

Soar-EpMem Tutorial

Soar Workshop 31 – Nate Derbinsky

While waiting...

1. Make sure you have internet access

2. Download Soar 9.3.1

soar.googlecode.com

3. Download Graphviz

www.graphviz.org

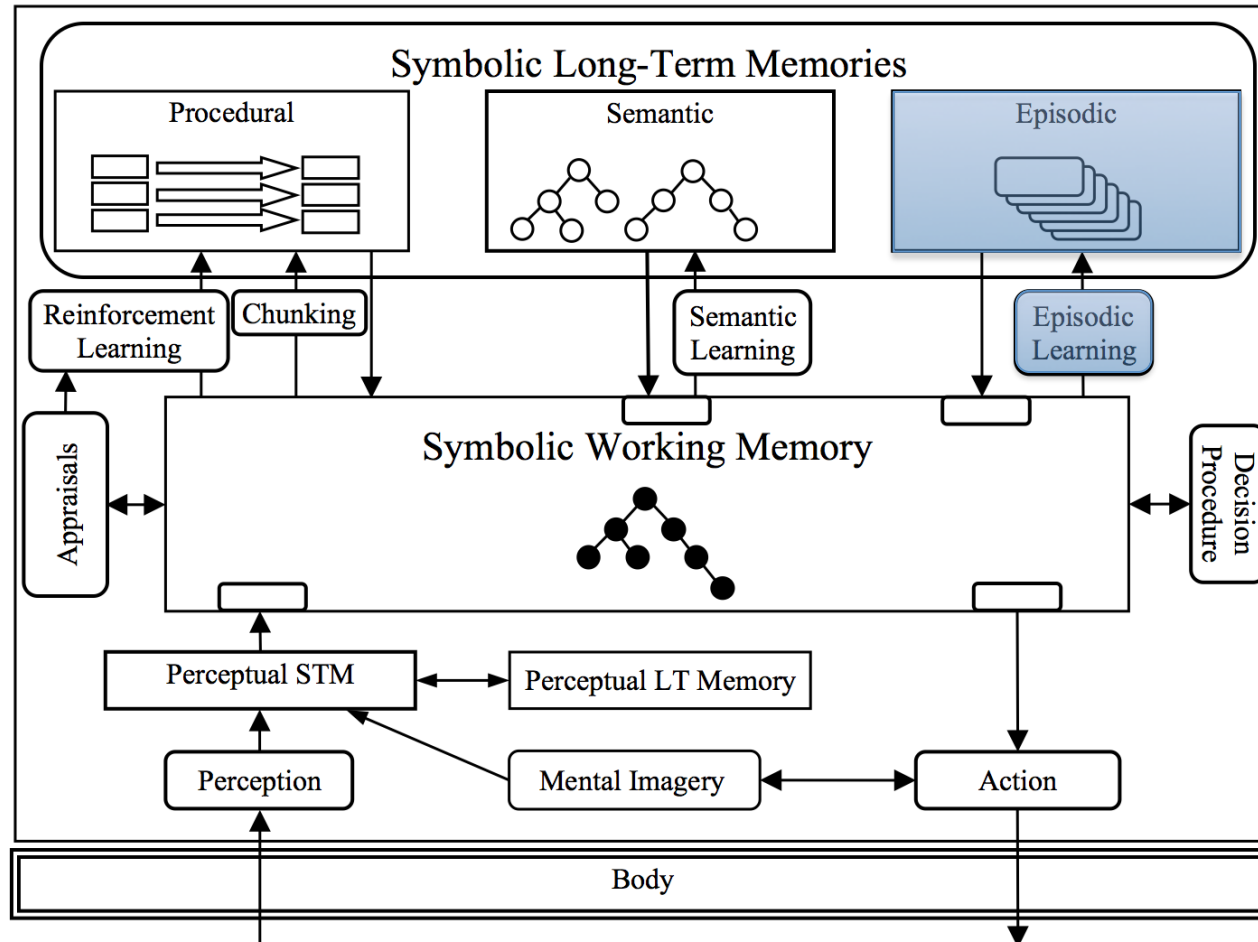
4. Download tutorial support files

www.eecs.umich.edu/~nlderbin/workshop31

Agenda

- Big picture
- Basic usage
- Demo task
- Additional resources

Soar 9

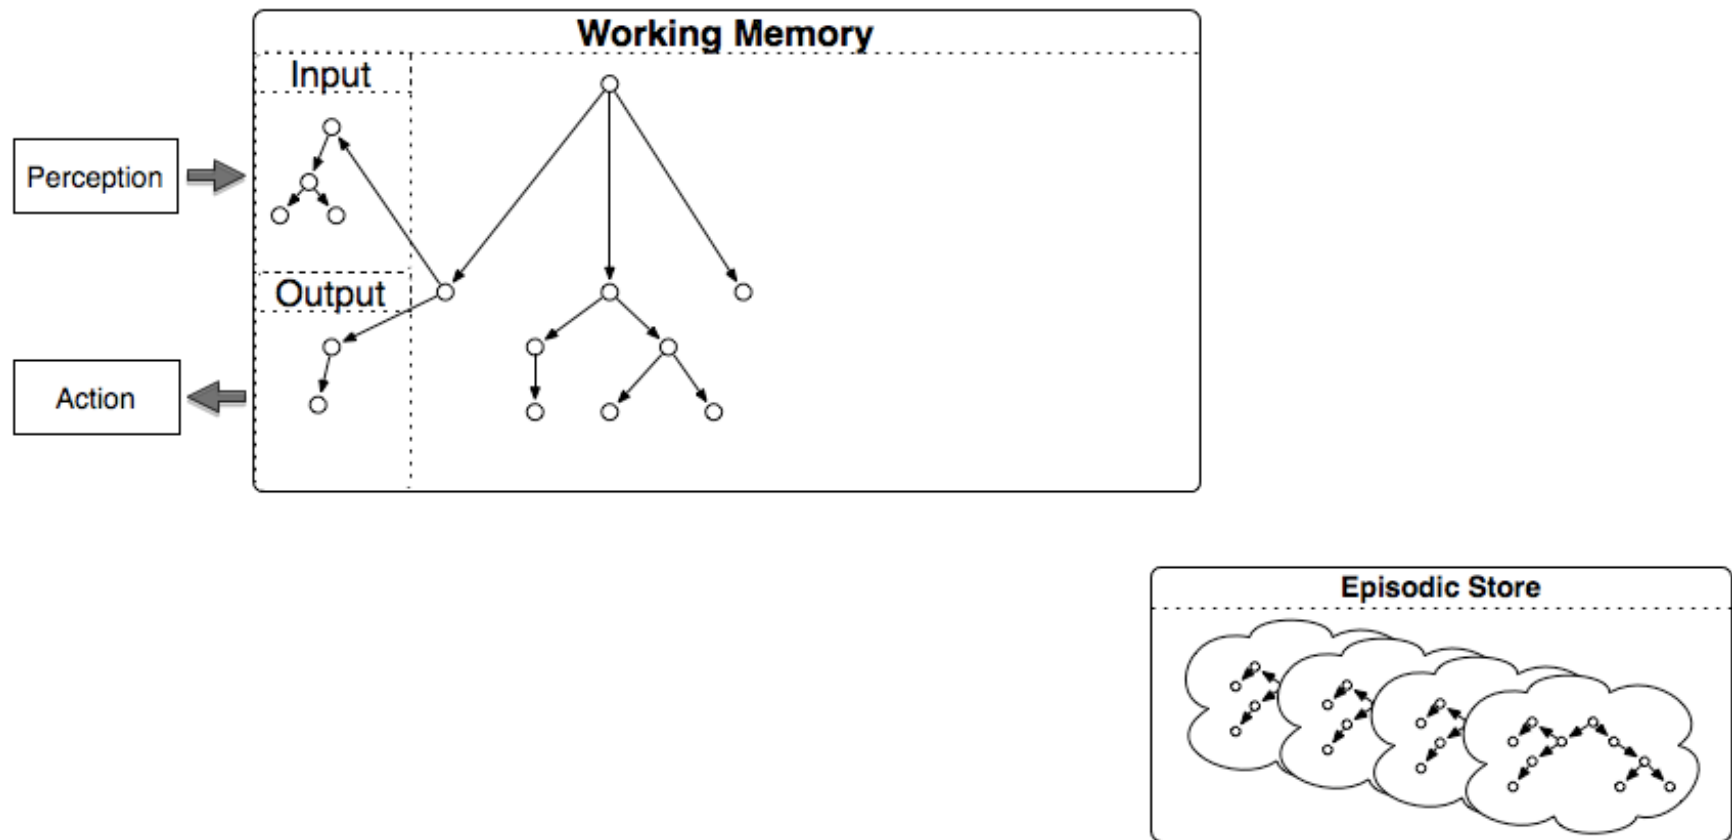


Episodic Memory: Big Picture

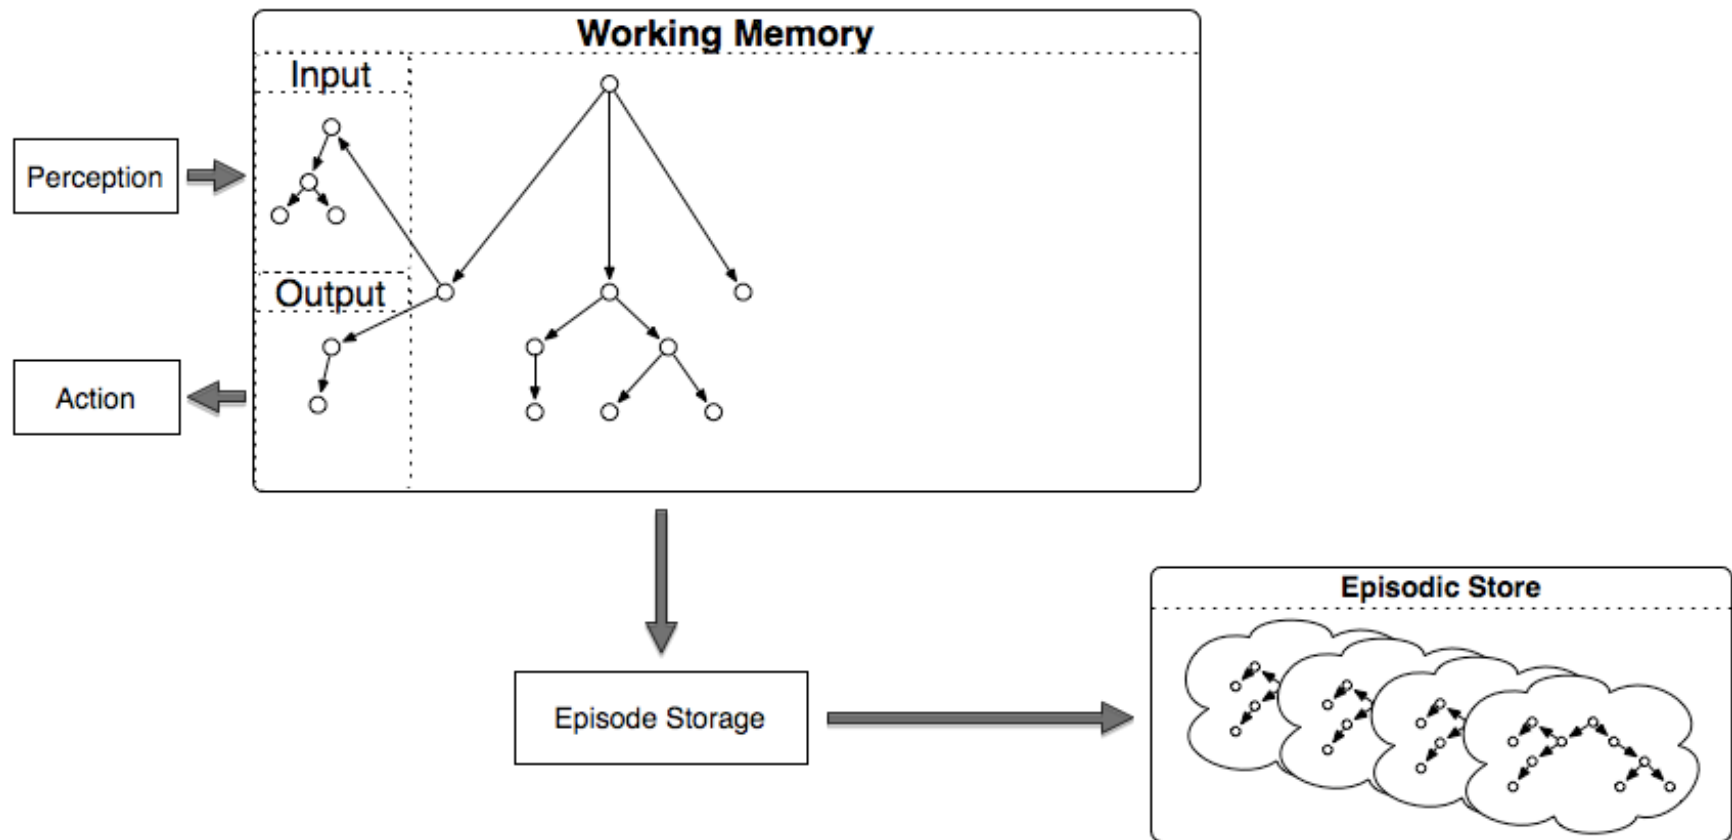
Episodic memory is a weak learning mechanism

- Automatically captures, stores, and temporally indexes agent state
- Supports content-addressable agent interface to autobiographical prior experience

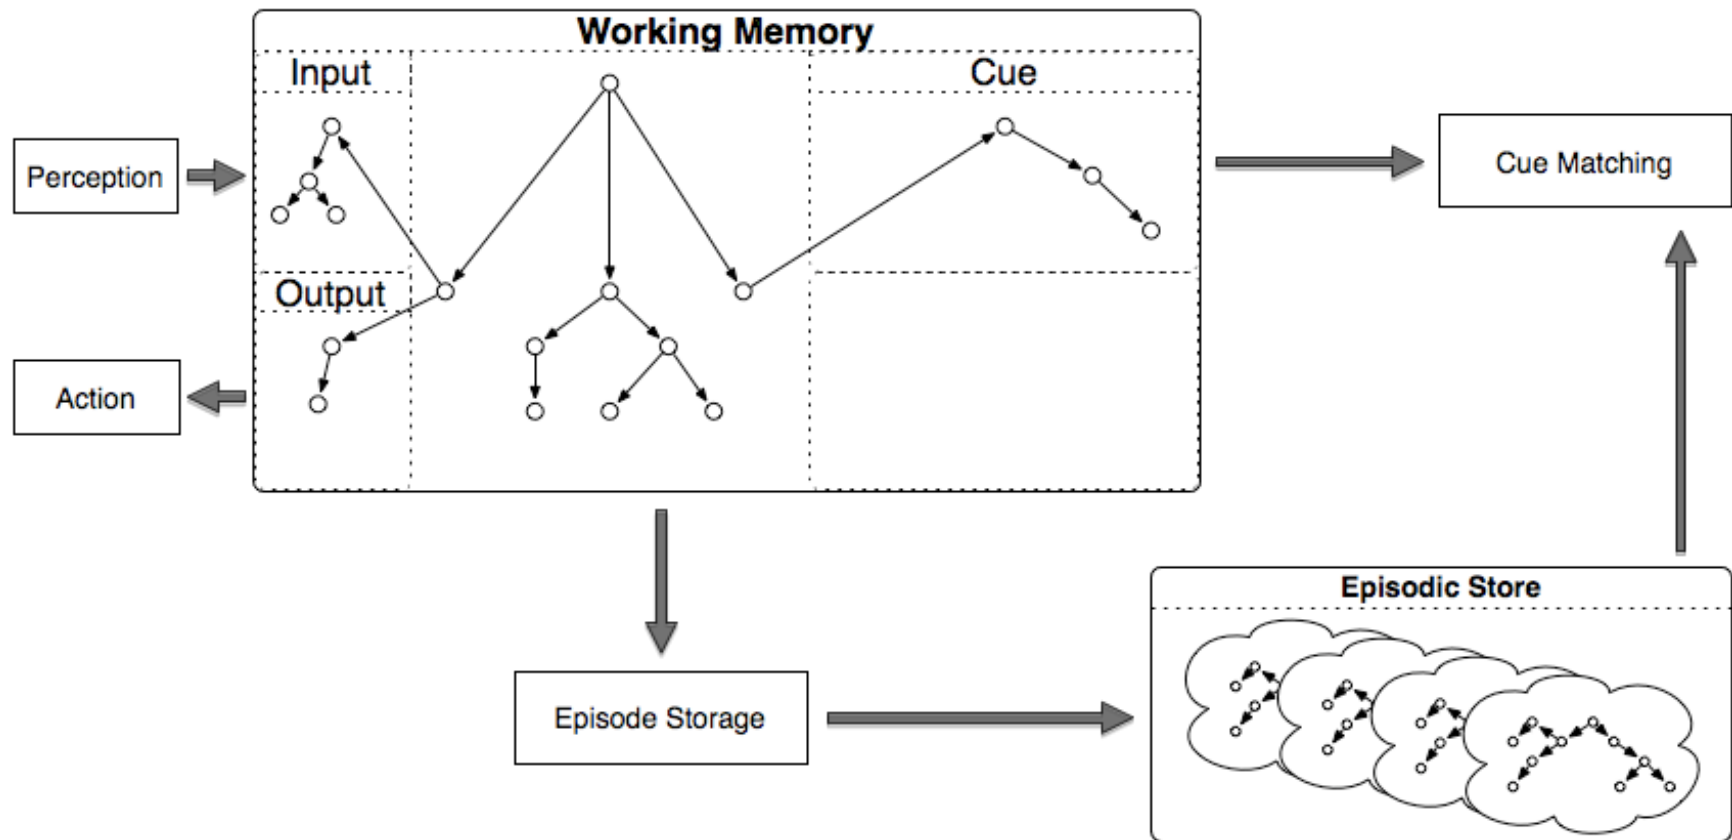
Architectural Integration



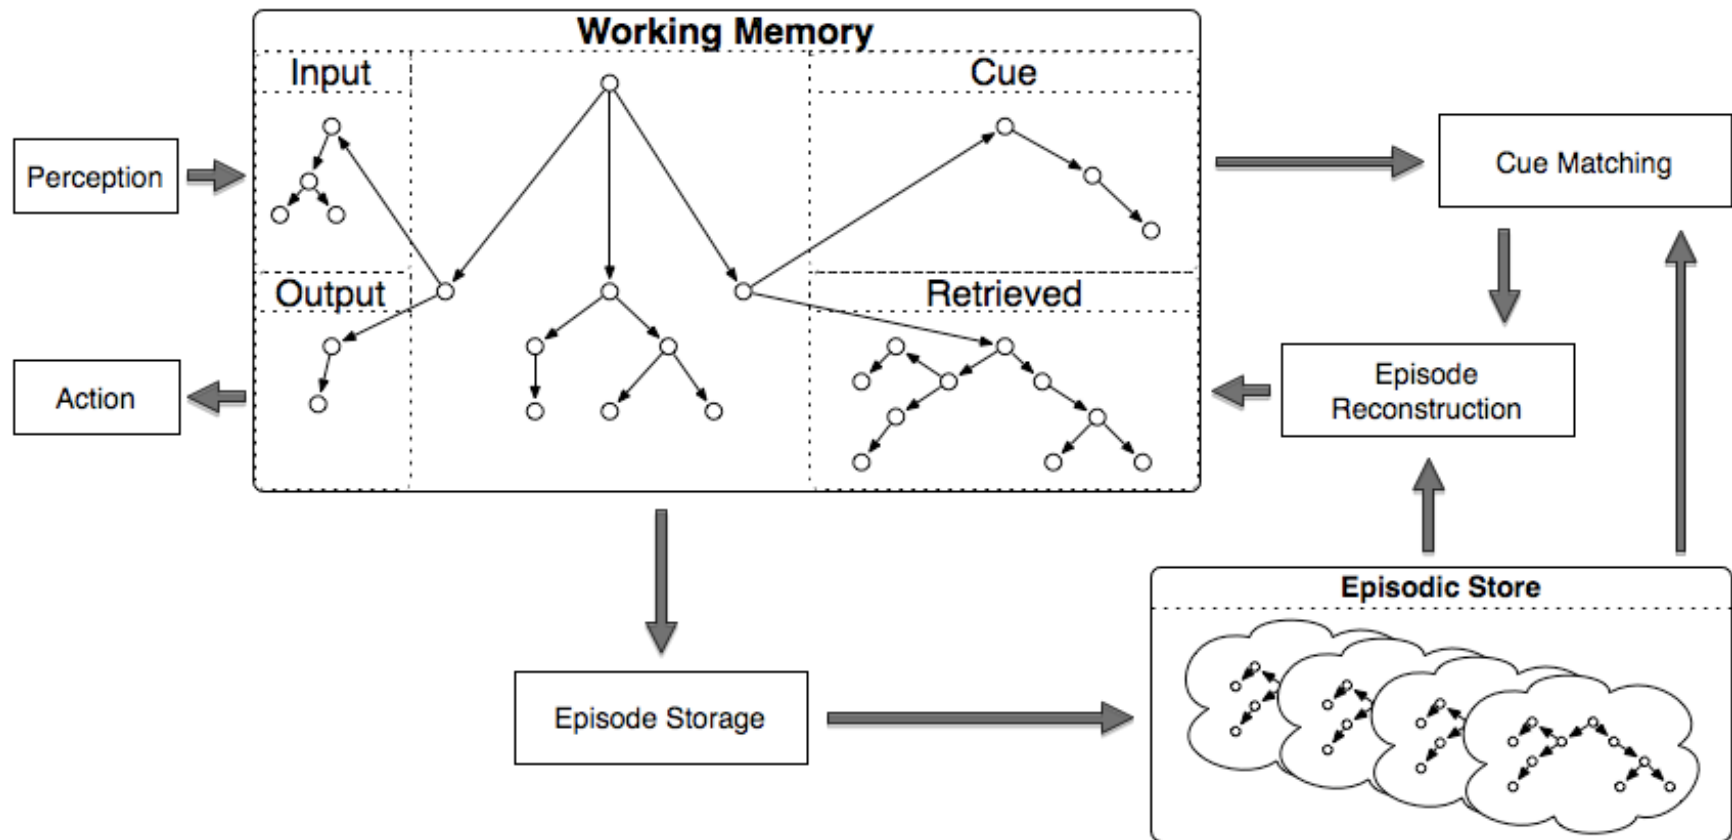
Architectural Integration



Architectural Integration



Architectural Integration



Basic Usage

- Working memory structure
- Episodic memory representation
- Controlling episodic memory
- Storing knowledge
- Retrieving knowledge

Working Memory Structure

Soar creates an `epmem` structure on each state

- Soar Java Debugger
 - `step 5`
 - `print --exact (* ^epmem *)`
 - `print s2`

Each `epmem` structure has specialized substructure

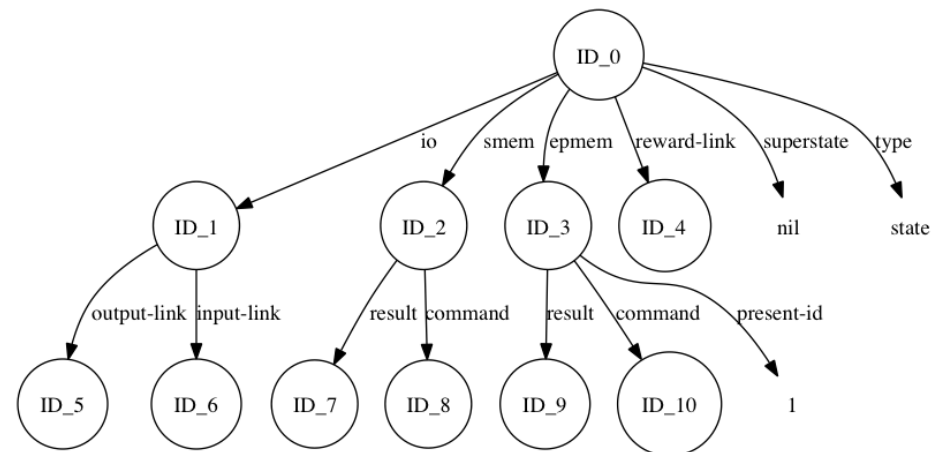
- `command`: agent-initiated actions
- `result`: architectural feedback
- `present-id`: current episode number (more later)

Episodic Memory Representation

Similar to working memory: symbolic triples

- Attributes cannot be identifiers (currently)
- Structures within an episode are connected; separate episodes are disconnected

```
(<id0> ^epmem <id3>  
  ^io <id1>  
  ^reward-link <id4>  
  ^smem <id2>  
  ^superstate nil  
  ^type state)  
(<id1> ^input-link <id6>  
  ^output-link <id5>)  
(<id2> ^command <id8> ^result <id7>)  
(<id3> ^command <id10> ^present-id 1  
  ^result <id9>)
```



Controlling Episodic Memory

Get/Set a parameter:

- `epmem [-g|--get] <name>`
- `epmem [-s|--set] <name> <value>`

EpMem is **disabled** by default. Try enabling it...


1. `epmem`
2. `epmem --set learning on`
3. `epmem`

Storing Knowledge

- Automatic storage requires EpMem to be **enabled** (see slide 12)
- Storage captures the top state of working memory
- Events trigger storage of new episodes
 - `epmem --set trigger << dc output >>`
 - `dc`: decision cycle
 - `output`: new identifier on output-link (default)
- Storage takes place at the end of a phase
 - `epmem --set phase << output selection >>`
 - `output` is default
 - `selection` may be useful for in-the-head agents

Automatic Storage: Example (1)

- Soar Java Debugger
 1. `epmem --set trigger dc`
 2. `epmem --set learning on`
 3. `watch --epmem`
 4. `run 5 -p`
 5. `epmem --print 1`
 6. `ctf ep.gv epmem --viz 1`
 7. `print e1`
 8. `epmem --stats`



Automatic Storage: Example (1)

Debrief

- What wasn't captured?
- Attributes can be excluded from encoding (and subsequent recursion)
 - `epmem --set exclusions <label>`
 - If <label> already excluded, now included
- Try previous example, add:
 - `epmem --set exclusions epmem`
 - `epmem --set exclusions smem`

Automatic Storage: Example (2)

- Eaters
 1. New agent (advanced-move.soar)
 - Spawn Debugger
 2. `epmem --set learning on`
 3. `epmem --stats`

Retrieving Knowledge

Cue-Based

Find the episode that best matches a cue and add it to working memory

Temporal Progression

Replace the currently retrieved episode with the next/previously encoded episode

Non-Cue-Based (not covered)

Add an episode to working memory from episode #

Common Constraints:

- Requires that EpMem is enabled (slide 12)
- Only one per state per decision
- Processed during phase (slide 13)
- Only re-processed if WM changes to commands
- Meta-data (status, etc) automatically cleaned by the architecture

Cue-Based Retrieval: Syntax

```
( <epmem> ^command <cmd> )  
( <cmd> ^query <q>  
      ^neg-query <nq> )
```

- The neg-query is optional
- Cues must be acyclic
- The <q> and <nq> identifiers form the roots of episode sub-graph cues
 - query represents desired structures
 - neg-query represents undesired structures

Cue-Based Retrieval: Cue Semantics

Values of cue WMEs are interpreted by type

- Constant: exact match
- Long-Term ID: exact match, stop
- Short-Term ID: Wildcard (but must be identifier)

Cue-Based Retrieval: Episode Scoring

- **Leaf WME**, either...
 - Cue WME whose value is a constant OR
 - Cue WME whose value is an identifier and that identifier has no augmentations
- A leaf wme is *satisfied* (w.r.t. an episode) if...
 - The episode contains that WME AND
 - The episode contains a path from root to that WME
- Episode scoring
 - $(\text{balance})(\text{cardinality}) + (1 - \text{balance})(\text{activation})$
 - balance: parameter=[0,1], default=1
 - cardinality: # satisfied leaf WMEs
 - activation: \sum satisfied leaf WME activation (see Manual)
 - cardinality/activation negated for neg-query

Cue-Based Retrieval: Cue Matching

Graph matching

`epmem --set graph-match << on off >>`

- on by default

Candidate episode

Defined as satisfying at least one leaf WME

Cue matching will return the most recent graph-matched episode, or the most recent non-graph-matched candidate episode with the maximal episode score

Cue-Based Retrieval: Result

Augmentation	Meaning
<code>^retrieved <retrieval-root></code>	Root of the retrieved memory
<code>^<< success failure >> <query> <neg-query></code>	Query status
<code>^match-score #</code>	Float, episode score (slide 19)
<code>^cue-size #</code>	Integer, number of leaf WMEs
<code>^normalized-match-score #</code>	match-score/cue-size
<code>^match-cardinality #</code>	Integer, number of satisfied leaf WMEs ($ \text{query} - \text{neg-query} $)
<code>^memory-id #</code>	Integer, episode # retrieved
<code>^present-id #</code>	Integer, current episode #
<code>^graph-match << 0 1 >></code>	Integer, 1 if graph match succeeded
<code>^mapping <mapping-root></code>	A mapping from the cue to episode


Cue-Based Retrieval: Example (1)

query-1.soar (find superstate nil)

- Soar Java Debugger
 1. epmem --set trigger dc
 2. epmem --set learning on
 3. watch --epmem
 4. sp {query1
 (state <s> ^superstate nil
 ^epmem.command <cmd>)
 -->
 (<cmd> ^query.superstate nil)}
 5. run 5 -p
 6. print -d 10 e1

Cue-Based Retrieval: Example (1)

Result



Pop Quiz: how did I make this?

Cue-Based Retrieval: Example (1)

Trace

```
CONSIDERING EPISODE (time, cardinality, score): (1, 1, 1.000000)  
NEW KING (perfect, graph-match): (true, true)
```

Cue-Based Retrieval: Example (2)

query-2.soar (find when max is defined and first is true)

- Soar Java Debugger
 1. `source query-2.soar`
 2. `run`

```
CONSIDERING EPISODE (time, cardinality, score): (11, 1, 1.000000)
NEW KING (perfect, graph-match): (false, false)
CONSIDERING EPISODE (time, cardinality, score): (1, 2, 2.000000)
NEW KING (perfect, graph-match): (true, true)
```

Cue-Based Retrieval: Example (2b)

query-2b.soar (longer version of 2)

- Soar Java Debugger
 1. `source query-2b.soar`
 2. `run`
 3. `epmem --stats`
 4. `epmem --timers`

```
CONSIDERING EPISODE (time, cardinality, score): (10001, 1, 1.000000)
NEW KING (perfect, graph-match): (false, false)
CONSIDERING EPISODE (time, cardinality, score): (1, 2, 2.000000)
NEW KING (perfect, graph-match): (true, true)
```

Cue-Based Retrieval: Example (2b)

Stats

Time: 10002
SQLite Version: 3.6.12
Memory Usage: 2957920
Memory Highwater: 2988656
Queries: 1
Nexts: 0
Prevs: 0
Last Retrieval WMEs: 14
Last Query Positive: 2
Last Query Negative: 0
Last Query Retrieved: 1
Last Query Cardinality: 2
Last Query Literals: 2

Cue-Based Retrieval: Example (2b)

Timers

```
_total: 0.147855
epmem_api: 1.9e-05
epmem_hash: 0.010883
epmem_init: 0.000355
epmem_ncb_retrieval: 2.6e-05
epmem_next: 0
epmem_prev: 0
epmem_query: 4.2e-05
epmem_storage: 0.134101
epmem_trigger: 6e-06
epmem_wm_phase: 1e-06
ncb_edge: 5e-06
ncb_edge_rit: 2e-06
ncb_node: 5e-06
ncb_node_rit: 0
query_dnf: 3.6e-05
query_graph_match: 0
query_neg_end_ep: 0
query_neg_end_now: 0
query_neg_end_point: 0
query_neg_start_ep: 0
query_neg_start_now: 0
query_neg_start_point: 0
query_pos_end_ep: 0
query_pos_end_now: 0
query_pos_end_point: 0
query_pos_start_ep: 0
query_pos_start_now: 0
query_pos_start_point: 0
```

Cue-Based Retrieval: Example (3)

query-3.soar (find a number that is even and odd)

- Soar Java Debugger
 - source query-3.soar
 - run
 - epmem --stats
 - epmem --timers

```
CONSIDERING EPISODE (time, cardinality, score): (11, 1, 1.000000)
NEW KING (perfect, graph-match): (false, false)
CONSIDERING EPISODE (time, cardinality, score): (9, 1, 1.000000)
CONSIDERING EPISODE (time, cardinality, score): (8, 1, 1.000000)
CONSIDERING EPISODE (time, cardinality, score): (7, 1, 1.000000)
CONSIDERING EPISODE (time, cardinality, score): (6, 1, 1.000000)
CONSIDERING EPISODE (time, cardinality, score): (5, 1, 1.000000)
CONSIDERING EPISODE (time, cardinality, score): (4, 1, 1.000000)
CONSIDERING EPISODE (time, cardinality, score): (3, 1, 1.000000)
CONSIDERING EPISODE (time, cardinality, score): (2, 1, 1.000000)
CONSIDERING EPISODE (time, cardinality, score): (1, 1, 1.000000)
```

Cue-Based Retrieval

Optional Modifiers

(`<cmd> ^before time-id`)

(`<cmd> ^after time-id`)

(`<cmd> ^prohibit time-id1 time-id2 ...`)

Task. Modify `query-3.soar` to find an episode with an even count, before episode 10, that is not a power of 2.

1. Using a neg-query (`query-3b.soar`)
2. Using modifiers (`query-3c.soar`)

Cue-Based Retrieval

Debrief: neg-query vs. modifiers

neg-query

```
CONSIDERING EPISODE (time, cardinality, score): (11, 0, 0.000000)
NEW KING (perfect, graph-match): (false, false)
CONSIDERING EPISODE (time, cardinality, score): (8, 0, 0.000000)
CONSIDERING EPISODE (time, cardinality, score): (6, 1, 1.000000)
NEW KING (perfect, graph-match): (true, true)
```

modifiers

```
CONSIDERING EPISODE (time, cardinality, score): (6, 1, 1.000000)
NEW KING (perfect, graph-match): (true, true)
```

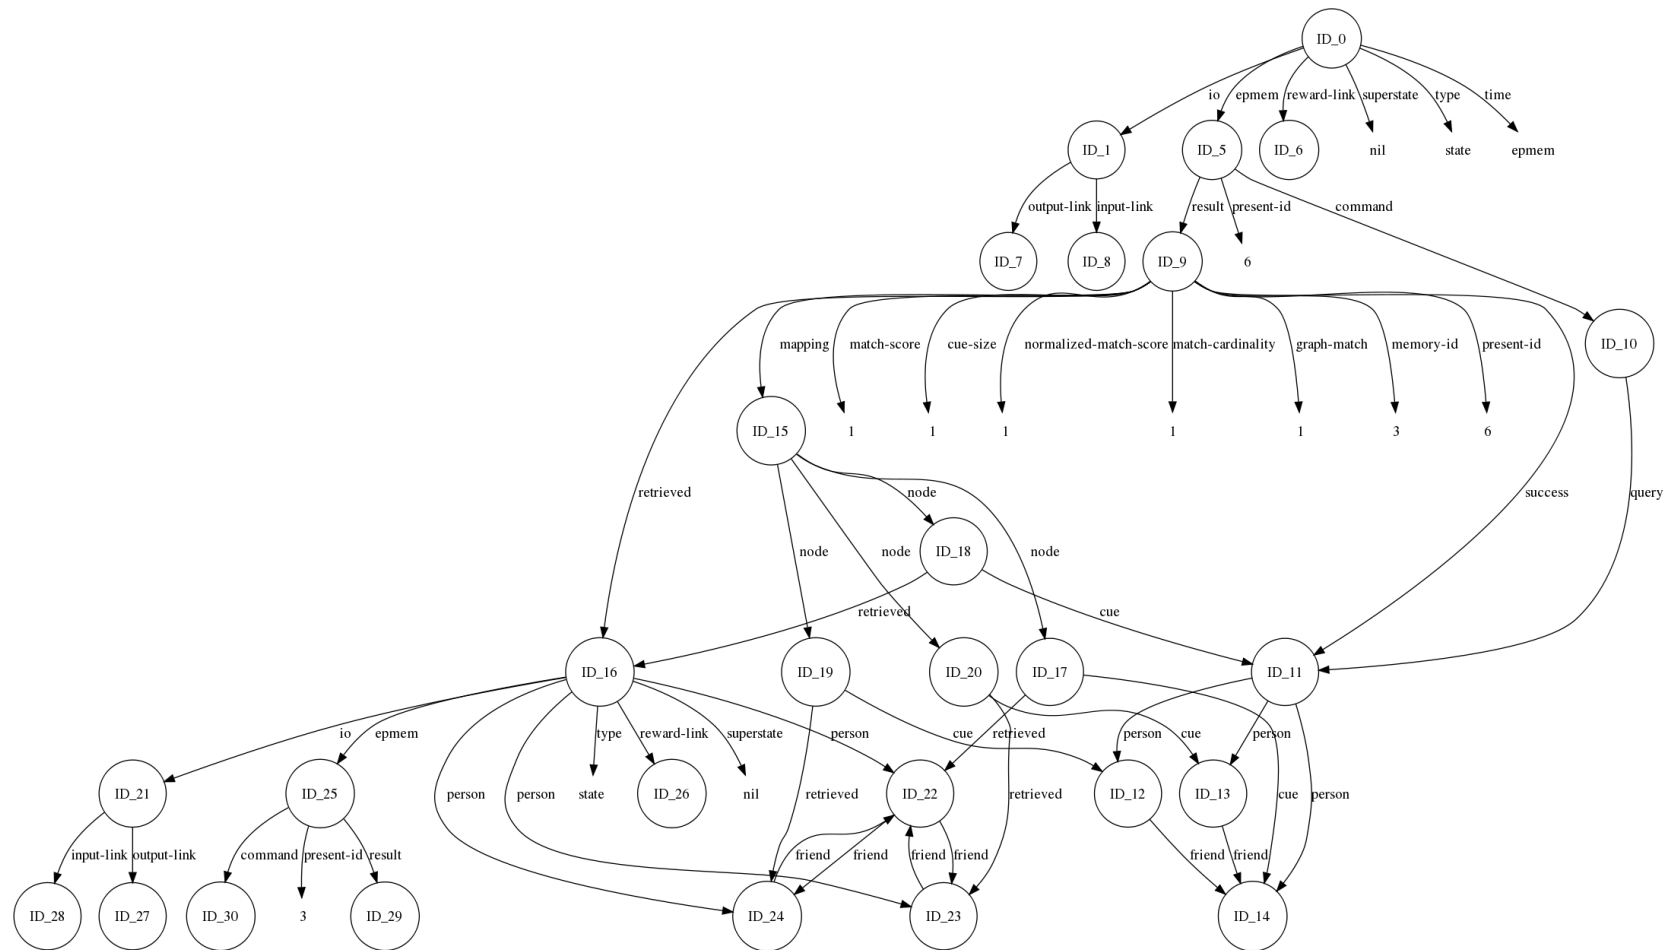

Cue-Based Retrieval: Example (4)

friends.soar

- Soar Java Debugger
 - source friends.soar
 - run

```
CONSIDERING EPISODE (time, cardinality, score): (4, 1, 1.000000)
NEW KING (perfect, graph-match): (true, false)
CONSIDERING EPISODE (time, cardinality, score): (3, 1, 1.000000)
NEW KING (perfect, graph-match): (true, true)
```

Cue-Based Retrieval: Example (4)



Temporal Progression

(<cmd> ^next <new-id>)

(<cmd> ^previous <new-id>)



Retrieves the next/previous episode, temporally, with respect to the last that was retrieved

Task. Modify query-3c.soar (slide 31) to find the episode *after* (query-3c-after.soar).

Demo: Virtual Sensing Task

demo.soar

1. Produce a random number in WM
EpMem automatically records this episode
(demo-start.soar)
2. Remove the number from WM
Write to the trace (for later verification)
3. Query episodic memory
When did I last see a random number?
4. Reason about the retrieved episode
Extract and print the number



Additional Resources

- Documentation
- Demo agents
- Readings

Documentation

Manual

`share/soar/Documentation`

Additional Topics

- Absolute non-cue-based retrievals
- Disk-based databases
- Performance
- Usage: commands, parameters, statistics, etc.
- ...

Demo Agents

`share/soar/Demos`

- kb
 - Demonstrates and unit tests the EpMem API
- count-epmem
 - Counting agent: stores then retrieves
 - Used for performance evaluation

Select Readings

code.google.com/p/soar/wiki/Publications

2004

- A Cognitive Model of Episodic Memory Integrated with a General Cognitive Architecture
Andrew M. Nuxoll, John E. Laird (ICCM)

2007

- Extending Cognitive Architecture with Episodic Memory
Andrew M. Nuxoll, John E. Laird (AAAI)
- Enhancing Intelligent Agents with Episodic Memory
Andrew M. Nuxoll (Dissertation)

2009

- Efficiently Implementing Episodic Memory
Nate Derbinsky, John E. Laird (ICCBR)
- A Year of Episodic Memory
John E. Laird, Nate Derbinsky (IJCAI Workshop)
- Learning to Use Episodic Memory
Nicholas A. Gorski, John E. Laird (ICCM)

2010

- Extending Soar with Dissociated Symbolic Memories
Nate Derbinsky, John E. Laird (AISB)
- Instance-Based Online Learning of Deterministic Relational Action Models
Joseph Xu, John E. Laird (AAAI)