Knowledge Representation and Reasoning

Chapters 10.1-10.3, 10.6, 10.9

Some material adopted from notes
by Andreas Geyer-Schulz
and Chuck Dyer

Overview

- Approaches to knowledge representation
- Deductive/logical methods
  - Forward-chaining production rule systems
  - Semantic networks
  - Frame-based systems
  - Description logics
- Abductive/uncertain methods
  - What's abduction?
  - Why do we need uncertainty?
  - Bayesian reasoning
  - Other methods: Default reasoning, rule-based methods, Dempster-Shafer theory, fuzzy reasoning

Introduction

- Real knowledge representation and reasoning systems come in several major varieties
- These differ in their intended use, expressivity, features,…
- Some major families are
  - Logic programming languages
  - Theorem provers
  - Rule-based or production systems
  - Semantic networks
  - Frame-based representation languages
  - Databases (deductive, relational, object-oriented, etc.)
  - Constraint reasoning systems
  - Description logics
  - Bayesian networks
  - Evidential reasoning

Semantic Networks

- A semantic network is a simple representation scheme that uses a graph of labeled nodes and labeled, directed arcs to encode knowledge.
  - Usually used to represent static, taxonomic, concept dictionaries
- Semantic networks are typically used with a special set of accessing procedures that perform "reasoning"
  - e.g., inheritance of values and relationships
- Semantic networks were very popular in the '60s and '70s but less used in the '80s and '90s. Back in the '00s as RDF
  - Much less expressive than other KR formalisms: both a feature and a bug!
- The **graphical depiction** associated with a semantic network is a significant reason for their popularity.

Nodes and Arcs

Arcs define binary relationships that hold between objects denoted by the nodes.

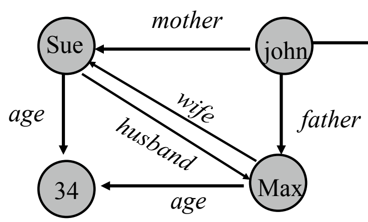

mother(john,sue)
age(john,5)
wife(sue,max)
age(max,34)
...

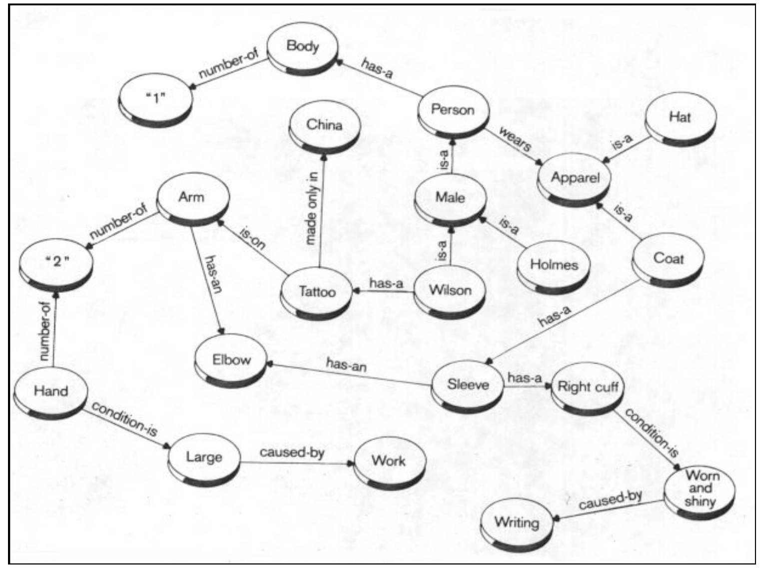

Semantic Networks

- The ISA (is-a) or AKO (a-kind-of) relation is often used to link instances to classes, classes to superclasses
- Some links (e.g. hasPart) are inherited along ISA paths.
- The *semantics* of a semantic net can be relatively informal or very formal
  - often defined at the implementation level

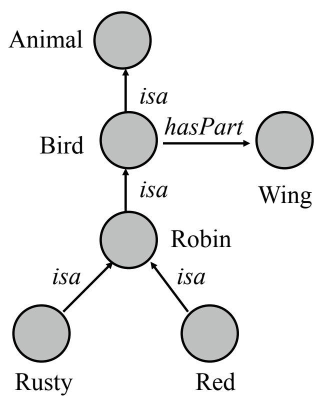

Reification

- Non-binary relationships can be represented by "turning the relationship into an object"
- This is an example of what logicians call "reification"
  - reify v : consider an abstract concept to be real
- We might want to represent the generic give event as a relation involving three things: a giver, a recipient and an object, give(john,mary,book32)

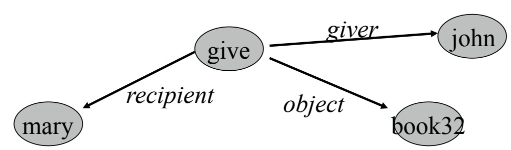

Individuals and Classes

Many semantic networks distinguish

- nodes representing individuals and those representing classes
- the “subclass” relation from the “instance-of” relation

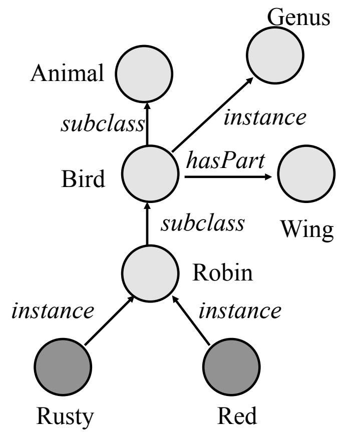

Link types

Link Type	Semantics	Example
$A \xrightarrow{Subset} B$	$A \subset B$	$Cats \subset Mammals$
$A \xrightarrow{Member} B$	$A \in B$	$Bill \in Cats$
$A \xrightarrow{R} B$	$R(A,B)$	$Bill \xrightarrow{Age} 12$
$A \xrightarrow{\boxed{R}} B$	$\forall x \; x \in A \Rightarrow R(x,B)$	$Birds \xrightarrow{\boxed{Legs}} 2$
$A \xrightarrow{\boxed{\boxed{R}}} B$	$\forall x \; \exists y \; x \in A \Rightarrow y \in B \wedge R(x,y)$	$Birds \xrightarrow{\boxed{\boxed{Parent}}} Birds$

Inference by Inheritance

- One of the main kinds of reasoning done in a semantic net is the inheritance of values along subclass and instance links
- Semantic networks differ in how they handle the case of inheriting multiple different values.
  - All possible values are inherited, ***or***
  - Only the “lowest” value or values are inherited

Conflicting inherited values

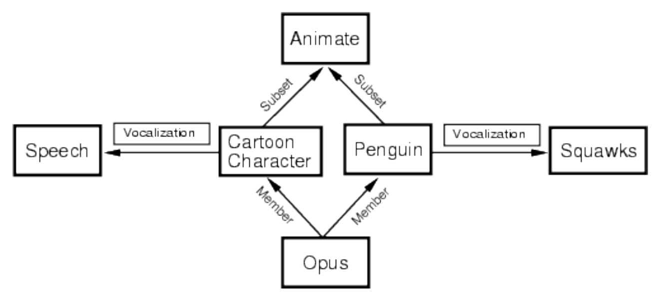

Multiple inheritance

- A node can have any number of superclasses that contain it, enabling a node to inherit properties from multiple “parent” nodes and their ancestors in the network
- These rules are often used to determine inheritance in such “tangled” networks where multiple inheritance is allowed:
  - If X<A<B and both A and B have property P, then X inherits A’s property.
  - If X<A and X<B but neither A<B nor B<A, and A and B have property P with different and inconsistent values, then X does not inherit property P at all.

Nixon Diamond

- This was the classic example circa 1980

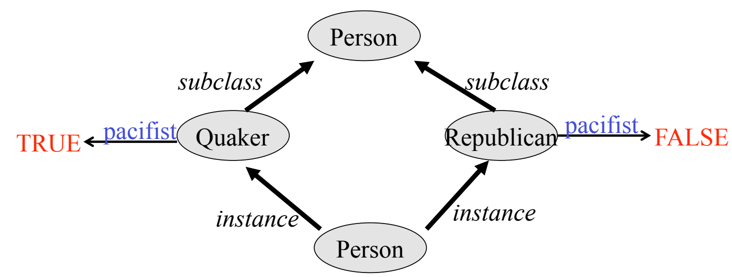

From Semantic Nets to Frames

- Semantic networks morphed into Frame Representation Languages in the ‘70s and ’80s
- A frame is a lot like the notion of an object in OOP, but has more meta-data
- A **frame** has a set of **slots**
- A **slot** represents a relation to another frame (or value)
- A slot has one or more **facets**
- A **facet** represents some aspect of the relation

Facets

- A slot in a frame holds more than a value.
- Other facets might include:
  - **Value**: current fillers
  - **Default:** default fillers
  - **Cardinality:** minimum and maximum number of fillers
  - **Type:** type restriction on fillers (usually expressed as another frame object)
  - **Proceedures:** attached procedures (if-needed, if-added, if-removed)
  - **Salience:** measure on the slot’s importance
  - **Constraints:** attached constraints or axioms
- In some systems, the slots themselves are instances of frames.

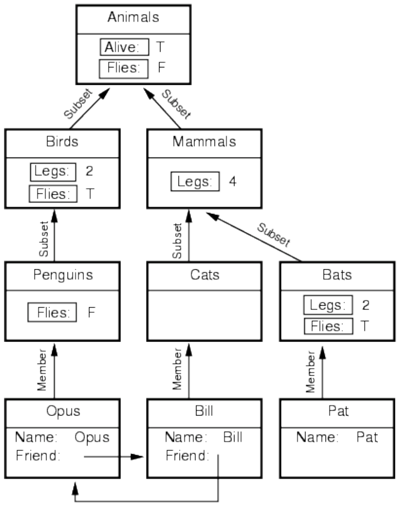

(a) A frame-based knowledge base

Rel(Alive,Animals,T)
Rel(Flies,Animals,F)

Birds ⊂ Animals
Mammals ⊂ Animals

Rel(Flies,Birds,T)
Rel(Legs,Birds,2)
Rel(Legs,Mammals,4)

Penguins ⊂ Birds
Cats ⊂ Mammals
Bats ⊂ Mammals
Rel(Flies,Penguins,F)
Rel(Legs,Bats,2)
Rel(Flies,Bats,T)

Opus ∈ Penguins
Bill ∈ Cats
Pat ∈ Bats
Name(Opus,"Opus")
Name(Bill,"Bill")
Friend(Opus,Bill)
Friend(Bill,Opus)
Name(Pat,"Pat")

(b) Translation into first-order logic

Description Logics

- Description logics provide a family of frame-like KR systems with a formal semantics.
  - E.g., KL-ONE, LOOM, Classic, …
- An additional kind of inference done by these systems is automatic **classification**
  - finding the right place in a hierarchy of objects for a new description
- Current systems take care to keep the languages simple, so that all inference can be done in polynomial time (in the number of objects)
  - ensuring tractability of inference
- The Semantic Web language OWL is based on description logic

Abduction

- **Abduction** is a reasoning process that tries to form plausible explanations for observations
  - Distinctly different from deduction and induction
  - Inherently unsound and uncertain
- Uncertainty is an important issue in abductive reasoning
- Some major formalisms for representing and reasoning about uncertainty
  - Mycin's certainty factors (an early representative)
  - **Probability theory (esp. Bayesian belief networks)**
  - Dempster-Shafer theory
  - Fuzzy logic
  - Truth maintenance systems
  - Nonmonotonic reasoning

Abductive reasoning

- **Definition** (Encyclopedia Britannica): reasoning that derives an explanatory hypothesis from a given set of facts
  - The inference result is a **hypothesis** that, if true, could **explain** the occurrence of the given facts
- **Examples**
  - Dendral, an expert system to construct 3D structure of chemical compounds
    - Fact: mass spectrometer data of the compound and its chemical formula
    - KB: chemistry, esp. strength of different types of bounds
    - Reasoning: form a hypothetical 3D structure that satisfies the chemical formula, and that would most likely produce the given mass spectrum

Abduction examples (cont.)

- Medical diagnosis
  - Facts: symptoms, lab test results, and other observed findings (called manifestations)
  - KB: causal associations between diseases and manifestations
  - Reasoning: one or more diseases whose presence would causally explain the occurrence of the given manifestations
- Many other reasoning processes (e.g., word sense disambiguation in natural language process, image understanding, criminal investigation) can also been seen as abductive reasoning

abduction, deduction and induction

Deduction: major premise: All balls in the box are black
minor premise: These balls are from the box
conclusion: These balls are black

A => B
A

B

Abduction: rule: All balls in the box are black
observation: These balls are black
explanation: These balls are from the box

A => B
B

Possibly A

Induction: case: These balls are from the box
observation: These balls are black
hypothesized rule: All ball in the box are black

Whenever A then B

Possibly A => B

Deduction reasons from causes to effects
Abduction reasons from effects to causes
Induction reasons from specific cases to general rules

Characteristics of abductive reasoning

- "Conclusions" are **hypotheses**, not theorems (may be false *even if* rules and facts are true)
  - E.g., misdiagnosis in medicine
- There may be multiple plausible hypotheses
  - Given rules A => B and C => B, and fact B, both A and C are plausible hypotheses
  - Abduction is inherently uncertain
  - Hypotheses can be ranked by their plausibility (if it can be determined)

Reasoning as a hypothesize-and-test cycle

- **Hypothesize**: Postulate possible hypotheses, any of which would explain the given facts (or at least most of the important facts)
- **Test**: Test the plausibility of all or some of these hypotheses
- One way to test a hypothesis H is to ask whether something that is currently unknown–but can be predicted from H–is actually true
  - If we also know A => D and C => E, then ask if D and E are true
  - If D is true and E is false, then hypothesis A becomes more plausible (**support** for A is increased; **support** for C is decreased)

Abduction is non-monotonic

- That is, the plausibility of hypotheses can increase/ decrease as new facts are collected
- In contrast, deductive inference is **monotonic:** it never change a sentence's truth value, once known
- In abductive (and inductive) reasoning, some hypotheses may be discarded, and new ones formed, when new observations are made

Sources of uncertainty

- Uncertain **inputs**
  - Missing data
  - Noisy data
- Uncertain **knowledge**
  - Multiple causes lead to multiple effects
  - Incomplete enumeration of conditions or effects
  - Incomplete knowledge of causality in the domain
  - Probabilistic/stochastic effects
- Uncertain **outputs**
  - Abduction and induction are inherently uncertain
  - Default reasoning, even in deductive fashion, is uncertain
  - Incomplete deductive inference may be uncertain

‣ Probabilistic reasoning only gives probabilistic results (summarizes uncertainty from various sources)

Decision making with uncertainty

- **Rational** behavior:
  - For each possible action, identify the possible outcomes
  - Compute the **probability** of each outcome
  - Compute the **utility** of each outcome
  - Compute the probability-weighted **(expected) utility** over possible outcomes for each action
  - Select the action with the highest expected utility (principle of **Maximum Expected Utility**)

Bayesian reasoning

- Probability theory
- Bayesian inference
  - Use probability theory and information about independence
  - Reason diagnostically (from evidence (effects) to conclusions (causes)) or causally (from causes to effects)
- Bayesian networks
  - Compact representation of probability distribution over a set of propositional random variables
  - Take advantage of independence relationships

Other uncertainty representations

- Default reasoning
  - Nonmonotonic logic: Allow the retraction of default beliefs if they prove to be false
- Rule-based methods
  - Certainty factors (Mycin): propagate simple models of belief through causal or diagnostic rules
- Evidential reasoning
  - Dempster-Shafer theory: Bel(P) is a measure of the evidence for P; Bel(¬P) is a measure of the evidence against P; together they define a belief interval (lower and upper bounds on confidence)
- Fuzzy reasoning
  - Fuzzy sets: How well does an object satisfy a vague property?
  - Fuzzy logic: "How true" is a logical statement?

Uncertainty tradeoffs

- **Bayesian networks:** Nice theoretical properties combined with efficient reasoning make BNs very popular; limited expressiveness, knowledge engineering challenges may limit uses
- **Nonmonotonic logic:** Represent commonsense reasoning, but can be computationally very expensive
- **Certainty factors:** Not semantically well founded
- **Dempster-Shafer theory:** Has nice formal properties, but can be computationally expensive, and intervals tend to grow towards [0,1] (not a very useful conclusion)
- **Fuzzy reasoning:** Semantics are unclear (fuzzy!), but has proved very useful for commercial applications