

Issues in Knowledge Engineering

Definition of Artificial Intelligence

Ginsberg: the enterprise of constructing an intelligent artifact
(a "physical symbol system" that can pass the Turing test)

Disciplines:

software engineering + computer science + philosophy = cognitive science

In reality:

neat explorations into advanced (theoretical, formal) programming techniques

Neats vs Scruffies:

formal logic and proof (the Stanford approach)

vs

experimental programming (the MIT approach)

neural nets, genetic algorithms, simulation

AI Subject Matter (logic + graphs)

Knowledge Representation

formal logic

proof theory

functional and declarative programming

the "Knowledge Level"

Search

non-polynomial algorithms (NP-complete)

state space

brute force (depth first, breadth first)

heuristics

objective functions (hill climbing, best first)

adversary search (taking turns)

History of AI

Aristotle	first formal symbol system	c325 BC
Copernicus	Earth is not the center of the universe	1543 AD
Descartes	Mind \neq Body	1641
Galileo	Math is a world model	1642
Euler	state space	1735
Boole	formal logic	1847
Babbage	analytic engine	1854
Frege	formal computation	1879
Russell & Whitehead	symbolic math foundation	1910
Tarski	theory of reference and meaning	1944
Turing	test of computational intelligence	1950

Formal Knowledge

A **conceptual model** consists of

- discrete objects, presumed to exist: the Universe of Discourse
- interrelations between objects

- functions: compound names for objects and for unnamed objects

- relations: truth statements about objects

No matter how the world is conceptualized, other conceptualizations that are just as useful.

Information Processing Systems

Information = representation + transformation
(program = data structure + processes)

All representations are content free. A theory of meaning must link representation to reality.

The Physical Symbol System Hypothesis

Newell and Simon, 1972: To resolve the mind-body problem of Descartes:
Minds and computers are physical systems which manipulate symbols.

The **knowledge level** is an abstraction layer for Computer Science which unites symbolic computation with world modeling

- hardware

- assembly, microcode, machine instructions

- programming languages

- algorithms and data-structures

- symbol level

- knowledge level

Principles of Representation

Symbol systems = patterns + process

Qualitative (symbolic) information

- not numerical (although this is just a stylistic difference)

Inference of "new" knowledge from a base of facts and rules

General principles of representation

- variables, quantification, dynamic binding

Interaction and semantics

- inheritance, context, theory of meaning

Meta-reasoning

- knowledge about what is known or unknown

New control structures

- learning from examples, explanation

Declarative Style

An **AI program** consists of
a set of objects
a set of functions (names for compound objects)
a set of relations (facts)
a set of permissible transformations

Objects and relations form a **state**.
Transformations move between states.
Algorithms explore/search the **state space**.
Programmers control the search.

State Space

The collection of facts (the database) at one given time defines the **state** of the world.
All possible state configurations define the **state space**.
To move from one state to another, apply a permitted **transformation rule**.
The state space and the moves between states form a **graph**.

Predicate Calculus

A **general purpose** language for describing objects, facts, and transformations for particular domains. Also called **First Order Logic**.

logic	{and, or, if, not, iff} inference, proof
object domains	{<unique atoms>}
quantification	{all.x, exists.x}
predicates	classes and properties
relations	connections between objects
functions	indirect names

Knowledge Representation

Objects: Names are intended to point to actual concrete finite objects in the application domain. The choice of names is a part of interface/documentation. The actual names don't matter.

Variables: The **patterns** of linkage between variables with the same name determines the meaning of the variable.

Functions: Functions are **indirect names**. They name objects by telling how to get to them. Functions are compound names.

Knowledge Representation Labels

Constants:

names of specific objects: John, Tuesday, My-Phone-Number
names of specific functions: House-of[x], Phone-of[x], Truth-of[p]
names of specific relations: Likes[Mary, Tom], Phone-Number[Tom, x]

Variables:

refer to sets/classes/domains of objects
always scoped/introduced by a quantifier

Knowledge Representation Atoms

Named objects (object constants)
Indirect/compound named objects (functions)
Relations between objects (facts)

Logical connectives (and, or, not, if, iff) connect atoms. They cannot be used inside atoms.

yes: eyes-of[John] AND hair-of[John]
no: (eyes-of AND hair-of)[John]
no: hair-of[John AND Mary]
yes: hair-of[John] AND hair-of[Mary]

Knowledge Representation Quantification

Variables name **classes of objects** (all.x) and **arbitrary objects** from a class (exists.x).

Variables are bound to specific objects by the act of **instantiation**.

Quantification provides a mechanism to refer to entire classes and to arbitrary objects.

all.x P[x] every x for which P[x] is True
exists.x P[x] some arbitrary x for which P[x] is True

Model Theory

Given an object domain and a collection of functions and relations on objects in that domain, a **model** of the domain is defined by the facts:

all atoms (atom-names) are True
all atoms not in the domain are False

Eg: Domain = {Mary, Tom, John} Relation: {Likes}

all possible states: all possible models:

Likes[Mary, Tom]	1	empty (no relations true)
Likes[Mary, John]	6	one Likes relation is True
Likes[John, Mary]	15	two Likes relations True
Likes[Tom, Mary]	20	three True
Likes[John, Tom]	15	four True
Likes[Tom, John]	6	five True
	1	six True
	64	possible models in total

Entailment (implication) and Computability

$P \models Q$ (double turnstile)

All models for which the collection of facts in P are True imply that the collection of facts in Q are True for every model.

$P \vdash Q$ (single turnstile)

Using the rules of a given system, we can compute Q from P.

The central issue (1920-1970): Just because we can compute it ($P \vdash Q$), does that mean that what we compute matches our model ($P \models Q$), the ways in which the world can actually be?

Correctness

Sound: if $P \vdash Q$, then $P \models Q$
A sound computation always supports the world model.

Complete: if $P \models Q$, then $P \vdash Q$
A complete computation always generates all possible models.

Sound and Complete: $P \vdash Q = P \models Q$
The model and the computation represent the same world.

Decidability

Universal: if it can be stated formally, then it can be stated in First Order Logic.

Decidable: the computational procedure will terminate with a Yes/No result.

Semi-decidable: The computation might halt, but you don't know when. It may never halt if you ask the wrong kind of question. What we can't do is ask if questions which depend on the **failure** to prove something:

No: "Check to see if nothing is wrong"
No: "Prove that this search will fail to find X"

Representation =?= Reality

We have seen that

| - = | =

that is, that computation and formal models of reality are the same thing.

And that the LISP programming language is the flexible substrate for both representation and computation. When the field of AI began to grow in 1980s, the challenge was to actually express our models of reality in code. This turned out to be exceedingly difficult. It is not that computation is a particularly weak point, it is that our formal languages do not let us speak easily of reality.

Jerome Bruner, in his book *Actual Minds, Possible Worlds*, shows that the relationships between actual reality, language, and understanding are complex and not formal. He notes that understanding is always in context, meaning is always ambiguous. There is no reality independent of mental activity and symbolic language. We know the world and construct meaning through multiple perspectives, including emotion, culture, language, and stories.

The failure of AI to be able to parse and translate natural languages underscores this problem. What then can we model formally? The field of AI believes that sooner or later, everything can be modelled. What are the technical difficulties in trying to model anything (like making simple stacks of blocks, for example)?

Technical Difficulties in Modeling and Knowledge Representaiton (using blocks world in LISP as an example)

1. What is important to describe?

Build little theories of little worlds.

(Block A) (OnTable A) (Hand Empty)

2. How should descriptions be partitioned?

Functions or Relations, special or general objects?

(OnTable A) (On A Table) (not (OnTable Table))

3. How do we talk about groups and classes of objects?

Quantification and abstraction

(All (x) (Block x))

4. How do we address things with no names?

Functions as indirect, compound names.

(House-of John)

5. How do we handle things with more than one name?
unique name hypothesis, unification

(Uncle John) = (Brother (Father John)) = Bob

6. How do we make general rules which define the structure of relations?
quantification

(All (x) (iff (Uncle x) (Brother (Father x))))

7. How are typing and filters on domains represented?
predicates in conjunction

(All (x) (and (Person x) (Father x y)))

8. How do we join more than one fact?
conjunction

(and (F x) (G x))

9. How do we compute with logic?
inference as natural deduction and as resolution

(if (and (P x) (if (P x) (Q x))) (Q x))

10. How do we compute with quantifiers and classes of objects?
implicit universal quantification, Skolemization

(Exist (x) (P x)) ==> (P (Sk-1 x))

11. What is the difference between a fact and a query?
query combination rules

A. conjunction with negated query

(and (P x) (not (Q ?)))

B. Skolemization of query variables

(Q ?) ==> (Q Sk-1)

C. Facts imply Query

(if (P x) (Q ?))

D. The answer predicate

```
(if (P x) (Answer x))
```

12. What kinds of rules do we need for query answering?

A. definitions

```
(iff (P x) (Q (R x)))
```

B. mathematical structures (symmetry, transitivity, etc)

```
(if (and (if (P x) (Q x)) (if (Q x) (R x))) (R x))
```

C. permissible state transformations

```
(Pick-up x) = (Assert (not (onTable x)))
```

13. How can we control the inference/search procedure?

A. Pre and Post conditions

B. Compound queries

C. Searching databases of rules and facts

14. How do we steer the resolution process?

A. set of support

B. ordered resolution

C. static vs dynamic approaches (compiled vs run-time)

D. lookahead, cheapest first, dependency directed search

15. How do we express meta-level reasoning (rules about rules)
measure the savings vs brute force

16. What is the appropriate reasoning strategy?

look-up tables
natural deduction
resolution
forward or backward chaining
simulation
boundary logic

17. What do we do about contradictions and inconsistencies?

A. Forbid them

Artificial Intelligence

- B. Default reasoning
- C. Exception handlers
- D. Three-valued logic (True, False, Inconsistent)
- E. Multi-valued logic (eg True, False, Contradictory, Meaningless)
- F. Contradiction maintenance