CS 2740 Knowledge Representation
Lecture 13

Hierarchies and inheritance.

Milos Hauskrecht
milos@cs.pitt.edu
5329 Sennott Square

Hierarchy and taxonomy

- Hierarchy or taxonomy is a natural way to view the world
  - It is used in frames (IS-A relation) and in DL
- importance of *abstraction* in remembering and reasoning
  - groups of things share properties in the world
  - we do not have to repeat representations

Example:

- Saying "elephants are mammals" is sufficient to know a lot about them

Inheritance is the result of reasoning over paths in a hierarchy

- "does *a* inherit from *b*?" is the same as "is *b* in the transitive closure of :**IS-A** (or subsumption) from *a*?"

Graphical representation of inheritance

- IS relations:
- Clyde is an Elephant, Elephant is Gray

Grey

↓

Elephant

↓

Clyde

- Reasoning with paths and conclusions they represent:
  - Transitive relations
- Transitive closure:
- Clyde is Gray, Elephant is Gray, Clyde is Elephant

Graphical representation of inheritance

- IS relations:
- Clyde is an Elephant, Elephant is Gray

Grey

↓

Elephant

↓

Clyde

- Reasoning with paths and conclusions they represent:
  - Transitive relations
- Transitive closure:
- Clyde is Gray, Elephant is Gray, Clyde is Elephant

CS 2740 Knowledge Representation M. Hauskrecht

Inheritance networks

(1) Tree structures with strict inheritance:

– as in description logics
– conclusions produced by complete transitive closure on all paths (any traversal procedure will do);
– all reachable nodes are implied

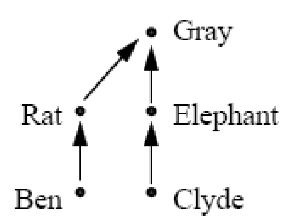

Inheritance networks

(2) Lattice structures with strict inheritance:

– as in DL's with multiple AND parents (= multiple inheritance)
– same as in trees: all conclusions you can reach by any paths are supported

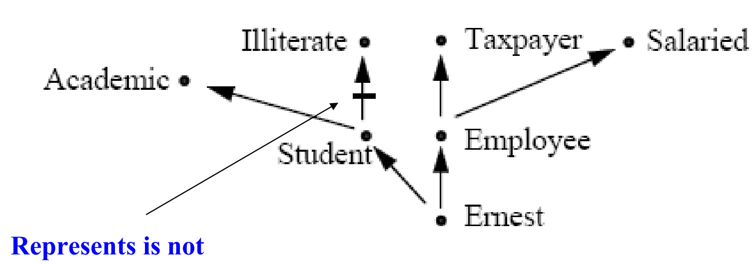

Inheritance networks

(3) Defeasible inheritance

– as in frame systems
– inherited properties do not always hold, and can be *overridden* (defeated)
– conclusions determined by searching upward from "focus node" and selecting **first version of property** you want

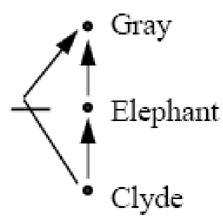

Elephants are gray but Clyde is not

Inheritance networks

(3) Defeasible inheritance

– as in frame systems
– inherited properties do not always hold, and can be *overridden* (defeated)
– conclusions determined by searching upward from "focus node" and selectin first version of property you want

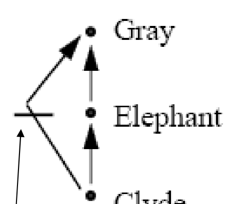

Elephants are gray but Clyde is not

Problem:

- ambiguity

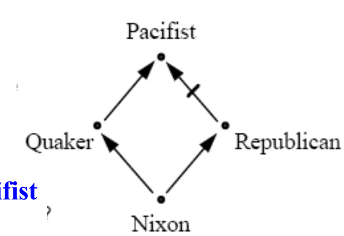

Is Nixon a pacifist or not ?

Inheritance networks

(3) Defeasible inheritance

- links have *polarity* (positive or negative)
- use **shortest path heuristic** to determine which polarity counts
- as a result, not all paths count in generating conclusions some are "preempted" but some are "admissible"
- think of paths as *arguments* in support of conclusions

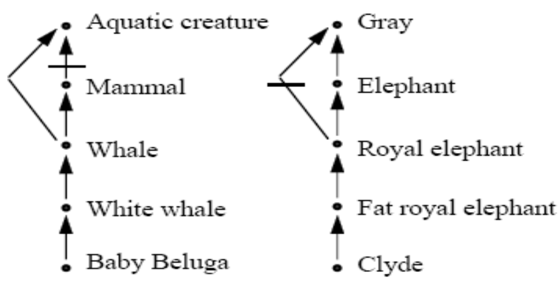

Problems with the shortest path

(3) Defeasible inheritance

Problem 1: redundant edges

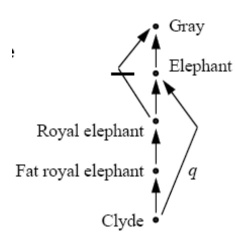

Problem 2: conclusion is changed
by adding additional categories, edges

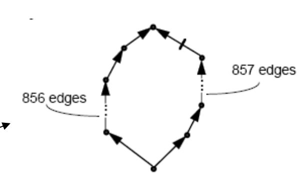

Addition of 2 edges switches the conclusion

Formal: Inheritance hierarchy

An **inheritance hierarchy** G = <*V*,*E*> is a directed, acyclic graph (DAG) with positive and negative edges, intended to denote "(normally) is-a" and "(normally) is-not-a", respectively.

- positive edges are written $a \bullet x$
- negative edges are written $a \bullet \neg x$

A sequence of edges is a **path**:

- **a positive path** is a sequence of one or more positive edges $a \bullet \ldots \bullet x$
- **a negative path** is a sequence of positive edges followed by a single negative
- edge $a \bullet \ldots \bullet v \bullet \neg x$

Note: there are no paths with more than 1 negative edge.

- Also: there might be 0 positive edges.
- A path (or argument) supports a conclusion:
  - $a \bullet \ldots \bullet x$ supports the conclusion "*a* is an *x*"
  - $a \bullet \ldots \bullet v \bullet \neg x$ supports "*a* is not an *x*"

Note: a conclusion may be supported by many arguments

However: not all arguments are equally believable...

Support and Admissibility

G **supports a path** $a \bullet s1 \bullet \ldots \bullet sn \bullet (\neg)x$ if the corresponding set of edges $\{a \bullet s1 \bullet \ldots \bullet sn \bullet (\neg)x\}$ is in E, and the path is admissible.

The **hierarchy G supports a conclusion *a* is *x*** (or *a* is not *x*) if it supports some corresponding path

A **path is admissible** if every edge in it is admissible.

An **edge $v \bullet x$ is admissible** in G wrt *a* if there is a positive path $a \bullet s1 \bullet \ldots \bullet sn \bullet v$ ($n \geq 0$) in E and

1. each edge in $a \bullet s1 \bullet \ldots \bullet sn \bullet v$ is admissible in G wrt *a* (recursively);
2. no edge in $a \bullet s1 \bullet \ldots \bullet sn \bullet v$ is redundant in G wrt *a* (see below);
3. no intermediate node *a,s1,…,sn* is a preemptor of $v \bullet x$ wrt *a* (see below).

A negative edge $v \bullet \neg x$ is handled analogously.

Preemptor

A node y along path $a \bullet ... y ... \bullet v$ is a **preemptor of the edge** $v \bullet x$ wrt a

- if $y \bullet \neg x \in E$ (or analogously for $v \bullet \neg x$)

the node Whale preempts the negative edge from Mammal to Aquatic creature wrt both Whale and Blue whale

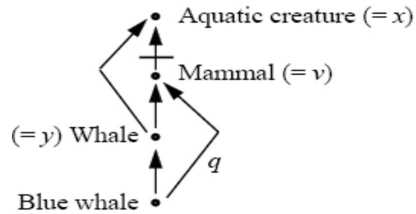

Redundancy

A positive edge $b \bullet w$ is redundant in G wrt node a if there is some positive path $b \bullet t1 ... tm \bullet w \in E$ ($m \geq 1$), for which

1. each edge in $b \bullet t1 ... tm$ is admissible in G wrt a;
2. there are no c and i such that $c \bullet \neg ti$ is admissible in G wrt a;
3. there is no c such that $c \bullet \neg$ w is admissible in G wrt a.

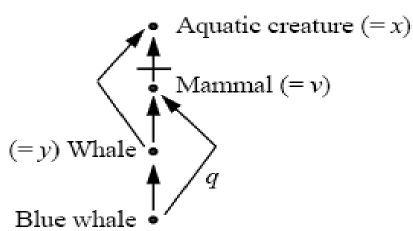

- The definition for a negative edge $b \bullet \neg w$ is analogous

Credulous extensions

G is ***a*-connected** iff for every node x in G, there is a path from a to x, and for every edge $v \bullet (\neg)x$ in G, there is a *positive* path from a to v.

- In other words, every node and edge is reachable from a

G is (potentially) **ambiguous** wrt a node a if there is some node $x \in$ V

- such that both $a \bullet s1 \ldots sn \bullet x$ and $a \bullet t1 \ldots tm \bullet \neg x$ are paths in G

A **credulous extension** of G wrt node a is a maximal unambiguous a-connected subhierarchy of G wrt a

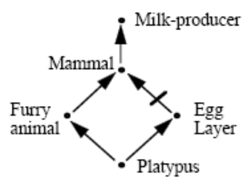

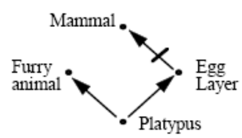

Extension 1

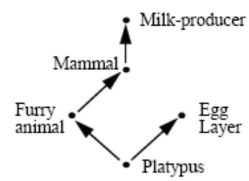

Extension 2

Preferred extensions

Credulous extensions do not incorporate any notion of admissibility or preemption.

Let X and Y be credulous extensions of G wrt node a. ***X* is preferred to *Y*** iff there are nodes v and x such that:

- X and Y agree on all edges whose endpoints precede v topologically,
- there is an edge $v \bullet x$ (or $v \bullet \neg x$) that is *inadmissible* in G,
- this edge is in Y, but not in X

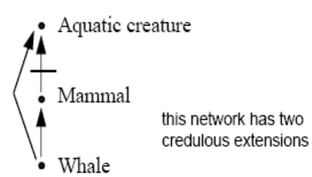

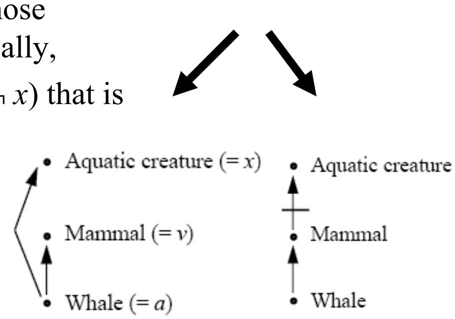

Preferred extensions

Credulous extensions do not incorporate any notion of admissibility or preemption.

Let X and Y be credulous extensions of G wrt node a. ***X* is preferred to *Y*** iff there are nodes v and x such that:

- X and Y agree on all edges whose endpoints precede v topologically,
- there is an edge $v \bullet x$ (or $v \bullet \neg x$) that is *inadmissible* in G,
- this edge is in Y, but not in X.

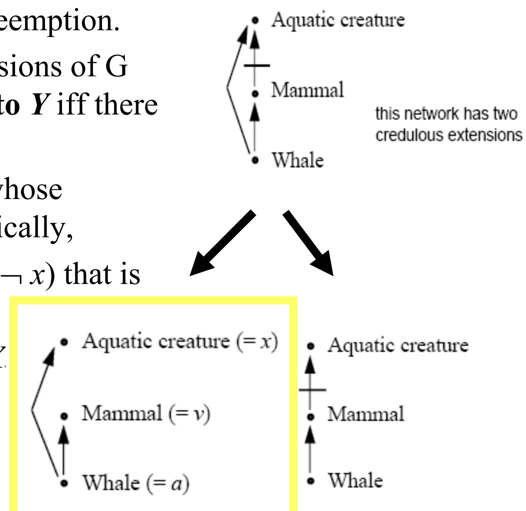

Subtleties

What to believe?

- **"credulous" reasoning**: choose a preferred extension and believe all the conclusions supported
- "**skeptical" reasoning:** believe the conclusions from any path that is supported by all preferred extensions
- **"ideally skeptical" reasoning**: believe the conclusions that are supported by all preferred extensions

Note: ideally skeptical reasoning cannot be computed in a path-based way (conclusions may be supported by different paths in each extension)

We've been doing "upwards" reasoning

- start at a node and see what can be inherited from its ancestor nodes
- there are many variations on this definition; none has emerged as the agreed upon, or "correct" one
- an alternative looks from the top and sees what propagates down upwards is more efficient