

# Levels of Learning in Natural and Artificial Agents

(FA9550-19-0180)

**PI: John Laird (University of Michigan)**

**Co-PI: Shiwali Mohan (PARC)**

**Graduate Student: Bryan Stearns (UM)**



**AFOSR Program Review:  
Computational Cognition and Machine Intelligence Program  
(October 8, 2020; Cyberspace)**

# How is Learning Integrated with Performance for Autonomous Agents in Dynamic Environments?

- Key Challenge: Performance & learning must be real-time
  - Must keep up with environmental dynamics with bounded computation.
  - Even as long-term knowledge grows.
  - Across the breadth learning we find in humans.
- Surprisingly little analysis on complexity of online learning algorithms in AI.

# Possible Types of Learning

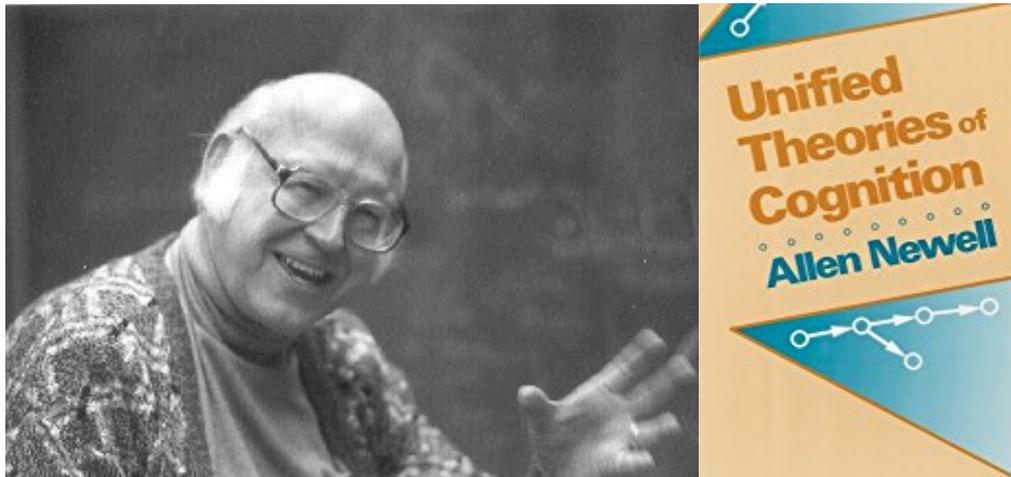
Self-Explanation  
Recognition  
Discovery  
Episodic Learning  
Learning by Analogy  
Category and Concept Learning  
Learning by Instruction  
Sequence Learning

Learning by Demonstration  
Rehearsal  
Procedure Learning  
Meta-Learning  
Temporal-Difference Learning  
Experimentation  
Imitation Learning  
Perceptual Learning  
Practice & Rehearsal

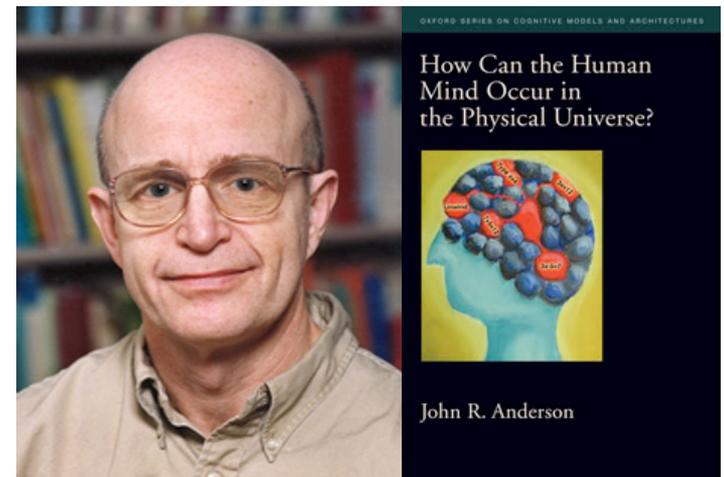
**How are these types of learning integrated in agent architecture?**

# Cognitive Architecture Hypothesis

- Complex cognition arises from a combination of:
  - a fixed set of computational building blocks (memories, processes, representations, learning mechanisms); and
  - knowledge (innate and learned through experience).



Allen Newell



John Anderson

# Our Hypothesis: Two Levels of Learning

## 1. Architecture mechanisms (L1) for *basic* learning

- Innate, automatic, continuous, online, not under agent control.
- Parasitic process on top of task performance.
- Small fixed number of mechanisms.
- Examples: production composition, chunking, episodic memory storage.

## 2. Knowledge-based strategy (L2) for *complex* learning

- Hijacks task processing via metacognition and deliberation.
- Create experiences for L1s: no additional learning mechanisms.
- Can be learned – no fixed number.
- Examples: learn from instruction; retrospective analysis; deliberate training; experimentation or exploration; ...

# Level 1

Self-Explanation  
Recognition  
Discovery  
Episodic Learning  
Learning by Analogy  
Category and Concept Learning  
Learning by Instruction  
Sequence Learning

# Level 2

Learning by Demonstration  
Rehearsal  
Procedure Learning  
Meta-Learning  
Temporal-Difference Learning  
Experimentation  
Imitation Learning  
Perceptual Learning  
Practice & Rehearsal

# Test Hypothesis: Common Model of Cognition

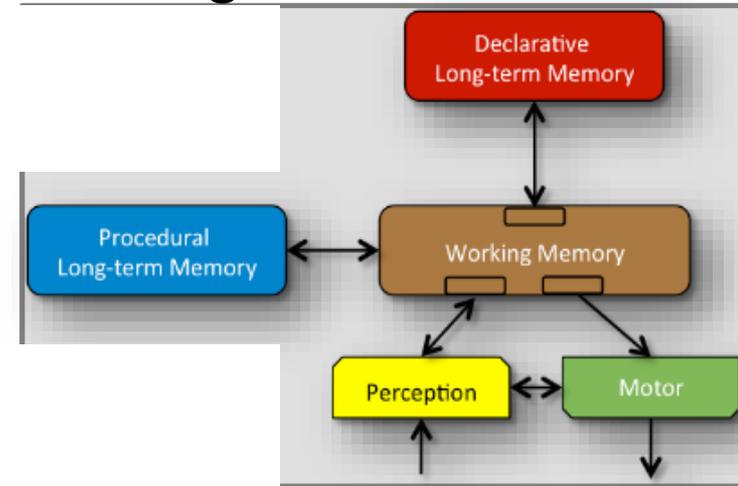
Dig into how performance and learning are integrated in Common Model for architectural learning (L1).

- What processing is done during task performance?
- How is learning integrated with performance?
  1. What data (and *metadata*) is used for learning?
  2. What processing is performed with that data?
  3. What is the computational complexity of learning?

# Common Model of Cognition

- An abstract specification of human-like cognitive architecture

- Community consensus/agreement
- Not itself a cognitive architecture
- Still missing many components
  - Perceptual and motor learning



- A set of communicating processing and memory modules

- Processing is parallel across modules
- Processing is parallel within modules
- Learning mechanisms associated with many modules

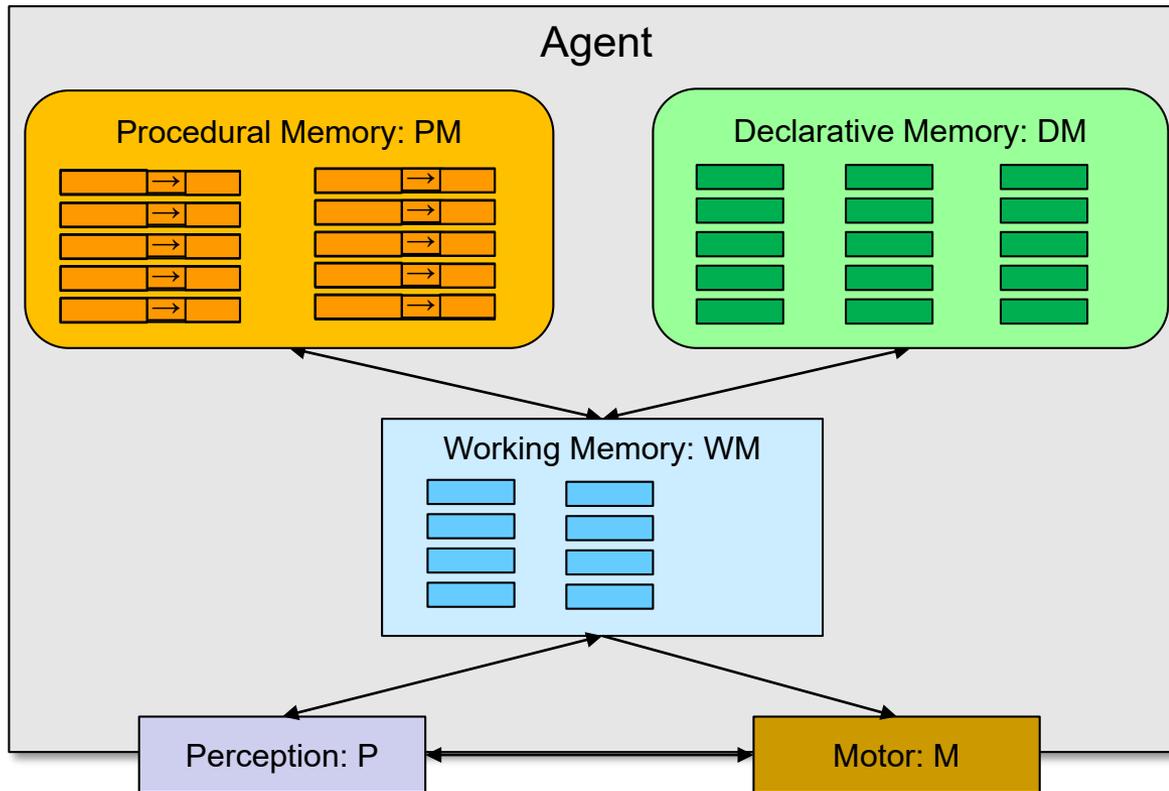
- Cognitive cycle: sequential actions

- Complex behavior arises from a sequence of cognitive cycles:
  - Each cycle is 50msec in humans
  - No additional modules for complex cognition

# Why Common Model?

1. Covers a range of implemented, well-researched cognitive architectures: Soar, ACT-R, Sigma, Spaun, LIDA, ...
2. Includes multiple architectural learning mechanisms.
3. Has strong connections to the human mind and brain.
  - *Andrea Stocco presentation @2:15pm.*

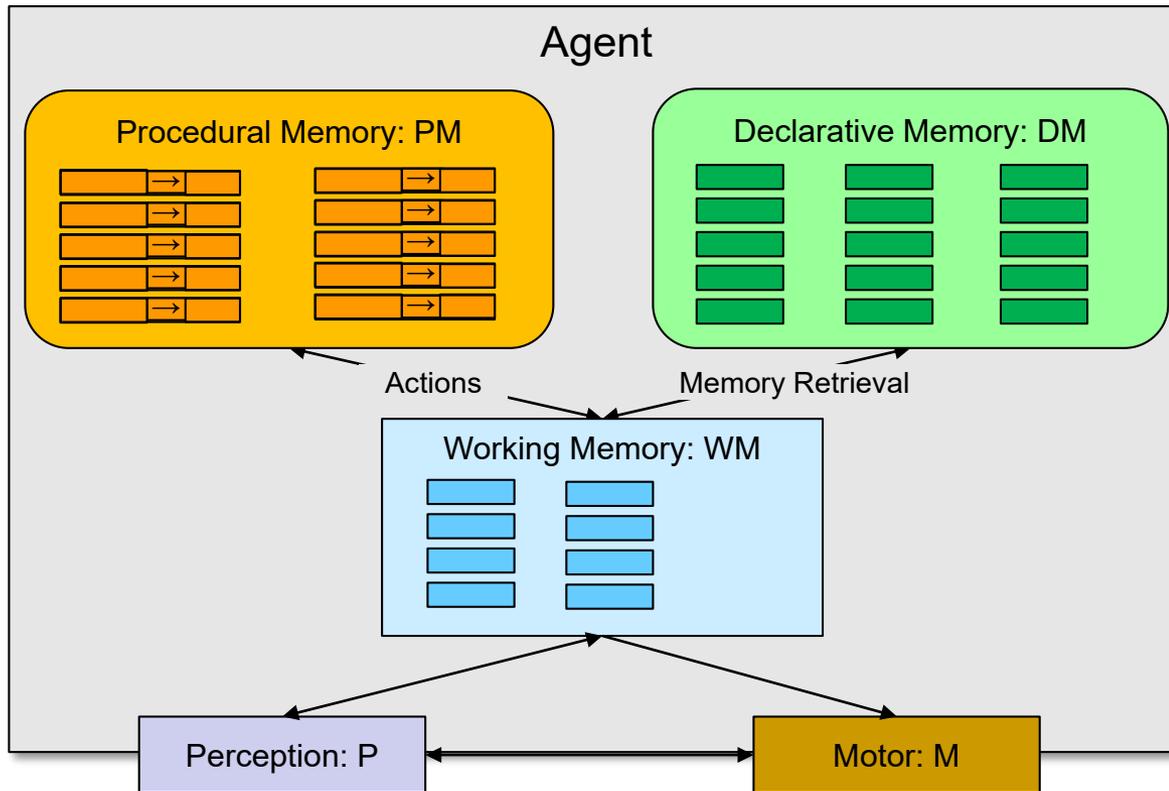
# Agent Architecture: Data/Knowledge



- Architecture processing and structure is fixed, not open to learning
- Memories contain independent knowledge elements.

- Working Memory [WM]: Perception, situational awareness, goals, intentions, hypotheticals, ...
- Declarative Memory [DM]: Facts, beliefs, experiences, ...
- Procedural Memory [PM]: Skills, procedures, goal structures, ...

# Agent Architecture: Data/Knowledge

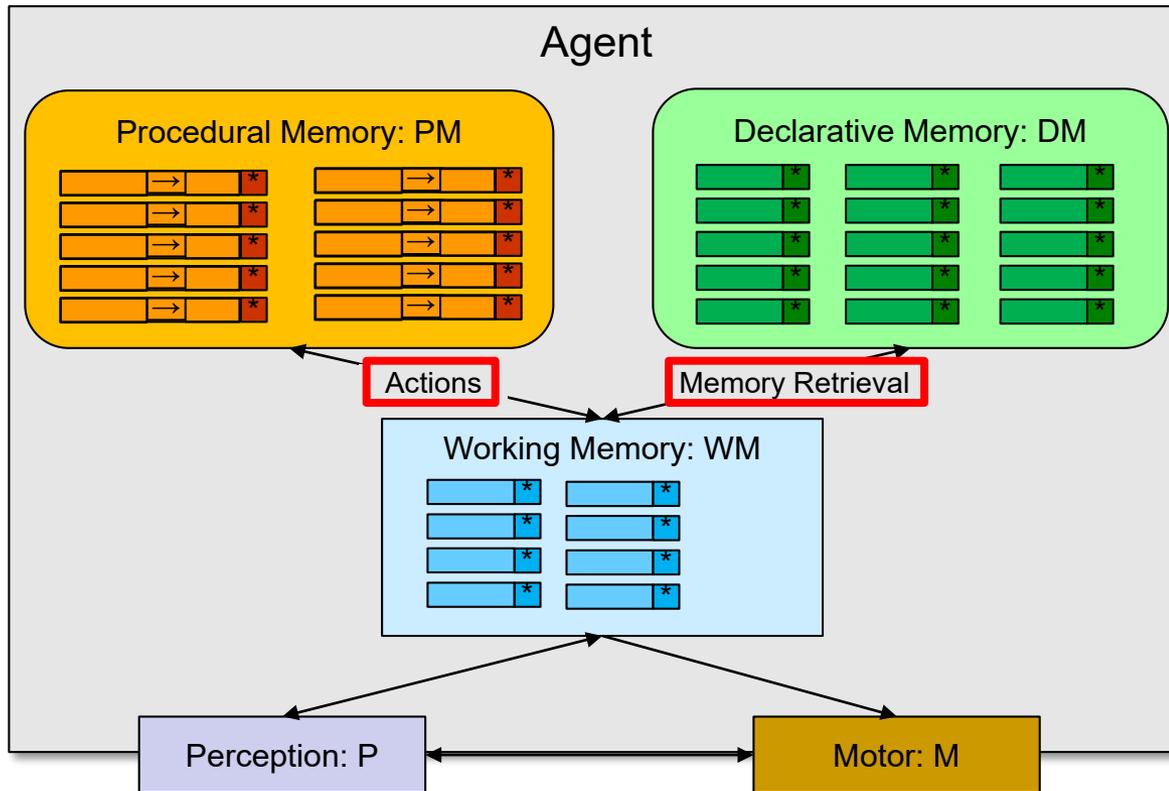


Main computational cost is accessing PM and DM relative to WM.

Everything else is cheap.

- Perception: Change to WM from sensors
  - $\Delta P \rightarrow \Delta WM$
- Internal Action: Changes to WM initiated by procedural memory (PM)
  - $\Delta WM; WM; PM \rightarrow \Delta WM$
- DM Retrieval: Cue in WM leads to retrieval from declarative memory (DM) into WM
  - $\Delta WM; WM; DM \rightarrow \Delta WM$
- Motor Action: Command in WM is sent to motor system
  - $\Delta WM \rightarrow \Delta M$

# Agent Data & Meta Data

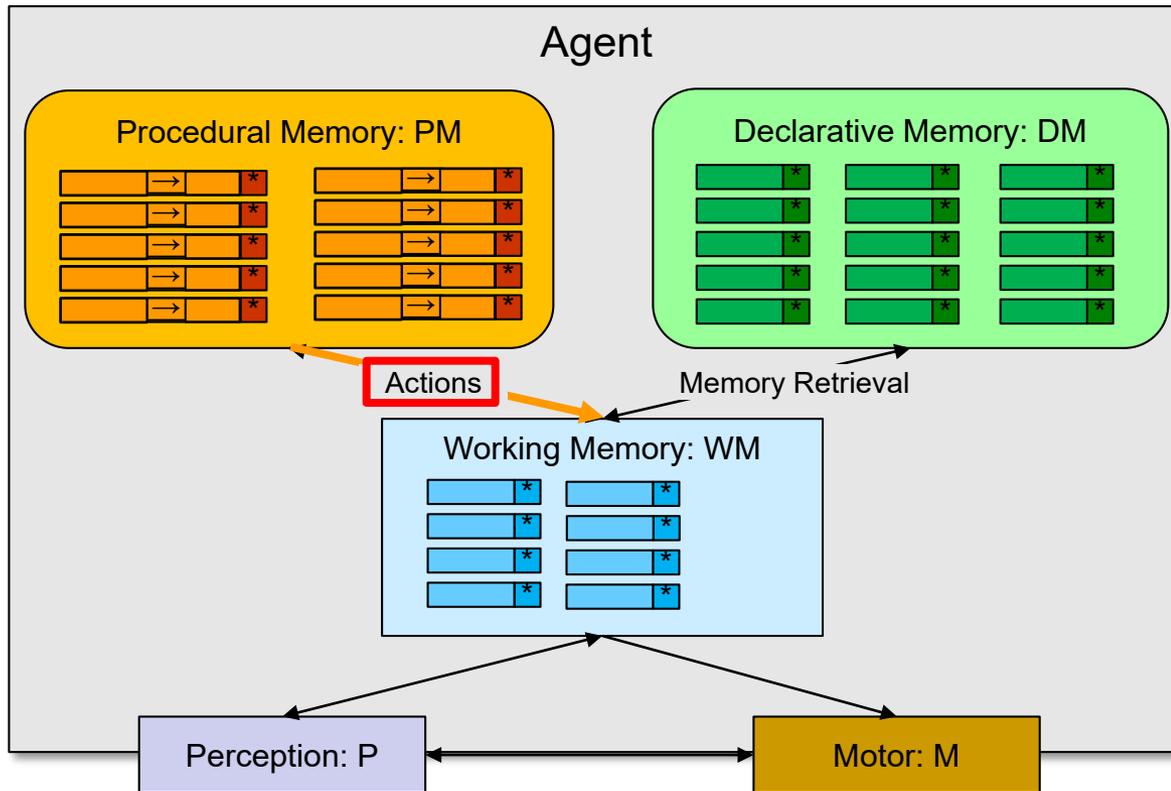


Metadata =  
numeric/statistical data  
associated with agent data.

Not directly testable by  
procedural memory.

- Working Memory Metadata [WM\*]: History of creation & access → activation.
- Declarative Memory Metadata [DM\*]: History of creation & access → activation.
- Procedural Memory Metadata [PM\*]: History of access, expected utility, ...

# Agent Internal Action: with Metadata



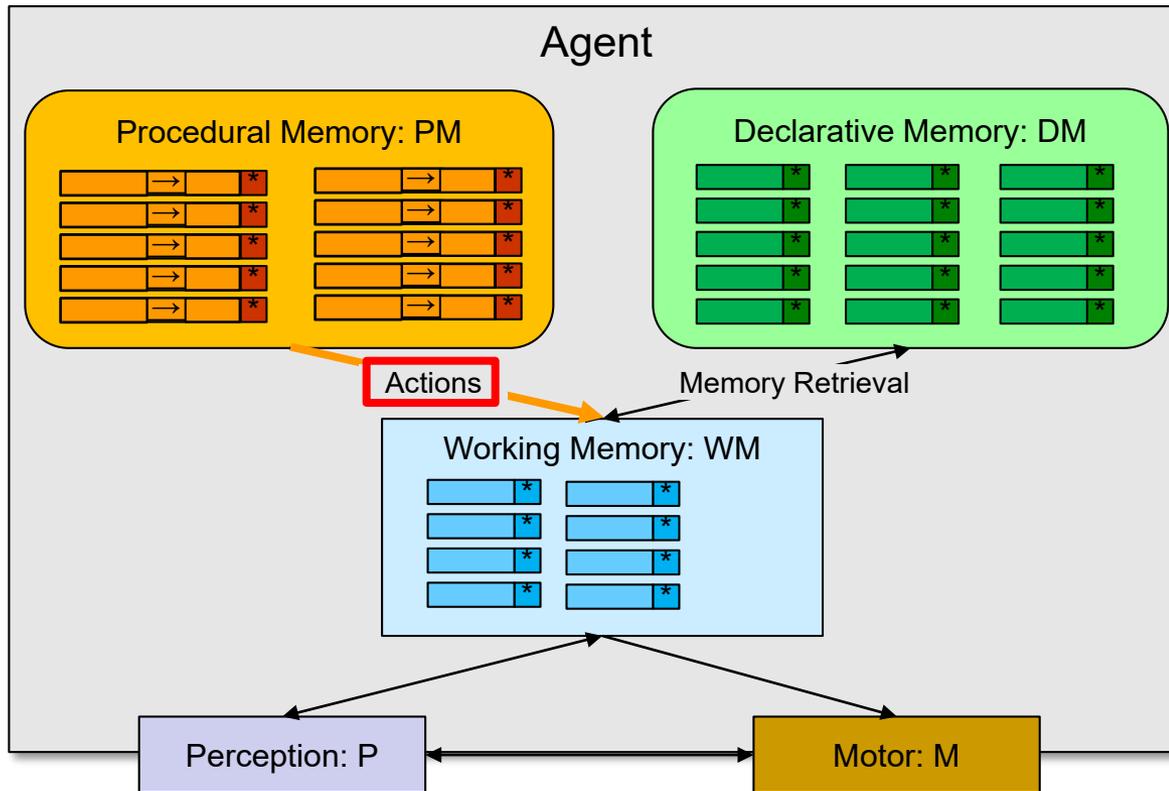
Learning = change to long-term memory structures

Learning! [RL]

Action:  $\Delta WM; \underline{WM}; PM \rightarrow \Delta WM$   
 with  $PM^*$  &  $WM^*$   $\Delta WM; \underline{WM}; WM^*; PM; PM^* \rightarrow \Delta WM; \Delta WM^*; \Delta PM^*$

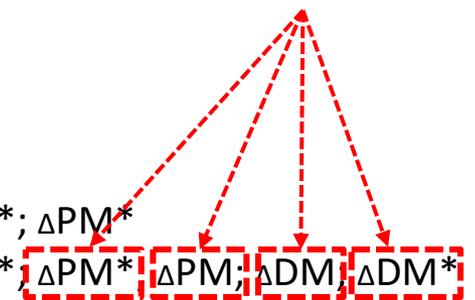
- Procedural metadata [ $PM^*$ ] and working memory metadata [ $WM^*$ ] influence which available actions [ $PM$ ] modify  $WM$ .
- Working memory [ $WM^*$ ] and procedural memory metadata [ $PM^*$ ] are updated.

# Agent Action with Learning



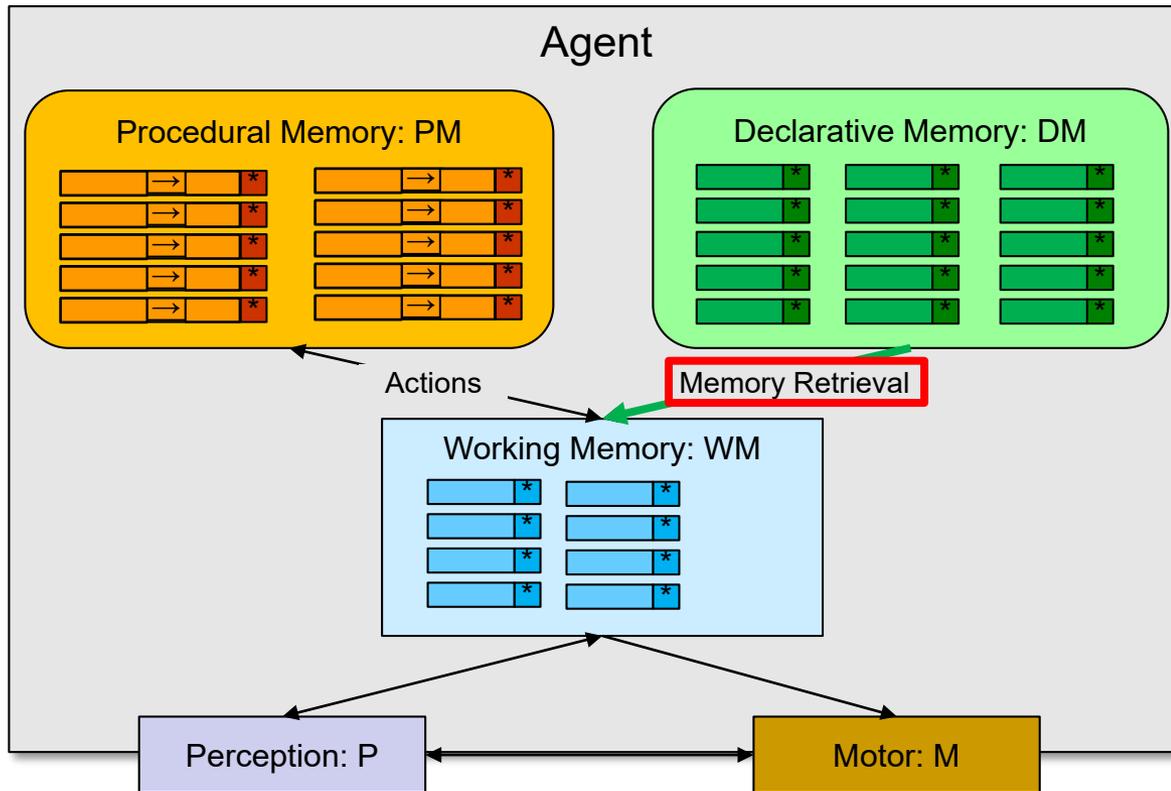
Action:	$\Delta WM; WM; PM$	$\rightarrow \Delta WM$
with $PM^*$ & $WM^*$	$\Delta WM; WM; WM^*; PM; PM^*$	$\rightarrow \Delta WM; \Delta WM^*; \Delta PM^*$
with learning	$\Delta WM; WM; WM^*; PM; PM^*$	$\rightarrow \Delta WM; \Delta WM^*; \Delta PM^*; \Delta PM; \Delta DM; \Delta DM^*$

Learning



- Learn new procedural knowledge (chunking, procedure composition)
- Learn new declarative knowledge (declarative learning – semantic/episodic)
- Update declarative memory metadata (recency, frequency, ...)

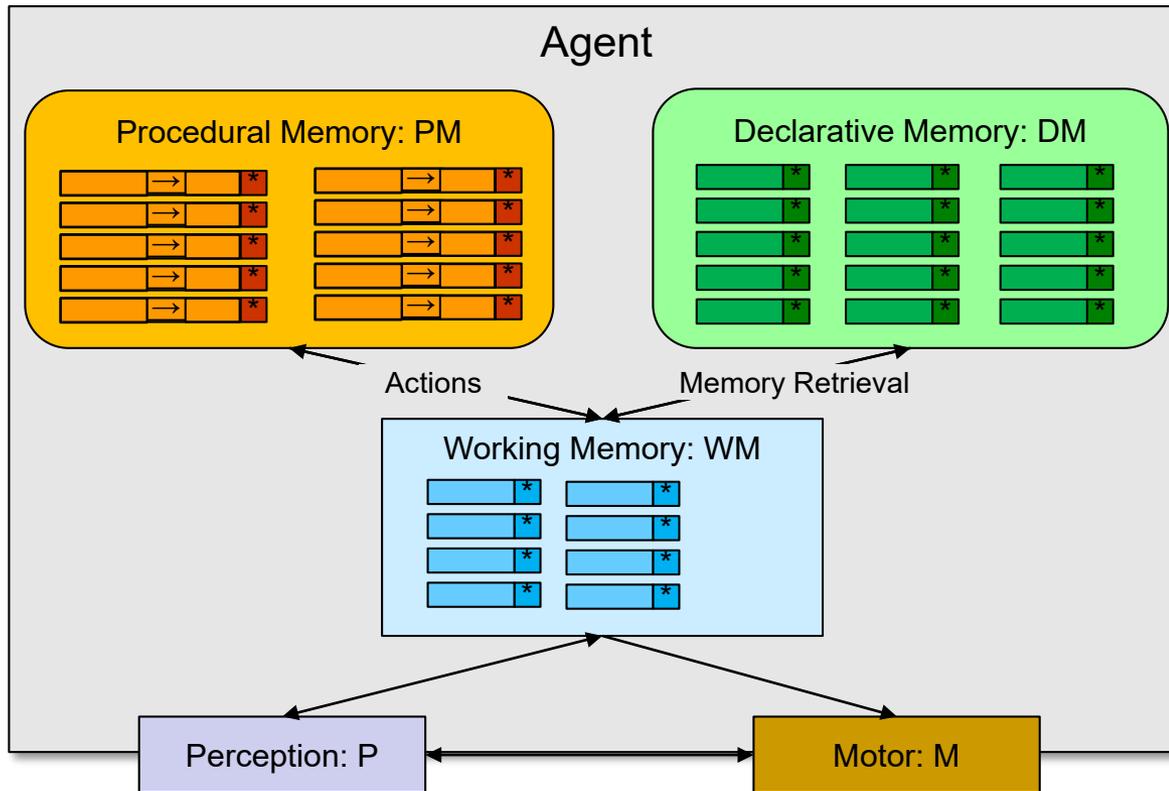
# Agent DM Retrieval with Learning



DM Retrieval:  $\Delta WM; WM; DM$   $\rightarrow \Delta WM$   
 with  $DM^*$  &  $WM^*$ :  $\Delta WM; WM; WM^*; DM; DM^*$   $\rightarrow \Delta WM; \Delta WM^*; \Delta DM^*$

Declarative memory metadata [ $DM^*$ ] and working memory metadata [ $WM^*$ ] influence which structure is retrieved from declarative memory [DM].

# Agent Metadata Maintenance & Forgetting



Local calculations that are easily parallelizable

- WM decay:  $WM^* \rightarrow \Delta WM^*$
  - PM decay:  $PM^* \rightarrow \Delta PM^*$
  - DM decay:  $DM^* \rightarrow \Delta DM^*$
- Update recency/frequency metadata of unchanged data.
- WM Forgetting:  $\Delta WM^* \rightarrow -WM$
  - PM Forgetting:  $\Delta PM^* \rightarrow -PM$
  - DM Forgetting:  $\Delta DM^* \rightarrow -DM$
- Removed data that has decayed beyond threshold.

# Summary: Performance and Learning

## 1. Task Performance: Changes to working memory

- $\Delta WM; WM; WM^*; PM; PM^* \rightarrow \Delta WM$
- $\Delta WM; WM; WM^*; DM; DM^* \rightarrow \Delta WM$
- Requires access to PM and DM relative to WM.

*Main computational expense is accessing PM/DM*

## 2. Short-term Metadata:

- $\Delta WM \rightarrow \Delta WM^*$
- Update accessed WM metadata

*Avoid additional PM/DM accesses; use accessed*

## 3. Learning: Changes to Long-term Memory

- $\Delta WM; WM; WM^*; DM; DM^* \rightarrow \Delta DM^*, \Delta DM, \Delta PM^*, \Delta PM$
- Update metadata based on data and metadata accessed by task performance.
- Create structures based on data and metadata accessed by task performance.

## 4. Metadata Decay and Forgetting:

- $WM^* \rightarrow \Delta WM^*; WM^* \rightarrow \Delta WM; \Delta PM^* \rightarrow -PM; \Delta DM^* \rightarrow -DM$
- Local update of WM, PM, DM metadata

# Limits of Pure Architectural Learning

1. No “deliberate” learning
  - Can’t decide to learn something.
  - Suggests need for meta-cognitive control.
2. Constrained to incremental, constant-time learning algorithms
  - Suggests need for unconstrained learning analysis.
3. Learns from data recently *accessed* by task processing.
  - Limited temporal horizon.
  - Difficult to quickly learn temporal regularities across longer time scales.
  - Suggests need for access to DM during learning.
  - Difficult to learn *implications* of current knowledge.
  - Suggests need for access to PM during learning for planning/etc.

# Hypothesis: Second Level of Learning

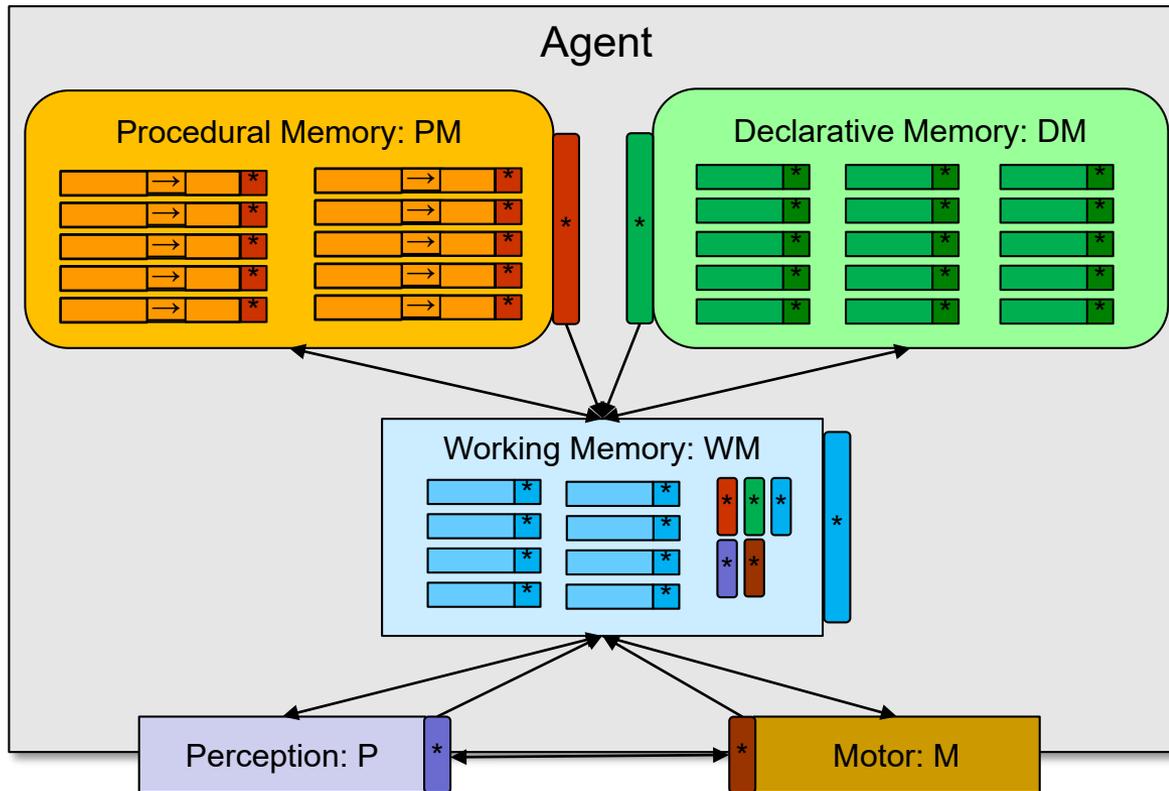
## Knowledge-based strategies

- Occurs during task slack time (but interruptible)
  - Another task the agent pursues
  - Finesses real time, bounded computation issues.
- Metacognitive reasoning across multiple cycles
  - Accesses DM, PM, M to bring in temporally distant data of L1.
  - Creates “experiences” for L1 mechanisms to learn from (no new mechanisms)
    - Create historical, hypothetical/counterfactual experience from memory
    - Create novel experiences through environmental exploration
    - ...
- Allows unconstrained processing of experiences when there was insufficient time during task performance.

# Contrasts between L1 and L2

- L1 continual architectural learning
  - Fast, bounded processing with constrained task knowledge and task metadata all the time.
- L2 intermittent learning strategy
  - Arbitrary processing with arbitrary task knowledge at limited times.
- What about metadata for L2 learning?
  1. Episodic memory
  2. Deliberate maintenance of metadata in semantic memory
  3. Architectural process data that is made explicit (in WM).
    - Failures to retrieve data from long-term memories
    - Failures in perception and motor
    - Innate appraisals: surprise, goal failure, loss of control, ...

# Architectural Process Metadata



Available in working memory for metacognition, such as reasoning about failure.

Available to architecture for learning: episodic memory and RL.

Procedural Memory Process Metadata [PMP\*]:

- Retrieval failure (impasse in Soar)

Declarative Memory Process Metadata [DMP\*]:

- Retrieval failure

Working Memory Process Metadata [WMP\*]:

- Surprise and other appraisals

Perception Process Metadata [PP\*]

- Failure, ...

Motor Process Metadata [MP\*]

- Completion, fault, ...

# Summary and Conclusion

- Analysis of how learning can be integrated with performance for autonomous agents under time constraints
- Analysis of L1 learning in Common Model
  - Identify importance of metadata
  - Identify shortcomings that lead to necessity of L2
  - Identify need for metadata access in L2
- Next talk shows how L1 and L2 mechanisms support Interactive Task Learning
- Still long way to go to completely understand this.

# List of Publications, Awards, Honors, etc. Attributed to the Grant

- Mohan, S., Klenk, M., Shreve, M., Evans, K., Ang, A., and Maxwell, J. (2020). Characterizing an Analogical Concept Memory for Newellian Cognitive Architectures, *Advances in Cognitive Systems*.
- Laird, J. E. (2020). Intelligence, Knowledge & Human-like Intelligence, *Journal of Artificial General Intelligence* 11(2), 41-44. doi:10.2478/jagi-2020-0003.
- Laird, J. E. (2019). Introduction to *Sciences of the Artificial*, 4<sup>th</sup> Edition, Herbert A. Simon, MIT Press. - Not attributed
- Laird, J. E., & Mohan, S. (2018). Learning Fast and Slow: Levels of Learning in General Autonomous Intelligent Agents , *National Conference on Artificial Intelligence, AAAI-2018. Senior Track, Winner of Blue Sky Award*.
- John E. Laird was co-winner (with Paul S. Rosenbloom) of the Herbert A. Simon Award for Cognitive Systems.
- John E. Laird was a member of the Games, Exercises, Modeling and Simulation (GEMS) Defense Science Board (2018-2019).
- John E. Laird is co-organizing a DoD (Basic Research Office - OSD), Future Directions Workshop on the Design of General, Integrated Artificial Systems, July 2019.
- John E. Laird organized the first Virtual International Conference on Cognitive Architecture, June 2020.