

# **What is Knowledge Representation?**

Davis, Shrobe and Szolovitz

AI Magazine, 1993

# Five roles that KR plays

- 1 A surrogate for some part of the real world
- 2 A set of ontological commitments
- 3 A fragmentary theory of intelligent reasoning
- 4 A medium for pragmatically efficient computation
- 5 A medium of human expression

# 1 KR as a surrogate

- **Agents “reason” about models of the world to**
  - Deduce properties without having to directly gather information from the world
  - Predict consequences of potential actions rather than performing the actions directly
- **Given a KR, there are two questions to ask:**
  - Semantics -- *For what* is it a surrogate?
  - Fidelity -- How accurate is it?
- **When modeling the natural world, KRs are always imperfect**
  - Consequently, even with a sound reasoning system, incorrect conclusions are inevitable!

## 2 Ontological commitments

- A KR is a set of ontological commitments
- An ontology is a theory of *what exists* in the world
  - Classes, objects, relations, attributes, properties, constraints, special individuals, etc.
  - We could also view this as providing a vocabulary.
- A KR makes a commitment to a particular ontology – i.e., to describing the world with particular terms.
- “The commitments are in effect a strong pair of glasses that determine what we can see, bringing some part of the world into sharp focus, at the expense of blurring other parts.”

# Example of ontological commitments

## Problem: representing electronic circuits

- In a “lumped element model”
  - a circuit consists of
    - Components with terminals that have certain I/O behaviors
    - Connections between terminals
  - signals flow instantaneously along the connections.
- Another model may represent the electrodynamics of the situation
  - Signals propagate with finite speed

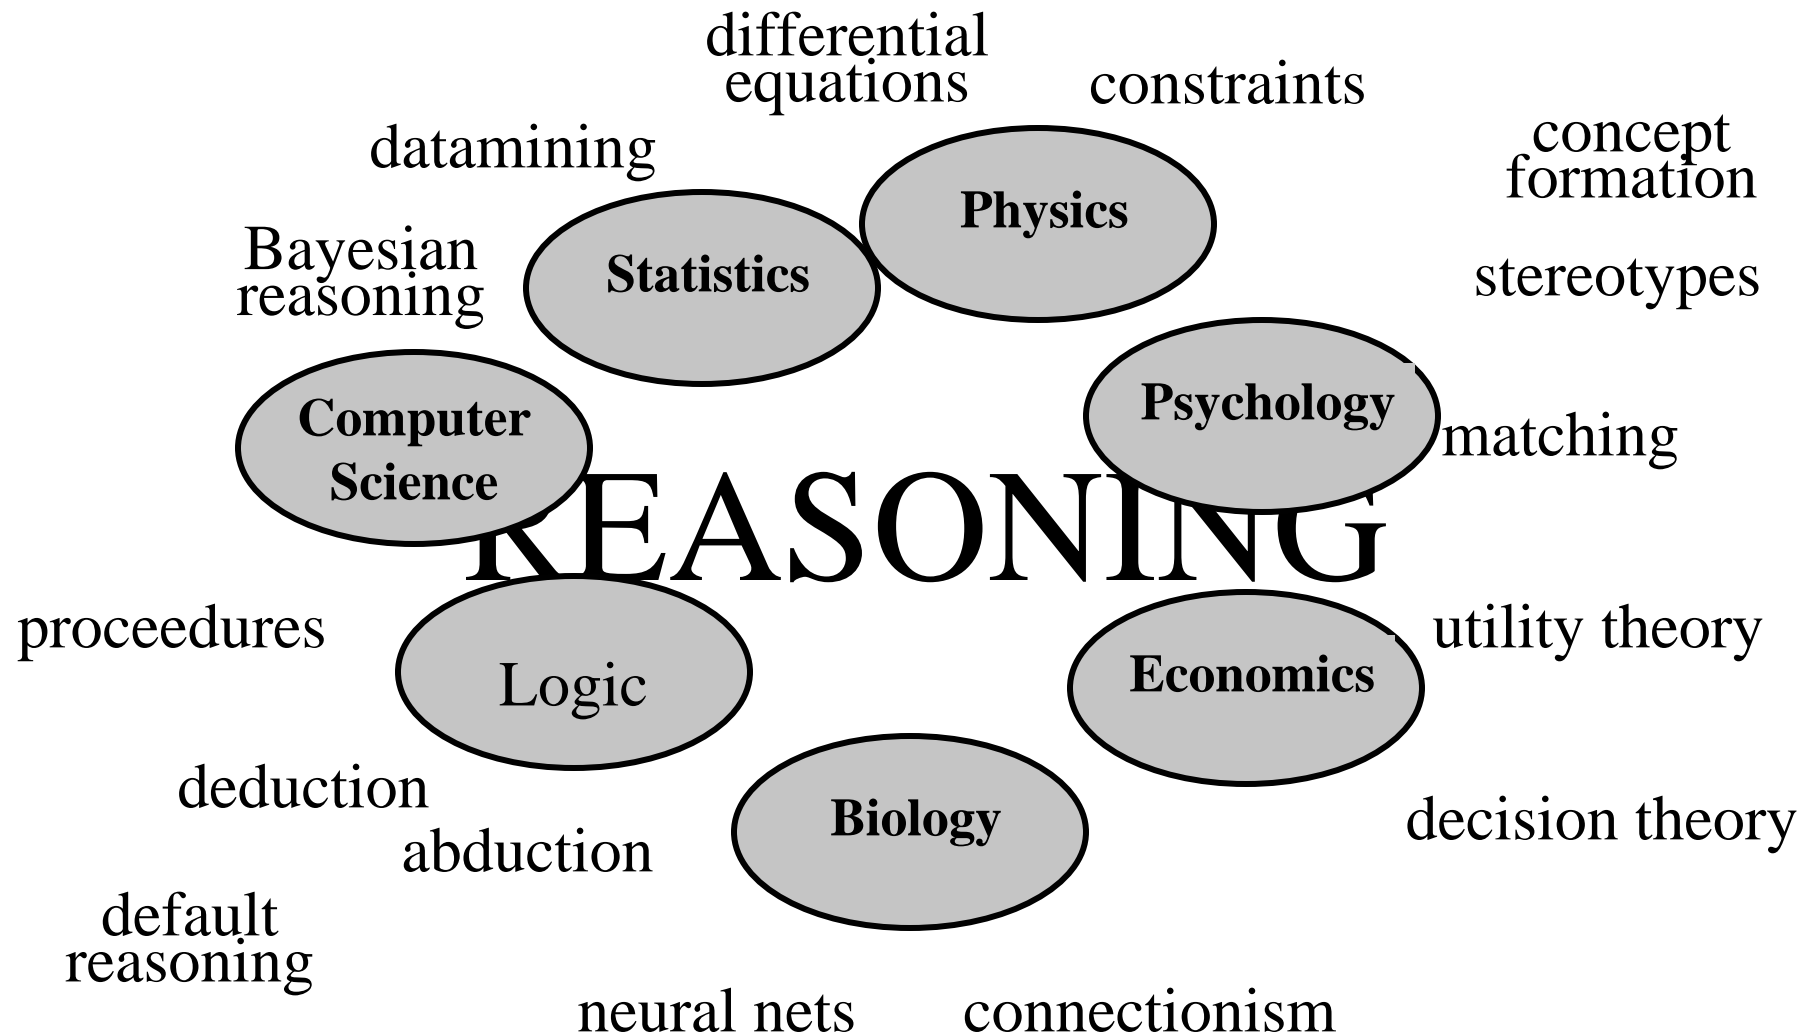
# Note

- Every representation ignores *something* about the world
- A KR is not just a data structure.
  - “ Part of what makes a language representational is that it carries meaning, i.e., there is a correspondence between its constructs and things in the external world”.

# 3 KR as a theory of reasoning

- Many knowledge representations offer fragmentary theories of *intelligent reasoning*.
- Most agree that humans (and animals) employ multiple strategies for representing and reasoning about the world:
  - E.g., deductive, abductive, inductive, Bayesian, case-based, etc.
- Three components of a theory:
  - (1) The fundamental *conception* of intelligent inference
  - (2) *Sanctioned* inferences
  - (3) *Recommended* inferences

# (1) What counts as inference?





## (2) & (3) A theory of reasoning

- Sanctioned vs. recommended reasoning.
- Most, but not all, logical systems assume a sound reasoner.
  - An example of non-sound reasoning is *abductive reasoning*.
  - Humans do pretty well with non-sound reasoning
- More interesting is having a theory of what inference steps to make
- There's been lots of work done on different reasoning strategies.

## 4 Efficient computation

- Some KR languages and frameworks have focused on “heuristic adequacy” – providing a representation which supports adequately efficient problem solving. E.g.
  - Early heuristic systems
  - Any-time computations
- Other work has focused on the knowledge content and what could, in principle, be derived from it without concern for efficiency except in fairly abstract measures (e.g., complexity), for example
  - Naïve physics
- Both perspectives are important, so avoiding either pole (or including both) is a good idea.

# 5 Human expression

- Knowledge representations can be intended for humans as well as machines.
- Some KR languages attempt to be easily generated and understood by people and others do not.
  - E.g., neural networks vs. rule based systems
  - E.g., OO representations vs. logical formalisms
- Individual talents, experience and training makes a difference, of course.

# Summary

- We're talking about KR broadly, including KR frameworks, KR languages, and representations of some particular knowledge.
- These have different roles or aspects.
- Any particular approach or piece of work will focus on some of the roles and not others.
- Keeping this in mind will help you in understanding the work and its importance.