Algorithms for Scaling in a General Episodic Memory

Nate Derbinsky, Justin Li, John E. Laird

University of Michigan

supported by

PROBLEM STATEMENT

Agents with episodic memory are…

- more capable in problem solving
- better able to account for human psychological phenomena
- more believable as virtual characters and long-term companions

How to effectively and efficiently support task-independent episodic memory for real-time agents for long periods of time?

AGENT INTEGRATION

The Soar cognitive architecture (Laird, 2012)

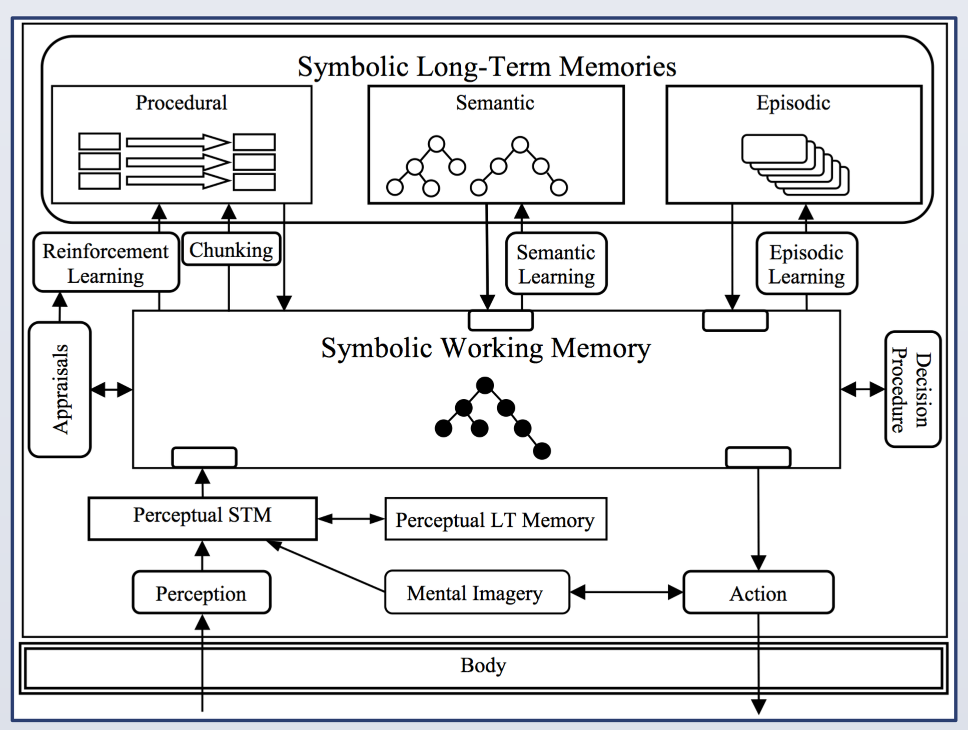

Episodic operations:

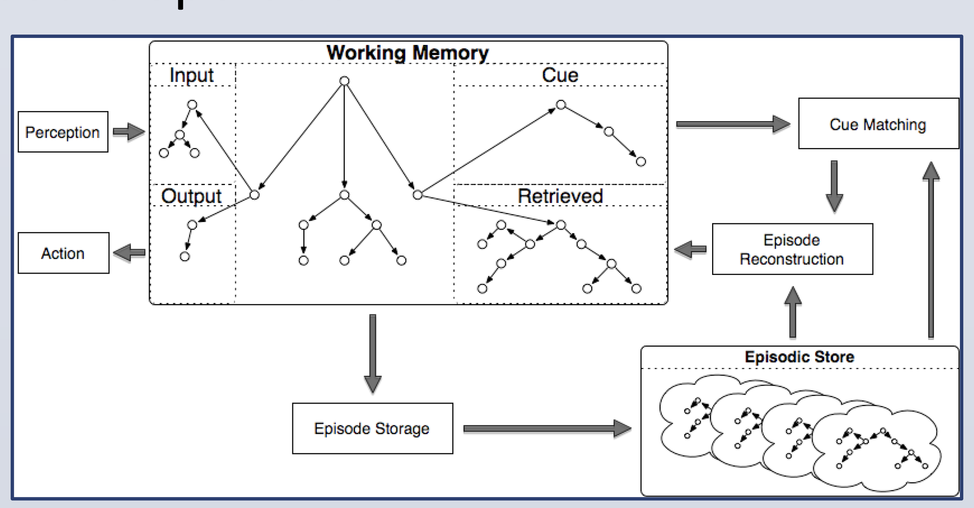

Representation

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

Storage

- Automatic, no dynamics (e.g. forgetting)

Cue Matching

- Cue: acyclic di-graph
- Find the most recent episode that shares the most leaf nodes in common with cue

CHALLENGES

- Arbitrary, dynamic state
- Scaling potential
  - agent state (10^3 nodes/edges)
  - agent lifetime (10^6 – 10^9 episodes/day)
- Cue-matching optimality
  - constrained sub-graph isomorphism
  - search: O(# episodes)

OPPORTUNITY: STRUCTURAL REGULARITY

Agents reuse representational structure

- |distinct structures| << |all structures|
- Exploited via alpha/beta-node sharing in rule-based systems (Doorenbos, 1995)

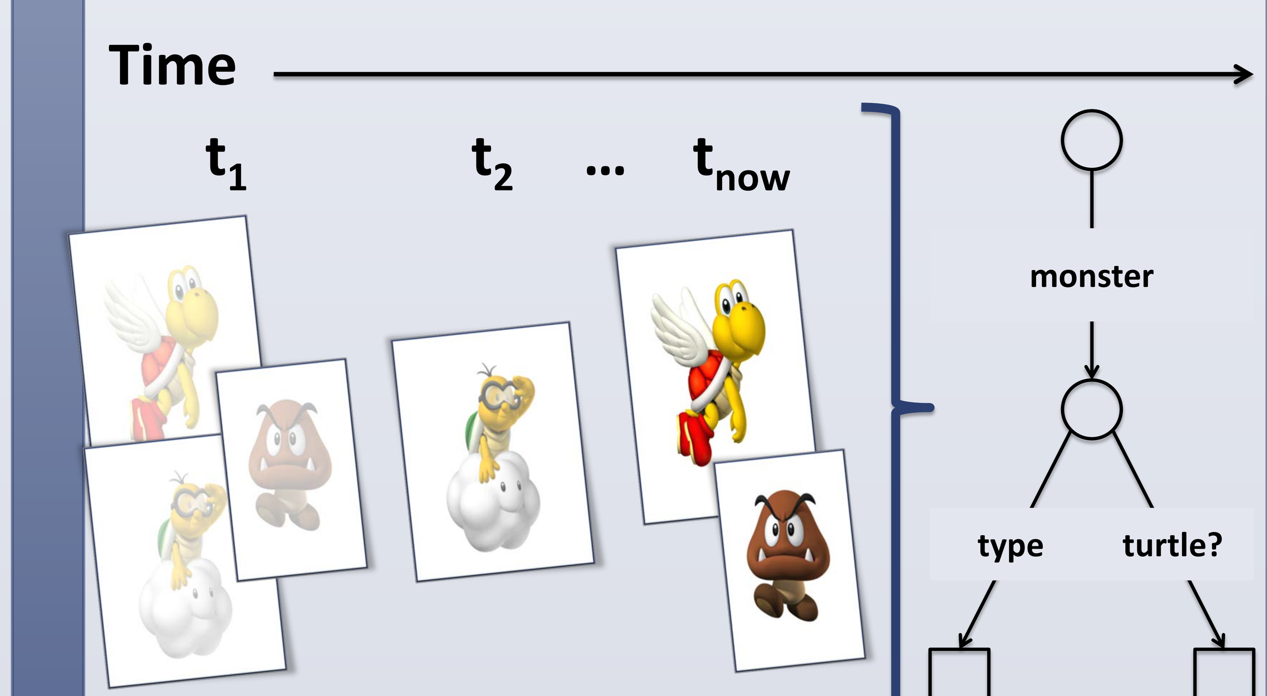

OPPORTUNITY: TEMPORAL CONTIGUITY

The world changes slowly

- |state changes| << |state|
- Exploited via state saving in rule-based systems (Forgy, 1982)

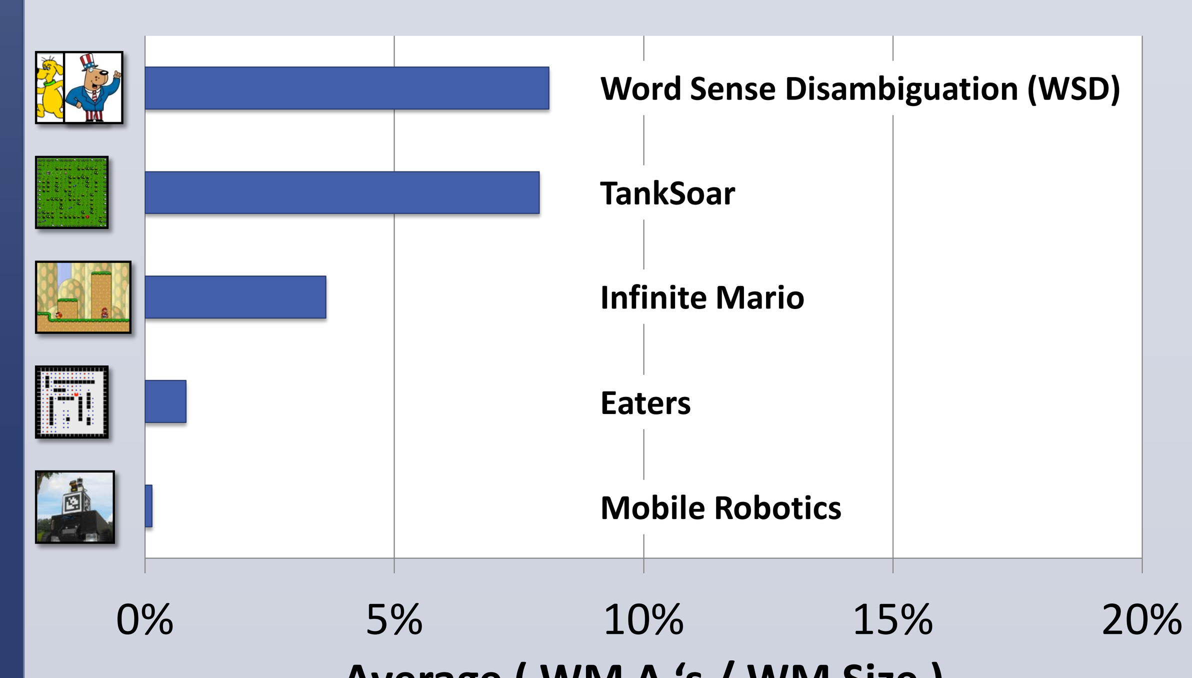

DYNAMIC GRAPH INDEX

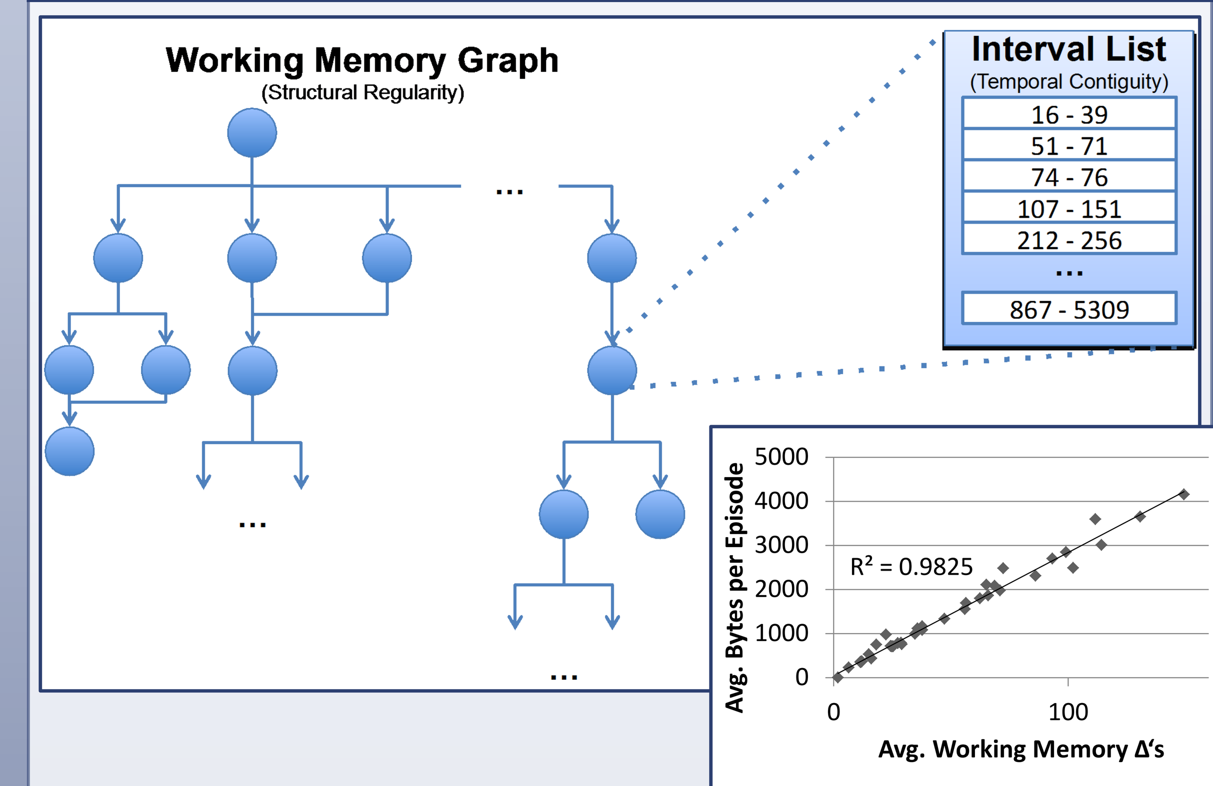

CUE-MATCHING ALGORITHM

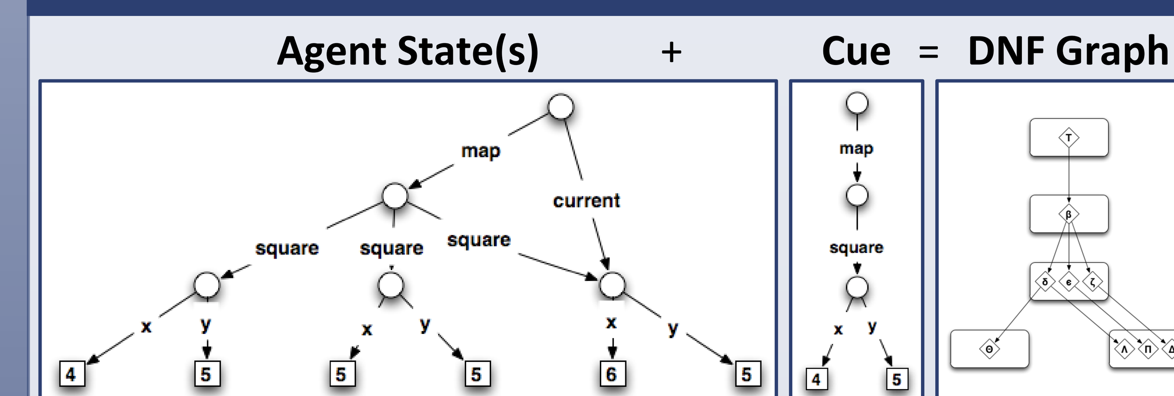

Incremental cue matching maps to incremental DNF SAT; DNF Graph is a novel discrimination network for scoring via Δ's

Temporal Interval Search via **Pr. Queue of B+ Tree Pointers**

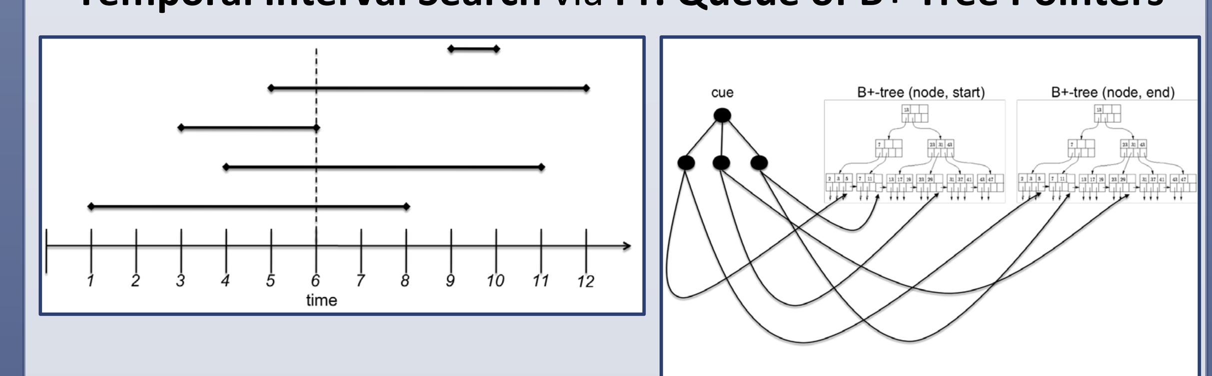

Optimizations: (a) 2-phase search (only graph match eps with all features independently), (b) only score eps with feature Δ's

PERFORMANCE CHARACTERIZATION

Cue matching scales as a function of general properties of tasks and cues:

Temporal Selectivity & **Co-Occurrence**

- Efficient *if* cue features occur rarely overall, but together within episodes
- O(Search Distance)

Structural Selectivity

- Efficient *if* cue features are distinctive within each individual episode
- O(Episode Hyper-edges)

EVALUATION

Setup

- 49 domains: WSD, 12 PDDL planning domains, mobile robotics, games
- 10^5 – 10^8 episodes (~ days of real time)
- >100 cues, 7 general capabilities
  - virtual sensing, action modeling, …

Summary of Results

- <50 msec. storage (except)
  - 0.18 – 4 kb/episode (days – months)
- <50 msec. cue matching for many cues
  - No cue-matching time growth
  - Validated performance models

Cue-Matching Time (msec) vs. 10^6 Episodes