

Exploring the Space of Computational Memory Models

Nate Derbinsky and Nicholas A. Gorski
University of Michigan

Lots of Knowledge

Learning agents gain knowledge that may be pertinent to making intelligent decisions

Type	Source	Size
Language	WordNet	>800K assertions
Common Sense	Cyc	>5M facts
Personal History	Episodic Memory	>42GB/year
Skills		
...		

Memory Systems

To harness this experience, while remaining reactive, cognitive architectures employ one or more *memory systems*

Mechanisms to encode, store, and retrieve agent knowledge

Prior psychological and computational work evidences dissociation

Question: for a task, what is the *functionally* optimal set?

Exploring the Space

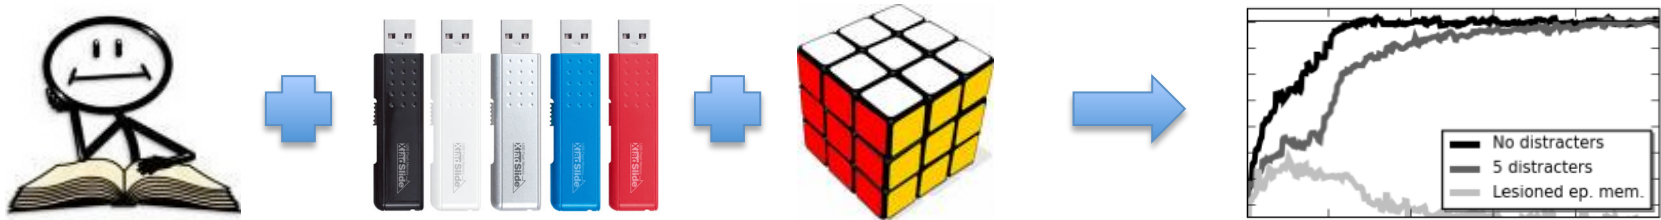
Hypothesis

For different classes of problem spaces, different classes of memory models will result in qualitatively different behavior

Goals

1. Understand how memory system requirements and efficacy change along with parameterized properties of task
2. Develop computational structure and constraints for studies of memory system dissociation

Proposal Outline

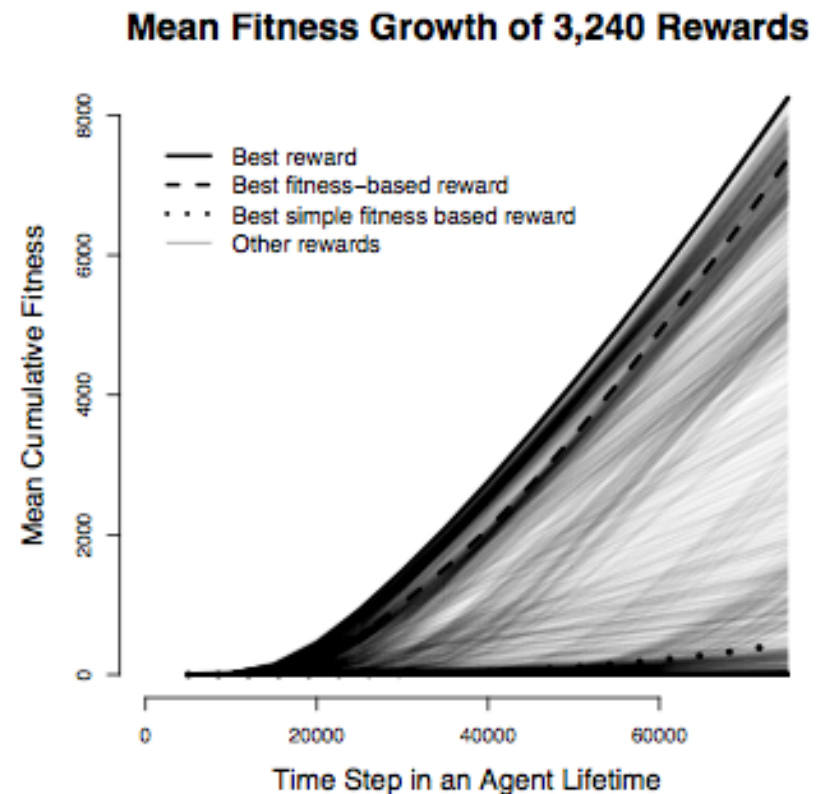


- Assume adaptive (RL) agent with fixed initial knowledge
- Endowed with a set of memory systems
 \in **Memory space**
- Situated in a domain
 \in **Task space**
- Quantitative and qualitative analysis
 - **Evaluation metrics**

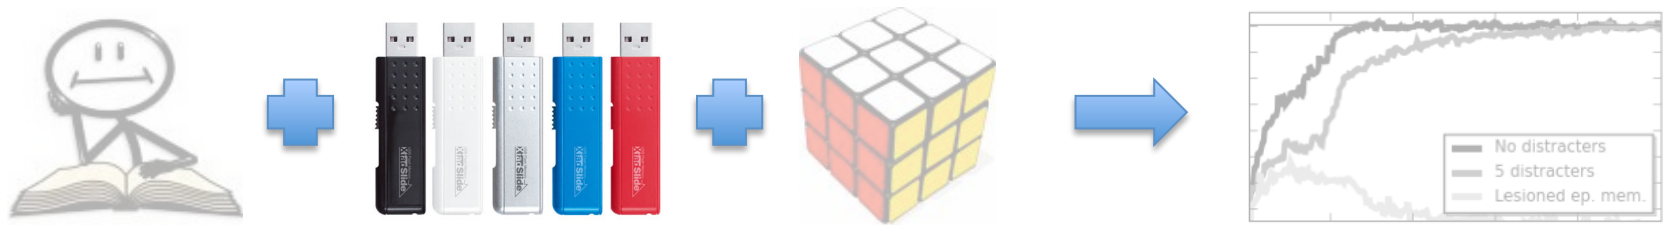
Large Design Space

Comprehensive empirical study of a large design space has been done before

- Reward functions in RL
Singh, Lewis (2009)
- Cognitive architecture
Howes, Lewis, Vera (2009)
- Game modeling
Schvartzman, Wellman (2009)



Progress: Memory Space



A Memory System



We define a memory system implementation as a commitment to features in the space of...

Encoding, Storage, Retrieval

Captures multiple memory systems along common dimensions to facilitate principled exploration

Memory Dimensions



Encoding

Initiation
Determination

Experience Selection
When and what to store

Storage

Granularity
Dynamics

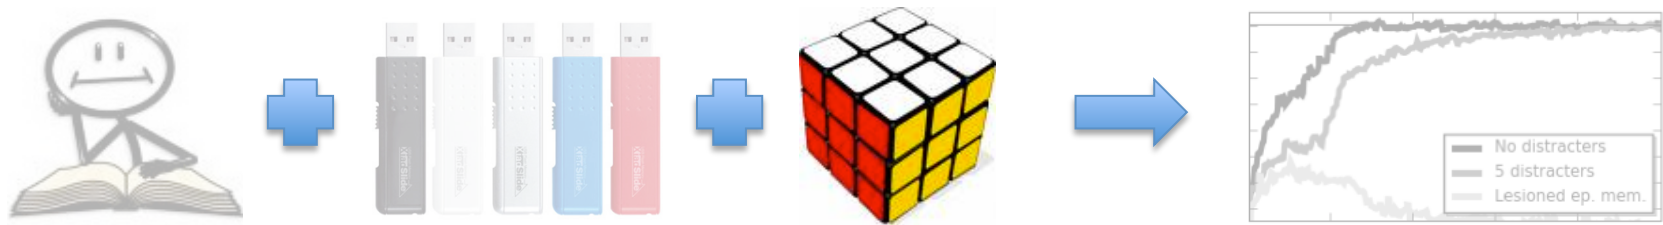
Experience -> Knowledge
How knowledge is stored

Retrieval

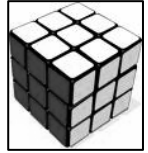
Accessibility
Initiation
Cue Determination
Selection
Result

Knowledge -> Agent
How stored knowledge is brought
to bear

Progress: Task Space



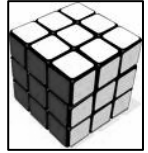
What is Task Space?



A preliminary taxonomy of environmental characteristics pertinent to learning control over internal memory systems

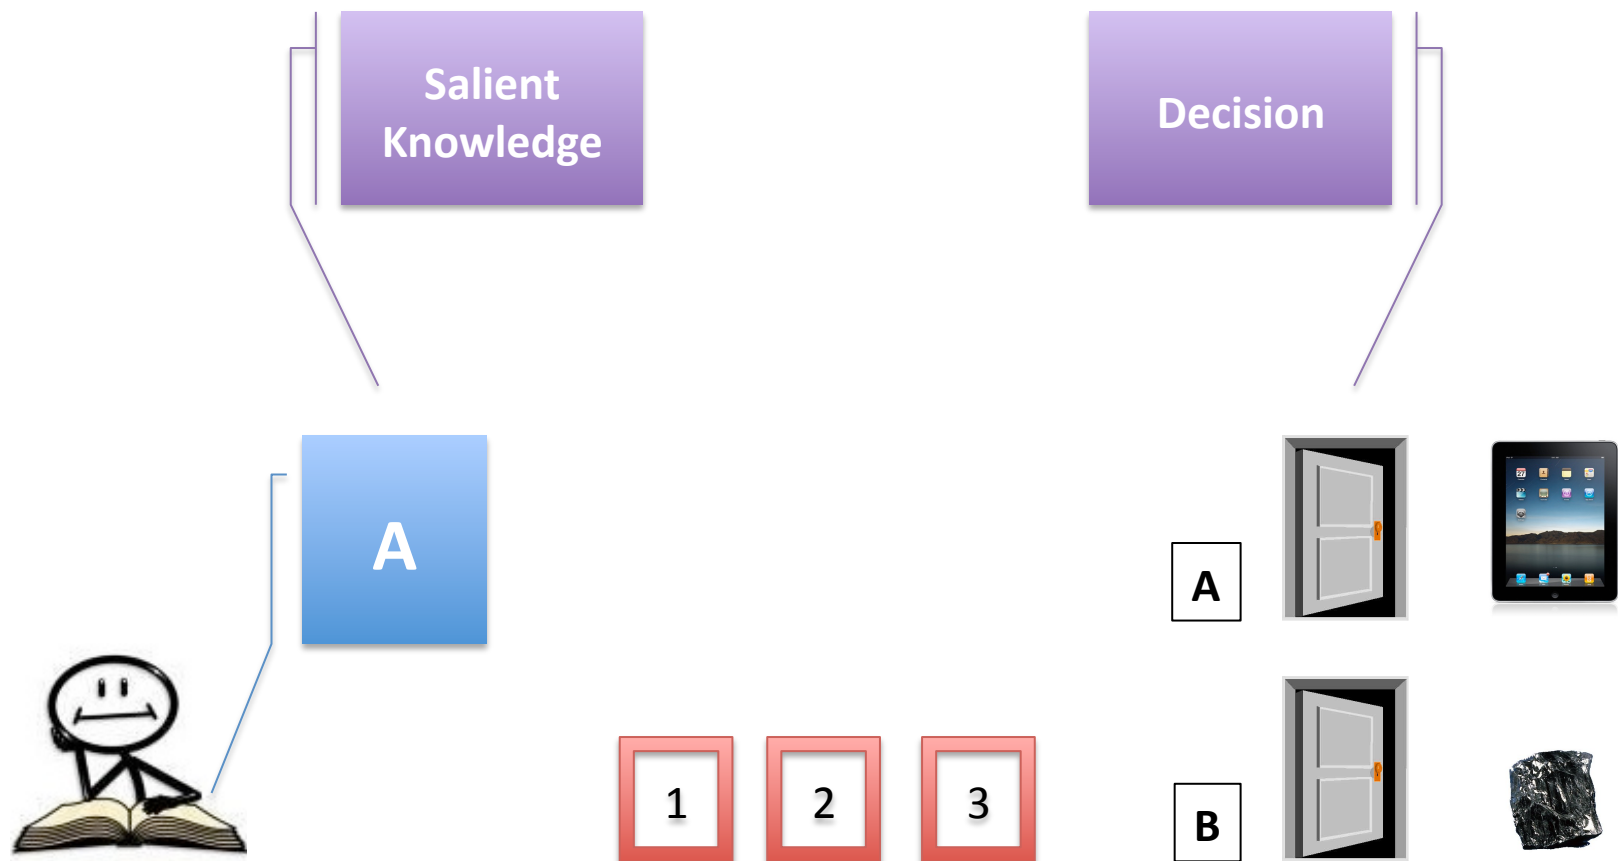
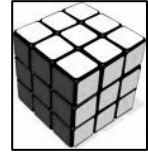
Can be quantitatively parameterized and are independent of each other

Task Space Dimensions

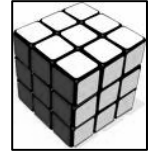


- Temporal Distance to Salient Knowledge
- Categories of Salient Knowledge
- Quantity of Salient Knowledge
- Quantity of Distracting Observations
- Sparseness of Reward
- Relative Cost of Acting vs. Reasoning
- Size of Action Space
- Stochasticity of Actions

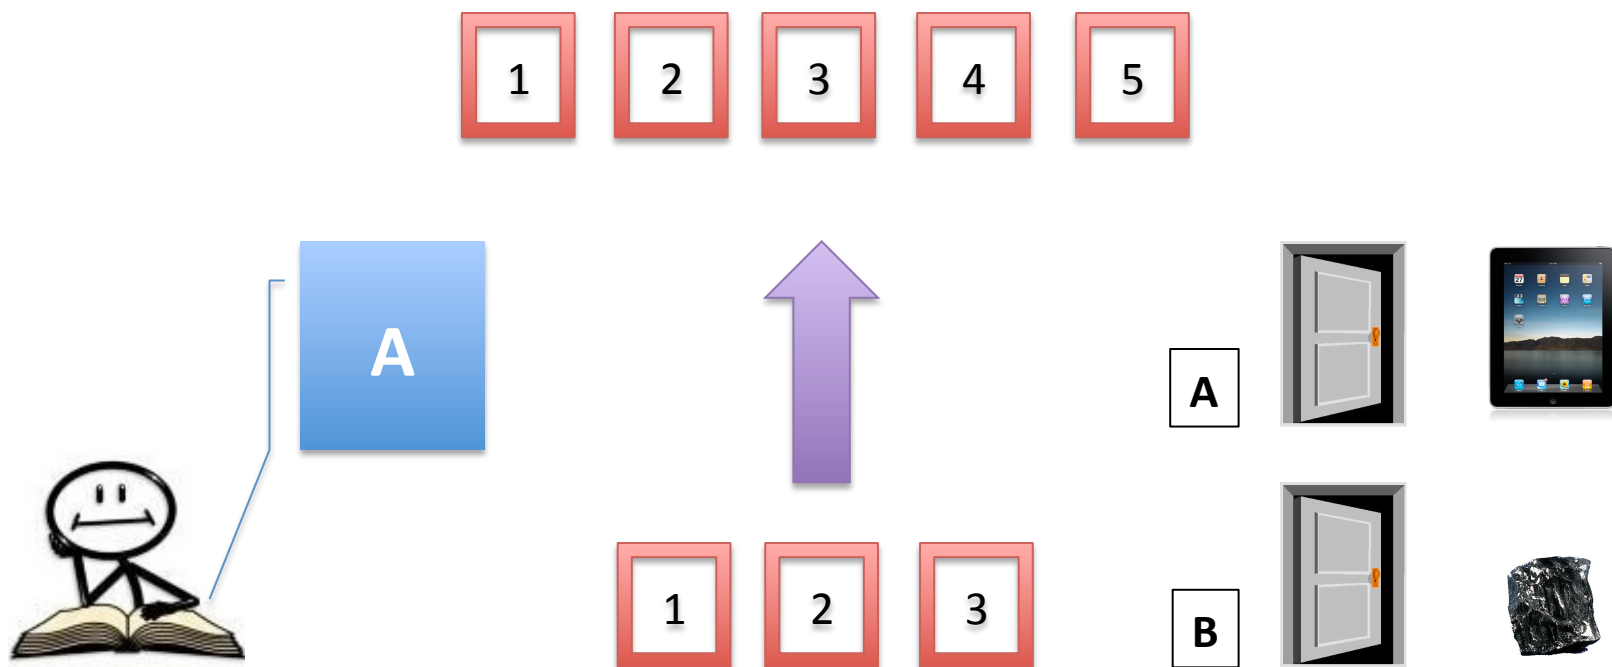
Task Example



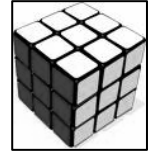
Temporal Distance to Salient Knowledge



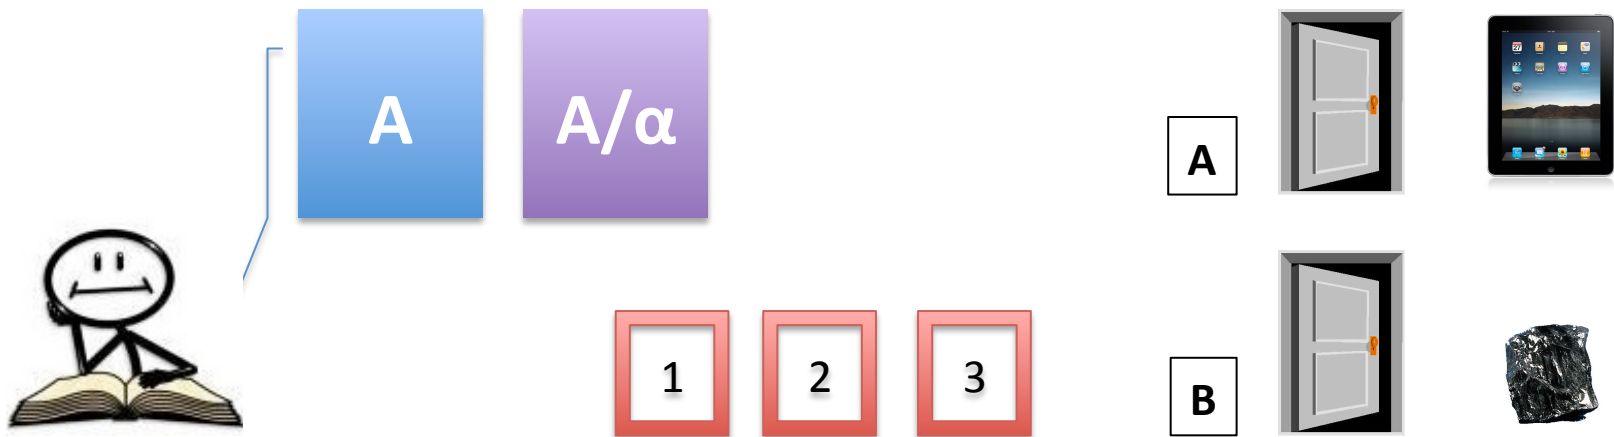
Time between salient knowledge and a decision



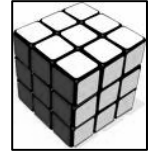
Categories of Salient Knowledge



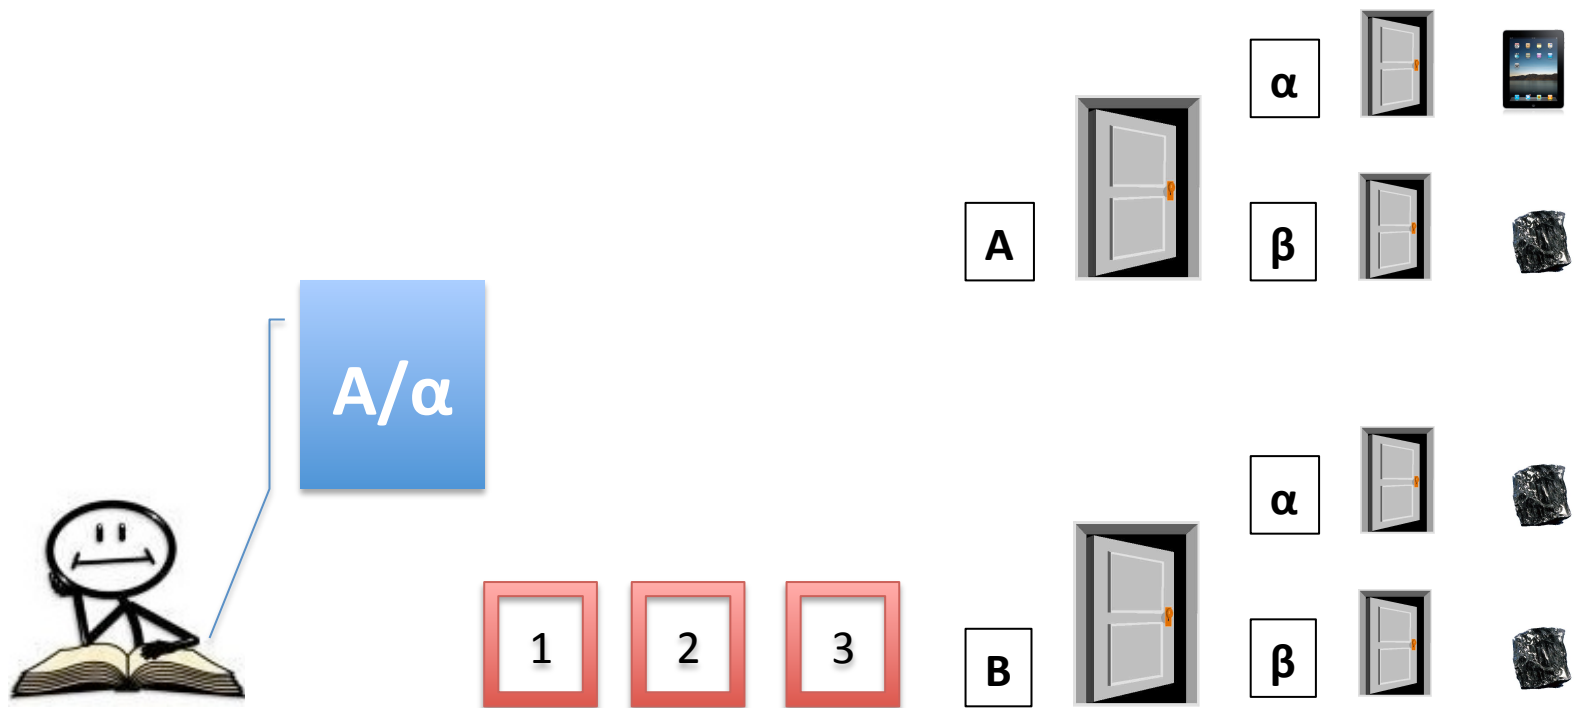
Number of distinct elements of salient knowledge that must be maintained



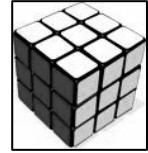
Categories of Salient Knowledge



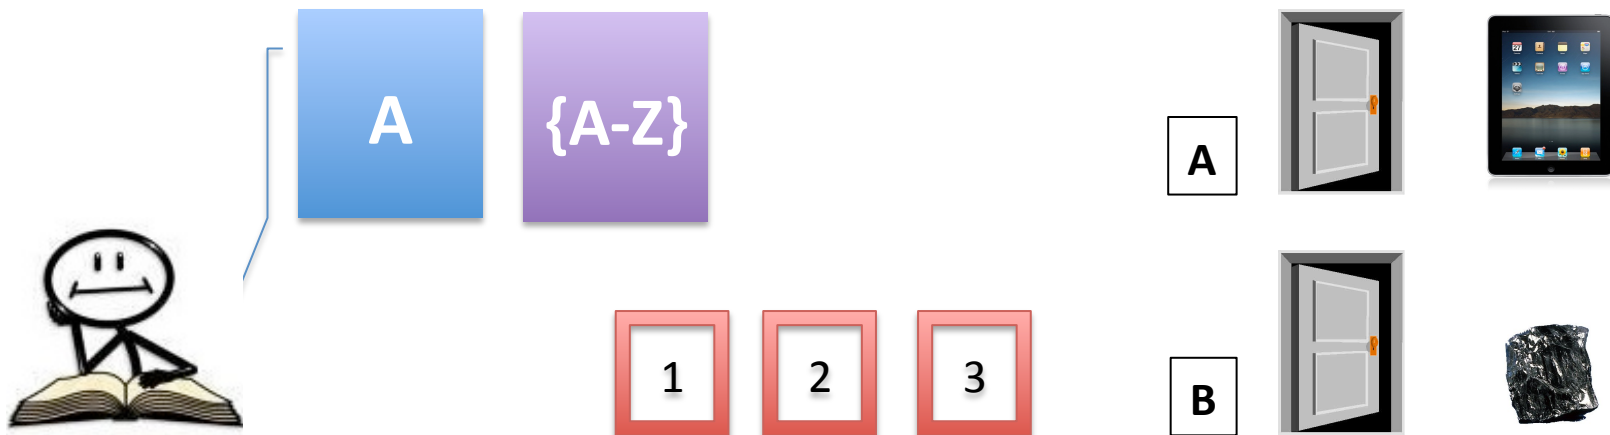
Number of distinct elements of salient knowledge that must be maintained



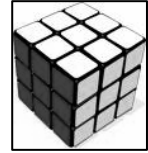
Quantity of Salient Knowledge



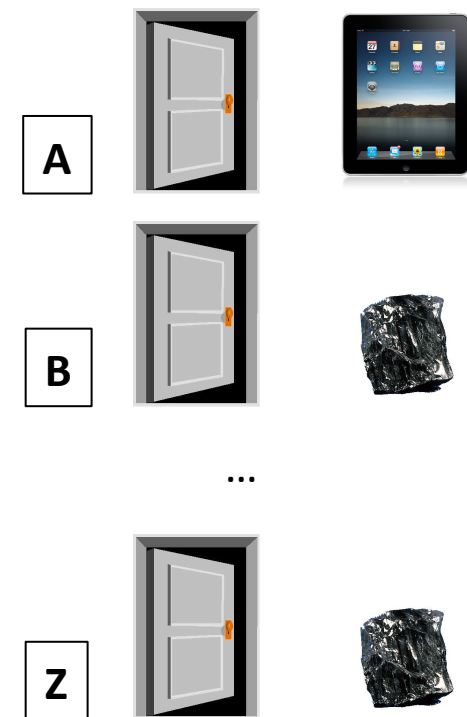
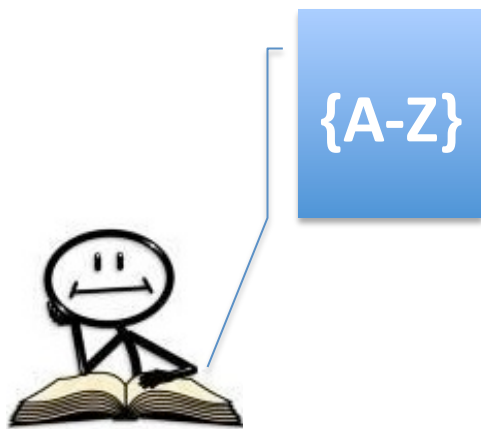
Number of possible values for each category of salient knowledge



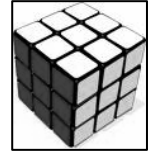
Quantity of Salient Knowledge



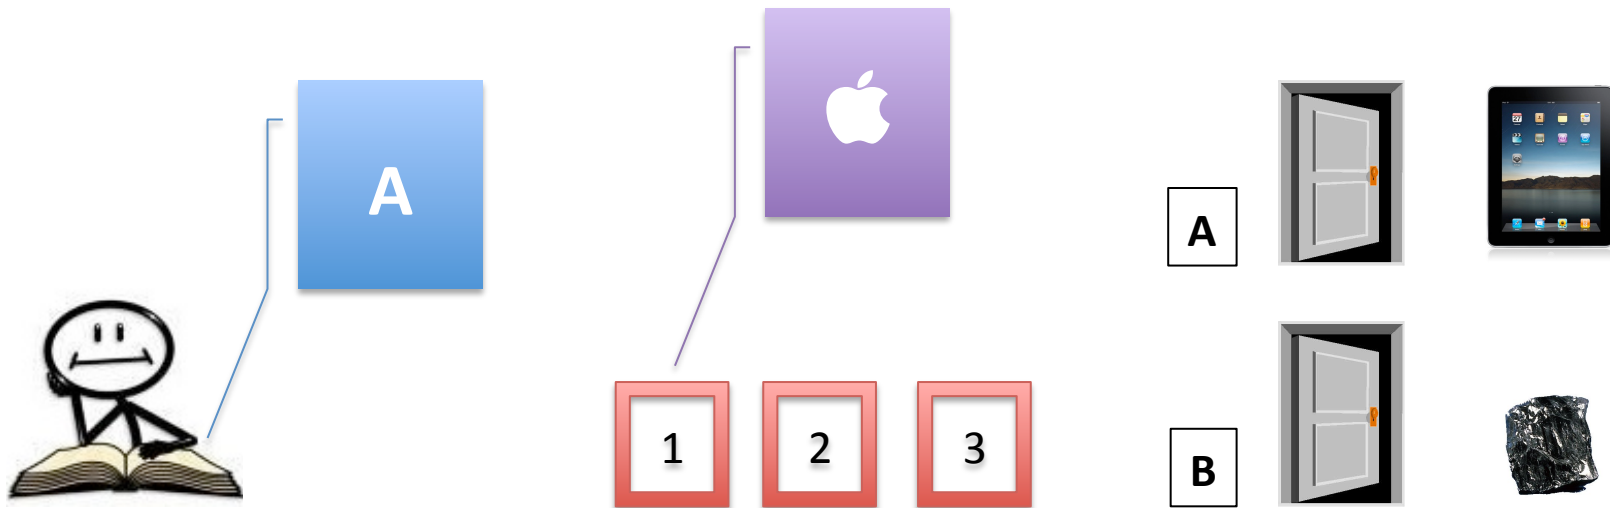
Number of possible values for each category of salient knowledge



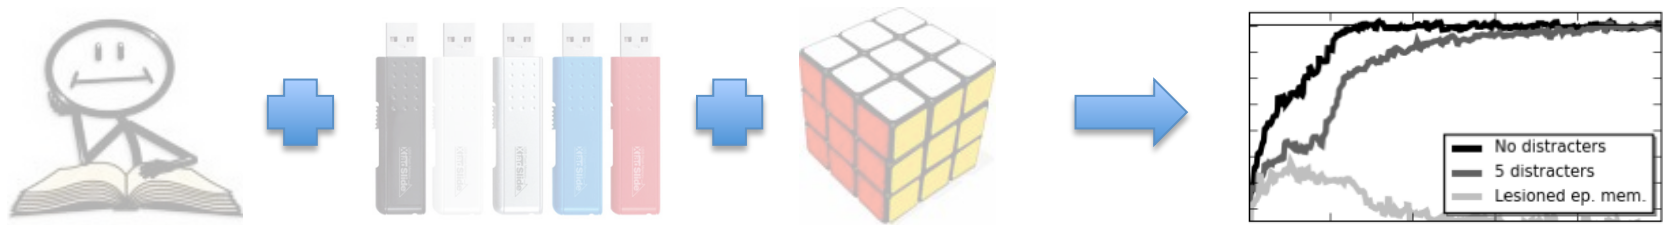
Quantity of Distracting Observations



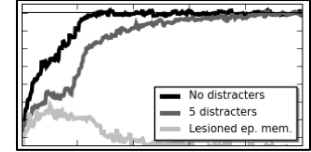
Number of observations irrelevant to decisions



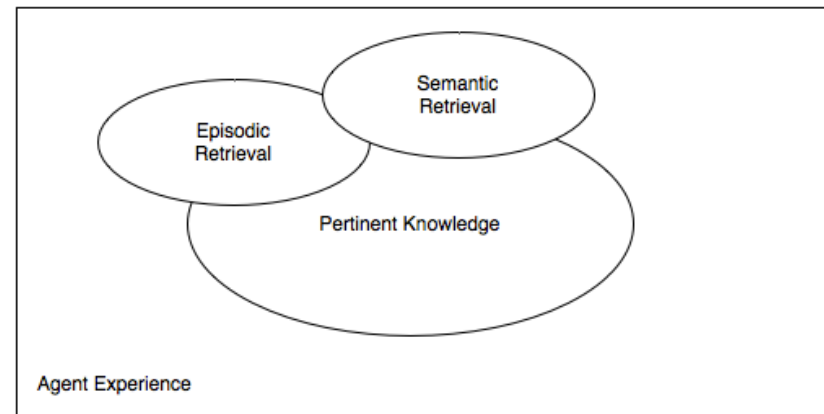
Progress: Evaluation Metrics



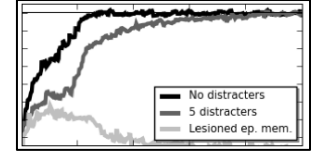
Evaluation: Quantitative



- Average reward-per-step
- Maximum reward
- Speed of convergence
- Amount of memory
- Knowledge Coverage



Evaluation: Qualitative



Aggregate

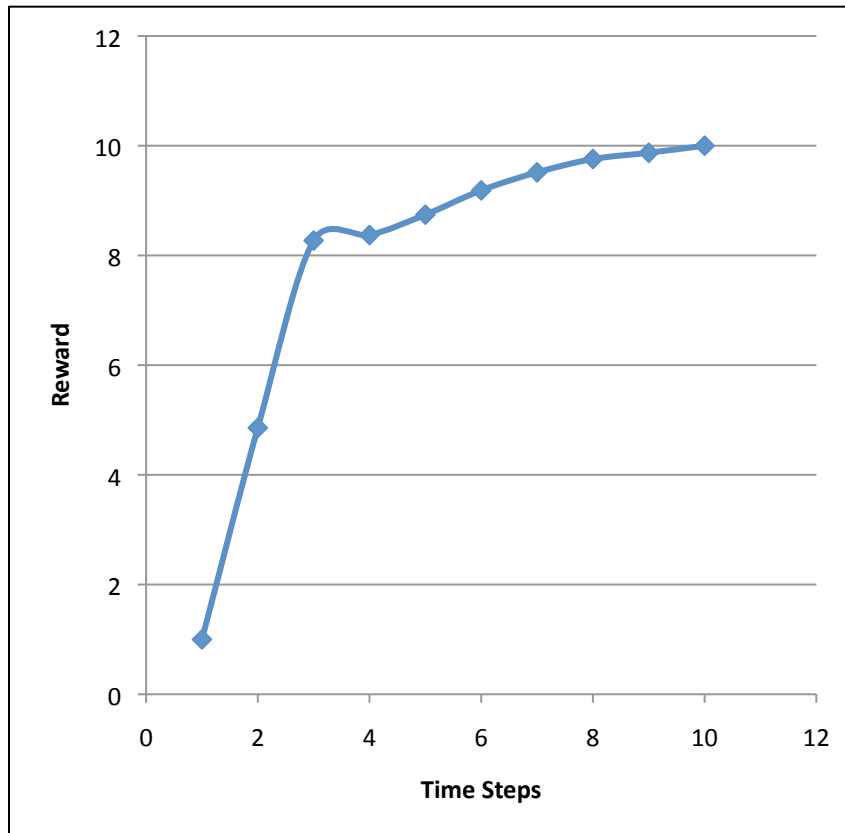
Compare agent performance across different points in memory space

Individual

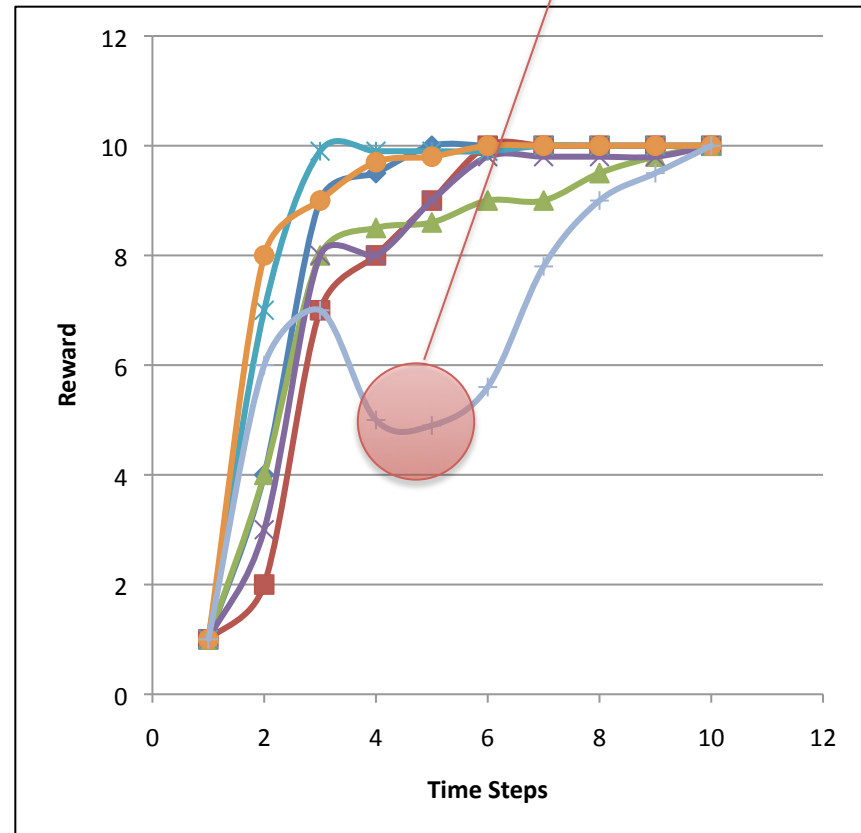
Categorizing behavior supported by memory in accordance with exhibited cognitive capabilities

Evaluation Example

Aggregate



Individual



Strategy Shift?

Example Studies

Episodic Memory

- Fix task
- Vary encoding policy as a function of activation, emotion, and other elements of architecture state

Learning Control #1

- Fix task
- Vary available memories: episodic, semantic, bit

Learning Control #2

- Vary task along temporal distance
- Fix available memories

Unresolved Issues

- Incomplete memory and task taxonomies
- Automatically instantiating computational memory systems and tasks
- Optimally adaptive agent
- Memory space is truly, truly vast
 - So how to guarantee sufficient coverage of space?

Thank You

Questions?