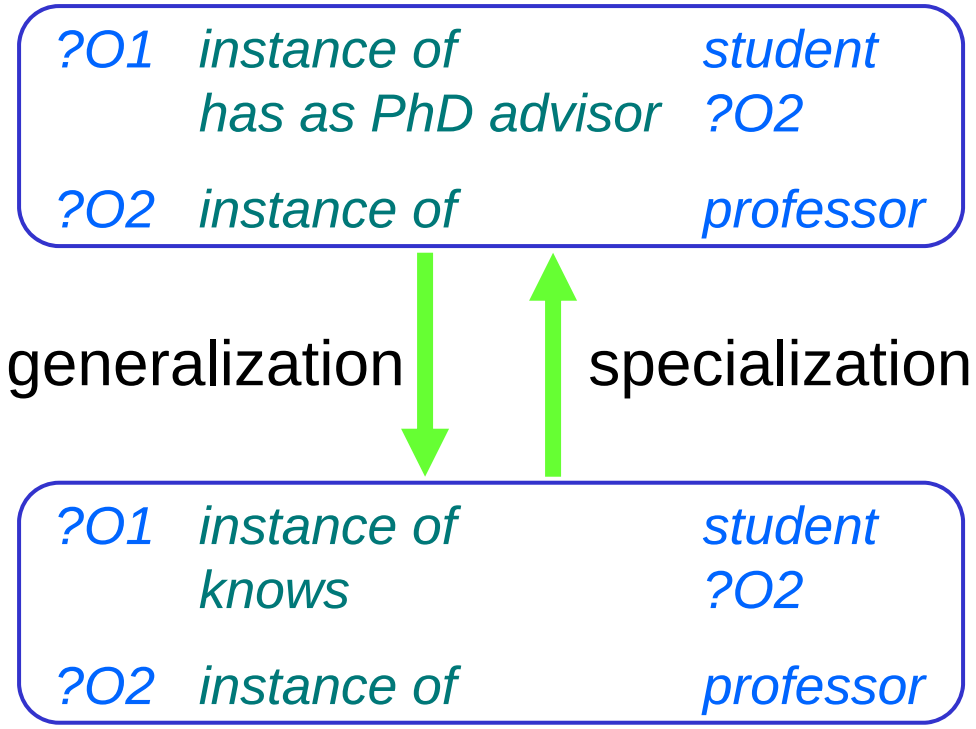

Overview

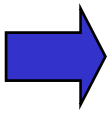

- Concept Representation
- Generalization and Specialization Rules
- Types of Generalizations and Specializations
- Partially Learned Knowledge
- Readings

Operational Definition of Generalization

Non-operational definition:

*A concept P is said to be **more general than** another concept Q if and only if the set of instances represented by P includes the set of instances represented by Q.*

Why isn't this an operational definition?

Because it requires to show that each instance I from a potential infinite set Q is also in the set P.

Operational definition:

*A concept P is said to be **more general than** another concept Q if and only if Q can be transformed into P by applying a sequence of generalization rules (assuming a complete set of rules).*

Generalization of a Concept: Illustration

C1:

?O1 is assistant professor
number of publications 10
is employed by George Mason University

Demonstrate that C is more general than C1

C:

?O1 is professor
number of publications ?N1
?N1 is in [10 .. 35]

Generalization of a Concept: Illustration

C1:

?O1	is	assistant professor
	number of publications	10
	is employed by	George Mason University

Generalize assistant professor to professor
Generalize 10 to [10 .. 35]
Drop “?O1 is employed by George Mason University”

C:

?O1	is	professor
	number of publications	?N1
?N1	is in	[10 .. 35]

Generalization of Two Concepts

How would you define this?

Definition:

The concept Cg is a generalization of the concepts C1 and C2 if and only if Cg is more general than C1 and Cg is more general than C2.

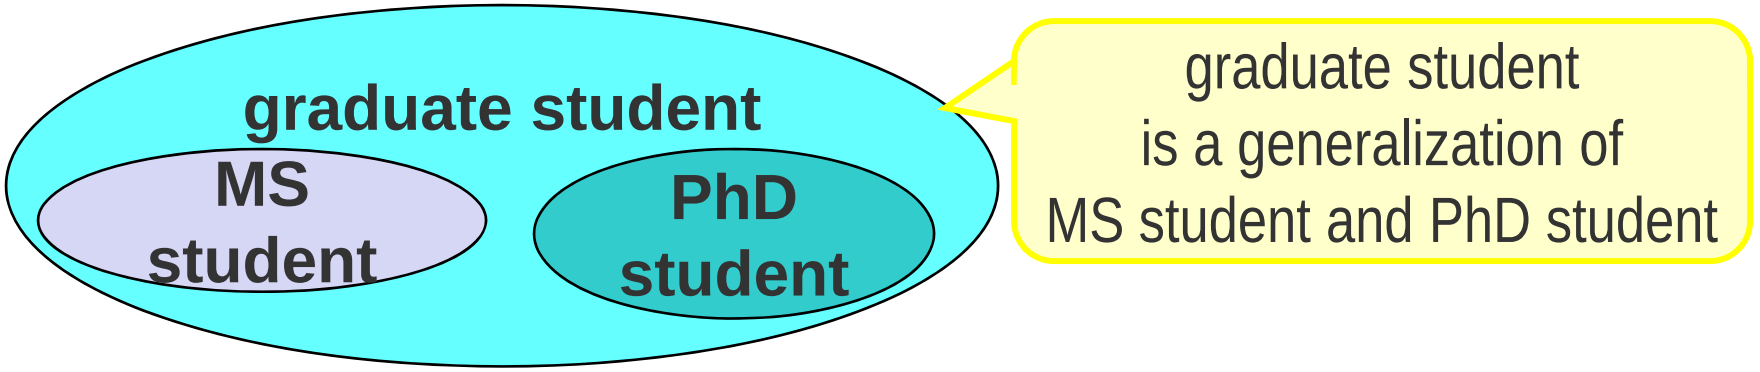

Is the above definition operational?

Which would be an operational definition?

Generalization of Two Concepts

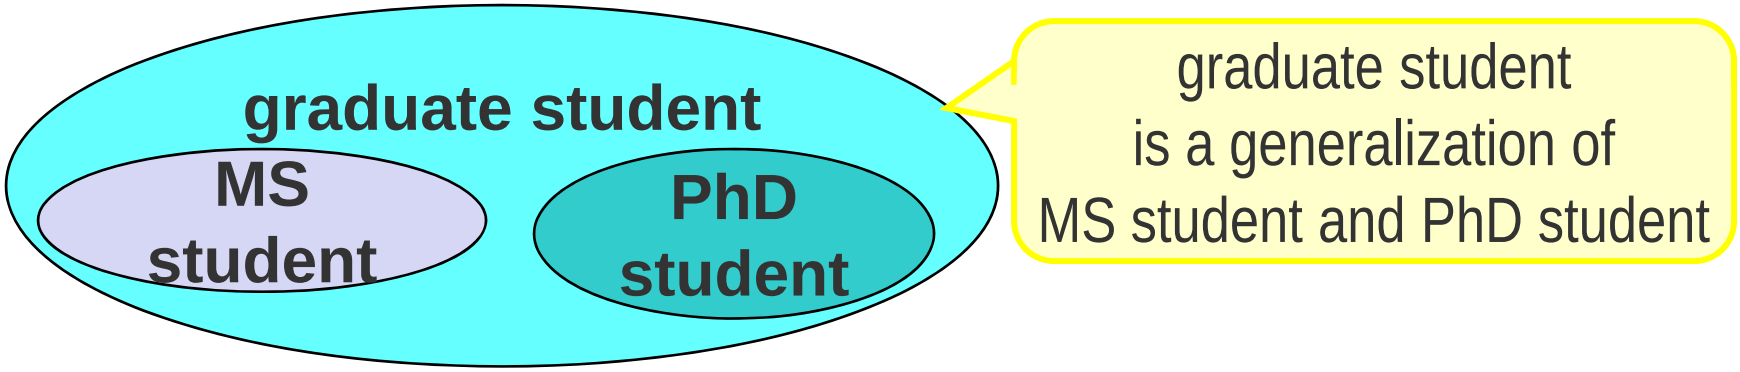

Is the above definition operational? No

Operational definition:

The concept Cg is a generalization of the concepts C1 and C2 if and only if both C1 and C2 can be transformed into Cg by applying generalization rules (assuming the existence of a complete set of rules).

Generalization of Two Concepts: Illustration

C1:

?O1 is assistant professor
number of publications 10
is employed by George Mason University

C2:

?O1 is associate professor
number of publications 35

Demonstrate that C is a generalization of C1 and C2

C:

?O1 is professor
number of publications ?N1
?N1 is in [10 .. 35]

Generalization of Two Concepts: Illustration

C1:
?O1 is assistant professor
number of publications 10
is employed by George Mason University

C2:
?O1 is associate professor
number of publications 35

Generalize assistant professor to professor
Generalize 10 to [10 .. 35]
Drop “?O1 is employed by George Mason University”

Generalize associate professor to professor
Generalize 35 to [10 .. 35]

C:
?O1 is professor
number of publications ?N1
?N1 is in [10 .. 35]

Specialization of Two Concepts

How would you define this?

Definition:

The concept Cs is a specialization of the concepts C1 and C2 if and only if Cs is less general than C1 and Cs is less general than C2.

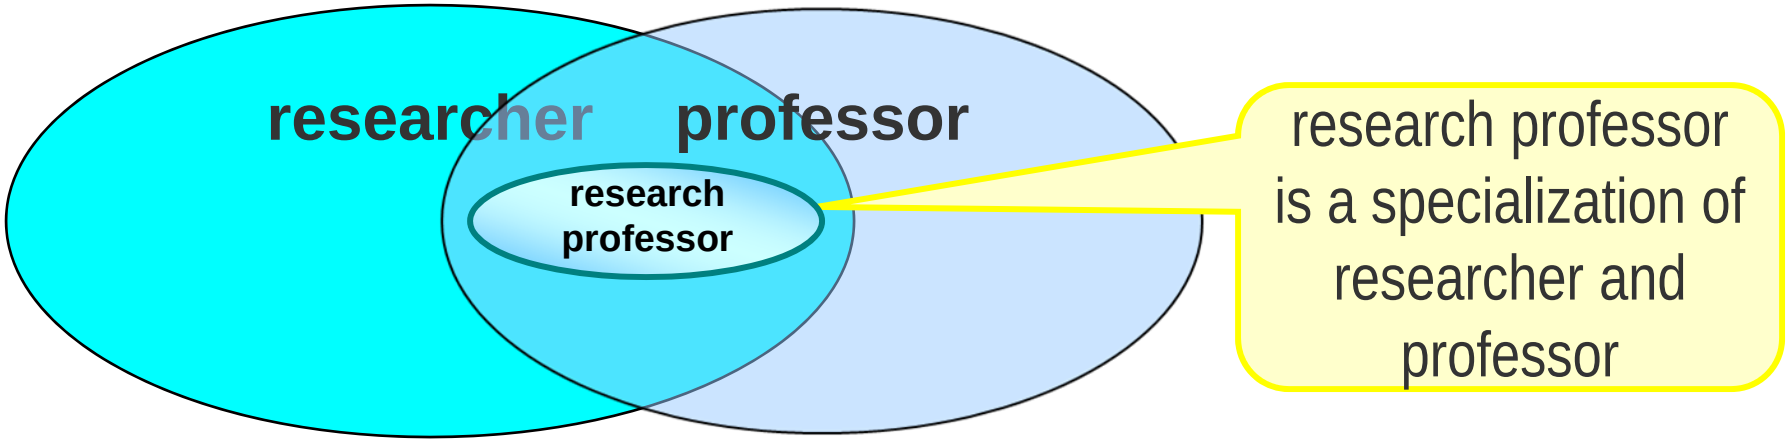

Which would be an operational definition?

Specialization of Two Concepts

Definition:

The concept Cs is a specialization of the concepts C1 and C2 if and only if Cs is less general than C1 and Cs is less general than C2.

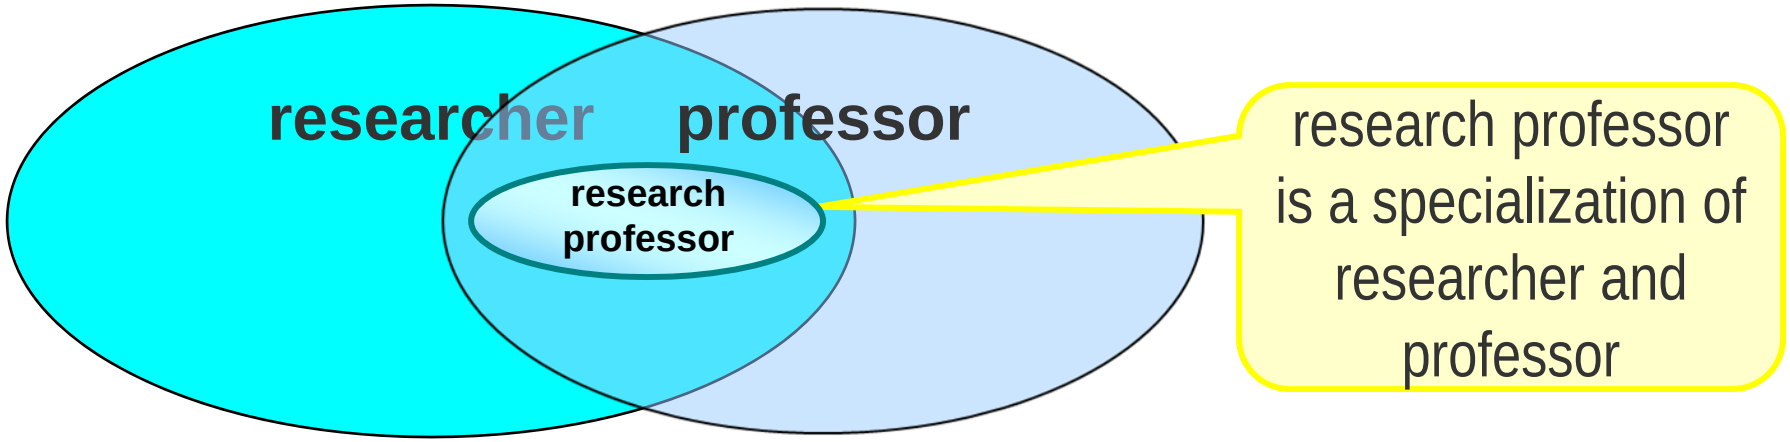

Operational definition:

The concept Cs is a specialization of the concepts C1 and C2 if and only if both C1 and C2 can be transformed into Cs by applying specialization rules (or Cs can be transformed into C1 and into C2 by applying generalization rules).
This assumes a complete set of rules.

Generalization Hierarchy (for the next examples)

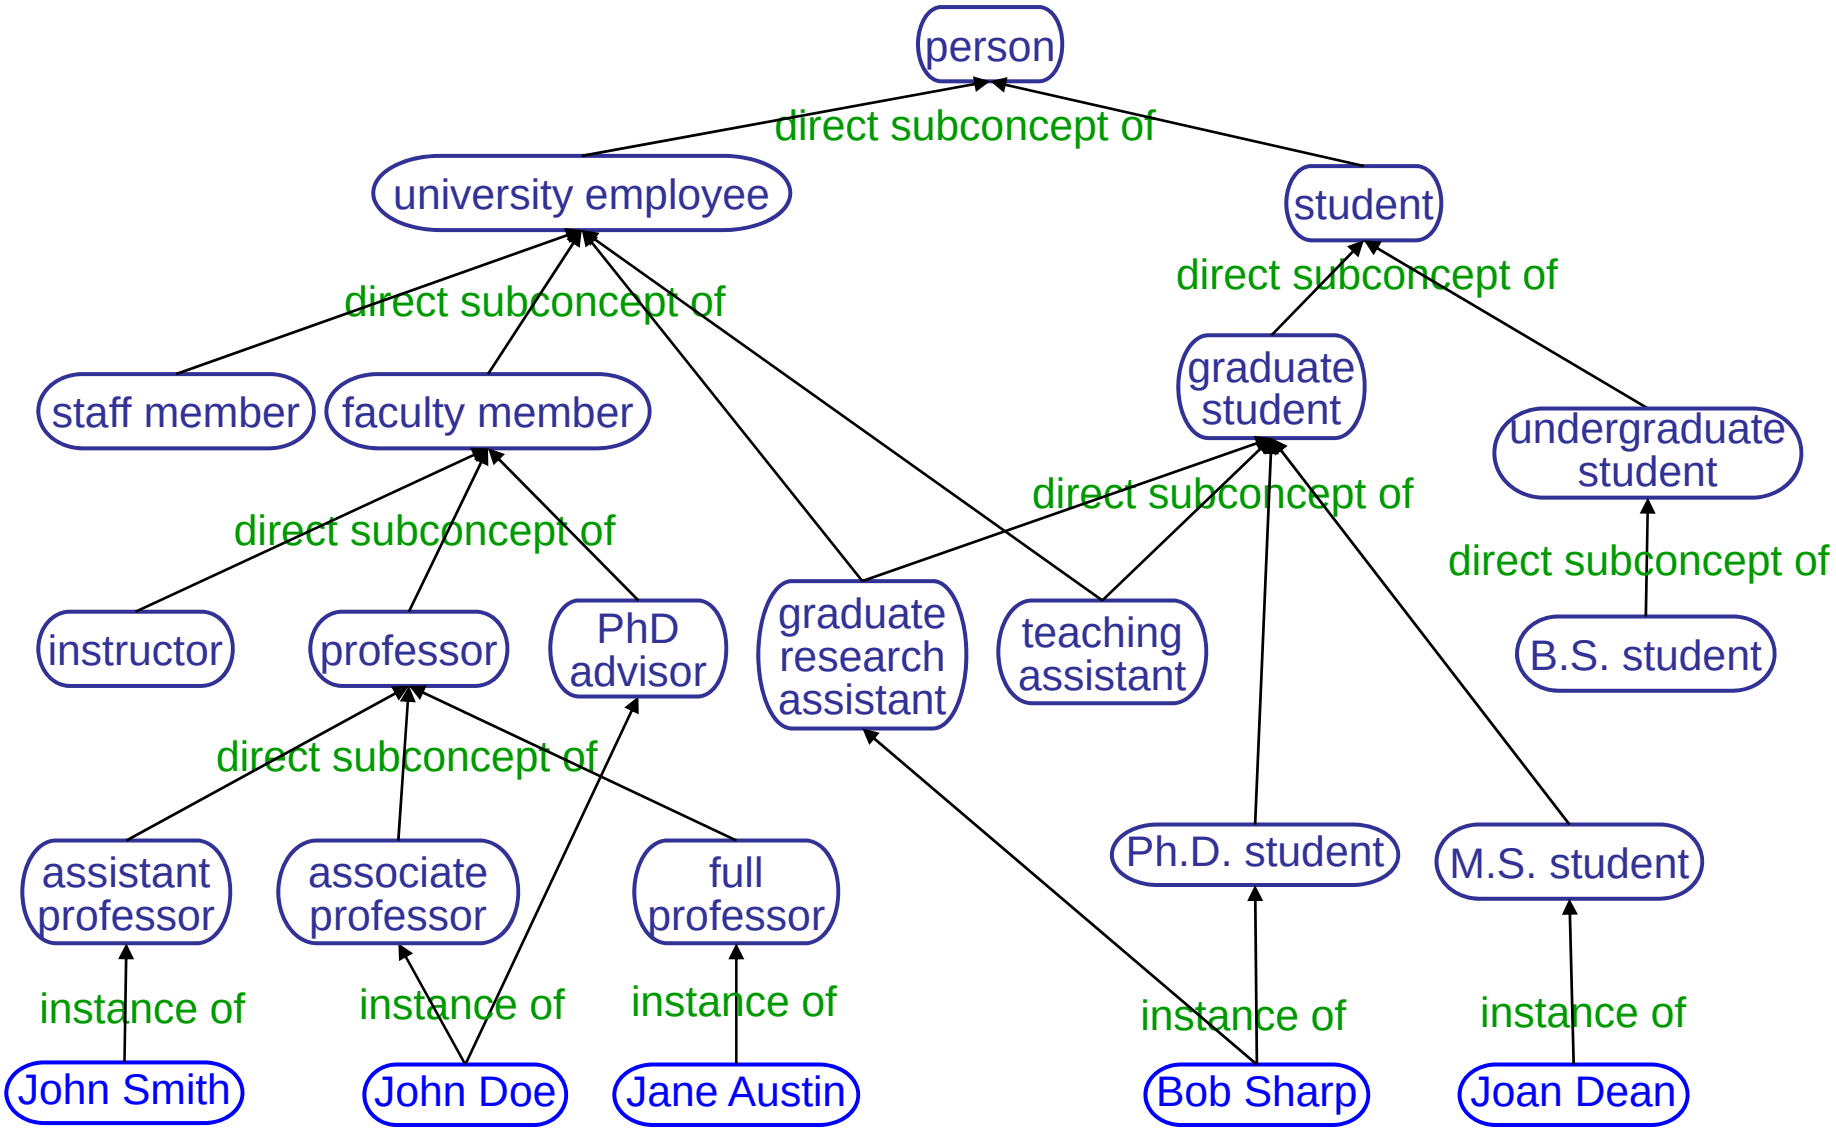

Minimally General Generalization

The concept G is a **minimally general generalization** of A and B if and only if G is a generalization of A and B, and G is not more general than any other generalization of A and B.

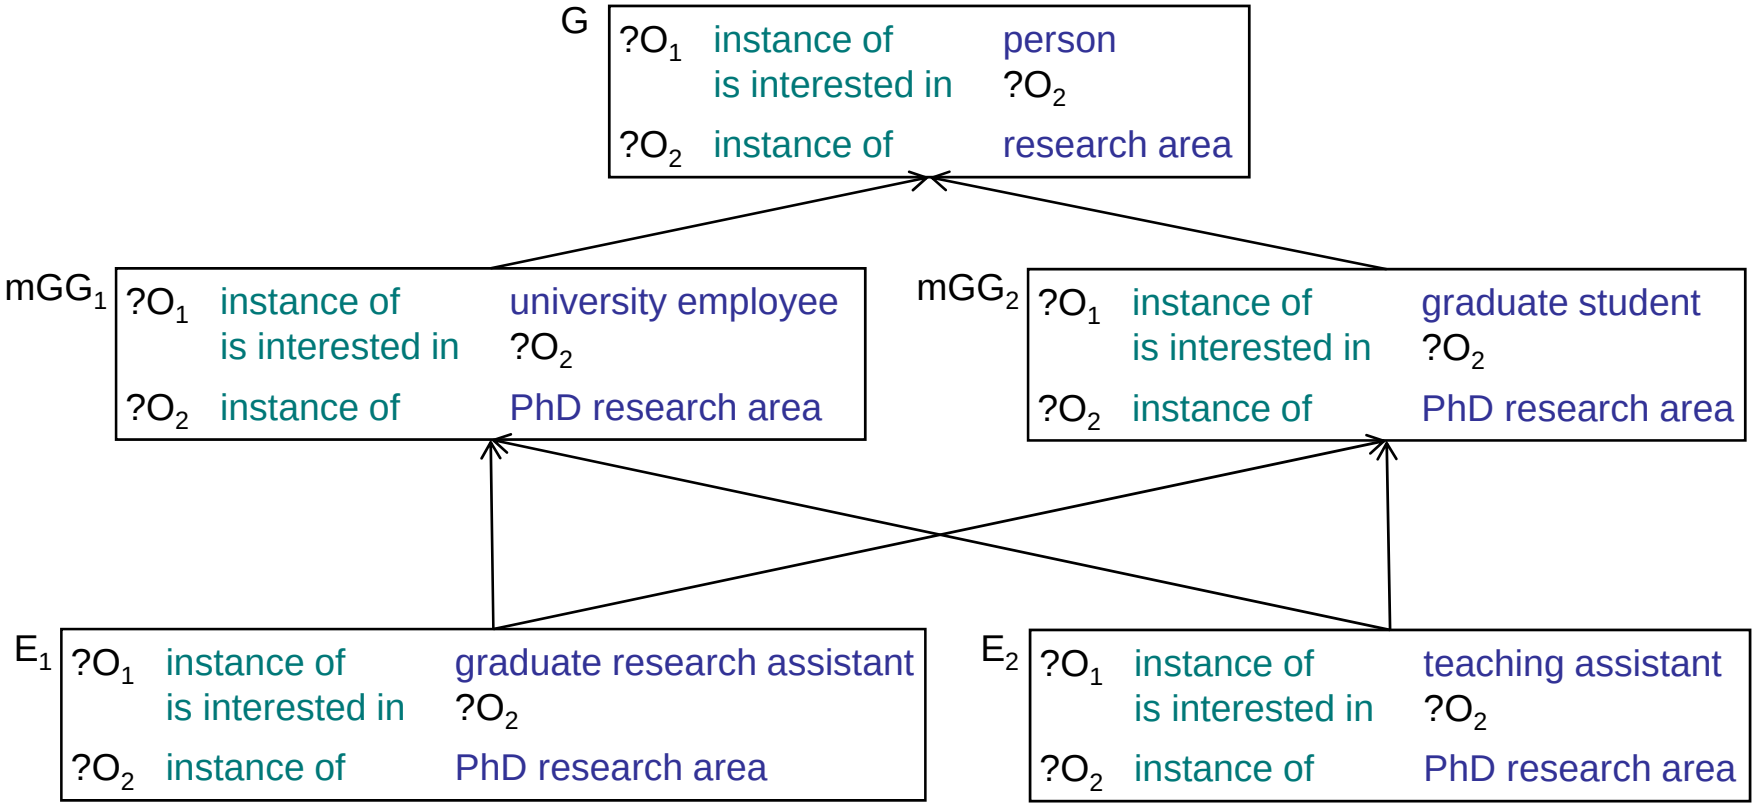

If there is only one minimally general generalization then this generalization is called **the least general generalization**.

Maximally General Specialization

The concept C is a **maximally general specialization** of two concepts A and B if and only if C is a specialization of A and B and no other specialization of A and B is more general than C.

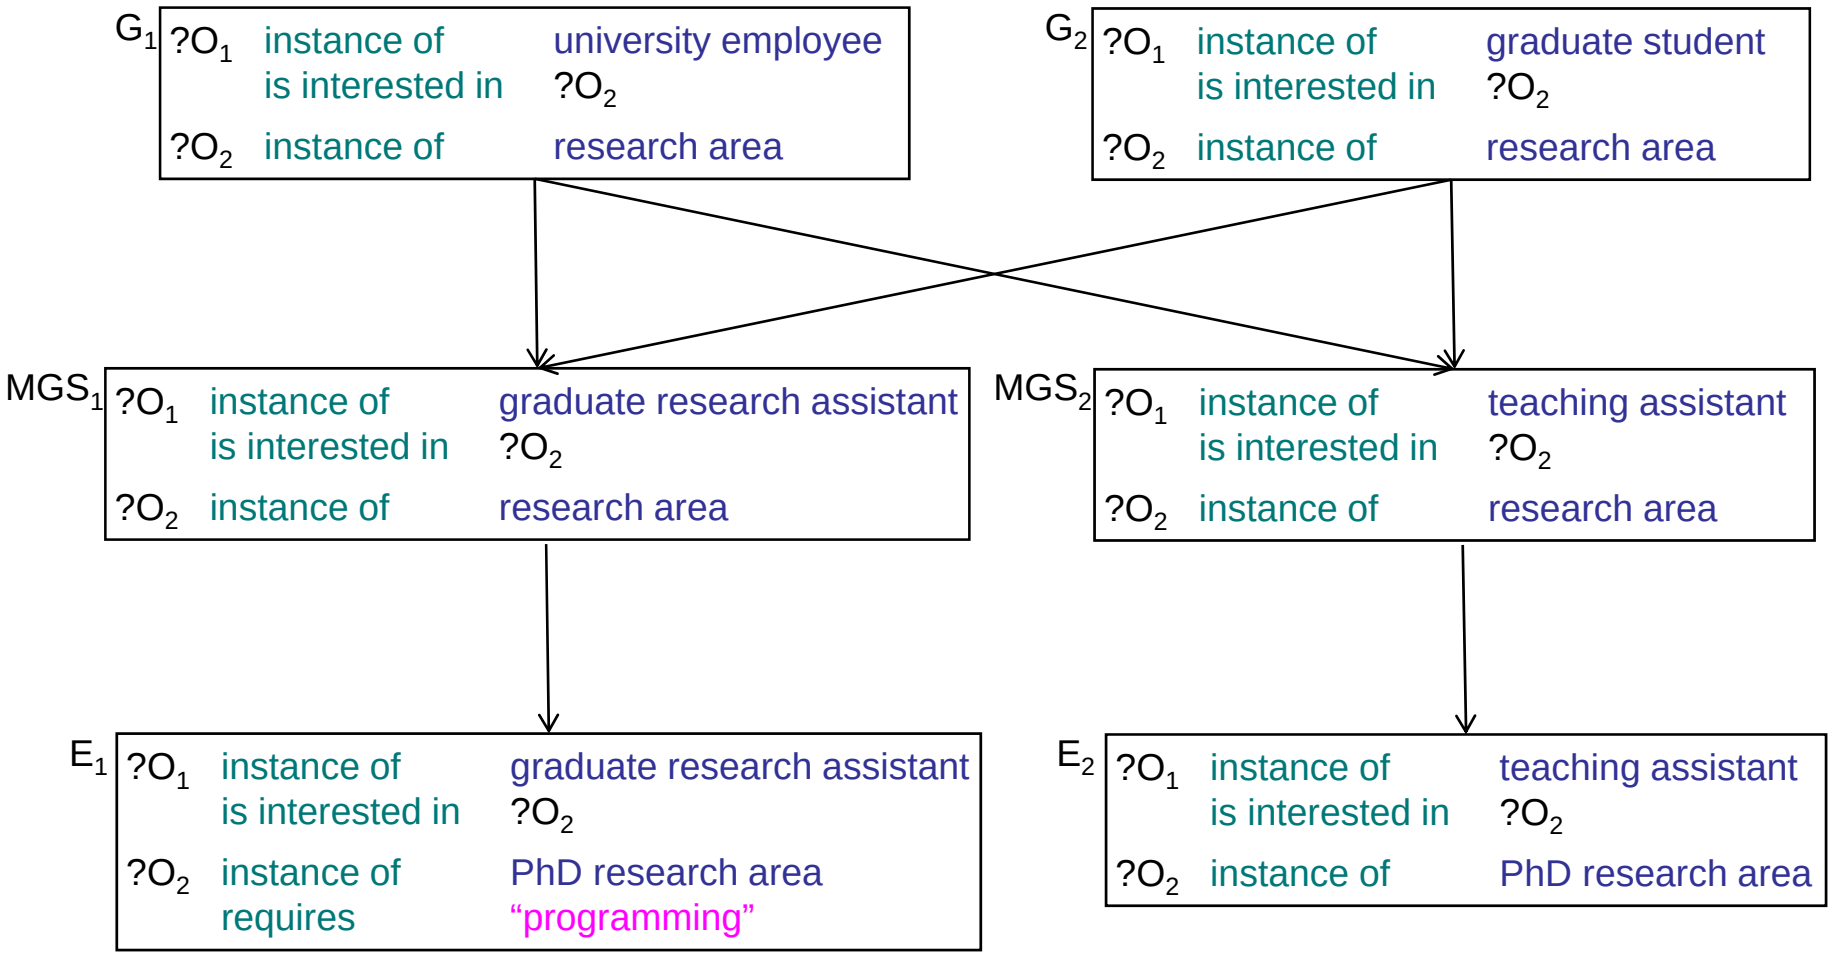

The Problem of Learning Concepts from Examples

Given

- a representation language for instances and concepts;
- a set of positive examples (E1, ..., En) of the concept;
- a set of negative (counter) examples (C1, ... , Cm) of the concept;
- background knowledge

Determine

- a concept description which is a generalization of the positive examples that does not cover any of the negative examples

Purpose of concept learning

Predict if an instance is an example of the learned concept.

Generalization of a Concept with a Positive Example

An example is the description of a cell consisting of two bodies, each with two attributes: color (yellow or green) and number of nuclei (1 or 2). The position of the bodies is not relevant because they can move inside the cell.

Find the minimally general generalizations of the concept C that cover the example:

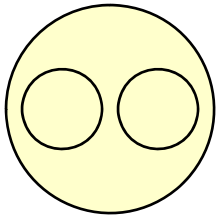

Representation of the generalization

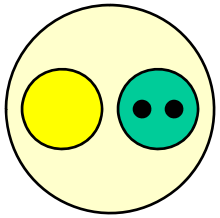

Concept C **((?x yellow) (2 green))**

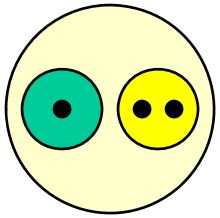

⊕ **((1 green) (2 yellow))**

Generalization of a Concept with a Positive Example

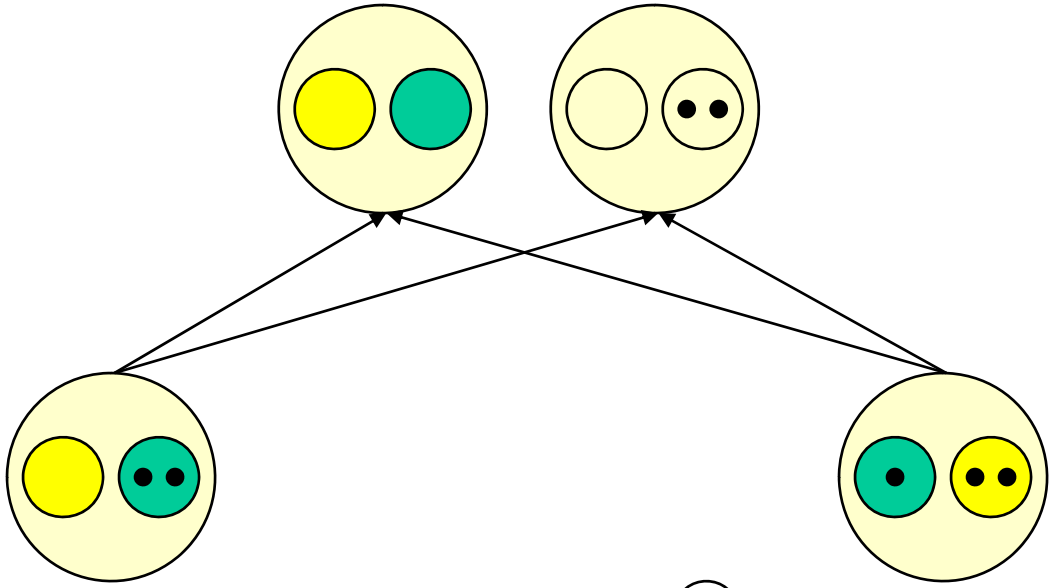

concept ((?x yellow) (2 green))

+ ((1 green) (2 yellow))

A difficulty in learning is that there are many ways in which a concept can be generalized to cover a new positive example.

Specialization of a Concept with a Negative Example

Find the maximally general specializations of the concept C
that do not cover the negative example:

Concept C **((?x yellow) (2 green))**

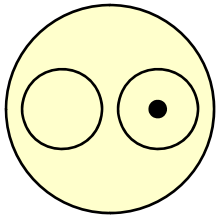

\- **((1 green) (2 yellow))**

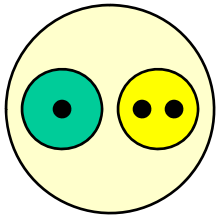

Representation of
the specialization:

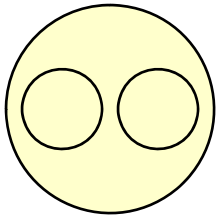

Specialization of a Concept with a Negative Example

Another difficulty in learning is that there are many ways in which a concept can be specialized to uncover a negative example.

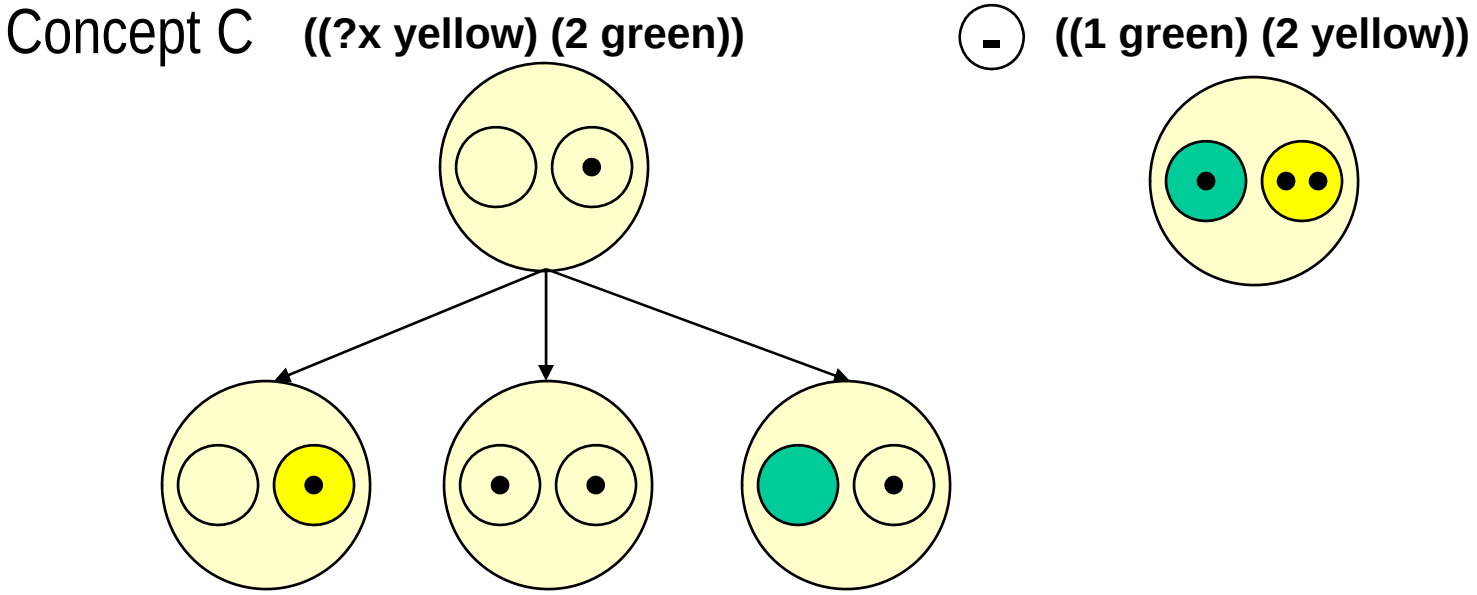

Concept Learning: Another Illustration

Positive examples:

Mark White	instance of	associate professor
	is employed by	George Mason University
Janet Rice	instance of	assistant professor
	is employed by	University of Virginia

Negative examples:

George Dean	instance of	computer technician
	is employed by	Stanford University

What concept would be learned by a cautious learner?

Concept Learning: Another Illustration

Positive examples:

Mark White	instance of	associate professor
	is employed by	George Mason University
Janet Rice	instance of	assistant professor
	is employed by	University of Virginia

Negative examples:

George Dean	instance of	computer technician
	is employed by	Stanford University

Cautious learner

Learned concept:

?O1	instance of	professor
	is employed by	?O2
?O2	instance of	state university

A professor employed by a state university.

Discussion

What could be said about the predictions of a cautious learner?

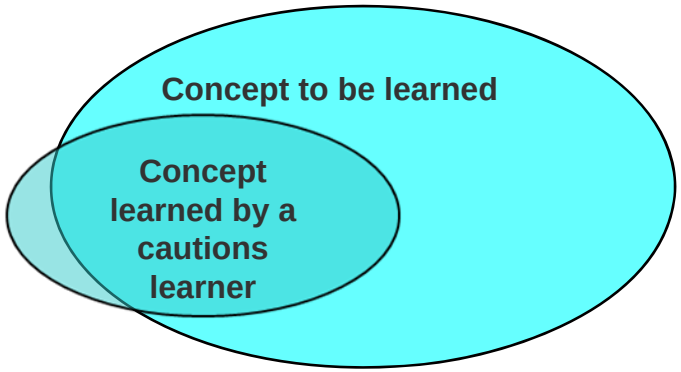

Explanation

There are many different generalizations of the positive examples that do not cover the negative examples.

For instance, a cautious learner might attempt to learn the most specific generalization.

When such a learner classifies an instance as a positive example of a concept, this classification is most likely to be correct.

However, the learner may more easily make mistakes when classifying an instance as a negative example (this type of error is called “error of omission” because some positive examples are omitted – are classified as negative examples).

Concept Learning: Yet Another Illustration

Positive examples:

Mark White	instance of	associate professor
	is employed by	George Mason University
Janet Rice	instance of	assistant professor
	is employed by	University of Virginia

Negative examples:

George Dean	instance of	computer technician
	is employed by	Stanford University

What concept would be learned by an aggressive learner?

Concept Learning: Yet Another Illustration

Positive examples:

Mark White	instance of	associate professor
	is employed by	George Mason University
Janet Rice	instance of	assistant professor
	is employed by	University of Virginia

Negative examples:

George Dean	instance of	computer technician
	is employed by	Stanford University

Aggressive learner

Learned concept:

?O1	instance of	person
	is employed by	?O2
?O2	instance of	state-university

A person employed by a state university.

What other concept might have been learned by an aggressive learner?

Discussion

What could be said about the predictions of an aggressive learner?

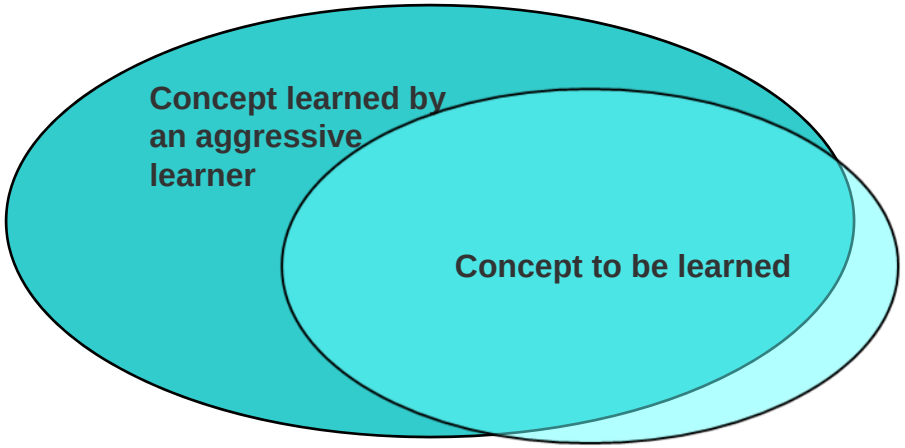

Explanation

A more aggressive learner, on the other hand, might attempt to learn the most general generalization.

When such a learner classifies an instance as a negative example of a concept, this classification is most likely to be correct.

However, the learner may more easily make mistakes when classifying an instance as a positive example (this type of error is called “error of commission” because some negative examples are committed – are classified as positive examples).

Discussion

How could one synergistically integrate a cautious learner with an aggressive learner to take advantage of their strengths to compensate for each other's weaknesses?

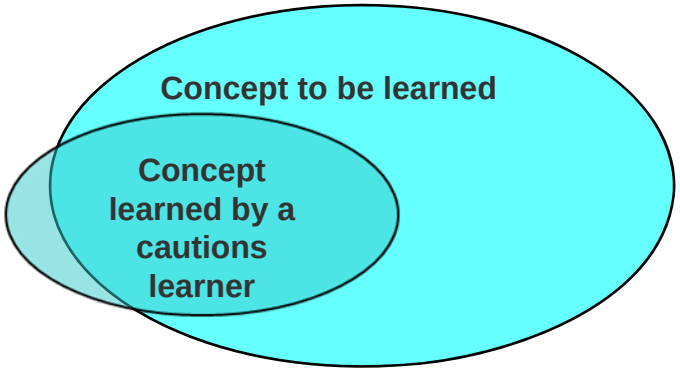

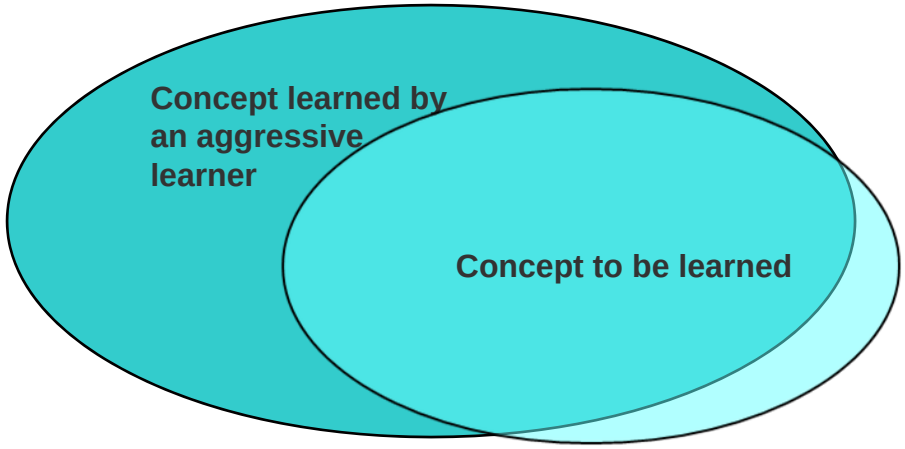

Discussion

How could one synergistically integrate a cautious learner with an aggressive learner to take advantage of their strengths to compensate for each other's weaknesses?

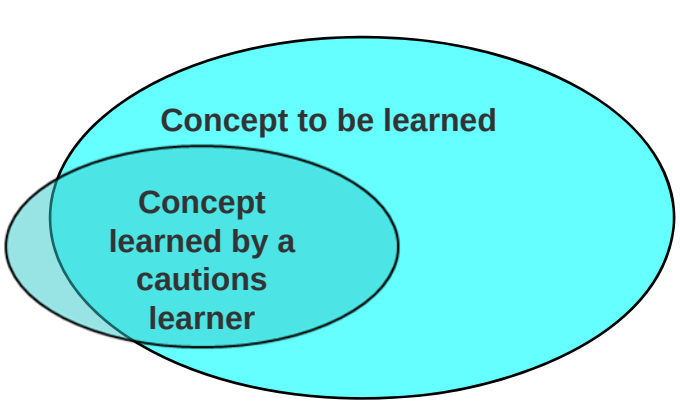

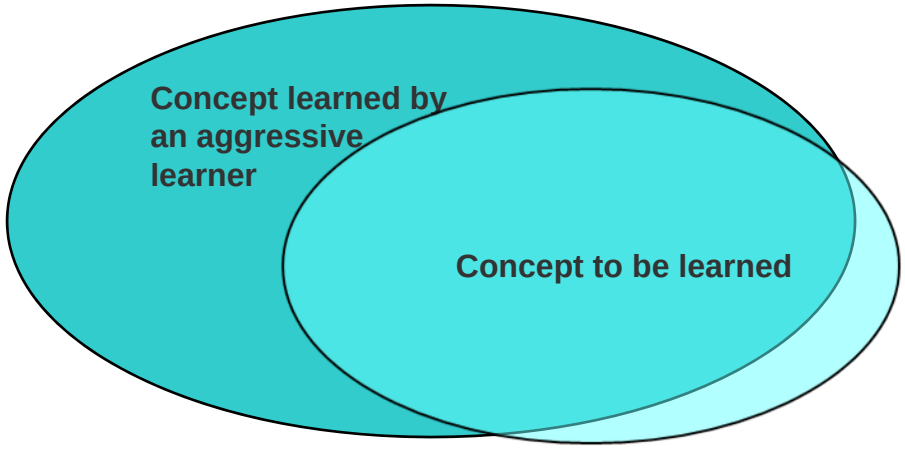

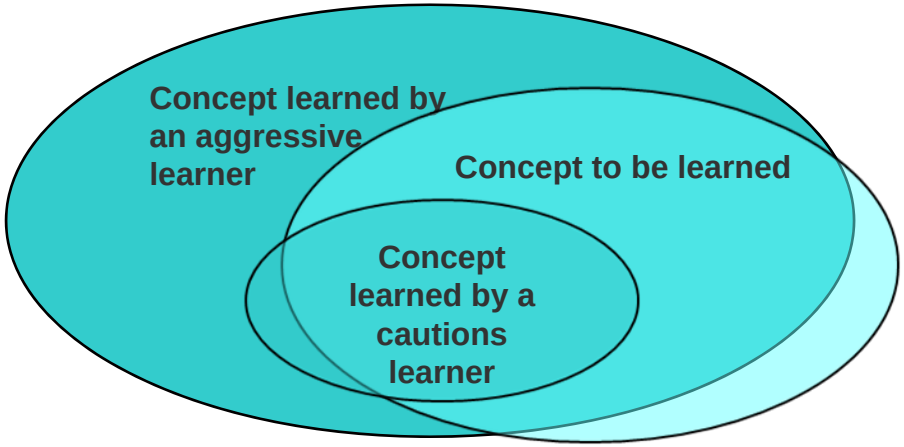

Concept Learning based on Version Spaces

Consider the examples E1, … , En in sequence.

Initialize the lower bound to the first positive example (LB=E1) and the upper bound (UB) to the most general generalization of E1.

If the next example is a positive one, then generalize LB as little as possible to cover it.

If the next example is a negative one, then specialize UB as little as possible to uncover it and to remain more general than LB.

Repeat the above two steps with the rest of examples until UB=LB.
This is the learned concept.

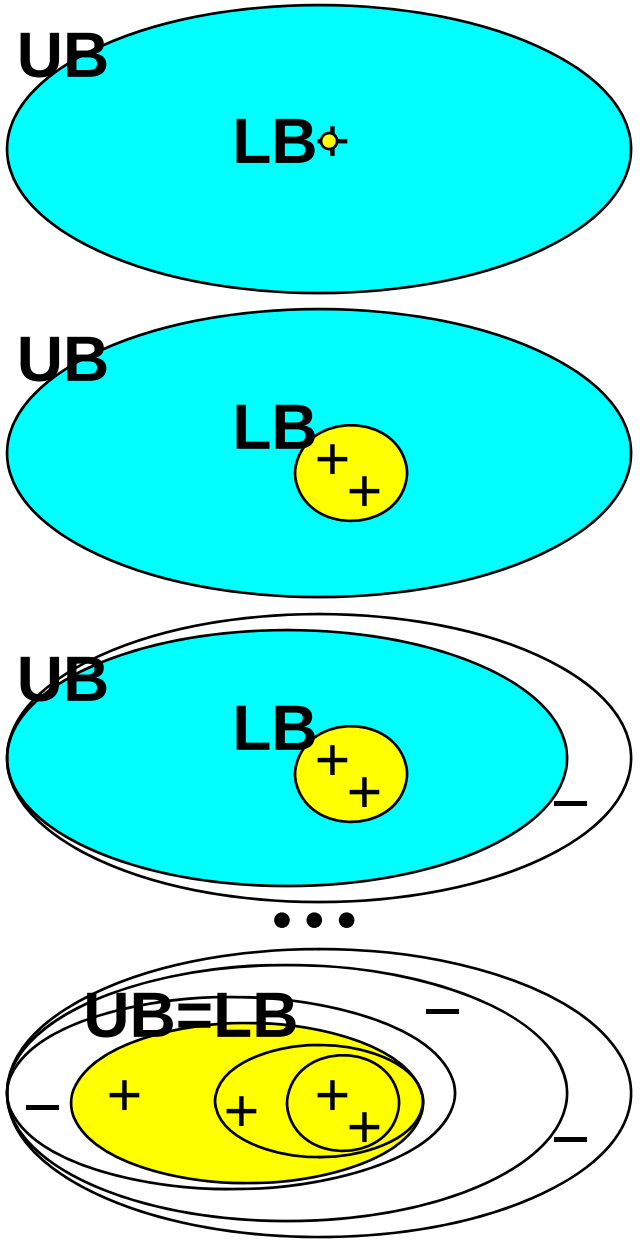

Overview

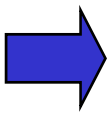

- Concept Representation
- Generalization and Specialization Rules
- Types of Generalizations and Specializations
- Partially Learned Knowledge
- Readings

Learning with Incomplete Representation Language

Which is the most general generalization of the given examples?

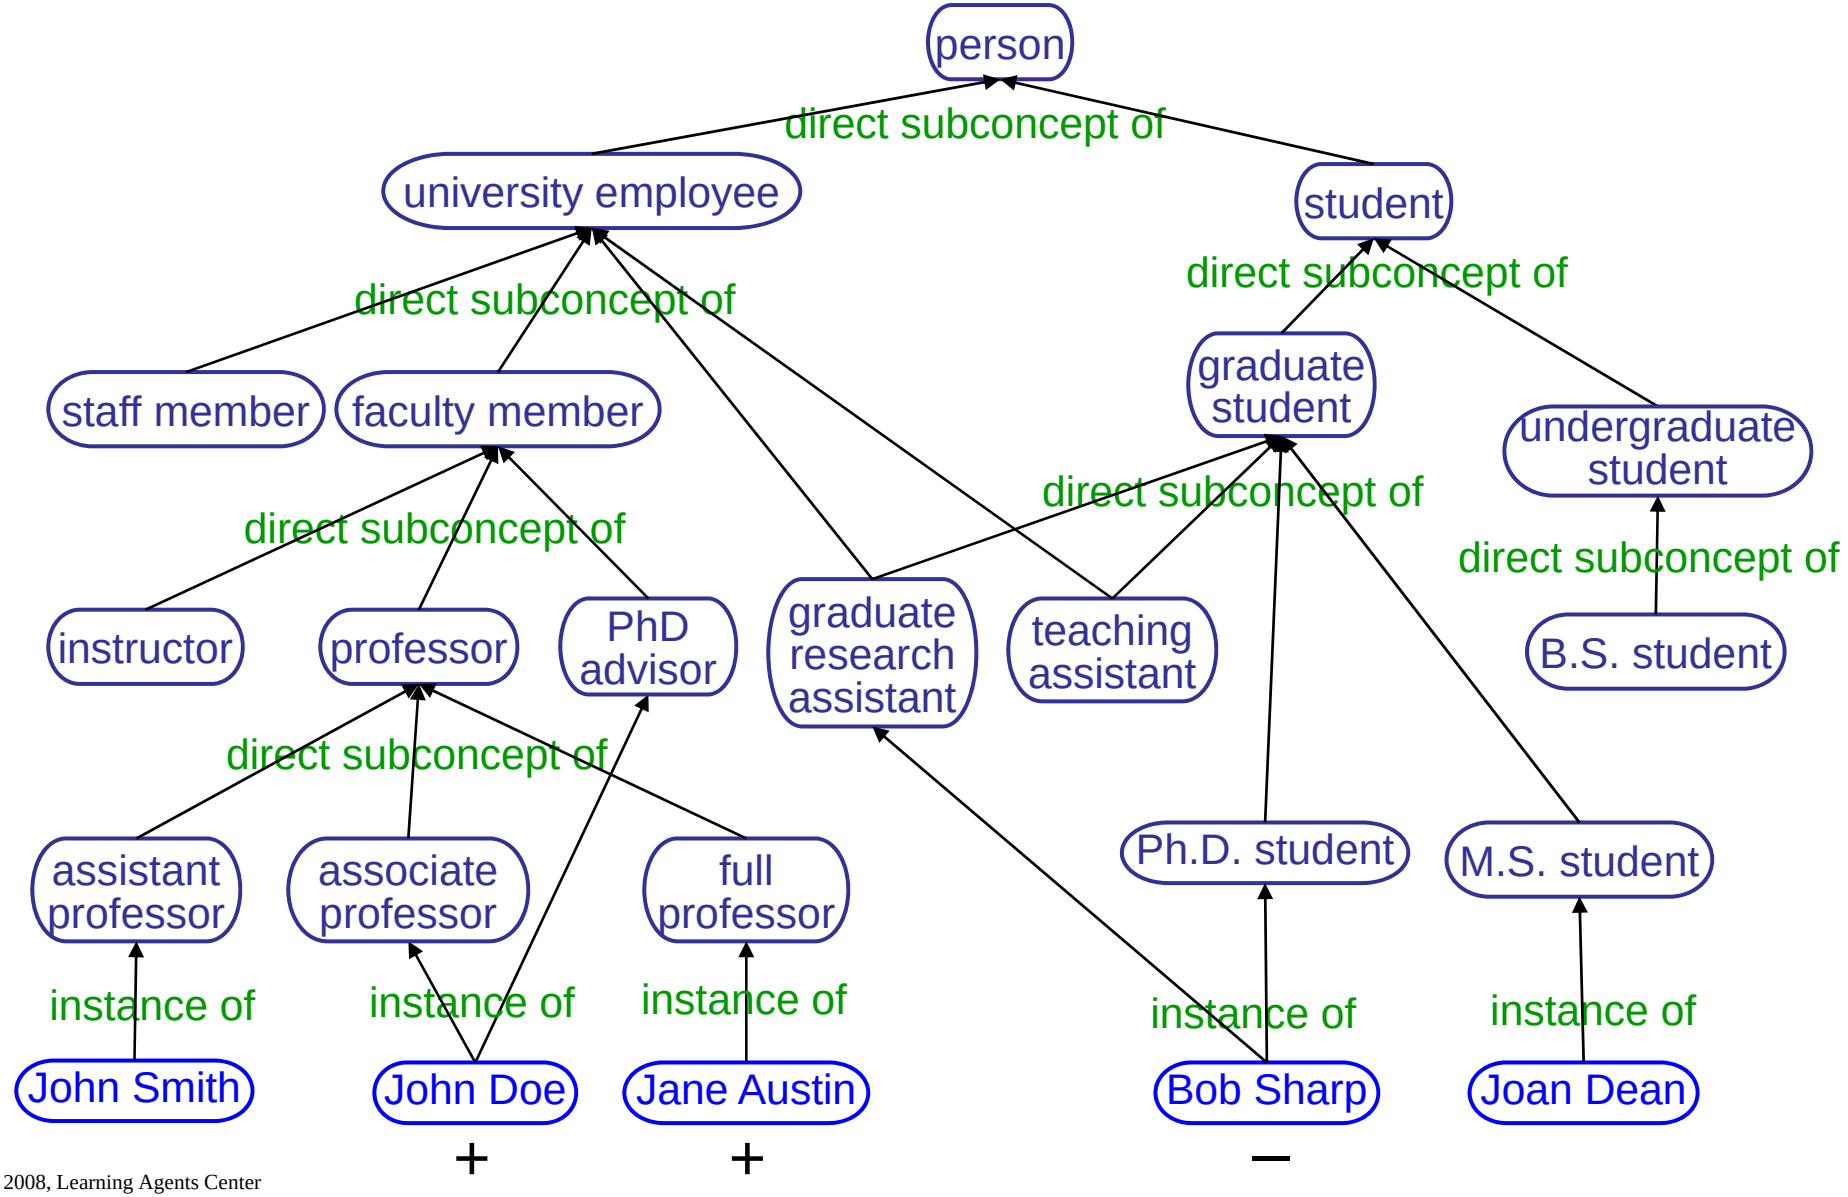

Learning with Incomplete Representation Language

Which is the most general generalization of the given examples when "graduate research assistant" is missing?

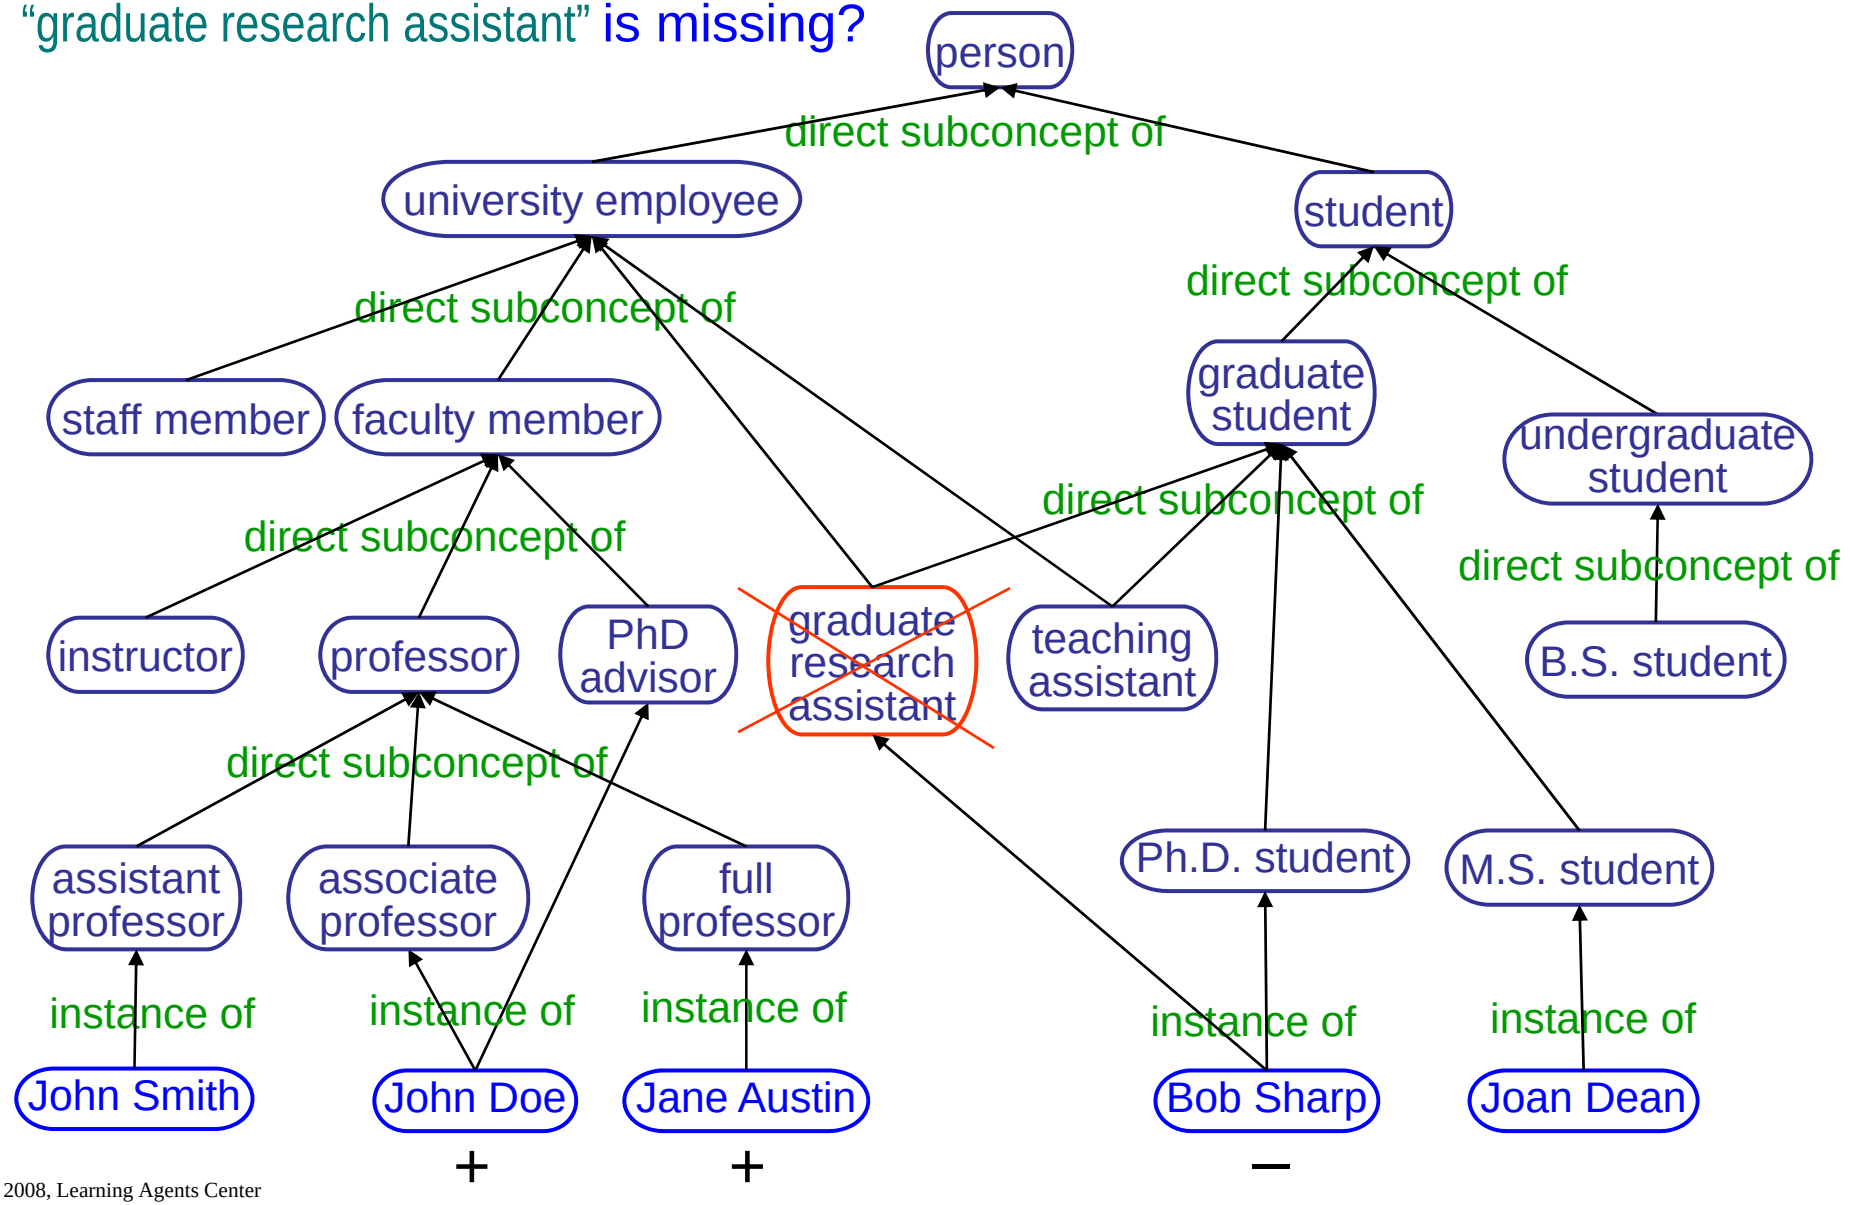

Discussion

What can be said about the relationships between the least general concepts and the most general concepts learned from examples, when the generalization language is incomplete?

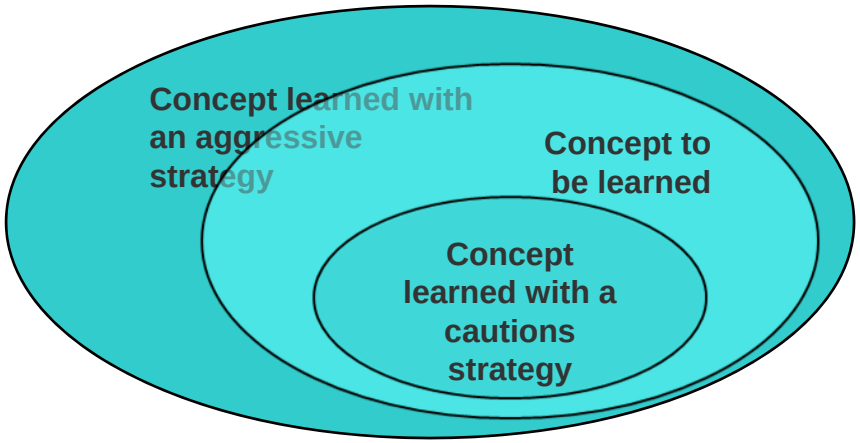

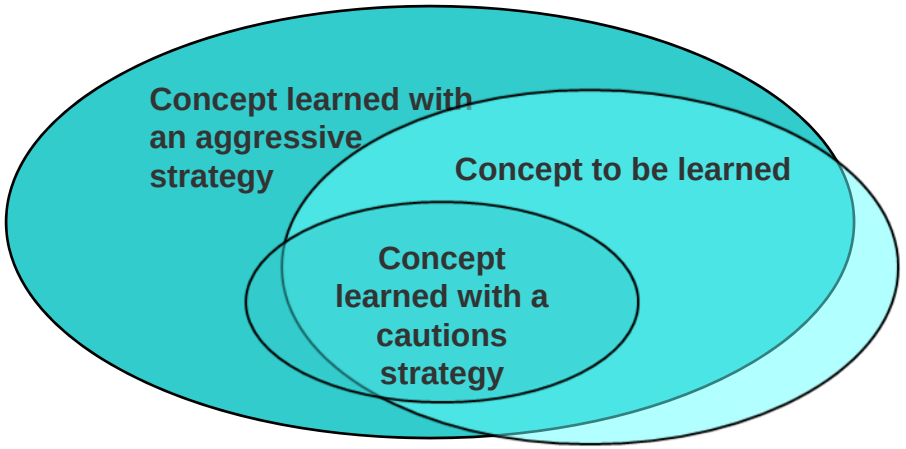

Why could we assume that the concept learned with an aggressive strategy is more general than the one learned with a cautious strategy?

Plausible Version Space

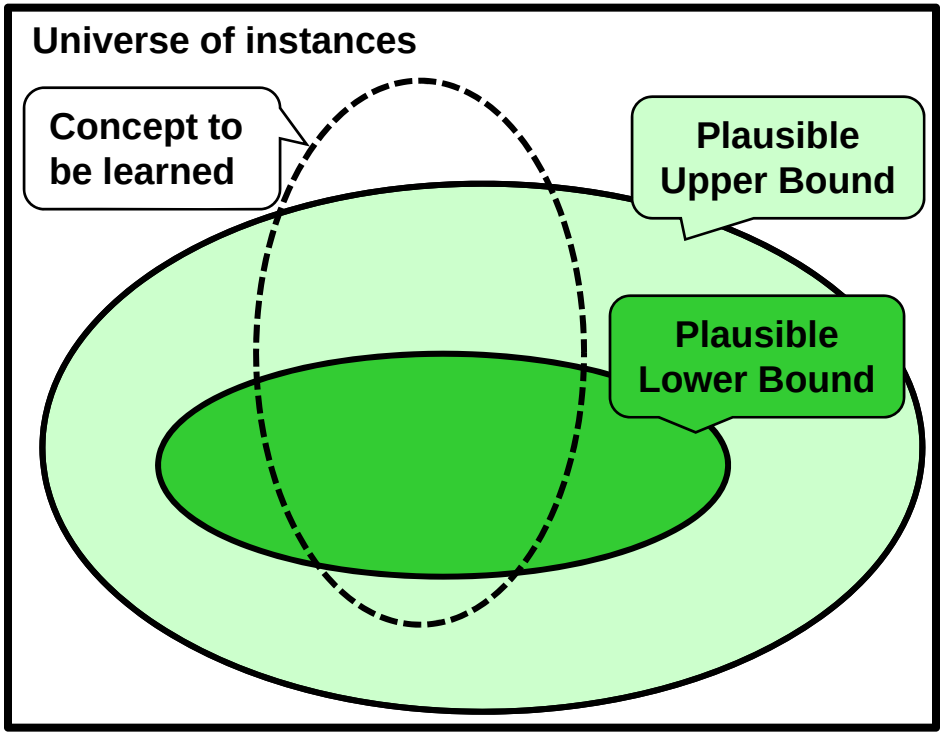

Plausible Upper Bound

$?O_1$	instance of	{faculty member, student}
	is interested in	$?O_2$
$?O_2$	instance of	research area

Plausible Lower Bound

$?O_1$	instance of	{associate professor, graduate student}
	is interested in	$?O_2$
$?O_2$	instance of	PhD research area

Which are some concepts included in this version space?

Examples/Exceptions of Partially Learned Concept

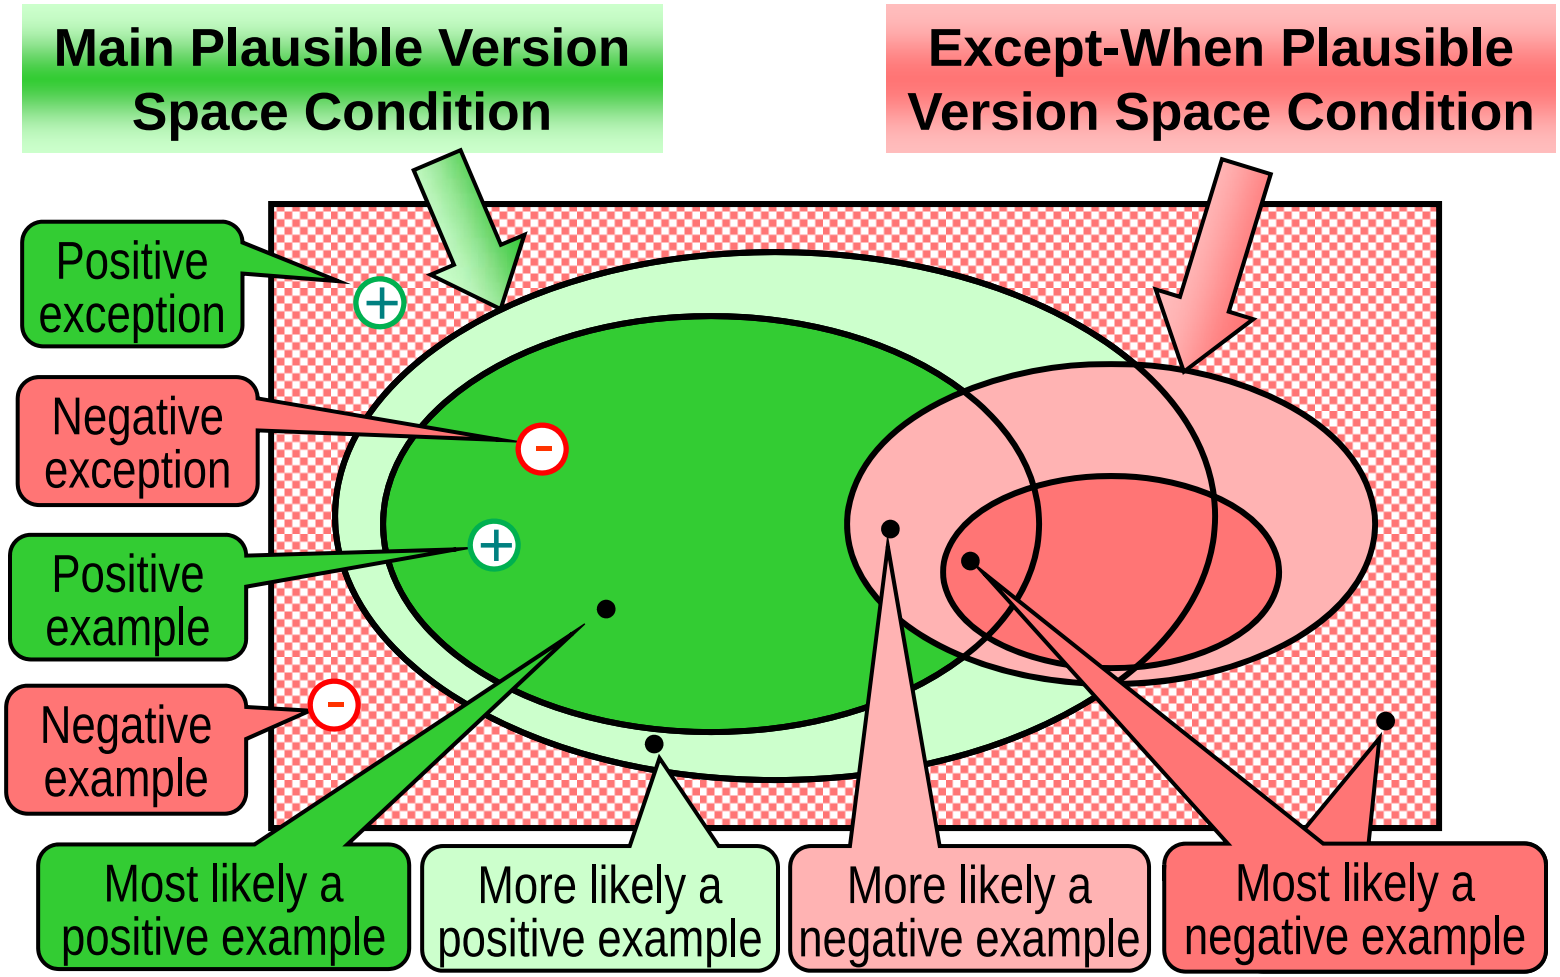

Partially Learned Feature

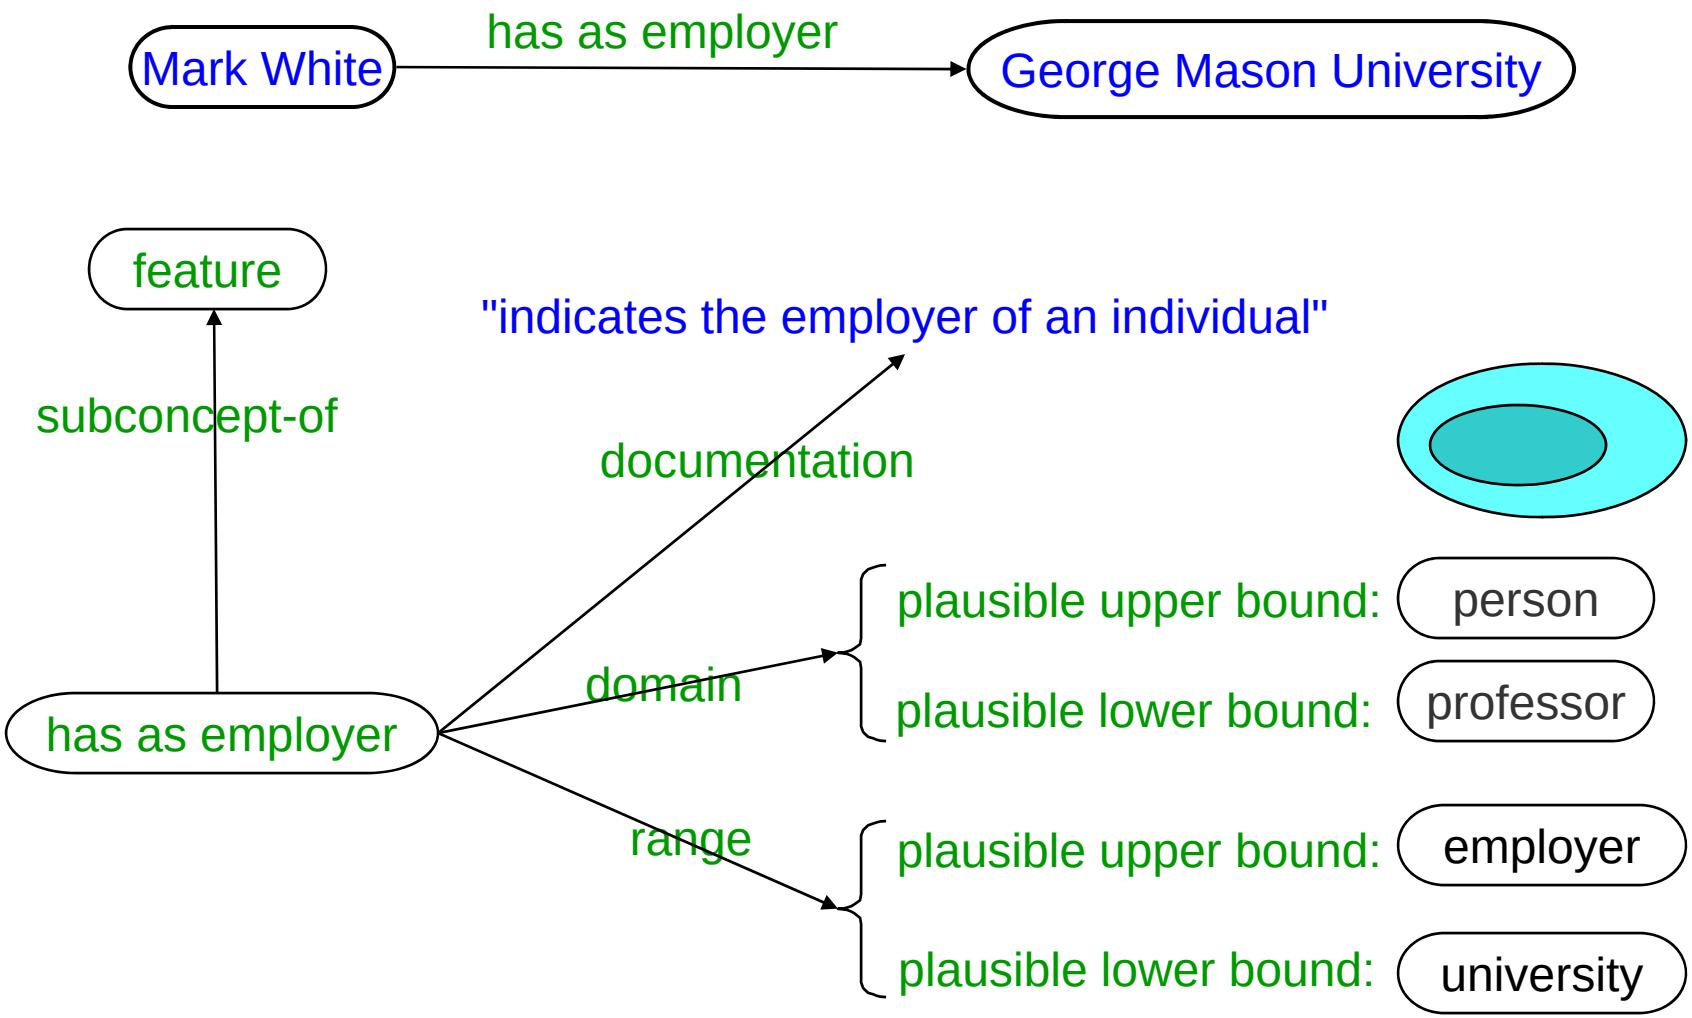

Problem Solving through Problem Reduction

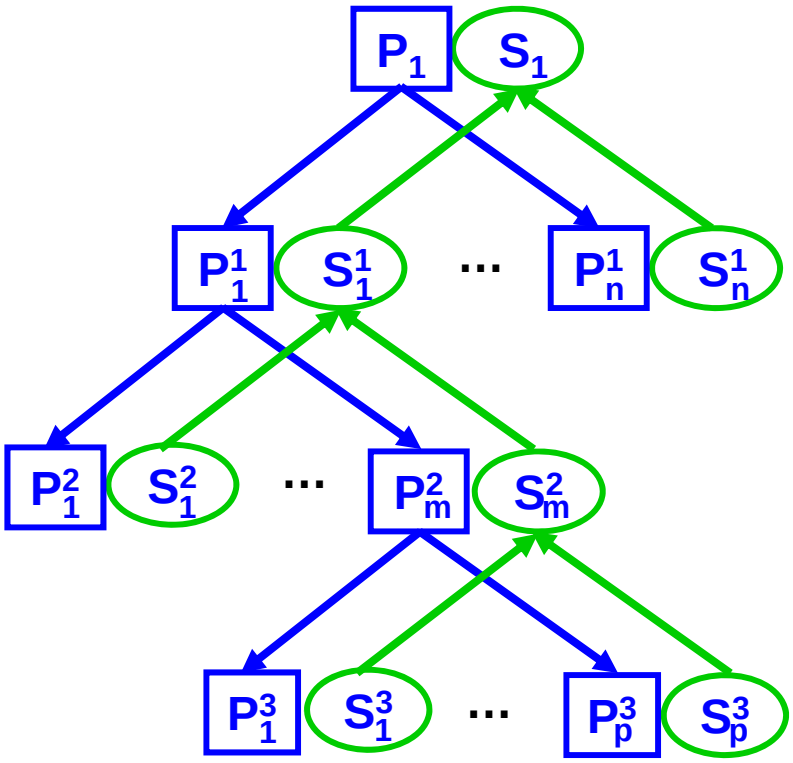

Reduction Rule

IF we need to solve <Problem>

and <Condition> is true

THEN solve
<Subproblem 1>
<Subproblem 2>
…
<Subproblem n>

Sample Reduction Rule

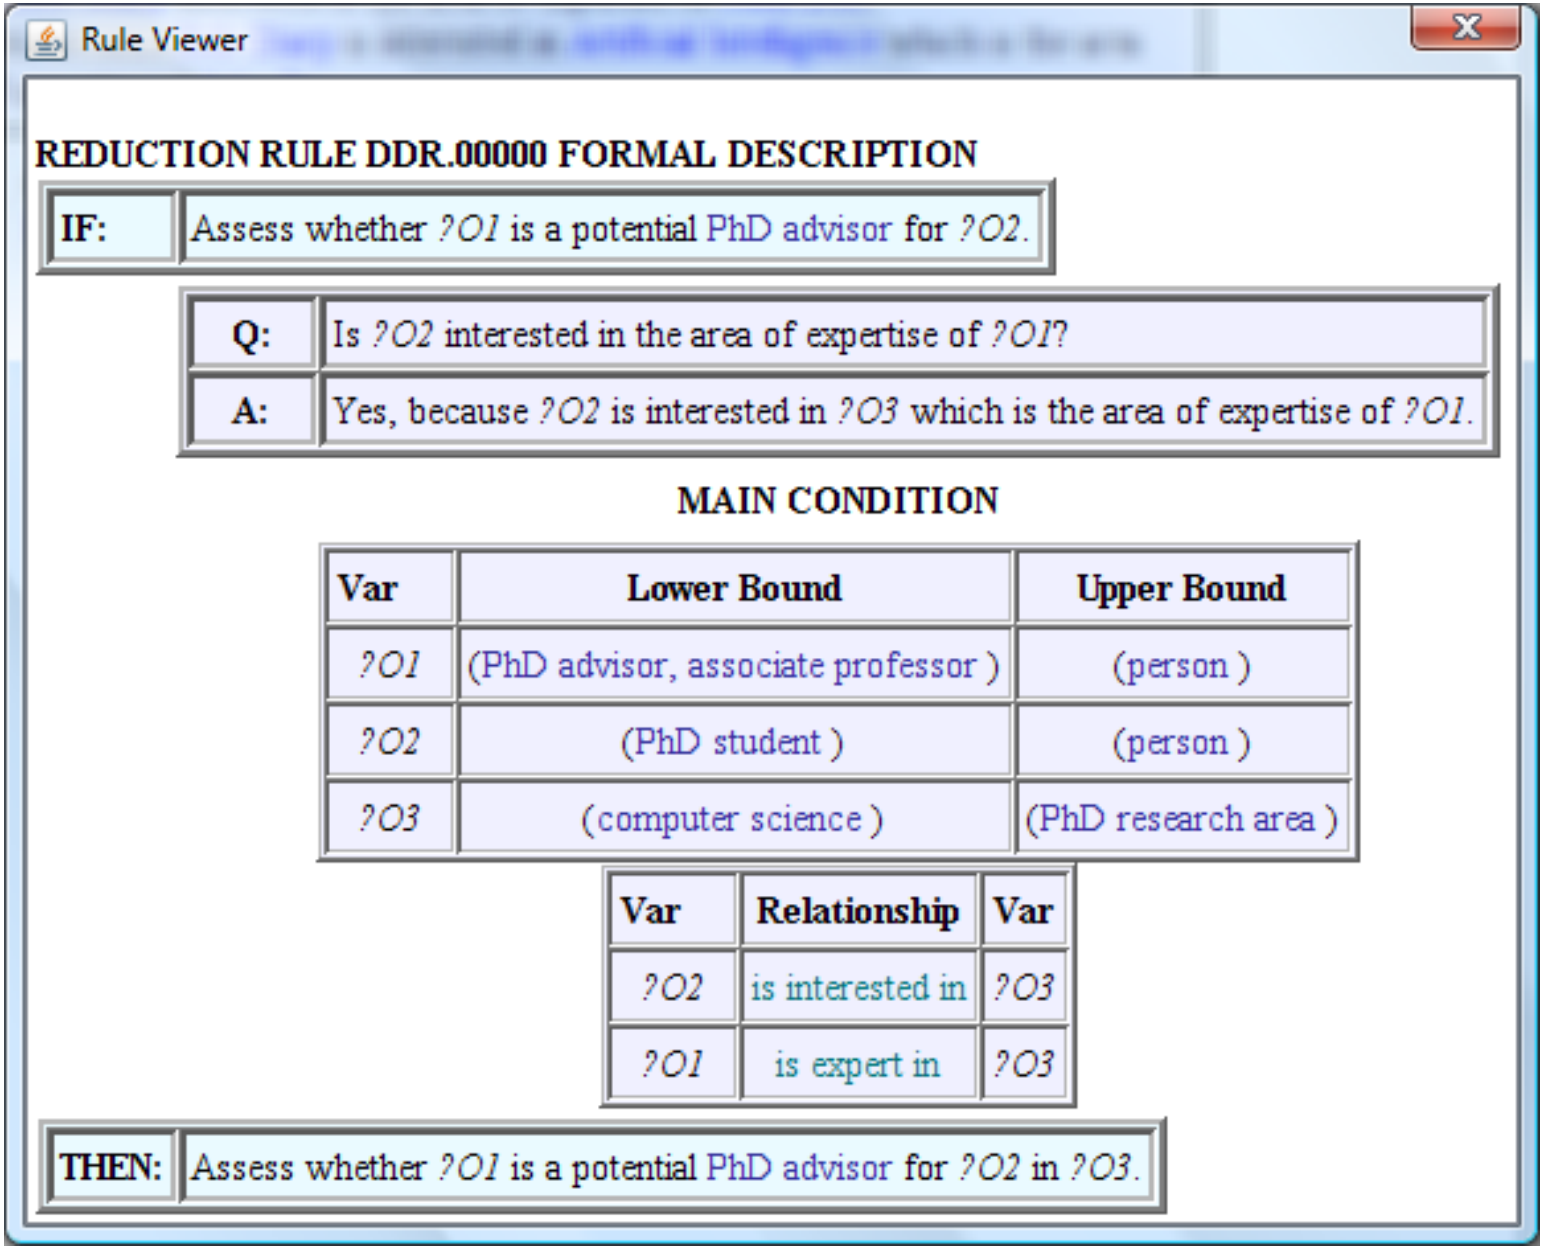

Plausible Reasoning with a Partially Learned Rule

IF <Problem>

PVS Condition

Except-When PVS Condition

THEN <Subproblem 1>
<Subproblem 2>
…
<Subproblem m>

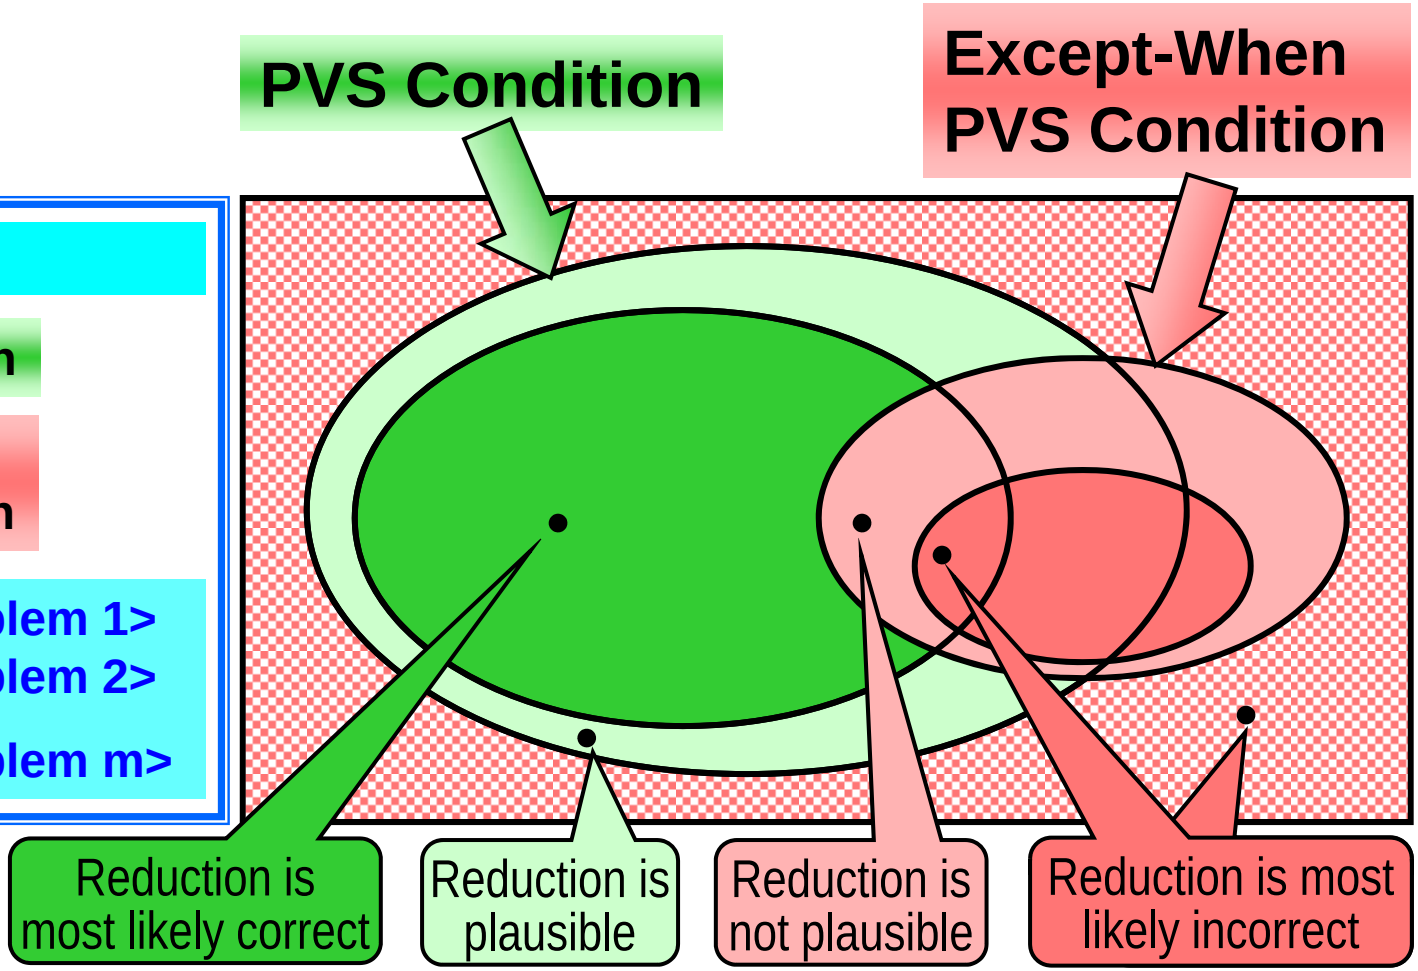

Reading

Tecuci G., Lecture Notes on Learning-based Knowledge Representation, 2008 (required).

Tecuci G., Boicu M., Learning-based Knowledge Representation, *Research Report 4,* Learning Agents Center, George Mason University, 2008 (required).

Tecuci G., *Building Intelligent Agents,* Academic Press, 1998, pp.50-65 (3.2 Generalization in the representation language of the agent) (recommended).