Extending Soar with Dissociated Symbolic Memories

Nate Derbinsky and John E. Laird
University of Michigan

Goal: Human-Level AI

- Autonomous
- Long-Living
  - Months, Years
- Multiple, complex tasks

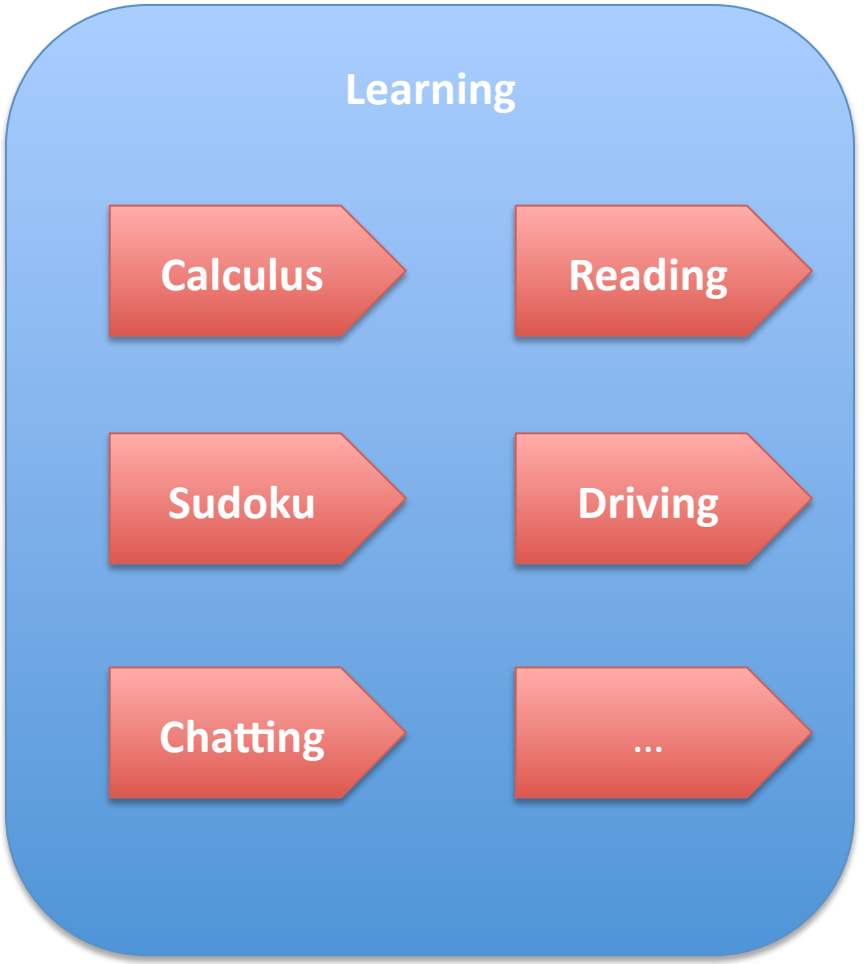

Challenge: Lots of Knowledge

Over long lifetimes, learning agents accumulate large stores of knowledge

- Lexical (WordNet)
- Common sense (Cyc)
- Personal history
- etc.

Given limited computation, they must remain reactive to a complex, dynamic environment

Memory Systems

Data structures and algorithms to *encode*, *store*, and *retrieve* agent knowledge

- Fixed computational profile

Design decision: which system(s)?

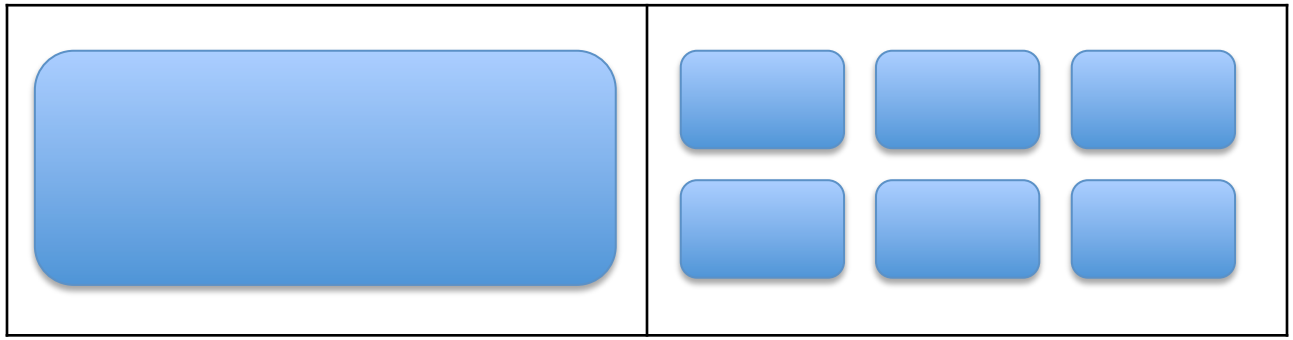

An Argument for Dissociation

A memory system will solve some problems

Problems may require memory systems that are *functionally incompatible*

Evidence

- Evolutionary (Sherry and Shacter, 1987)
- Computational (O'Reilly, 2006)

Our Work

1. Analysis of symbolic memory systems in Soar
   - Functionality
   - Storage and retrieval breakdown
   - Computational arguments for dissociating long-term memory systems

2. Integration issue: persistent object identity
   - See paper for detailed evaluation

Soar Memory Systems Overview

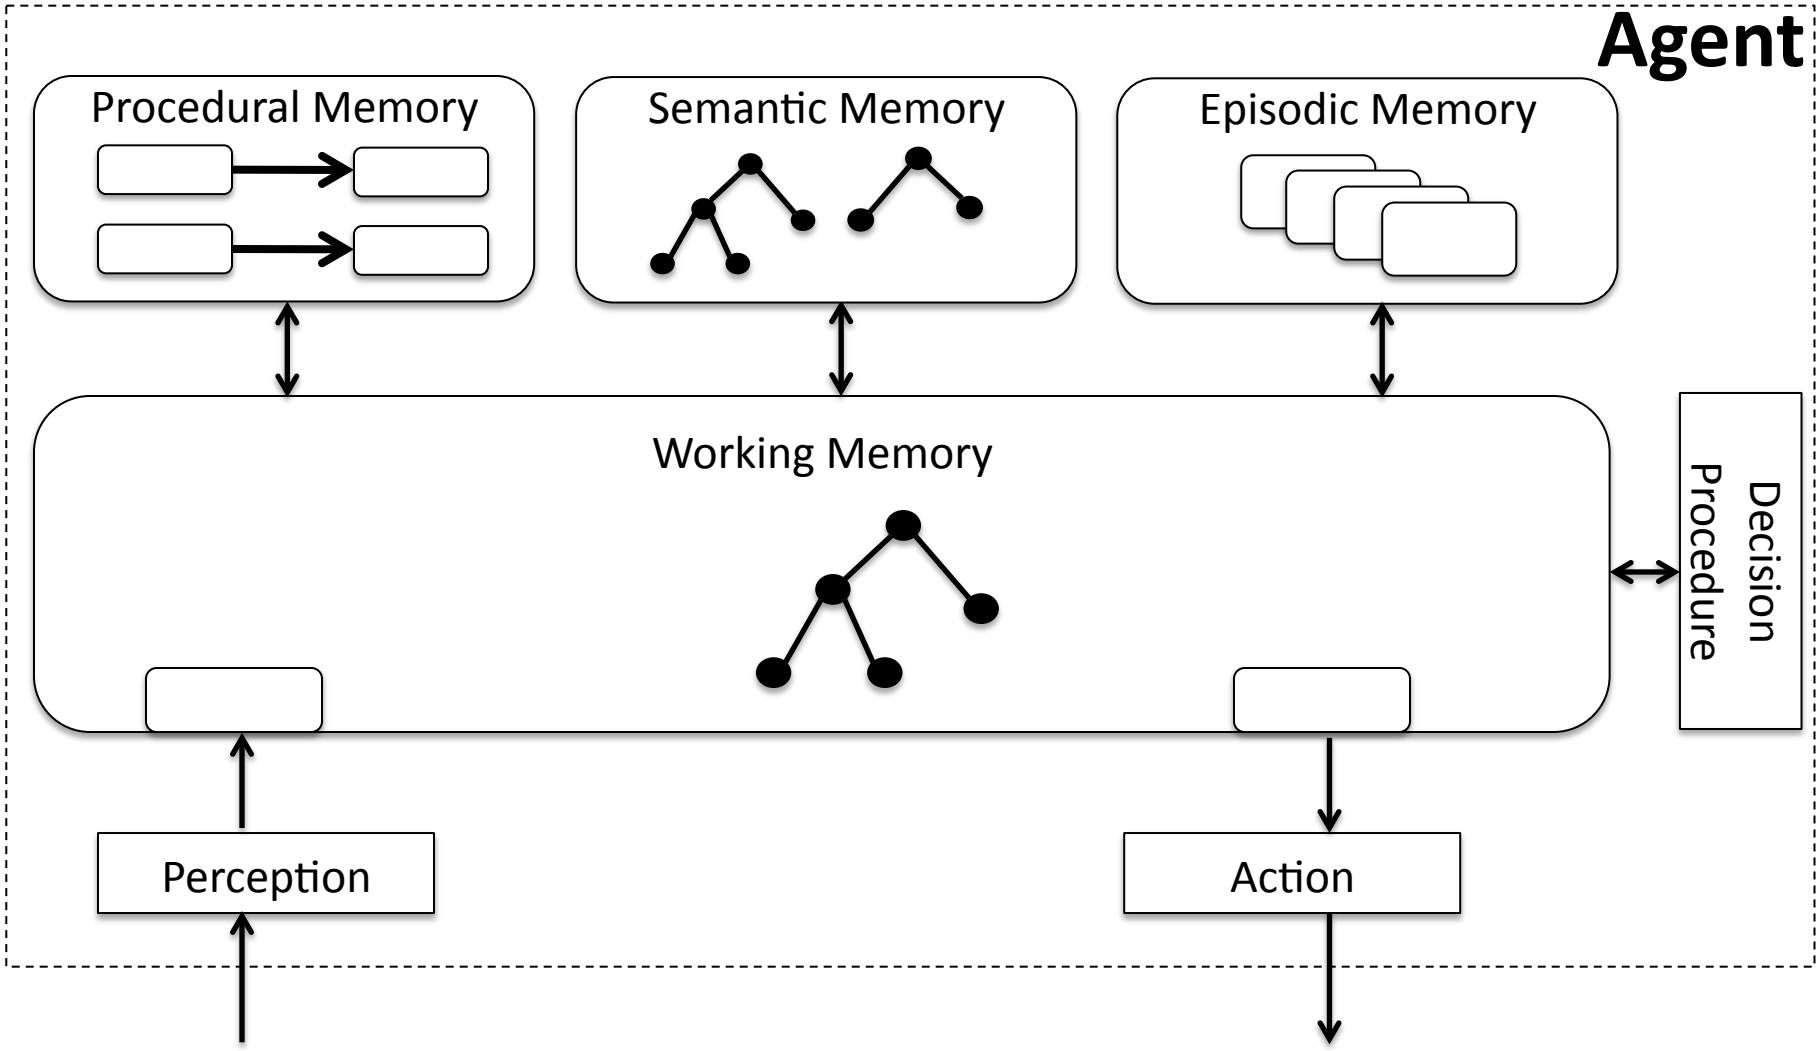

Working Memory

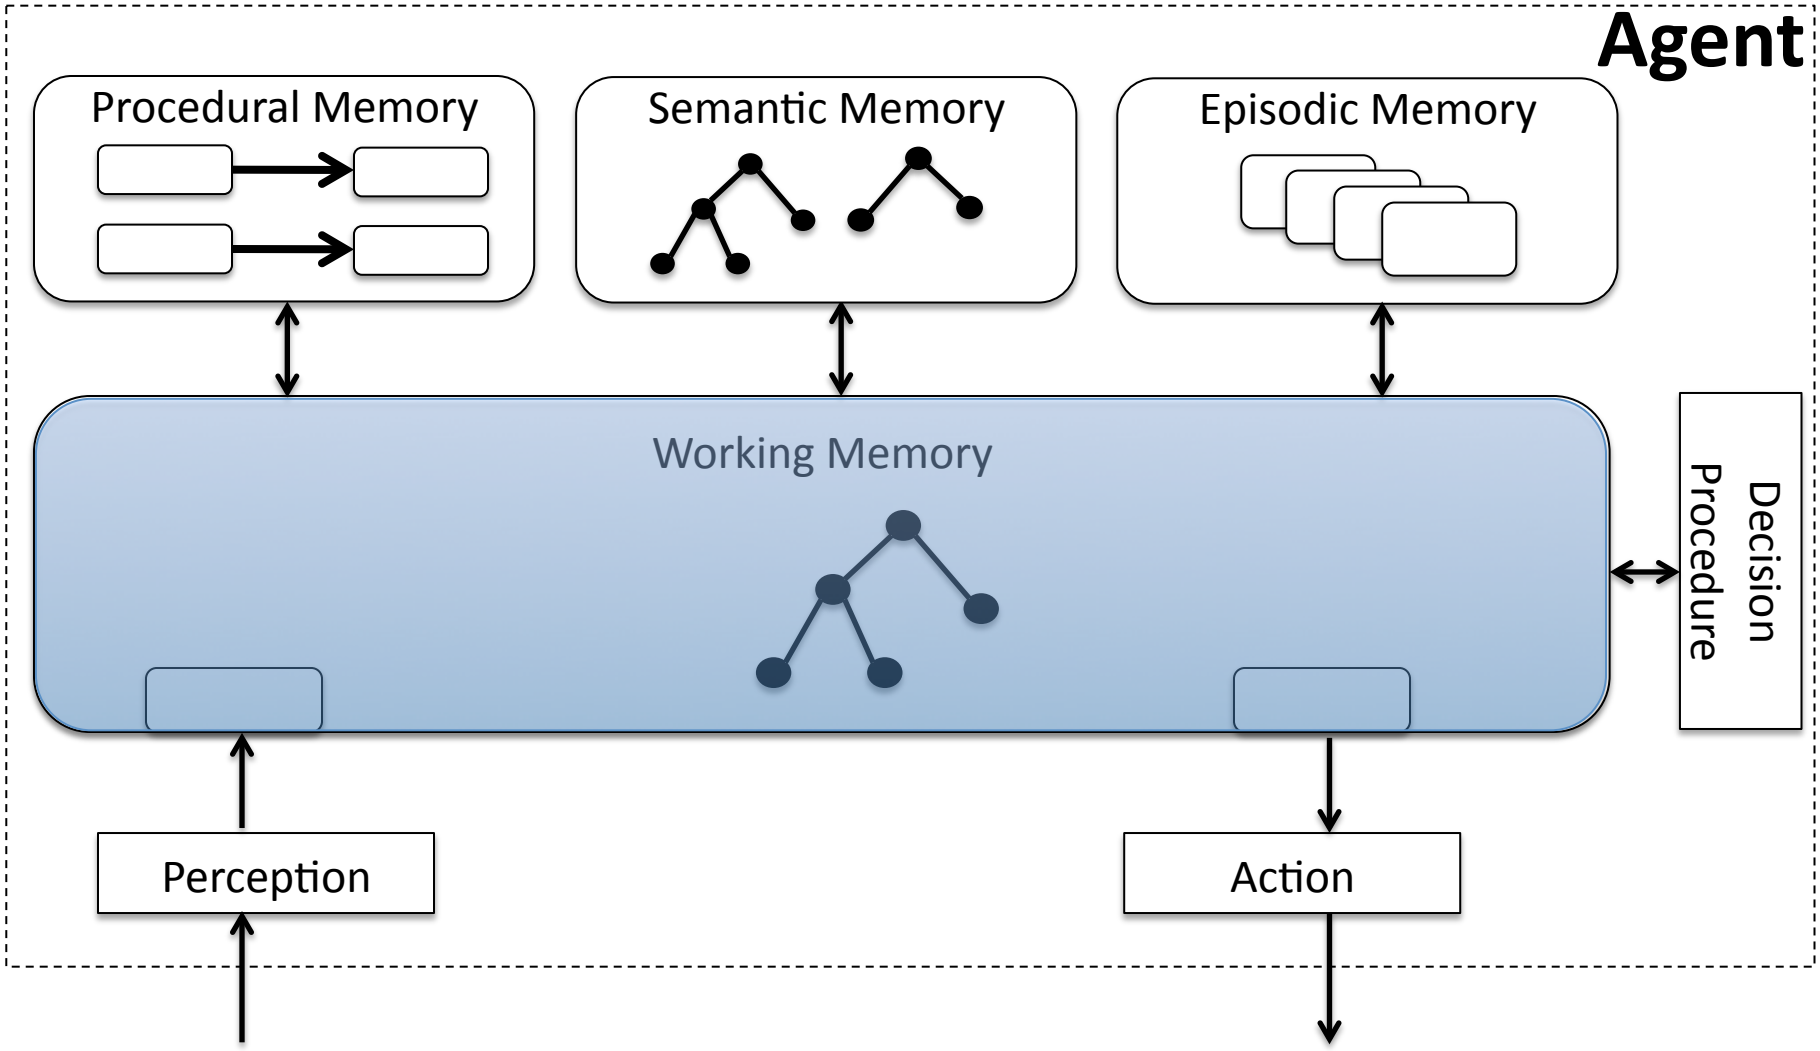

Working Memory Function

Capture current agent state

- No active processing

Represent arbitrary and novel combinations and compositions of symbols

- Separate store (vs. activated long-term)

Provide common representation substrate

- Procedural reasoning
- Initiate external action
- Cue long-term memories

Working Memory Representation

Directed, connected graph of *mental entities* and associated augmentations

- Object in the world (past, present, hypothetical) or idea
- Originate from external environment (perception, feedback) or memory systems

Working Memory Example

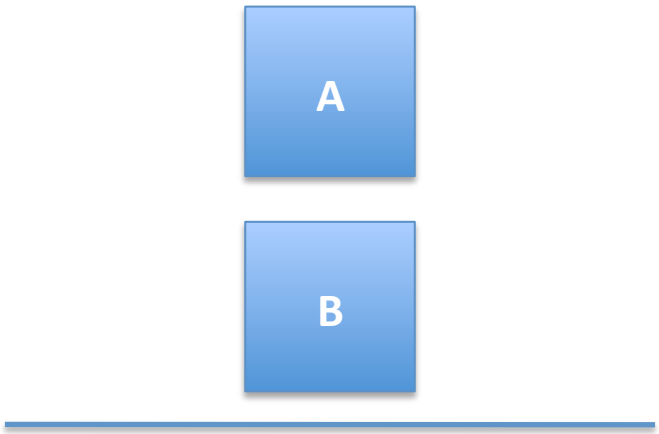

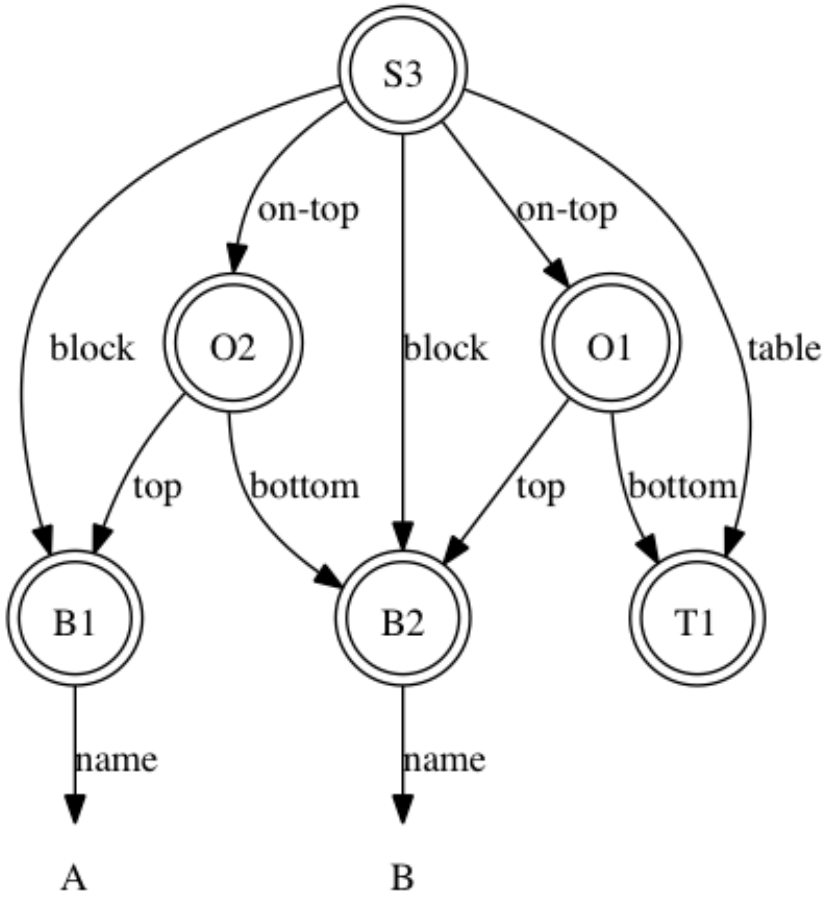

Procedural Memory

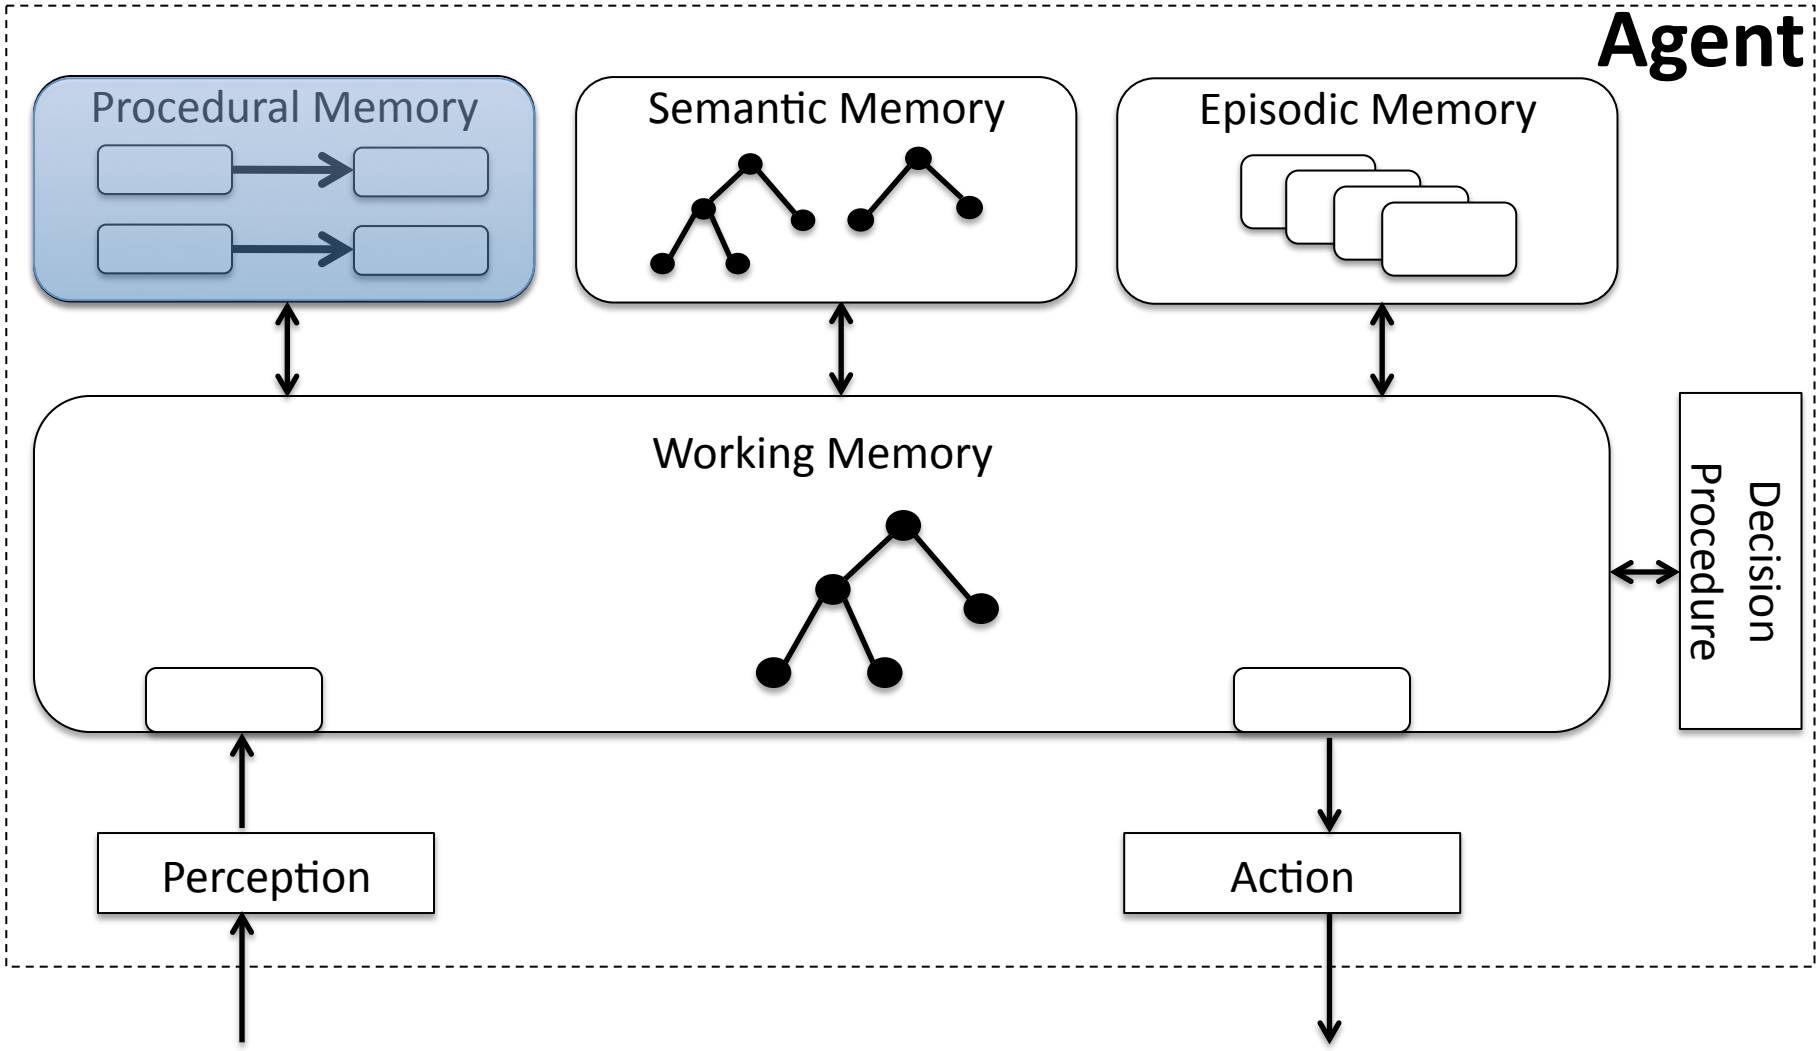

Procedural Memory Function

Knowledge of when and how to perform internal and external actions

- Initiate actions in the external environment
- Cue internal memory systems
- Deliberate reasoning

Procedural Memories

Knowledge is encoded as production rules

- *Antecedent (LHS)* - conjunctive set of variablized working memory structures
- *Consequent (RHS)* – conjunctive set of variablized working memory modifications

Production Example

If

- There is a block <b1> named A *and*
- There is a block <b2> named B *and*
- Block <b1> is on <b2>

Then

- Move <b1> to the table

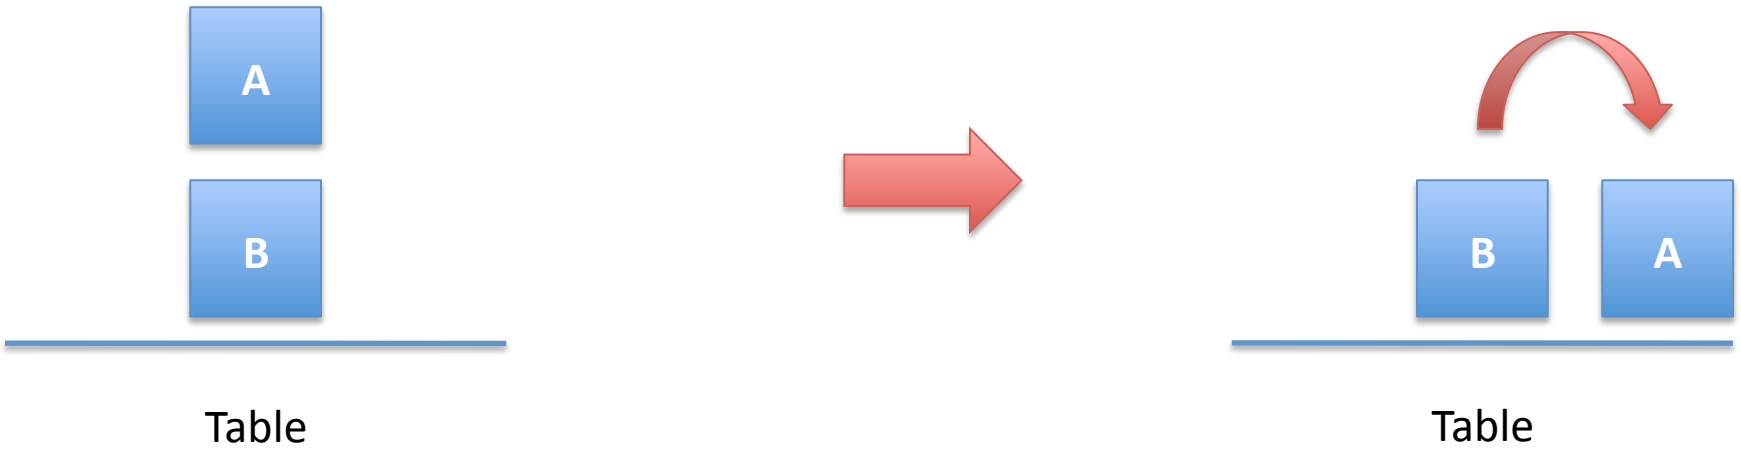

Procedural Retrievals

Given

- Productions
- Working Memory

Find

- ALL production instantiations that match working memory

Requires specialized data structures and precise algorithms for reactivity over time given many complex rules and/or large working memory

Procedural Storage

Chunking

Compiles sub-task processing to convert deliberation to reaction

Reinforcement Learning

Incrementally tunes production actions to reflect an expectation of action performance

Status & Agenda

At this point, Soar functionally…

- represents state
- has long-term knowledge of when and how to perform internal and external action

Is there missing functionality?

If so, can an existing memory system efficiently provide this functionality?

Semantic Memory

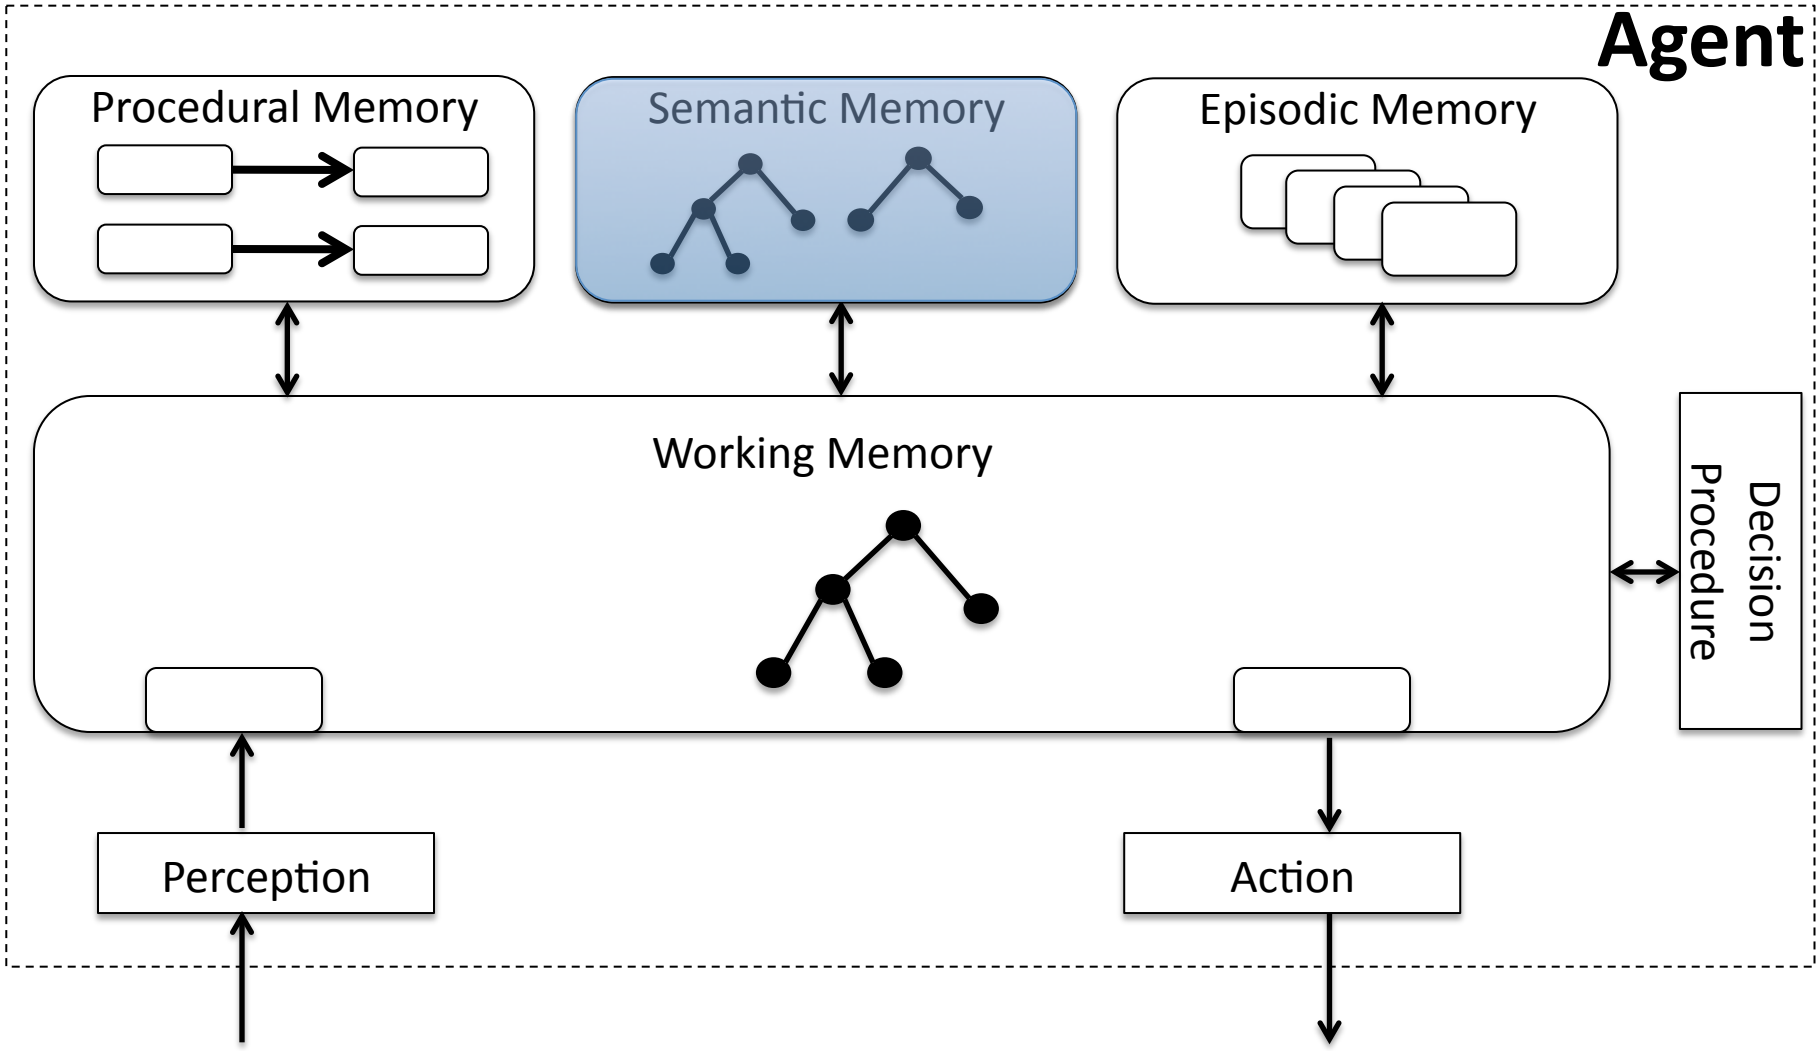

Semantic Motivation

- Who is the current President of the United States?
- When/where/how did you first learn this fact?
- Is context always necessary/advantageous?

Semantic Memory Function

Efficiently retrieve facts about a mental entity, independent of original context

- Assumes some experience is useful in differing situations (vocabulary, arithmetic, etc.)

Semantic Memories

Mental entities and augmentations

- Same representation as working memory

Similar to ACT-R declarative chunks

- No pre-specified chunk-types

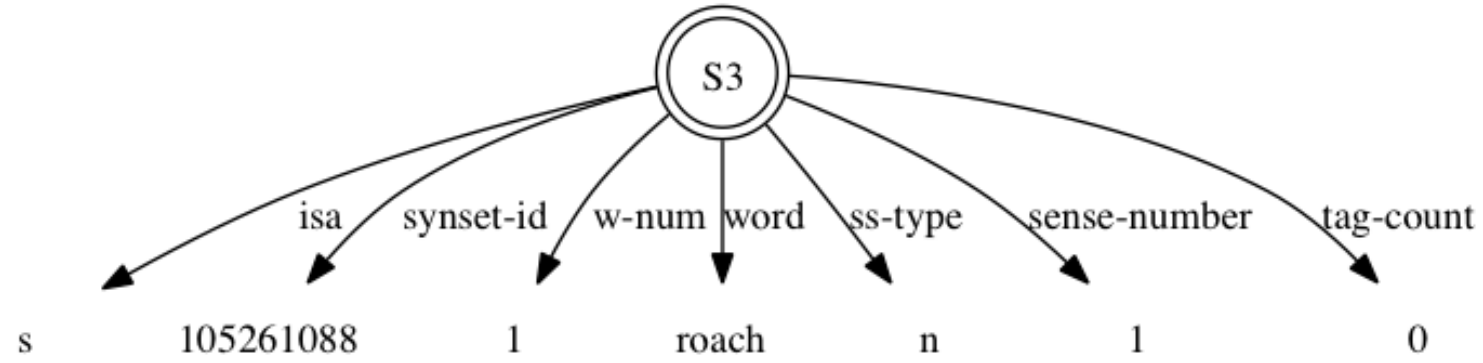

Semantic Retrievals

Given

- Conjunctive set of attribute/value pairs
- Semantic memories

Find

- Single mental entity that contains cue structures

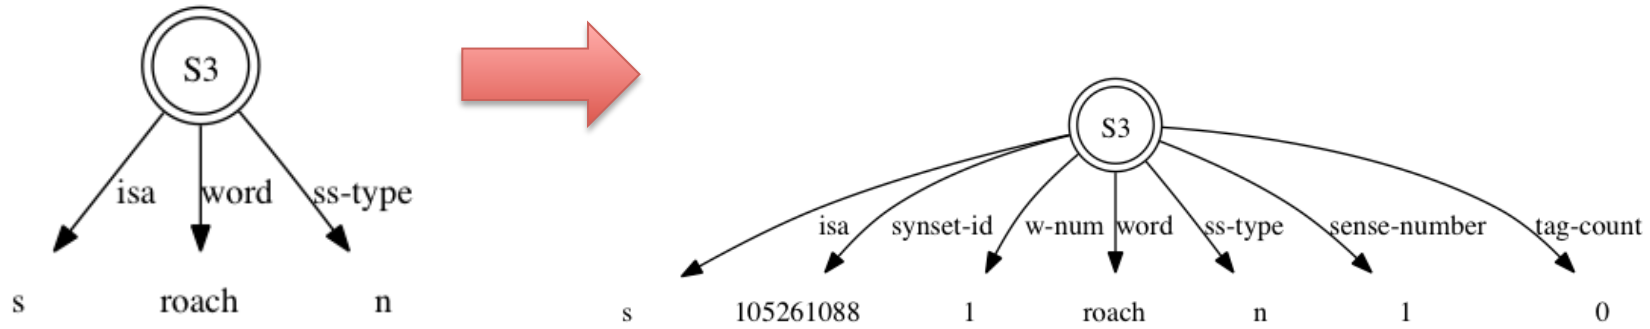

Semantic Dissociation

Can the *Procedural* memory system efficiently support *Semantic* retrievals?

Approach #1: Data Chunking

Approach

Productions condition upon all possible combinations of cue structures

Results in creation, matching, and updating an exponential number of rules

Example

1. (a AND b AND c AND d)
2. (a AND b AND c)
3. (a AND b AND d)
4. (a AND c AND d)
5. (b AND c AND d)

…

Approach #2: Working Memory

On-demand procedural rule matching per-cue

Match time grows with the number of mental entities

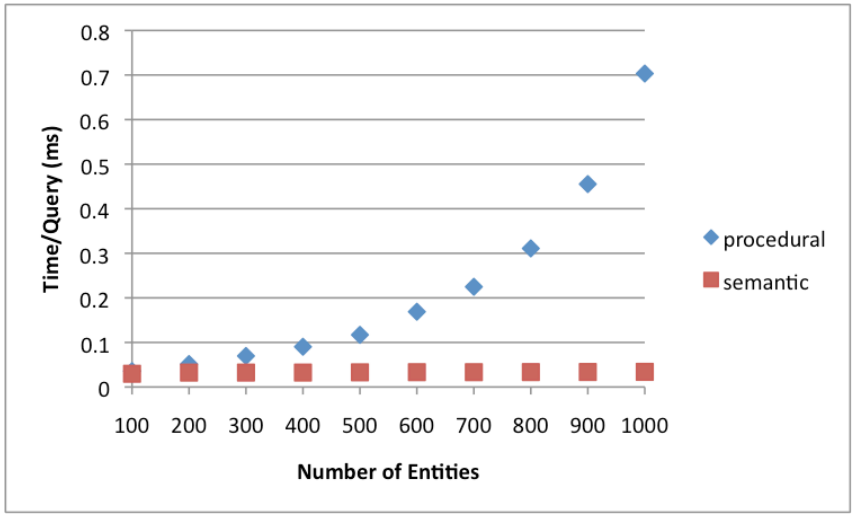

Scaling to Large Semantic Stores

Over long lifetimes, it is conceivable that a learning agent will accumulate large amounts of semantic knowledge

- WordNet > 800K mental entities
- Cyc > 5M facts

Argument for Semantic Dissociation

To maintain reactivity, must examine possibly functionally incompatible mechanism

Procedural	Semantic
Known cues (static)	Unknown cues (dynamic)
All matches	One match
Global cache	Local optimization

Status & Agenda (2)

At this point, Soar functionally…

- represents state
- has long-term knowledge of when and how to perform internal and external action
- stores and retrieves context-less facts

Is there missing functionality?

If so, can an existing memory system efficiently provide this functionality?

Episodic Memory

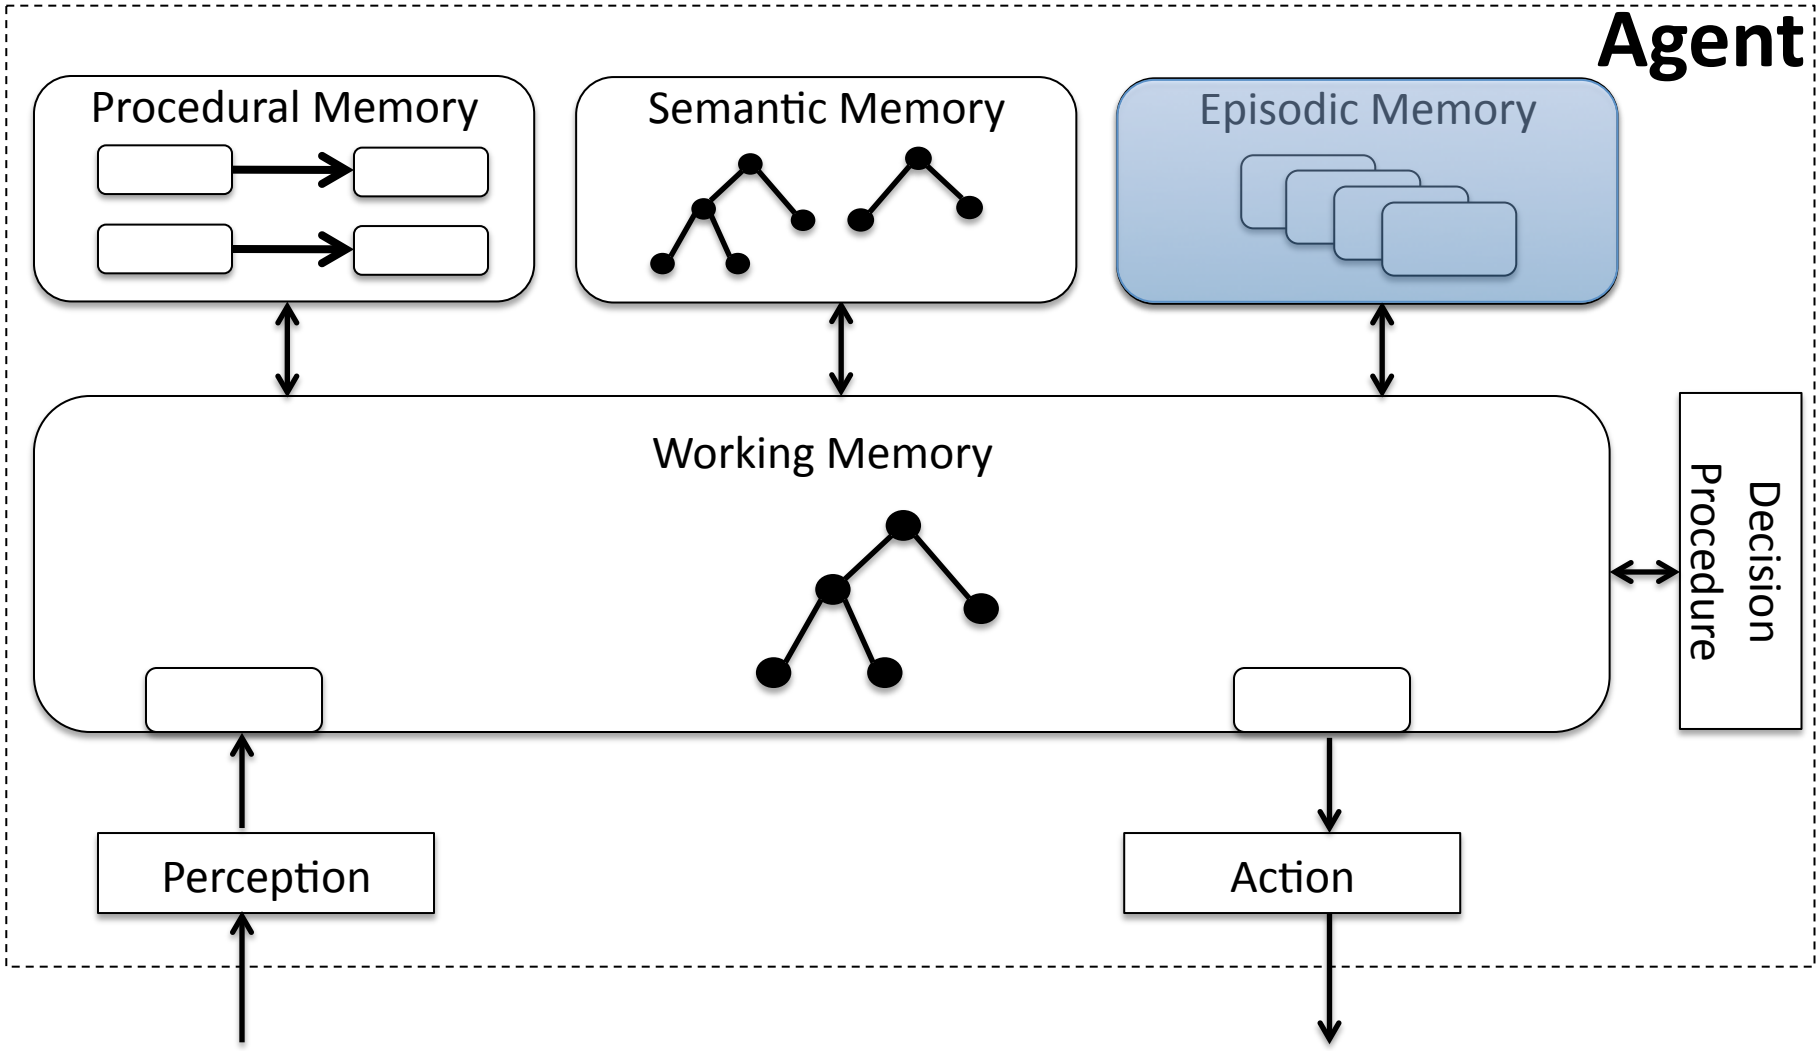

Episodic Motivation

Recall the first time you failed a school assignment…

- What was the subject?
- How did you feel?
- How did your parents react?
  - Lessons learned?

Is personal history advantageous?

Episodic Memory Function

Efficiently capture and retrieve what an agent "remembers"

- Architectural, automatic, autonoetic, autobiographical, temporally indexed

Many possible functional roles

- Virtual sensing, action modeling, retroactive learning, …

Episodic Integration in Soar

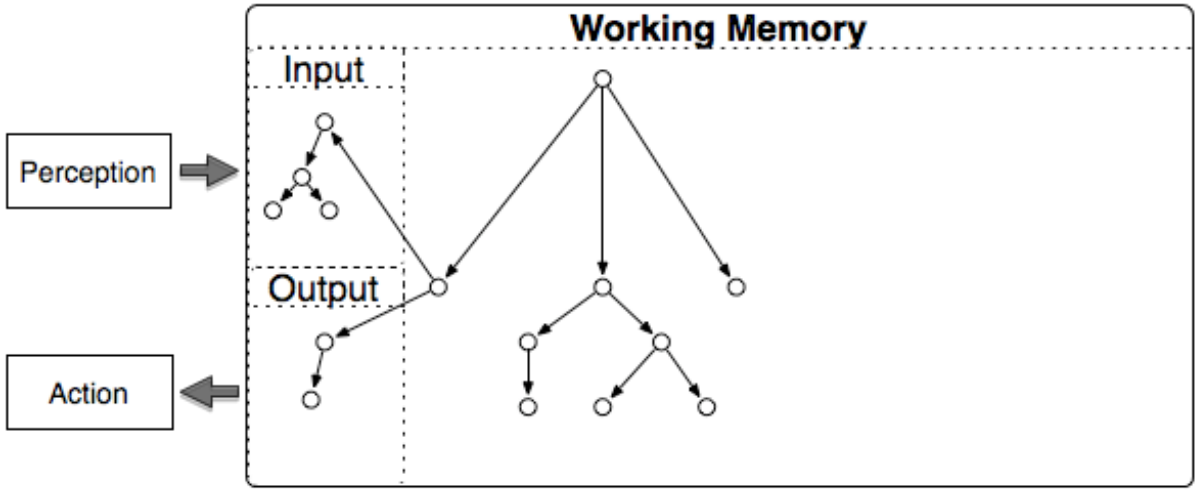

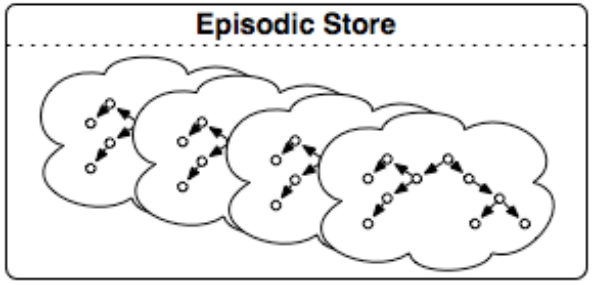

Episodic Integration in Soar

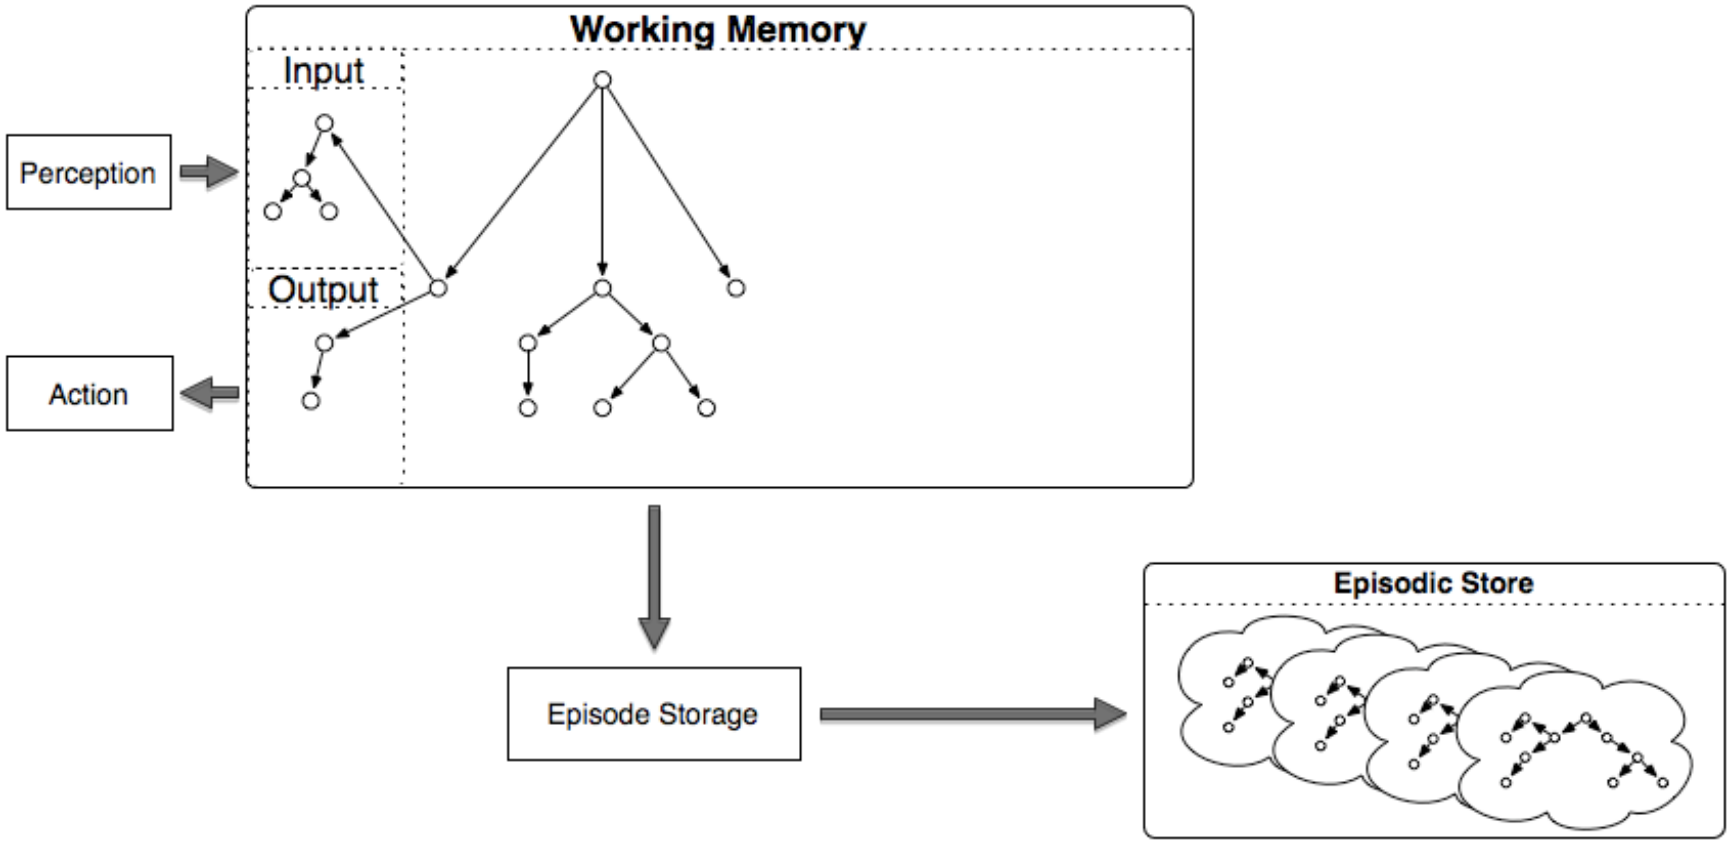

Episodic Integration in Soar

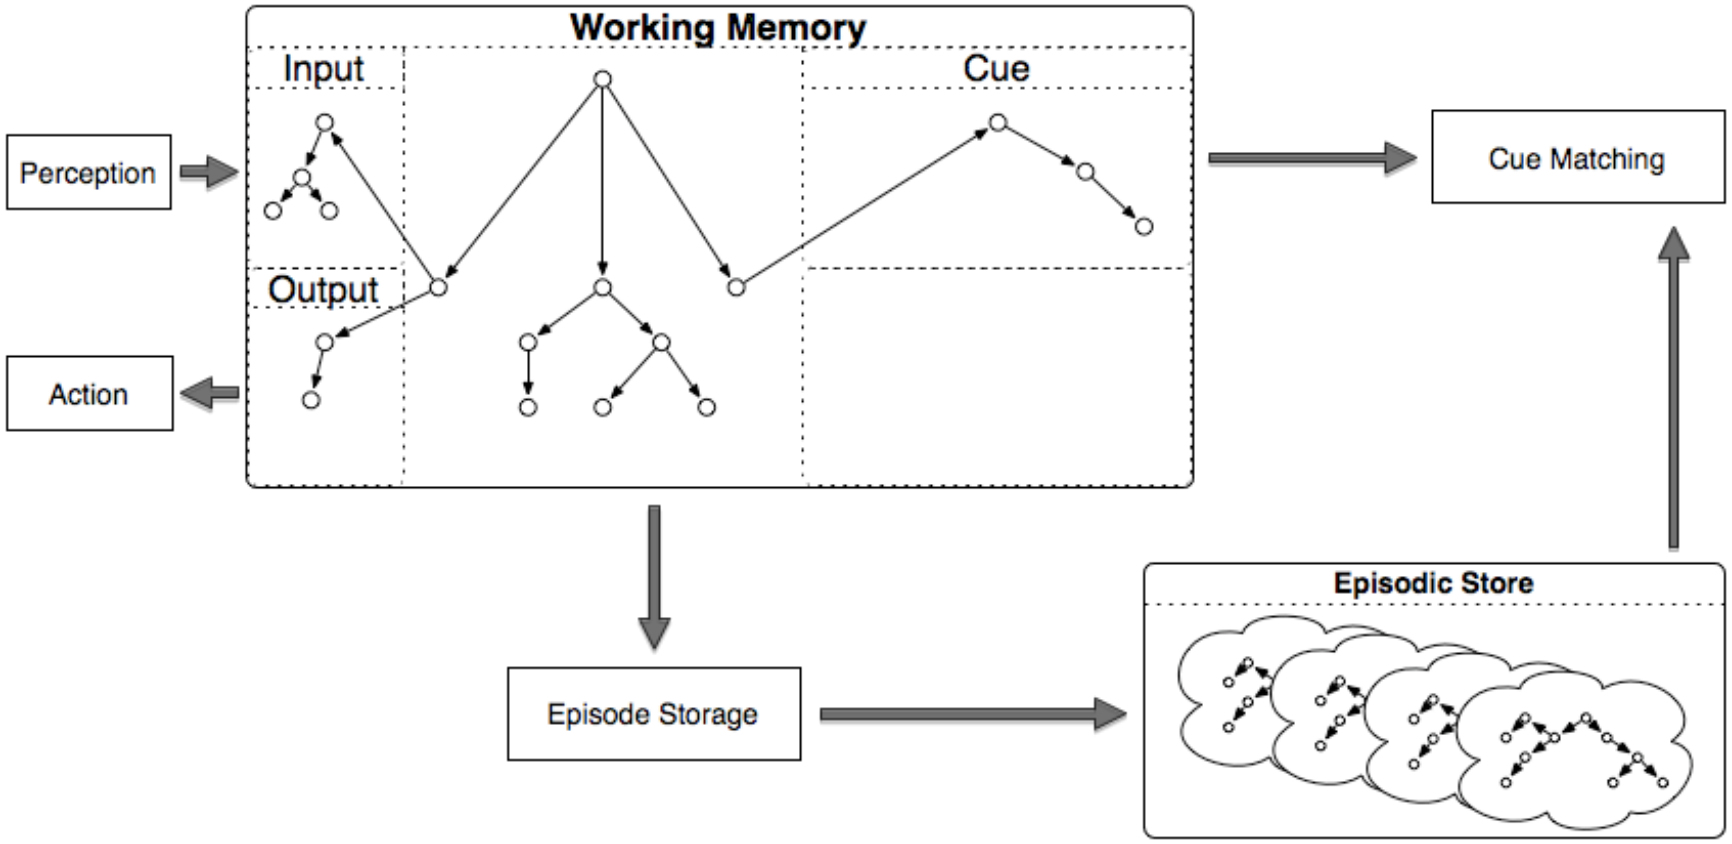

Episodic Integration in Soar

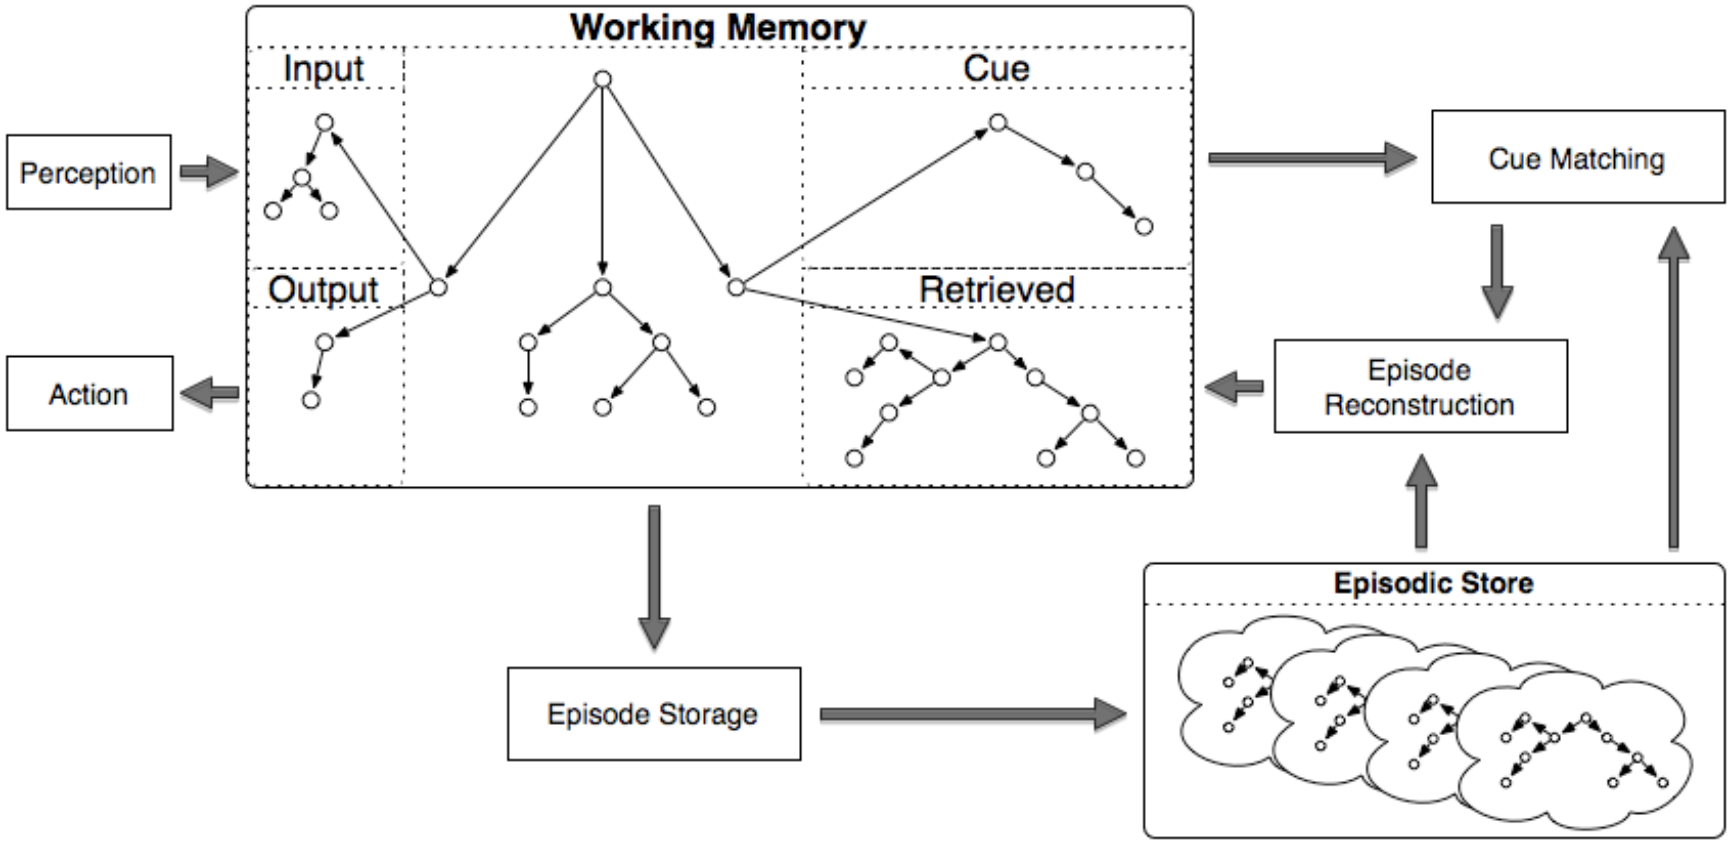

Episodic Retrievals

Given

- Contextualized cue of working memory structures
- Episodic store

Retrieve

- Single, best episode

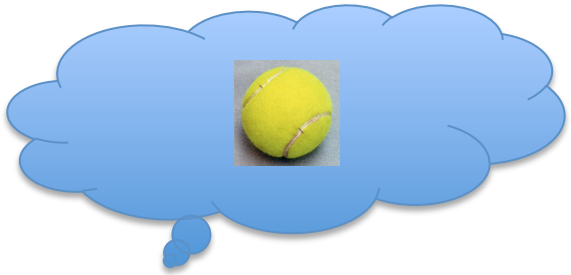

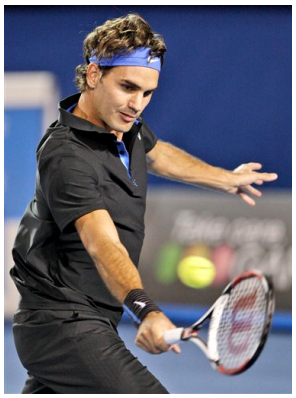

Episodic Dissociation

Can the *Procedural* or *Semantic* memory systems efficiently support *Episodic* retrievals?

Argument for Episodic Dissociation

Semantic

Discounts knowledge context

Procedural

Unlikely any approach could scale to long agent lifetime

Soar Memory Systems

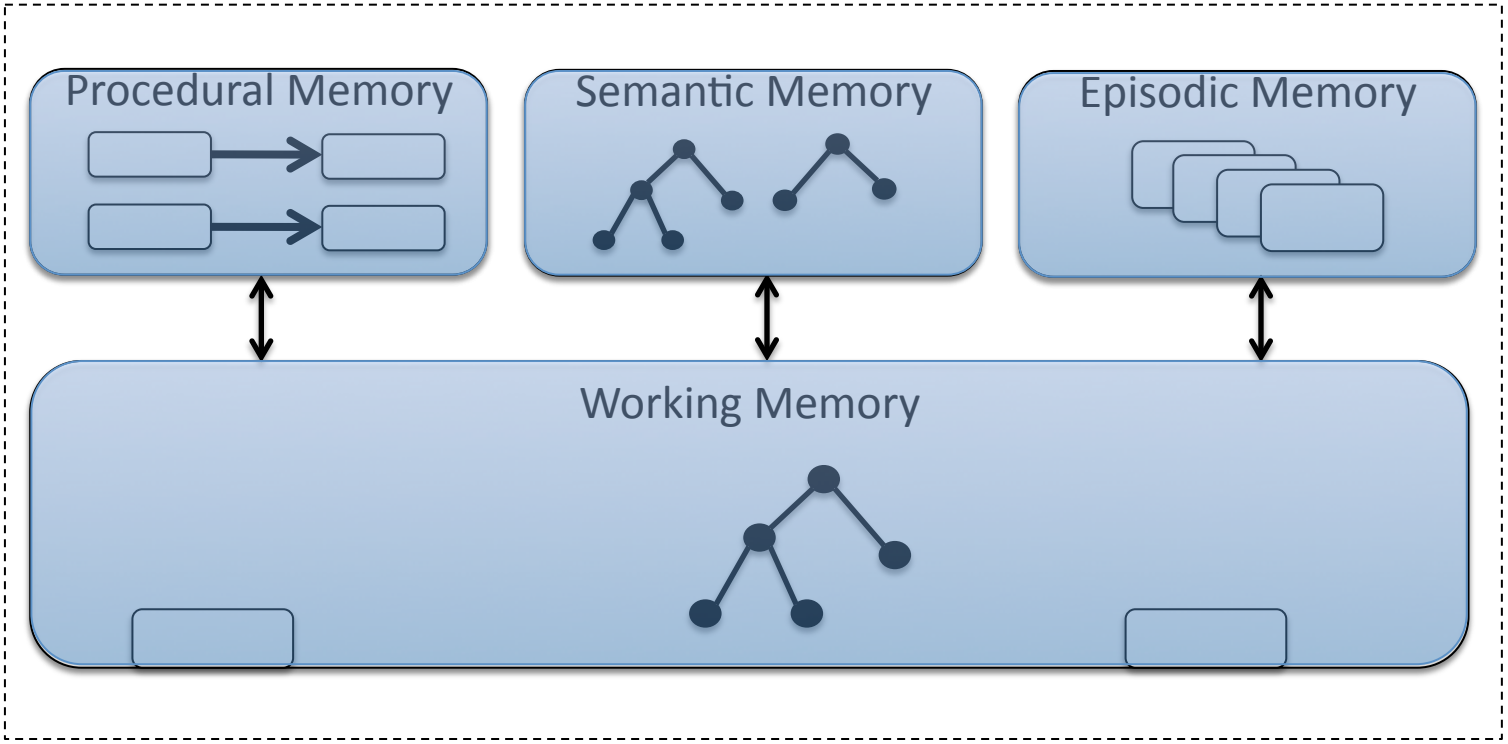

about arbitrary and novel combinations of mental entities

Inter-Memory Object Identity

The problem of managing distinct, persistent objects over multiple memory systems

- A spectrum from contextual/relational identity (weak) to globally unique identity (strong)

Introduces tradeoffs in memory system implementation between learning generality/ correctness and retrieval efficiency

- See paper for detailed evaluation

Object Identity in Soar

Learned object identity

- Initially, mental entities are *weakly* identified
- Once stored to semantic memory, they gain *strong*, globally unique identification

Learning starts general and becomes increasingly specific

Places a strong burden on the agent and architecture to implement an effective policy for learning persistence

Future Work 

- Learning object identity
  - Functionally optimal policy
- Symbol learning from regularities in perceptual data
- Incorporating non-symbolic modalities
- Developing agents that *learn* effective control over multiple memory systems

Thank You!

Questions?