

Artificial Intelligence

Knowledge Representation & KB

09/25/2002

Knowledge Representation

- John McCarthy, 1958:
 - “In order for a program to be capable of learning something it must first be capable of being told it.”
 - “We shall say that a program has a common sense if it automatically deduces for itself a sufficiently wide class of immediate consequence of anything it is told and what it already knows.”

“In the Knowledge is the Power”

- It is not enough to be smart and clever
- Knowledge is a core enabler of intelligent behavior
 - reasoning about goals, environment, other agents, one’s self
- Declarative representations are the focus of KR research
 - knowledge that is domain-specific but task-independent
 - reasoning that is task-specific but domain-independent
- A declarative representation is called a **‘knowledge base’**

What is knowledge representation?

- Surrogate
 - description of something else
 - inside the agent
 - requires specification of mapping to intended referent
 - contains simplifying assumptions and inaccuracies
 - susceptible to supporting incorrect reasoning results
- Set of ontological commitments
 - What to consider in thinking about the world?
 - concepts, relations, objects (e.g., representing an electronic circuit)
 - How to categorize objects and relations?
 - e.g., diseases: alcoholism? chronic fatigue syndrome?
 - Not about data structures

What is knowledge representation? (cont'd)

- Fragmentary theory of intelligent reasoning
 - Provides a conception of intelligent inference
 - What does it mean to reason intelligently?
 - Sanctions a set of inferences
 - What can we infer from what we know?
 - Recommends a set of inferences
 - What ought we to infer from what we know?
- Medium for efficient computation
 - Reasoning in machines is a computational process
 - Computational efficiency is a central design goal
 - Expressivity and tractability of reasoning are traded off
- How useful is it as a medium of communication?
 - Can we easily ‘talk’ or think in the representation language?
 - What kinds of things are easily said in the language?
 - What kinds of things are so difficult to say in the language as to be pragmatically impossible?

Sample issues in KR Research

- What knowledge needs to be represented to answer given questions?
- How is incomplete or noisy information represented?
- How is qualitative or abstract knowledge represented?
- How can knowledge be encoded so that it is reusable?
- How are assumptions represented and reasoned with?
- How can knowledge be reformulated for a given purpose?
- How can effective automatic reasoning be done with large-scale knowledge bases?
- How can computer-interpretable knowledge be extracted from documents?
- How can knowledge from multiple sources be obtained and used?

Knowledge Representation & Reasoning

- Focused on designing forms for expressing information
 - mostly general-purpose languages and systems
- Must consider reasoning to be done with representation
 - What queries are to be answerable from the represented knowledge?
 - What algorithms are to be used to reason with the knowledge?
- Expressiveness vs. tractability tradeoff
 - How to express what we know?
 - How to reason with what we express?

KR and Data Base Research

- Both ‘represent’ knowledge
- Databases contain only ‘ground literals’
 - No disjunctions (e.g., ‘The ball is either red or blue.’)
 - No quantifiers (e.g., ‘Every person has two parents.’)
- Database schema provide some quantified information
- Deductive databases include implications
- Database concerns:
 - Efficient access and management of large databases
 - Concurrent updating
- KR concerns:
 - Expressivity
 - Effective reasoning

History of KR

- 60s – 70s
 - origins
 - problem solving work primarily at CMU and MIT
 - natural language understanding
 - many ad hoc formalisms
 - ‘procedural’ vs. ‘declarative’ knowledge
 - no formal semantics
- 70s – 80s
 - Semantic nets
 - unstructured node-link graphs
 - no semantics to support interpretation
 - no axioms to support reasoning
 - Frames
 - structured semantic nets
 - object-oriented descriptions
 - prototypes
 - class-subclass taxonomies

History of KR (cont'd)

- 60s – 70s
 - origins
 - problem solving work primarily at CMU and MIT
 - natural language understanding
 - many ad hoc formalisms
 - ‘procedural’ vs. ‘declarative’ knowledge
 - no formal semantics
- 70s – 80s
 - Semantic nets
 - unstructured node-link graphs
 - no semantics to support interpretation
 - no axioms to support reasoning
 - Frames
 - structured semantic nets
 - object-oriented descriptions
 - prototypes
 - class-subclass taxonomies

Example Class-Subclass Taxonomy

Class hierarchy (20 classes defined):

Author

Document

 Book

 Edited-Book

 Miscellaneous-Publication

 Artwork

 ...

 Technical-Manual

 Periodical-Publication

 Journal

 Magazine

 Newspaper

 Proceedings

 Technical-Report

 Thesis

 Doctoral-Thesis

 Masters-Thesis

Title

History of KR (cont'd)

- 70s – 80s
 - Production rule systems
 - if-then inference rules
 - if (warning-light on) then (engine overheating)
 - if (warning-light on) then ((engine overheating) 0.95)
 - situation-action rules
 - if (warning-light on) then (turn-off engine)
 - Hybrid procedural-declarative representation
 - Basis for expert systems
 - Predicate calculus
 - primarily first order logic
 - ‘Everybody loves somebody sometime.’
(for all ?p (implies (Person ?p1) (exists (?p2 ?t) (and (Person ?p2)
(Time ?t) (Loves ?p1 ?p2 ?t))))))

History of KR (cont'd)

- 90s – 00s
 - Declarative representations
 - Easier to change
 - Multi-use
 - Extendable by reasoning
 - Accessible for introspection
 - Formal semantics
 - Defines what the representation means
 - Specifies correct reasoning
 - Allows comparison of representation/algorithms
 - KR rooted in the study of logics
 - temporal, context, modal, default, nonmonotonic, ...
 - Rigorous theoretical analysis

Representing Uncertain Information

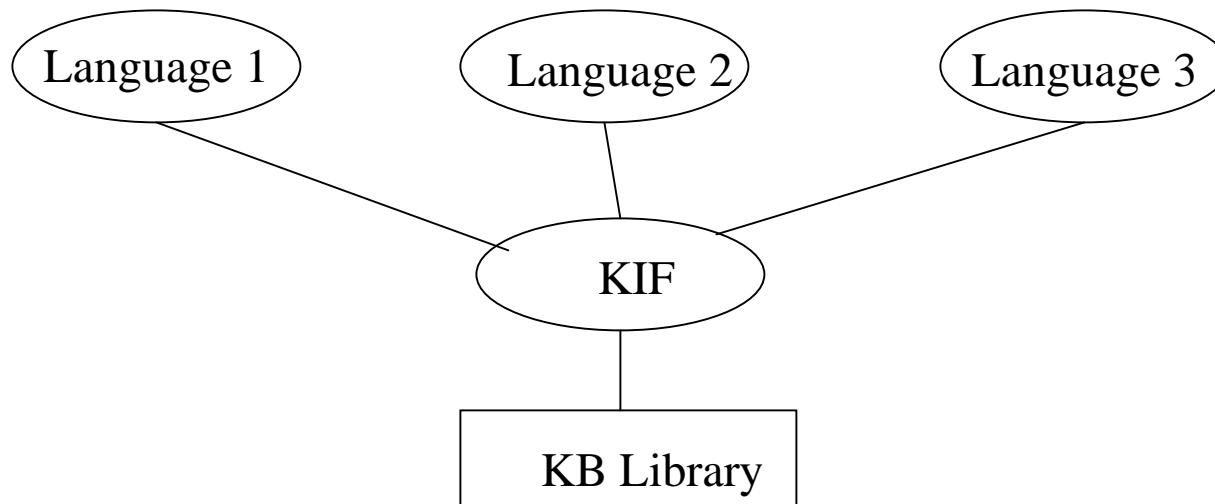
- Agent cannot access the whole truth about its environment
- Represent the degree of beliefs of knowledge
- Core problem in knowledge representation
- Many competing formalisms in the past
 - Certainty factors
 - Dempster-Shafer Theory
 - Fuzzy logic
- Probability theory is now the dominant formalism
 - reasoning based on Belief Networks (a.k.a. Bayes Nets)

KR Language Components

- A logical formalism
 - Syntax for Well-Formed Formulae (WFF)
 - Vocabulary of logical symbols (e.g., AND, OR, NOT, \Rightarrow , \Leftrightarrow)
 - Interpretation semantics for the logical symbols
 - e.g., ' $\Rightarrow A B$ ' is true if and only if B is true or A is false
- An ontology
 - Vocabulary of non-logical symbols
 - relations, functions, constants
 - Definition of non-primitive symbols
 - e.g., (\Leftrightarrow (Person ?x) (Gender (Mother ?x) Female))
- A proof theory
 - Specification of the reasoning steps that are logically sound
 - e.g., From ' $\Rightarrow S1 S2$ ' and ' $S1$ ' conclude ' $S2$ '

Interlingua for Multi-Use Knowledge

- Knowledge Interchange Format (KIF)
 - First order predicate logic
 - includes numbers, lists, and strings
 - linear ASCII syntax
 - in the process of becoming an ANSI standard



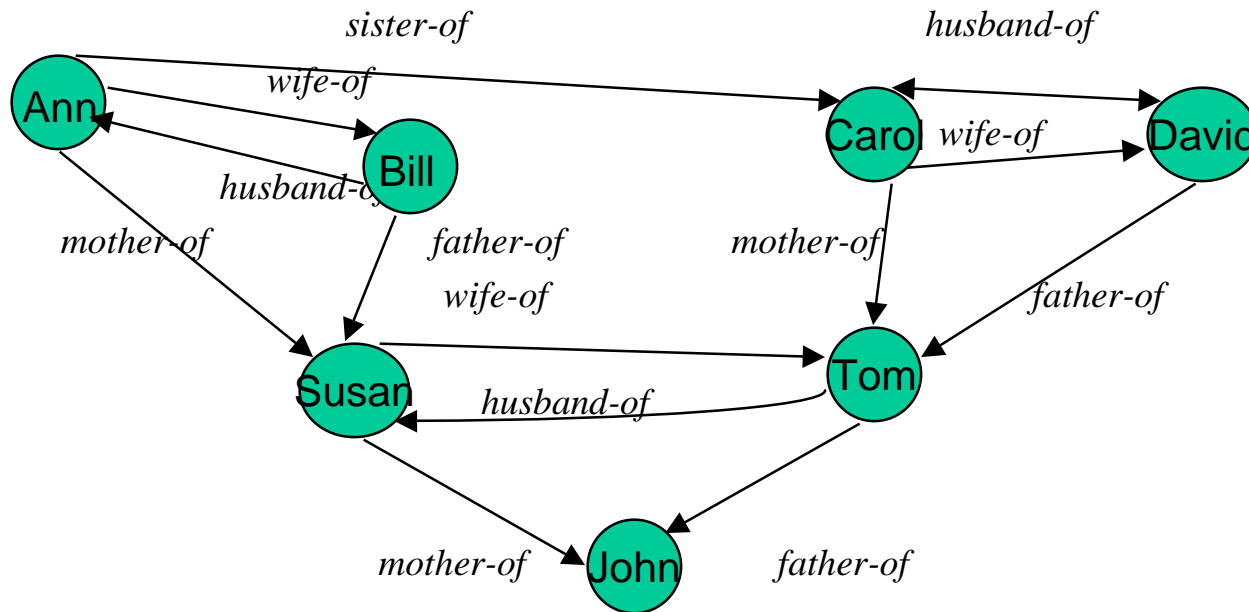
KBs, Sentences, Terms, Words

- Knowledge Base – Collection of sentences
- Sentence – Expression denoting a statement
- Term – Expression denoting an object
- Words
 - Constant
word not beginning with ‘?’ or ‘@’
e.g., Fred, Block-A, Justice
 - Individual Variable
word beginning with ‘?’
e.g., ?x, ?The-First-Murderer
 - Sequence Variable
word beginning with ‘@’
e.g., @x, @The-Other-Murderers

Semantic Networks

A semantic net is a labeled directed graph, where each node represents an object (a proposition), and each link represents a relationship between two objects.

Example:



Semantic Networks (cont'd)

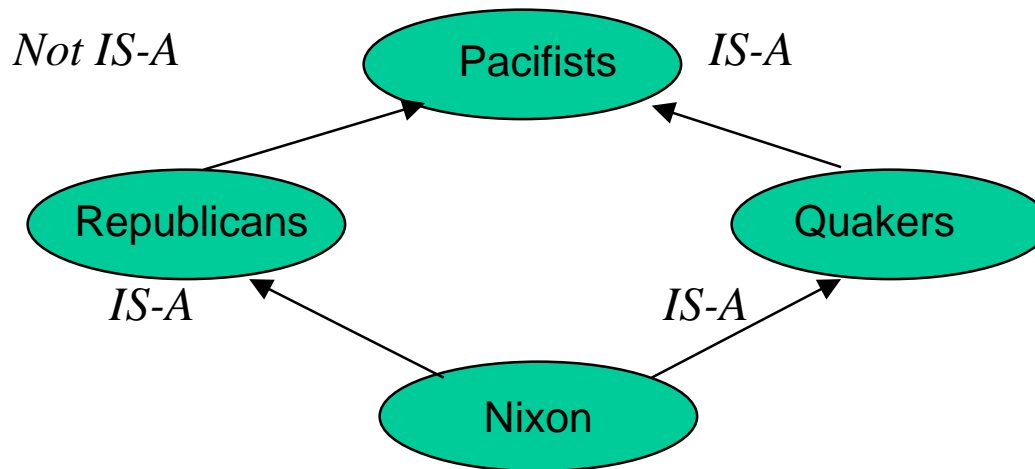
- Semantic nets can represent only propositional information.
 - Relations between propositions are of primary interest because they provide the basic structure for organizing knowledge. Some important relations are
 - “IS-A” (is an instance of). Refers to a member of a class, where a class is a group of objects with one or more common attributes (properties). For example, “Tom IS-A bird”.
 - “A-KIND-OF”. Relates one class to another, for example “Birds are A-KIND-OF animals”.
 - “HAS-A”. Relates attributes to objects, for example “Mary HAS-A cat”.
 - “CAUSE”. Expresses a causal relationship, for example “Fire CAUSES smoke”.
- Semantic nets can be easily converted into a set of FOL formulas, and vice versa. Semantic nets, however, have two important advantages, which makes them a representation of choice in some applications (for example, in natural language understanding):
 - A very simple execution model.
 - Very readable representation, which makes its easy to visualize inference steps.

Inference in Semantic Networks

- The inference procedure in semantic nets is called inheritance, and it allows one node's characteristics to be duplicated by a descendent node.
 - e.g., consider a class “aircraft”, and assume that “balloons”, “propeller-driven objects” and “jets” are subclasses of it, i.e.
 - “Balloons are A-KIND-OF aircrafts”
 - “Propeller-driven objects are A-KIND-OF aircrafts”, etc.
 - Assume that the following attributes are assigned to aircrafts: “Aircraft IS-A flying object”, “Aircraft HAS-A wings”, “Aircraft HAS-A engines”
- All properties assigned to the superclass, “aircraft”, will be inherited by its subclasses, unless there is an “exception” link capturing a non-monotonic inference relation.

Multiple inheritance may result in a conflicting inference

- In some semantic networks, one class can inherit properties of more than one superclass.
 - The “Nixon diamond” example: It is widely accepted that Quakers tend to be pacifists, and Republicans tend not to be. Nixon is known to be both - a Quaker, and a Republican.



- The resulting conflict can be resolved only if additional information stating a preference to one of the conflicting inferences is provided.

Object-Attribute-Value (OAV) triplets

- One problem with semantic nets is that there is no standard definition of link names
 - To avoid this ambiguity, we can restrict this formalism to a very simple kind of a semantic network, which has only two types of links, “HAS-A” and “IS-A”. Such a formalism is called Object-Attribute-Value (OAV) triplets, and it is a widely used mode of knowledge representation (especially for representing declarative knowledge)
 - e.g., Consider **object** “airplane”. Some of its **attributes** are:
 - number of engines;
 - type of engines;
 - type of wing design.
 - **Possible values** of these attributes are:
 - number of engines: 2, 3, 4.
 - type of engines: jet, propeller-driven.
 - type of wing design: conventional, swept-back.

Difficulties with semantic nets and OAV triplets

- There is no standard definition of link and node names. This makes it difficult to understand the network, and whether or not it is designed in a consistent manner.
- Inheritance is a combinatorially explosive search, especially if the response to a query is negative. Plus, it is the only inference mechanism built in semantic nets, which may be insufficient for some applications.
- Initially, semantic nets were proposed as a model of human associative memory (Quinllian, 1968). But, are they an adequate model? It is believed that human brain contains about 10^{10} neurons, and 10^{15} links. Consider how long it takes for a human to answer “NO” to a query “Are there trees on the moon?” Obviously, humans process information in a very different way, not as suggested by the proponents of semantic networks.
- Semantic nets are logically and heuristically very weak. Statements such as “Some books are more interesting than others”, “No book is available on this subject”, “If a fiction book is requested, do not consider books on history, health and mathematics” **cannot** be represented in a semantic network.

Frames (Minsky, 1975)

- Semantic nets represent shallow knowledge, because all of the information must be represented in terms of nodes and links which are propositions
 - What if the objects in the domain are, in turn, complex structures? For example, consider object “animal”. We may want to incorporate as part of the object’s description all of the important properties of this object.
 - e.g., AIMA, page 318, Figure 10.7
 - Note that “frames” comprising the nodes of this frame-based network represent “typical examples” or “stereotypes” of described objects. Such typical example are called **concepts**, and they like data structures where data, in turn, contain data.
 - The underlying assumption of the frame theory is that when one encounters a new situation (or a substantial change in one’s view of the situation occurs), one selects from his/her memory the frame representing a given concept and changes it to reflect the new reality.

Frame Language (KIF)

- Object-oriented representation of languages
- Class-subclass taxonomies
- Prototype descriptions of class instances
- Frame systems perform standard set of inferences
 - inheritance of attribute values and constants
 - type checking of attribute values
 - checking number of attribute values
- Open Knowledge Base Connectivity (OKBC)
 - An API for KR systems
 - Assumes a frame language representation

Example Class Frame and Instance (KIF)

- Male-Person
 - Gender: Male
 - Value-Type: Gender
 - Slot-Cardinality: 1
 - Age
 - Unit: Year
 - Value-Type: Integer
 - Slot-Cardinality: 1
 - Parent:
 - Value-Type: Person
 - Slot-Cardinality: 2
 - Inverse: Child
 - Subset-Of-Values:
 - (Child Grandparent)
- Woojin
 - Instance-Of: Male-Person
 - ...
 - Parent: JK
 - Value-Type: Person
 - Slot-Cardinality: 2
 - Inverse: Child
 - Subset-Of-Values:
 - (Child Grandparent)
 - Child: XY
 - JK
 - Child: Woojin
 - XY
 - Grandparent: JK
 - Parent: Woojin