Knowledge Representation and Reasoning

1.

Introduction

What is knowledge?

Easier question: how do we talk about it?

We say "John knows that ..." and fill the blank with a proposition

– can be true / false, right / wrong

Contrast: "John fears that ..."

– same content, different attitude

Other forms of knowledge:

- know how, who, what, when, ...
- sensorimotor: typing, riding a bicycle
- affective: deep understanding

Belief: not necessarily true and/or held for appropriate reasons

and weaker yet: "John suspects that ..."

Here: no distinction

the main idea — taking the world to be one way and not another

What is representation?

Symbols standing for things in the world

"John" ⟶ John

"John loves Mary" ⟶ the proposition that John loves Mary

Knowledge representation:

symbolic encoding of propositions believed (by some agent)

What is reasoning?

Manipulation of symbols encoding propositions to produce representations of new propositions

Analogy: arithmetic

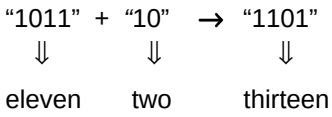

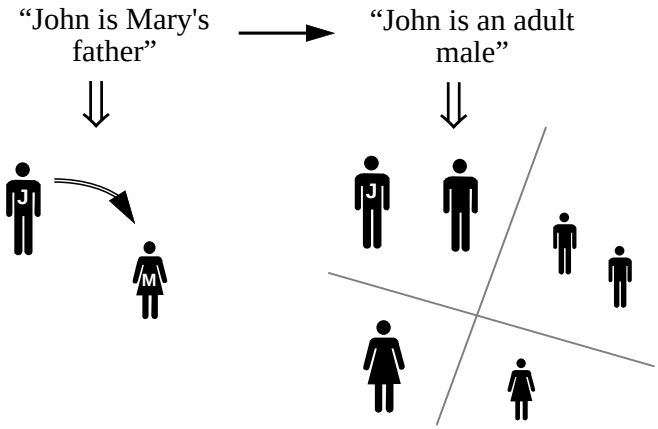

Why knowledge?

For sufficiently complex systems, it is sometimes useful to describe systems in terms of beliefs, goals, fears, intentions

e.g. in a game-playing program

"because it believed its queen was in danger, but wanted to still control the center of the board."

more useful than description about actual techniques used for deciding how to move

"because evaluation procedure P using minimax returned a value of +7 for this position

= taking an intentional stance (Dan Dennett)

Is KR just a convenient way of talking about complex systems?

- sometimes anthropomorphizing is inappropriate

  e.g. thermostats

- can also be very misleading!

  fooling users into thinking a system knows more than it does

Why representation?

Note: intentional stance says nothing about what is or is not represented symbolically

> e.g. in game playing, perhaps the board position is represented, but the goal of getting a knight out early is not

KR Hypothesis: (Brian Smith)

> "Any mechanically embodied intelligent process will be comprised of structural ingredients that a) we as external observers naturally take to represent a propositional account of the knowledge that the overall process exhibits, and b) independent of such external semantic attribution, play a formal but causal and essential role in engendering the behaviour that manifests that knowledge."

Two issues: existence of structures that

- we can interpret propositionally
- determine how the system behaves

Knowledge-based system: one designed this way!

Two examples

Example 1

```
printColour(snow) :- !, write("It's white.").
printColour(grass) :- !, write("It's green.").
printColour(sky) :- !, write("It's yellow.").
printColour(X) :- write("Beats me.").
```

Example 2

```
printColour(X) :- colour(X,Y), !,
       write("It's "), write(Y), write(".").
printColour(X) :- write("Beats me.").

colour(snow,white).
colour(sky,yellow).
colour(X,Y) :- madeof(X,Z), colour(Z,Y).
madeof(grass,vegetation).
colour(vegetation,green).
```

Both systems can be described intentionally.

Only the 2nd has a separate collection of symbolic structures à la KR Hypothesis

its knowledge base (or KB)

∴ a small knowledge-based system

KR and AI

Much of AI involves building systems that are knowledge-based

ability derives in part from reasoning over explicitly represented knowledge

- language understanding,
- planning,
- diagnosis,
- "expert systems", etc.

Some, to a certain extent

game-playing, vision, etc.

Some, to a much lesser extent

speech, motor control, etc.

Current research question:

how much of intelligent behaviour is knowledge-based?

Challenges: connectionism, others

Why bother?

Why not "compile out" knowledge into specialized procedures?

- distribute KB to procedures that need it

  (as in Example 1)
- almost always achieves better performance

No need to think. *Just do it!*

- riding a bike
- driving a car
- playing chess?
- doing math?
- staying alive??

Skills (Hubert Dreyfus)

- novices think; experts *react*
- compare to current "expert systems":

  knowledge-based !

Advantage

Knowledge-based system most suitable for *open-ended* tasks

can structurally isolate *reasons* for particular behaviour

Good for

- explanation and justification
  - "Because grass is a form of vegetation."
- informability: debugging the KB
  - "No the sky is not yellow. It's blue."
- extensibility: new relations
  - "Canaries are yellow."
- extensibility: new applications
  - returning a list of all the white things
  - painting pictures

Cognitive penetrability

Hallmark of knowledge-based system:

the ability to be <u>told</u> facts about the world and adjust our behaviour correspondingly

for example: read a book about canaries or rare coins

Cognitive penetrability (Zenon Pylyshyn)

actions that are conditioned by what is currently believed

an example:

we normally leave the room if we hear a fire alarm

we do not leave the room on hearing a fire alarm
if we believe that the alarm is being tested / tampered

can come to this belief in very many ways

so this action is cognitively penetrable

a non-example:

blinking reflex

Why reasoning?

Want knowledge to affect action

not do action A if sentence P is in KB

but do action A if world believed in satisfies P

Difference:

P may not be *explicitly* represented

Need to apply what is known in general to the particulars of a given situation

Example:

"Patient x is allergic to medication m."

"Anybody allergic to medication m is also allergic to m'."

Is it OK to prescribe m' for x ?

Usually need more than just DB-style retrieval of facts in the KB

Entailment

Sentences $P_1, P_2, ..., P_n$ entail sentence P iff the truth of P is implicit in the truth of $P_1, P_2, ..., P_n$.

If the world is such that it satisfies the P_i then it must also satisfy P.

Applies to a variety of languages (languages with truth theories)

Inference: the process of calculating entailments

- sound: get only entailments
- complete: get all entailments

Sometimes want unsound / incomplete reasoning

for reasons to be discussed later

Logic: study of entailment relations

- languages
- truth conditions
- rules of inference

Using logic

No universal language / semantics

- Why not English?
- Different tasks / worlds
- Different ways to carve up the world

No universal reasoning scheme

- Geared to language
- Sometimes want “extralogical” reasoning

Start with first-order predicate calculus (FOL)

- invented by philosopher Frege for the formalization of mathematics
- but will consider subsets / supersets and very different looking representation languages

Knowledge level

Allen Newell's analysis:

- Knowledge level: deals with language, entailment
- Symbol level: deals with representation, inference

Picking a logic has issues at each level

- Knowledge level:
  expressive adequacy,
  theoretical complexity, ...
- Symbol level:
  architectures,
  data structures,
  algorithmic complexity, ...

Next: we begin with FOL at the knowledge level