Cognitive Architecture

Knowledge, Patterns, and Reasoning

Eric Saund, Ph.D.
October, 2019

www.saund.org

Architecture

Brains and Minds

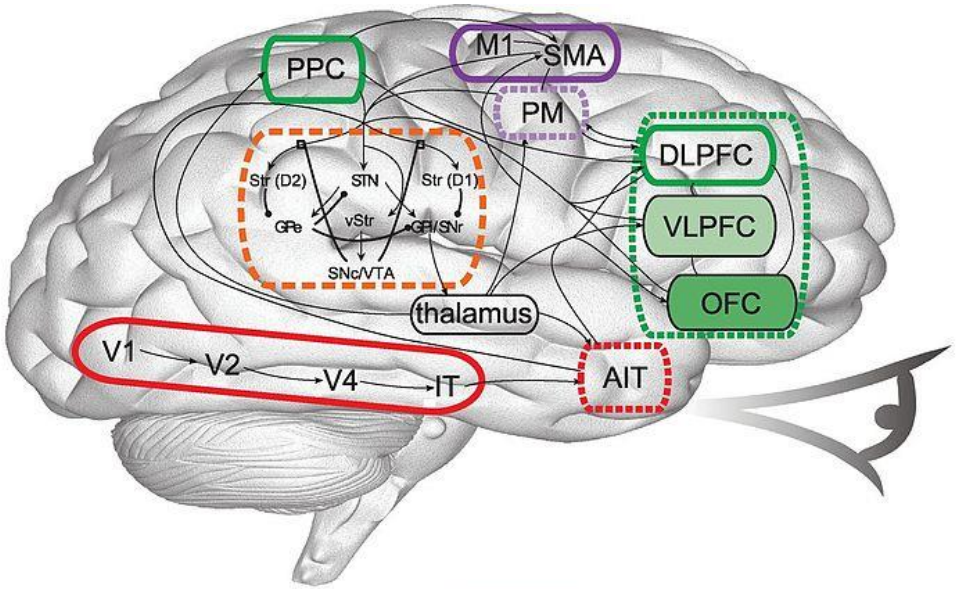

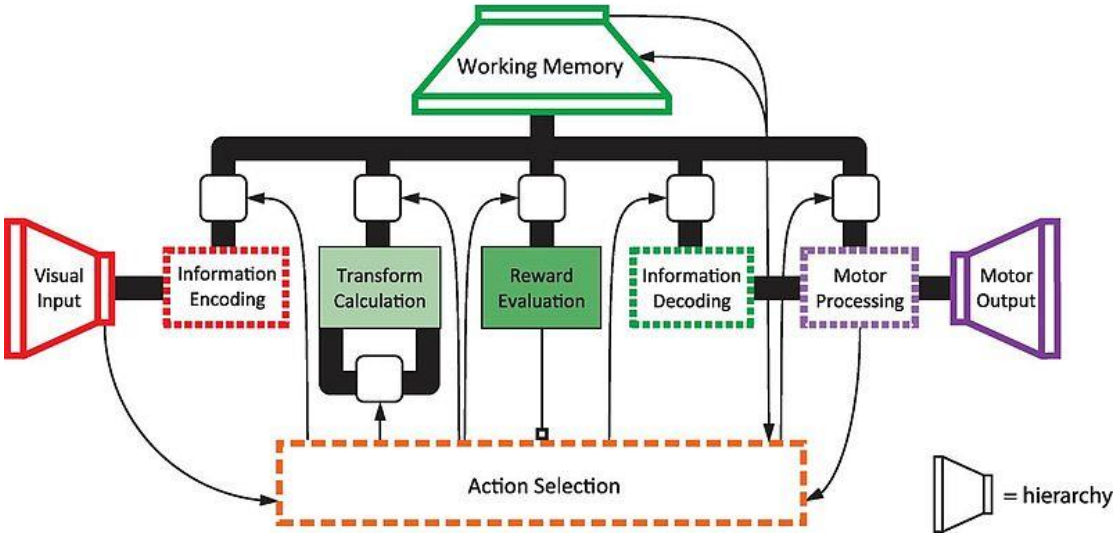

Cognitive Architecture: Outline

- Architecture in Information Systems
- History
  - Computational theory of mind
- The Standard Model Cognitive Architecture
  - Example: Soar
- Important Concepts
  - Marr's Three Levels
  - Reactive vs Deliberative
- Architecture in NN / Deep Learning Networks
- Flight of Imagination
  - LIDA, CopyCat
- Architecture of a Conversational Agent

Information Architecture: A Book

Table of Contents

page numbers

logical sections

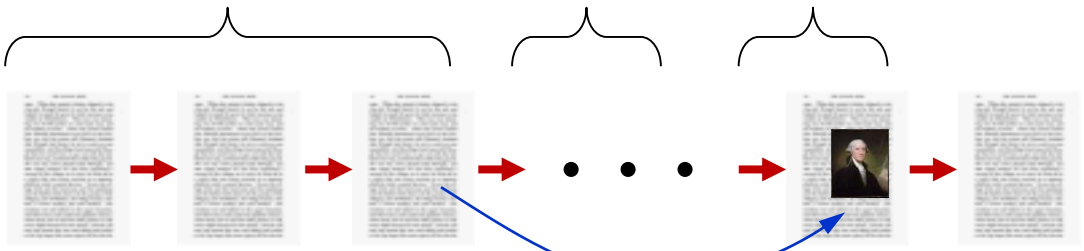

content Pages

cross-references

page numbers

content items

Index

Information Architecture: A Web Site

HTML1.0

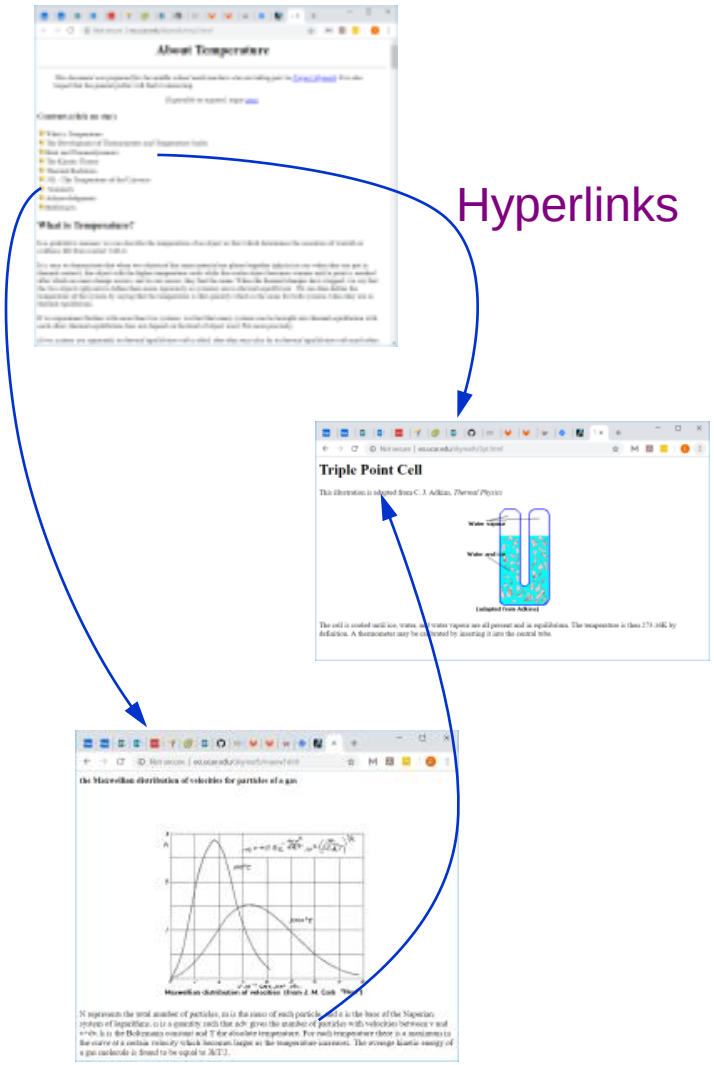

HTML5/CSS/Javascript

- active nav widgets
- guided tours
- search
- state
  - cookies
  - navigation history
- contingent access (passwords)
- personalization

Information Architecture: A Computer

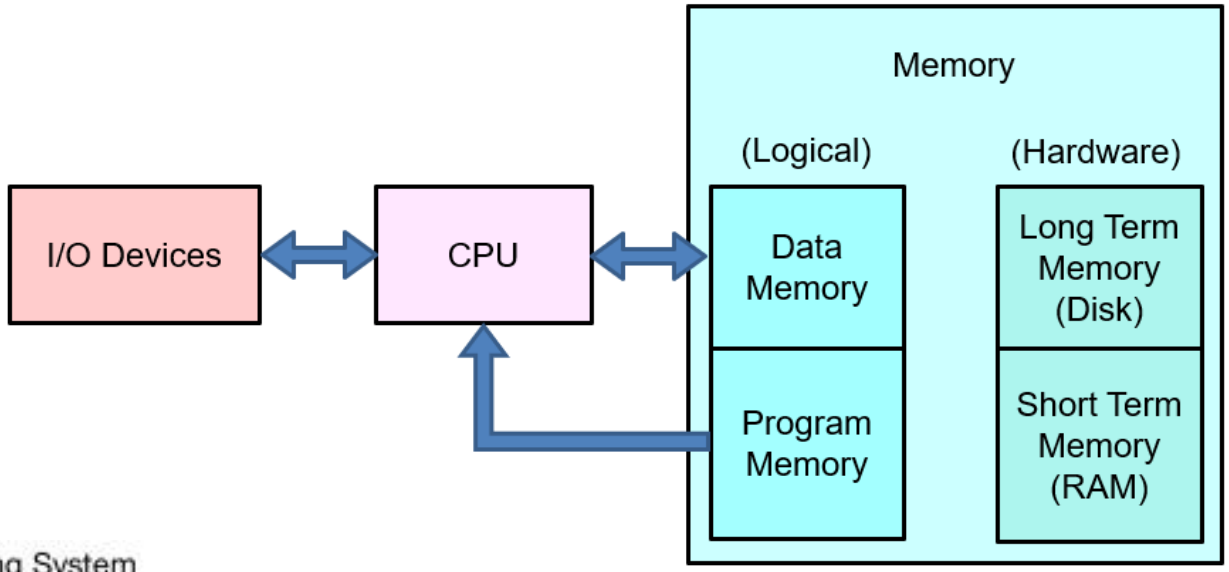

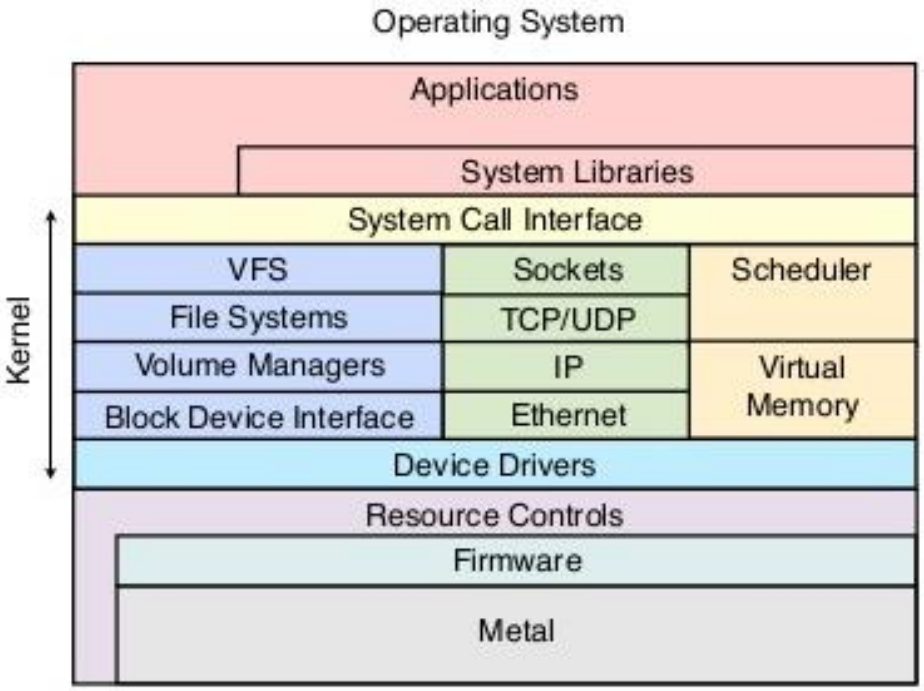

Cognitive Architecture: Basic Agent

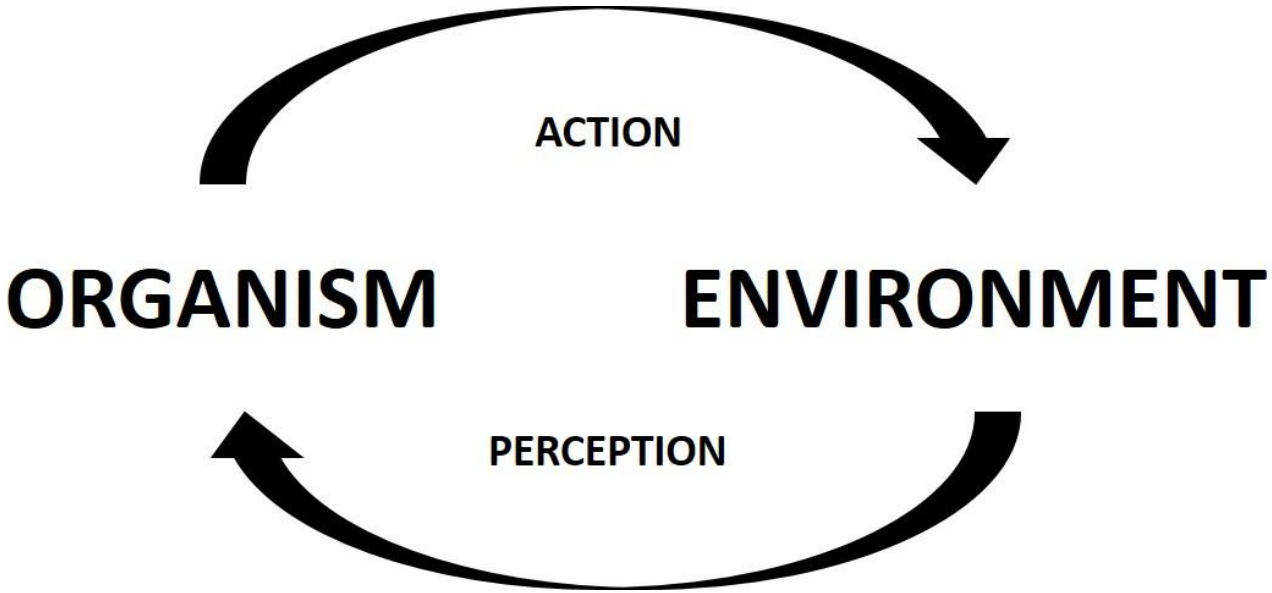

http://psychsciencenotes.blogspot.com/2018/07/you-cannot-perceive-relational.html

Cognitive Architecture: Basic Agent

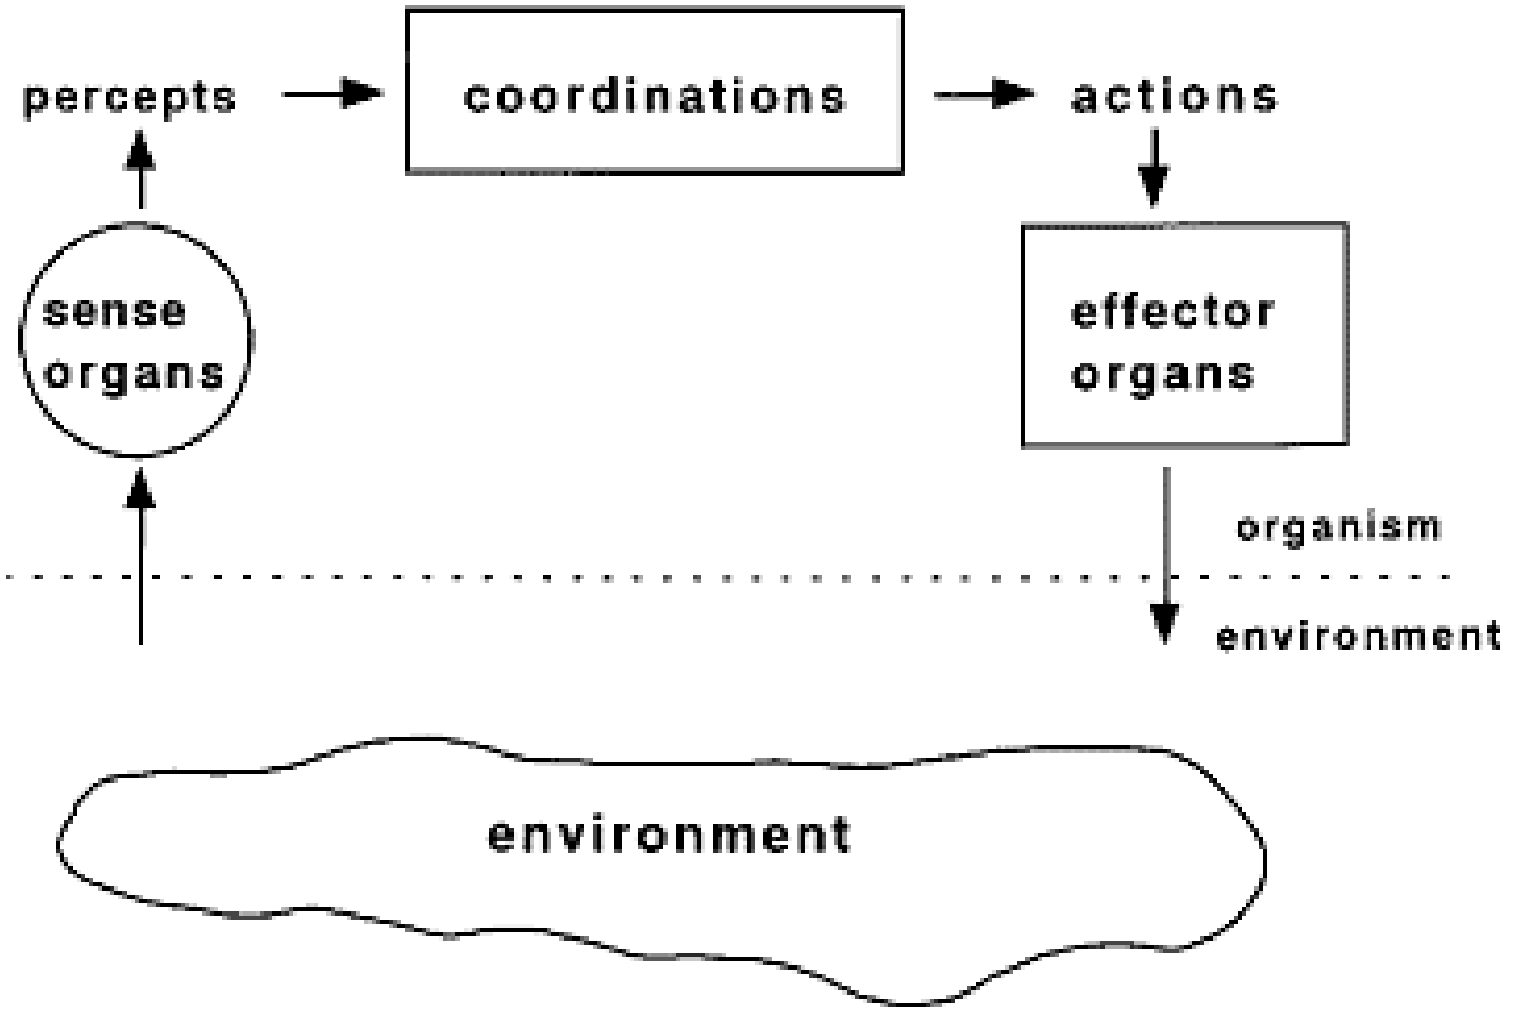

Cognitive Architecture: Historical Roots

Psychoanalytic Theories
Freud, Jung

Perceptual &
Phenomenal Psychology
Helmholtz, William James

Cybernetics
Norbert Wiener

Artificial Intelligence
McCarthy, Minsky

Theory of Computation
Turing, von Neumann

Computational Theory of Mind
McCulloch and Pitts
Chomsky
Newell & Simon

Behaviorism
Pavlov, B.F. Skinner

Guiding Metaphor

- engines and hydraulics
- signal transmission
- signal processing
- calculating machines
- computers

https://plato.stanford.edu/entries/computational-mind/

Cognitive Architecture: Standard Model

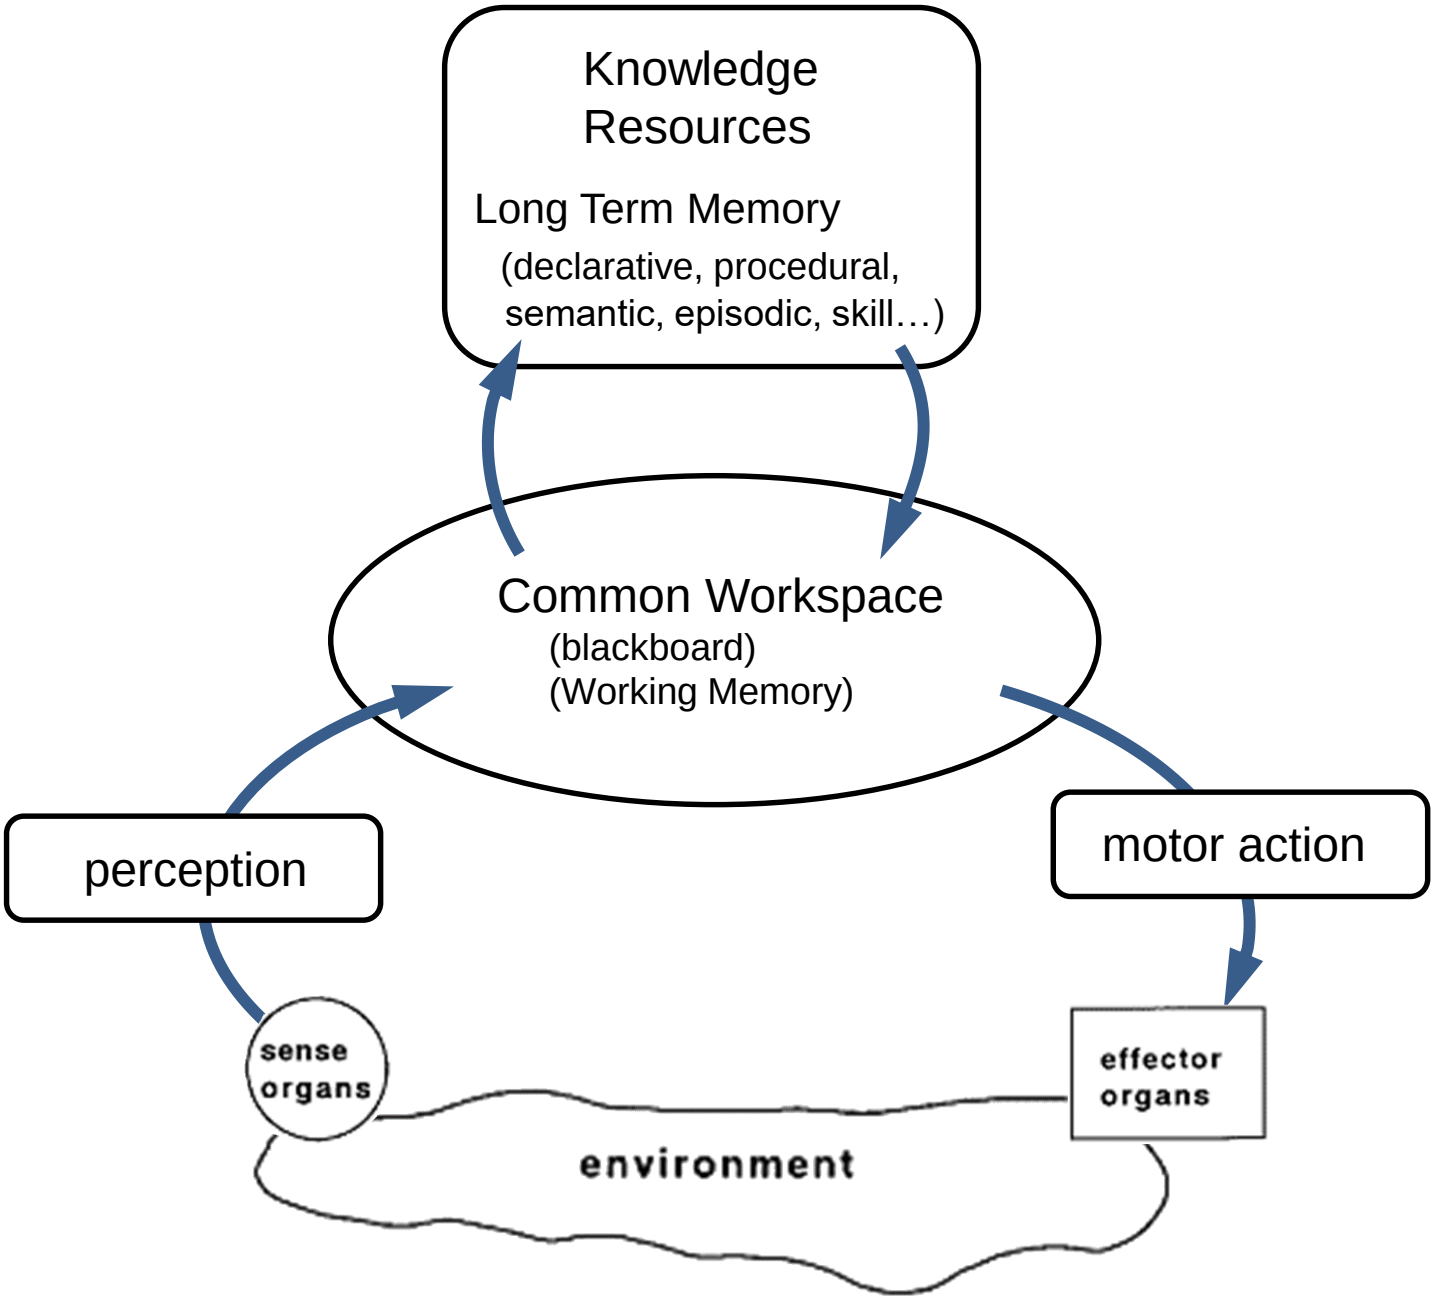

Cognitive Architecture: Big Questions

- What are the *types of content* held in the workspace?
  - percepts
  - beliefs
  - memories
  - goals
  - intentions & plans
  - emotions
  - attitudes
- What are the *representations* for state and knowledge?
  - activation patterns over fixed vectors
  - graphs of objects and relations
  - frequencies and phases of waveforms
- How is processing *controlled*?
  - automatic processes
  - conscious deliberation
  - selection among choices

Soar

(Newell, Laird,
1983 -> present)

- Definition of intelligence:
  - problem states and transitions
  - solutions found through *search* in state space
- Representation:
  - graphs of objects and relations
- Control: production system
  - Working Memory blackboard
  - procedural knowledge
  - declarative knowledge

} Long-Term Memory

Soar: Water Jug Problem Example

Definition of intelligence:

- problem states and transitions
- solutions found through *search* in state space

5 gal.

3 gal.

Start state: both jugs empty.

Goal state: 3-gallon jug contains 1 gallon of water.

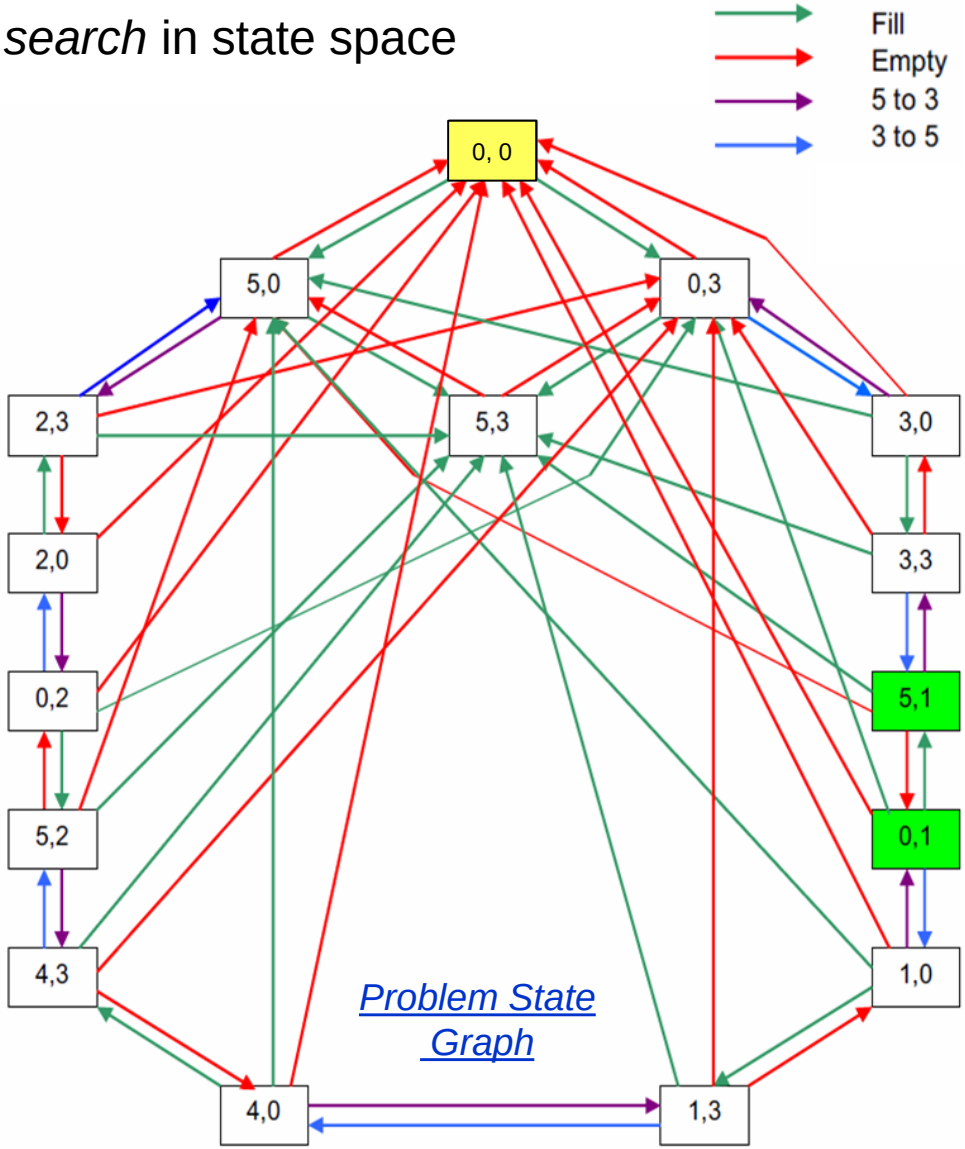

Representation in Soar

Workspace State Graph

- data objects
- attributes & relations
- operators
- Working Memory (state)
- Long-Term Memory (knowledge)

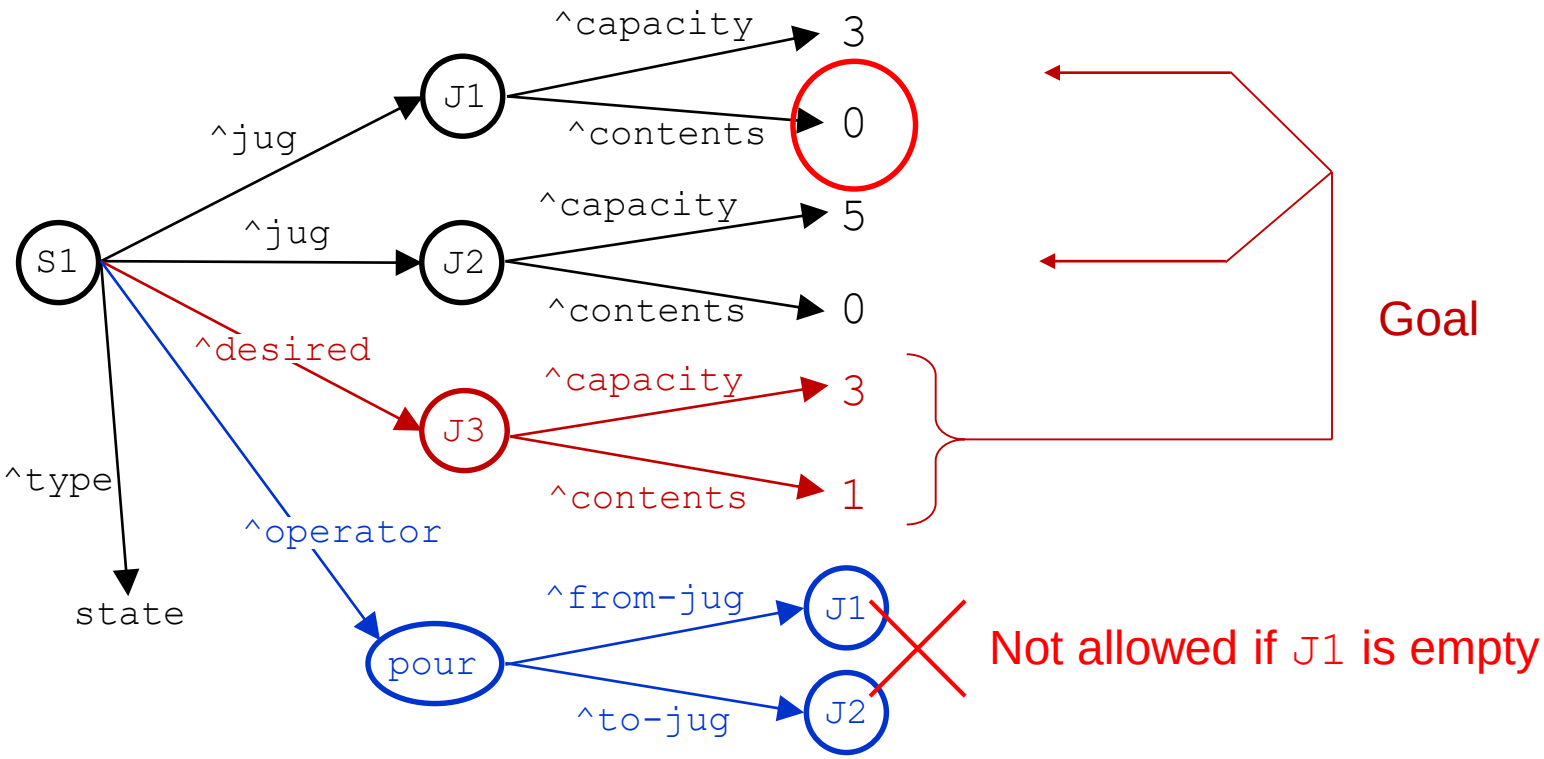

Representation in Soar

Workspace State Graph

- data objects
- attributes & relations
- operators
- Working Memory (state)
- Long-Term Memory (knowledge)

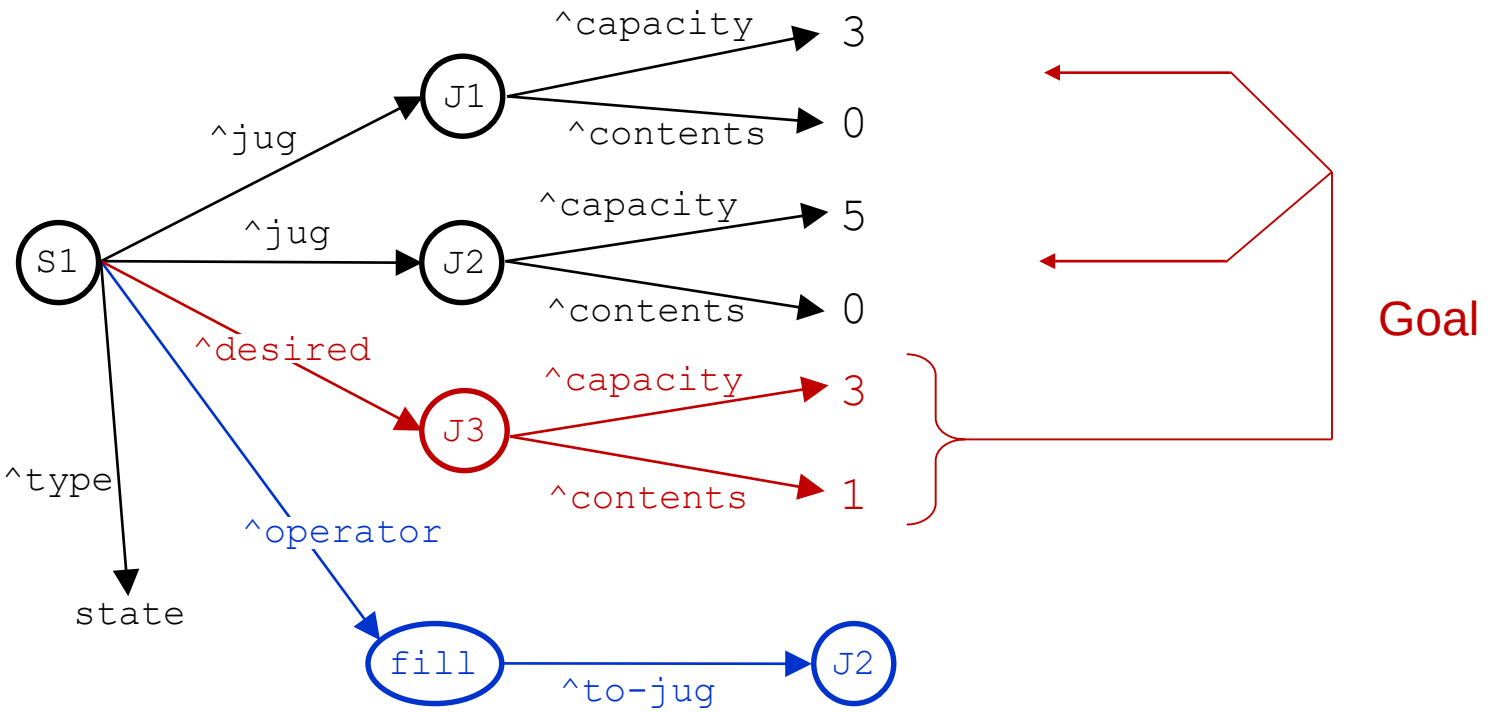

Control in Soar

Production System

- Working Memory blackboard
- declarative knowledge - what
- procedural knowledge – how
  - rules
  - operators
- subgoal states

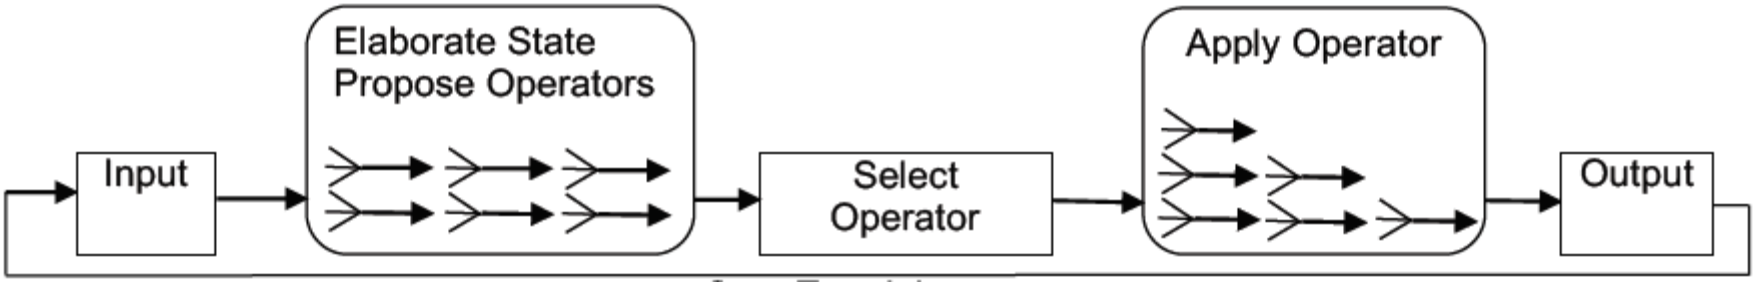

Processing Cycle

Executive Function (Psychology, Cognitive Neuroscience):

- update Working Memory from sensory and Long-Term Memory resources
- focus attention, inhibit distractors
- shift task context

Cognitive Architecture: Outline

- Architecture in Information Systems

- History
  - Computational theory of mind

- The Standard Model Cognitive Architecture
  - Example: Soar
- Important Concepts
  - Marr's Three Levels
  - Reactive vs Deliberative
- Architecture in NN / Deep Learning Networks
- Flight of Imagination
  - LIDA, CopyCat
- Architecture of a Conversational Agent

Production Rules vs. Neurons

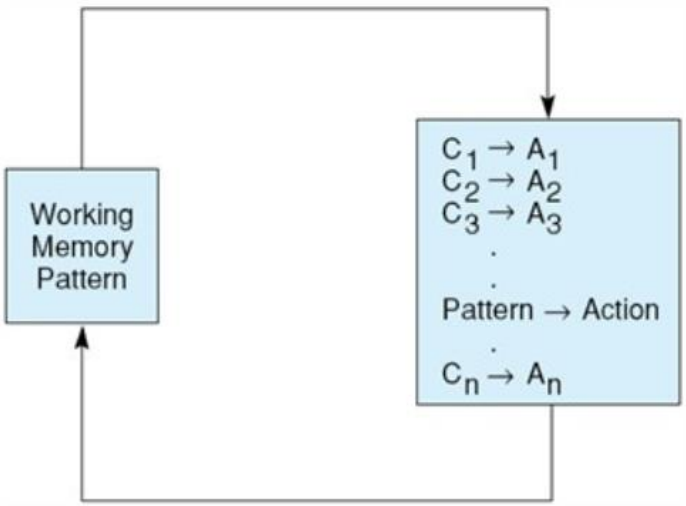

?

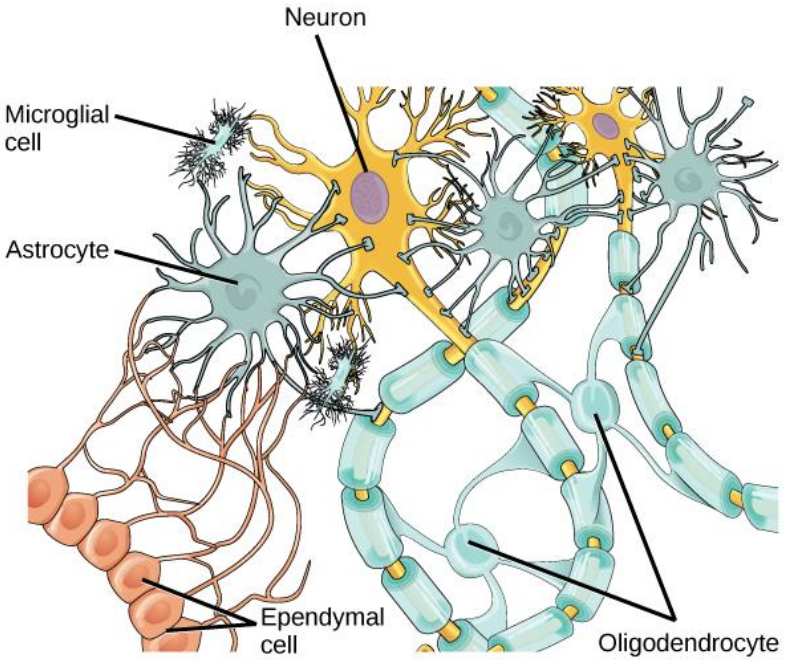

Marr's Three Levels of Abstraction

David Marr: *Theoretical Neuroscience* → *Computational Intelligence*

what? why?

Example

- **Computational Theory**

  What is the computation and by what principles is it accomplished?

  Textbook:
  ToC, body text, index

- **Algorithm**

  What representations and algorithms are used to carry forth computation?

  Scrolling text roll, manual + automatic positioning

- **Implementation**

  On what physical hardware and firmware is the algorithm run?

  Projection of markings printed on acetate film, keyboard controls scrolling per ToC and index.

Cognitive Architecture: Reactive Agent

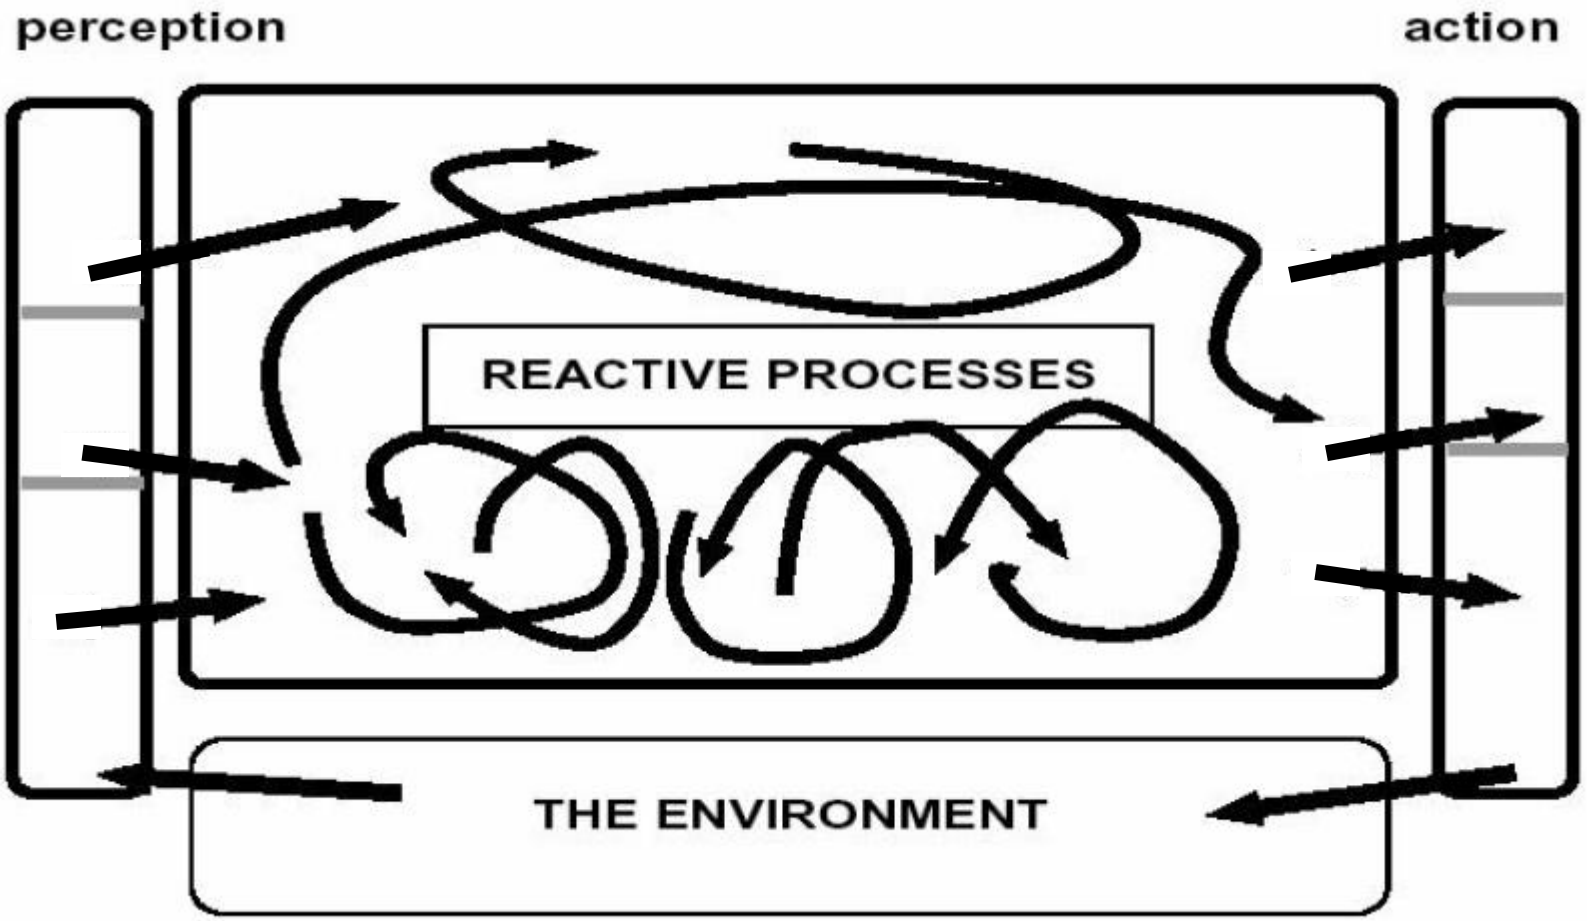

Cognitive Architecture: Deliberative Agent

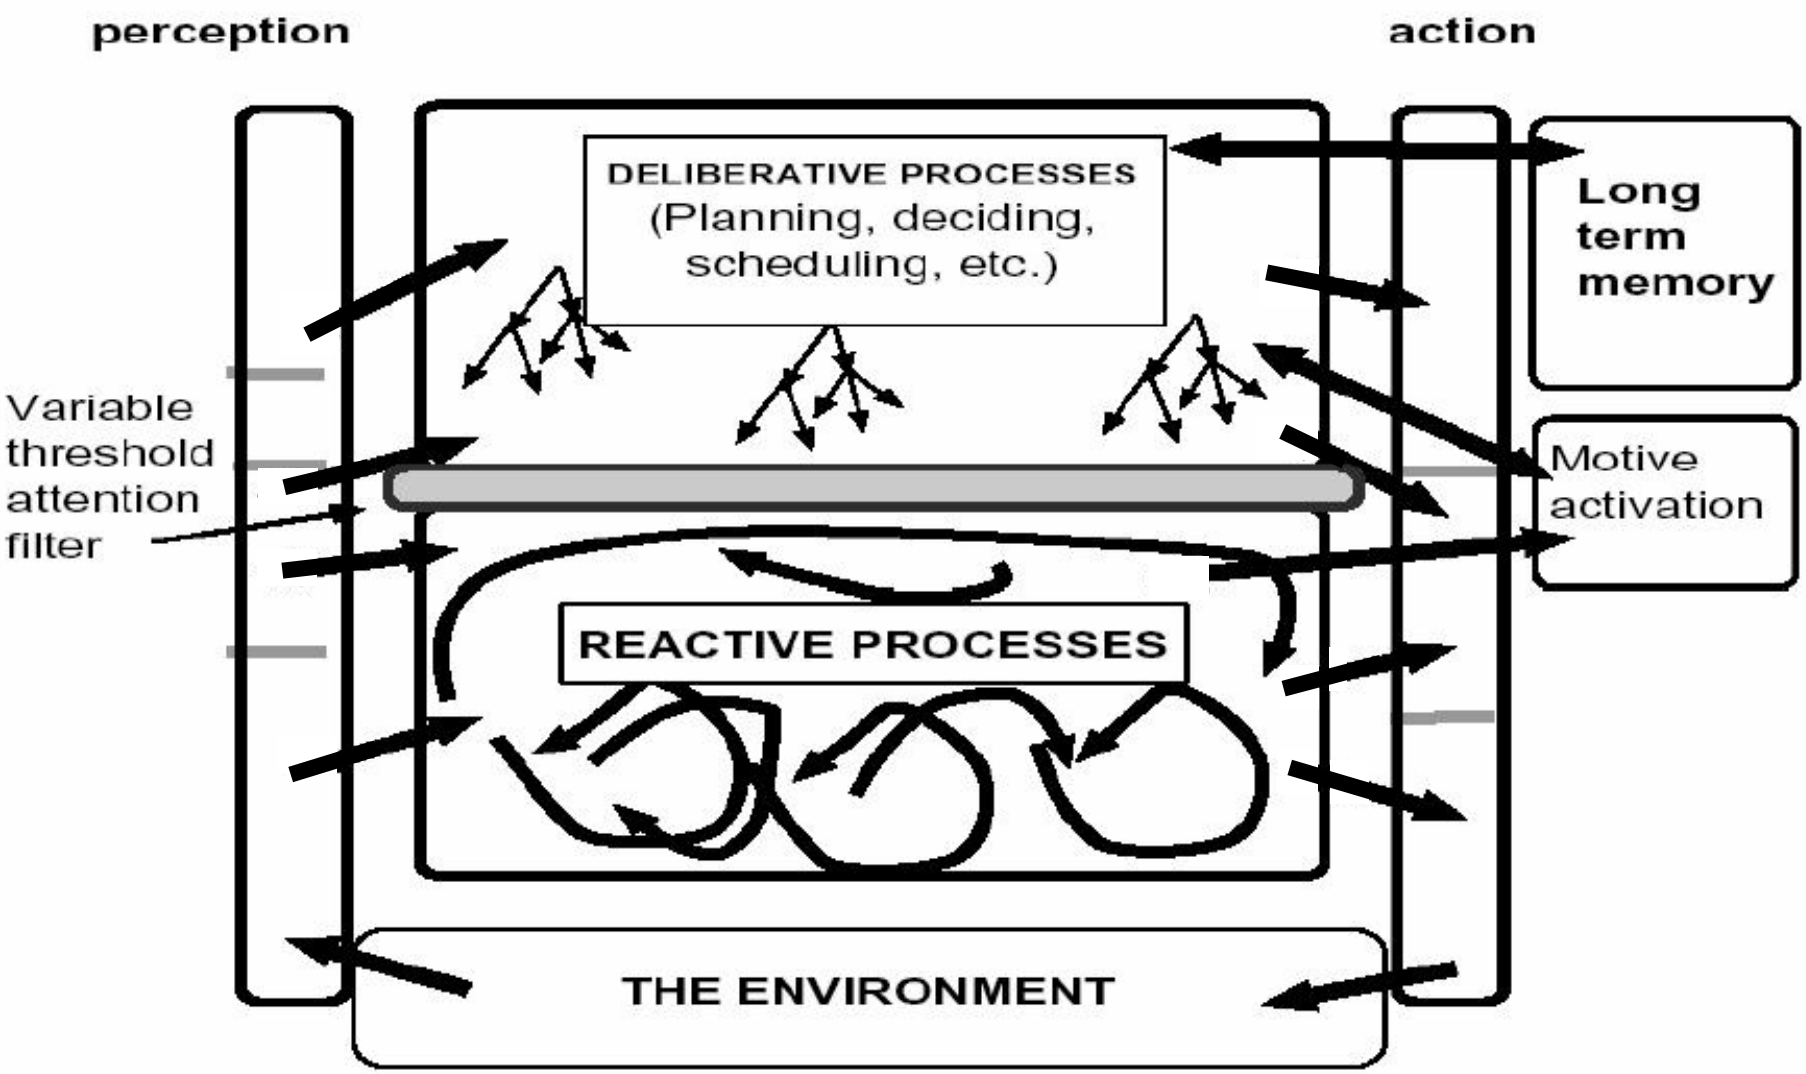

Cognitive Architecture: Reflective Agent

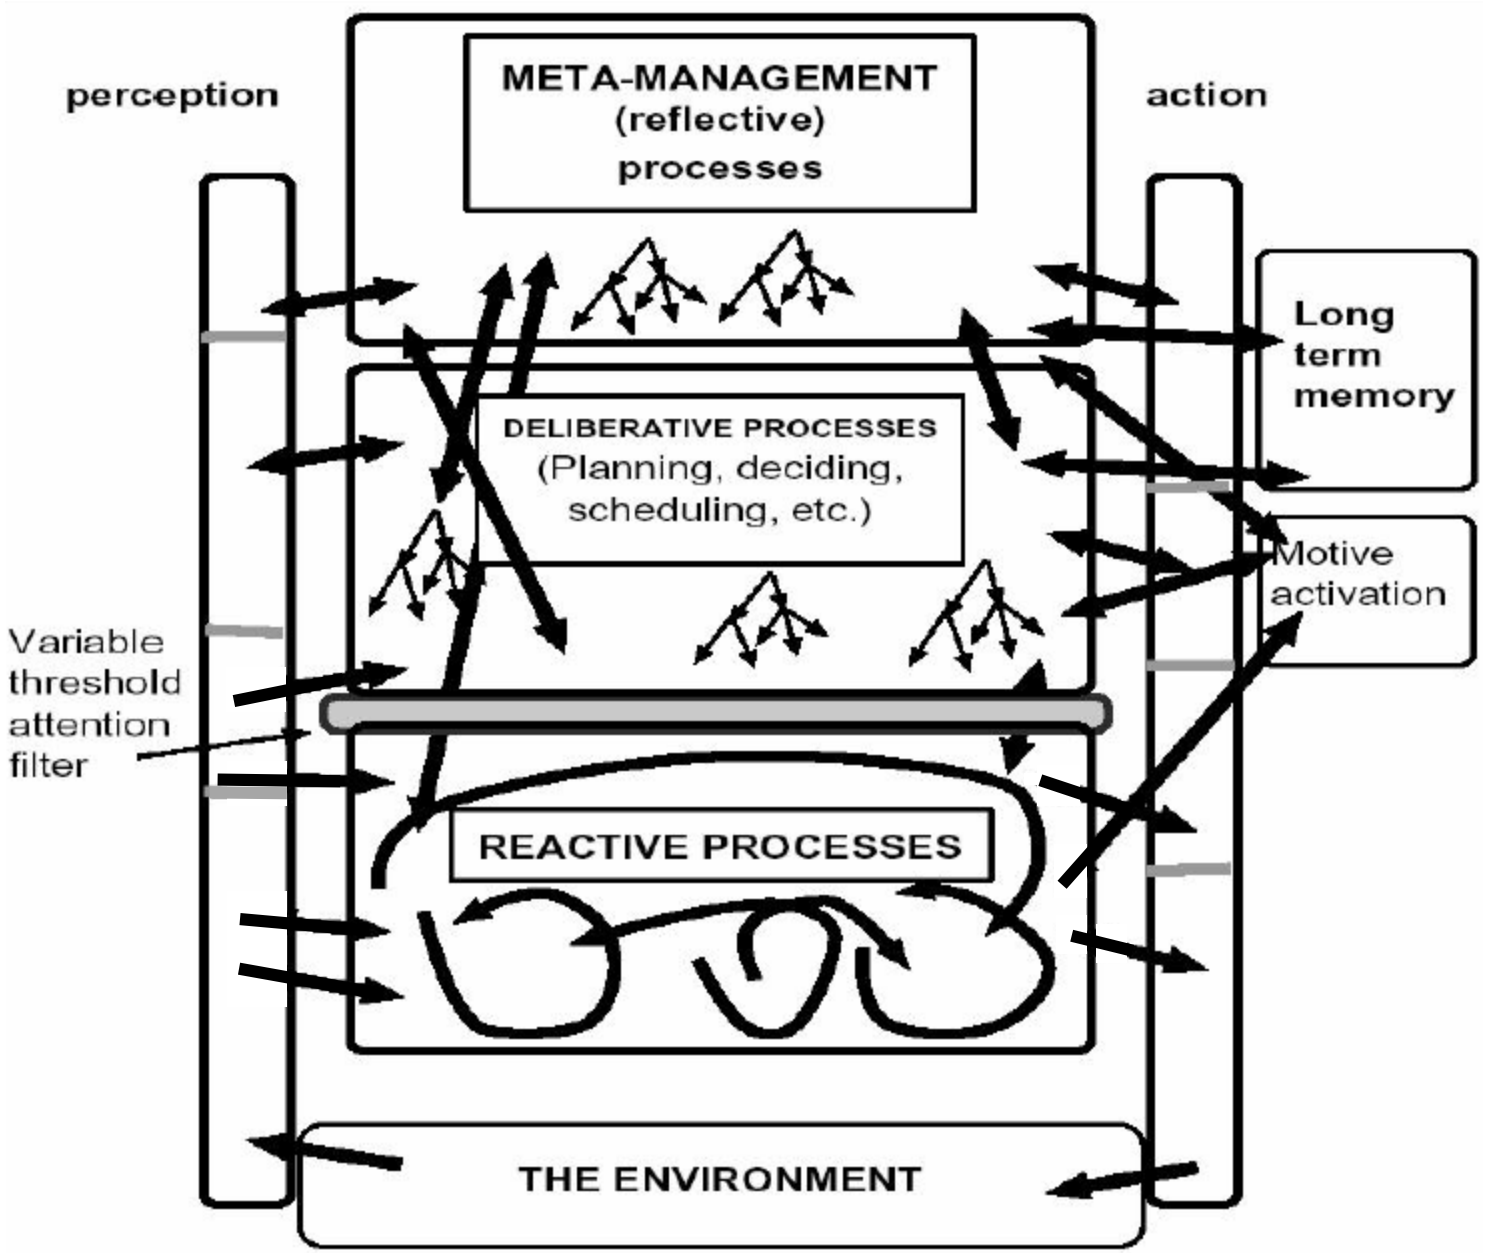

Aaron Sloman

Reactive vs. Deliberative

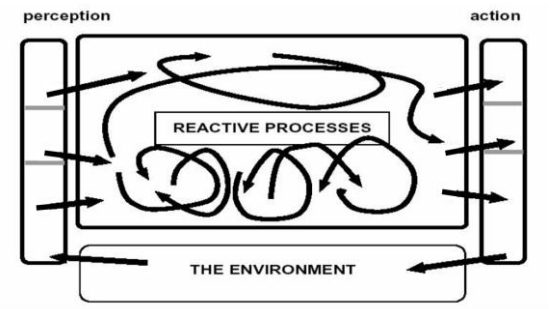

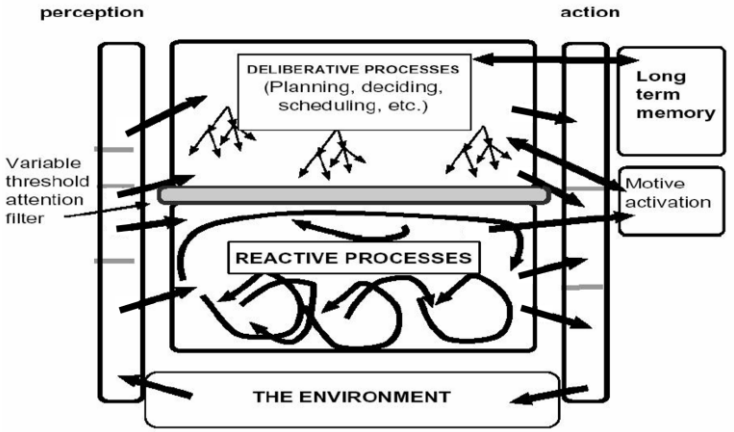

Reactive

automatic & strictly determined

modest internal state

implicit representations

Kahneman System 1

Example: thermostat

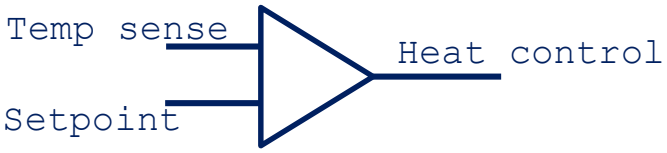

Deliberative

makes choices

rich internal state

explicit world models

Kahneman System 2

Example: building temperature management system

Reactive vs. Deliberative Building Temperature Controller

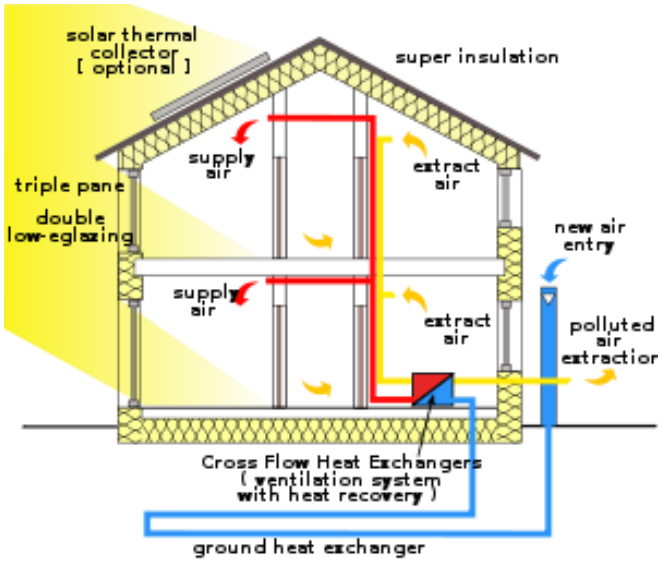

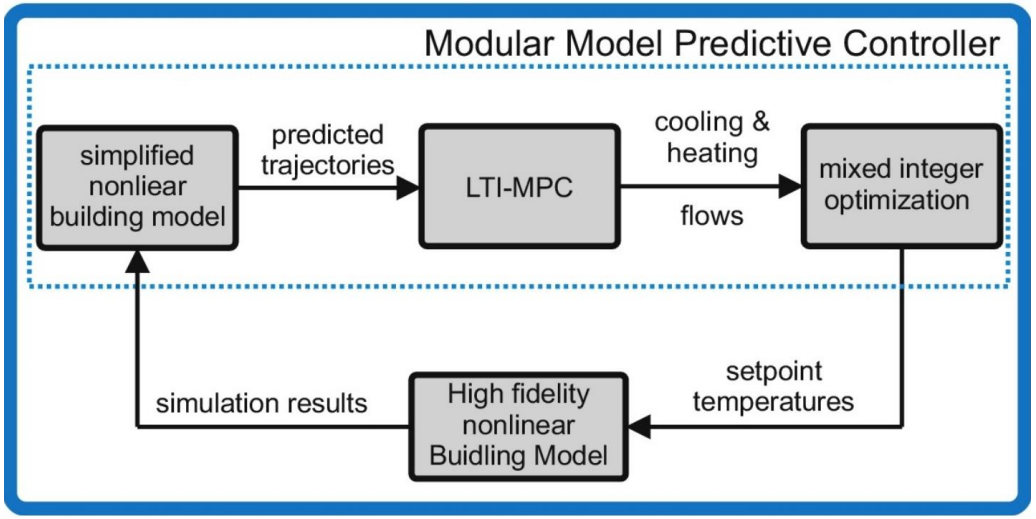

Example: thermostat

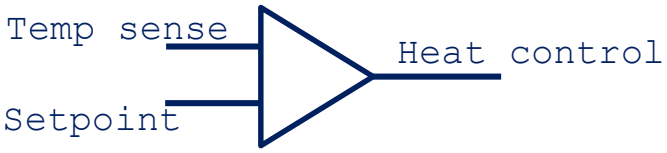

Example: building temperature management system

Reactive vs. Deliberative

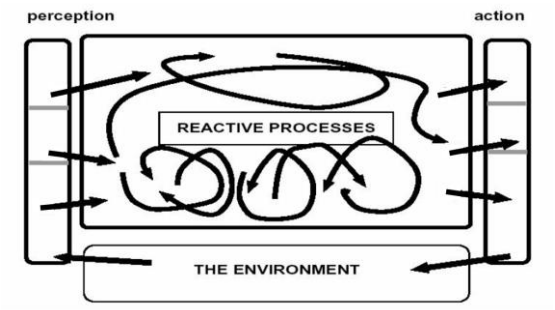

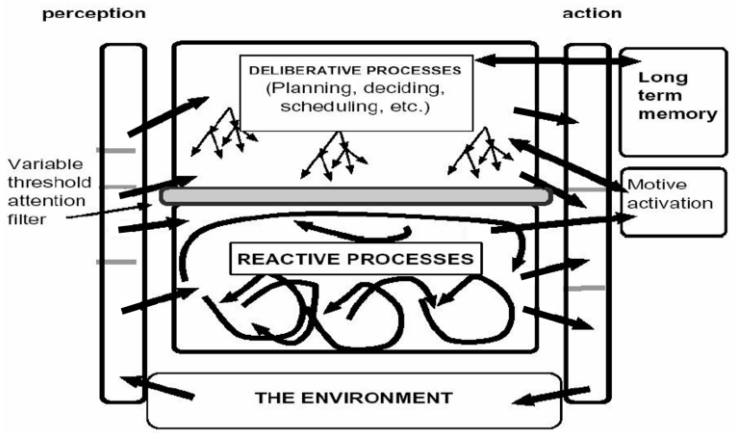

Reactive

automatic & strictly determined

modest internal state

implicit representations

Deliberative

makes choices

rich internal state

explicit world models

?

Soar

A Reactive Water Jug Solution in Soar

Deliberative search space

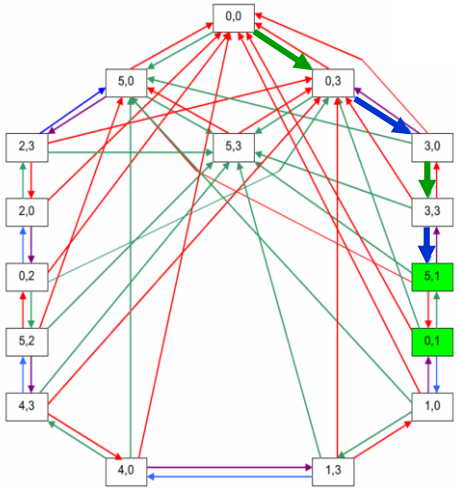

jug-5, jug-3

Reactive program rules

```
If   [0, 0]   Then   FILL(3)      =>   [0, 3]
If   [0, 3]   Then   Pour(3,5)    =>   [3, 0]
If   [3, 0]   Then   Fill(3)      =>   [3, 3]
If   [3, 3]   Then   Pour(3,5)    =>   [5, 1]   ✓
```

Neural Network Architectures

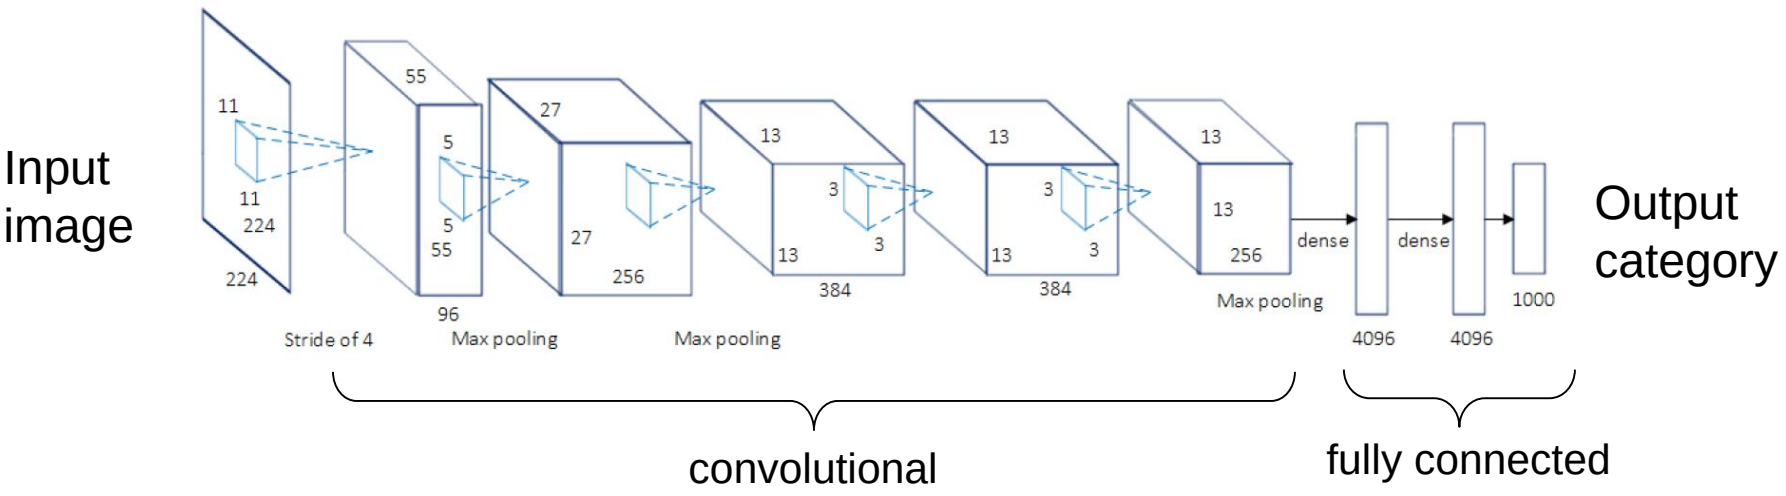

Alexnet

- Architectural Elements — Layer dimensions, weights, nonlinearities
- Organization — Layer connectivity
- Purpose — Function approximation

Neural Network Architectures

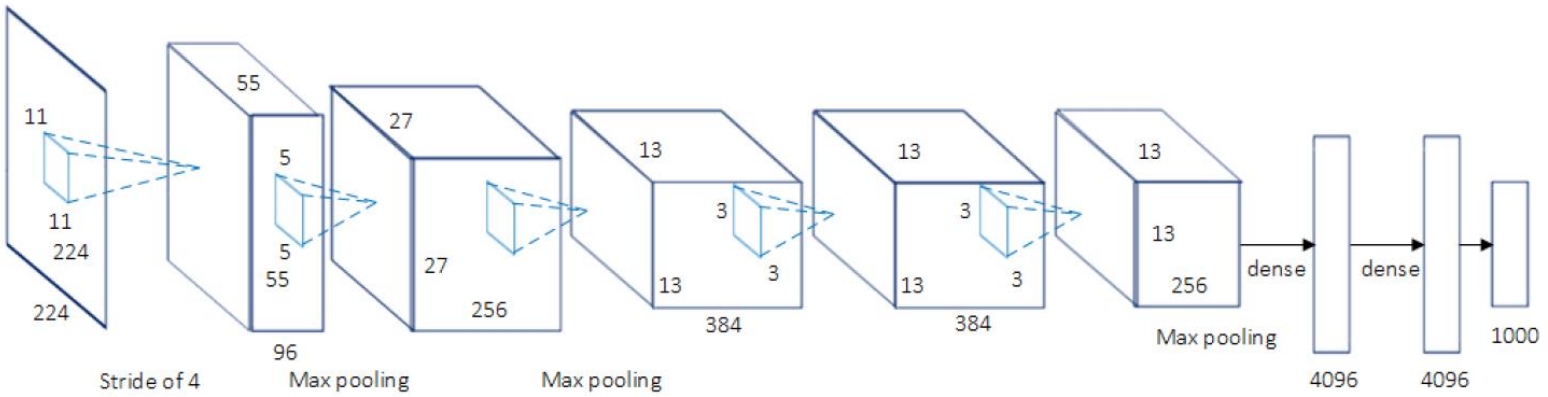

Reactive or deliberative?

choices?
state?
implicit/explicit

The Ingredients of Intelligence

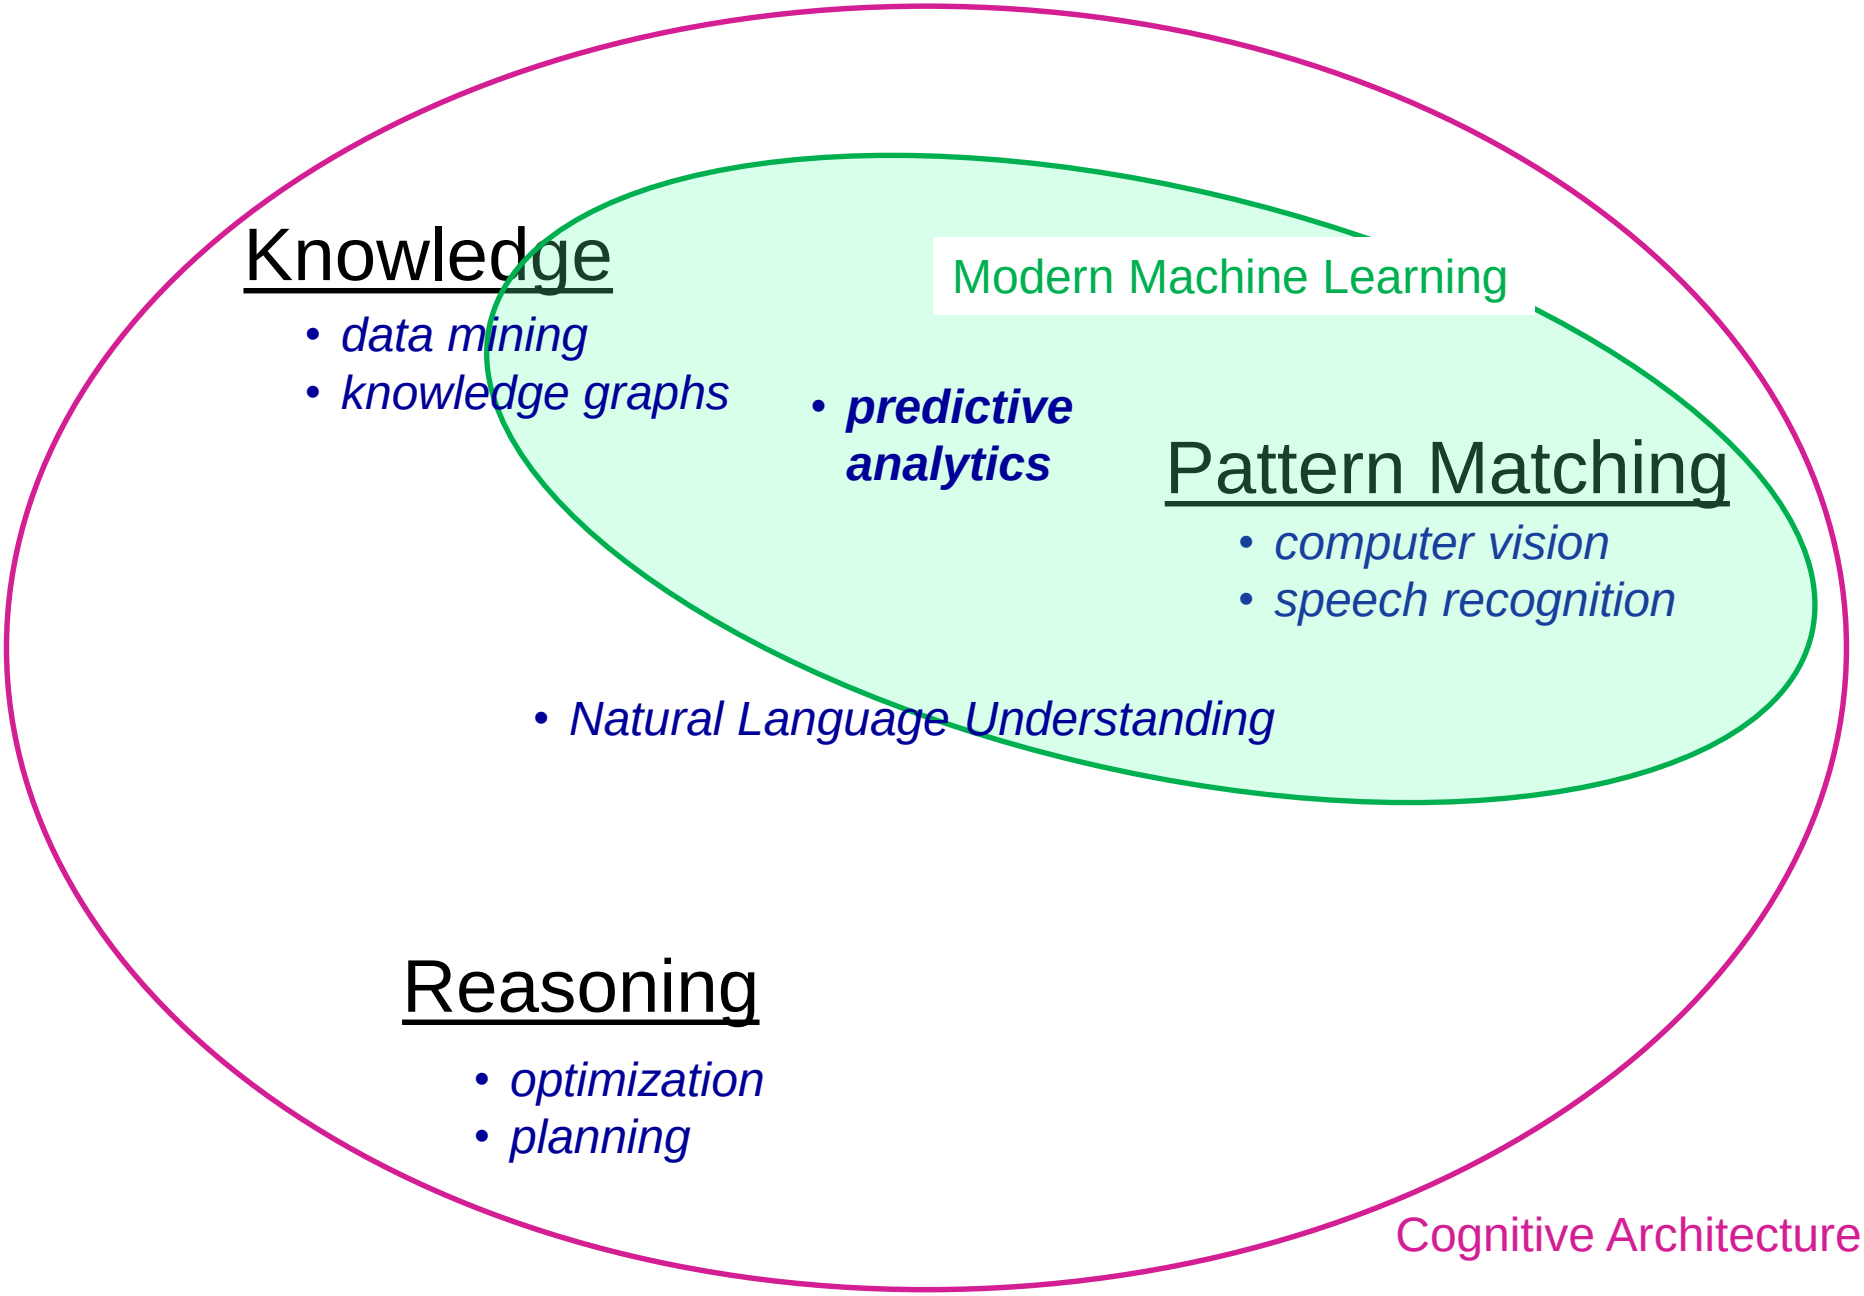

Cognitive Architecture: Outline

- Architecture in Information Systems

- History
  - Computational theory of mind

- The Standard Model Cognitive Architecture
  - Example: Soar

- Important Concepts
  - Marr's Three Levels
  - Reactive vs Deliberative

- Architecture in NN / Deep Learning Networks

- Flight of Imagination
  - LIDA, CopyCat

- Architecture of a Conversational Agent

Cognitive Architecture: Fundamental Parameters

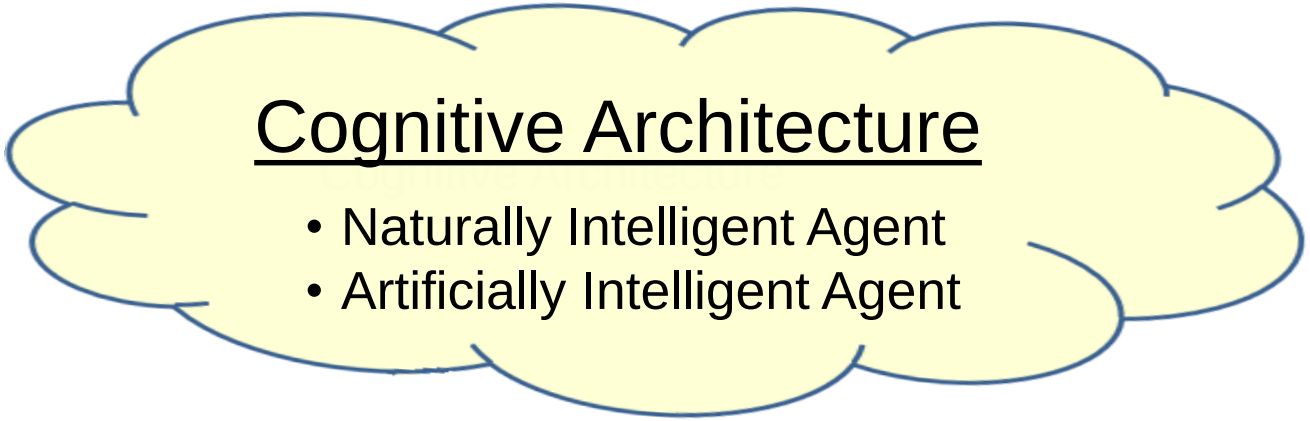

Where does knowledge reside?

How is state represented and utilized in decisions?

Forms of Memory

Localist vs. Distributed Representations

→ What is the control mechanism: what to do & think next?

Distinction between Program Control and Data

How Learning Happens

LIDA Cognitive Architecture

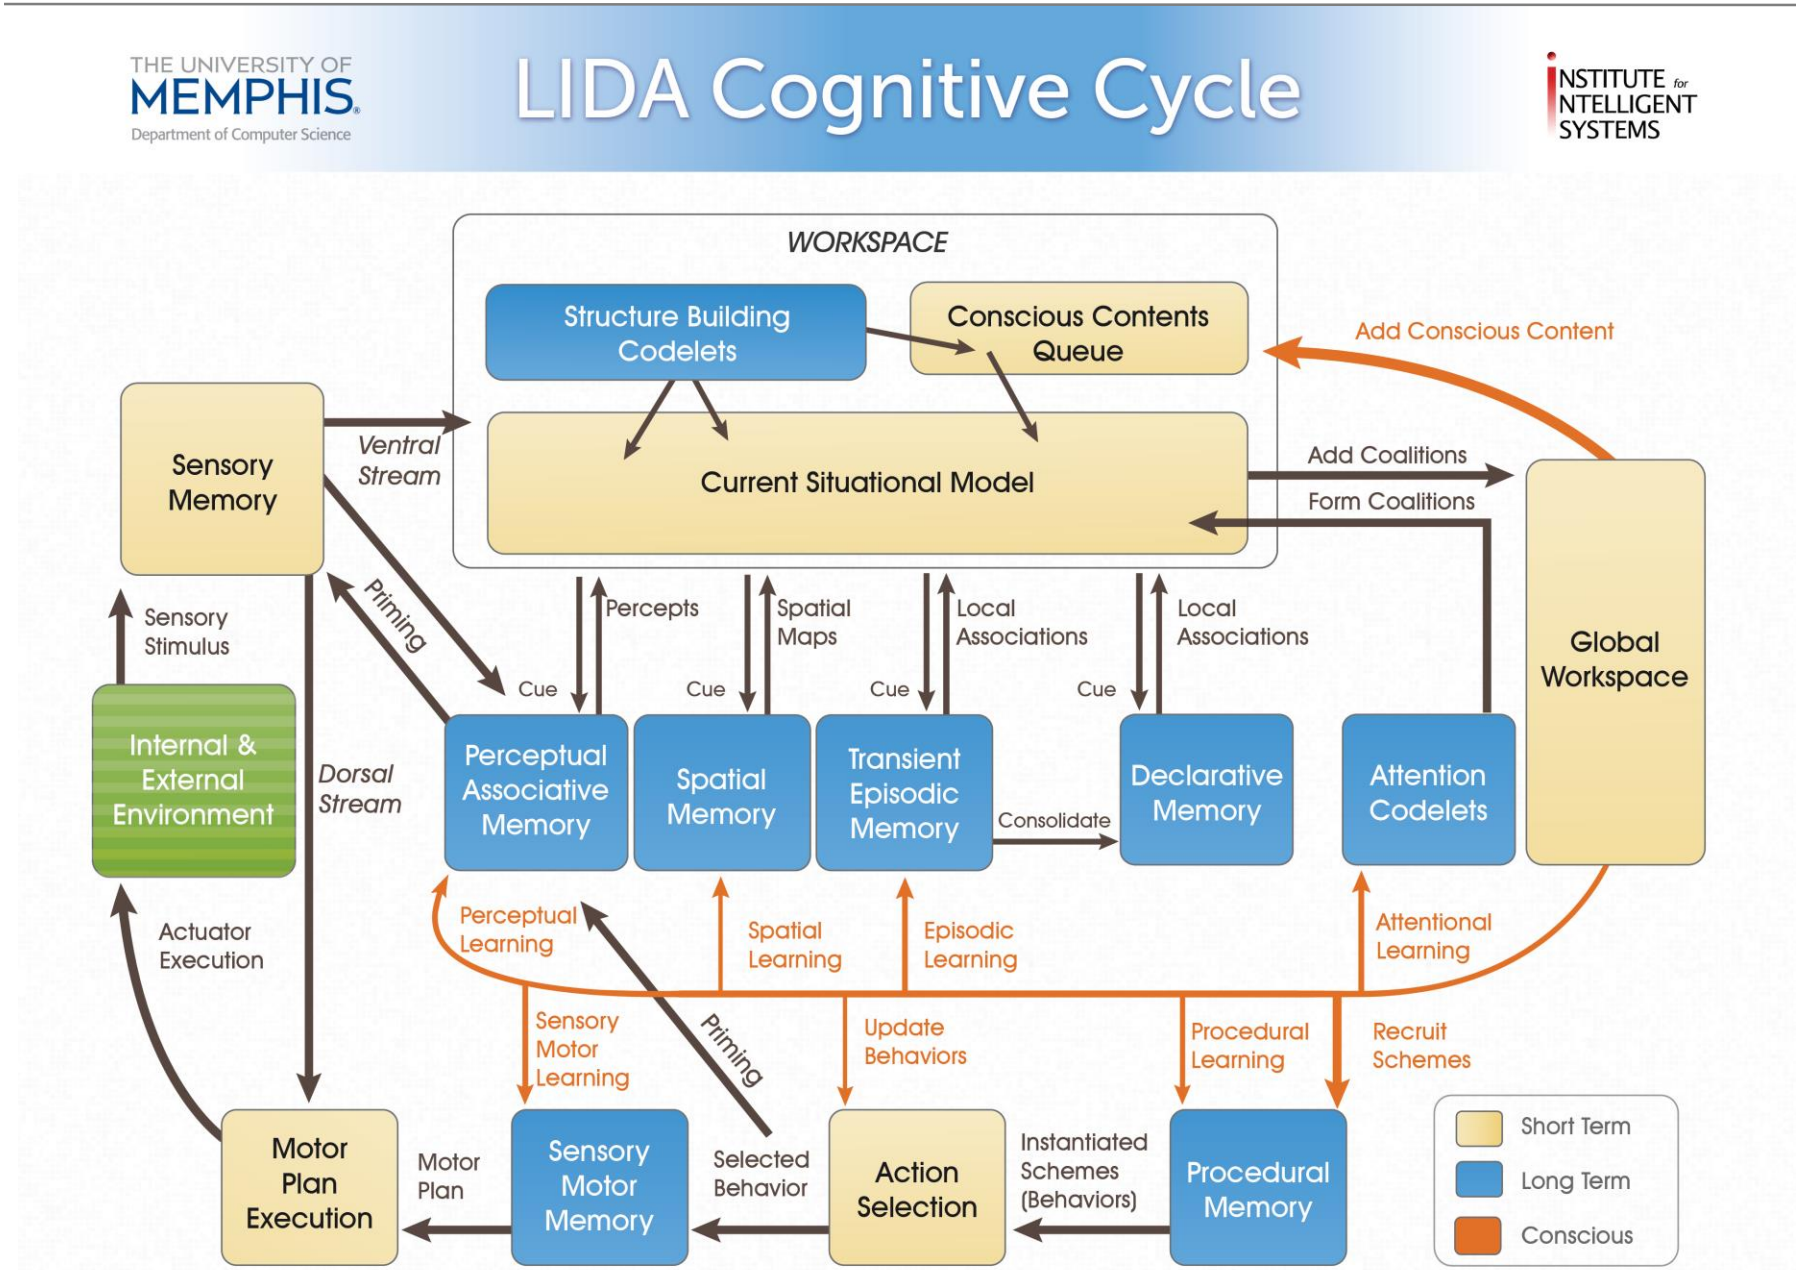

Stan Franklin

Copycat/Metacat

Task: Letter string analogy

$$abc \Rightarrow abd$$

$$mrrjjj \Rightarrow ?$$

Douglas Hofstadter, Melanie Mitchell, James Marshall

Copycat/Metacat

Task: Letter string analogy

$$abc \Rightarrow abd$$

$$xyz \Rightarrow ?$$

Douglas Hofstadter, Melanie Mitchell, James Marshall

Copycat Architecture

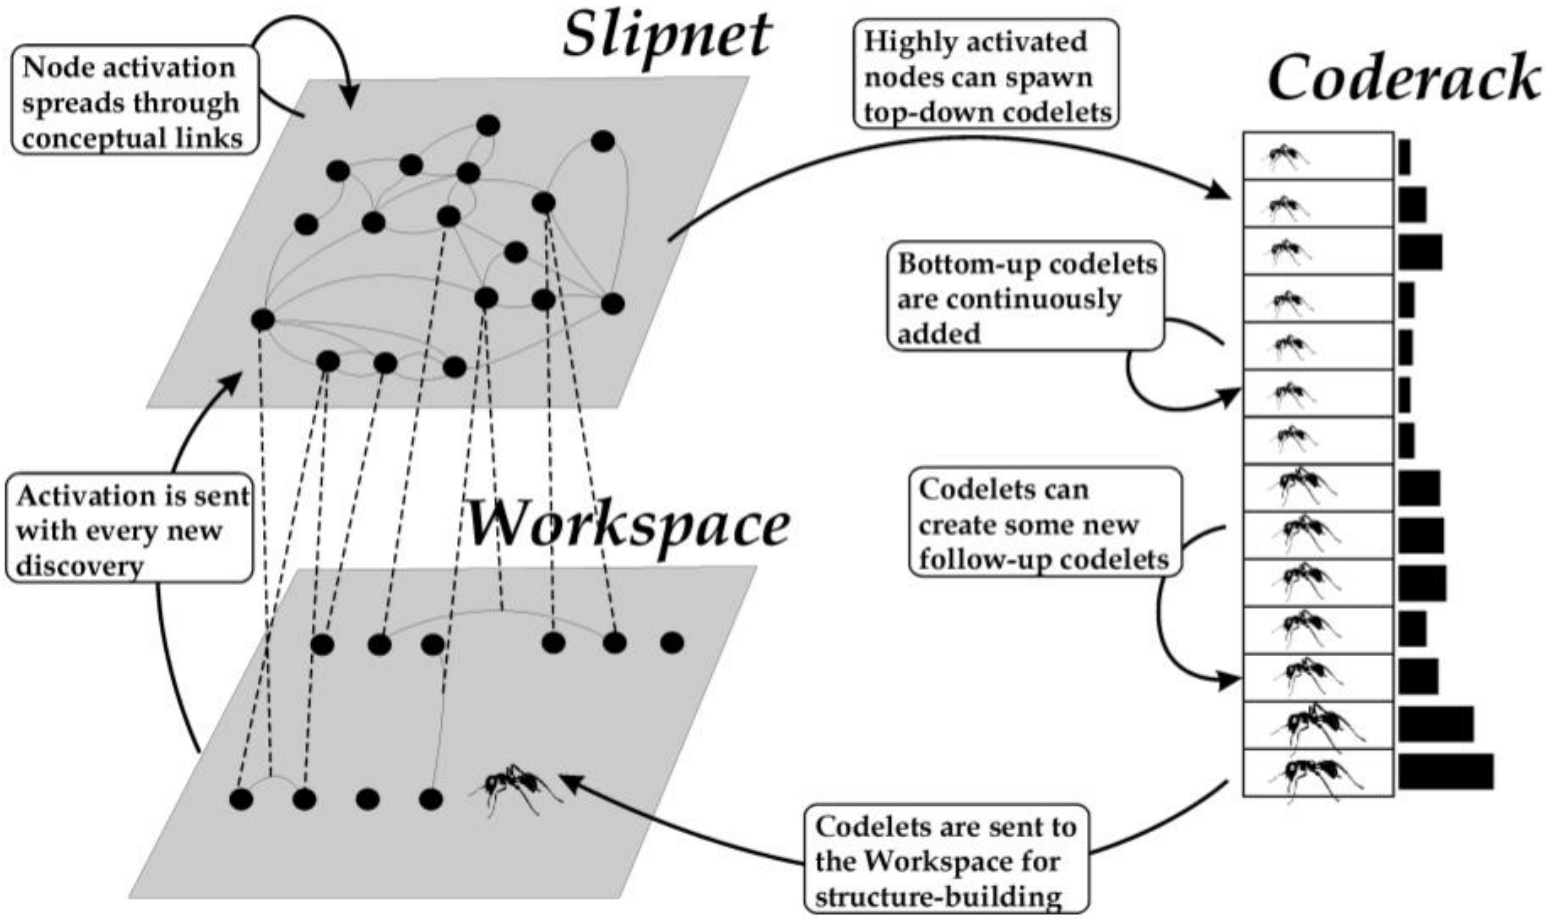

Figure 7: A feedback loop between perceptual and conceptual activity.

Copycat/Metacat Slipnet

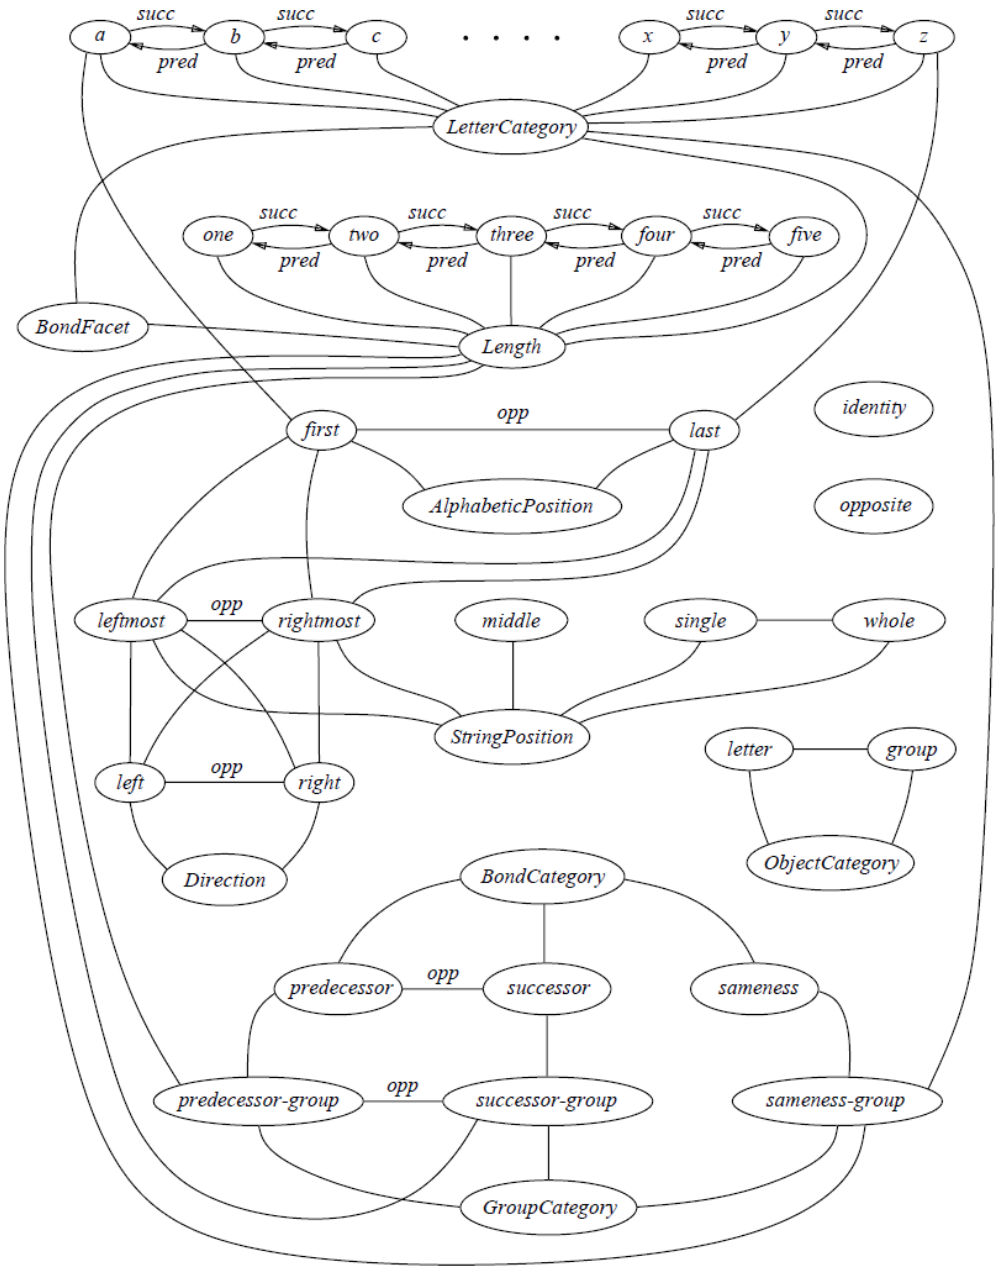

Running Metacat

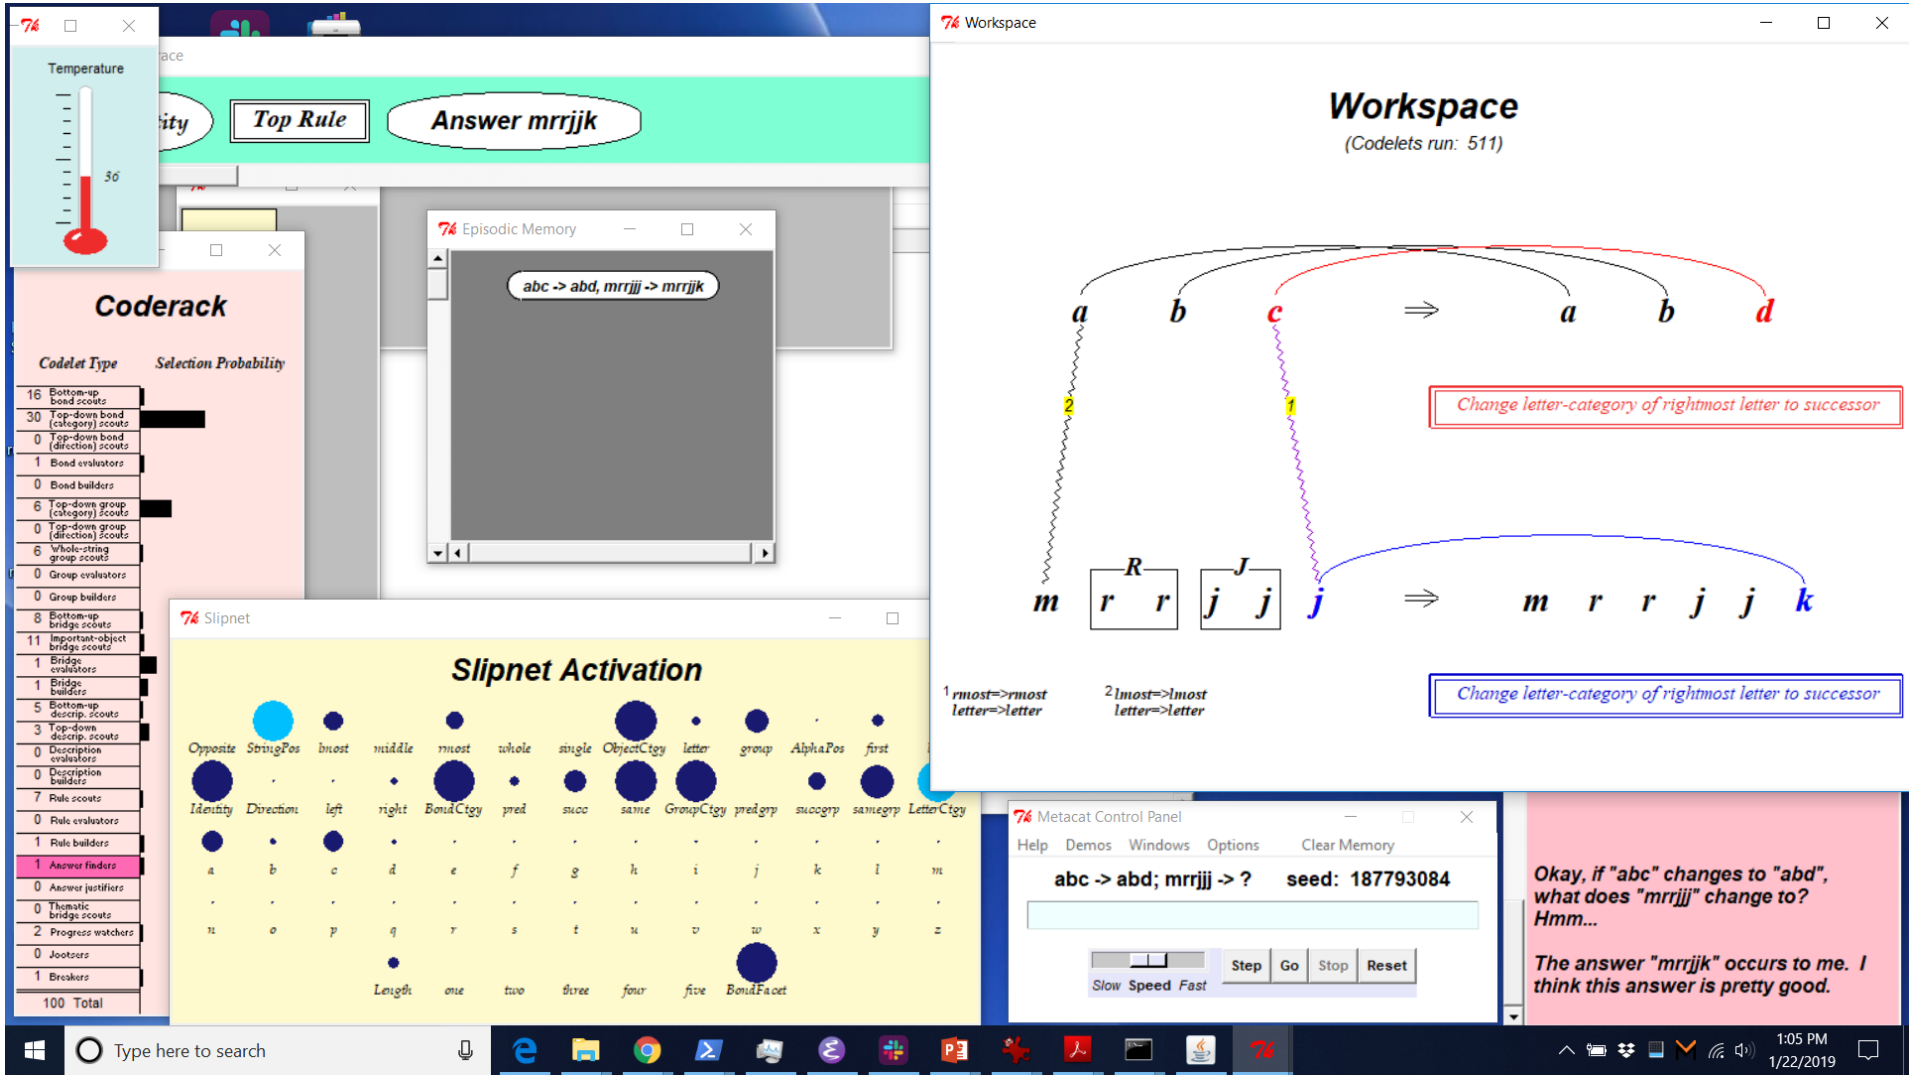

Cognitive Architecture: Outline

- Architecture in Information Systems

- History
  - Computational theory of mind

- The Standard Model Cognitive Architecture
  - Example: Soar

- Important Concepts
  - Marr's Three Levels
  - Reactive vs Deliberative

- Architecture in NN / Deep Learning Networks

- Flight of Imagination
  - LIDA, CopyCat

- Architecture of a Conversational Agent

Conversational Agent

Question answering task

"Alexa, who won the 1934 world series?

"The Saint Louis Cardinals beat the Detroit Tigers 4-3 in the 1934 World Series."

Architecture of a Conversational Agent

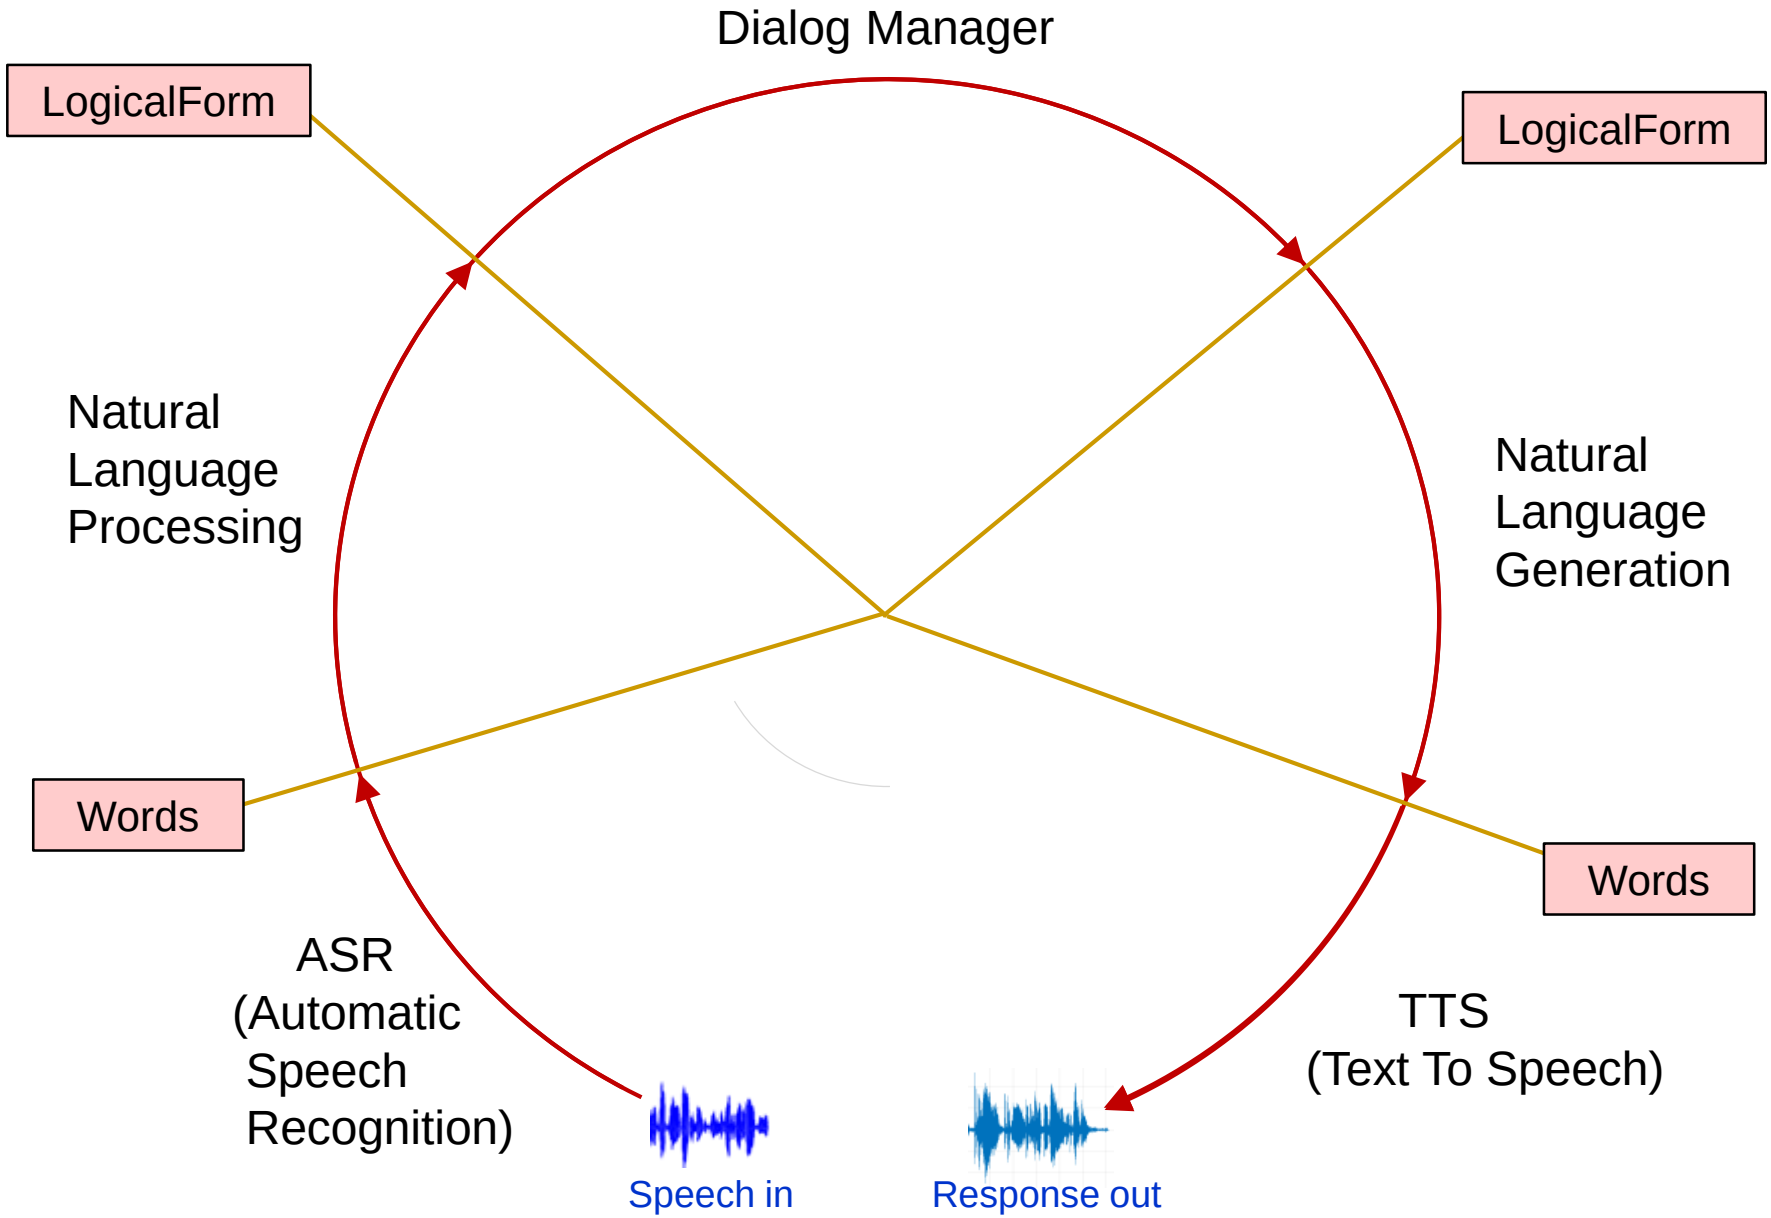

Architecture of a Conversational Agent

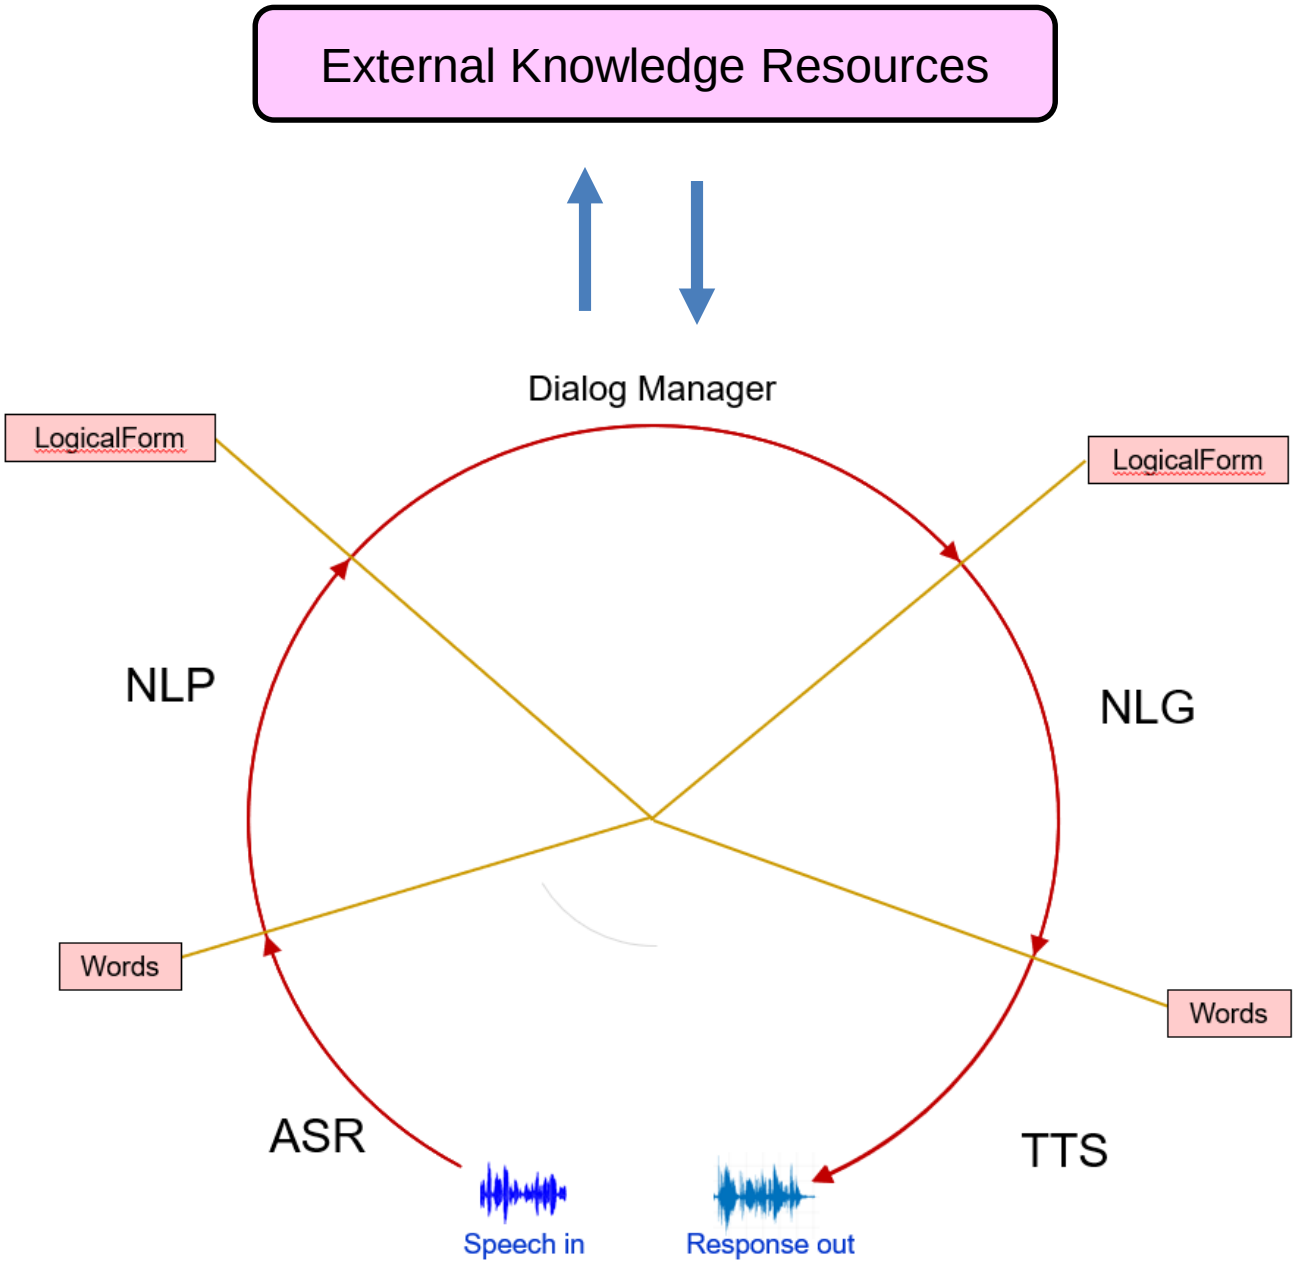

Cognitive Architecture: Standard Model

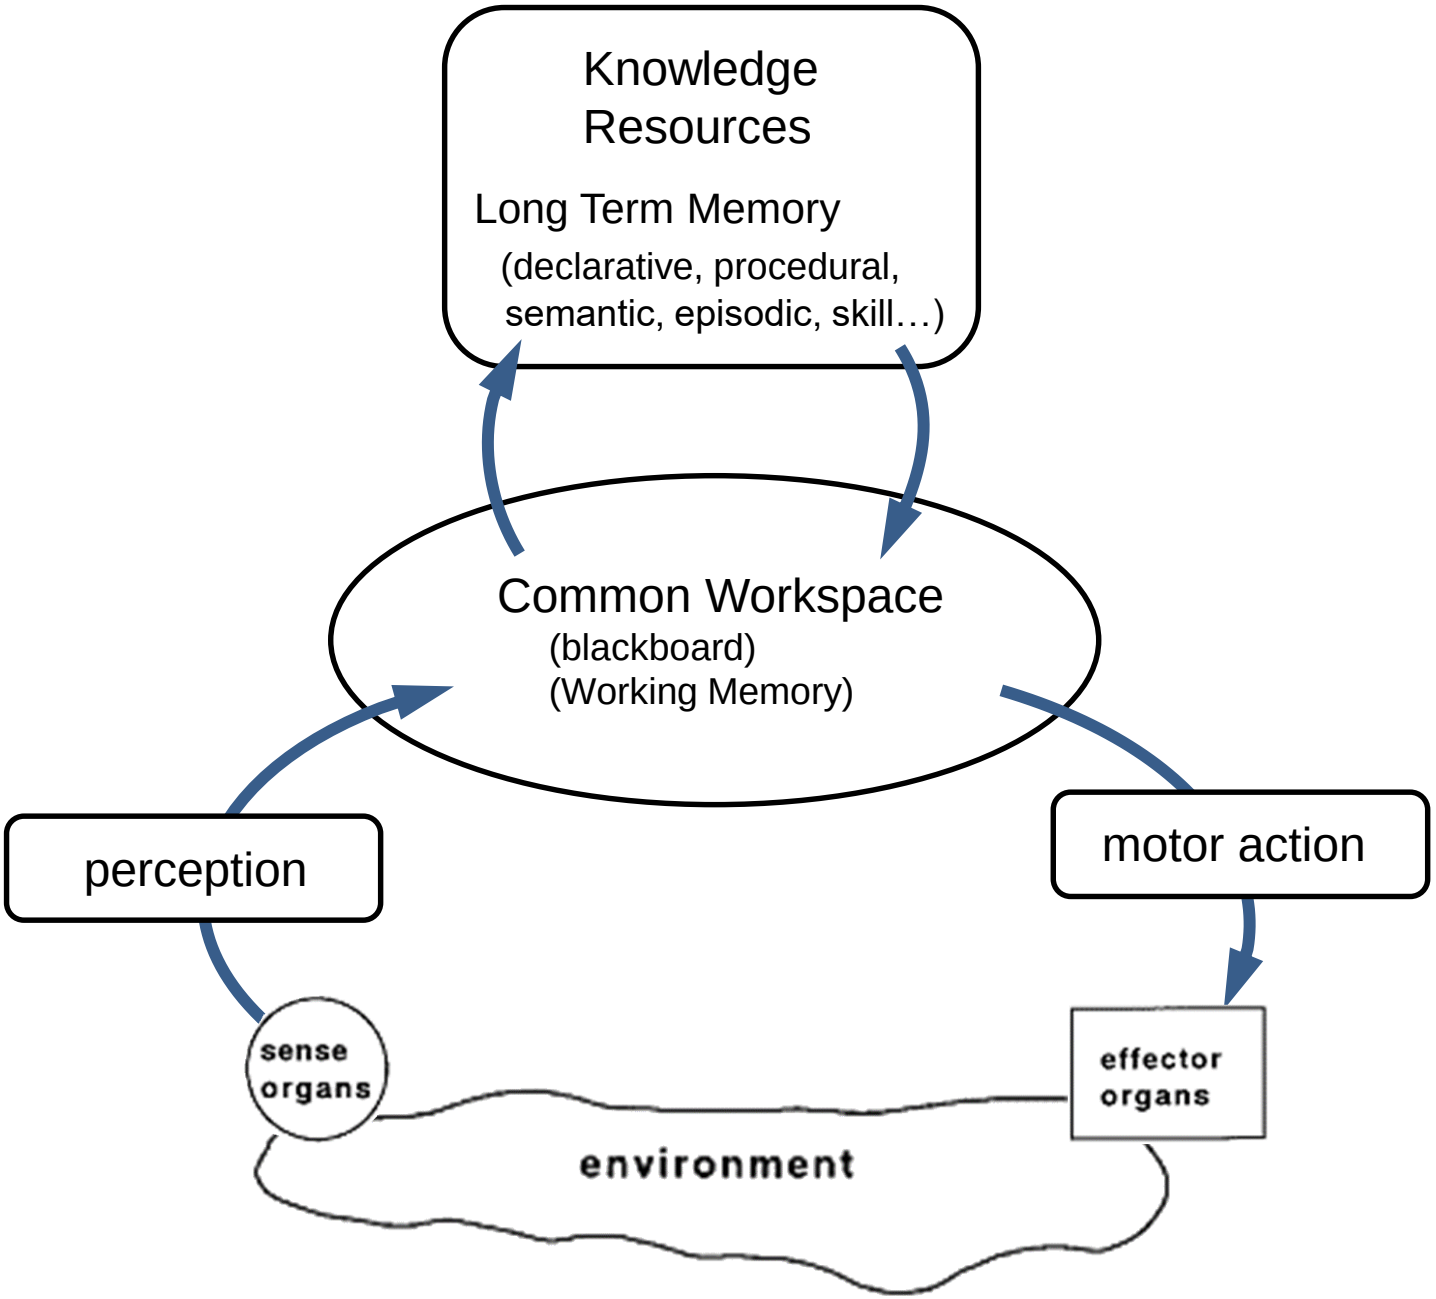

Why Are Conversational Agents So Dumb?

Question answering task

"Alexa, who won the 1934 world series?

"The Saint Louis Cardinals beat the Detroit Tigers 4-3 in the 1934 World Series."

"Alexa, who was the president then?"

"This might answer your question. The president of the United States is Donald Trump."

Cognitive Architecture for a Conversational Agent

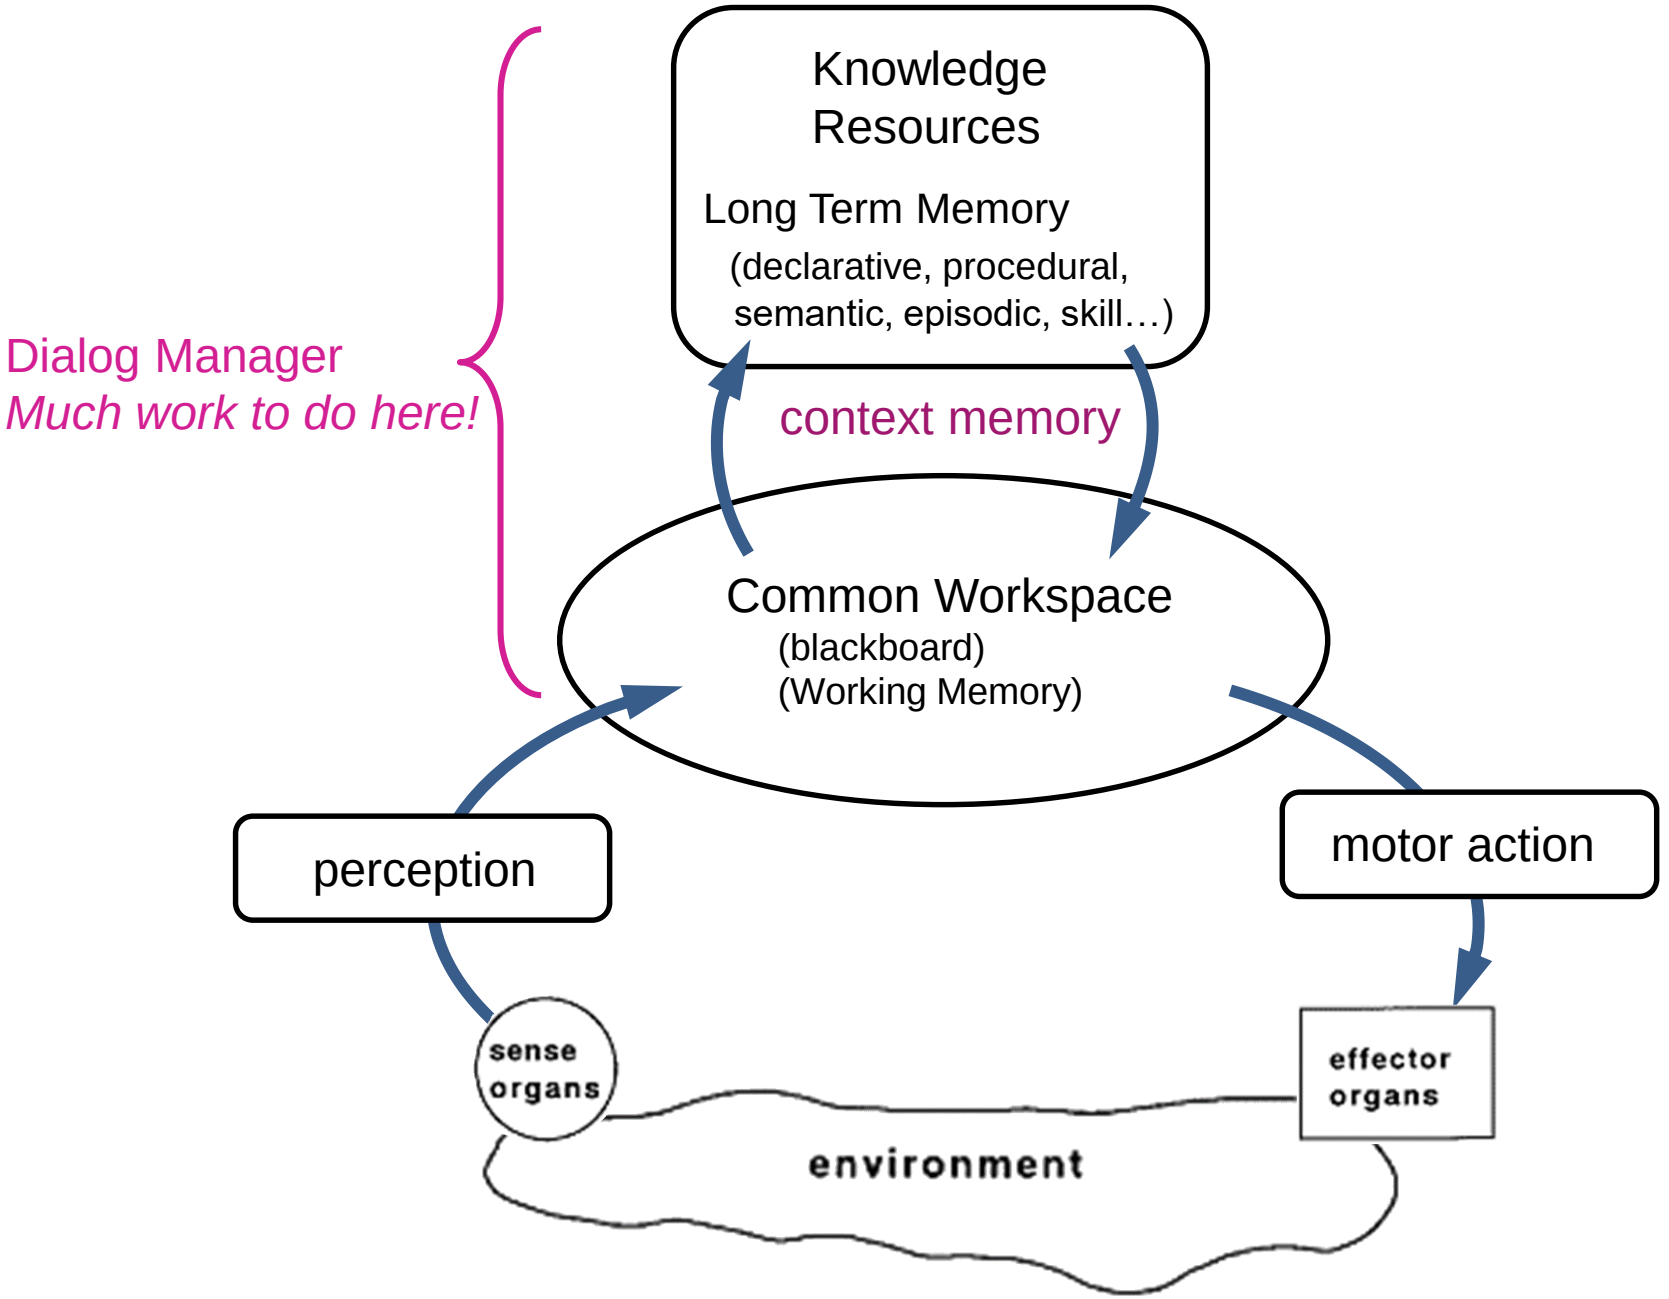

The Ingredients of Intelligence

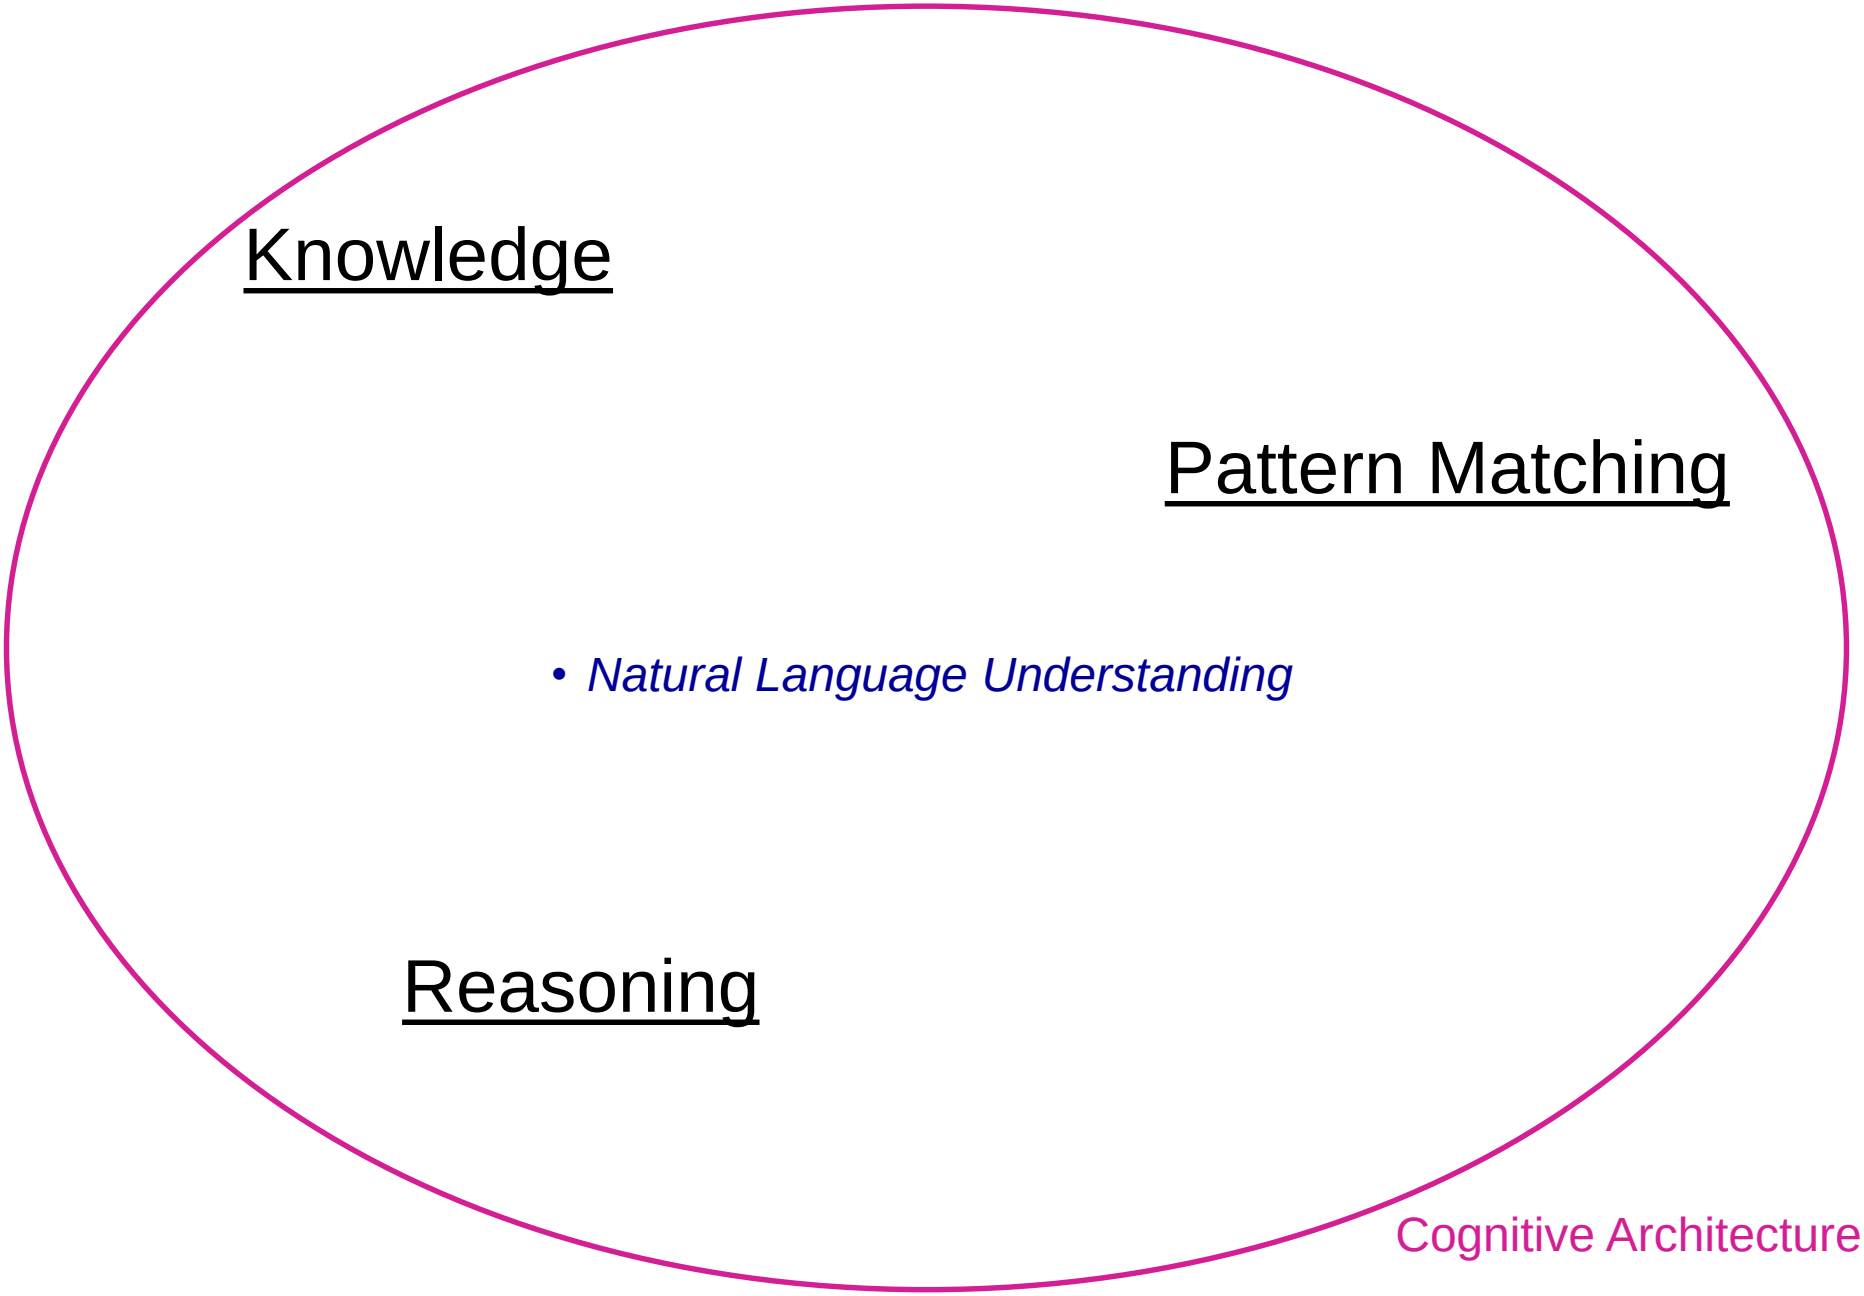

Summary: Taxonomy

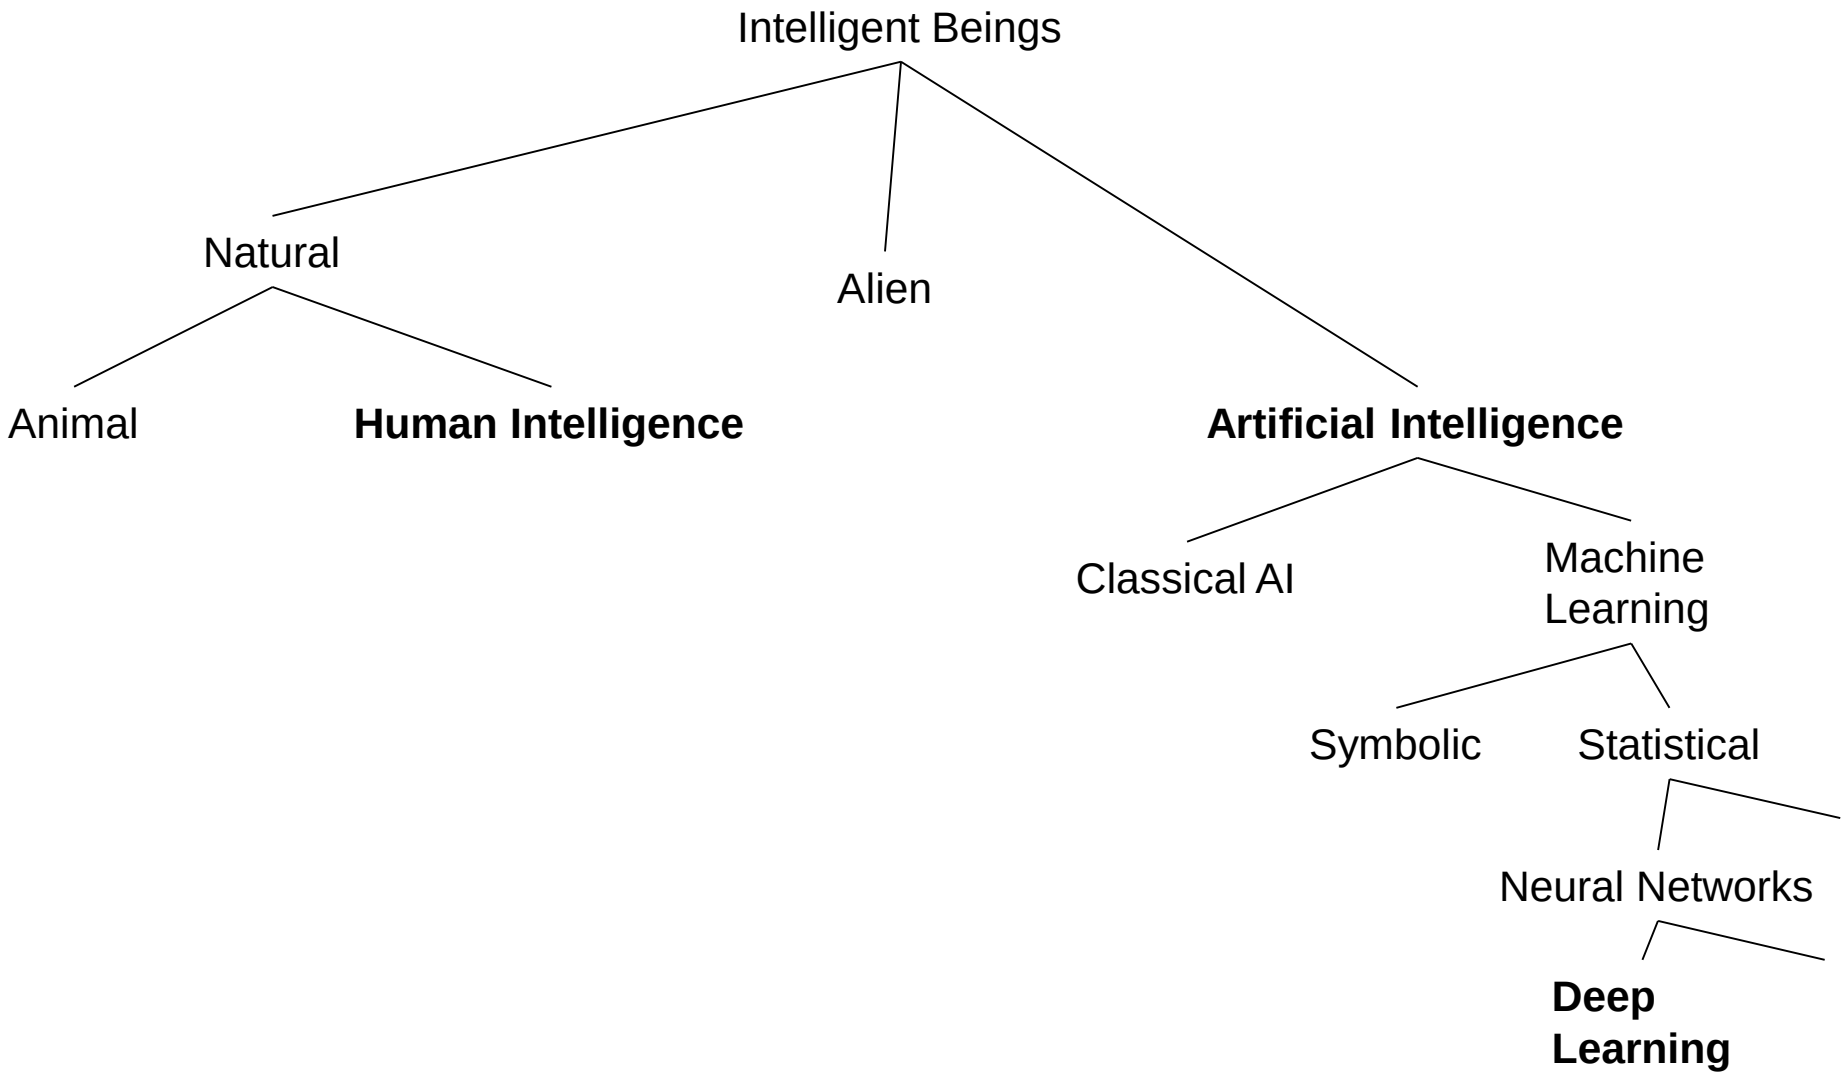

Cognitive Architecture

NN Architecture

Eric Saund

- *Research scientist in Cognitive Science and AI.*
- *Conversational Agents, Visual Perception, Cognitive Architectures.*
- *I build stuff.*

Projects

Papers

Curiosities

Links

Contact

http://www.saund.org

saund@alum.mit.edu

Conversation