10.

Inheritance

Hierarchy and inheritance

As we noticed with both frames and description logics, hierarchy or taxonomy is a natural way to view the world

importance of *abstraction* in remembering and reasoning

- groups of things share properties in the world
- do not have to repeat representations

  e.g. sufficient to say that "elephants are mammals" to know a lot about them

Inheritance is the result of transitivity reasoning over paths in a network

- for strict networks, *modus ponens* (if-then reasoning) in graphical form
- "does *a* inherit from *b*?" is the same as "is *b* in the transitive closure of :**IS-A** (or subsumption) from *a*?"

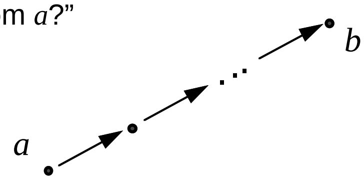

graphically, is there a path of :**IS-A** connections from *a* to *b*?

Path-based reasoning

Focus just on inheritance and transitivity

- many interesting considerations in looking just at where information comes from in a network representation
- abstract frames/descriptions, and properties into <u>nodes</u> in graphs, and just look at reasoning with paths and the conclusions they lead us to

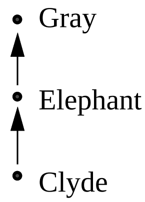

note the translation of property, Gray, and the constant Clyde into a node

- <u>edges</u> in the network: Clyde·Elephant, Elephant·Gray
- <u>paths</u> included in this network: edges plus {Clyde·Elephant·Gray}

  in general, a path is a sequence of 1 or more edges
- <u>conclusions</u> supported by the paths:

  Clyde → Elephant; Elephant → Gray; Clyde → Gray

Inheritance networks

(1) Strict inheritance in <u>trees</u>

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

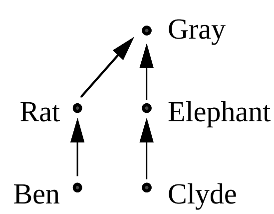

(2) Strict inheritance in <u>DAGs</u>

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

Note: negative edge from Student: "is not a"

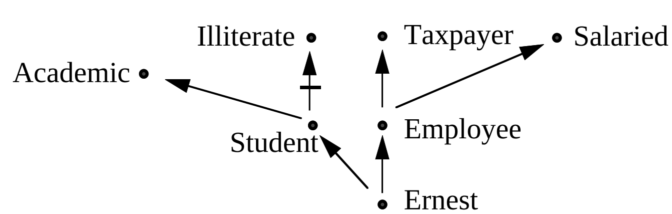

Inheritance with defeasibility

(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

while elephants in general are gray, Clyde is not

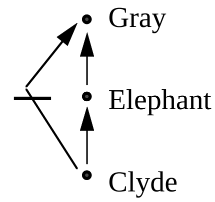

A key problem: *ambiguity*

- *credulous* accounts choose arbitrarily
- *skeptical* accounts are more conservative

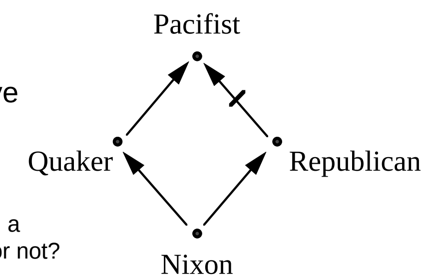

Is Nixon a pacifist or not?

Shortest path heuristic

Defeasible inheritance in DAGs

- links have *polarity* (positive or negative)
- use <u>shortest path heuristic</u> to determine which polarity counts

Intuition: inherit from the most specific subsuming class

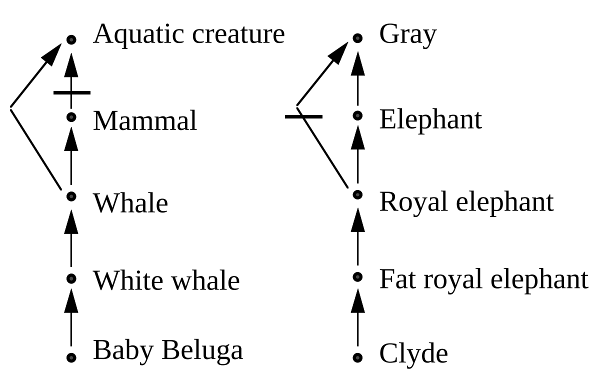

- as a result, not all paths count in generating conclusions
  - some are "preempted"
  - but some are "admissible"

think of paths as *arguments* in support of conclusions

⇒ the inheritance problem = what are the admissible conclusions?

Problems with shortest path

1. Shortest path heuristic produces incorrect answers in the presence of redundant edges (which are already implied!)

   the redundant edge *q*, expressing that Clyde is an Elephant changes polarity of conclusion about color

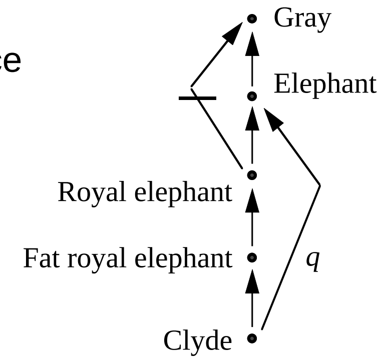

2. Anomalous behavior with ambiguity

   adding 2 edges to the left side changes the conclusion!

   Why should length be a factor?
   This network should be ambiguous…

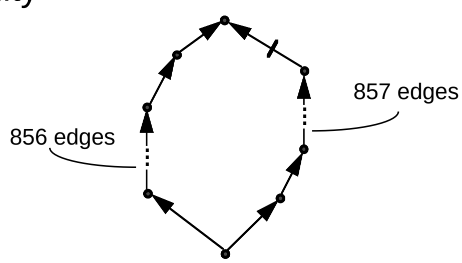

Specificity criteria

Shortest path is a <u>specificity criterion</u> (sometimes called a <u>preemption strategy</u>) which allows us to make admissibility choices among competing paths

- It's not the only possible one
- Consider "*inferential distance*": not linear distance, but topologically based
  - a node *a* is nearer to node *b* than to node *c* if there is a path from *a* to *c* through *b*
  - idea: conclusions from *b* preempt those from c

  This handles Clyde → ¬Gray just fine, as well as redundant links

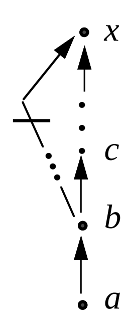

- But what if path from *b* to *c* has some of its edges preempted? what if some are redundant?

A formalization (Stein)

An <u>inheritance hierarchy</u> $\Gamma = <V,E>$ 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 $a \cdot x$
- negative edges are written $a \cdot \neg x$

A sequence of edges is a <u>path</u>:

- a <u>positive path</u> is a sequence of one or more positive edges $a \cdot \ldots \cdot x$
- a <u>negative path</u> is a sequence of positive edges followed by a single negative edge $a \cdot \ldots \cdot v \cdot \neg x$

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

  Also: there might be 0 positive edges.

A path (or <u>argument</u>) supports a <u>conclusion</u>:

- $a \cdot \ldots \cdot x$ supports the conclusion $a \rightarrow x$ (a is an x)
- $a \cdot \ldots \cdot \neg x$ supports $a \not\rightarrow x$ (a is not an x)

  Note: a conclusion may be supported by many arguments

However: not all arguments are equally believable...

Support and admissibility

Γ supports a path $a \cdot s_1 \cdot \ldots \cdot s_n \cdot (\neg)x$ if the corresponding set of edges $\{a \cdot s_1, \ldots, s_n \cdot (\neg)x\}$ is in E, and the path is <u>admissible</u> according to specificity (see below).

the hierarchy supports a conclusion $a \rightarrow x$ (or $a \not\rightarrow x$)
if it supports some corresponding path

A path is <u>admissible</u> if every edge in it is admissible.

An edge $v \cdot x$ is <u>admissible</u> in Γ wrt a if there is a positive path $a \cdot s_1 \ldots s_n \cdot v$ ($n \geq 0$) in E and

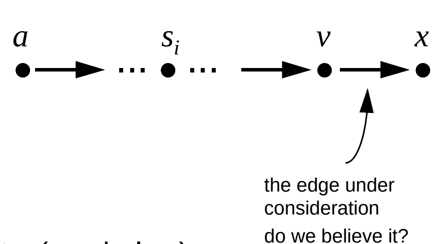

1. each edge in $a \cdot s_1 \ldots s_n \cdot v$ is admissible in Γ wrt a (recursively);
2. no edge in $a \cdot s_1 \ldots s_n \cdot v$ is redundant in Γ wrt a (see below);
3. no intermediate node $a, s_1, \ldots, s_n$ is a preemptor of $v \cdot x$ wrt a (see below).

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

Preemption and redundancy

A node y along path $a \cdot \ldots y \ldots \cdot v$ is a <u>preemptor</u> of the edge $v \cdot x$ wrt a if $y \cdot\neg x \in E$ (and analogously for $v \cdot \neg x$)

for example, in this figure the node Whale preempts the negative edge from Mammal to Aquatic creature wrt both Whale and Blue whale

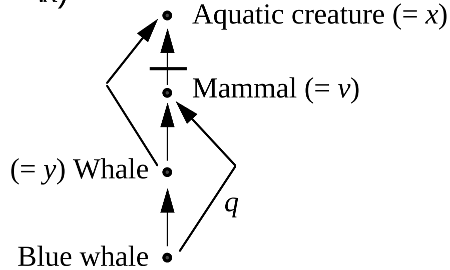

A positive edge $b \cdot w$ is <u>redundant</u> in Γ wrt node a if there is some positive path $b \cdot t_1 \ldots t_m \cdot w \in E$ ($m \geq 1$), for which

1. each edge in $b \cdot t_1 \ldots \cdot t_m$ is admissible in Γ wrt a;
2. there are no c and i such that $c \cdot \neg t_i$ is admissible in Γ wrt a;
3. there is no c such that $c \cdot \neg w$ is admissible in Γ wrt a.

The edge labelled q above is redundant

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

Credulous extensions

Γ is <u>a-connected</u> iff for every node x in Γ, there is a path from a to x, and for every edge $v \cdot (\neg) x$ in Γ, there is a *positive* path from a to v.

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

Γ is (potentially) <u>ambiguous</u> wrt a node a if there is some node $x \in V$ such that both $a \cdot s_1 \ldots s_n \cdot x$ and $a \cdot t_1 \ldots t_m \cdot \neg x$ are paths in Γ

A <u>credulous extension</u> of Γ wrt node a is a maximal unambiguous a-connected subhierarchy of Γ wrt a

If X is a credulous extension of Γ, then adding an edge of Γ to X makes X either ambiguous or not a-connected

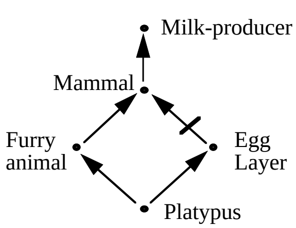

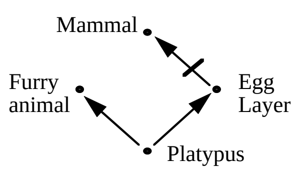

Extension 1

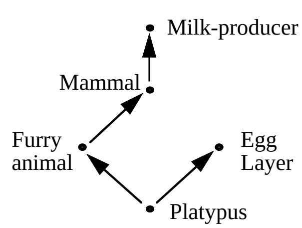

Extension 2

Preferred extensions

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

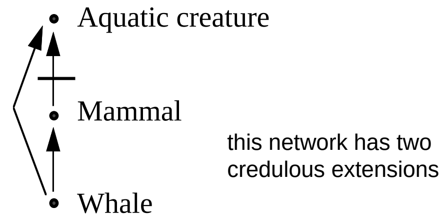

Let X and Y be credulous extensions of Γ wrt node a. X is preferred to Y iff there are nodes v and x such that:

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

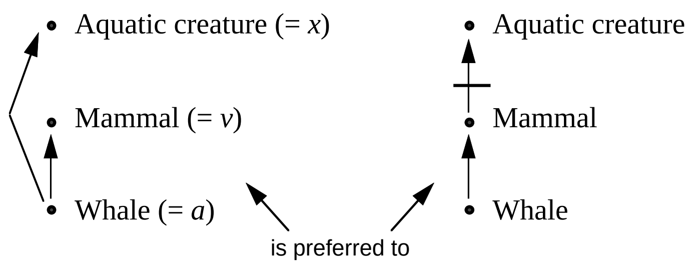

A credulous extension is a preferred extension if there is no other extension that is preferred to it.

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