Effective Scaling of Long-term Memory for Reactive Rule-based Agents

Nate Derbinsky, Ph.D. (ABD)

Computer Scientist

University of Michigan

Intelligent Agents

- Autonomous, Persistent
- Observes and acts upon an environment
- Uses and learns knowledge
- Directs activity towards achieving **goals**

Example: Ground Robotics

Environment

- Multiple terrains, other agents, weather
- Movement, obstructions

Tasks

- Patrols, search-and-rescue, exploration, experiments, …
- Terrain, topological relations, traffic patterns, …

Agent

- Days – years
- Autonomy and interaction with other agents, handlers

Scenario #1

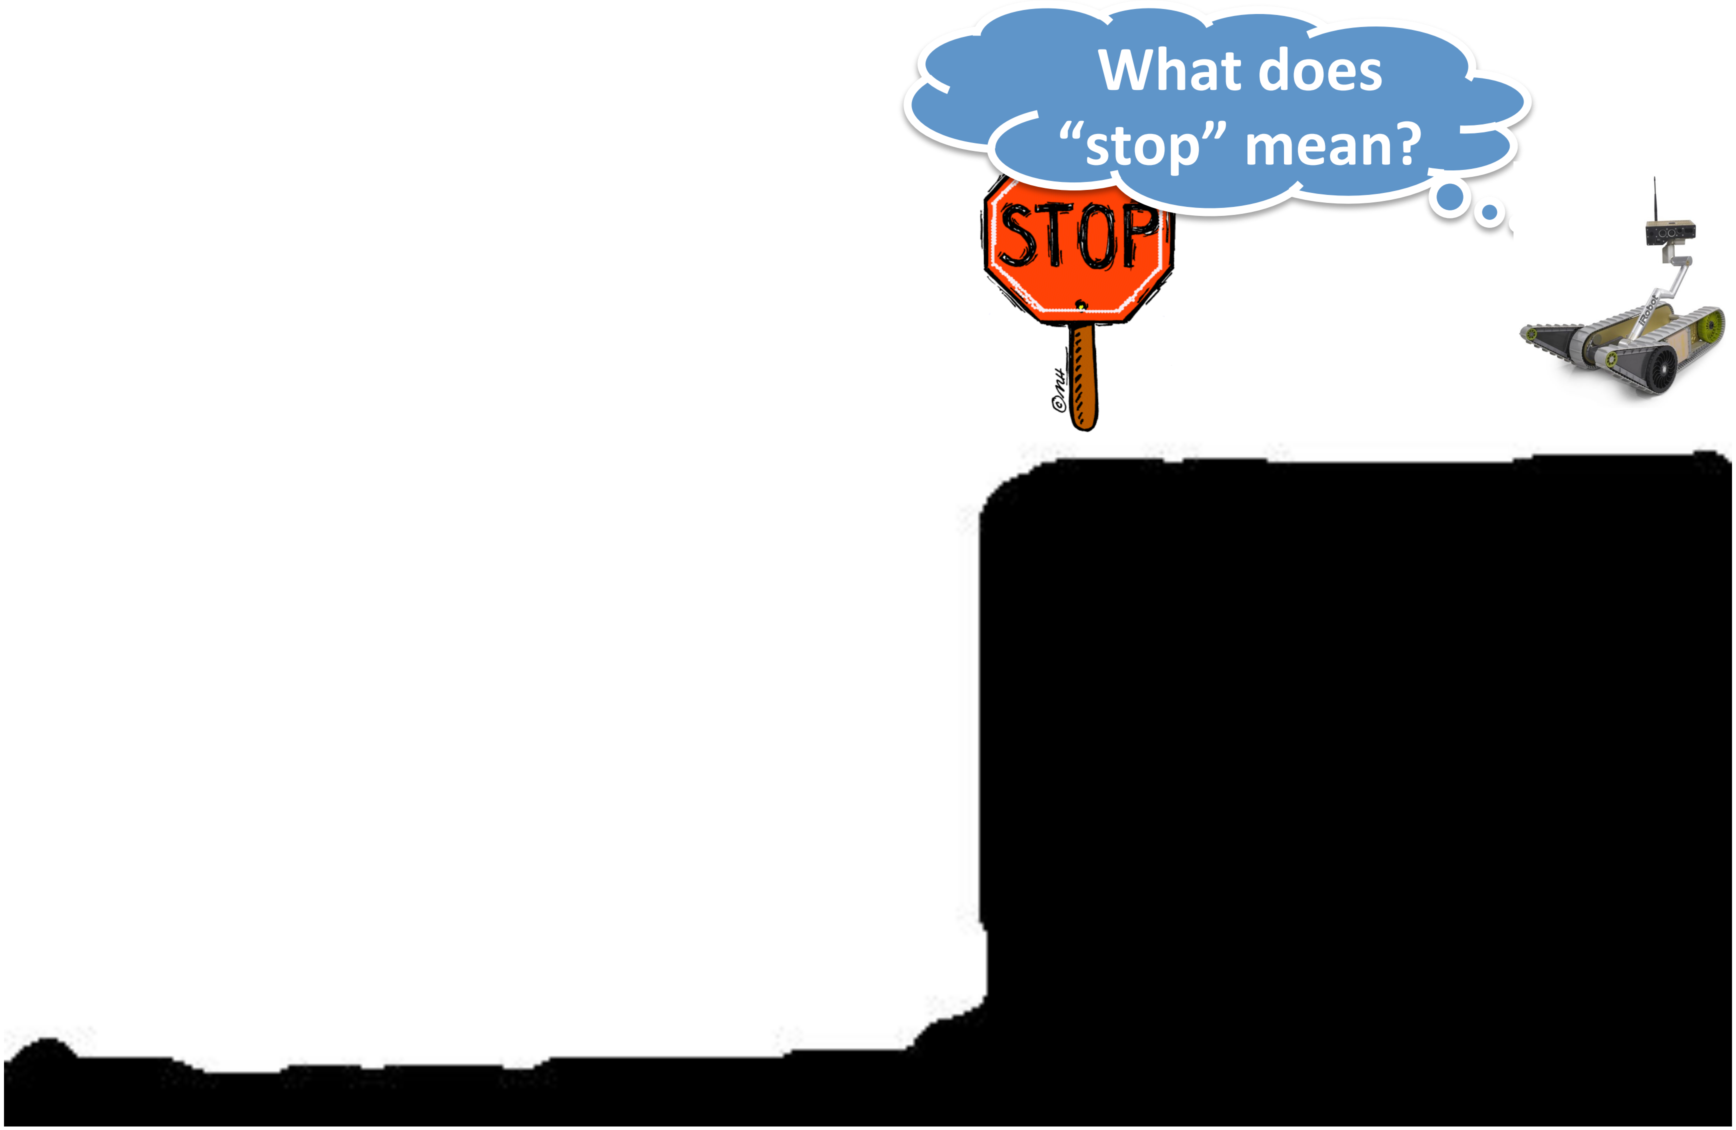

Scenario #2

What happened last time?

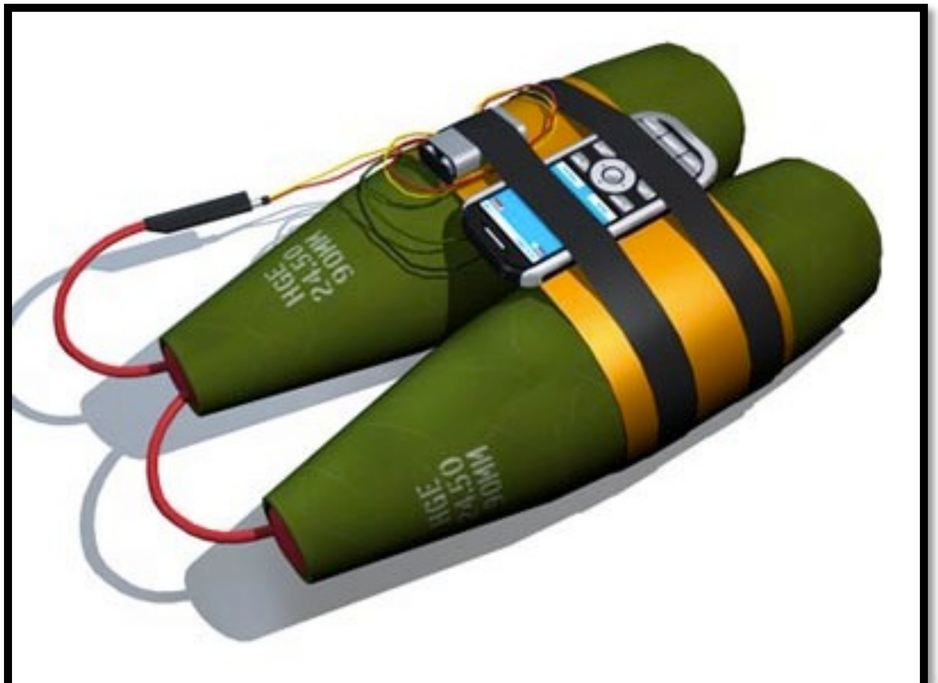

A Common Problem

Agents need <u>effective</u> access to diverse information

- Factual
- Experiential

Agents need to maintain <u>real-time reactivity</u> in dynamic environments

< 50 msec.

Approach: RBS

Rules

Combinatorial set of possible conditions

Working Memory

Match time scales with WM size

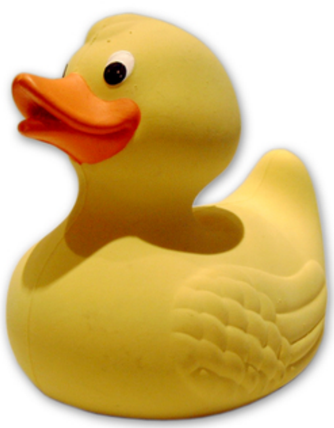

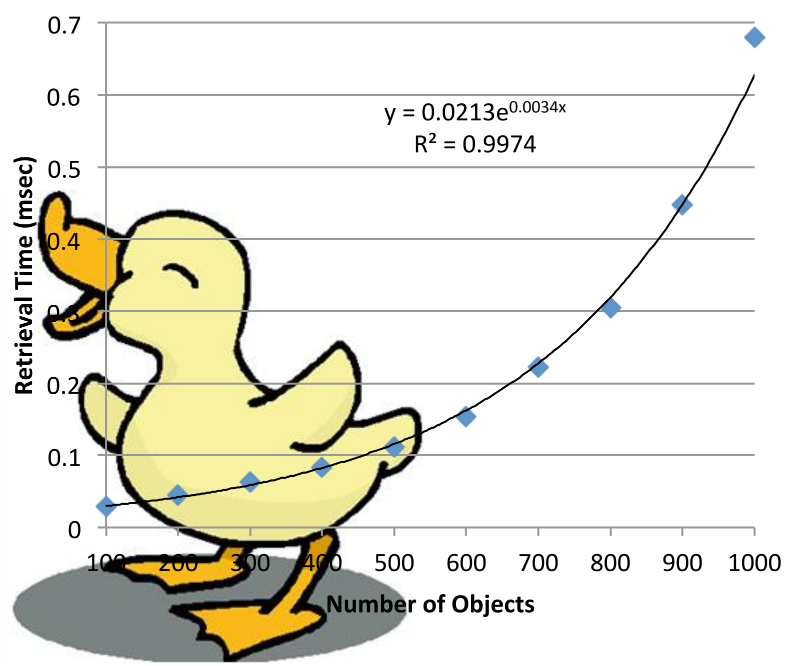

Approach: Remote Search

Factual

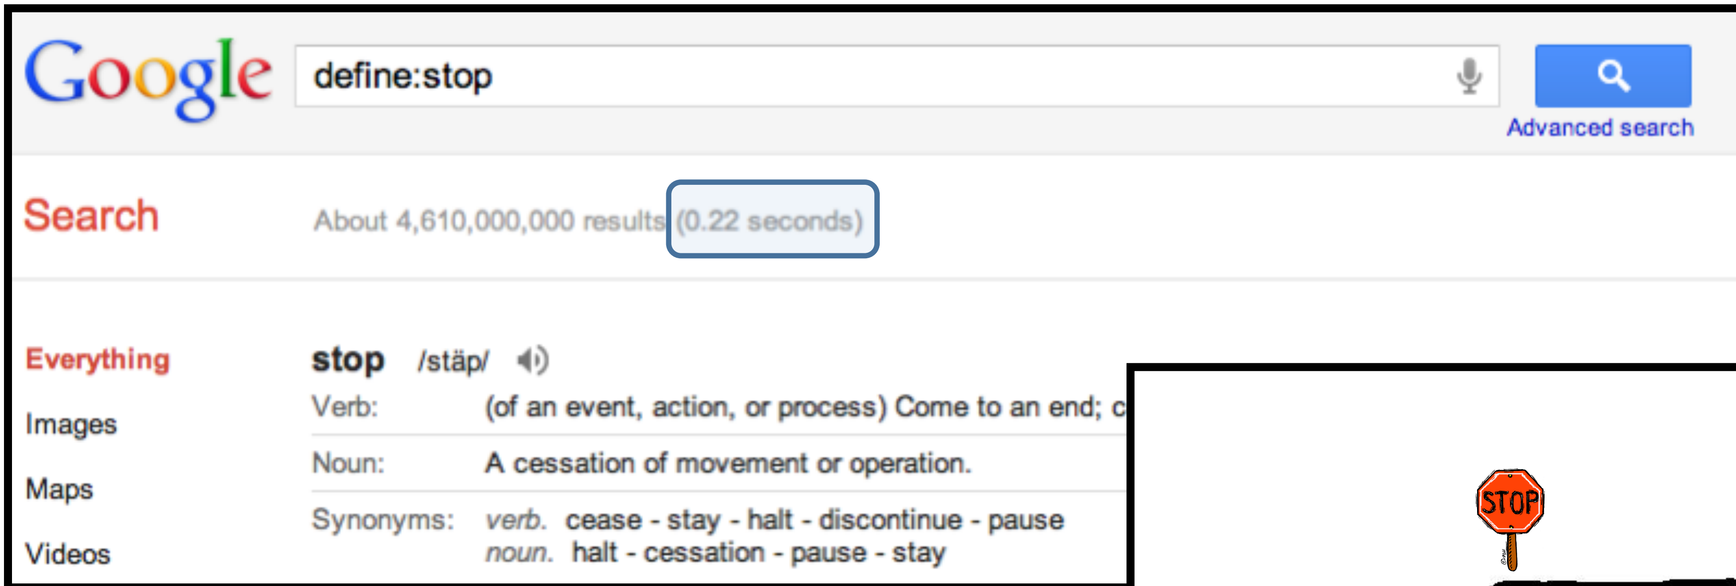

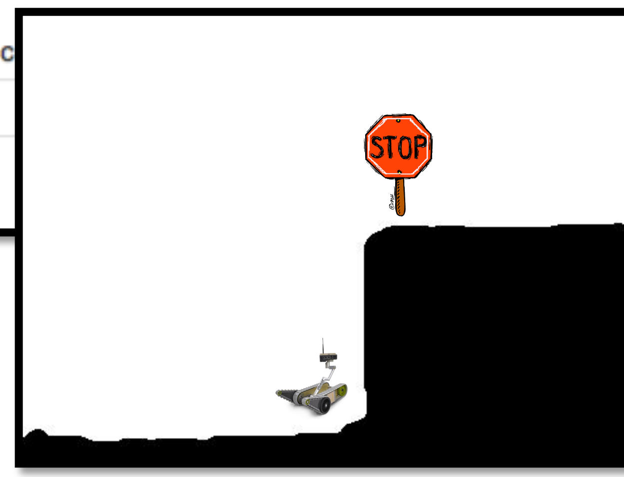

Approach: Remote Search

Experiential

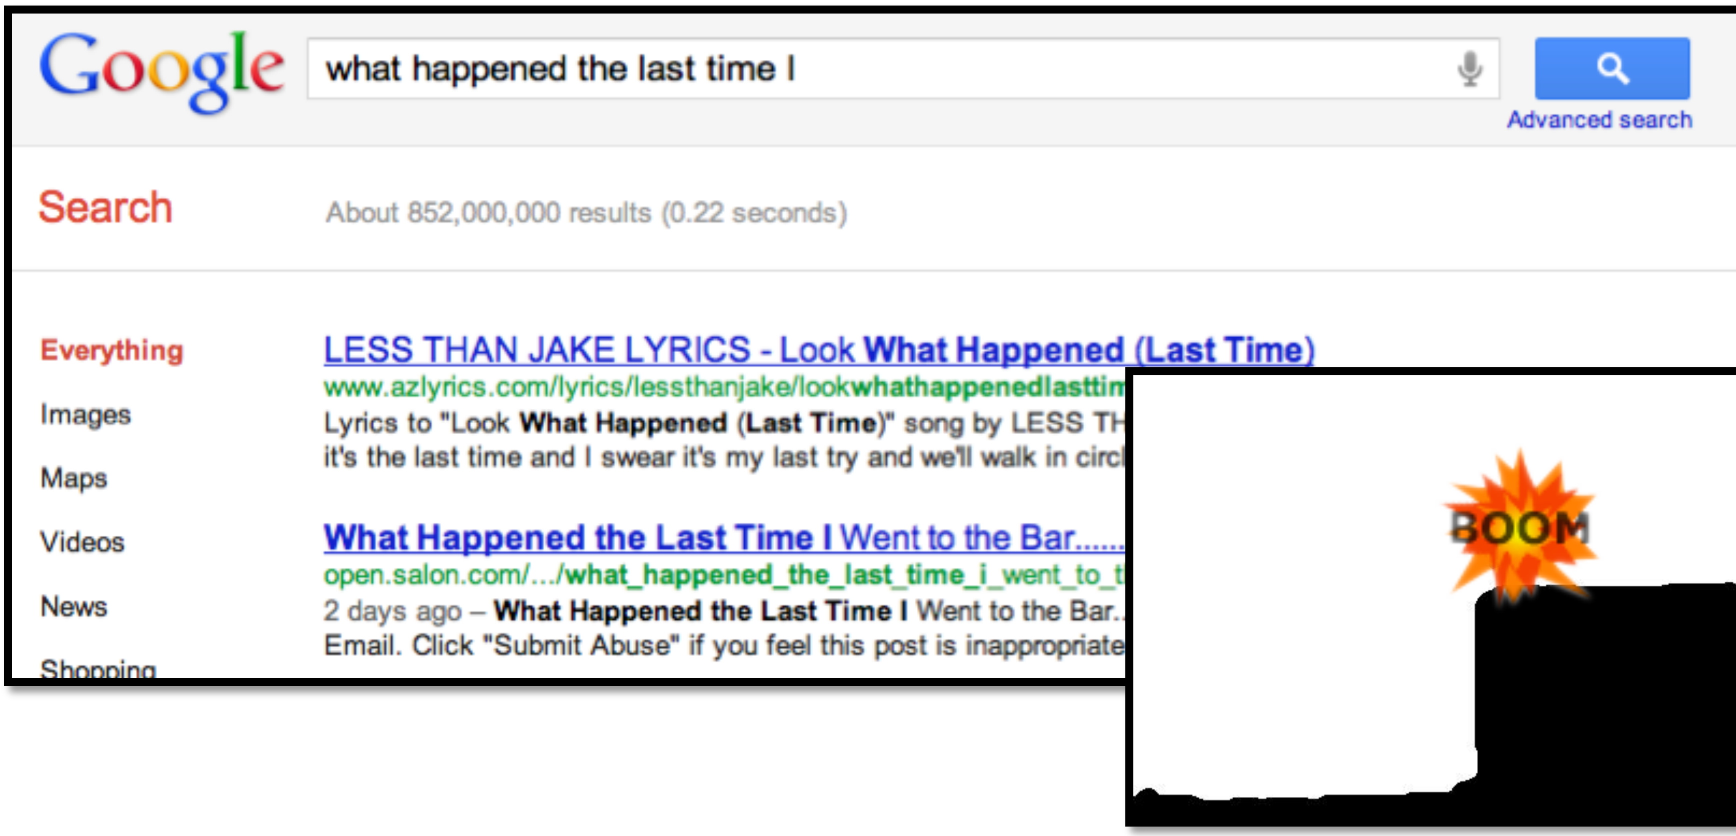

Long-Term Memory (LTM)

Class of mechanism to help agents cope with dynamic, partially-observable environments

- **Encodes** experience
- **Stores** internally
- Supports **retrieval**

An Interesting Dichotomy:
Stored Context

Semantic
"knowing"

Episodic
"remembering"

Agents with LTM are functionally enhanced across a variety of problems

The Problem
LTM for Reactive Agents

Support…
- incremental encoding and storage of experience
- access to stored knowledge

Requirements
- Reactivity: decisions < 50msec.
- Scalability: support large amounts of knowledge
- Generality: effective across a variety of tasks

This Work

Development and evaluation of two LTMs

- Integration within the Soar cognitive architecture
- Efficient algorithms and data structures
- Formal analysis & empirical evaluation

<u>Claims</u>

- Effective and efficient across a variety of tasks
- Scale computationally to…
  - Large amounts of knowledge
  - Long agent lifetimes

Outline

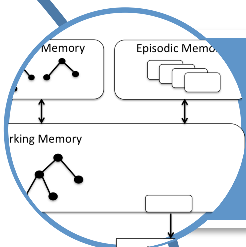

Cognitive Architecture

Semantic Memory

Episodic Memory

Cognitive Architecture

Specification of those aspects of cognition that remain constant across the lifetime of an agent

- Memory systems of agent's beliefs, goals, experience
- Knowledge representation
- Functional processes that lead to behavior
- Learning mechanisms

Goal. Develop and understand intelligence across a diverse set of tasks and domains

Research Focus

Biological Plausibility		Leabra
Psychological Plausibility	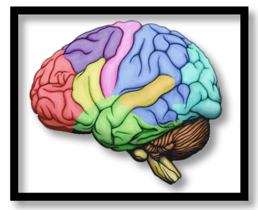	ACT-R CLARION EPIC
Agent Functionality	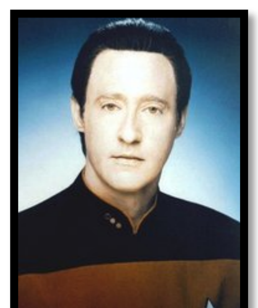	Companions ICARUS LIDA Graphical Soar

Soar: Distinctive Characteristics
[Laird, Newell, Rosenbloom 1987]

- Diverse processing and learning mechanisms that support general problem solving methods

- Efficiently brings to bear large amounts of knowledge

- Applied to many application domains
  - Language, cognitive modeling, games, tactics, robotics, …

- Public distribution and documentation
  - Major operating systems (Windows, OS X, Linux, iOS)
  - Many languages (C++, Java, Python)

Soar: Comparison to RBS Processing

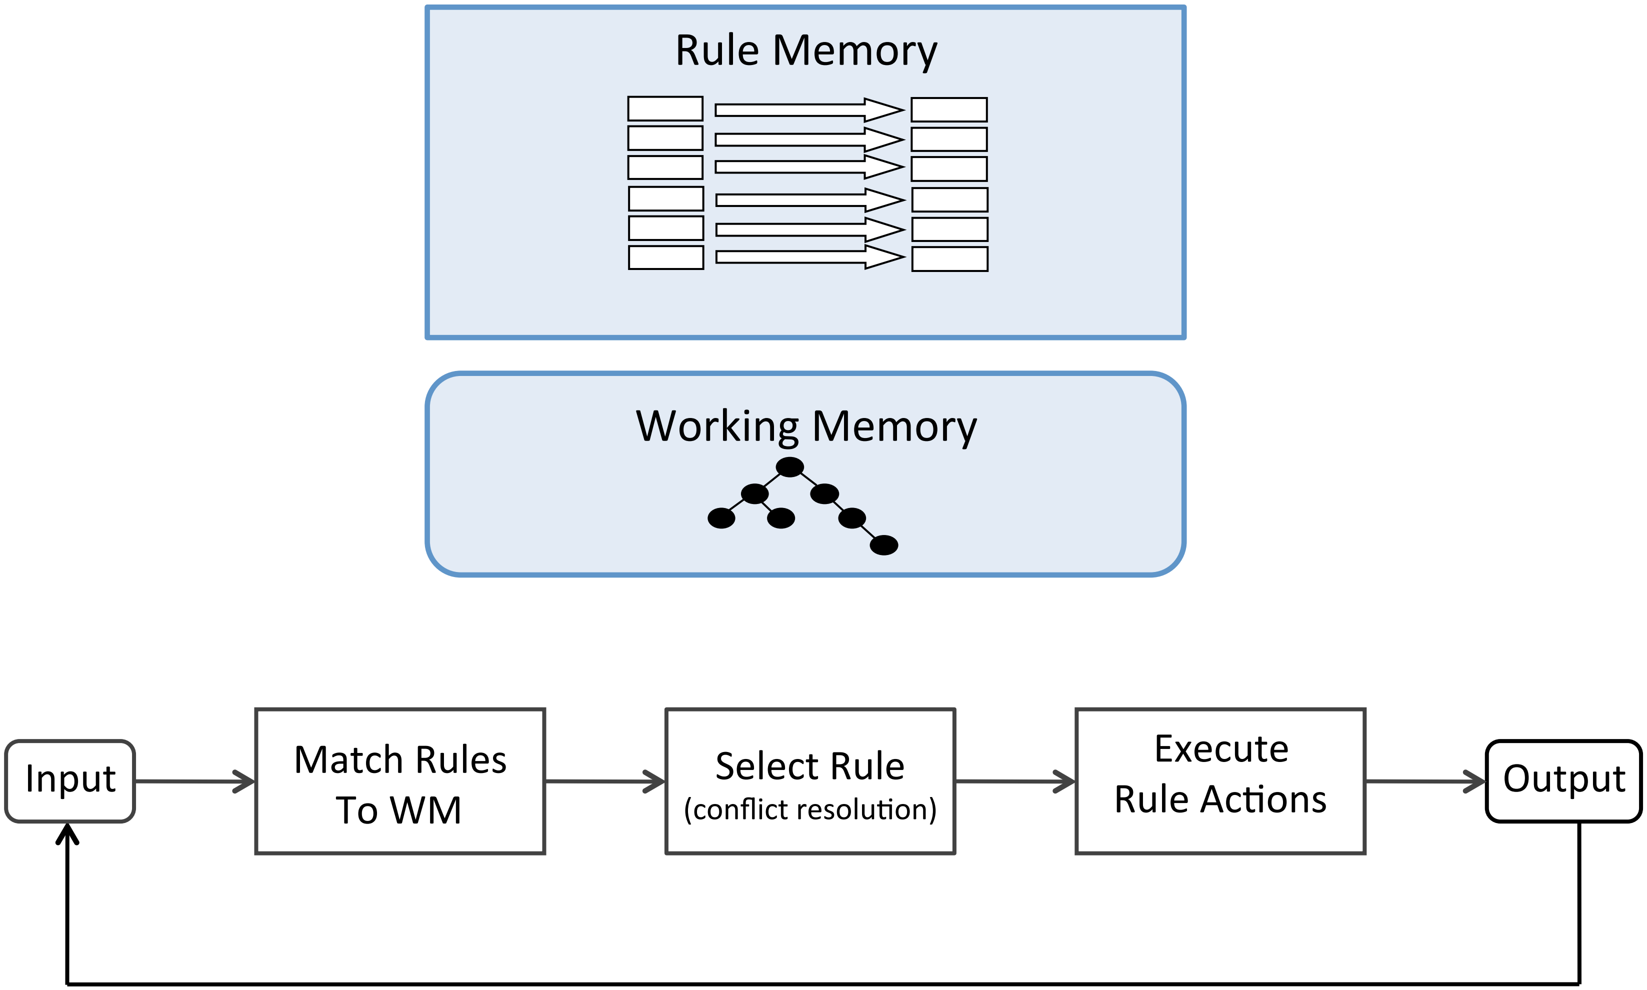

Soar: Comparison to RBS Processing

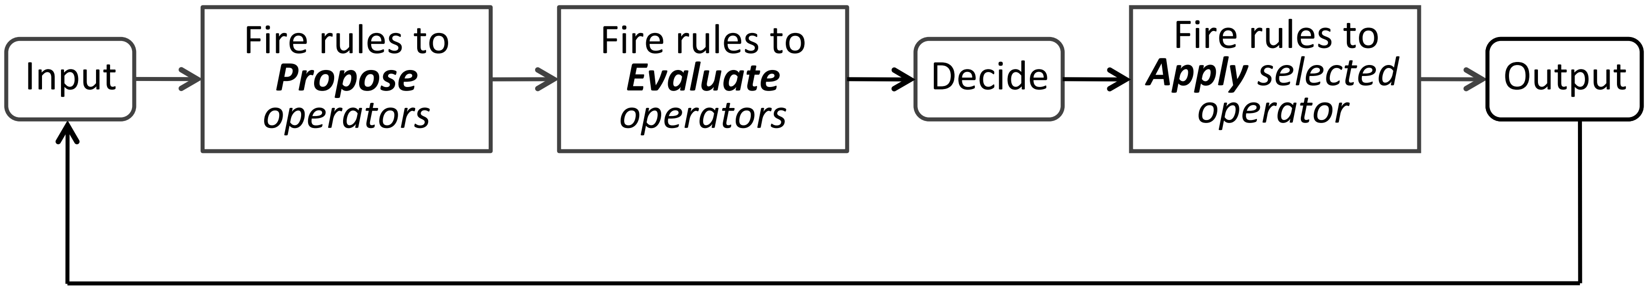

- **Operators**: basic unit of deliberation
  - Explicitly represent current operator
- Rules contain knowledge that
  - **Propose** Operators: what is possible?
  - **Evaluate** Operators: what is preferred?
  - **Apply** Operator: modify working memory
- All rules that match fire in **parallel**

Soar: Architecture

Focus on Memory

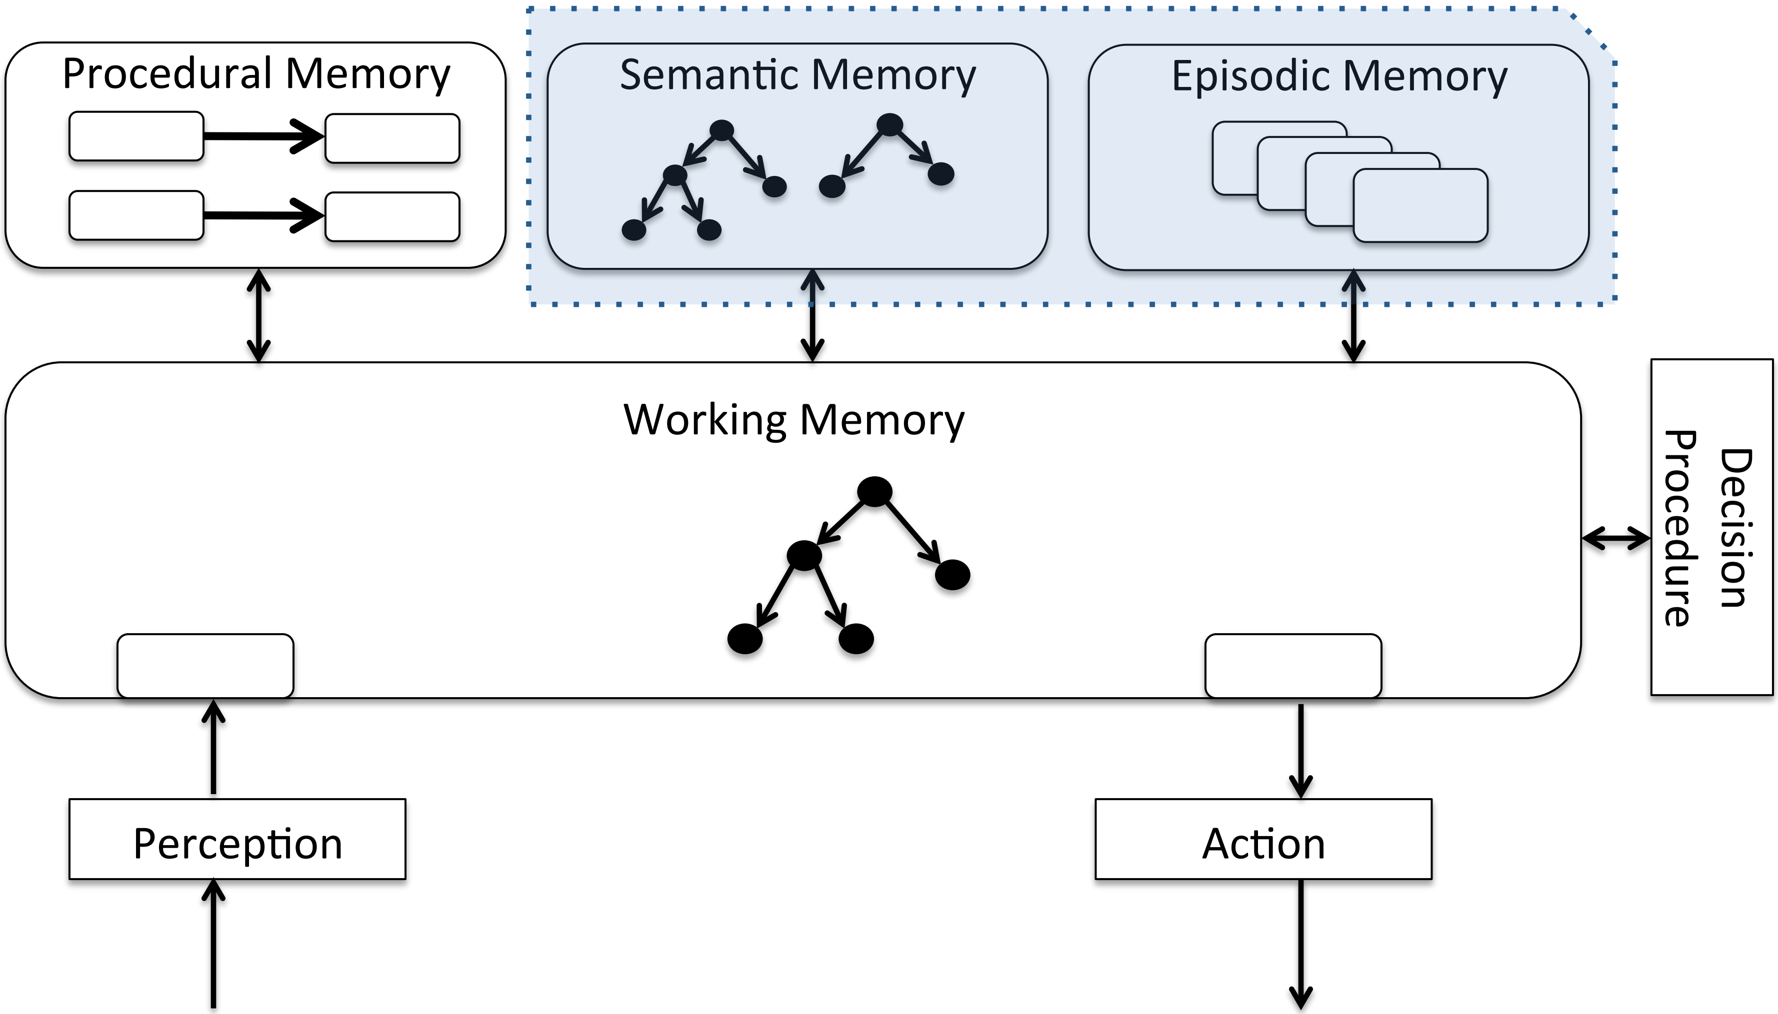

Soar: Memory Access

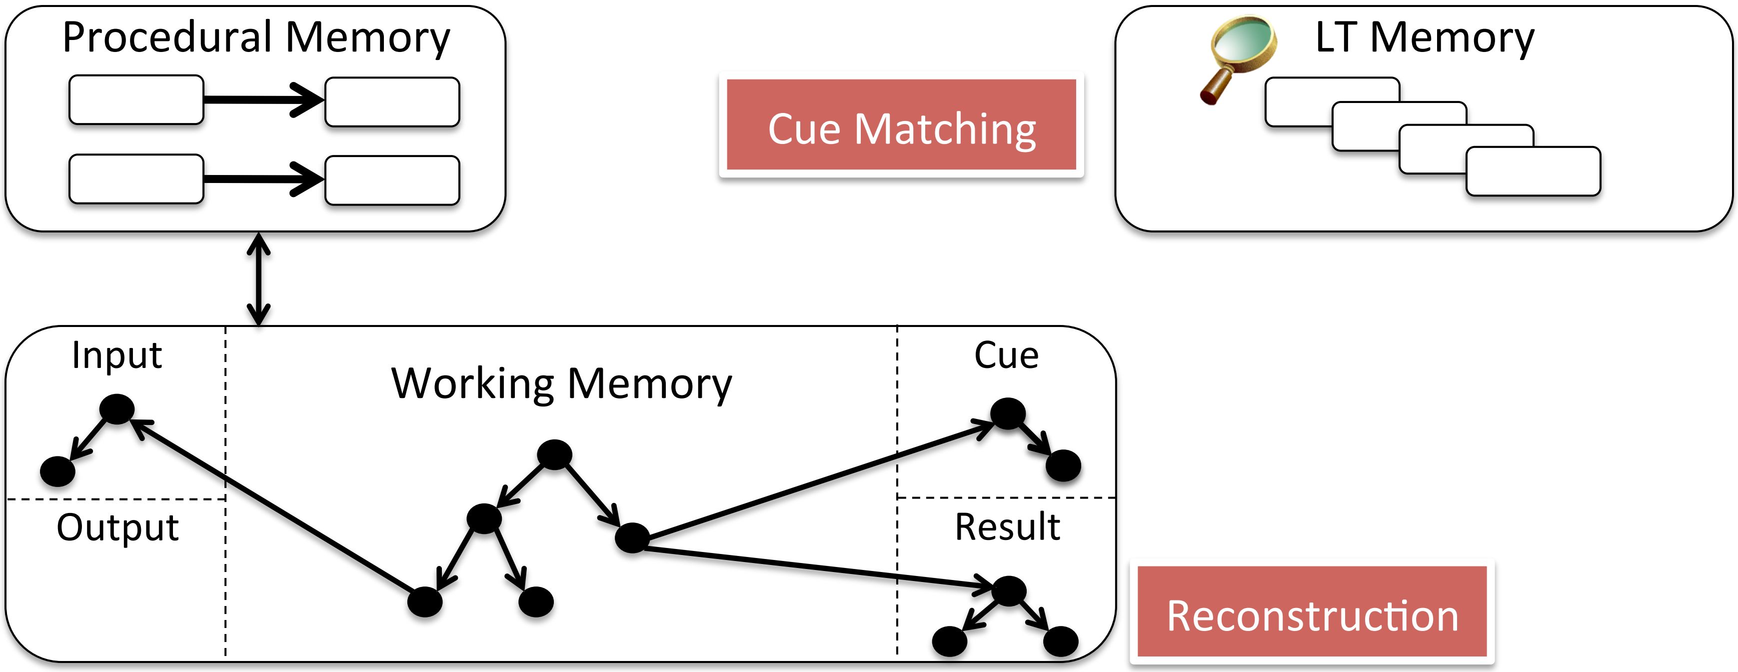

The reactivity of a Soar agent is the time required to make a decision, which includes accessing and modifying long-term memories

Soar: Memory Evaluation

Semantic Memory

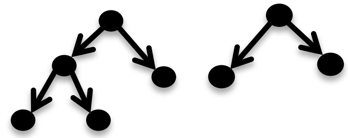

Episodic Memory

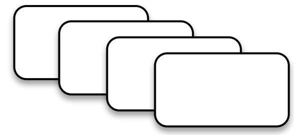

Metrics

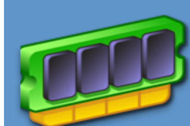

Memory Usage

Max Decision Time

Task Performance

Domains

Linguistics

Mobile Robotics

Games

Semantic Memory

Functional Analysis

- Access to large KBs
- Retrieval bias as a reasoning heuristic

SUMO

- Ontology
- 4.5K classes, 250K facts

WordNet

- Lexicon
- 212K senses, 820K assertions

Cyc

- “Common Sense”
- 500K concepts, 5M facts

Semantic Memory

Integration

Representation

- Symbolic triples

Encoding

- Deliberate

Cue Semantics

- Feature subset

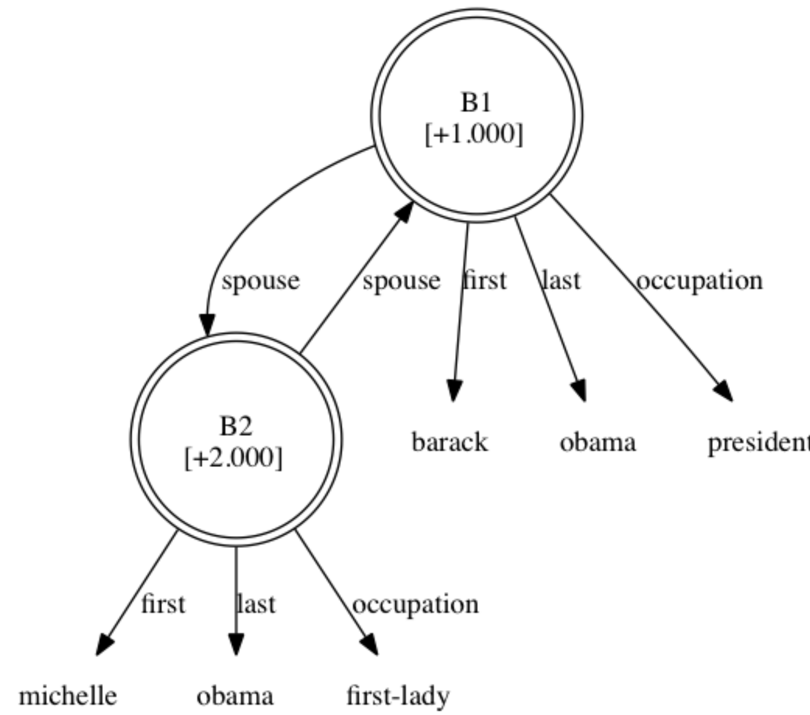

Semantic Memory

Efficient Implementation – [ICCM '10; AAAI '11]

Mapping to Set-Valued Stores

- Incremental inverted index
- Statistical query optimization
- Heuristic search

Locally Efficient Bias Functions

- Computation takes O(1) time, affects O(1) memories
- Class includes *f*(useful historical properties)

Semantic Memory

Empirical Evaluation – [ICCM '10]

Synthetic Scaling Study

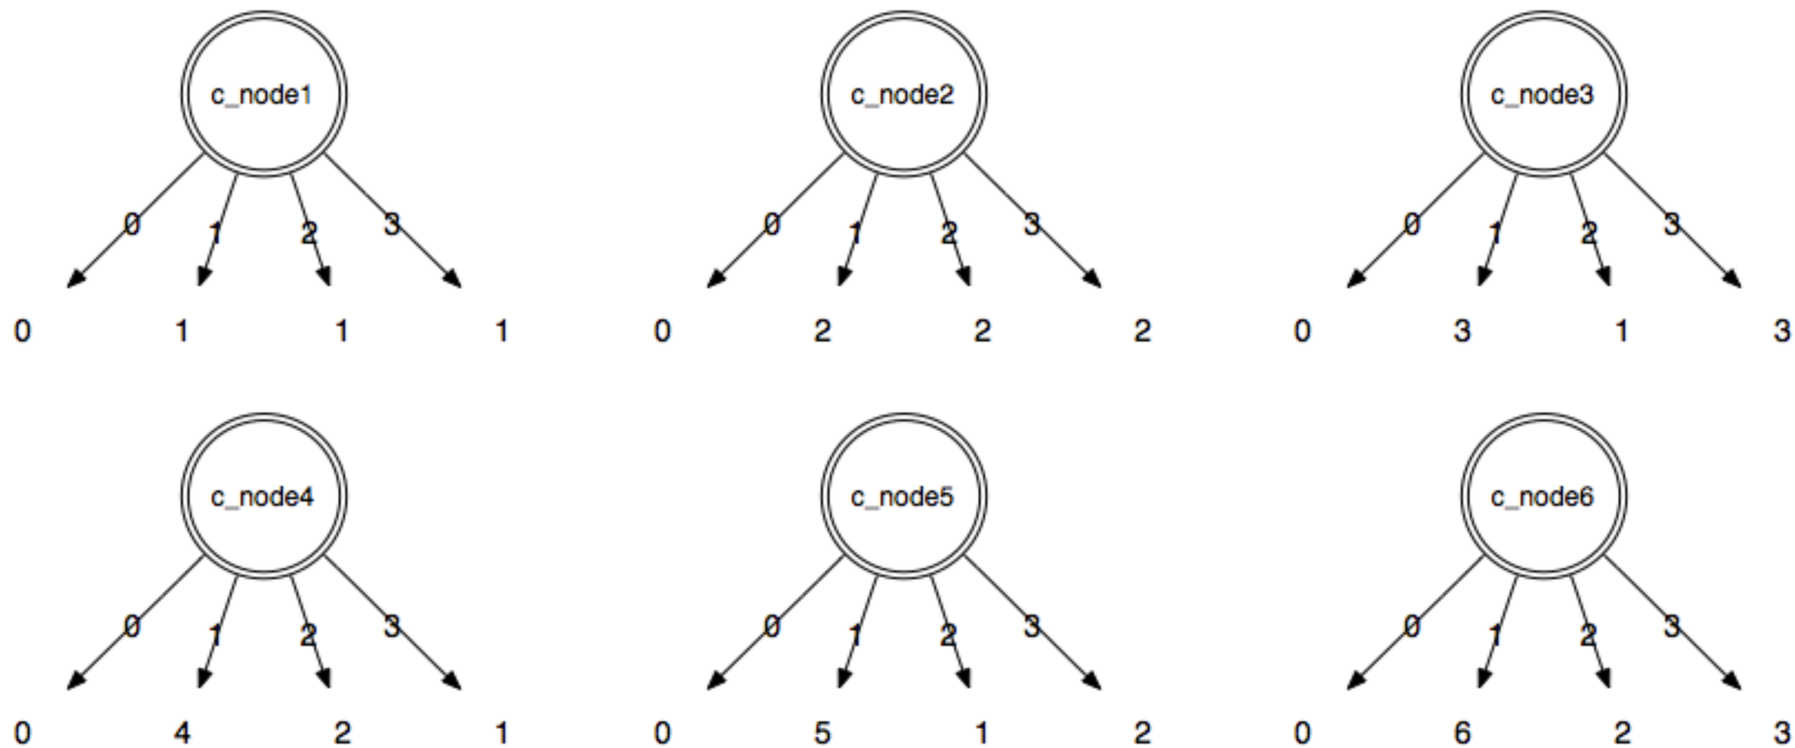

Semantic Memory

Empirical Evaluation – [ICCM ‘10]

Synthetic Scaling Study

- Scaling parameter: k
- Nodes = k!, Edges = [k+1]!

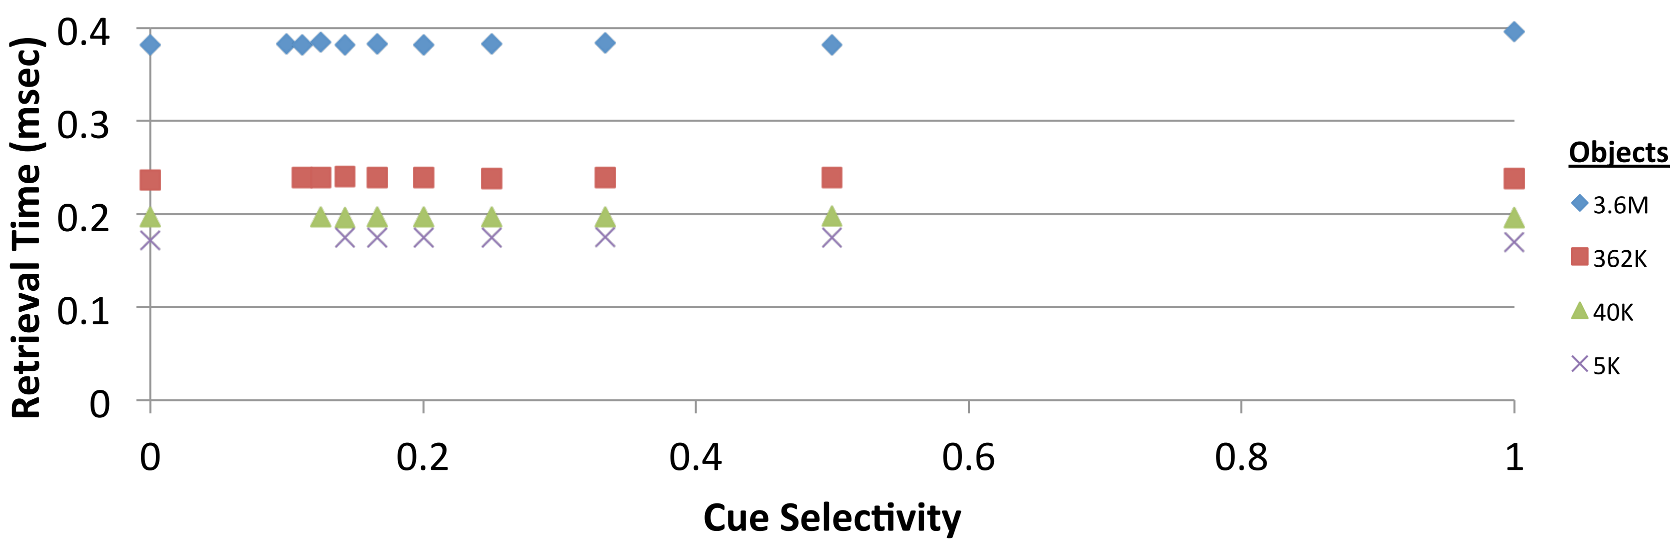

Semantic Memory
Empirical Evaluation – [ICCM '10]

Synthetic Scaling Study

- Scaling parameter: k
- Nodes = k!, Edges = [k+1]!

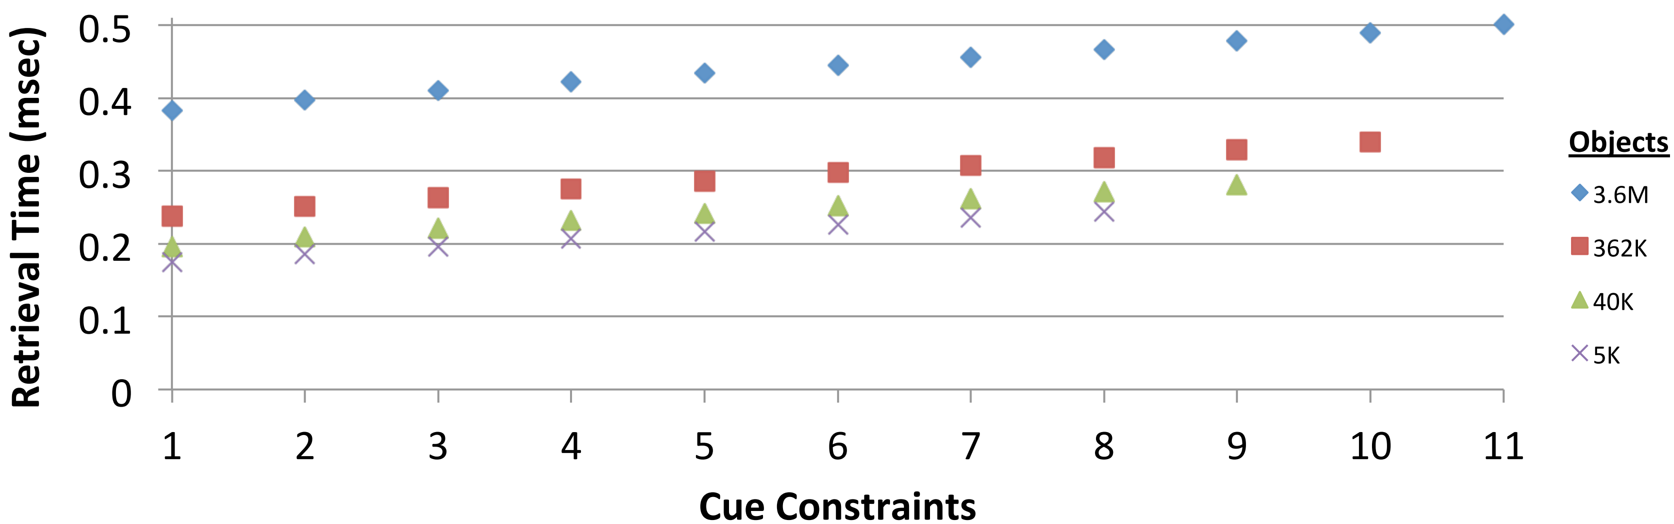

Semantic Memory
Empirical Evaluation – [ICCM '10]

Lexicon Queries: WordNet [Douglass et al. 2009]

<u>Experimental Setup</u>

- 10 random nouns
- Full sense (7 feat's)
- 10 trials

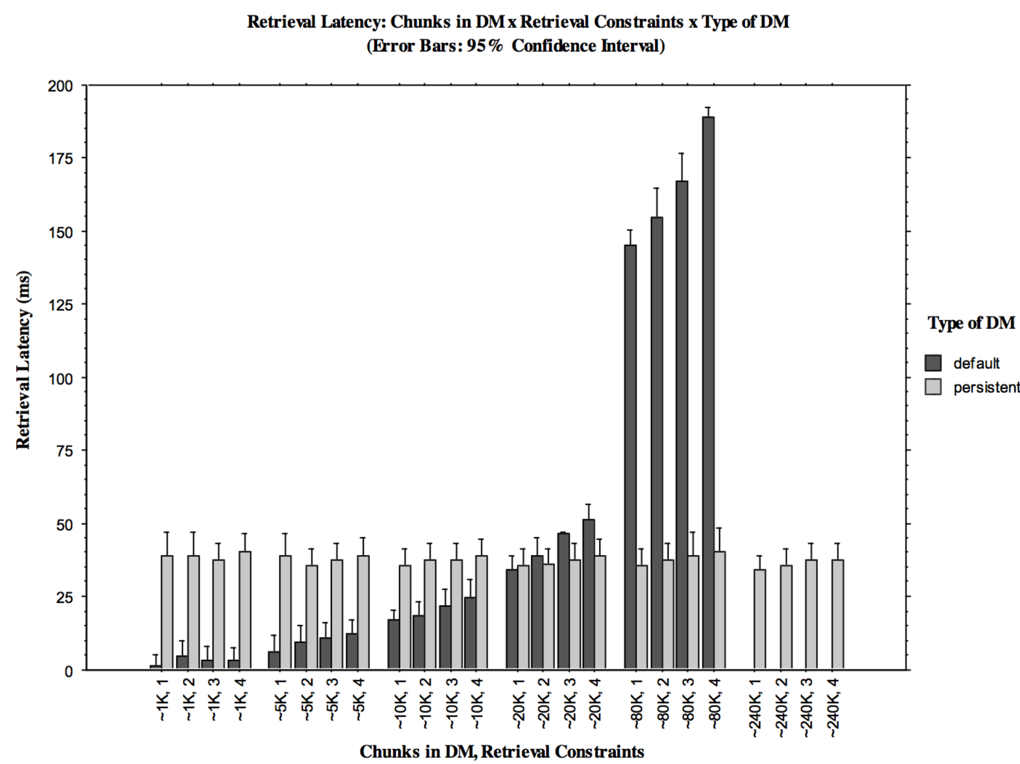

<u>Results</u>

≤ 0.3 msec (σ=0.0108)

>100x faster

>3x more data

Semantic Memory

Empirical Evaluation – [AAAI '11]

<u>**Long-Term Memory**</u> <u>**Agent**</u>

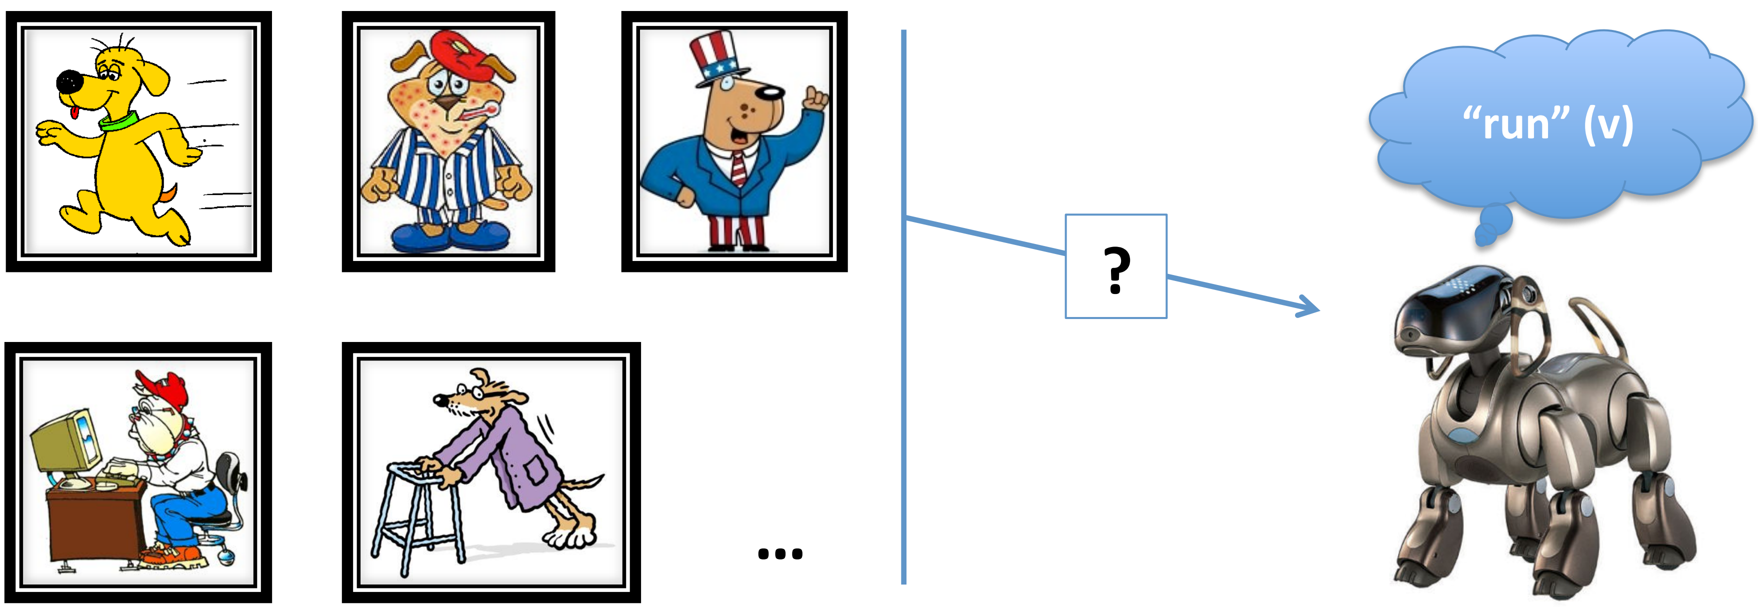

Problem. Ambiguous Cues
Hypothesis. Retrieval History is Useful
Application. Word Sense Disambiguation

Semantic Memory

Empirical Evaluation – [AAAI '11]

Word Sense Disambiguation

Experimental Setup

- Input: "word", POS
- Given: WordNet v3
- Correct sense(s) after each attempt

Efficiency

≤ 1.34 msec

Task Performance (2 corpus exp.)

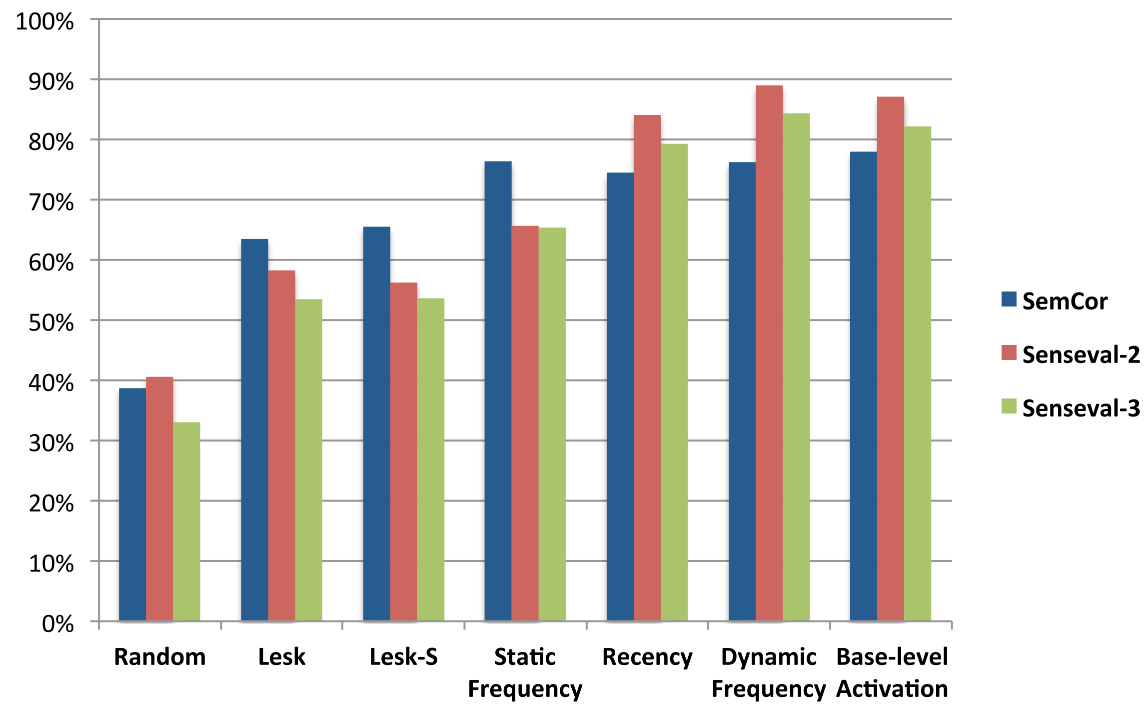

Semantic Memory

Empirical Evaluation – [BRIMS '11]

Mobile Robotics

- Incremental map learning
- Navigation and planning

Map in Working Memory

Map in Semantic Memory (< 1MB)

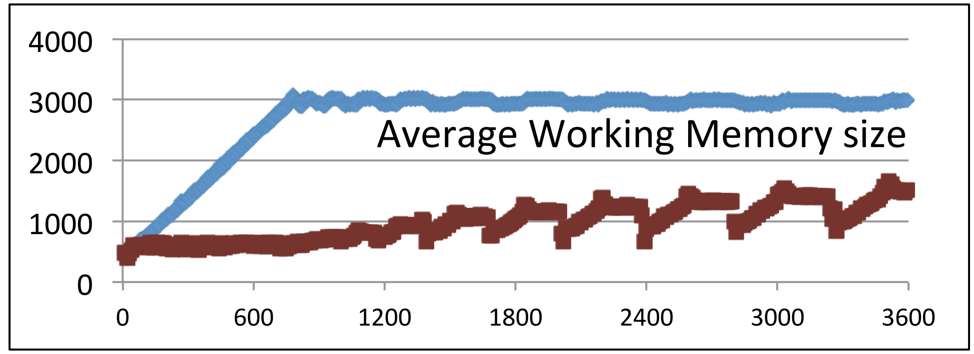

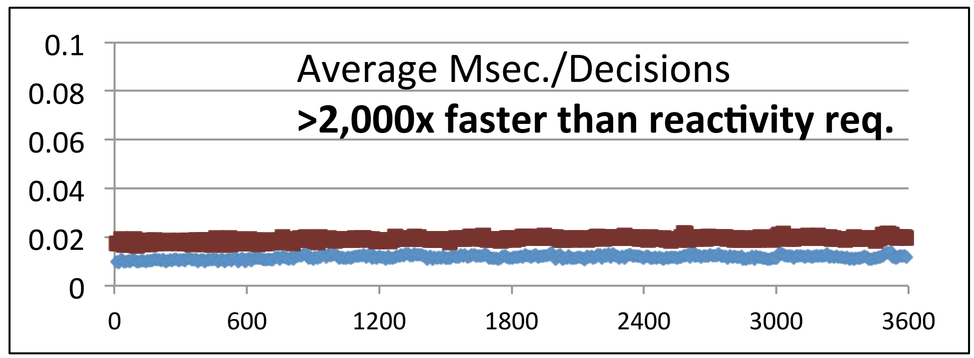

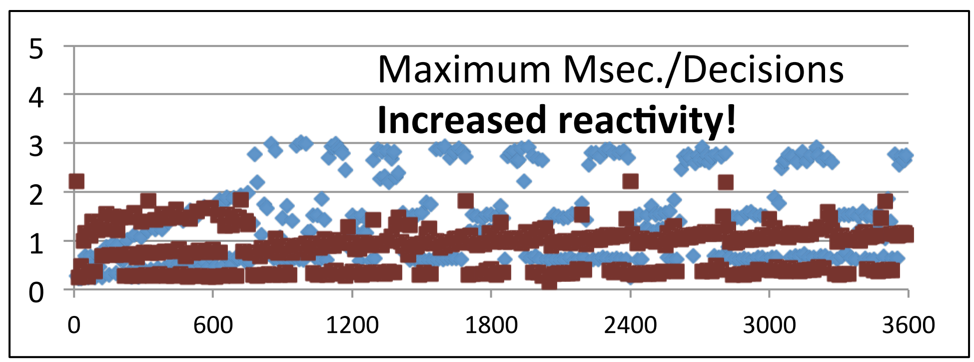

Semantic Memory

Summary

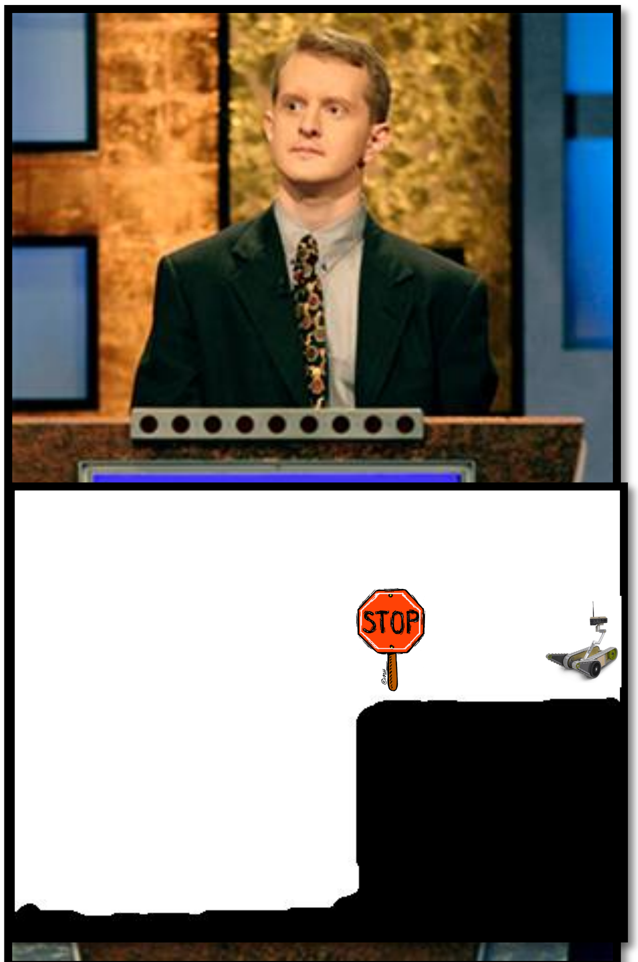

Reactivity

- More than an order of magnitude faster than reactivity requirement in practice

Scalability

- Synthetic: millions of objects
- WordNet: >820K objects

Generality

- Useful in linguistics and robotics

Episodic Memory

Humans

Long-term, contextualized store of specific events [Tulving 1983]

What you "remember" vs. what you "know"

Episodic Memory

Functional Analysis

Episodic Memory

Integration

Representation

- Episode: connected di-graph
- Store: temporal sequence

Encoding

- Automatic

Cue Semantics

- Partial graph-match
- Recency biased

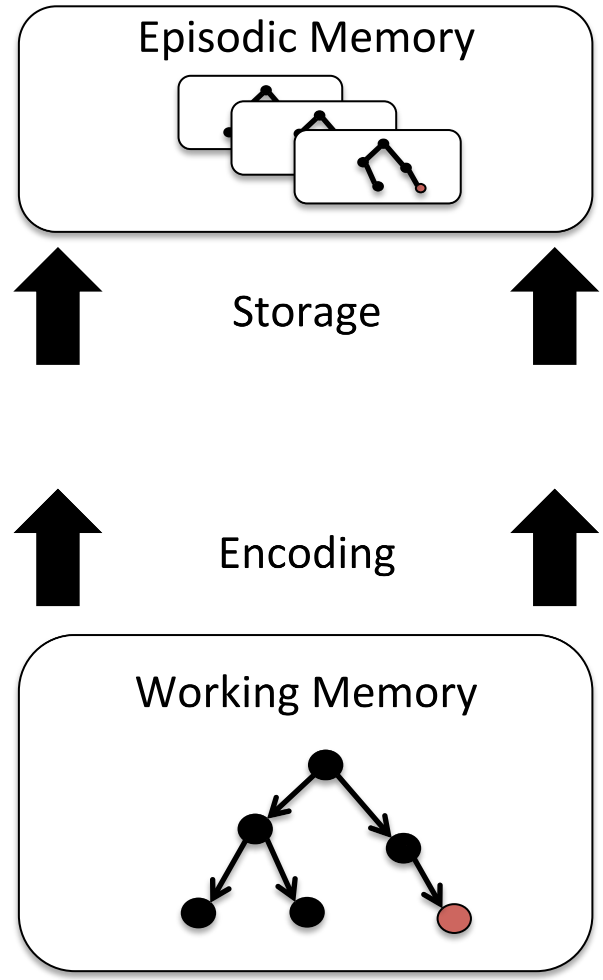

Episodic Memory

Efficient Implementation – [ICCBR '09]

Temporal Contiguity

- Interval-based representation, encoding, search, and reconstruction
- Scale with state *changes* (discrete edge +/-)

Structural Regularity

- Temporally-global structure index
- Scale with structural distinctions

Efficient Encoding & Storage

Incremental Dynamic-Graph Index

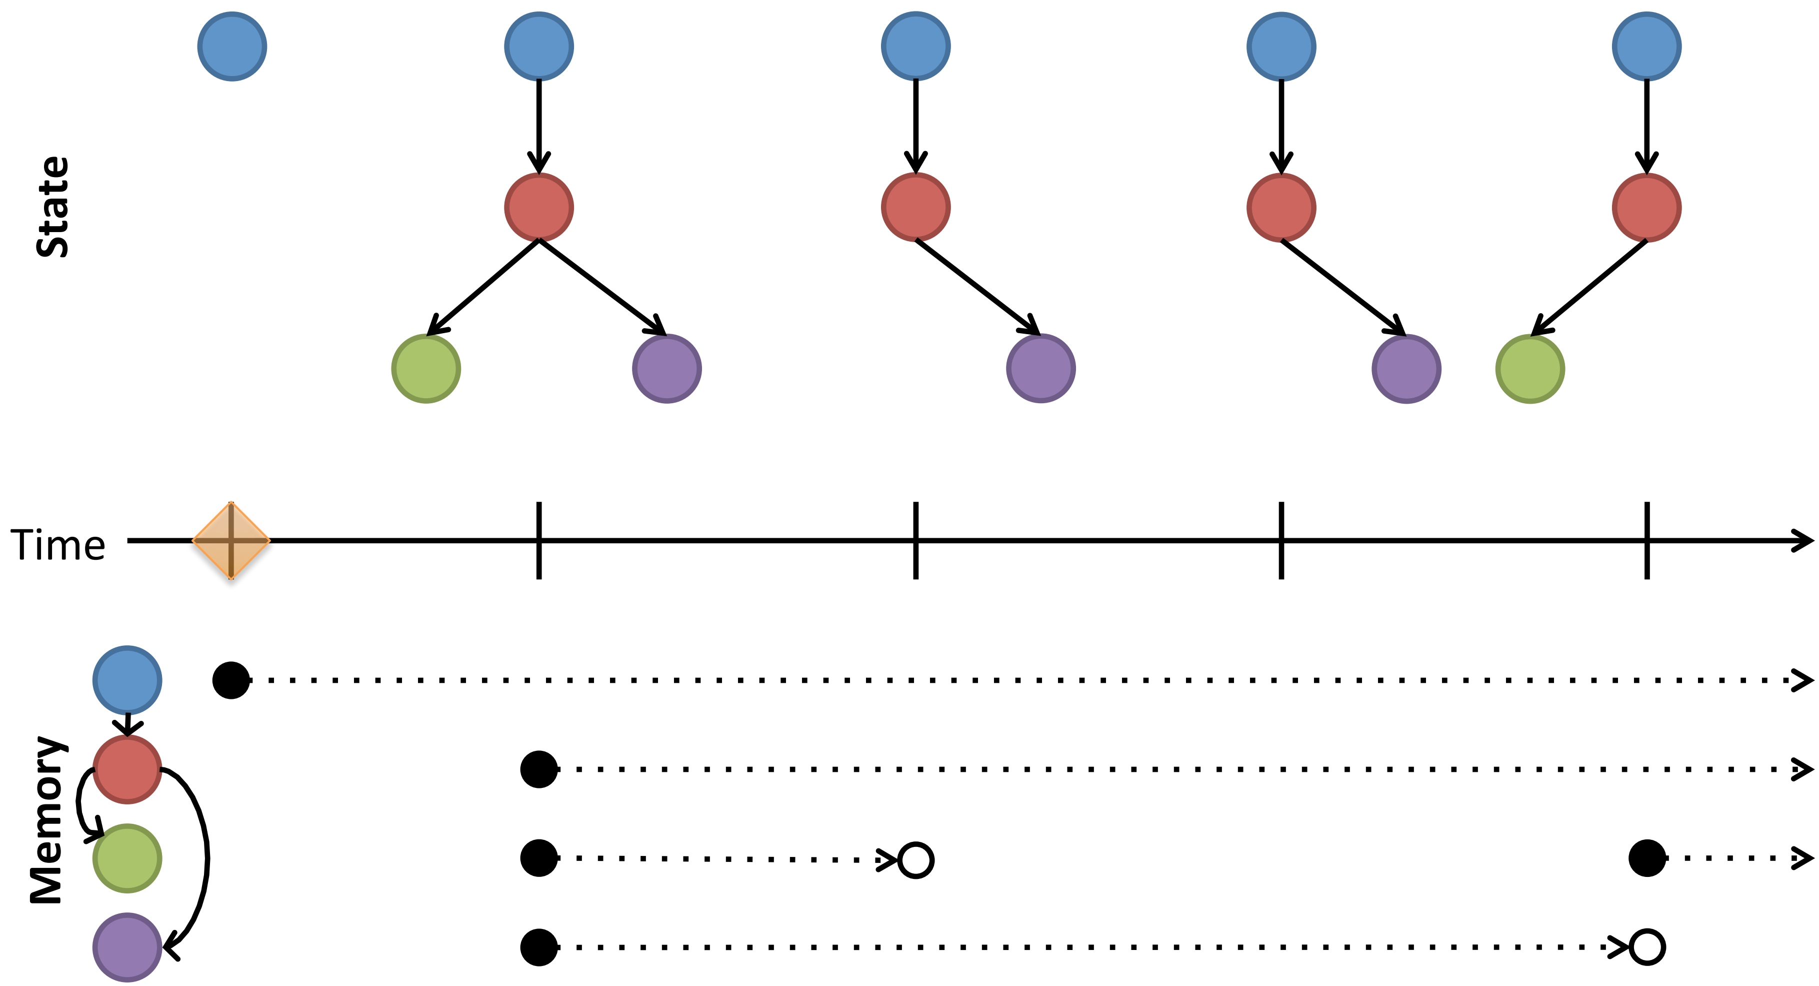

Efficient Retrieval

Overview

Cue matching is a constrained form of subgraph isomorphism

– Unify two rooted graphs with labeled edges

We utilize 2-phase matching to avoid expensive search [Forbus, Gentner, Law 1995]

– Surface: novel search algorithm (interval walk), discrimination network (DNF graph)

– Structure: standard heuristics (MCV, DFS)

Efficient Retrieval

Retrieval Algorithms

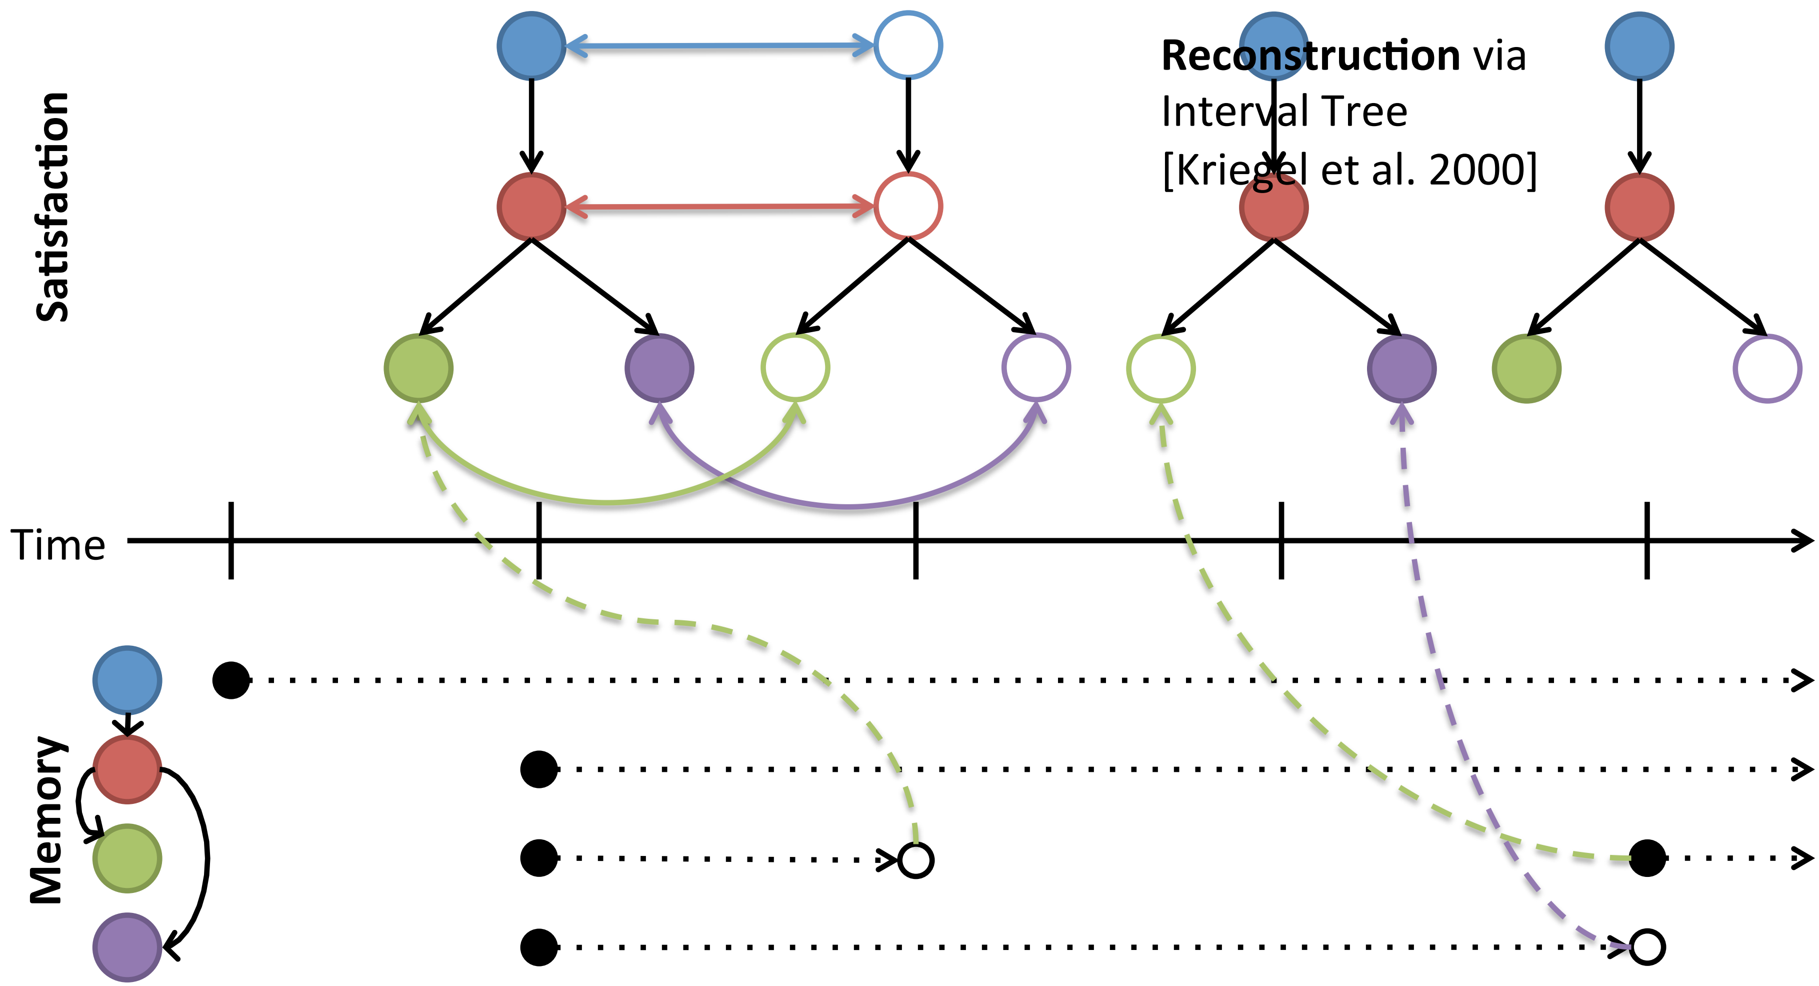

Efficient Retrieval

Surface Match Data Structures

1. **Interval Walk**
   - Maintain interval endpoint sorting via b+-trees
   - On cue, add leaf pointers to time keyed priority queue
     - Pop as necessary to process ● or ○

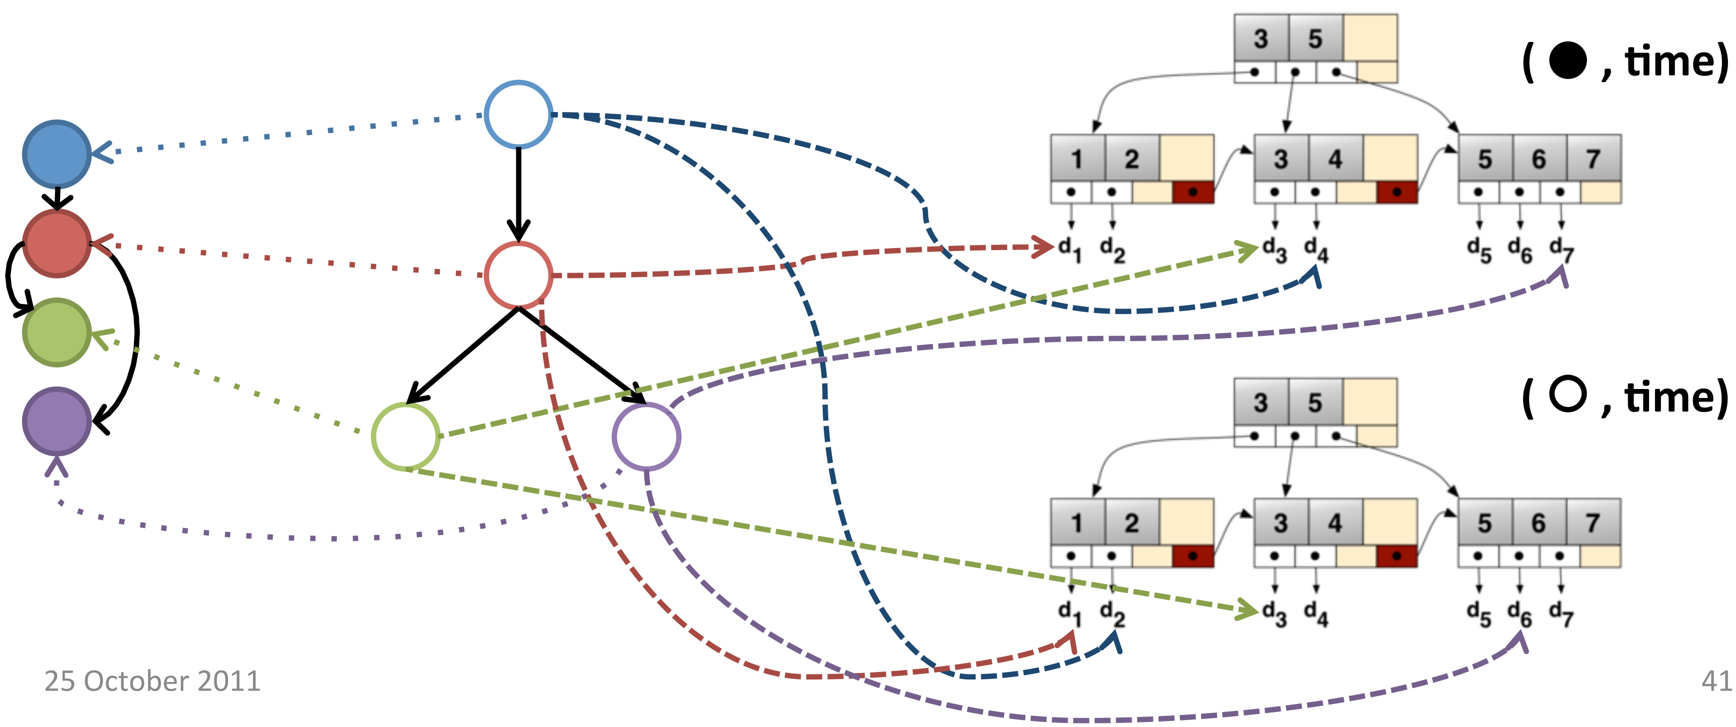

Efficient Retrieval

Surface Match Data Structures

2. Incremental Episode Scoring via DNF Graph

- Cue edges serve as minimal propagation directives
  - Maps to DNF SAT: sat(n) := sat(n) ^ sat(par(n))
- On ●/○, update clause(s), possibly recurse

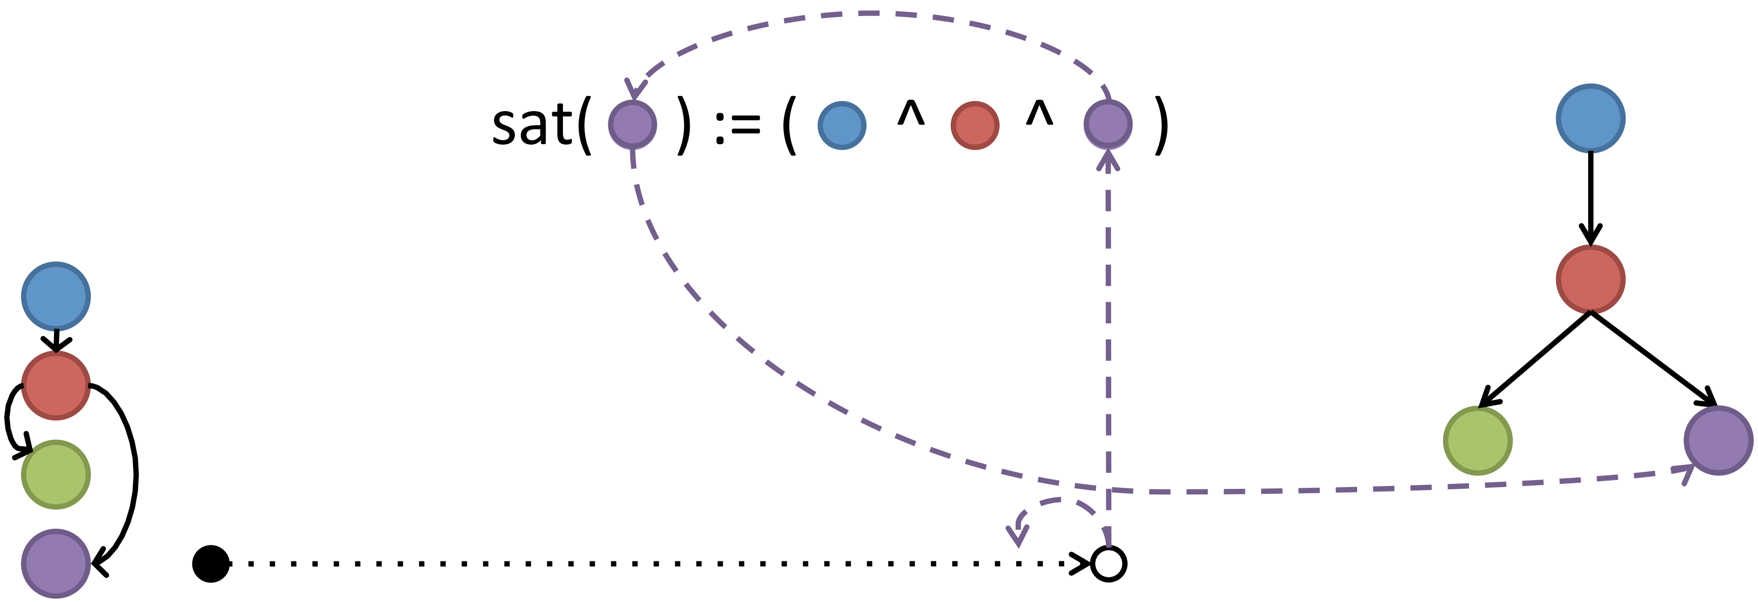

Efficient Retrieval

Scaling – [ICCBR ’09]

Interval Walk

O(|Δ| * Temporal Selectivity)

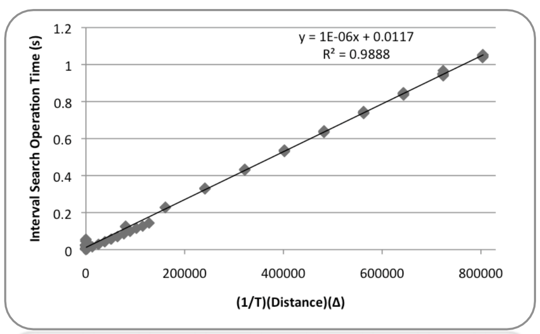

Incremental Episode Scoring

O(|Δ| * Structural Selectivity)

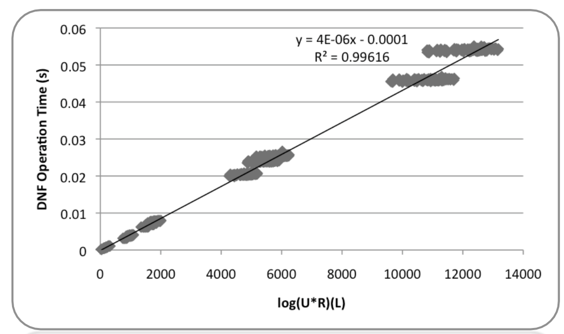

Episodic Memory

Empirical Evaluation

Games

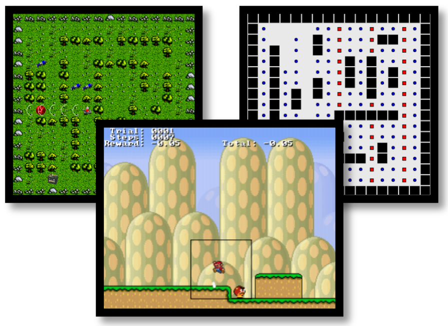

PDDL

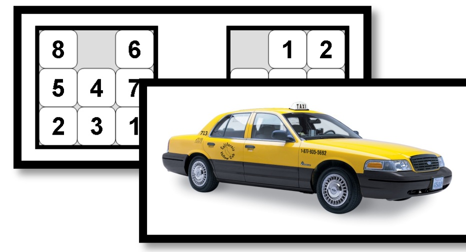

Robotics

WSD

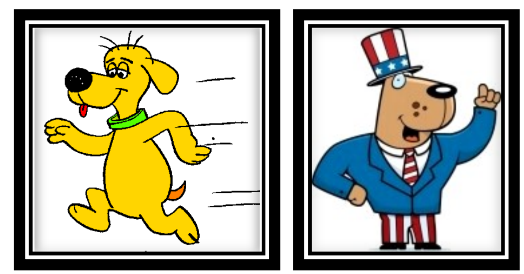

Useful: 7 general capabilities

Efficient: >100 cues, <50 msec.

Scalable: >48 hours, ~150 bytes – 2.5 kb/episode

Episodic Memory

Summary

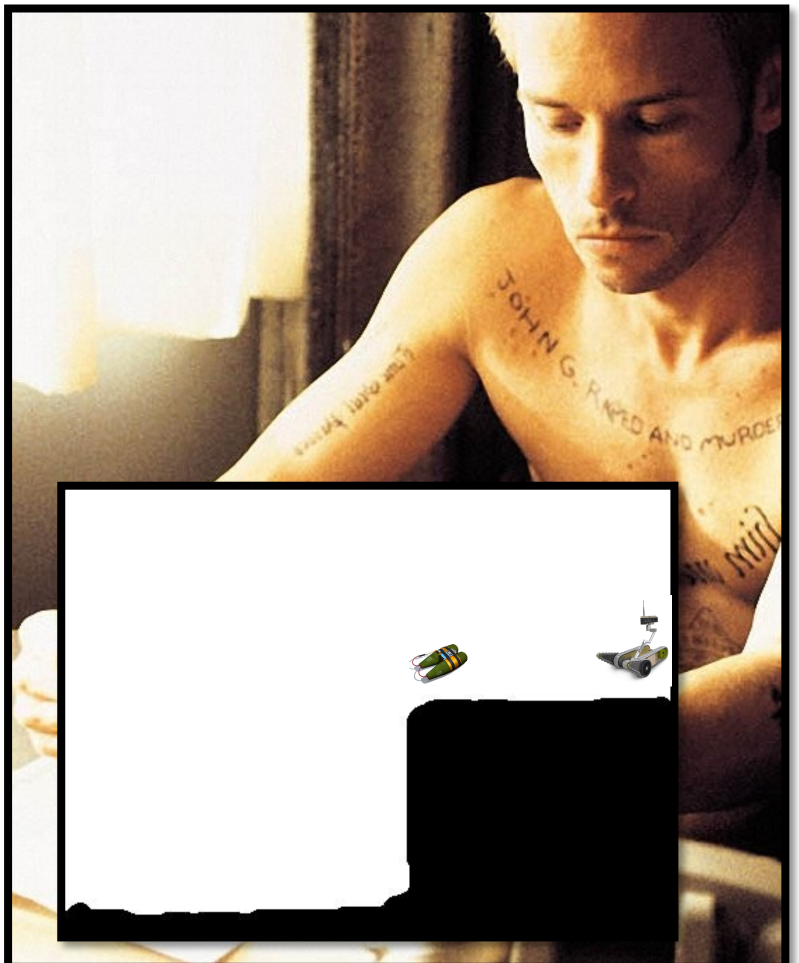

Reactivity

- Faster than reactivity requirement for many tasks/ queries in practice

Scalability

- Days of RT (millions of episodes)

Generality

- Useful in games, robotics, planning, linguistics

LTM for Intelligent Agents

Contributions

- Integrated effective and efficient semantic and episodic memories with Soar
- Novel methods that scale to large amounts of knowledge and long agent lifetimes
- Empirically evaluated on numerous tasks
  - Linguistics, robotics, games, planning
  - Desktop platforms, robotics hardware, (and mobile!)

LTM for Generally Intelligent Agents
Looking Forward

Future Directions

- Integrating context
- Automatic structure learning
- Reasoning with multiple sources of knowledge

LTMs make possible today…

- Robust decision-making
  - Improves with exploration and interaction
- Human-agent interaction
  - Complexity
  - Believability

Thank You :-)

Questions?

Nate Derbinsky

Soar Group, University of Michigan