Episodic Memory
and the
Complementary Learning Systems (CLS)
Hypothesis

Classic Taxonomy

Squire & Zola-Morgan (1988)

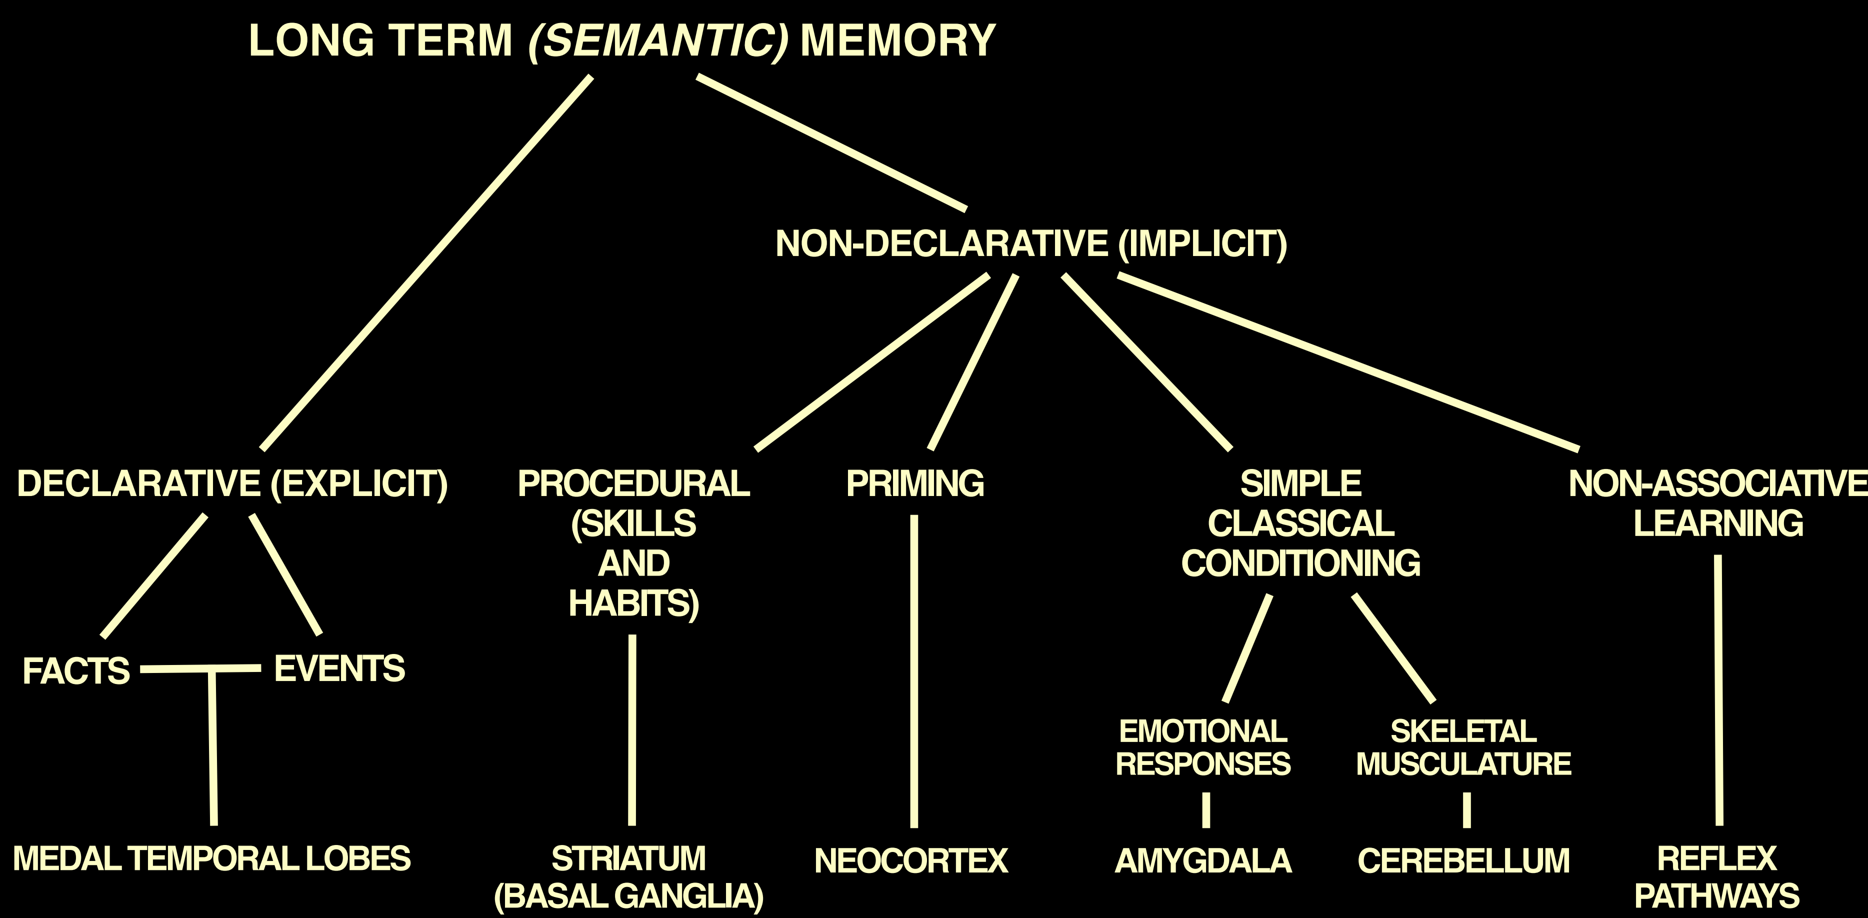

Classic Taxonomy

Squire & Zola-Morgan (1988)

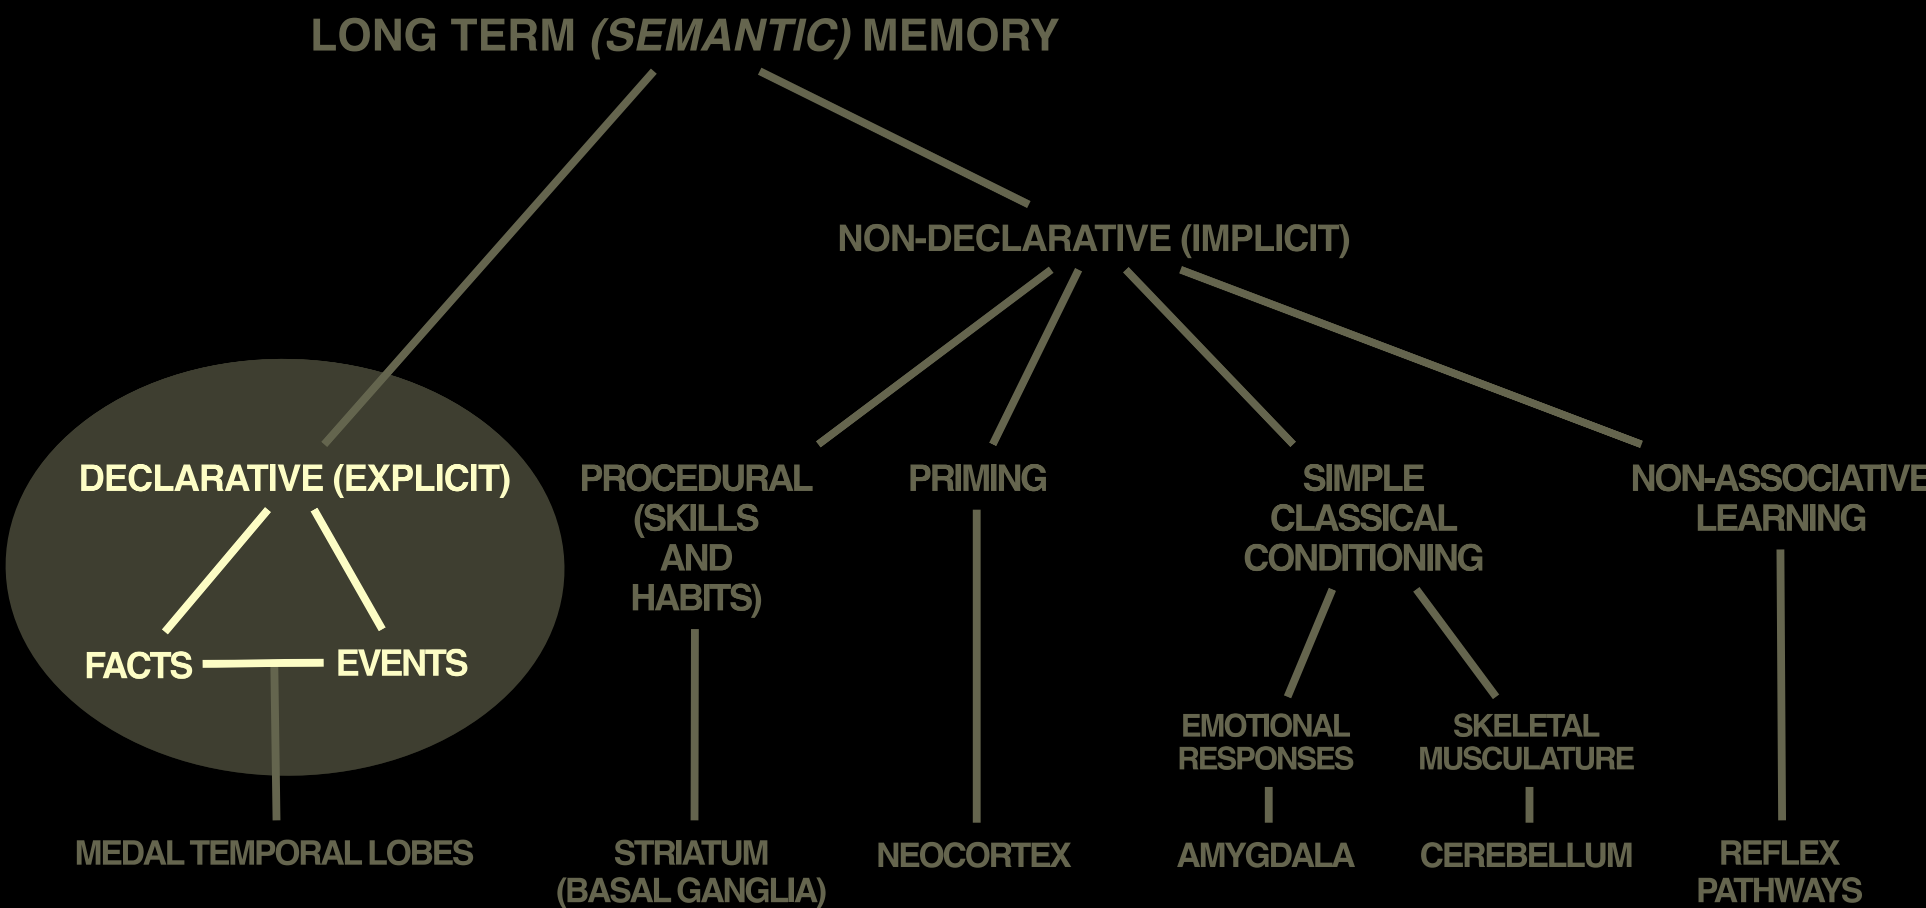

Memory

Memory

- **Memory = any persistent effect of experience**
  (not just memorization of facts, events, names, etc.)

- ***State-based*** **memory (active maint., "short term" memory)**

  vs.

- ***Weight-based*** **memory (long-term memory):**
  - **associative learning:**
    gradual, integrative cortical learning, and priming effects

Memory

- **Memory = any persistent effect of experience**
  (not just memorization of facts, events, names, etc.)

- ***State-based*** **memory (active maint., “short term” memory)**

  vs.

- ***Weight-based*** **memory (long-term memory):**

  - **associative learning:**
    gradual, integrative cortical learning, and priming effects

☞ - but what about one-shot learning / rapid memorization?

Memory

- **Memory = any persistent effect of experience**
  (not just memorization of facts, events, names, etc.)

- ***State-based* memory (active maint., "short term" memory)**

vs.

- ***Weight-based* memory (long-term memory):**

  - **associative learning:**
    gradual, integrative cortical learning, and priming effects

☞ - but what about one-shot learning / rapid memorization?

there's a fundamental problem…

Paired Associates (AB-AC) Learning

Paired Associates (AB-AC) Learning

- AB-AC paired-associates learning paradigm:

Paired Associates (AB-AC) Learning

– **first learn set of paired associates (AB):**

(AB)

window-reason
bicycle-garbage

Paired Associates (AB-AC) Learning

– then learn new associate for 1st member of each old pair (AC)…

(AC)

window-telephone
bicycle-desk

Paired Associates (AB-AC) Learning

– then test on both sets of associations:

(A)

window-

bicycle-

Paired Associates (AB-AC) Learning

– **then test on both sets of associations:**

(A)	***(B)***
window-	*?*
bicycle-	*?*

Paired Associates (AB-AC) Learning

– **then test on both sets of associations:**

	(A)	**(B)**
	window-	*reason*
	bicycle-	*garbage*

Paired Associates (AB-AC) Learning

– **then test on both sets of associations:**

(A)	***(C)***
window-	*?*
bicycle-	*?*

Paired Associates (AB-AC) Learning

– **then test on both sets of associations:**

	(A)	(C)
	window-	*telephone*
	bicycle-	*desk*

Paired Associates (AB-AC) Learning

- **Human performance:**
  - **"retroactive" interference:**
    get *some* loss of memory for AB association

Paired Associates (AB-AC) Learning

- Human performance:
  - “retroactive” interference:
    get *some* loss of memory for AB association
  - however, loss is modest and gradual (“graceful” degradation)…

Paired Associates (AB-AC) Learning

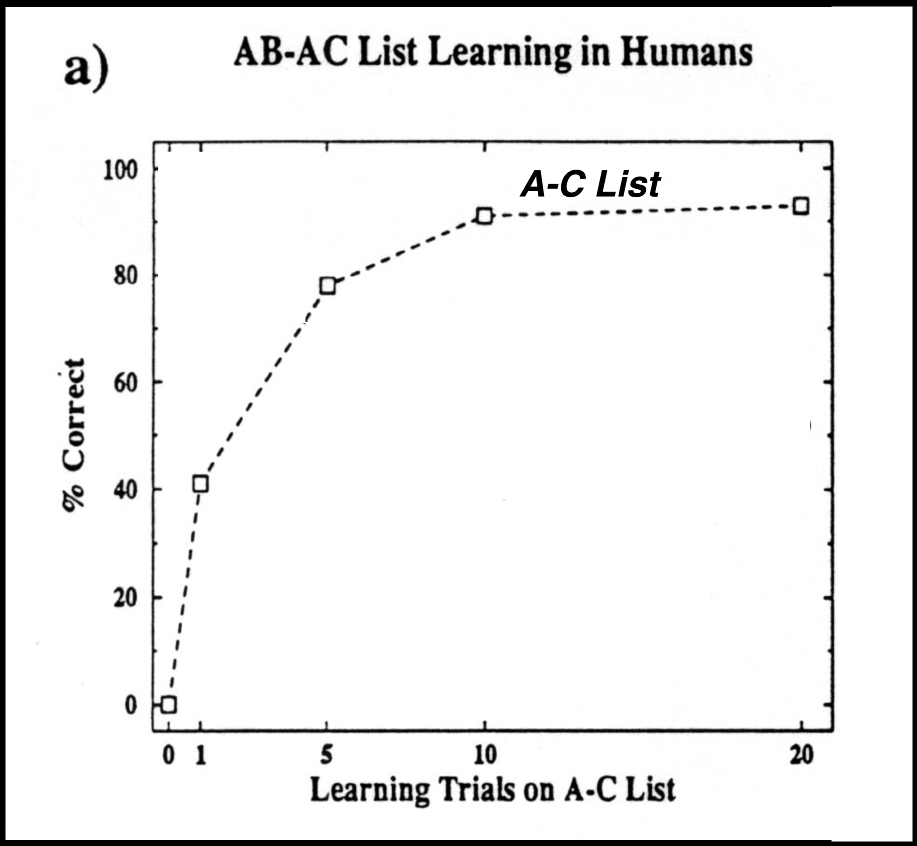

Paired Associates (AB-AC) Learning

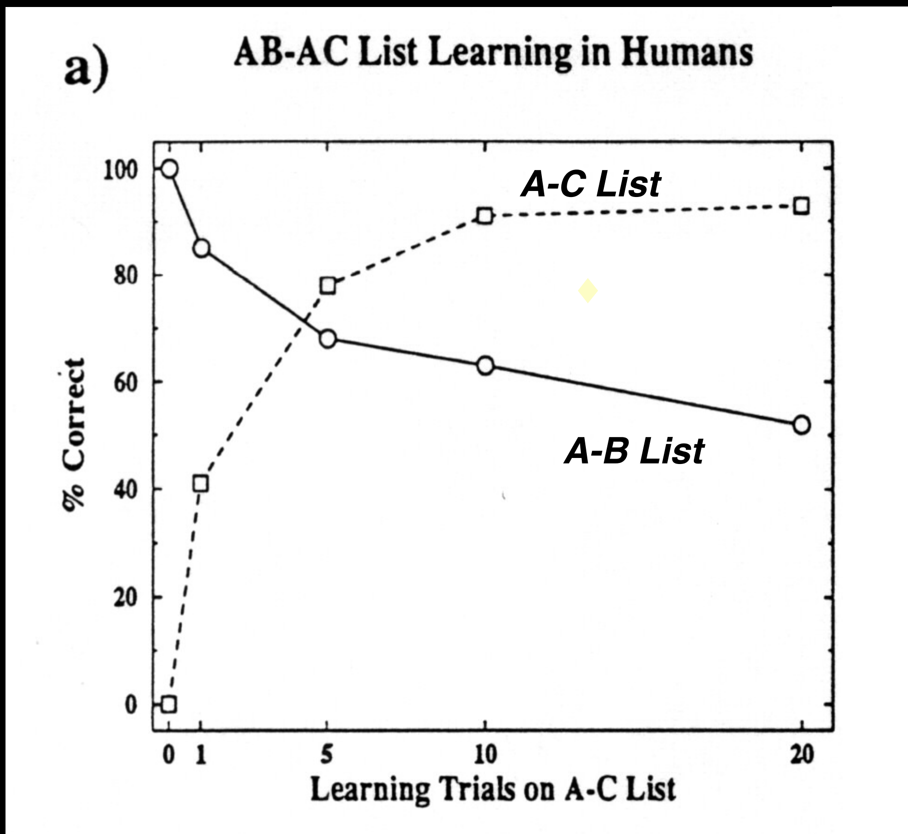

Model of AB-AC List Learning

McCloskey & NJ Cohen (1989)

- **Simple pattern associator**
  - **input:**
    - "A" stimuli
    - Context units (list label)
  - **output:**
    - "B" or "C" associate *(depending upon context)*
  - **trained with backprop**

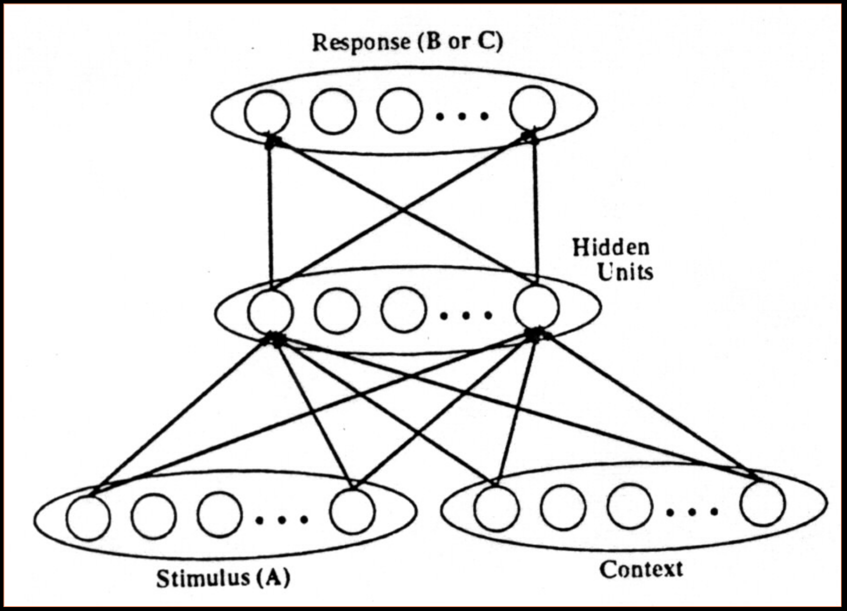

- **Model finding…**

Model of AB-AC List Learning

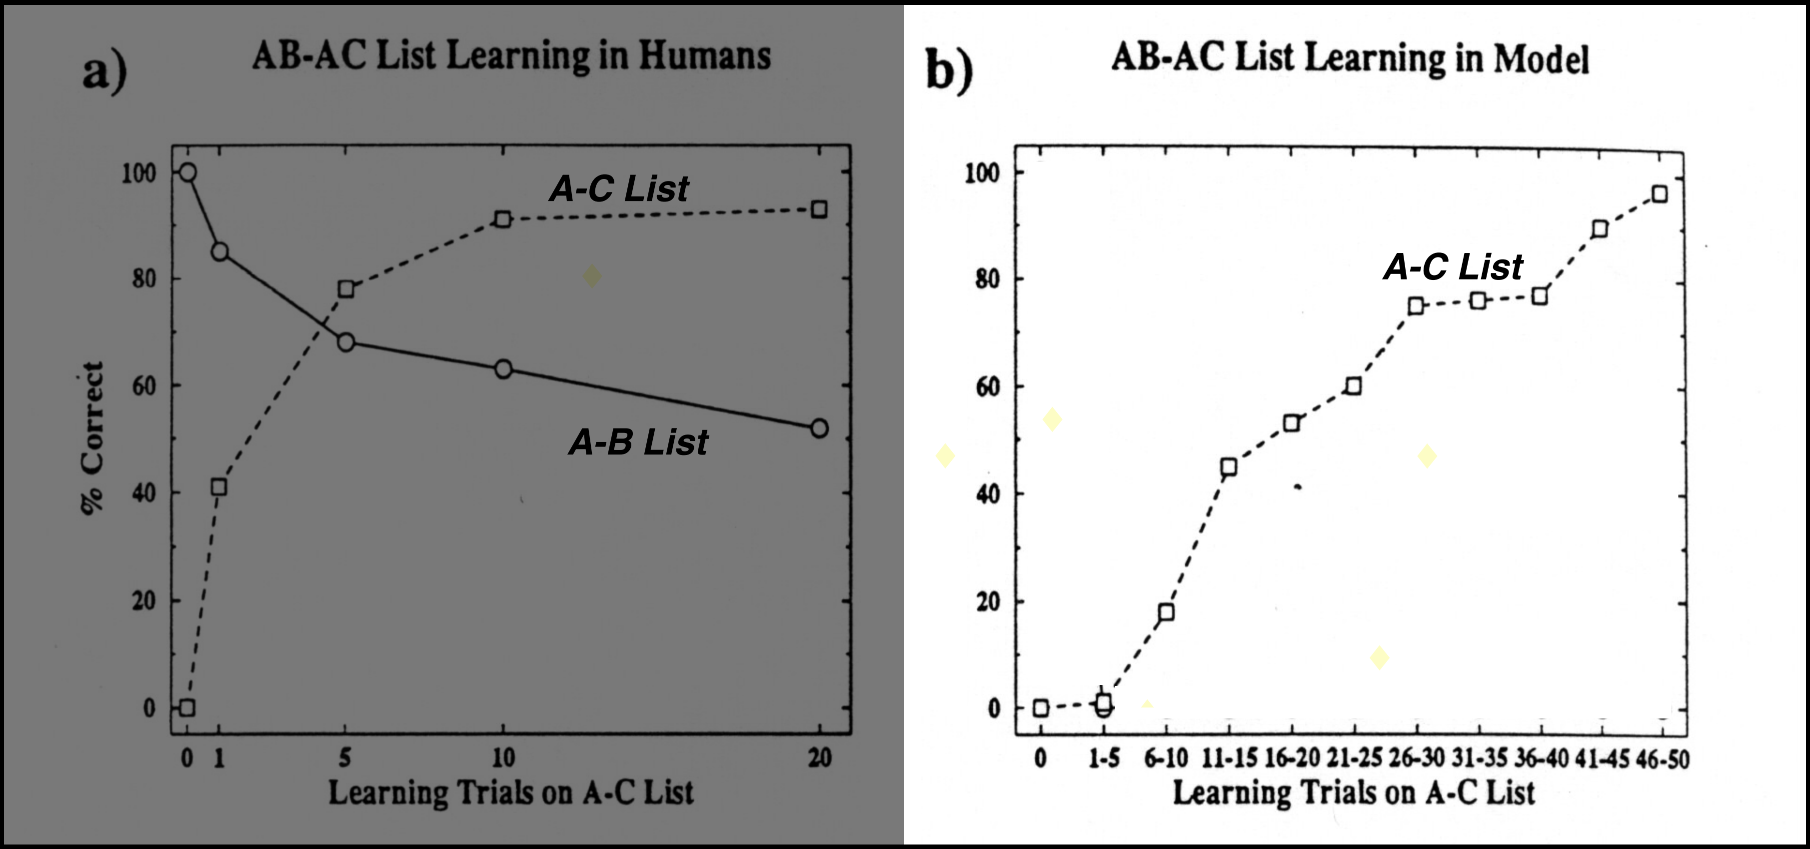

Model of AB-AC List Learning

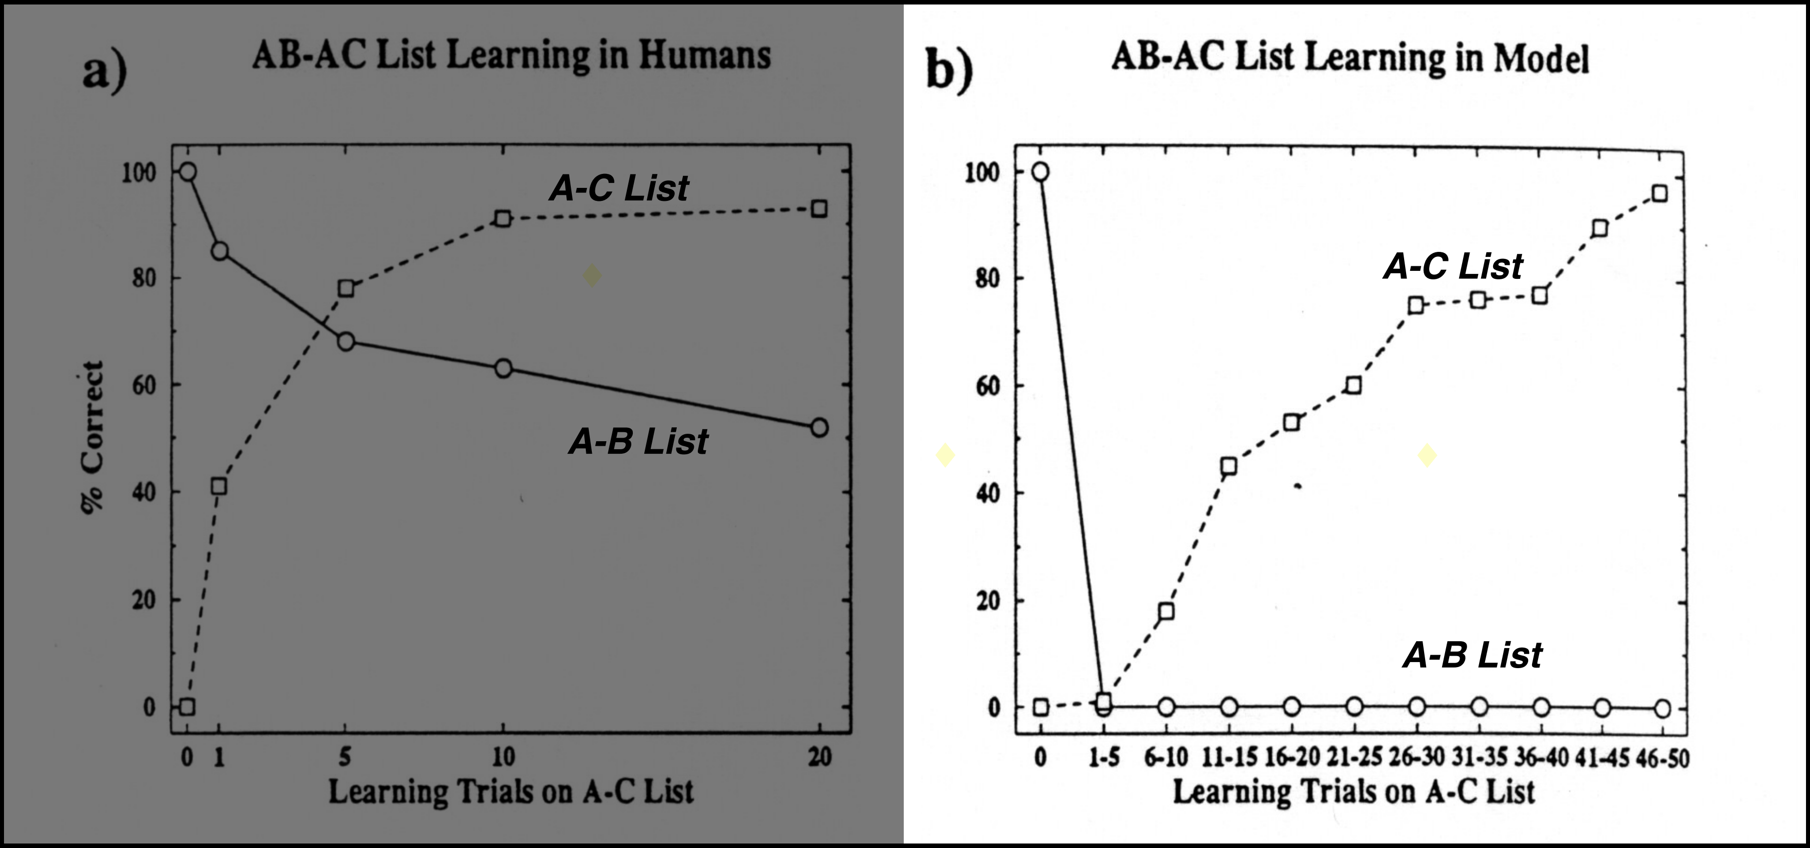

Model of AB-AC List Learning

- Learned first set of paired associates (AB) without any trouble

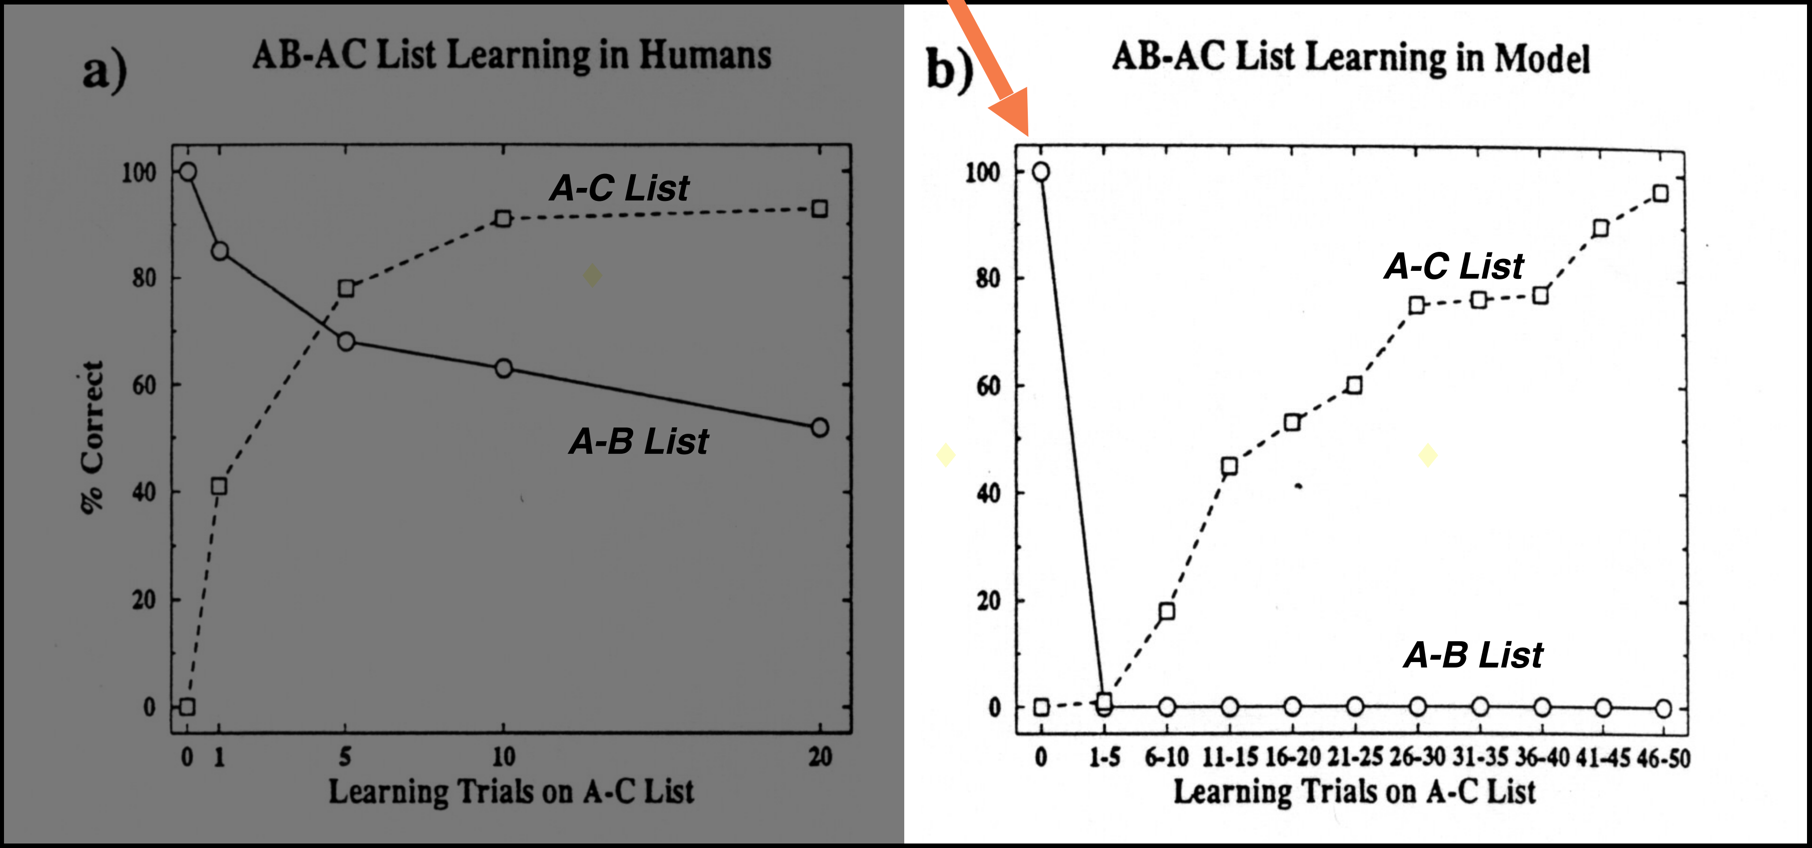

Model of AB-AC List Learning

- Took a bit longer to learn the second set (AC), but do could so pretty well; but…

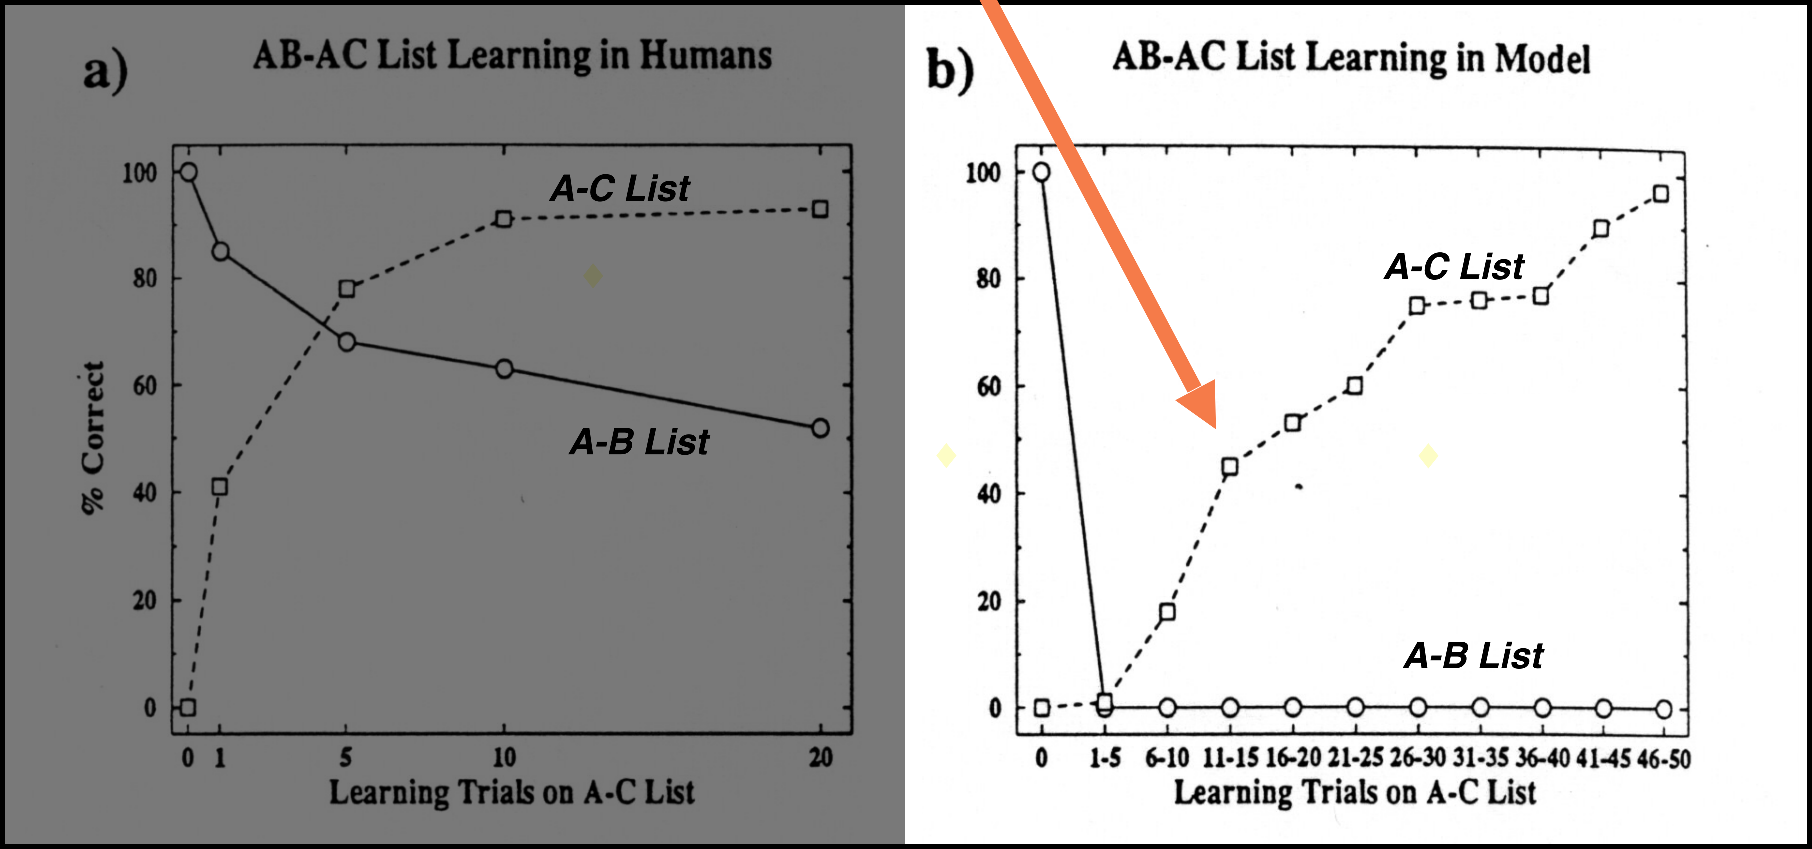

Model of AB-AC List Learning

- *Network loses AB association* even before it even begins to learn AC:
  Catastrophic interference!

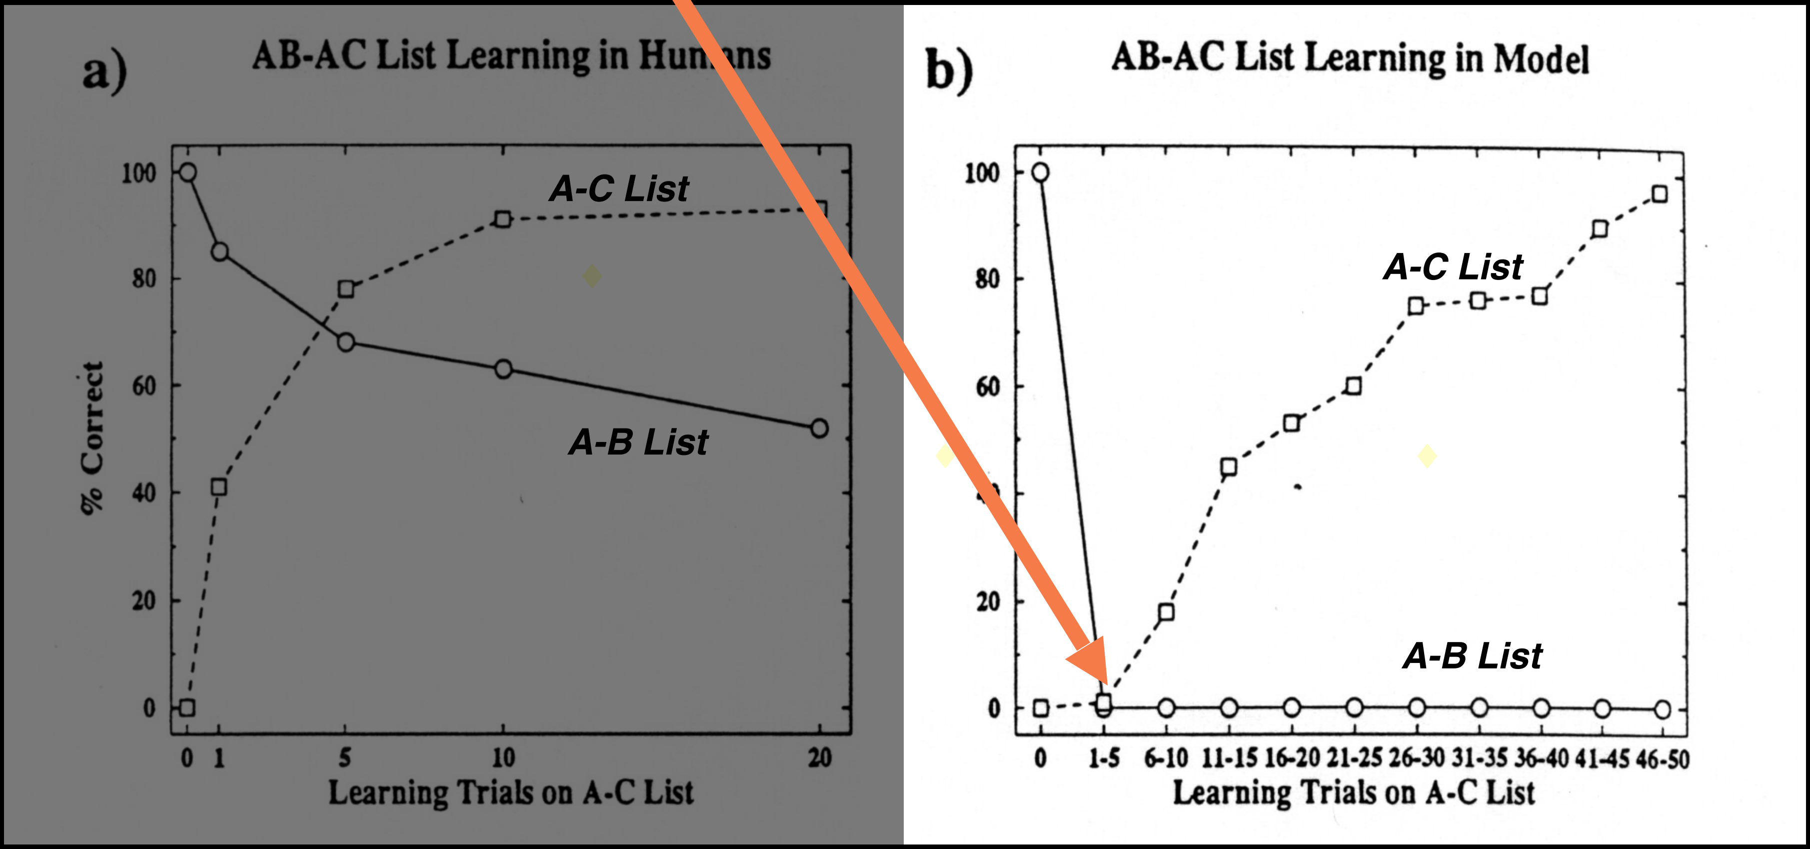

Episodic vs. Semantic Memory

Episodic vs. Semantic Memory

- **Distinction between:**
  - **memory for *actual details* of an item or event *("episodes")***
    - ♦ *acquired quickly;*

Episodic vs. Semantic Memory

- **Distinction between:**
  - **memory for *actual details* of an item or event *("episodes")***
    - *acquired quickly;* *highly specific,*

Episodic vs. Semantic Memory

- **Distinction between:**
  - **memory for *actual details* of an item or event *("episodes")***
    - *acquired quickly; highly specific, good for identification*

Episodic vs. Semantic Memory

- **Distinction between:**
  - **memory for *actual details* of an item or event *("episodes")***
    - ♦ *acquired quickly; highly specific, good for identification*
  - **its meaning, significance, and/or what it is related to *("semantics")***
    - ♦ *acquired slowly;*

Episodic vs. Semantic Memory

- **Distinction between:**
  - **memory for *actual details* of an item or event *("episodes")***
    - *acquired quickly; highly specific, good for identification*
  - **its meaning, significance, and/or what it is related to *("semantics")***
    - *acquired slowly; more abstract,*

Episodic vs. Semantic Memory

- **Distinction between:**
  - **memory for *actual details* of an item or event *("episodes")***
    - *acquired quickly; highly specific, good for identification*
  - **its meaning, significance, and/or what it is related to *("semantics")***
    - *acquired slowly; more abstract, good for generalization*

Episodic vs. Semantic Memory

- Distinction between:
  - memory for *actual details* of an item or event *("episodes")*
    - ♦ *acquired quickly; highly specific, good for identification*
  - its meaning, significance, and/or what it is related to *("semantics")*
    - ♦ *acquired slowly; more abstract, good for generalization*
- Sometimes we remember actual details quite well
  - "flashbulb memories": where were you on Nov. 5, 2024?

Episodic vs. Semantic Memory

- **Distinction between:**
  - **memory for *actual details* of an item or event *("episodes")***
    - *acquired quickly; highly specific, good for identification*
  - **its meaning, significance, and/or what it is related to *("semantics")***
    - *acquired slowly; more abstract, good for generalization*
- **Sometimes we remember actual details quite well**
  - **"flashbulb memories": where were you on Nov. 5, 2024?**
  - **even when they are not very important**

    (anyone remember the word I asked you to spell last class?)

Episodic vs. Semantic Memory

- **Distinction between:**
  - **memory for *actual details* of an item or event *("episodes")***
    - *acquired quickly; highly specific, good for identification*
  - **its meaning, significance, and/or what it is related to *("semantics")***
    - *acquired slowly; more abstract, good for generalization*
- **Sometimes we remember actual details quite well**
  - **"flashbulb memories": where were you on Nov. 5, 2024?**
  - **even when they are not very important**

    (anyone remember the word I asked you to spell last class?)
- **But *usually* not for very long** *(maybe days, but not months…)*

Episodic vs. Semantic Distinction

Episodic vs. Semantic Distinction

- **Episodic** *(Tulving, 1972)*:
  - **memory for events and/or details: "episodes"**
  - **"one trial learning"**

Episodic vs. Semantic Distinction

- **Episodic** *(Tulving, 1972)*:
  - **memory for events and/or details: "episodes"**
  - **"one trial learning"**
  - **examples:**
    - *where did you park your car today?*

Episodic vs. Semantic Distinction

- **Episodic** *(Tulving, 1972)*:
  - **memory for events and/or details: "episodes"**
  - **"one trial learning"**
  - **examples:**
    - *where did you park your car today?*
    - *a face you have seen only once…*

Observe these faces…

Which one did you see before?

Episodic vs. Semantic Distinction

Episodic vs. Semantic Distinction

- Episodic:
  - memory for events and/or details: "episodes"
  - "one trial learning"
  - examples:
    - *where did you park your car today?*
    - *a face you have seen only one…*
- *vs.* Semantic:
  – general knowledge, relationships: "meaning"

Episodic vs. Semantic Distinction

- **Episodic:**
  - **memory for events and/or details: "episodes"**
  - **"one trial learning"**
  - **examples:**
    - *where did you park your car today?*
    - *a face you have seen only one…*
- *vs.* **Semantic:**
  - **general knowledge, relationships: "meaning"**
  - **experience, repeated learning**

Episodic vs. Semantic Distinction

- **Episodic:**
  - **memory for events and/or details: "episodes"**
  - **"one trial learning"**
  - **examples:**
    - *where did you park your car today?*
    - *a face you have seen only one…*

- *vs.* **Semantic:**
  - **general knowledge, relationships: "meaning"**
  - **experience, repeated learning**
  - **examples:**
    - *where is the most popular place to park?*
    - *meanings of words…*

Observe these words…

Respond yes if you saw *exactly* the following words…

rats

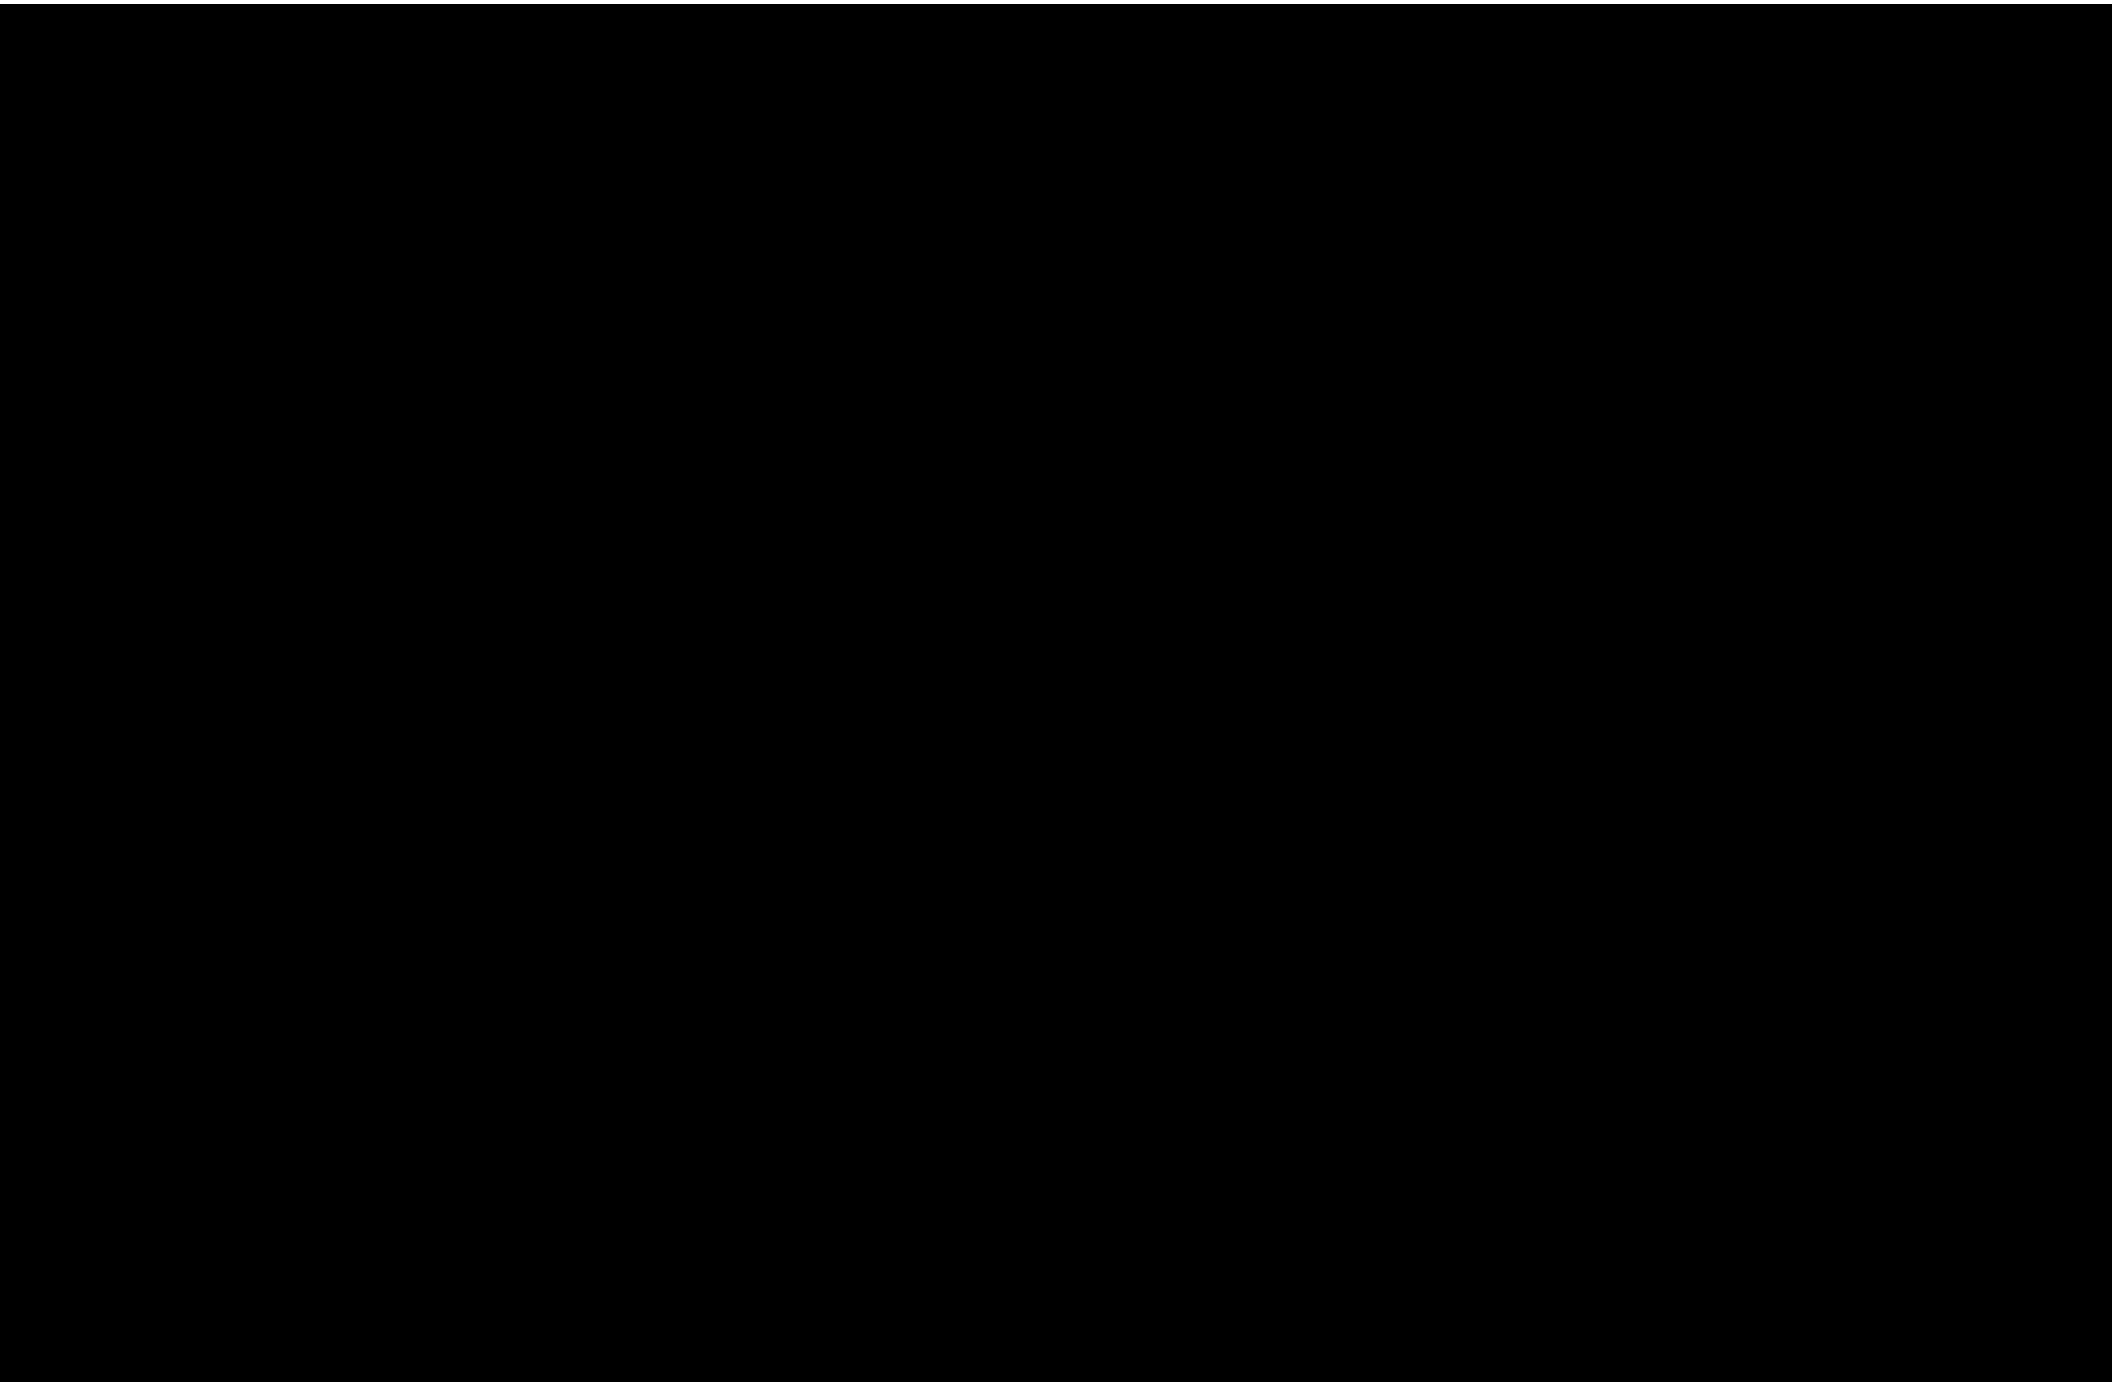

Episodic vs. Semantic Memory

Episodic vs. Semantic Memory

- **Task** *(Hintzman et al., 1992)*:
  - Study: house… rat… Nose… bee… friend… Door
  - Test: rats… bat… Bee…

Episodic vs. Semantic Memory

- **Task** *(Hintzman et al., 1992)*:
  - **Study: house… rat… Nose… bee… friend… Door**
  - **Test: rats… bat… Bee…**

- **Semantic memory:**
  - **distinguished:**
    - √ **rats vs. rat** → *different meaning:* ***semantic***

Episodic vs. Semantic Memory

- **Task** *(Hintzman et al., 1992)***:**
  - **Study: house… rat… Nose… bee… friend… Door**
  - **Test: rats… bat… Bee…**

- **Semantic memory:**
  - **distinguished:**
    - √ **rats vs. rat** → *different meaning:* ***semantic***
    - x **bee vs. Bee** → *same meaning:* ***episodic***

Episodic vs. Semantic Memory

- **Task** *(Hintzman et al., 1992)*:
  - **Study: house… rat… Nose… bee… friend… Door**
  - **Test: rats… bat… Bee…**

- **Semantic memory:**
  - **distinguished:**
    - √ **rats vs. rat** → *different meaning:* ***semantic***
    - x **bee vs. Bee** → *same meaning:* ***episodic***
  - **"rats" seemed familiar, but "bat" did not**

Episodic vs. Semantic Memory

- **Task** *(Hintzman et al., 1992)*:
  - Study: house… rat… Nose… bee… friend… Door
  - Test: rats… bat… Bee…
- **Semantic memory:**
  - distinguished:
    - √ rats vs. rat → *different meaning:* ***semantic***
    - x bee vs. Bee → *same meaning:* ***episodic***
  - "rats" seemed familiar, but "bat" did not
- **Episodic memory:**
  - "rats" could be rejected even though it was *familiar*

Episodic vs. Semantic Memory

- **Task** *(Hintzman et al., 1992)***:**
  - Study: house… rat… Nose… bee… friend… Door
  - Test: rats… bat… Bee…

- **Semantic memory:**
  - distinguished:
    - √ rats vs. rat → *different meaning:* ***semantic***
    - x bee vs. Bee → *same meaning:* ***episodic***
  - "rats" seemed familiar, but "bat" did not

- **Episodic memory:**
  - "rats" could be rejected even though it was *familiar*
    - familiarity → *semantic*

Episodic vs. Semantic Memory

- **Task** *(Hintzman et al., 1992)*:
  - Study: house… rat… Nose… bee… friend… Door
  - Test: rats… bat… Bee…
- **Semantic memory:**
  - distinguished:
    - √ rats vs. rat → *different meaning: **semantic***
    - x bee vs. Bee → *same meaning: **episodic***
  - "rats" seemed familiar, but "bat" did not
- **Episodic memory:**
  - "rats" could be rejected even though it was *familiar*
    - familiarity → ***semantic***
    - rejection → ***episodic***

A Problem

A Problem

- On the one hand:
  - it takes many exposures for the learning mechanisms we've considered so far to develop meaningful, reliable representations

A Problem

- **On the one hand:**
  - **it takes many exposures for the learning mechanisms we've considered so far to develop meaningful, reliable representations**
  - **they are *statistical:***
    - *where do you usually park your car?*
    - *what do birds have in common?*

A Problem

- **On the one hand:**
  - **it takes many exposures for the learning mechanisms we've considered so far to develop meaningful, reliable representations**
  - **they are *statistical:***
    - *where do you usually park your car?*
    - *what do birds have in common?*
- **On the other hand:**
  - **we *can* learn from / remember *single items* quickly, even "one-shot":**
    - *where did you park your car or bike today*
    - *which bird was it that I just saw?*

A Problem

- On the one hand:
  - it takes many exposures for the learning mechanisms we've considered so far to develop meaningful, reliable representations
  - they are ***statistical:***
    - *where do you usually park your car?*
    - *what do birds have in common?*
- On the other hand:
  - **we *can* learn from / remember *single items* quickly, even "one-shot":**
    - *where did you park your car or bike today*
    - *which bird was it that I just saw?*
  - **we can even quickly learn *exceptions* to statistical regularities:**
    - *a penguin is a bird that can't fly…*

A Problem

A Problem

- Solution?
  - turn up the learning rate to get single trial learning: episodic memory

A Problem

- Solution?
  - turn up the learning rate to get single trial learning: episodic memory
  - but that would defeat the purpose of extracting consistent (statistical) relationships from experience: semantic memory

A Problem

- Solution?
  - turn up the learning rate to get single trial learning: episodic memory
  - but that would defeat the purpose of extracting consistent (statistical) relationships from experience: semantic memory
  - and it is subject to catastrophic interference (remember AB-AC problem)

A Problem

- Solution?
  - turn up the learning rate to get single trial learning: episodic memory
  - but that would defeat the purpose of extracting consistent (statistical) relationships from experience: semantic memory
  - and it is subject to catastrophic interference (remember AB-AC problem)
- Another potential solution:
  - interleaved training…

Interference in Semantic Networks

(McClelland, McNaughton & O'Reilly 1995)

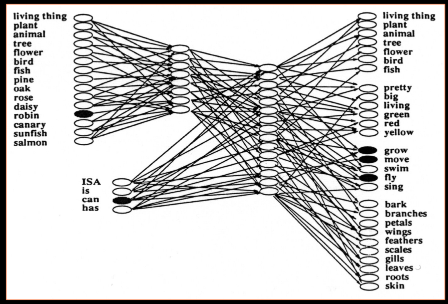

Interference in Semantic Networks

(McClelland, McNaughton & O'Reilly 1995)

- **Train network on new piece of inconsistent knowledge:**
  - **– penguin: bird that swims but doesn't fly**

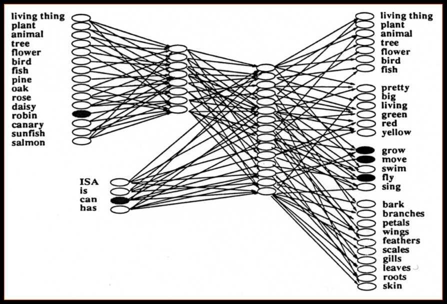

Interference in Semantic Networks

(McClelland, McNaughton & O'Reilly 1995)

- Train network on new piece of inconsistent knowledge:
  - penguin: bird that swims but doesn't fly

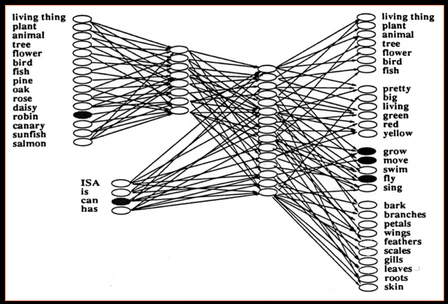

- *Focused* Training:
  - network learns new information quickly

Interference in Semantic Networks

(McClelland, McNaughton & O'Reilly 1995)

- **Train network on new piece of inconsistent knowledge:**
  - **– penguin: bird that swims but doesn't fly**

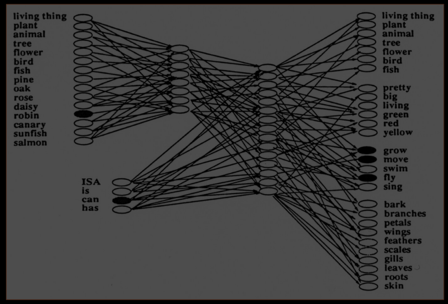

- ***Focused* Training:**
  - **– network learns new information quickly**
  - **– however, it interferes with old knowledge: catastrophic interference**

    begins to think all birds swim and don't fly!

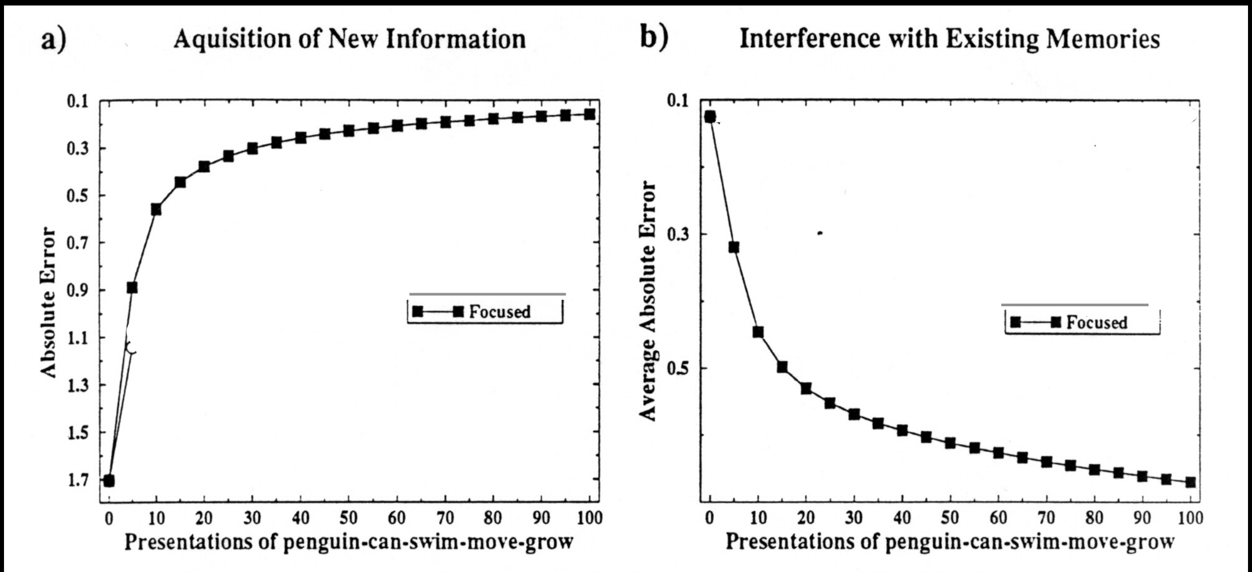

Interference in Semantic Networks

(McClelland, McNaughton & O'Reilly 1995)

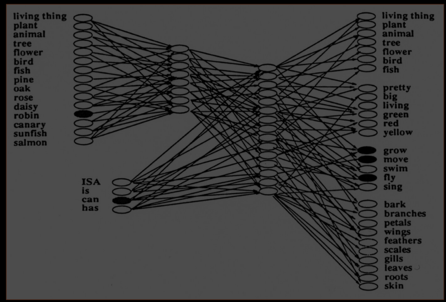

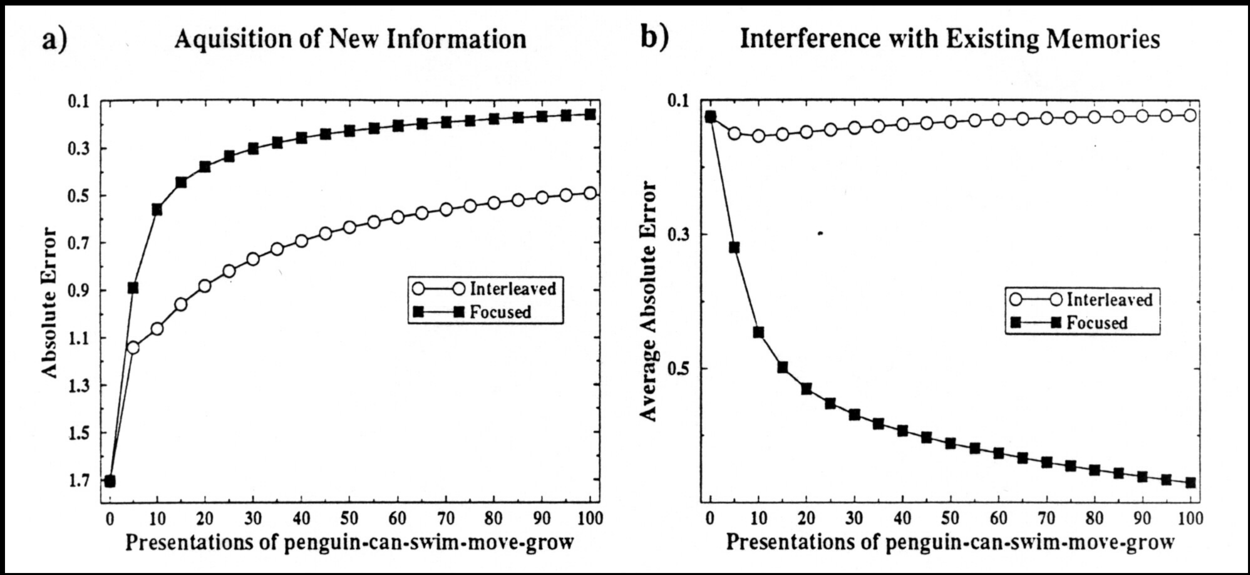

Interference in Semantic Networks

(McClelland, McNaughton & O'Reilly 1995)

- Train network on new piece of inconsistent knowledge:
  - penguin: bird that swims but doesn't fly

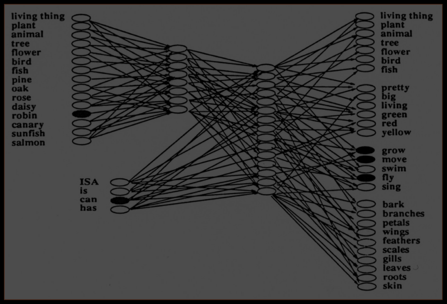

- ***Interleaved* Training:**
  - **new information can be learned without losing old information *(avoids catastrophic interference)***
    "carves out space" for penguin without disturbing other birds

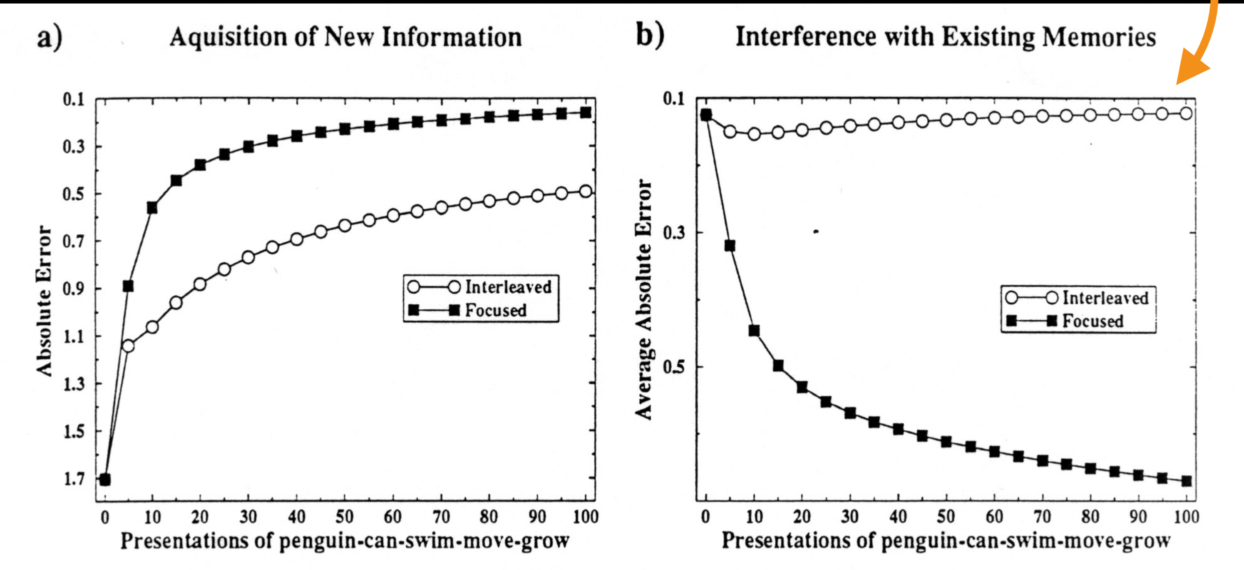

Interference in Semantic Networks

(McClelland, McNaughton & O'Reilly 1995)

- **Train network on new piece of inconsistent knowledge:**
  - **penguin: bird that swims but doesn't fly**

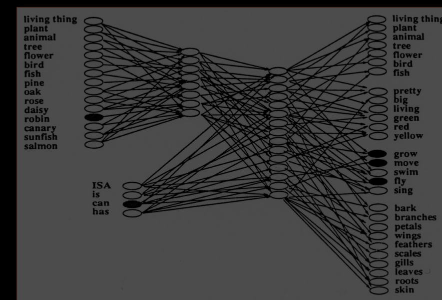

- ***Interleaved* Training:**
  - **new information can be learned without losing old information *(avoids catastrophic interference)***
    - *"carves out space" for penguin without disturbing other birds*
  - **however, still can't explain rapid (one-shot) learning**

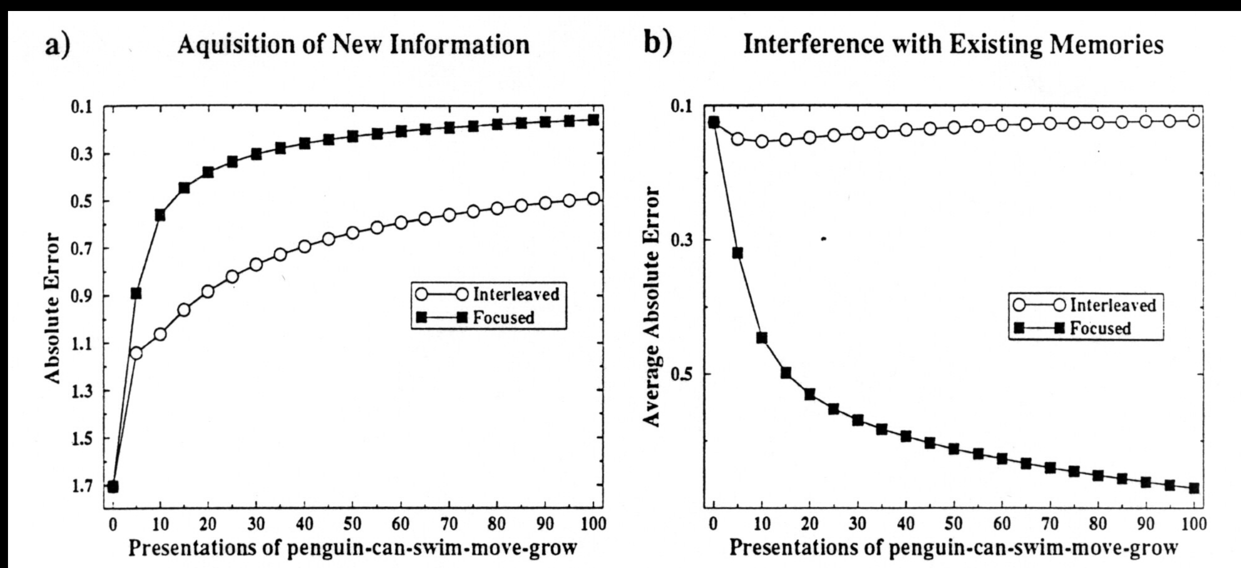

Interference in Semantic Networks

Interference in Semantic Networks

- **Fundamental tradeoff:**
  - **learn slow *(semantic)***
    - **acquire "statistical" knowledge** *(consistent relationships)*

Interference in Semantic Networks

- **Fundamental tradeoff:**
  - **learn slow** ***(semantic)***
    - **acquire "statistical" knowledge** *(consistent relationships)*
    - **then can't learn single events** *(no episodic memory)*

Interference in Semantic Networks

- Fundamental tradeoff:
  - learn slow *(semantic)*
    - acquire "statistical" knowledge *(consistent relationships)*
    - then can't learn single events *(no episodic memory)*
  - learn fast *(episodic)*
    - store individual events

Interference in Semantic Networks

- **Fundamental tradeoff:**
  - **learn slow *(semantic)***
    - **acquire "statistical" knowledge** *(consistent relationships)*
    - **then can't learn single events** *(no episodic memory)*
  - **learn fast *(episodic)***
    - **store individual events**
    - **but then lose old stuff (catastrophic interference)**

Interference in Semantic Networks

- **Fundamental tradeoff:**
  - **learn slow** ***(semantic)***
    - **acquire "statistical" knowledge** *(consistent relationships)*
    - **then can't learn single events** *(no episodic memory)*
  - **learn fast** ***(episodic)***
    - **store individual events**
    - **but then lose old stuff (catastrophic interference)**

- **Another possible solution to the tradeoff:**
  - ***separate patterns*** **(distinct episodic memories)**

Interference in Semantic Networks

- **Fundamental tradeoff:**
  - **learn slow *(semantic)***
    - **acquire "statistical" knowledge** *(consistent relationships)*
    - **then can't learn single events** *(no episodic memory)*
  - **learn fast *(episodic)***
    - **store individual events**
    - **but then lose old stuff (catastrophic interference)**

- **Another possible solution to the tradeoff:**
  - ***separate patterns* (distinct episodic memories)**
  - **however, then would lose benefit of *shared structure* (semantic knowledge)**

Two Incompatible Goals

Remember Specifics	**Extract Generalities**
e.g., Where did I park today?	*e.g., Where do I usually park?*
Goal: Avoid interference	**Goal: Accumulate experience**

Two Incompatible Goals

Remember Specifics	Extract Generalities
e.g., Where did I park today?	*e.g., Where do I usually park?*
Goal: Avoid interference	**Goal: Accumulate experience**
Solution:	**Solution:**
separated **representations** *(e.g., keep days separate)*	***overlapping*** **representations** *(integrate across days)*

Division of Labor

Remember Specifics	Extract Generalities
e.g., Where did I park today? **Goal: Avoid interference**	*e.g., Where do I usually park?* **Goal: Accumulate experience**
Solution: **separate representations (keep days separate)** 	**Solution:** **overlapping representations (integrate across days)** 
System: ***Hippocampus***	**System:** ***Neocortex***

The Hippocampus

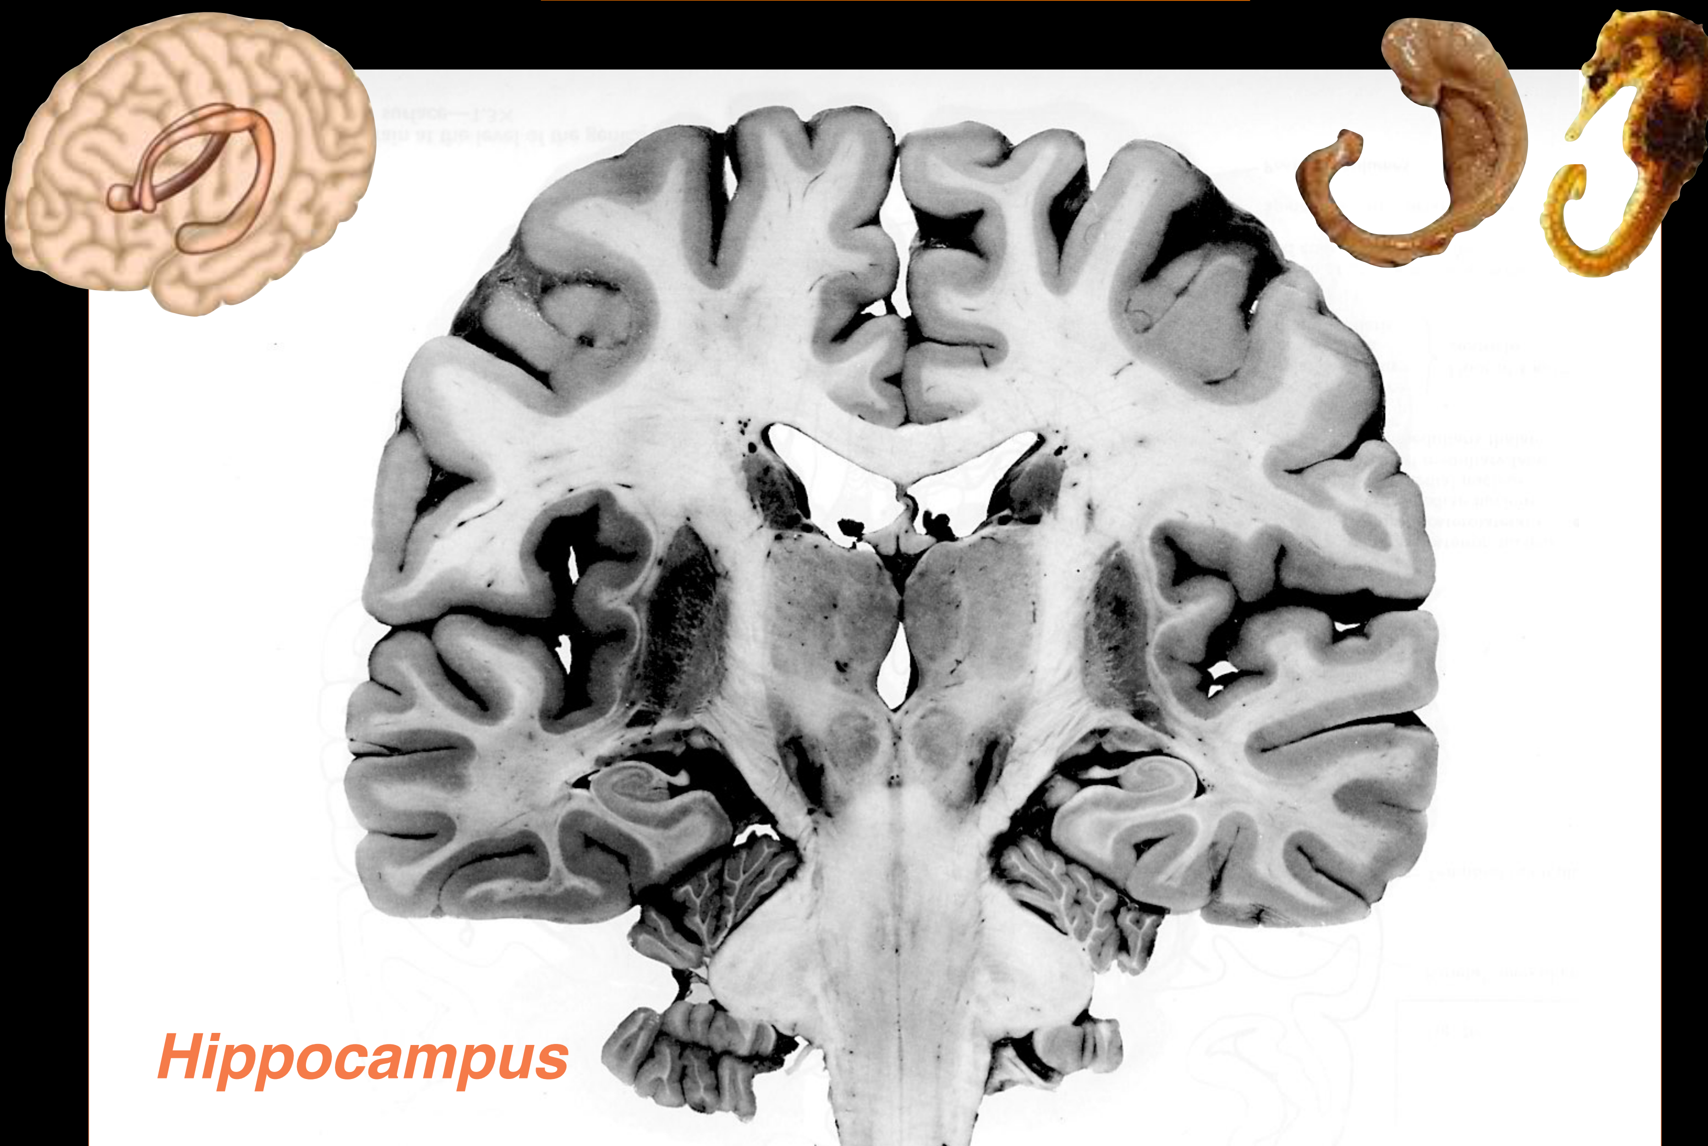

The Hippocampus

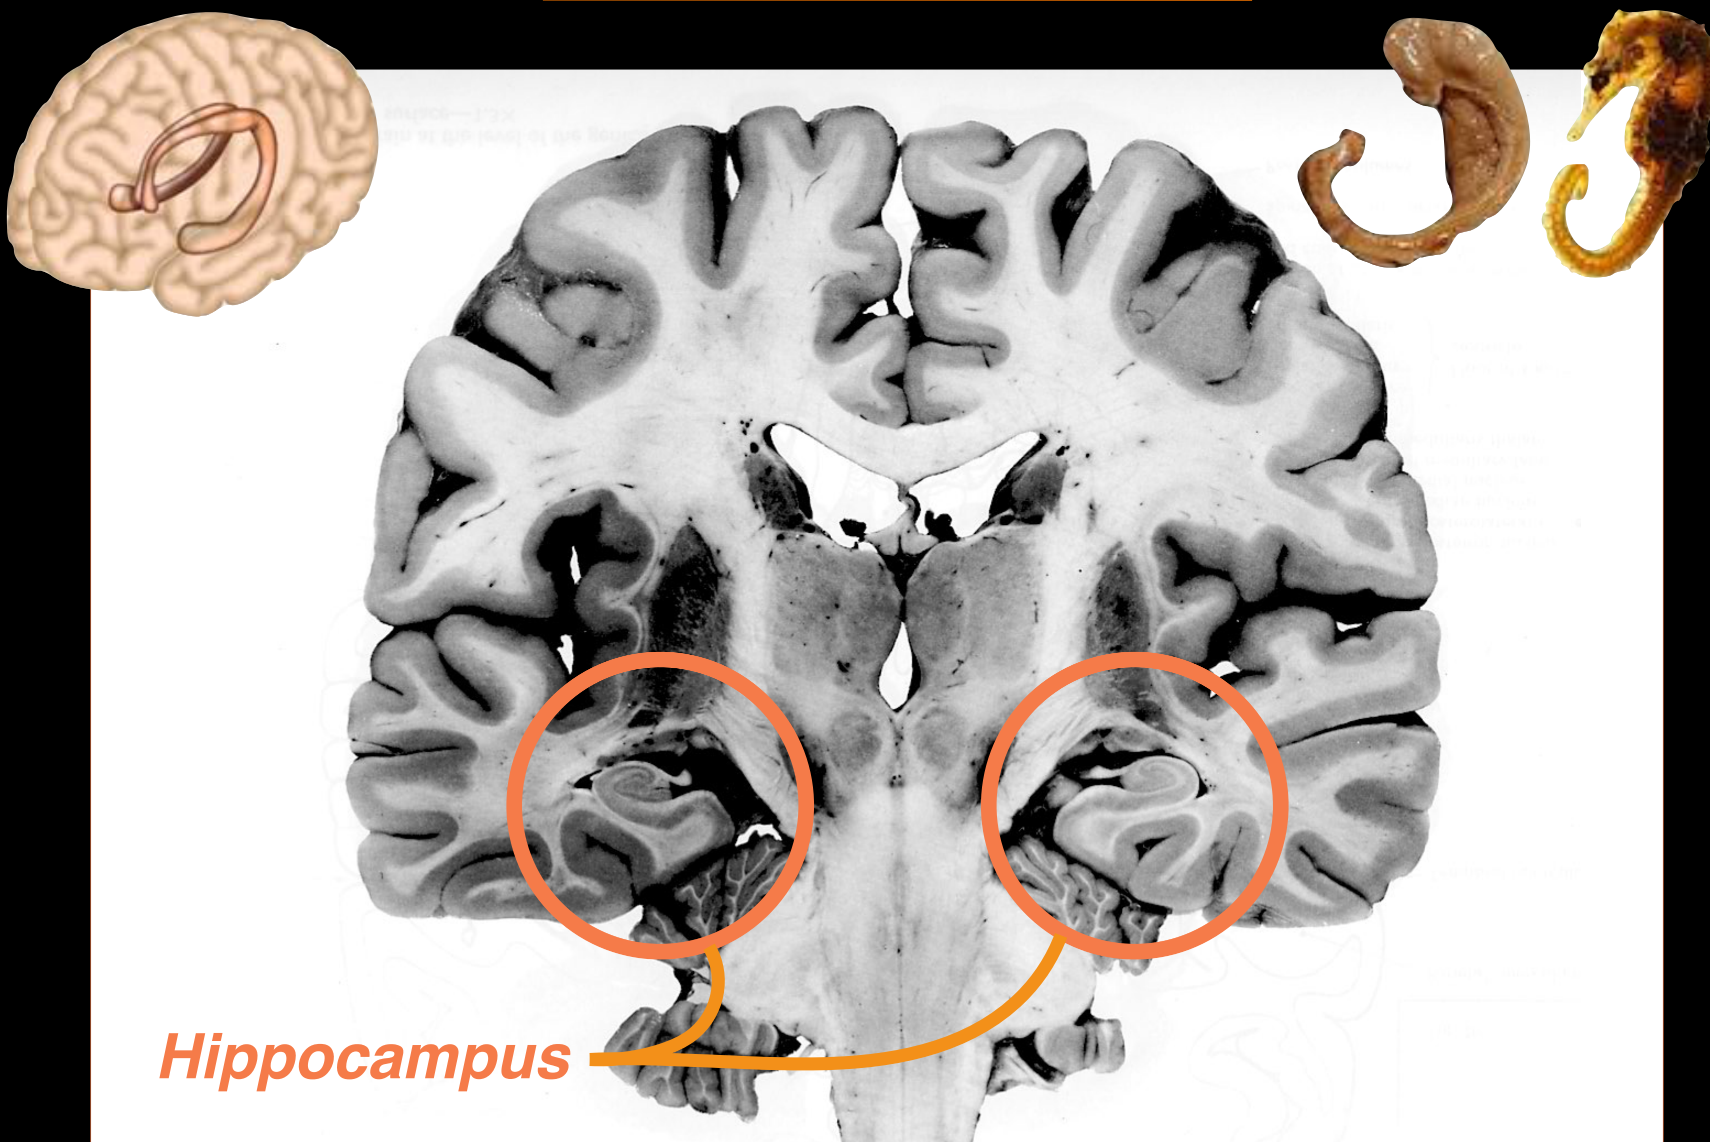

Hippocampus and Episodic Memory

Hippocampus and Episodic Memory

- **Anterograde amnesia** *(Scoville & Milner, 1957)*
  - **HM** *hippocampus removed to treat intractable epilepsy*

    impaired explicit / *episodic* memory

    paired associates

    one trial learning: recognition of a new face

Hippocampus and Episodic Memory

- **Anterograde amnesia** *(Scoville & Milner, 1957)*
  - **HM** *hippocampus removed to treat intractable epilepsy*

    impaired explicit / *episodic* memory

    paired associates

    one trial learning: recognition of a new face

    spared implicit / *semantic* memory

    category learning

    priming

    skill learning

    familiarity judgements

Hippocampus and Episodic Memory

- **Anterograde amnesia** *(Scoville & Milner, 1957)*
  - **HM** *hippocampus removed to treat intractable epilepsy*

    impaired explicit / *episodic* memory

    paired associates

    one trial learning: recognition of a new face

    spared implicit / *semantic* memory

    category learning

    priming

    skill learning

    familiarity judgements

- **Retrograde amnesia** *(Squire, 1986)*
  - **recent memories lost**

Hippocampus and Episodic Memory

- **Anterograde amnesia** *(Scoville & Milner, 1957)*
  - **HM** *hippocampus removed to treat intractable epilepsy*

    impaired explicit / *episodic* memory

    paired associates
    one trial learning: recognition of a new face

    spared implicit / *semantic* memory

    category learning
    priming
    skill learning
    familiarity judgements

- **Retrograde amnesia** *(Squire, 1986)*
  - **recent memories lost** ⇒ *episodic memories initially stored in hippocampus*

Hippocampus and Episodic Memory

- **Anterograde amnesia** *(Scoville & Milner, 1957)*
  - **HM** *hippocampus removed to treat intractable epilepsy*

    impaired explicit / *episodic* memory

    paired associates

    one trial learning: recognition of a new face

    spared implicit / *semantic* memory

    category learning

    priming

    skill learning

    familiarity judgements

- **Retrograde amnesia** *(Squire, 1986)*
  - **recent memories lost** $\Rightarrow$ *episodic memories initially stored in hippocampus*
  - **old ones spared**

Hippocampus and Episodic Memory

- **Anterograde amnesia** *(Scoville & Milner, 1957)*
  - **HM** *hippocampus removed to treat intractable epilepsy*

    impaired explicit / *episodic* memory

    paired associates

    one trial learning: recognition of a new face

    spared implicit / *semantic* memory

    category learning

    priming

    skill learning

    familiarity judgements

- **Retrograde amnesia** *(Squire, 1986)*
  - **recent memories lost** ⇒ *episodic memories initially stored in hippocampus*
  - **old ones spared** ⇒ *ends up somewhere else*

Complementary Learning Systems Hypothesis

Complementary Learning Systems Hypothesis

- Brain solves the problem of the fast/slow tradeoff in learning by having two learning systems:

Complementary Learning Systems Hypothesis

- Brain solves the problem of the fast/slow tradeoff in learning by having two learning systems:
  - One learns quickly, separates, and stores episodic information

hippocampus

Complementary Learning Systems Hypothesis

- Brain solves the problem of the fast/slow tradeoff in learning by having two learning systems:
  - One learns quickly, separates, and stores episodic information

    hippocampus

    

  - The other learns slowly, aggregates, and stores semantic information

    neocortex

    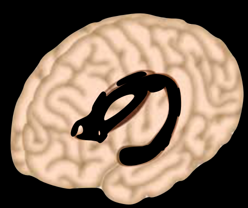

Functions of Hippocampus

Functions of Hippocampus

- ***Encoding*** **of arbitrary new associations:** *short (intermediate) term memory*
  - ***orthogonalization:*** **separation of representations to ensure specificity of association**
    - *isolate* items from their *semantic (statistical)* associations:
      where did I park my car *today*, irrespective of where I *usually* park it

Functions of Hippocampus

- ***Encoding*** **of arbitrary new associations:** *short (intermediate) term memory*
  - ***orthogonalization:*** **separation of representations to ensure specificity of association**
    - *isolate* items from their *semantic (statistical)* associations:
      where did I park my car *today*, irrespective of where I *usually* park it
  - ***binding:*** ***rapid*** **formation of associations**
    - one- (or “k”)-shot learning

Functions of Hippocampus

- ***Encoding*** **of arbitrary new associations:** *short (intermediate) term memory*
  - ***orthogonalization:*** **separation of representations to ensure specificity of association**
    - *isolate* items from their *semantic (statistical)* associations: where did I park my car *today*, irrespective of where I *usually* park it
  - ***binding:*** ***rapid*** **formation of associations**
    - one- (or "k")-shot learning
  - **Neural mechanisms**
    - ***orthonalization*****:** **projection of cortical representations into *high dimensionality* of dentate gyrus**
    - ***binding:*** **long term potentiation (LTP) — rapid synaptic plasticity**

Functions of Hippocampus

- ***Encoding*** **of arbitrary new associations:** *short (intermediate) term memory*
  - ***orthogonalization:*** **separation of representations to ensure specificity of association**
    - *isolate* items from their *semantic (statistical)* associations: where did I park my car *today*, irrespective of where I *usually* park it
  - ***binding:*** ***rapid*** **formation of associations**
    - one- (or "k")-shot learning
  - **Neural mechanisms**
    - ***orthonalization*****:** **projection of cortical representations into *high dimensionality* of dentate gyrus**
    - ***binding:*** **long term potentiation (LTP) — rapid synaptic plasticity**
  - **Computational mechanisms**
    - **rapid Hebbian (associational) learning**
    - **"neural dictionary:" key-value pairing**

Functions of Hippocampus

- ***Encoding*** **of arbitrary new associations:** *short (intermediate) term memory*
  - ***orthogonalization:*** **separation of representations to ensure specificity of association**
    - *isolate* items from their *semantic (statistical)* associations: where did I park my car *today*, irrespective of where I *usually* park it
  - ***binding:*** ***rapid*** **formation of associations**
    - one- (or "k")-shot learning
  - **Neural mechanisms**
    - ***orthonalization*****: projection of cortical representations into *high dimensionality* of dentate gyrus**
    - ***binding:*** **long term potentiation (LTP) — rapid synaptic plasticity**
  - **Computational mechanisms**
    - **rapid Hebbian (associational) learning**
    - **"neural dictionary:" key-value pairing**
- ***Consolidation*** **into neocortex:** *long term memory*
  - **slowly "sift flour" of new information into "dough" of old knowledge:**
    - **reinstatement through hippocampus replay (reheasrsal)**
    - **interleaved training of cortex**

Functions of Hippocampus

- ***Encoding*** **of arbitrary new associations:** *short (intermediate) term memory*
  - ***orthogonalization:*** **separation of representations to ensure specificity of association**
    - *isolate* items from their *semantic (statistical)* associations: where did I park my car *today*, irrespective of where I *usually* park it
  - ***binding:*** ***rapid*** **formation of associations**
    - one- (or "k")-shot learning
  - **Neural mechanisms**
    - ***orthonalization*****: projection of cortical representations into *high dimensionality* of dentate gyrus**
    - ***binding:*** **long term potentiation (LTP) — rapid synaptic plasticity**
  - **Computational mechanisms**
    - **rapid Hebbian (associational) learning**
    - **"neural dictionary:" key-value pairing**
- ***Consolidation*** **into neocortex:** *long term memory*
  - **slowly "sift flour" of new information into "dough" of old knowledge:**
    - **reinstatement through hippocampus replay (reheasrsal)**
    - **interleaved training of cortex**
  - **why must it be slow?**
    - **ensure it is relevant**
    - **minimize disruption of existing knowledge**

Complementary Learning Systems Hypothesis

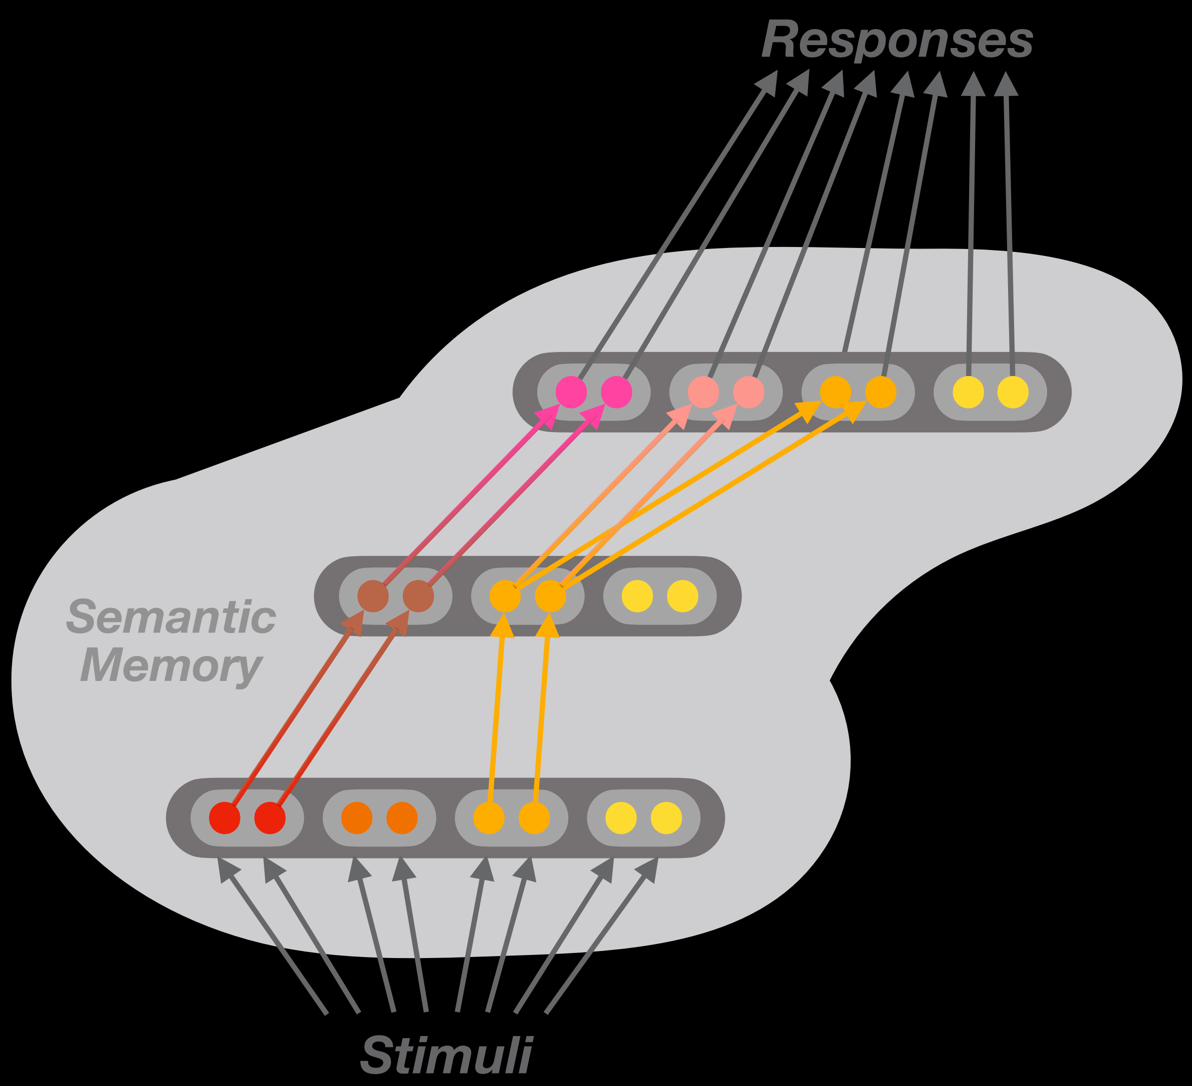

Complementary Learning Systems Hypothesis

Want associations

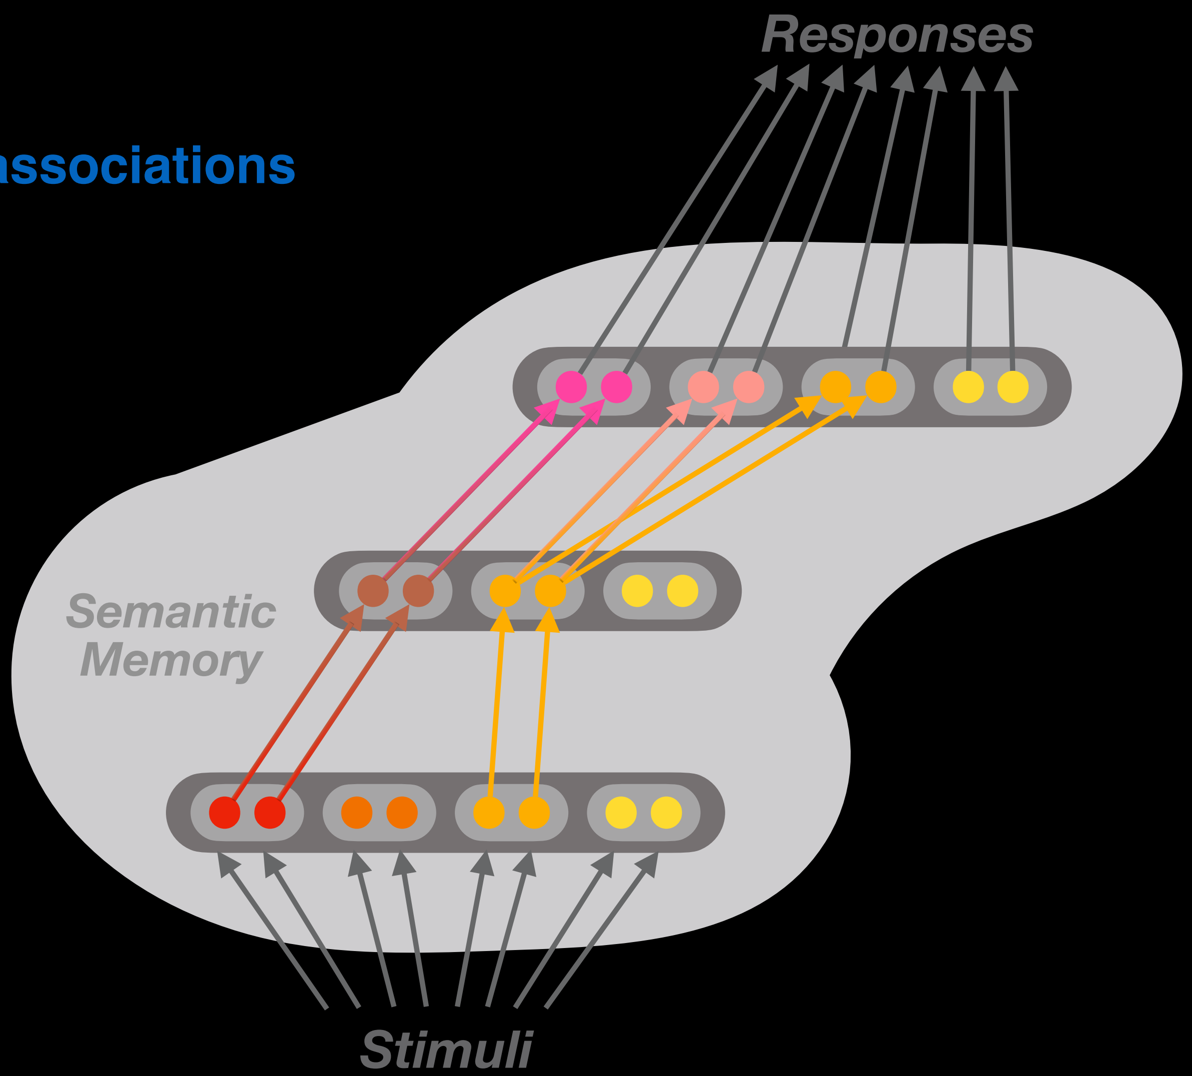

Complementary Learning Systems Hypothesis

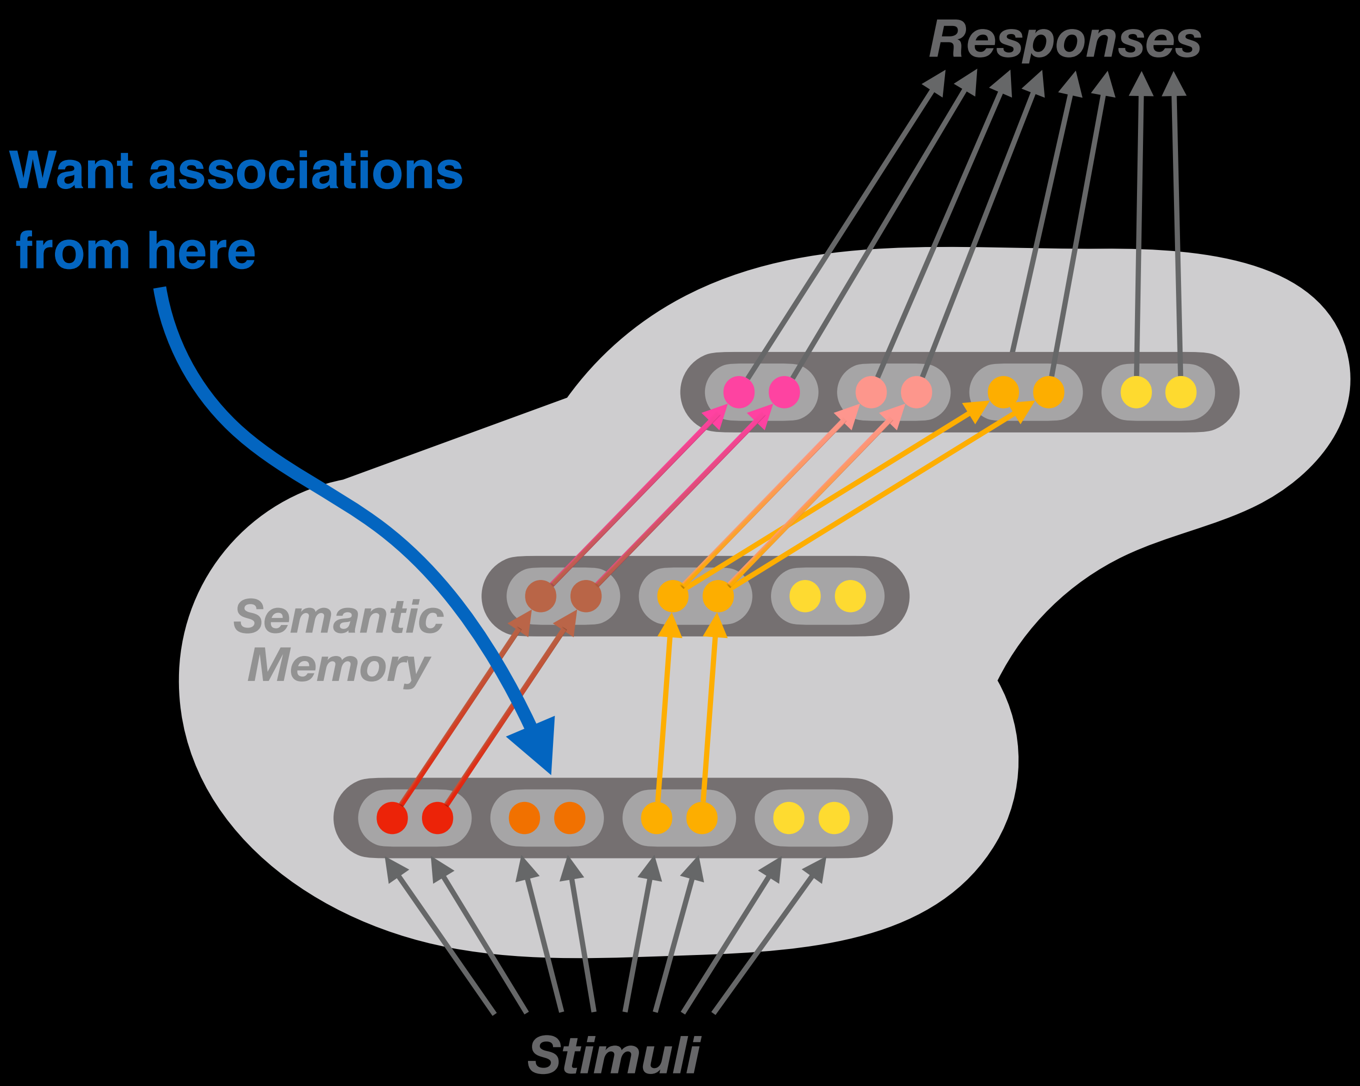

Complementary Learning Systems Hypothesis

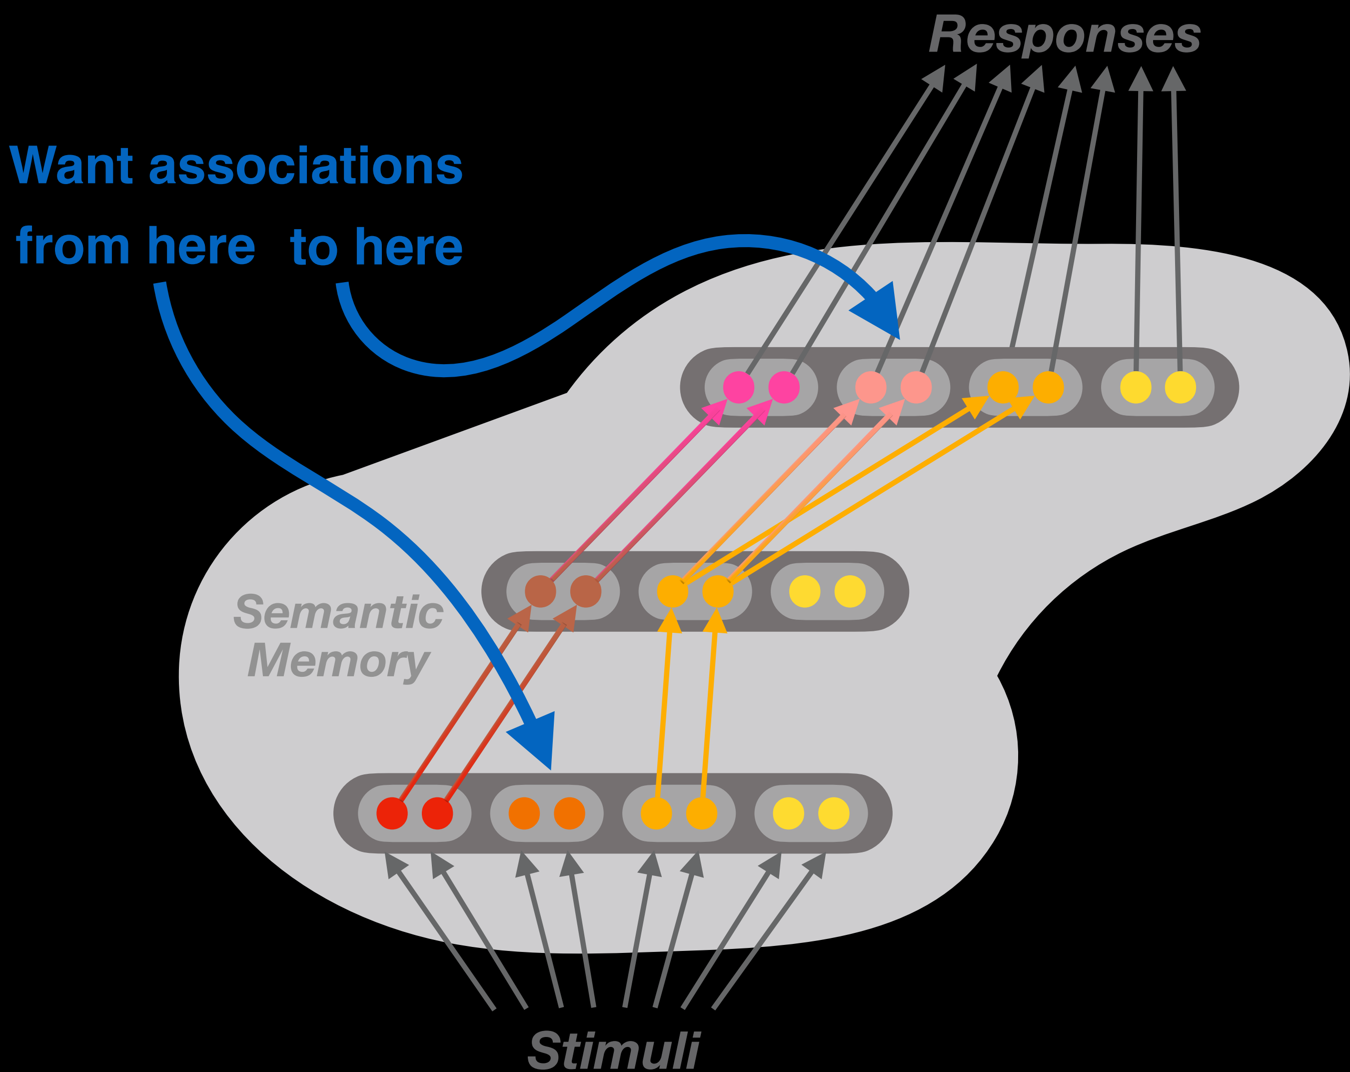

Complementary Learning Systems Hypothesis

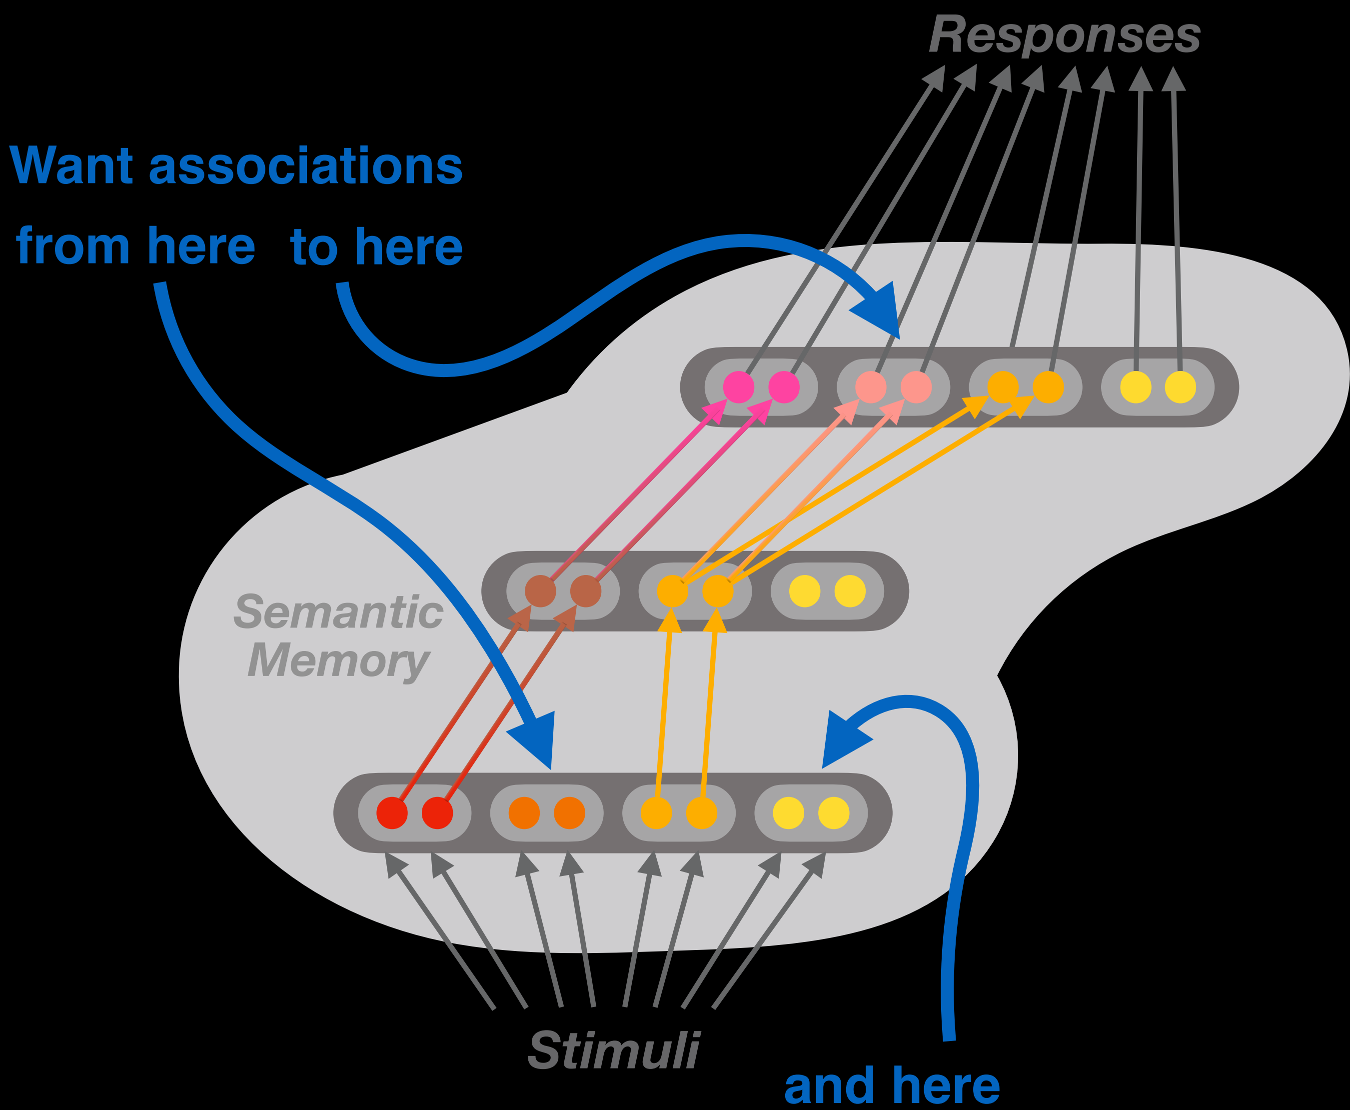

Complementary Learning Systems Hypothesis

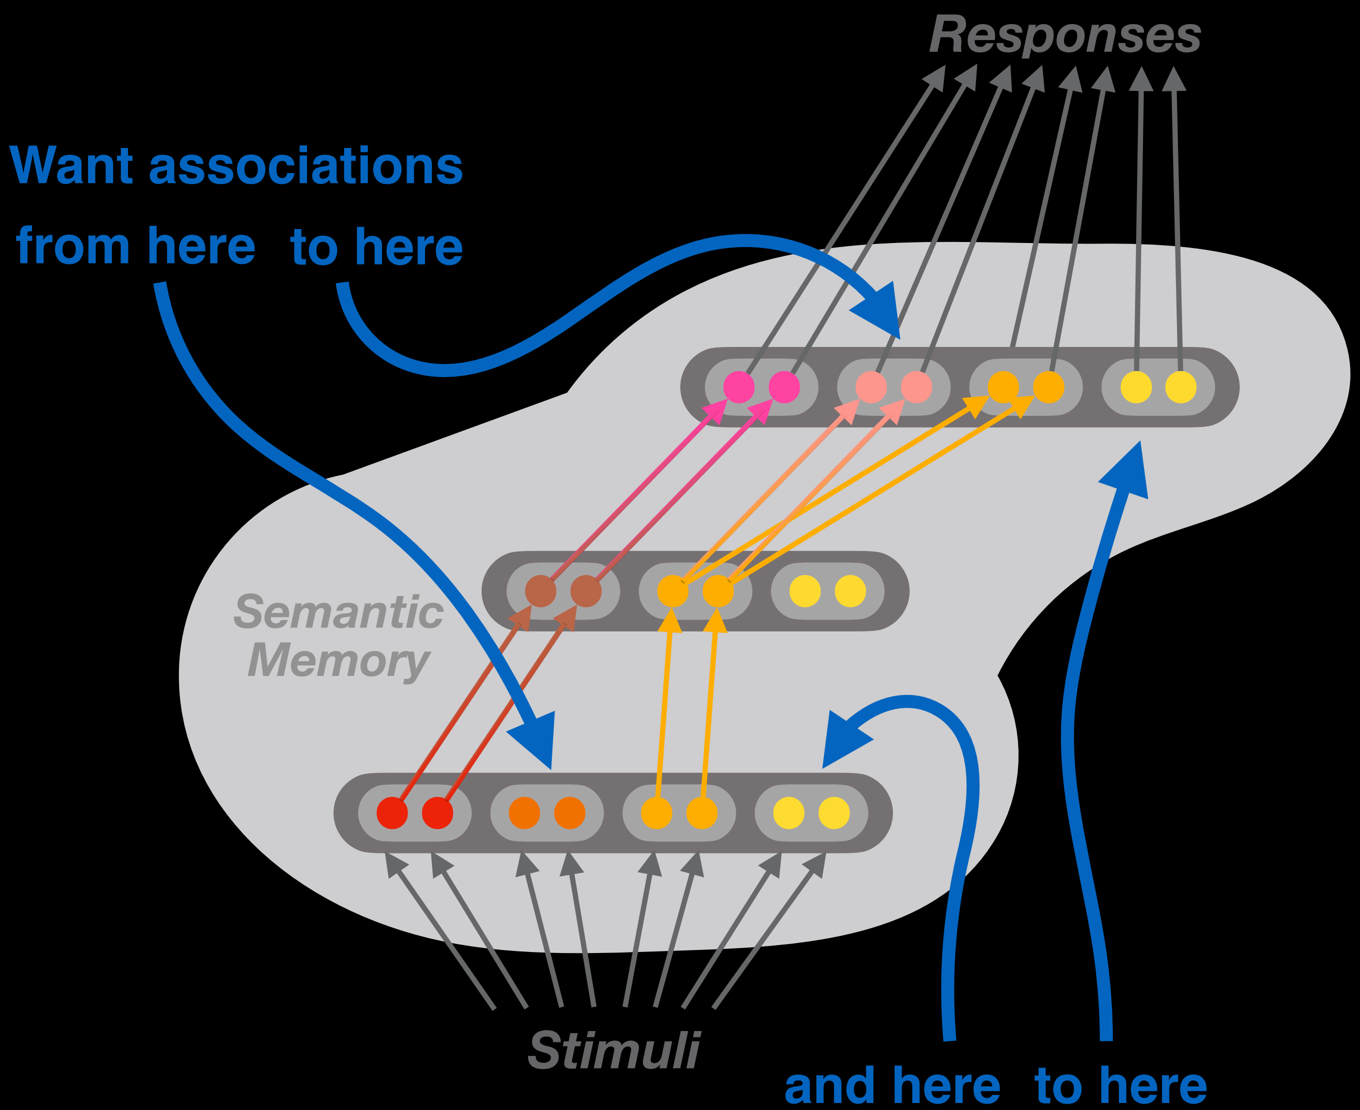

Complementary Learning Systems Hypothesis

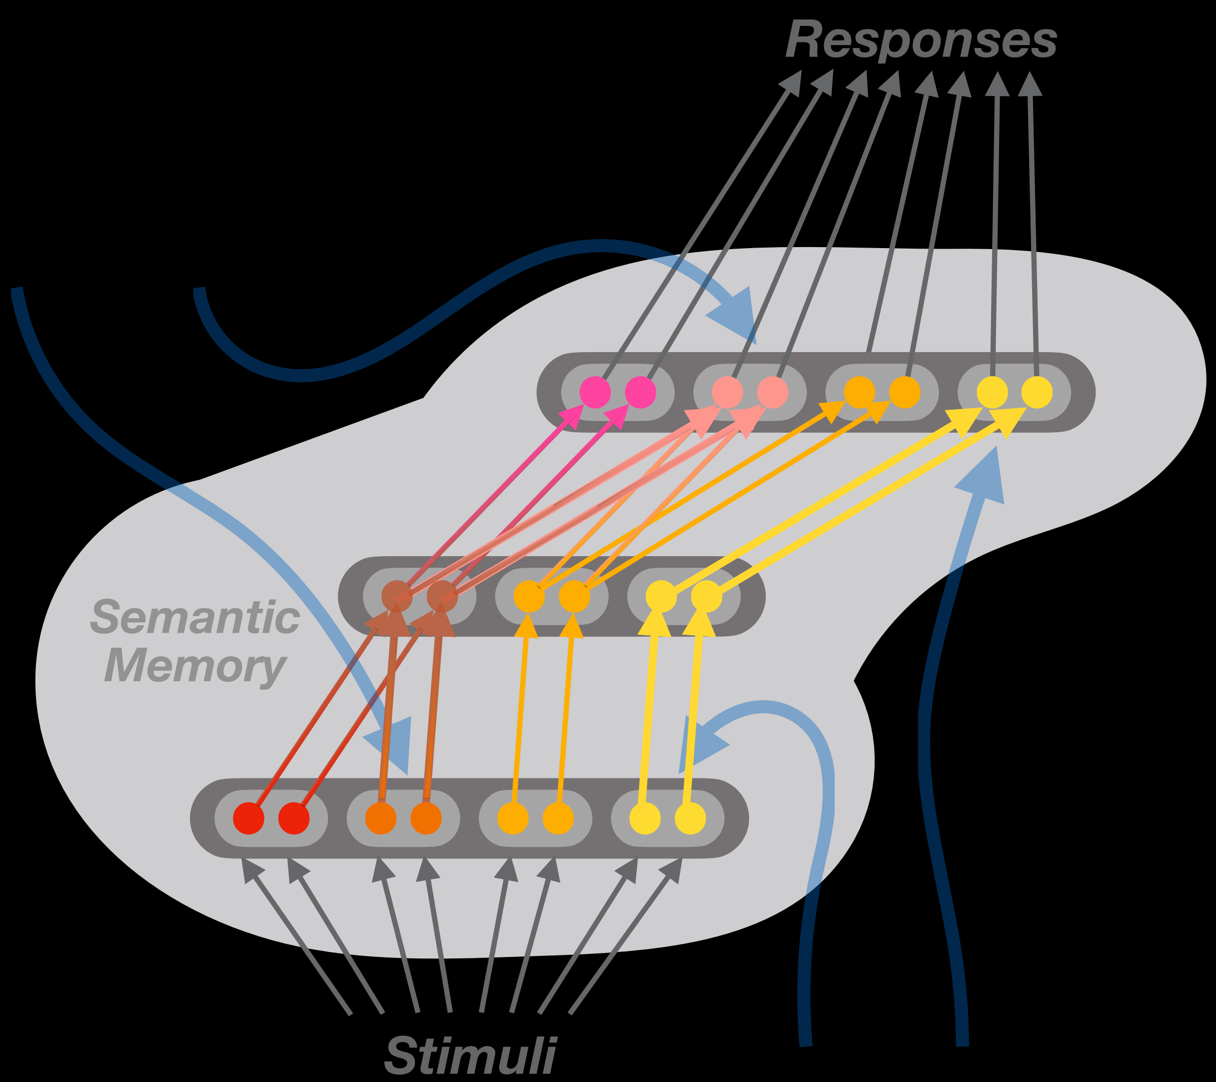

and with time,
consolidated

Functions of Hippocampal Reinstatement

Functions of Hippocampal Reinstatement

- **Recall**
  - –Retrieve “snapshots” of recent events

Functions of Hippocampal Reinstatement

- **Recall**
  - –Retrieve “snapshots” of recent events
- **Training**
  - –Replay to expose neocortex in interleaved fashion

Simulation Studies of Consolidation

(McClelland, McNaughton & O'Reilly 1995)

Simulation Studies of Consolidation

(McClelland, McNaughton & O'Reilly 1995)

- **Abstract model of hippocampus:**
  - Algorithm for reinstatement of learned items in a cortical network
    (cortical network comparable to McCloskey & Cohen — 3 layer backprop)

Simulation Studies of Consolidation

(McClelland, McNaughton & O'Reilly 1995)

- **Abstract model of hippocampus:**
  - Algorithm for reinstatement of learned items in a cortical network *(cortical network comparable to McCloskey & Cohen — 3 layer backprop)*
  - Reinstatement with probability:
    - ⬆ with salience/importance of original item

Simulation Studies of Consolidation

(McClelland, McNaughton & O'Reilly 1995)

- **Abstract model of hippocampus:**
  - Algorithm for reinstatement of learned items in a cortical network *(cortical network comparable to McCloskey & Cohen — 3 layer backprop)*
  - Reinstatement with probability:
    - ⬆ with salience/importance of original item
    - ⬇ with time

Simulation Studies of Consolidation

(McClelland, McNaughton & O'Reilly 1995)

- **Abstract model of hippocampus:**
  - Algorithm for reinstatement of learned items in a cortical network *(cortical network comparable to McCloskey & Cohen — 3 layer backprop)*
  - Reinstatement with probability:
    - ⬆ with salience/importance of original item
    - ⬇ with time
    - ⬆ for task-relevant vs. task-irrelevant contexts

Simulation Studies of Consolidation

(McClelland, McNaughton & O'Reilly 1995)

- **Abstract model of hippocampus:**
  - Algorithm for reinstatement of learned items in a cortical network *(cortical network comparable to McCloskey & Cohen — 3 layer backprop)*
  - Reinstatement with probability:
    - ⬆ with salience/importance of original item
    - ⬇ with time
    - ⬆ for task-relevant vs. task-irrelevant contexts
  - This captured the hypothesized dynamics of learning within the hippocampus, without actually simulating these in network form

    (Norman et al. 2006: similar results using actual neural network for hcmp)

Simulation Studies of Consolidation

(McClelland, McNaughton & O'Reilly 1995)

- **Abstract model of hippocampus:**
  - Algorithm for reinstatement of learned items in a cortical network *(cortical network comparable to McCloskey & Cohen — 3 layer backprop)*
  - Reinstatement with probability:
    - ⬆ with salience/importance of original item
    - ⬇ with time
    - ⬆ for task-relevant vs. task-irrelevant contexts
  - This captured the hypothesized dynamics of learning within the hippocampus, without actually simulating these in network form
    (Norman et al. 2006: similar results using actual neural network for hcmp)
- **Simulated effects of consolidation in animal studies**

  For example…

Memory Consolidation in Rodents

(Kim & Fanselow, 1992)

Memory Consolidation in Rodents

(Kim & Fanselow, 1992)

- **Empirical study**

  Training:

  15 pairings of tone & footshock

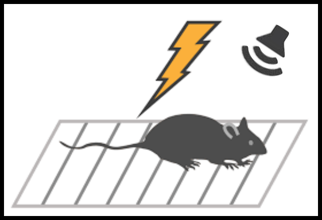

Memory Consolidation in Rodents

(Kim & Fanselow, 1992)

- **Empirical study**

  Training:

  15 pairings of tone & footshock

  Bilateral hippocampus lesion or sham

  1, 7, 14 or 28 days after training

Memory Consolidation in Rodents

(Kim & Fanselow, 1992)

- **Empirical study**
  - **Training:**
    - 15 pairings of tone & footshock
  - **Bilateral hippocampus lesion or sham**
    - 1, 7, 14 or 28 days after training
  - **Tested 7 days after lesion**
    - Tone presented
    - Fear response assessed

Memory Consolidation in Rodents

(Kim & Fanselow, 1992)

- **Empirical study**
  - **Training:**
    - 15 pairings of tone & footshock
  - **Bilateral hippocampus lesion or sham**
    - 1, 7, 14 or 28 days after training
  - **Tested 7 days after lesion**
    - Tone presented
    - Fear response assessed

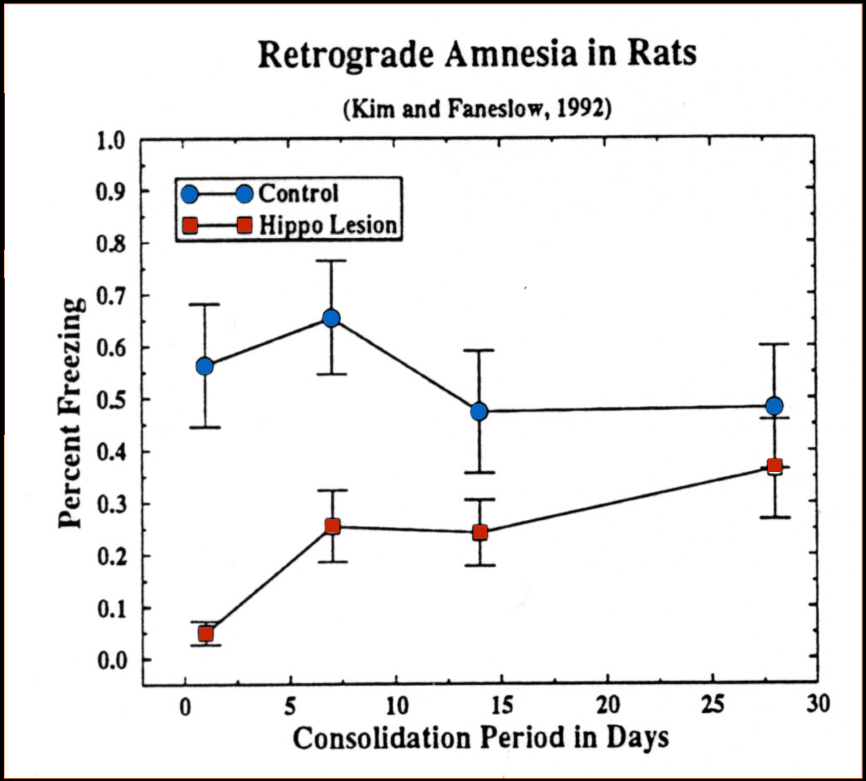

Memory Consolidation in Rodents

(Kim & Fanselow, 1992)

- **Empirical study**
  - **Training:**
    - 15 pairings of tone & footshock
  - **Bilateral hippocampus lesion or sham**
    - 1, 7, 14 or 28 days after training
  - **Tested 7 days after lesion**
    - Tone presented
    - Fear response assessed

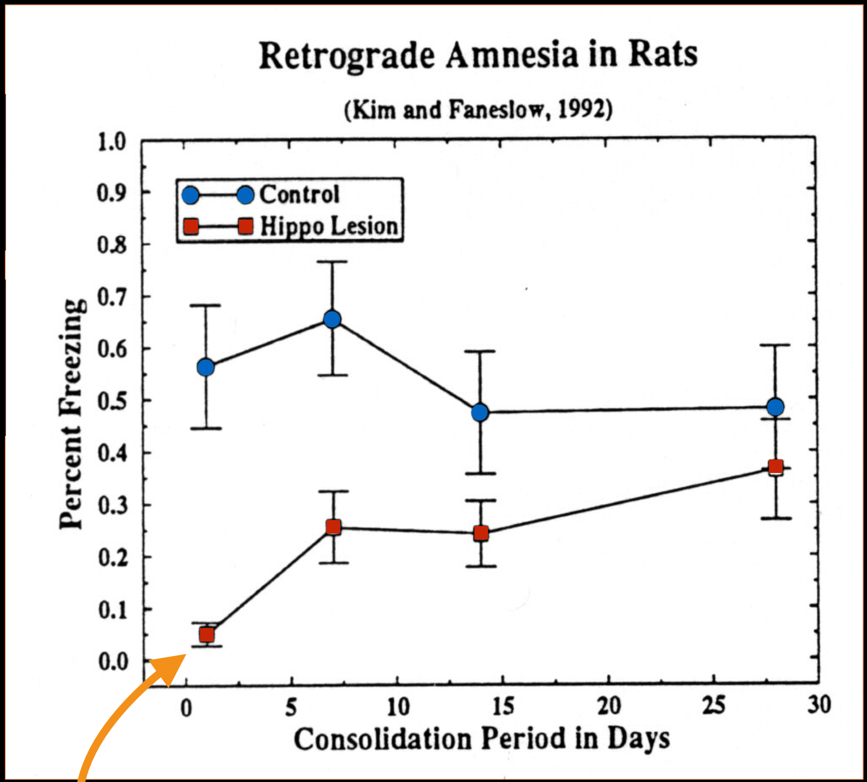

Lesioned immediately after training

Memory Consolidation in Rodents

(Kim & Fanselow, 1992)

- **Empirical study**
  - **Training:**
    - 15 pairings of tone & footshock
  - **Bilateral hippocampus lesion or sham**
    - 1, 7, 14 or 28 days after training
  - **Tested 7 days after lesion**
    - Tone presented
    - Fear response assessed

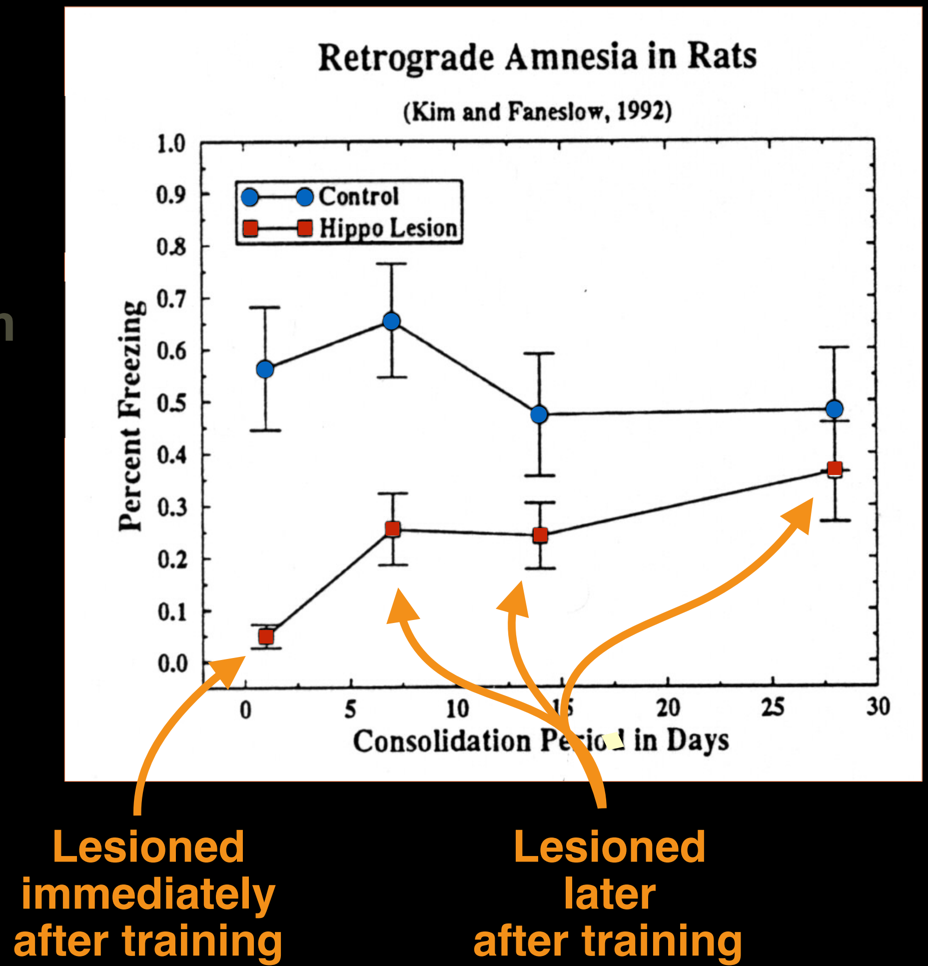

Model

(McClelland, McNaughton & O'Reilly, 1995)

Model

(McClelland, McNaughton & O'Reilly, 1995)

- **Simulation:** *3 layer network (16 units at each layer)*
  - **Pre-training:**

    20 input-output pairs (random patterns): network's "background knowledge"

Model

(McClelland, McNaughton & O'Reilly, 1995)

- **Simulation:** *3 layer network (16 units at each layer)*
  - **Pre-training:**
    - 20 input-output pairs (random patterns): network's "background knowledge"
    - Each of these patterns was reinstated once per simulated "day" of training

Model

(McClelland, McNaughton & O'Reilly, 1995)

- **Simulation:** *3 layer network (16 units at each layer)*
  - **- Pre-training:**
    20 input-output pairs (random patterns): network's "background knowledge"
    Each of these patterns was reinstated once per simulated "day" of training
  - **- Experimental training** *(simulation of tone ➠ footshock):*
    One additional pair (input ≈ tone; output ≈ fear response)

Model

(McClelland, McNaughton & O'Reilly, 1995)

- **Simulation:** *3 layer network (16 units at each layer)*
  - **Pre-training:**

    20 input-output pairs (random patterns): network's "background knowledge"

    Each of these patterns was reinstated once per simulated "day" of training

  - **Experimental training** *(simulation of tone ➠ footshock):*

    One additional pair (input ≈ tone; output ≈ fear response)

    Reinstated interleaved with other 20 patterns on each "day"
    (to simulate storage in the hippocampus until it was "lesioned"

Model

(McClelland, McNaughton & O'Reilly, 1995)

- **Simulation:** *3 layer network (16 units at each layer)*
  - **Pre-training:**

    20 input-output pairs (random patterns): network's "background knowledge"

    Each of these patterns was reinstated once per simulated "day" of training

  - **Experimental training** *(simulation of tone ➠ footshock):*

    One additional pair (input ≈ tone; output ≈ fear response)

    Reinstated interleaved with other 20 patterns on each "day"
    (to simulate storage in the hippocampus until it was "lesioned"

  - **Testing:**

    Error on test pair (tone ➠ foot shock)

Model

(McClelland, McNaughton & O'Reilly, 1995)

- **Simulation:** *3 layer network (16 units at each layer)*

  - Pre-training:

  20 input-output pairs (random patterns): network's "background knowledge"

  Each of these patterns was reinstated once per simulated "day" of training

  - Experimental training *(simulation of tone ➠ footshock):*

  One additional pair (input ≈ tone; output ≈ fear response)

  Reinstated interleaved with other 20 patterns on each "day"
  (to simulate storage in the hippocampus until it was "lesioned"

  - Testing:

  Error on test pair (tone ➠ foot shock)

  - Parameterization:

  Learning rate adjusted to fit empirical data

Simulation Results

(McClelland, McNaughton & O'Reilly, 1995)

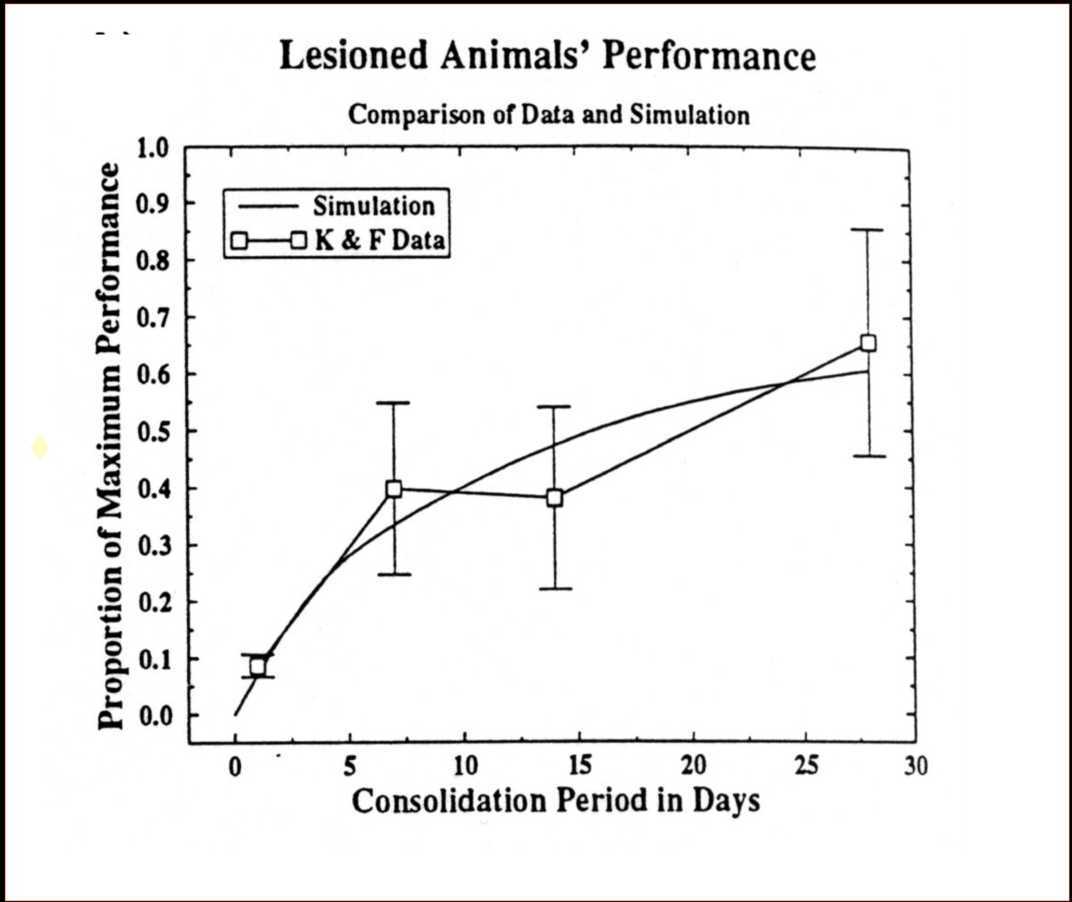

Simulation Results

(McClelland, McNaughton & O'Reilly, 1995)

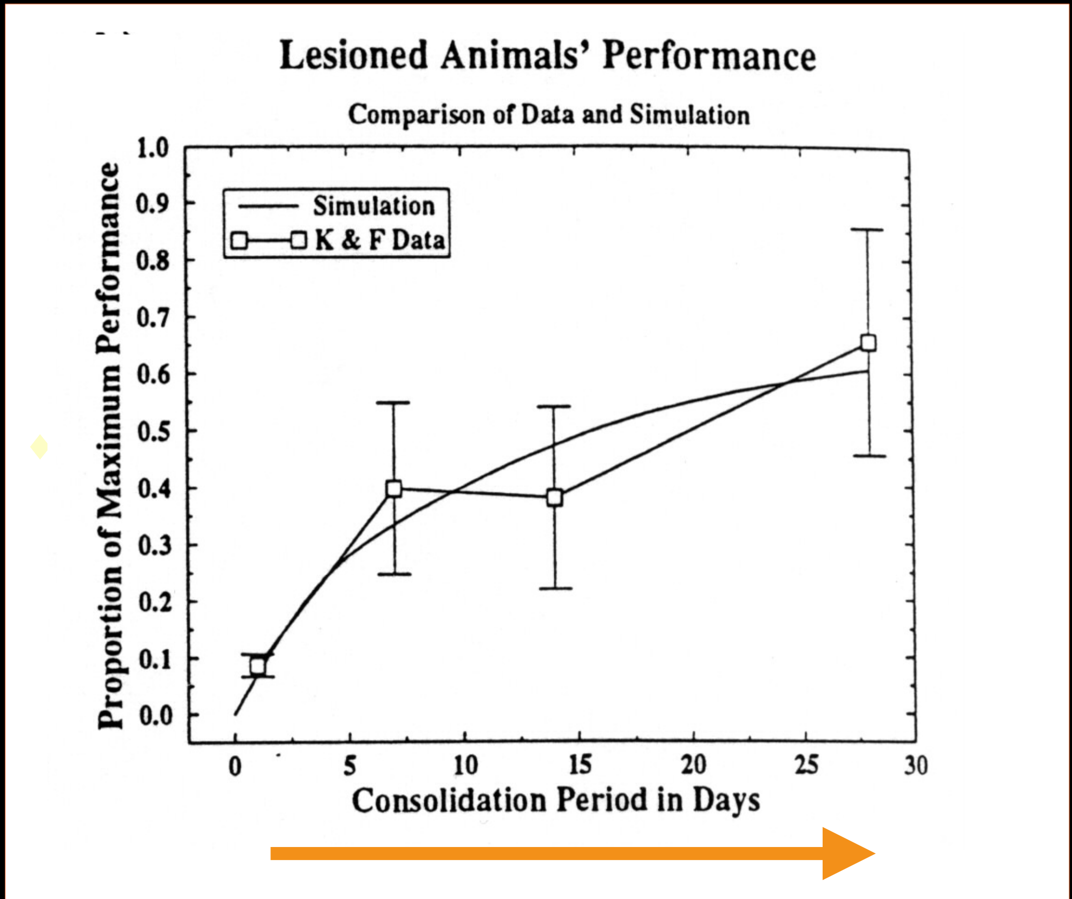

The later the lesion…

Simulation Results

(McClelland, McNaughton & O'Reilly, 1995)

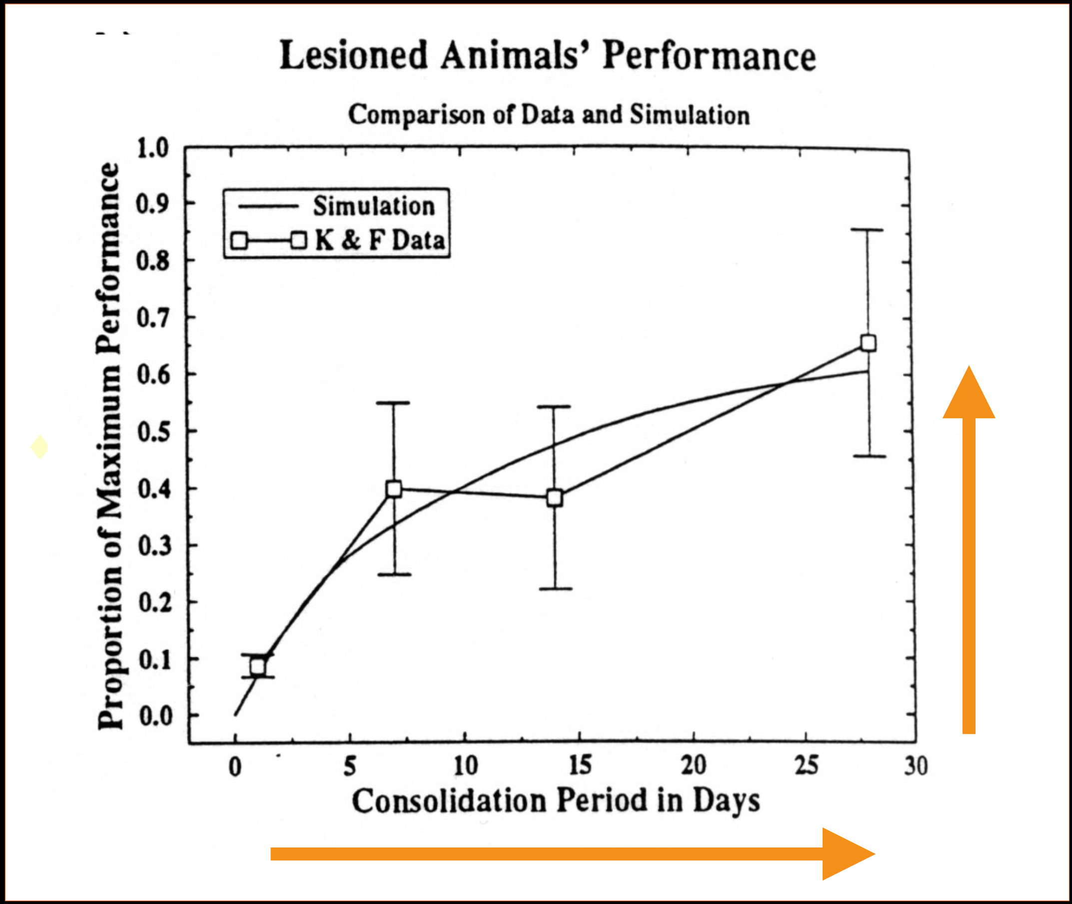

the better the retention

The later the lesion…

Memory Consolidation in Primates

(Squire & Zola-Morgan, 1990)

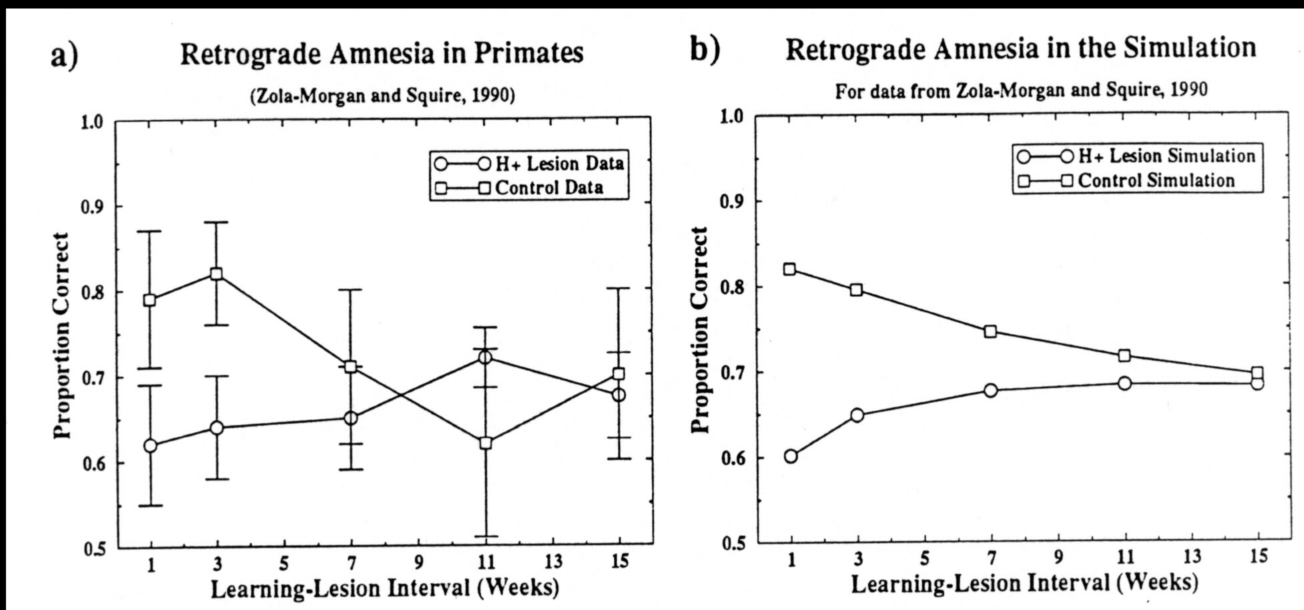

Memory Consolidation in Primates

(Squire & Zola-Morgan, 1990)

- Similar experiment in primates

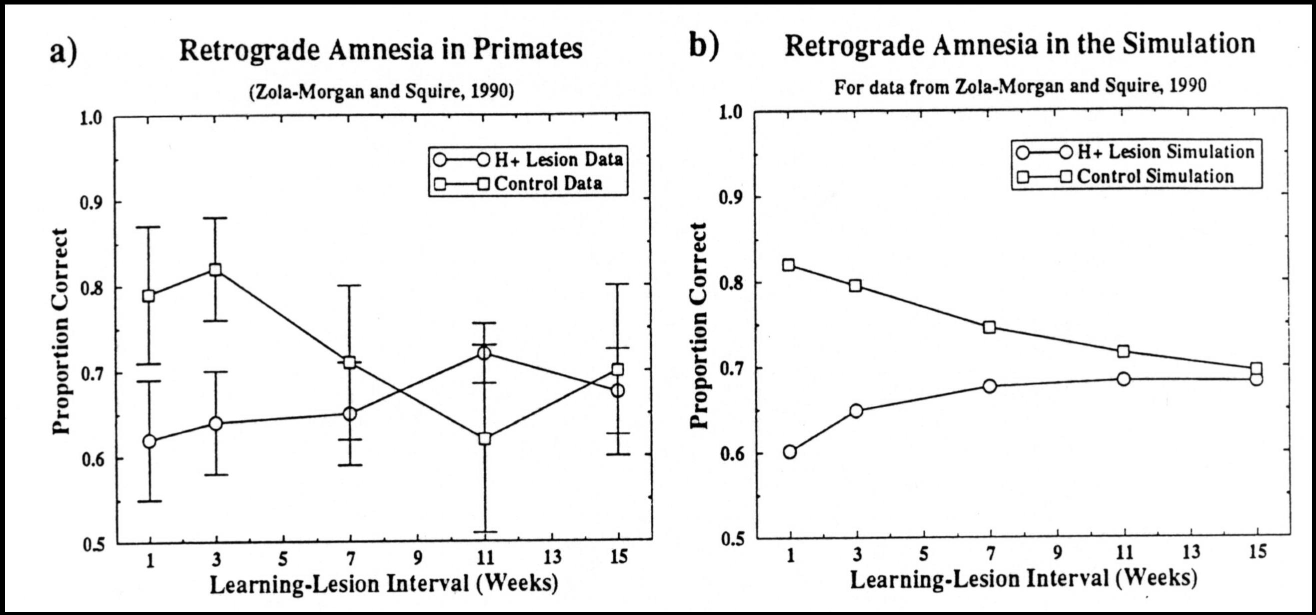

Memory Consolidation in Primates

(Squire & Zola-Morgan, 1990)

- Similar experiment in primates

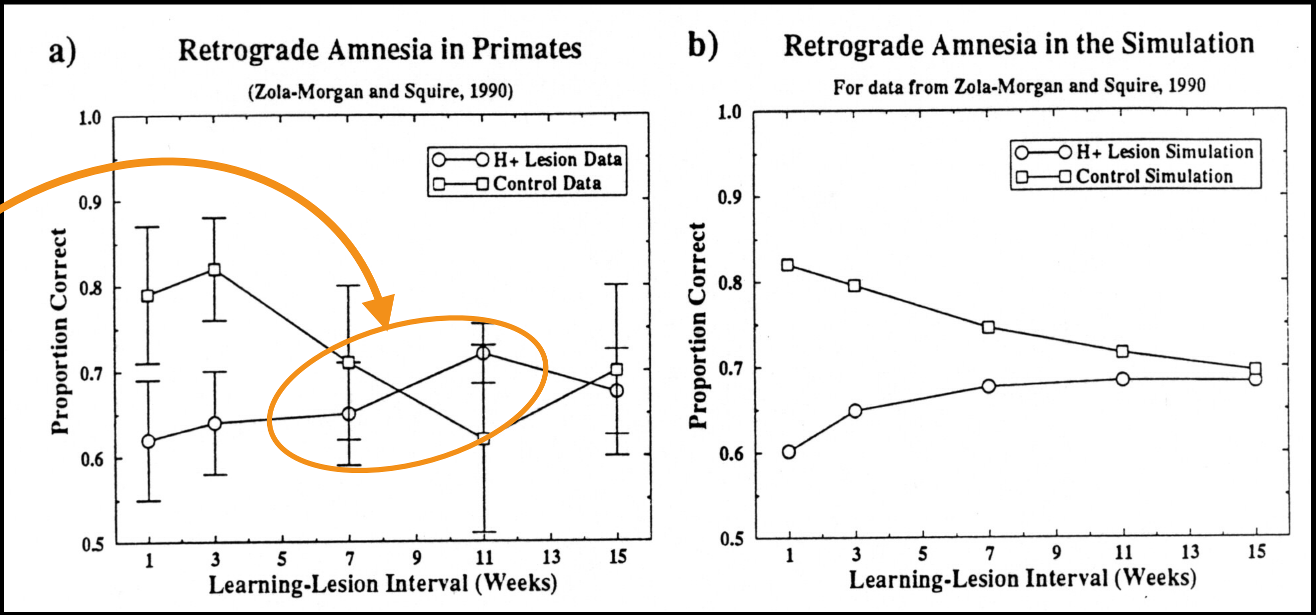

- Note:

  decay in control animals:

  hippocampal decay faster than neocortical learning

Memory Consolidation in Primates

(Squire & Zola-Morgan, 1990)

- Similar experiment in primates

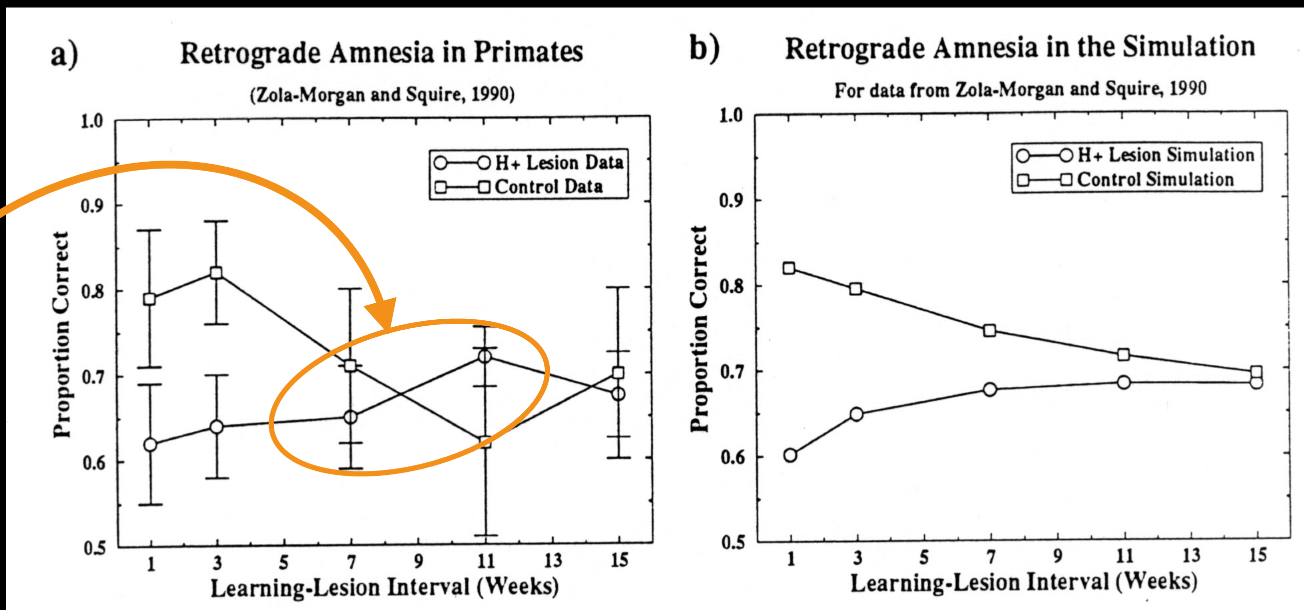

- Note:

  decay in control animals:

  hippocampal decay faster than neocortical learning

  ⇒ window of consolidation

Memory Consolidation in Primates

(Squire & Zola-Morgan, 1990)

- Similar experiment in primates

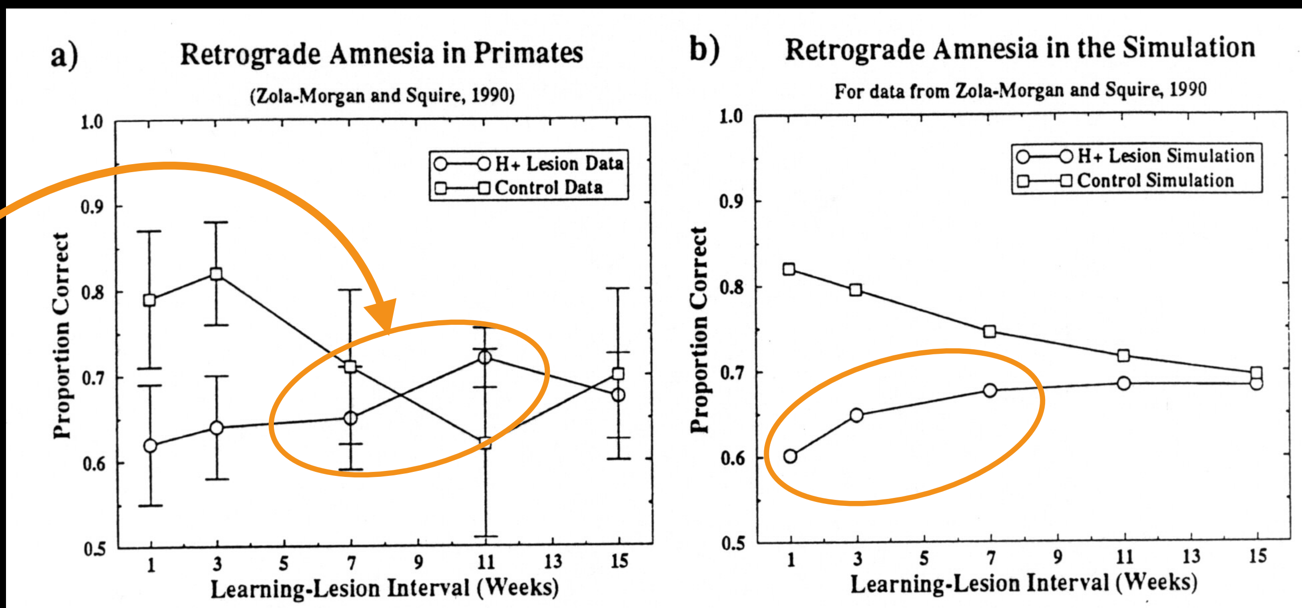

- Note:

  decay in control animals:

  hippocampal decay faster than neocortical learning

  ⇒ window of consolidation

Consolidation as "Replay"

Consolidation as "Replay"

- **Modeling:**
  - **machine learning:**
    - Dyna *(Sutton et al., 1991)*

Consolidation as “Replay”

- **Modeling:**
  - **machine learning:**
    - Dyna *(Sutton et al., 1991)*
  - **neuroscience:**
    - **model-based RL**
      Daw et al., 2005; Miller et al., 2017

Consolidation as "Replay"

- **Modeling:**
  - **machine learning:**
    - Dyna *(Sutton et al., 1991)*
  - **neuroscience:**
    - **model-based RL**

      Daw et al., 2005; Miller et al., 2017
    - **spiking neuron models**

      Azizi et al, 2013; Atherton et al., 2015; Chenkov et al., 2017

Consolidation as “Replay”

- **Modeling:**
  - **machine learning:**
    - Dyna *(Sutton et al., 1991)*
  - **neuroscience:**
    - **model-based RL**
      Daw et al., 2005; Miller et al., 2017
    - **spiking neuron models**
      Azizi et al, 2013; Atherton et al., 2015; Chenkov et al., 2017
- **Functions:**
  - **deliberation:** “vicarious trial and error” — fatigue
    (Tolman, 1939; Reddish, 2016; Agrawal, under review)

Consolidation as “Replay”

- **Modeling:**
  - **machine learning:**
    - Dyna *(Sutton et al., 1991)*
  - **neuroscience:**
    - **model-based RL**
      Daw et al., 2005; Miller et al., 2017
    - **spiking neuron models**
      Azizi et al, 2013; Atherton et al., 2015; Chenkov et al., 2017
- **Functions:**
  - **deliberation:** “vicarious trial and error” — fatigue
    (Tolman, 1939; Reddish, 2016; Agrawal, under review)
  - **planning and prospective memory:** “pre-play”
    (Cohen & O’Reilly, 1996; Einstein & McDaniel, 2005; Momennejad et al., 2020)

Consolidation as “Replay”

- **Modeling:**
  - **machine learning:**
    - Dyna *(Sutton et al., 1991)*
  - **neuroscience:**
    - **model-based RL**
      Daw et al., 2005; Miller et al., 2017
    - **spiking neuron models**
      Azizi et al, 2013; Atherton et al., 2015; Chenkov et al., 2017
- **Functions:**
  - **deliberation:** “vicarious trial and error” — fatigue
    (Tolman, 1939; Reddish, 2016; Agrawal, under review)
  - **planning and prospective memory:** “pre-play”
    (Cohen & O’Reilly, 1996; Einstein & McDaniel, 2005; Momennejad et al., 2020)
  - **working memory “assist”**
    (Hoskin et al., 2019; Beukers et al., in prep)

Consolidation as "Replay"

- **Modeling:**
  - **machine learning:**
    - Dyna *(Sutton et al., 1991)*
  - **neuroscience:**
    - **model-based RL**

      Daw et al., 2005; Miller et al., 2017
    - **spiking neuron models**

      Azizi et al, 2013; Atherton et al., 2015; Chenkov et al., 2017
- **Functions:**
  - **deliberation:** "vicarious trial and error" — fatigue

    (Tolman, 1939; Reddish, 2016; Agrawal, under review)
  - **planning and prospective memory:** "pre-play"

    (Cohen & O'Reilly, 1996; Einstein & McDaniel, 2005; Momennejad et al., 2020)
  - **working memory "assist"**

    (Hoskin et al., 2019; Beukers et al., in prep)
  - **sleep…**

Sleep

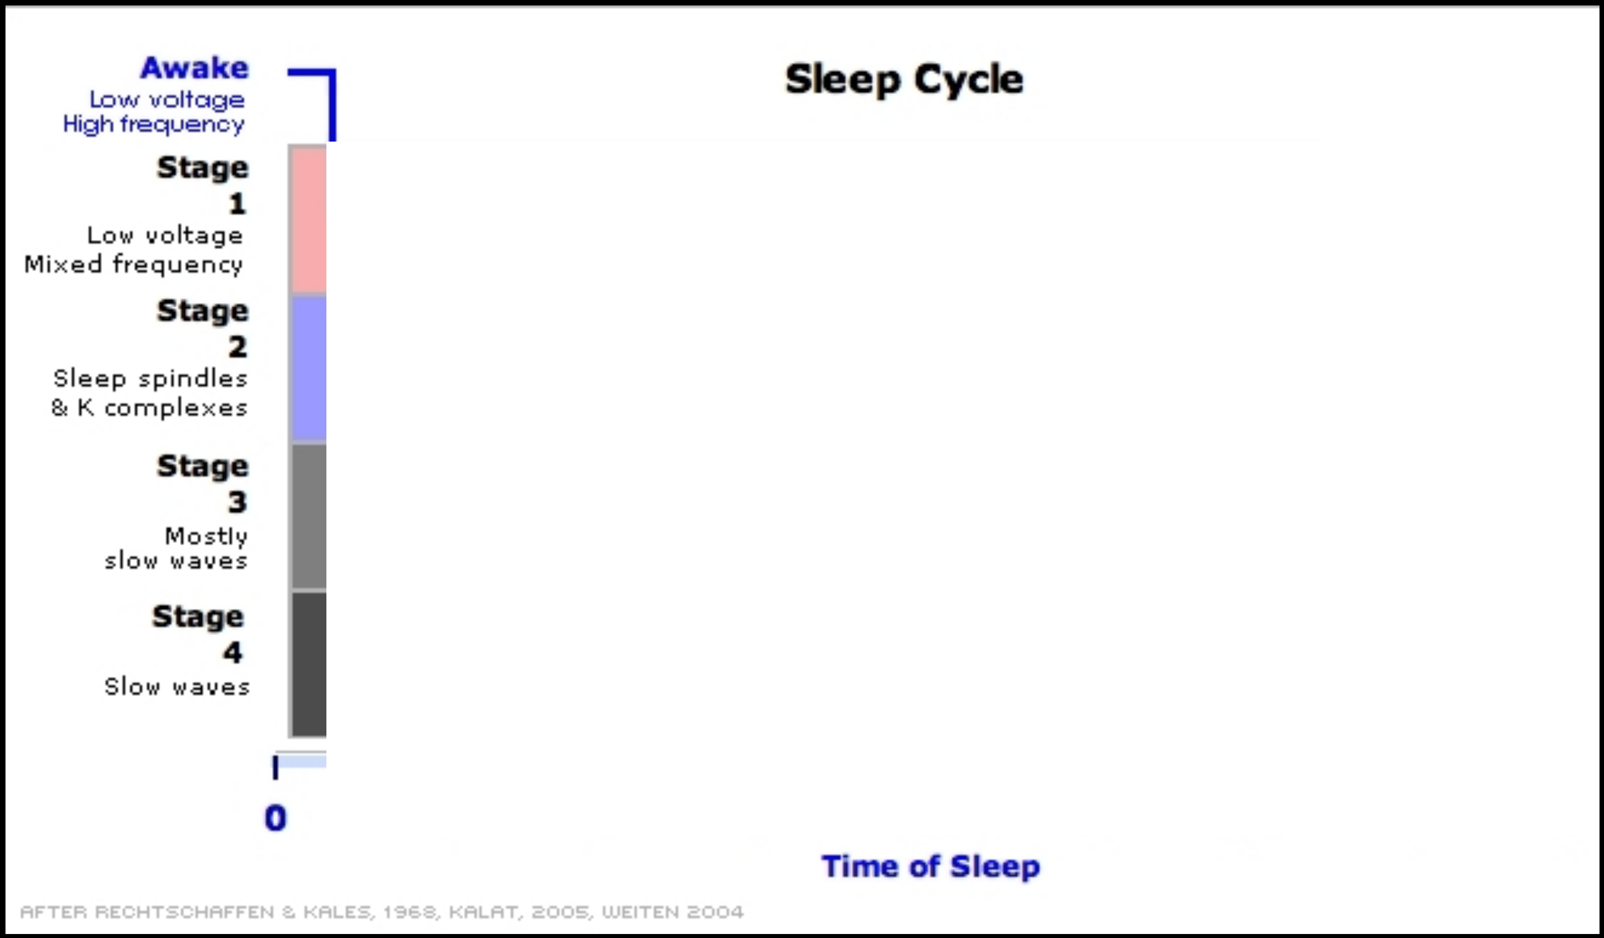

Sleep Cycle
Awake
Low voltage
High frequency
Stage 1
Low voltage
Mixed frequency
Stage 2
Sleep spindles
& K complexes
Stage 3
Mostly
slow waves
Stage 4
Slow waves
0
Time of Sleep
AFTER RECHTSCHAFFEN & KALES, 1968, KALAT, 2005, WEITEN 2004

Sleep

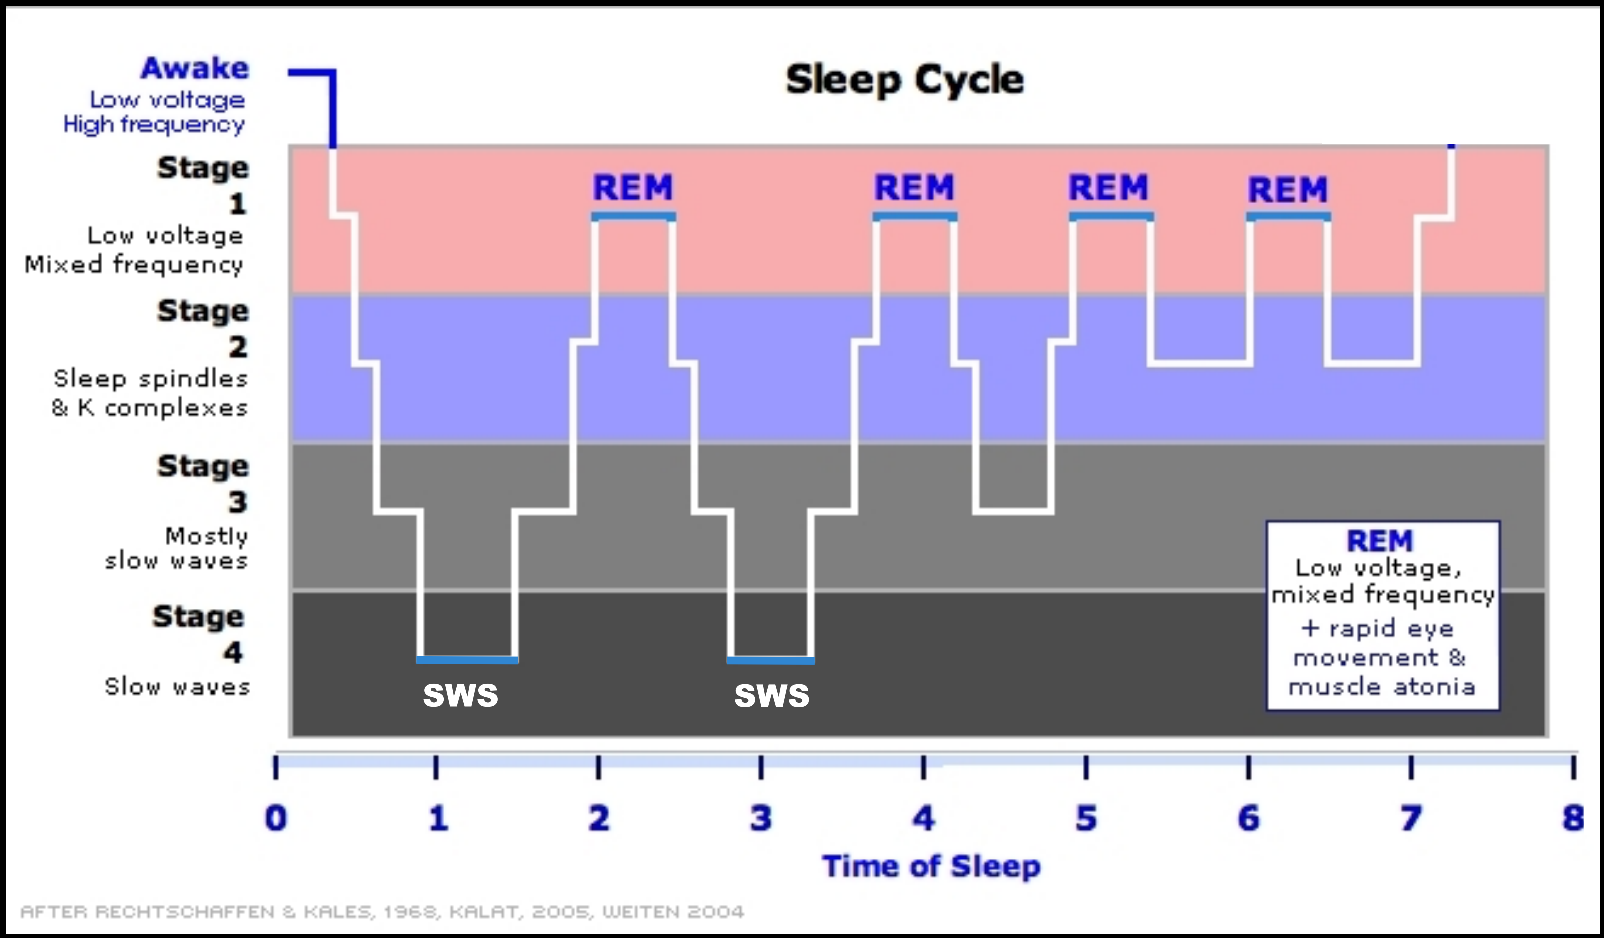
Sleep Cycle
Awake
Low voltage
High frequency
Stage 1
Low voltage
Mixed frequency
Stage 2
Sleep spindles
& K complexes
Stage 3
Mostly
slow waves
Stage 4
Slow waves
REM
REM
REM
REM
SWS
SWS
REM
Low voltage,
mixed frequency
+ rapid eye
movement &
muscle atonia
0
1
2
3
4
5
6
7
8
Time of Sleep
AFTER RECHTSCHAFFEN & KALES, 1968, KALAT, 2005, WEITEN 2004

Sleep and Dreaming

Sleep and Dreaming

Sleep and Dreaming

Sleep and Dreaming

Sleep and Dreaming

Why?