A Year of Episodic Memory

John E. Laird and Nate Derbinsky
University of Michigan

Outline

- Episodic Memory
  - What is It
  - Why is It Important
- The Challenge
  - Statement and Focus
  - Bounds and Requirements
- Conclusions
  - Future Work
  - Review and Discussion

Outline

- **Episodic Memory**
  - What is It
  - Why is It Important
- The Challenge
  - Statement and Focus
  - Bounds and Requirements
- Conclusions
  - Future Work
  - Review and Discussion

What is Episodic Memory?

- Long-term, contextualized store of specific events
  - Tulving, E.: Elements of Episodic Memory (1983)

- Functionally
  - Architectural
  - Automatic
  - Autonoetic
  - Temporally indexed

Architectural Integration

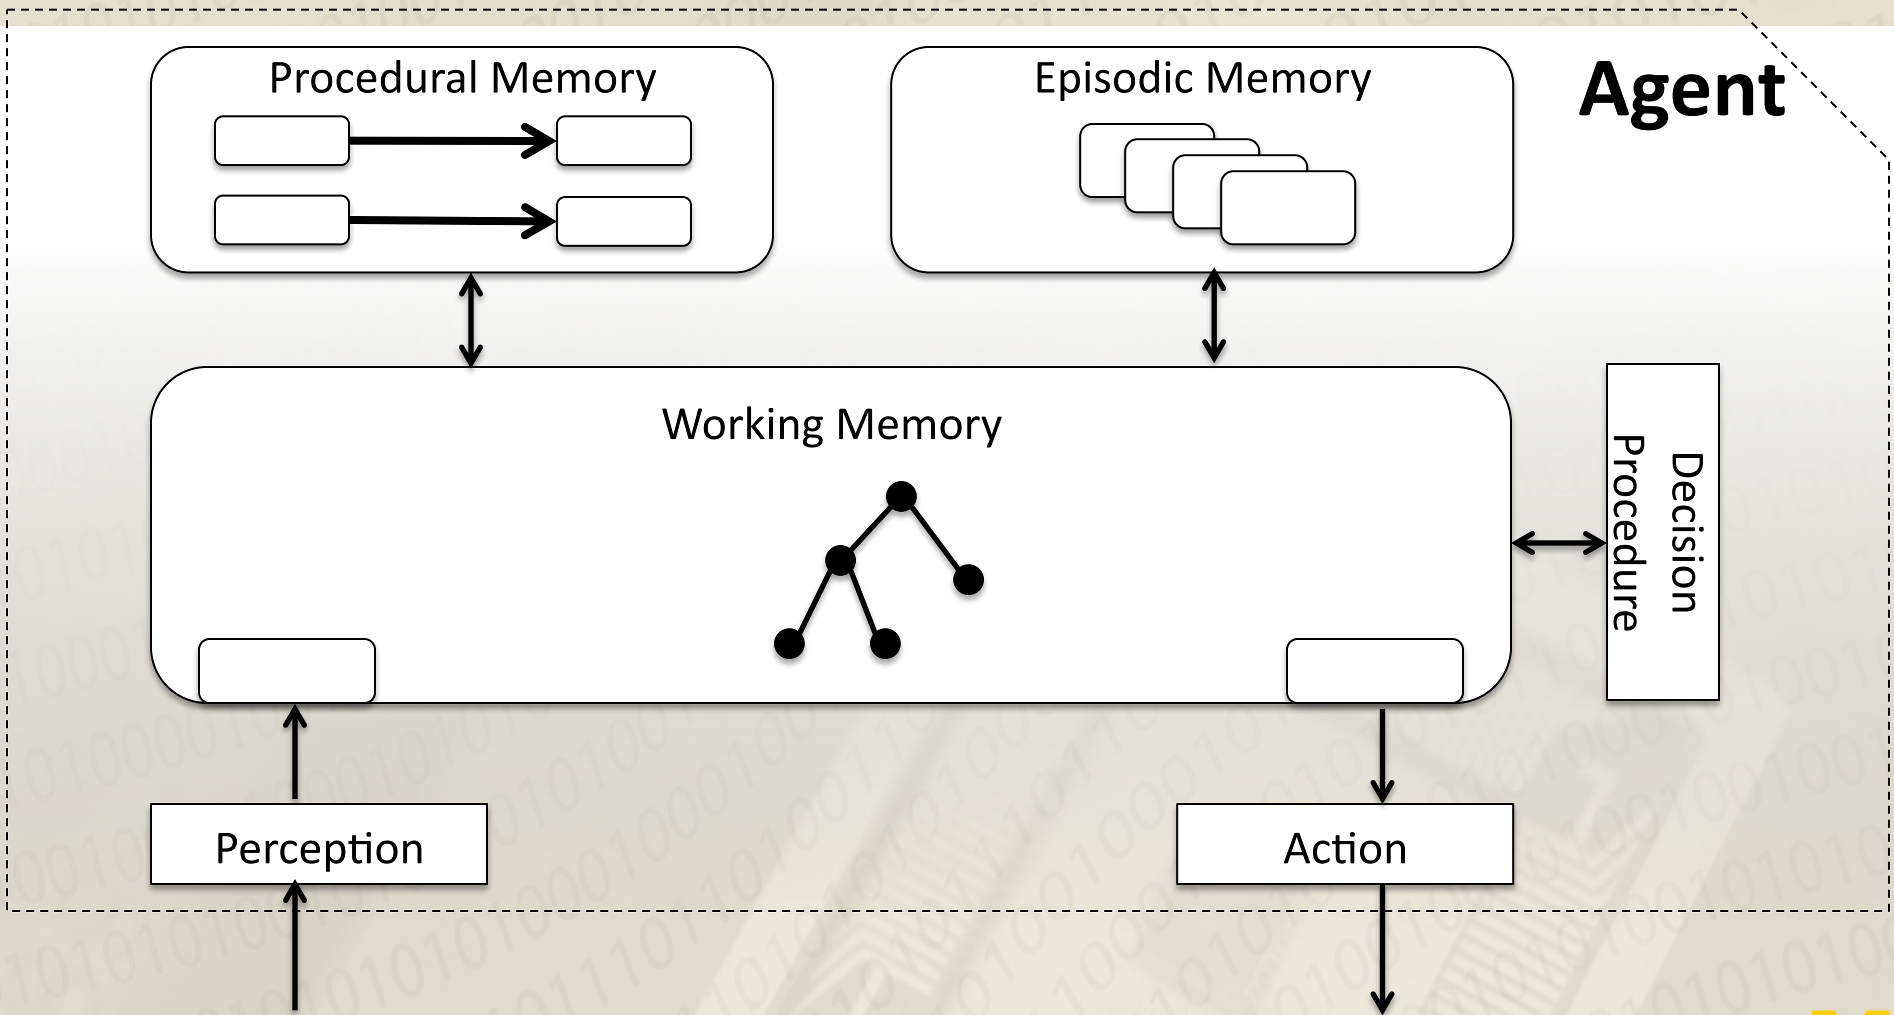

Episodic Operations

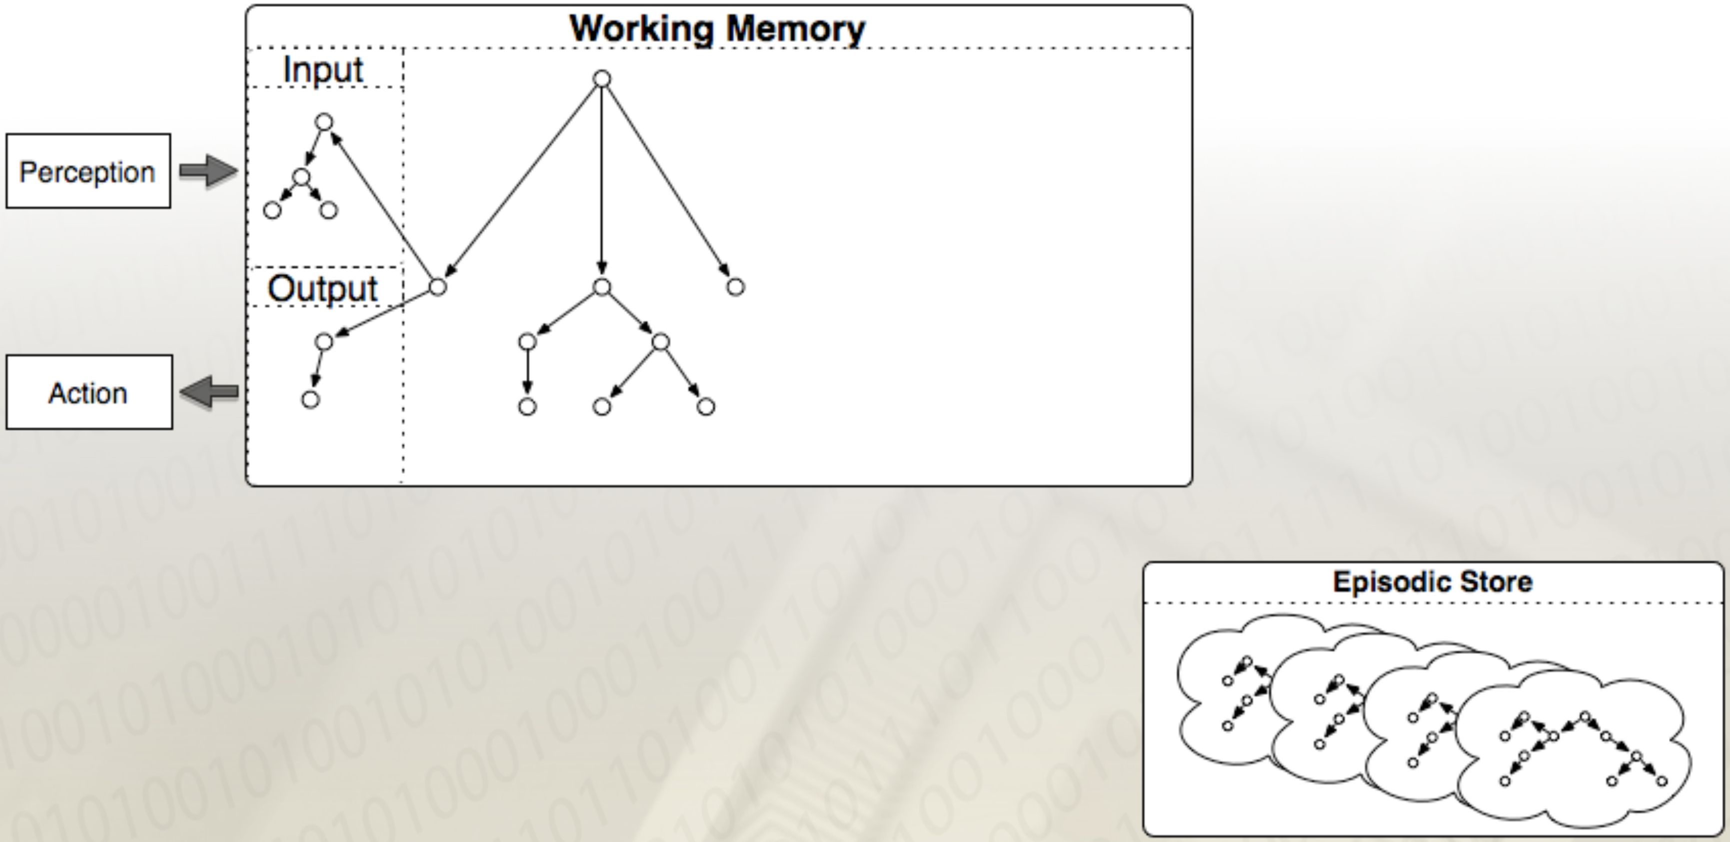

Episodic Operations

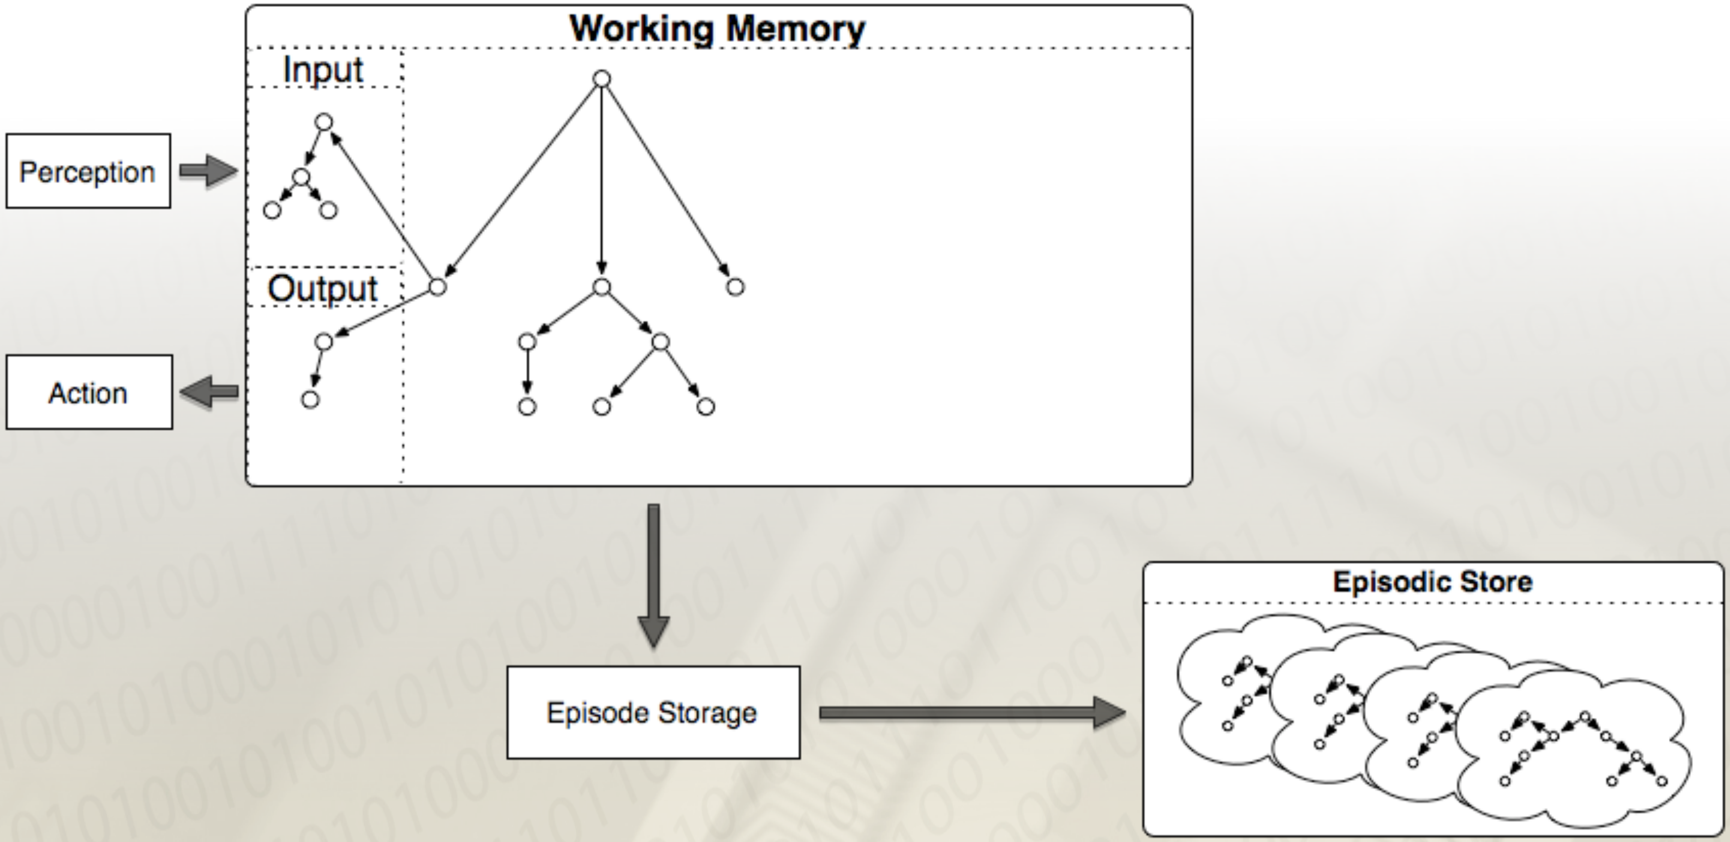

Episodic Operations

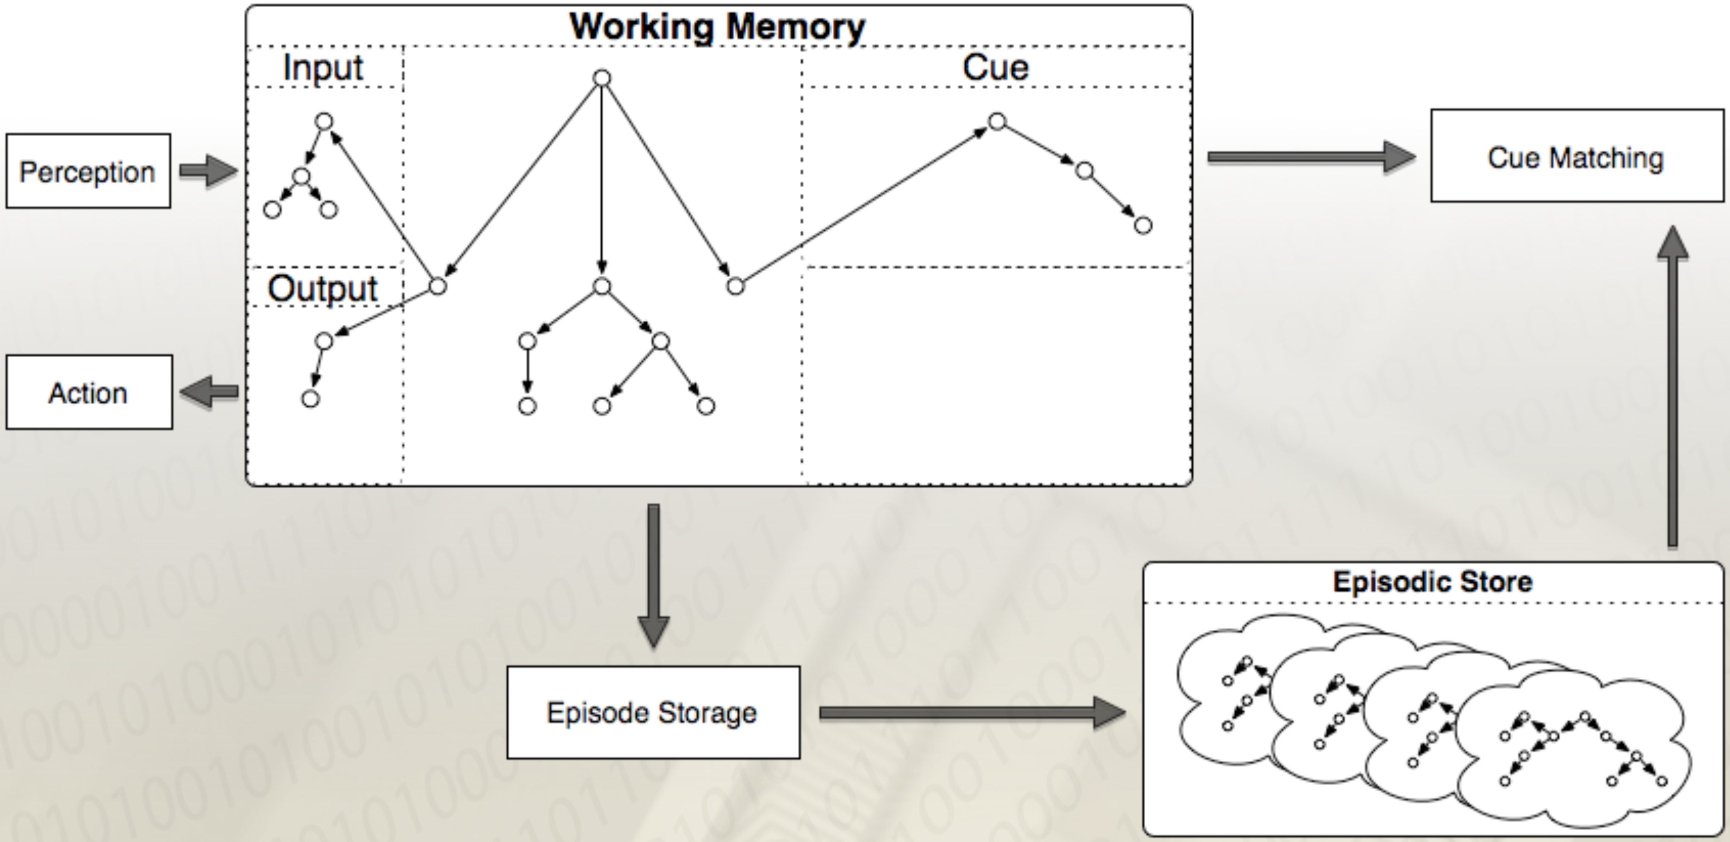

Episodic Operations

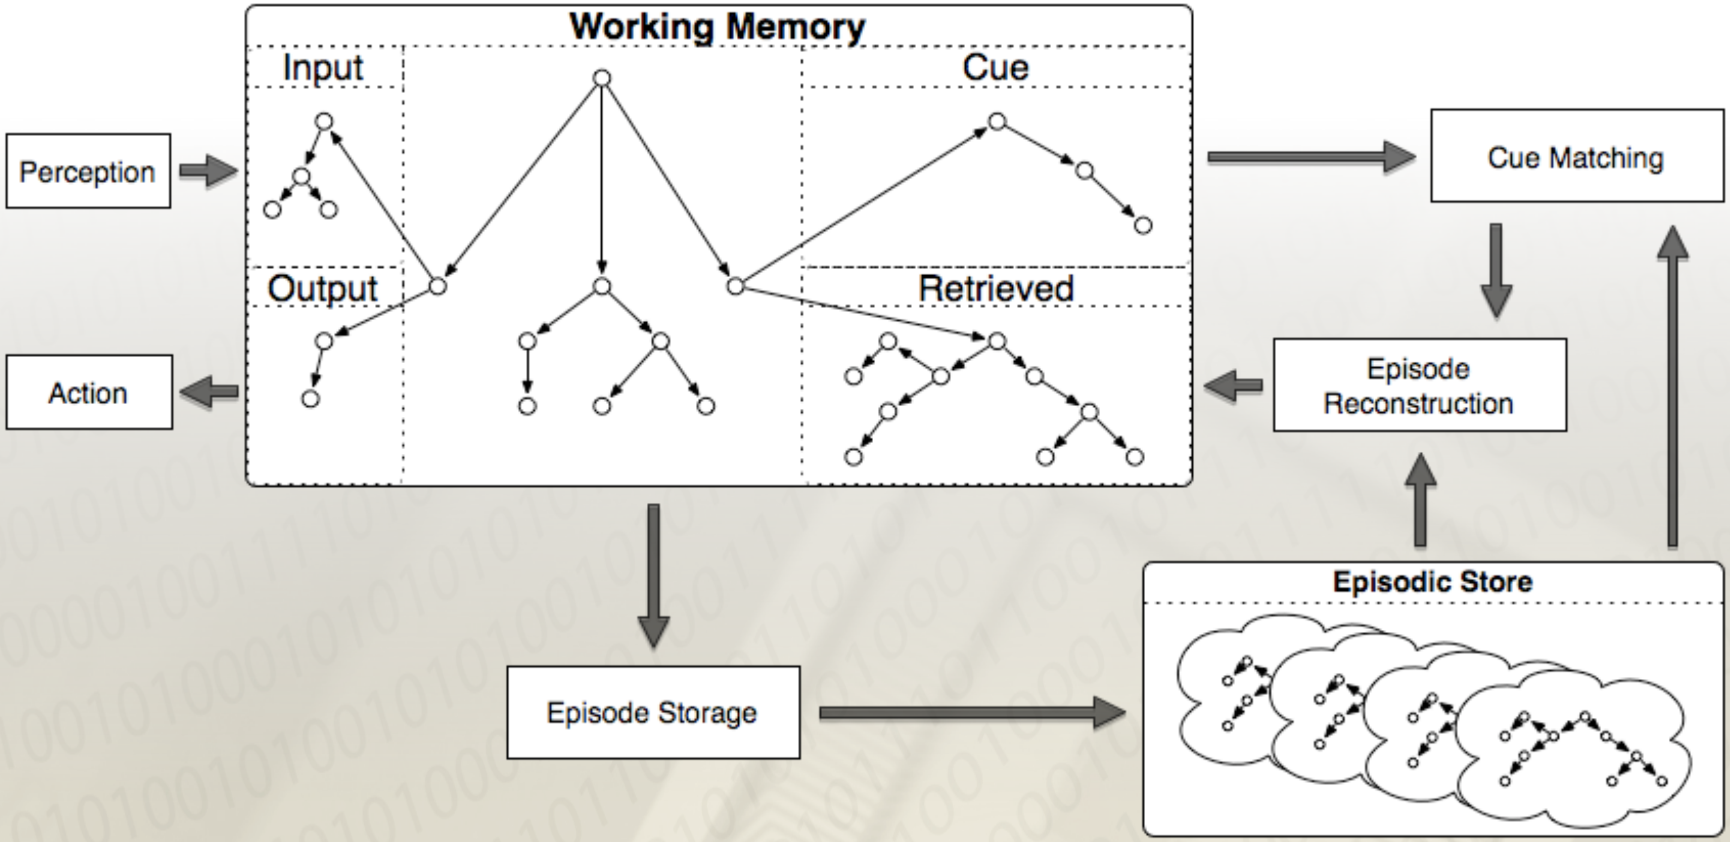

Comparison to CBR

CBR

Cases

- Contain problems and solutions
- Fields pre-specified

Case Base

- Fixed or slowly growing
- Deliberate updates
- No temporal relation between cases

EpMem

Episodes

- Structure and content reflect agent's experiences
  - Potentially fine-grain

Episodic Store

- Grows with experience
- Architectural & automatic storage
- Temporally structured

The Promise of EpMem

- Episodic memory has the potential to support cognitive capabilities across
  - Sensing
  - Reasoning
  - Learning

Nuxoll, A.: Enhancing Intelligent Agents with Episodic Memory. (2007)

Sensing

Virtual Sensing

- retrieving past sensing of features outside current perception that is relevant to the current task

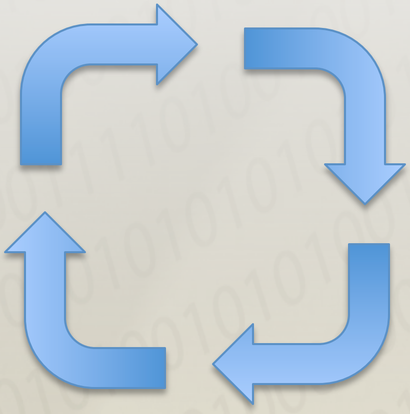

Detecting Repetition

- realizing when you are repeating the same series of actions and altering your behavior as a result

Reasoning

Action Modeling

- predicting the immediate outcome of your actions

Managing Long Term Goals

- keeping track of a plan and what steps in that plan have been accomplished so far

Learning

Retroactive Learning

- reviewing experiences and learning from them when sufficient resources become available

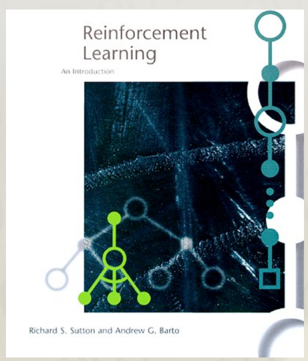

"Boost" Learning

- provide a database of knowledge that can be manipulated by other learning mechanisms

Outline

- Episodic Memory
  - What is It
  - Why is It Important
- **The Challenge**
  - Statement and Focus
  - Bounds and Requirements
- Conclusions
  - Future Work
  - Review and Discussion

The Challenge

- One year of continual use of episodic memory
  - Embedded within an agent living an "interesting" life
    - Learning about multiple challenging tasks
    - Dynamic environment

- Focus: efficiently support specific functional capabilities related to storing, maintaining, and retrieving experiences

Efficiency Issues

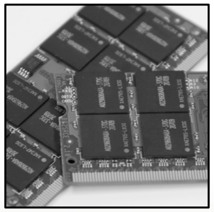

- Bounded Storage
  - Memory is cheap, plentiful, but not unlimited

- Bounded Retrievals
  - Complex, dynamic environments impose real-time constraints on agents
  - Episodic storage and retrievals must…
    - Not interfere with agent's ability to respond in the world
    - Retrieve information quickly enough to be useful

Bounded Storage

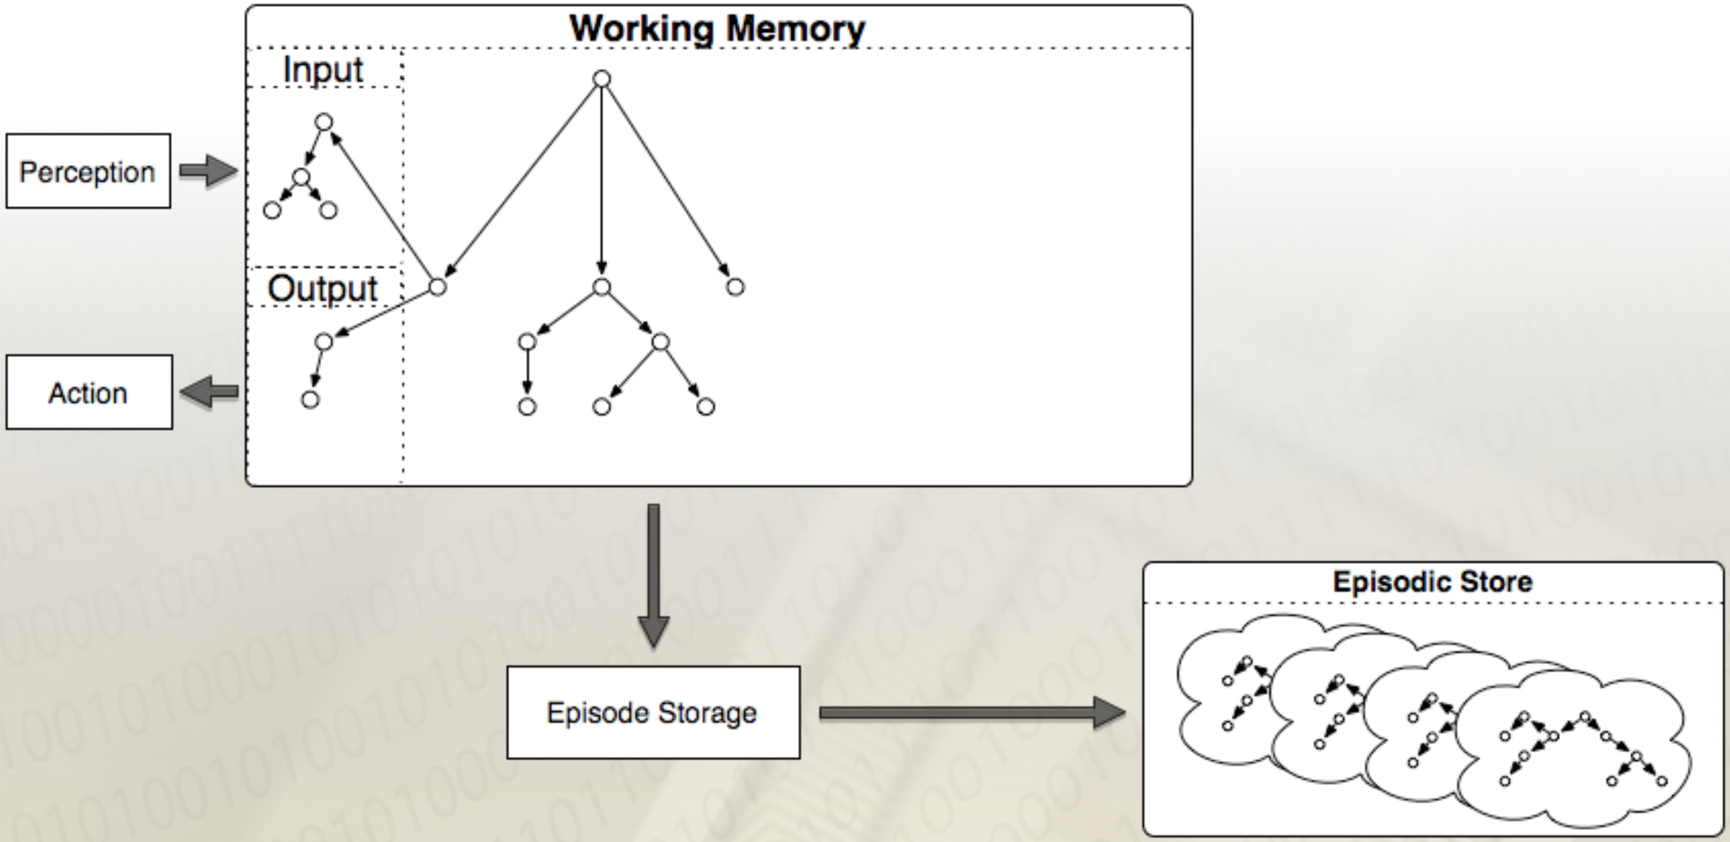

Naïve Implementation

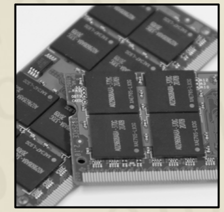

Time	Working Memory		Episodic Store		
1	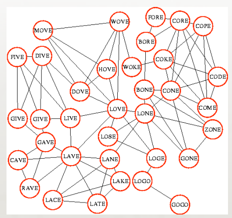		1		
			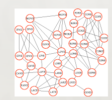		
2	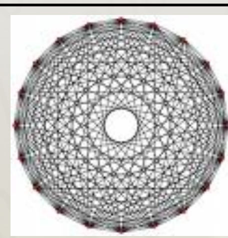		1	2	
			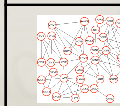	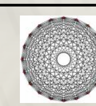	
3	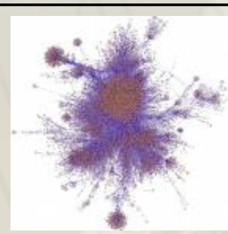		1	2	3
			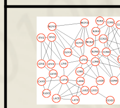	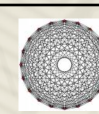	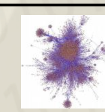

Cost of Storage

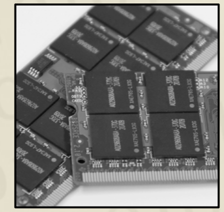

cost =	**(time)**	**(frequency)**	**(ep size)**	**(representation)**
bytes	sec	ep/sec	features/ep	bytes/feature

A Year of Storage

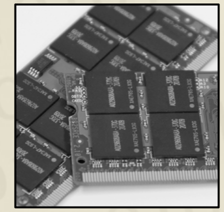

- Assume 16 hours waking activity/day
- (1 year)(365 day/year)(16 hour/day)(3600 sec/hour)

~ 21M sec

Frequency of Storage

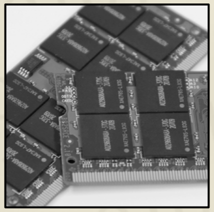

- Historical detail vs. efficiency
- An accurate episodic record demands capturing all structural and feature changes that have occurred since the last recorded episode
  - Worst case: reproduce all structure/features
  - Frequency is a linear multiplier of this cost
- Worst case: 50 ms (**20 ep/sec**)
- Best case: 500 ms (**2 ep/sec**)

CSE

Feature Storage

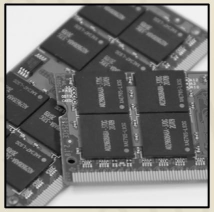

- A rich environment suggests a large representation of the current situation
  - Best case: **100 features**
  - Worst case: **1000 features**

- Real-world environments demand rich, relational descriptions to adequately express arbitrary, complex structures
  - Best case: **10 bytes/feature**
  - Worst case: **100 bytes/feature**

Cost of Storage: Revisited

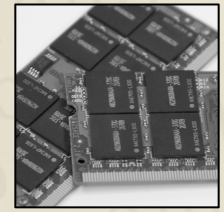

cost =	(time)	(frequency)	(ep size)	(representation)
bytes	sec	ep/sec	features/ep	bytes/feature

Best Case

- 21M sec
- 2 ep/sec
- 100 features/ep
- 10 bytes/feature
- **42GB**

Worst Case

- 21M sec
- 20 ep/sec
- 1000 features/ep
- 100 bytes/feature
- **42TB**

Storage Summary

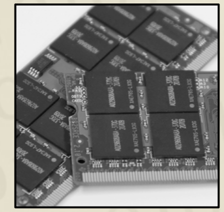

- **42GB – 42TB**

- Low end can fit entirely in commodity server main memory
  - Thus storage alone probably not limiting factor

- Due to simple encoding, 2 – 20 episodes/sec is well within capabilities of current processors

Bounded Retrievals

Match a cue against episodic store, retrieve *best* match

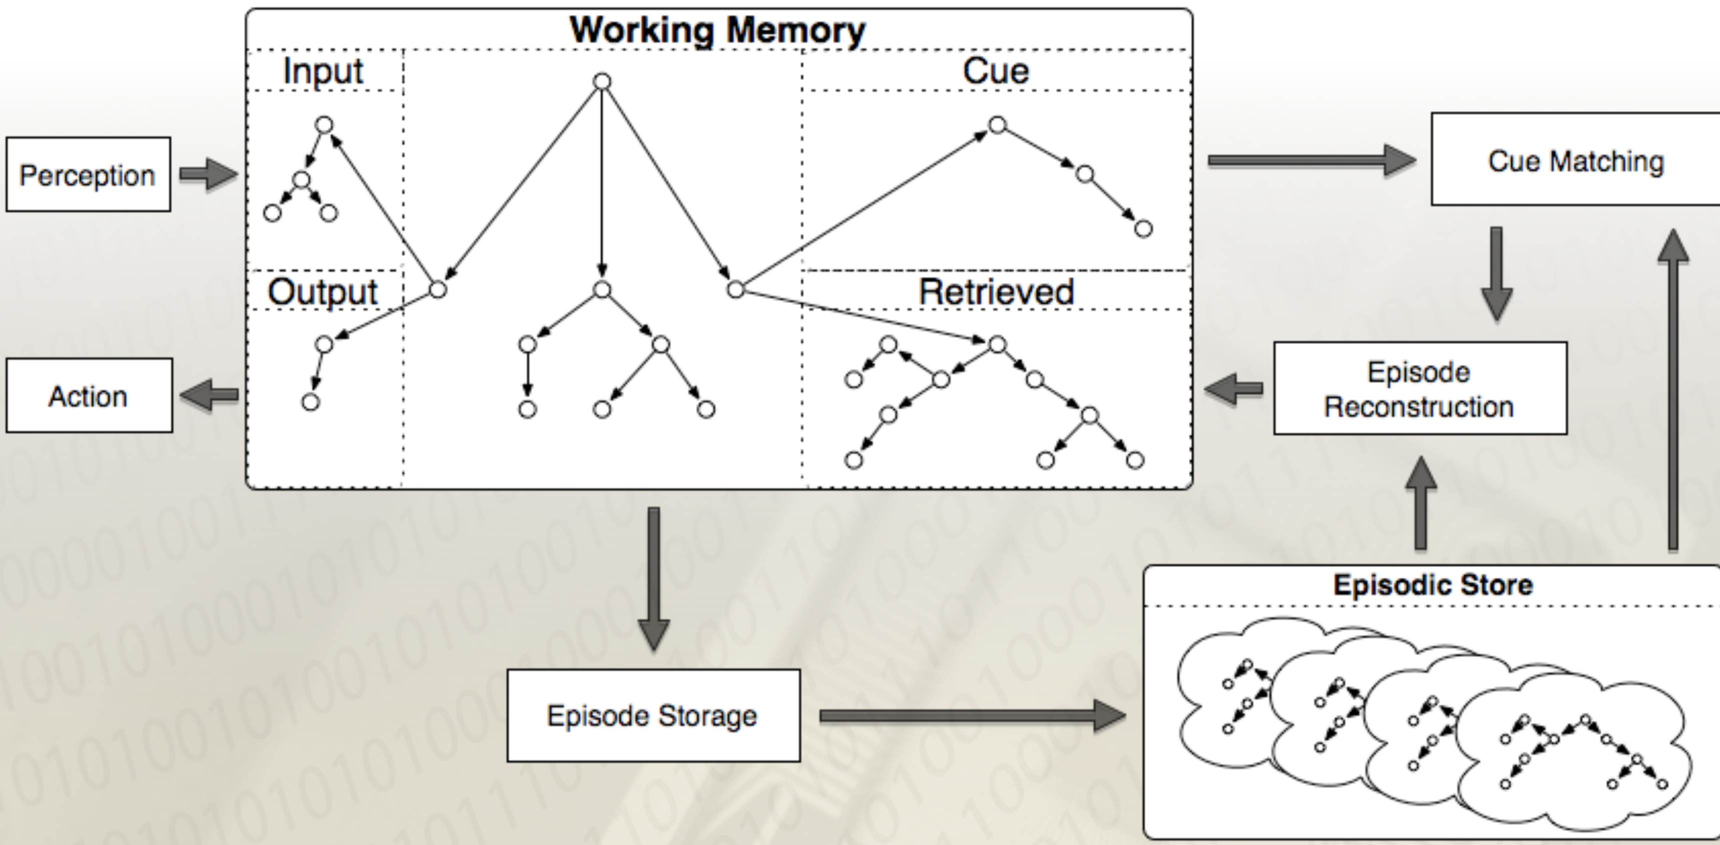

What Bounds Retrievals?

- In dynamic environments, a memory will lose its utility if not retrieved within some limited amount of time after cue initiation

- Assume episodic retrievals have a dedicated processor that can process a single retrieval in *parallel* to primitive decisions
  - Bound = (decision time) (utility w.r.t. world dynamics)
  - Assume decision time: 50 ms
  - Assume utility: 20 decisions
  - **Fixed bound: 1 sec**

Worst Case Retrieval Cost

- Linear scan = (data to scan) (time/datum)

- Best case
  - 42 GB of data
  - Assume 2GHz CPU
  - Time after 1 year: **20 sec**
    - **Worst case is 1000x**

Retrieval Summary

- Utility bound: **1 sec**
- Linear scan: **20 sec**

- <u>Real issue</u>
  – How to effectively organize episodic data, incrementally as it is learned over a year, such that it can be searched in <u>bounded time</u>

Outline

- Episodic Memory
  - What is It
  - Why is It Important
- The Challenge
  - Statement and Focus
  - Bounds and Requirements
- **Conclusions**
  - Future Work
  - Review and Discussion

Thoughts on Approaches

- Technology (hardware)
  - Custom, content-addressable memories
  - Massive parallelism

- Algorithms
  - Compression from exploiting repeated structure
  - Gains from exploiting temporal regularities
    - *Only process changes*
  - Clever data structures/algorithms to maintain best match
    - NN still linear in worst case

Derbinsky, N., Laird, J.E.: Efficiently Implementing Episodic Memory (2009)

Thoughts on Approaches (2)

- Heuristic retrieval strategies
  - History compression
  - Fast familiarity (via locally sensitive hashing)
  - Forgetting/consolidation
  - Query caching/optimization ala RDBMS

- Evaluation: efficiency vs. proficiency

CSE

Review the Challenge

- Endowing an agent with an episodic memory provides knowledge to support a vast array of cognitive capabilities crucial for intelligent behavior
- Our challenge: a year of continuous episodic memory in an agent living an *interesting* life
  - Storage: **42GB – 42TB**
    - Linear scan: **20 sec**
  - Retrieval: **1 sec**
- Much work to be done: will draw on and contribute to a variety of experience-based reasoning research