What's in a Link Revisited: a Semantic Odyssey

William A. Woods

william.woods@parsecraft.com

Background: Semantics for Question Answering

Before "What's in a Link," I worked on

Question Answering

A major concern was Meaning:

What is the meaning of a question?

What is the relationship between a question and an answer?

What is meaning anyway?

Meaning (Carnap)

"To know the truth conditions of a sentence is to know what is asserted by it -- in usual terms, its 'meaning' "

issue: how to capture truth conditions (an infinite set) in a finite reasoning automaton in a way that can be effectively used

Procedural Semantics

Meaning is defined by:

•abstract procedures for determining referents, verifying facts, computing values (including truth values), and carrying out actions

• built on computational operators -- cons, cond, loop, read, print, etc.

• provides a principled connection between "mental" symbols and what they denote or mean

Procedural Semantics in LUNAR (1971)

What is the average concentration of Aluminum in each breccia

```
(FOR EVERY X5 / (SEQ TYPECS) : T ;
   (PRINTOUT (AVGCOMP X5
        (QUOTE OVERALL) (QUOTE AL2O3))))
```

LUNAR's Meaning Representation Language

an extension of the Predicate Calculus with generalized quantifiers, computation, and imperative operators

(FOR <quant> <vbl> / <class> : <condition> ; <command>)

(FOR <quant> <vbl> / <class> : <condition> ; <condition>)

(TEST <condition>) (PRINTOUT <designator>)

Procedural Semantics

- permits a computer to understand, in a single, uniform way, the meanings of conditions to be tested, questions to be answered, and actions to be carried out
- permits a general-purpose system for language understanding to be used with different data bases that may have different representational conventions and different data structures
- can reason about meanings and can perceive and act on objects in the real world

Requirements for Semantic Representation

I wanted a representational system to satisfy two requirements:

- expressively adequate to represent all of the necessary elements of natural language questions, commands, assertions, conditions, and designators
- structured to support efficient semantic interpretation, retrieval, and inference

Semantic Networks

A network of concepts connected by links used for associative access

Questions:

What do the links mean?

What do they have to do with semantics?

Semantic Associations (Collins and Quillian)

A bird is an animal

An animal has skin

A canary is a bird

A bird has feathers

Does a canary have feathers?

Does a canary have skin?

Associative Distance in a Semantic Network

"A canary has feathers" is recognized faster than "a canary has skin" (because of an additional link in the inheritance chain?)

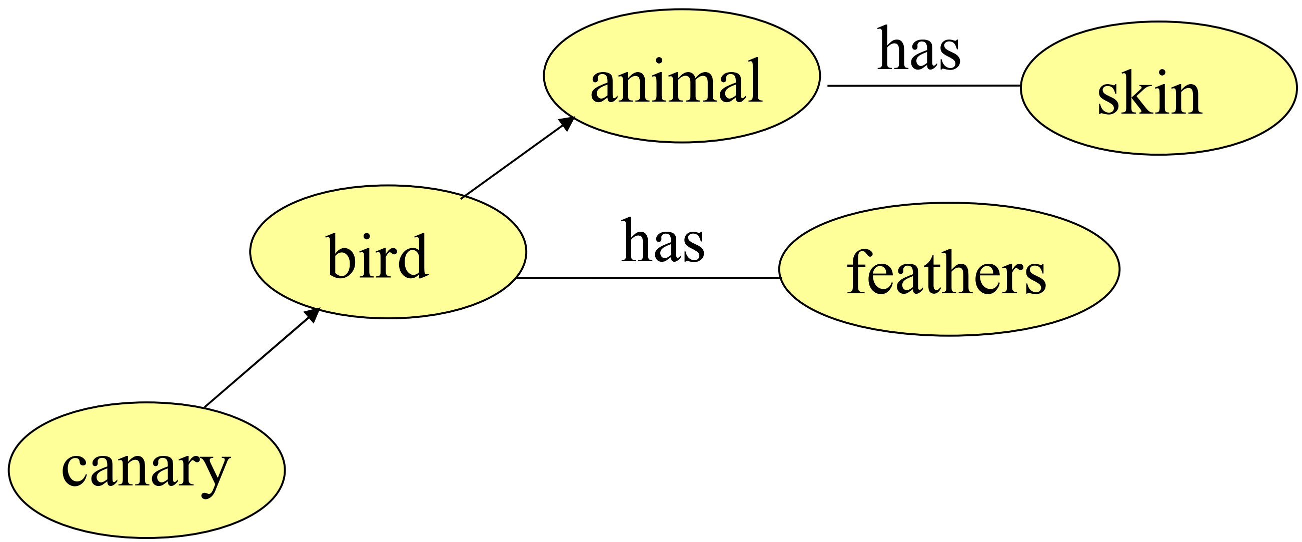

Semantics

The relationship between signs or symbols and the things that they denote or mean

Some Inadequate Notions of Semantics / Meaning

- The totality of the concepts connected to a concept are its semantics
- Whatever a system does with its input is its semantics
- There's no difference between syntax and semantics
- Semantics is in the eye of the beholder (e.g., in the node names)

Different Perspectives on Semantics

Linguist

- a way to represent different meanings of an ambiguous sentence
- mapping from sentences to different representations of meaning

Philosopher

- truth conditions in possible worlds (of an already unambiguous representation)

Different Perspectives on Semantics

The linguist is interested in Grammaticality

- I saw the man from the park.
    - Represent alternative interpretations:
    - I saw (the man) (from the park)
    - I saw (the man from the park)
- *The amoeba hit John with a hammer.
    - "Semantic" features:
    - Instrumental "with" requires + agent
    - Amoeba => - agent

Different Perspectives on Semantics

The philosopher is interested in Reasoning

- A model consists of an assignment of truth values to every possible instantiation of every predicate over a universe of individuals.
- A reasoning system ("logic") is "sound" if it can prove only "valid" conclusions (true in every possible model)
- A reasoning system is "complete" if it can prove every valid conclusion.

What's Still Missing

relationships between symbols and what they denote or mean

- how can abstract truth conditions be finitely represented and used by a reasoning automaton
- how can mental concepts be anchored to the physical world

Procedural Semantics as a Theory of Meaning

- meaning of a noun is a procedure for recognizing instances
- meaning of a proposition is a procedure for determining if it is true
- meaning of an action is the ability to do it and/or tell if it has been done
- can provide a principled sensorimotor connection to real-world objects

Semantics of Links

What does a link represent?

- attribute and value
- attribute and value-predicate
- relation and object-of-relation
- function and argument
- action and object
- role and constituent
- role and constituent-restriction

Attribute-Value Representations

John

height	6 feet
haircolor	brown
occupation	scientist

Value-Predicate and Role Relation Representations

"John's height is greater than 6 feet"

John

    height (greaterthan 6 feet)

"John's height is greater than Sue's"

(height John)

    greater (height Sue)

Links and Predication

John

hit Mary

Mary

hit* John

Mixed Conventions

John

	height	6 feet
	hit	Mary

Case Representation

sell

agent	John
recipient	Mary
patient	book

Case Frame Definition

sell

agent	person
recipient	person
patient	thing

Assertional Versus Structural Links

N12368

superc	telephone
color	black

a black telephone ?

all telephones are black ?

Intensions and Extensions

E.g., Venus

the morning star

the evening star

This example requires a notion of meaning in which distinct abstract entities (intensions) stand between the phrases and their common referent. Different intensions can have the same referent (extension).

Conclusions from What's in a Link

Need the capability to represent at least the functional equivalent of LUNAR's generalized quantification

- need to represent propositions without commitment to their truth
- need to represent descriptions of individuals without commitment to their existence
- need to represent intensional entities

Links and Logic

Can we combine best of two traditions:

Logical Deduction

rigorous, formal, but often counterintuitive

algorithms that match expressions, substitute values for variables, and invoke rules

Associative Networks

associative, intuitive, but typically informal

algorithms that follow paths through links

Enter KL-ONE

- Structured Inheritance Networks
- Primitive vs Defined Concepts
- Concepts and Roles
- Number Restrictions and Value Restrictions
- MSS and MGS classification algorithms

Original Goals of KL-ONE

- organizing structure for the knowledge of a reasoning system
- efficient associative access to knowledge
- complete expressivity
- support for high-level perception
- automatic classification
- inheritance of attributes and rules

Hypothesis: Most Reasoning is Recognize and React

Most reasoning consists of recognizing a current situation as one that you know what to do about and doing it.

Deductive reasoning happens less often, and it also uses recognize-and-react to pursue various reasoning strategies.

So it's important to be able to recognize known situations quickly and efficiently

Some Fundamental Issues

How does a reasoning system find relevant pieces of information and relevant rules of inference when it knows millions of things?

How does it acquire and organize millions of items of information?

How does it integrate new information with previously existing information?

How does it use its knowledge to impose structure on situations and decide what to do?

A Potential Answer: Taxonomic Subsumption

Organize concepts by subsumption (generality) in a conceptual taxonomy.

Use this taxonomic structure to organize all of the system's facts and rules.

Use taxonomic subsumption algorithms to build the taxonomy, assimilate new information into it, and find things in it.

For Example

Organize the pattern parts of inference rules in a conceptual taxonomy.

Classify the current subgoal against the taxonomy to determine the most specific applicable inference rules.

Choose a rule and a new subgoal and repeat

Automatic Classification and MSS (a la KL-ONE)

Automatic Classification:

- Taxonomic Subsumption algorithms use the meanings of concepts to determine when one concept is more general than another:
- A Most-Specific-Subsumer algorithm (MSS) can find where a new concept belongs in an existing taxonomy
- A Most-General-Subsumee algorithm (MGS) finds the children of the new concept

The evolution of KL-ONE

- Expressivity / Complexity Tradeoffs
- Many Theoretical Complexity Results
- The KL-ONE Family
- Terminological Logic / Description Logics

Understanding Subsumption and Taxonomy (Woods, 1991)

Revisit original goals of KL-ONE

- Efficient Principled Methodology for Organizing Knowledge
- Intensional rather than Extensional Subsumption Criterion

Intensional Subsumption Test

Every element of the parent concept subsumes some element of the child

[a person with a professional spouse]

subsumes

[a woman with a doctor husband]

Compare the sufficient conditions of the parent with the necessary conditions of the child

Many Links Have Implicit Quantificational Import

"birds have wings"

bird   have   wing   ?

"people need vitamins"

person   need   vitamin   ?

"people live in houses"

person   live-in   house   ?

"a person broke a window"

person   break   window   ?

Explicit Tags can Show Quantificational Import

birds have wings

-- Bird (AE have wing)

people need vitamins

-- Person (AA need vitamin)

children like candy

-- Child (TAA like candy)

John lives in Chicago

-- John (II live_in Chicago)

Quantificational Tags can be treated as Relation-Forming Operators

AE have =

(lambda (X Y) A x in X E y in Y : x have y)

AA need =

(lambda (X Y) A x in X A y in Y : x need y)

TAA like =

(lambda (X Y) A x in X A y in Y : (conclude likely(x like y)) unless known(not(x like y)))

Advantages of Explicit Quantificational Tags

- separates logical quantification from domain-specific elements of links
- provides the classifier and reasoner with needed information
- forms a contract between the knowledge engineer and the reasoner
- can be treated as relation-forming operators in link-oriented structures

Tagging the Kind/Instance Distinction

AI -- "everybody likes John"

EI -- "somebody likes John"

IA -- "John likes everybody"

IE -- "John likes somebody"

SRI -- "only sport fans like John"

IVR -- "John eats only vegetarian food."

II -- "John likes Mary"

Quantificational Tags can answer "What's in a Link?"

N12368

MI kind-of telephone

MI color black

"a black telephone"

II name telephone

AI color black

"all telephones are black"

Subsumption Depends on the Semantics of Link Tags

E.g.,:

[person] / (MR [has son] : [professional])

subsumes

[woman] / (MR [has child] : [doctor])

because of the semantics of the MR quantificational tag

MR Quantifier has inverse relational subsumption

[person whose sons are professionals]

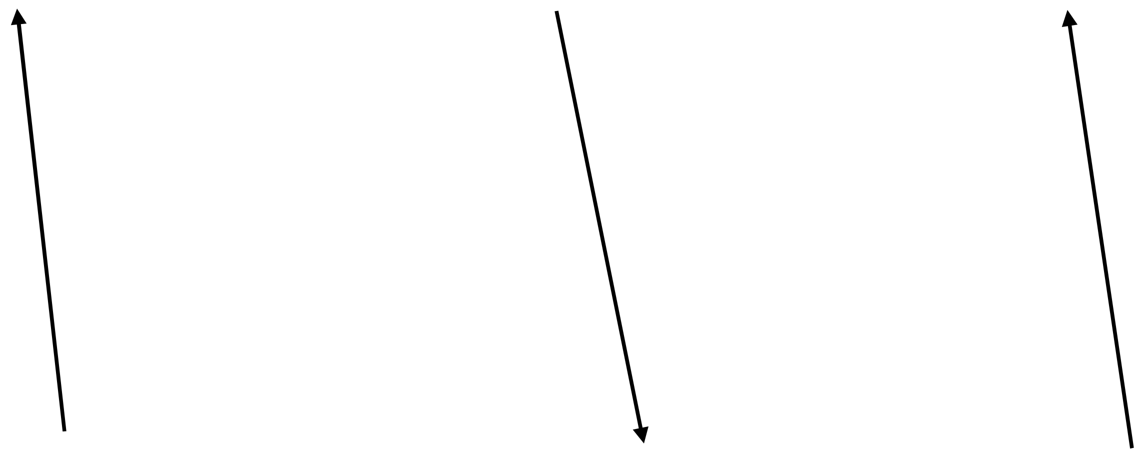

[woman whose children are doctors]

Conjecture: Subsumption Technology can Provide:

Efficient Knowledge Management for Large-Scale Automated Reasoning

- fluent at different levels of generality
- scalable to large knowledge bases
- automatically assimilating new information
- integrating different reasoning paradigms
  - probabilistic, plausibilistic, deductive, abductive
- with rapid reactive reasoning

Observation:

The range of open problems cannot be solved by mere application of one, or even a few, existing, well-understood techniques such as first-order logic and Bayesian statistics, although these are clearly parts of the solution.

We need to forge new tools sufficient to the tasks.

Practical Application:

Subsumption technology has been used to significantly improve human access to online information when combined with a specific passage retrieval algorithm that I developed at Sun Microsystems Laboratories.

But that's another talk.