## 1 Memory

- Memory = any persistent effect of experience (not just memorization of facts, events, names, etc.)
- Weights vs activations
- Gradual, integrative cortical learning and priming effects
- Rapid memorization: The hippocampus
- Active memory: prefrontal cortex

## 2 Memory: Weights vs Activations

Despite appearances, memory is not unitary.
(shoes; breakfast; sentence)

*Weights:*

- Long-lasting.
- Requires re-activation.
- Wts in diff't brain systems store different types of memories!

*Activations:*

- Short-term.
- Already active, can influence processing.

## 3 Weight-based Memories

- Cortex does gradual, integrative learning
- Cortex can learn arbitrary input-output mappings given:
  - multiple passes through the training set
  - a relatively small learning rate

## 4 Weight-based Memories

Rapid weight changes causes interference:

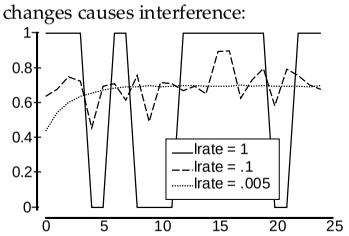


Two systems needed:

- Slow learning cortex.
- Rapid learning hippocampus (pattern sep avoids interference).

## 5 Activation-based Memories

Spreading activation causes interference.

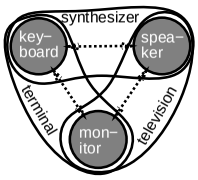


Two systems needed:

- Interconnected posterior cortex.
- Isolated prefrontal cortex.

## 6 Cortical Priming

Even slow cortical weight changes can yield one-trial learning effects..

win____

handle
winter
shower...
win____

Spell /*rēd*/.
Name a musical instrument that uses a reed.
Spell /*rēd*/.

## 7 Cortical Priming

- There are many, many types of priming effects:
  - Stem-completion & phonetic priming
  - Perceptual identification → faster, more accurate detection after recent exposure to words (even hours later)
  - Category generation priming: *"peach, kiwi"*; <many hrs later> → *"name some fruits"* (in absence of recall)
- Cortex is the key substrate for these priming effects
- Patients with hippocampus damage (sparing cortex) show impaired recall but intact priming
- These priming effects are long-lasting
  - This indicates that a weight change is involved (unlikely for activations to persist for long periods)

## 8 Simulations of Cortical Priming

- Train a network to learn input-output mappings
- Each input is associated with two valid outputs
- Analogous to:

win____ → window
win____ → winter
/rēd/ → "read"
/rēd/ → "reed"

## 9 Weight-based Priming Model

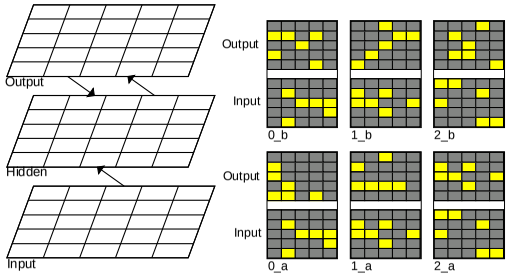


## 10 [wt_priming.proj]

## 11 Priming Simulations

- After training, the network is equally likely to produce the "a" or "b" output in response to a cue...
- Does not "blend" the two, but instead settles into one of the two valid attractors
- How does one additional study trial with the "a" input affect performance?
- Small weight changes (resulting from a single study trial) can "tip the balance" in favor of the recently studied response...

## 12 Priming Data

| trial | Event | sum_se | Outp_dist | ev_nm | sm_nm | both_err |
|---|---|---|---|---|---|---|
| 0 | 0_a | 5.22935 | 0 | 0_b | 1 | 0 |
| 1 | 1_a | 6.48608 | 0 | 1_b | 1 | 0 |
| 2 | 2_a | 7.77501 | 0.273233 | 2_b | 1 | 0 |
| 3 | 3_a | 7.64788 | 0 | 3_b | 1 | 0 |
| 4 | 4_a | 5.41569 | 0.551383 | 4_b | 1 | 0 |
| 5 | 5_a | 0 | 0 | 5_a | 0 | 0 |
| 6 | 6_a | 10.2454 | 0 | 6_b | 1 | 0 |
| 7 | 7_a | 8.33851 | 0 | 7_b | 1 | 0 |
| 8 | 8_a | 5.64973 | 2.61438 | 8_b | 1 | 0 |
| 9 | 9_a | 10.2408 | 0 | 9_b | 1 | 0 |
| 10 | 10_a | 3.21385 | 1.06278 | 10_b | 1 | 0 |
| 11 | 11_a | 2.82117 | 2.42077 | 11_b | 1 | 0 |
| 12 | 12_a | 4.69916 | 0.253711 | 12_b | 1 | 0 |
| 13 | 0_b | 6.68981 | 0 | 0_a | 1 | 0 |
| 14 | 1_b | 5.40769 | 0.330821 | 1_a | 1 | 0 |
| 15 | 2_b | 7.51547 | 0 | 2_a | 1 | 0 |
| 16 | 3_b | 7.73557 | 0 | 3_a | 1 | 0 |
| 17 | 4_b | 1.94789 | 1.94789 | 4_b | 0 | 0 |
| 18 | 5_b | 0.414954 | 0.414954 | 5_b | 0 | 0 |
| 19 | 6_b | 10.5514 | 0 | 6_a | 1 | 0 |
| 20 | 7_b | 8.79166 | 0 | 7_a | 1 | 0 |
| 21 | 8_b | 9.64561 | 0 | 8_a | 1 | 0 |
| 22 | 9_b | 10.2245 | 0 | 9_a | 1 | 0 |
| 23 | 10_b | 3.53423 | 0.766472 | 10_a | 1 | 0 |
| 24 | 11_b | 7.46935 | 0 | 11_a | 1 | 0 |
| 25 | 12_b | 5.72054 | 0 | 12_a | 1 | 0 |

## 13 Cortical Priming

Residual activation can also result in priming.
(Activation-based priming: later)

Three factors:

- Duration (short-term activations vs long-term weights).
- Content (visual, semantic, etc.)
- Similarity (repetition, semantic relation, etc).

## 14 Memory: Rapid Learning, Interference, & The Hippocampus

1. AB-AC List Learning
2. The Hippocampus.

## 15 AB-AC List Learning

Humans can rapidly learn overlapping associations without too much interference.

Example: learn one set of paired associates (the A-B list):
window-reason
bicycle-garbage
... Then, learn overlapping set (the A-C list):
window-locomotive
bicycle-dishtowel

## 16 AB-AC List Learning

Then test on AB list:
window- ?
bicycle- ?

and on AC list:
window- ?
bicycle- ?

## 17 AB-AC List Learning

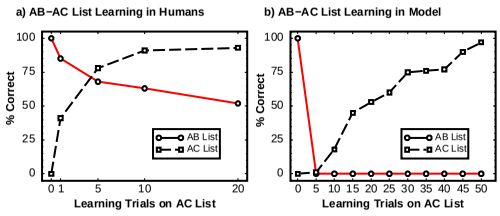


Standard network shows *catastrophic* interference (McCloskey & Cohen, 1989).

## 18 AB-AC Exploration

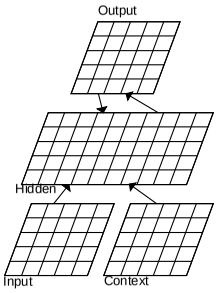


Input = A, Output = B,C

Context differentiates the lists: Each list is associated with a different context pattern

## 19 [ab_ac_interference.proj]

## 20 AB-AC Simulations: Summary

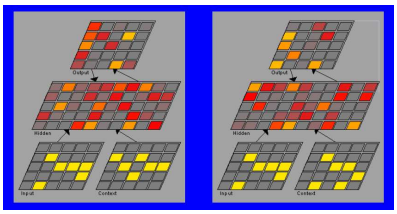

- There is overlap between the hidden units activated by an input pattern ("window") in the AB context and units activated by that same pattern in the AC context.
- This causes *interference* (changing weights for one changes weights for the other)
- Can this be fixed?

## 21 How to reduce Interference?

- How can we reduce overlap between hidden units activated by patterns in the AB and AC contexts?
- → Lower the number of units that are activated → increase inhibition (decrease $k$)...

But still need ***different*** units to be active for AB and AC inputs...

- → Increase relative weight scale of the context layer so that hidden units "pay more attention" to it
- → Also increase initial weight *variance*: Lowers the odds that a unit will "like" both the AB and AC version of a pattern

## 22 AB-AC Exploration: Summary

- Note that ***even with all these changes***, interference gets only slightly better...
- Also network learns much slower than people do... → Increase learning rate?
- This speeds up learning, but makes interference **worse**!
- Also, by changing all these parameters, cortex can no longer generalize (requires overlapping distributed representations)

→ *Trade-off:* Must need another brain system!

## 23 Memory

Memory is not unitary.

1. Weights versus activations.
2. Specialized neural systems: computational tradeoffs.

## 24 Memory

Memory is not unitary.

1. Weights (long-lasting, requires re-activation) versus activations (short-term, already active, can influence processing).
2. Specialized neural systems: computational tradeoffs. Cortex shows priming, but suffers catastrophic interference. Abandon neural network models?

## 25 Hippo To the Rescue

Two specialized, complementary systems resolve fundamental tradeoff:

The hippocampus can learn rapidly without interference by using sparse, pattern-separated representations!

Meanwhile, cortex slowly learns overlapping representations of similarity structure & regularities, semantic knowledge.

e.g. "one small step for man" 9/11, etc

## 26 Complementary Learning Systems

| Goals: | Remember Specifics | Extract Generalities |
|---|---|---|
| Example: | Where is car parked? | Best parking strategy? |
| Need to: | Avoid interference | Accumulate experience |
| | *Solution:* | |
| 1. | Separate reps (keep days separate) 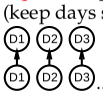  | Overlapping reps (integrate over days) 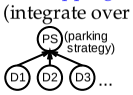  |
| 2. | Fast learning (encode immediately) | Slow learning (integrate over days) |
| 3. | Learn automatically (encode everything) | Task-driven learning (extract relevant stuff) |
| | *These are incompatible, need two different systems:* | |
| System: | Hippocampus | Neocortex |

## 27 Systematic Overlap Develops by Slowly Integrating over Experience

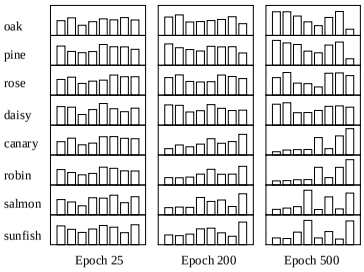


## 28 Effects of hippocampal damage in Amnesia

Amnesics show:

- spared implicit memory, skill learning (without recall)
  *small adaptive adjustments in synaptic weights*
- intact repetition priming for existing associations (table-chair) but not for arbitrary novel pairs of words (locomotive-spoon)
  *small cortical adjustments can prime existing reps but not novel conjunction*
- remote memories spared but recent ones completely forgotten
  *"Consolidation": reactivation of memories across multiple contexts, sleep, etc*

## 29

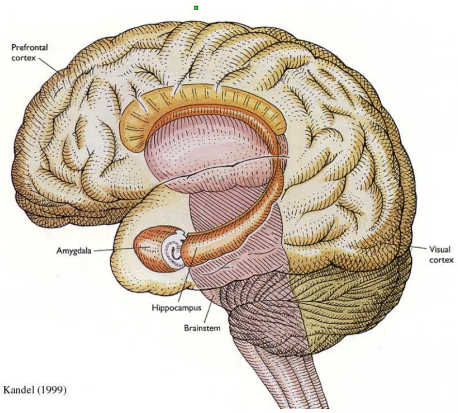


from Squire & Kandel (1999)

## 30

from DeArmond, Fusco, & Dewey (1989)

(Greek: hippo=horse, kampos=sea monster).

## 31 Hippo = King-of-the-Cortex

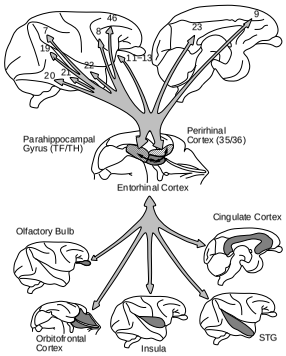


Hippo ***binds*** together multiple cortical representations into one coherent memory

## 32 Hippocampal Anatomy

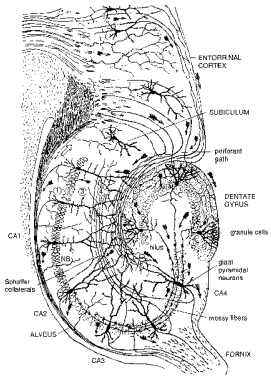


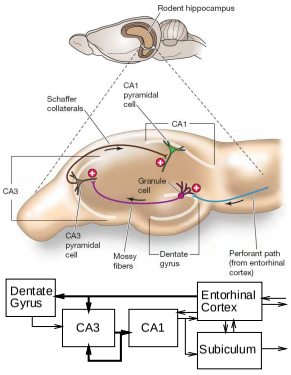


## 33 Hippocampal Anatomy

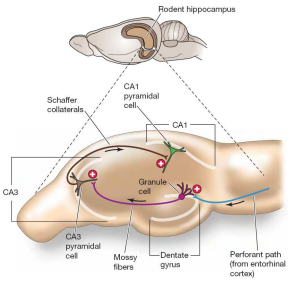


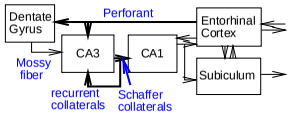


## 34 Pattern Separation & Conjunctions

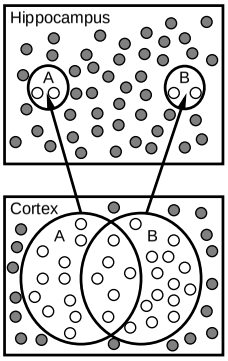


## 35 Explaining Pattern Separation

How does the hippocampus assign distinct representations to similar inputs?

- *Partial connectivity*: units are specialized for responding to a particular set of input features
- *Sparse activity*: fierce inhibitory competition
- Units only survive this competition if they receive a very large amount of excitatory input
- Units only fire if all features they detect are present in the input

→ Units represent *conjunctions* of features

## 36 Pattern Separation & Conjunctions: Space and episodes

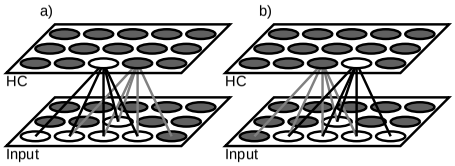


- Here each HC unit connected to 5 inputs; $k = 1$
- Changing one input unit causes a different HC unit to win!

## 37 Sparse Activity

| | Rat | | Model | |
|---|---|---|---|---|
| Area | Neurons | Pct Act | Units | Pct Act |
| EC | 200,000 | 7.0 | 144 | 25.0 |
| DG | 1,000,00 | 0.5 | 625 | 1.0 |
| CA3 | 160,000 | 2.5 | 240 | 5.0 |
| CA1 | 250,000 | 2.5 | 384 | 9.4 |

## 38 Sparse Activity

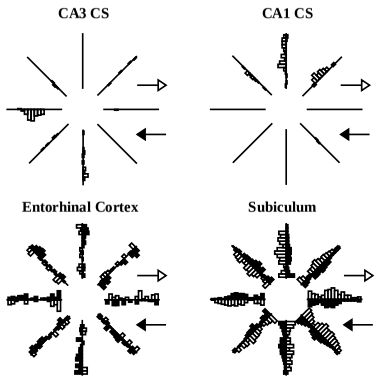


## 39 The Flip Side of Separation: Pattern Completion

College friend example: "This one time, at this one party..."

Pattern completion in CA3 activates corresponding CA1 rep, which reinstates original EC pattern...
→ "You told me this already!".

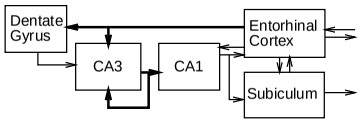


How does your hippo 'know' whether to store new memory and keep it separate, or instead complete to an existing memory?
→ hippo designed to minimize this tradeoff (LTP in CA3 supports pat complet while LTD supports pat sep; O'Reilly & McLelland '94).

## 40 The Model

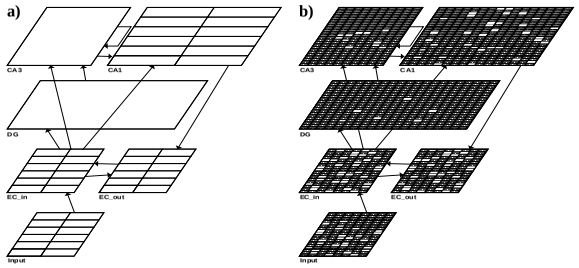


## 41 Hippo binding

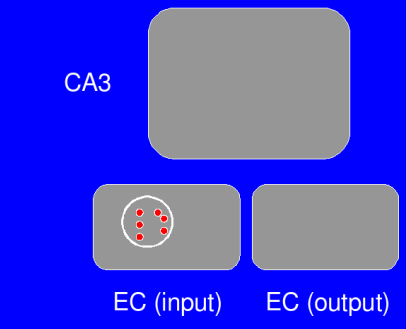

42

## Hippo binding

STUDY

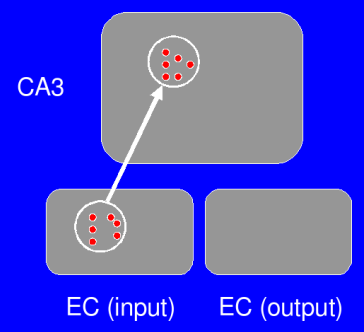


43

## Hippo binding

STUDY

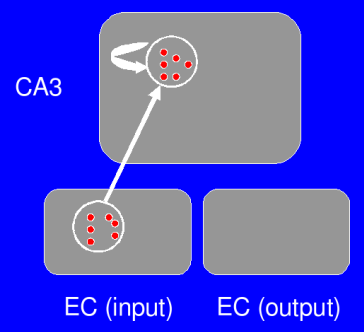


44

## Hippo binding

STUDY

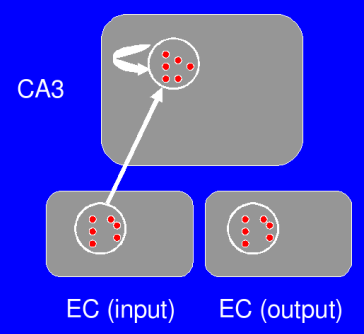


45

## Hippo binding

STUDY

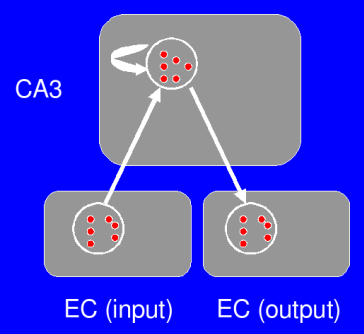


46

## Pattern Completion

TEST

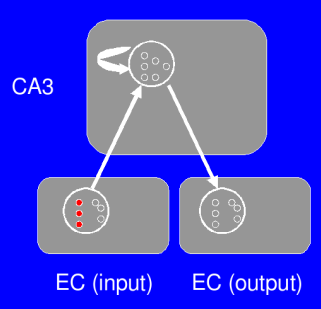


47

## Pattern Completion

TEST

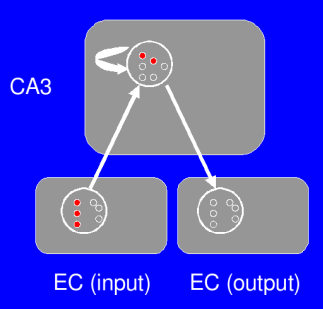

48

# Pattern Completion

TEST

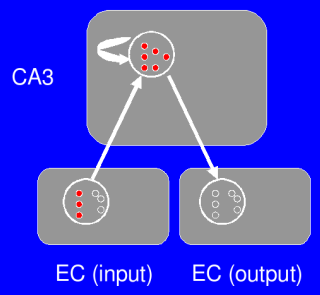


49

# Pattern Completion

TEST

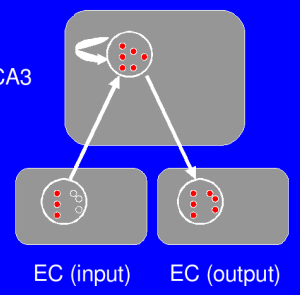


50

# Pattern Separation

STUDY

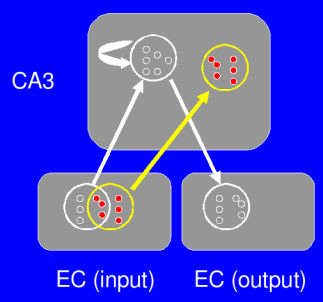


51

## DG = Pattern Separation Turbocharger

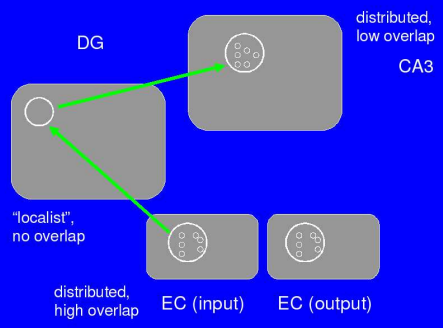


52

## DG = Pattern Separation Turbocharger

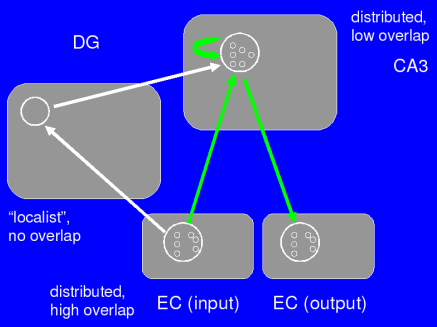


53

## DG = Pattern Separation Turbocharger

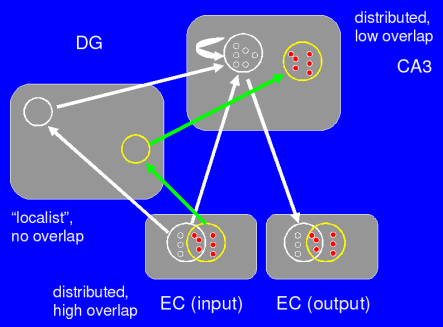

54

## DG = Pattern Separation Turbocharger

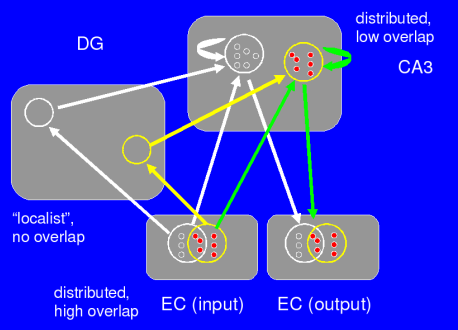


55

## DG = Pattern Separation Turbocharger

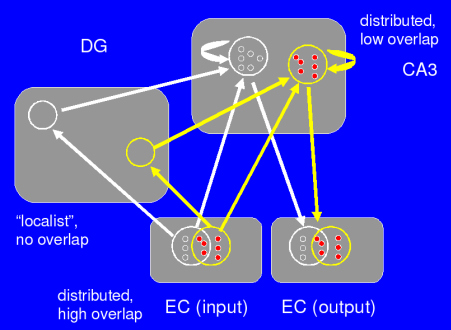


56

## DG = Pattern Separation Turbocharger

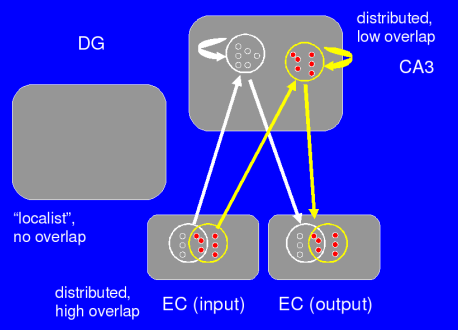


57

## DG is good for pattern separation, but it's bad for recall

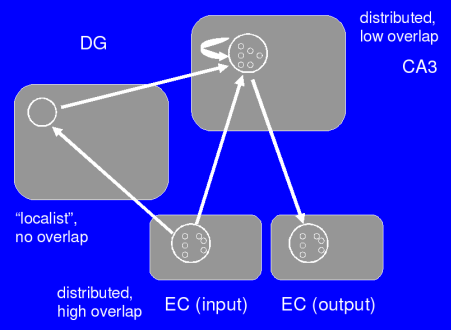


58

## DG is good for pattern separation, but it's bad for recall

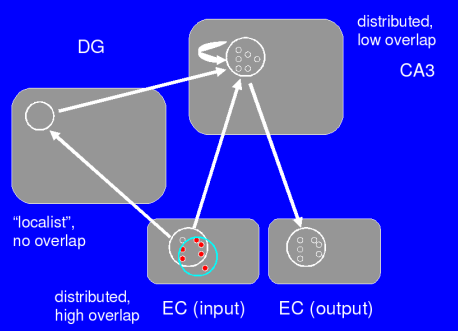


59

## DG is good for pattern separation, but it's bad for recall

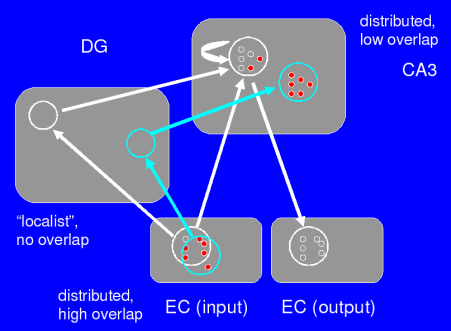

60

## Possible Solution: Two Modes

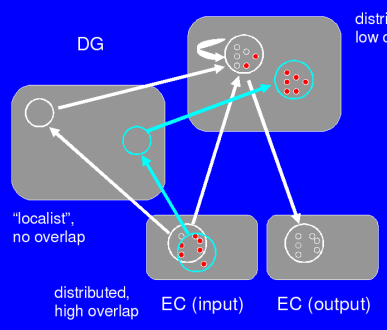


- Encoding mode: DG strong, facilitates pattern separation

61

## Possible Solution: Two Modes

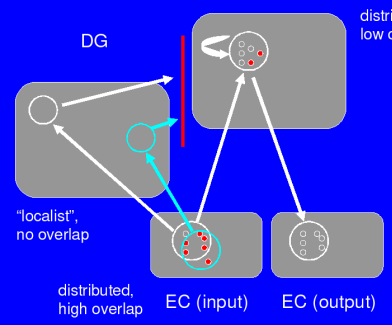


- Encoding mode: DG strong, facilitates pattern separation
- Retrieval mode: DG weak, facilitates recall

62

## Possible Solution: Two Modes

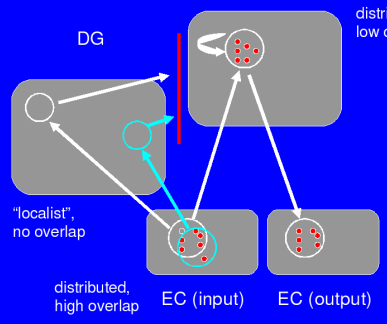


- Encoding mode: DG strong, facilitates pattern separation
- Retrieval mode: DG weak, facilitates recall

63

## Hippocampus: Summary

- CA3 stores sparse, pattern-separated representations of cortical input patterns
- Recurrent self projections in CA3 facilitate recall (pattern completion)
- Dentate Gyrus (DG) acts as a *removable pattern separation turbocharger*
  - DG uses super-sparse representations, helps increase pat separation at encoding
  - DG "steps aside" at retrieval
  - Evidence for two modes: theta cycle (eg. Hasselmo et al, 2002); neuromodulatory control over rel DG effect on CA3

64

## The Model

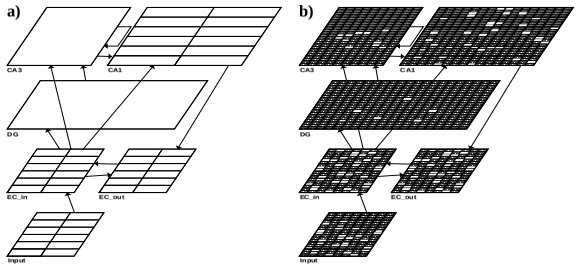


65

## An example of how Modeling informs Science

Model makes clear predictions about how different regions contribute to memory (not directly evident in experiments before)

Many of these have been subsequently confirmed! (note that model itself is incremental synthesis of many ideas in a coherent framework, ranging from Hebb to Marr to Nadel, McNaughton, O'Reilly...)

It has been applied to explain many different learning and memory phenomena in rats and humans.

## 66 An example of how Modeling informs Science

Pattern separation in Rat DG (*Leutgeb et al, 2007*, Science)
(change environment ever so slightly see new populations of correlated act)

Pattern separation in Human DG (*Bakker et al, 2008*, Science)
(encode new stims, some are similar to old but slightly diff).

Mouse genetic knockout of DG NMDA receptors impairs pat separation behaviorally; also CA3 becomes more biased toward completion than separation (*McHugh et al, 2007*, Science)

Monosynaptic route (EC→CA1→EC) sufficient on its own for incremental spatial learning, *but*
Trisynaptic route (EC→DG→CA3→CA1→EC) required for rapid, one-trial conjunctive learning, and for pat completion.
(Transgenic mouse *Nashashiba et al, 2008*, Science)

Neurogenesis in DG critical for behavioral pat sep *Clellan et al, 2009*, Science
(discriminate between items spatially close together but not far apart)

## 67 Role of CA1

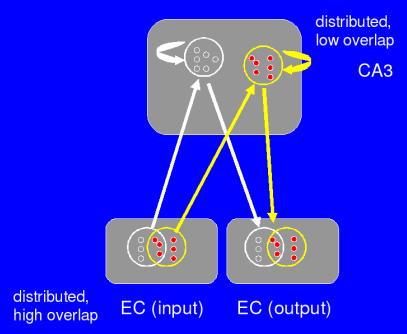


## 68 Role of CA1

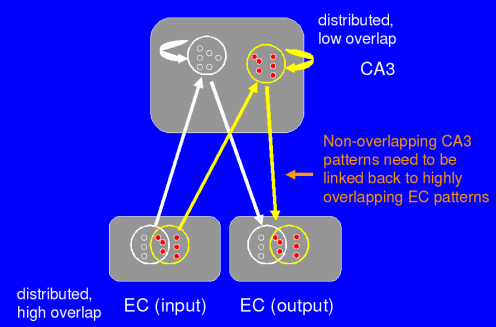


## 69 Role of CA1

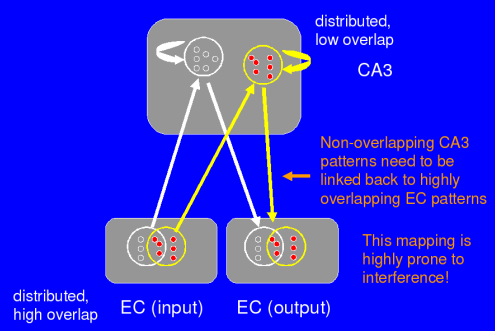


## 70 Role of CA1

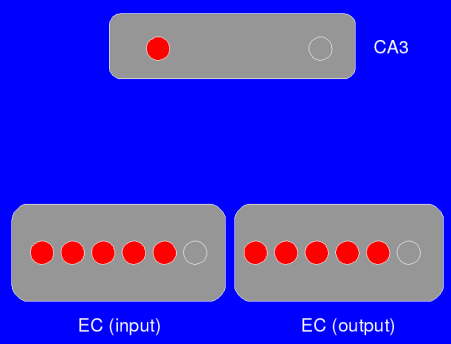

71

## Role of CA1

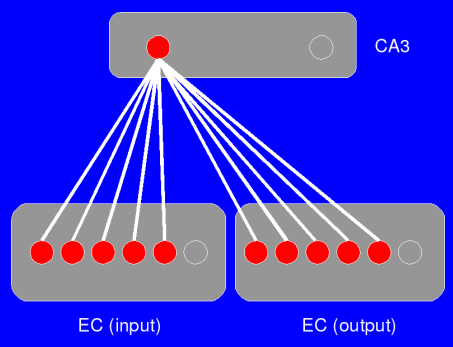


72

## Role of CA1

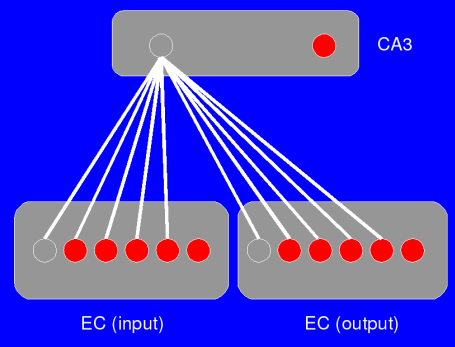


73

## Role of CA1

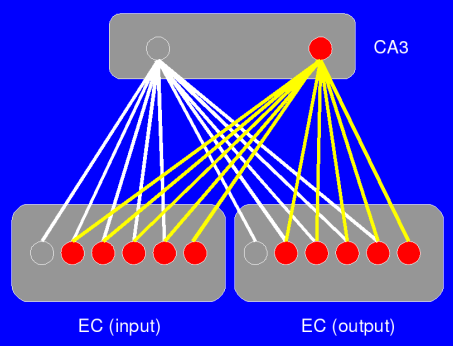


74

## Role of CA1

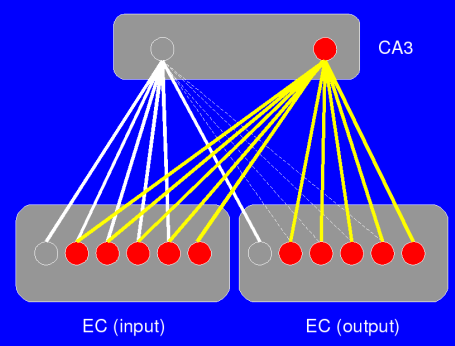


75

## Role of CA1

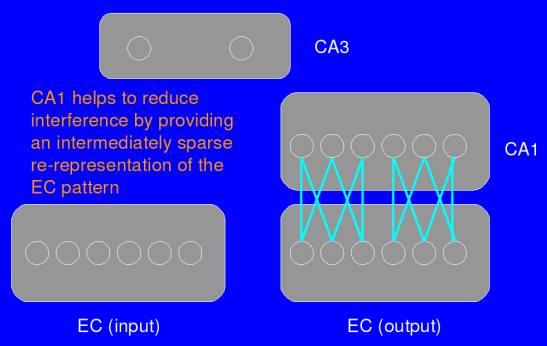


76

## Role of CA1

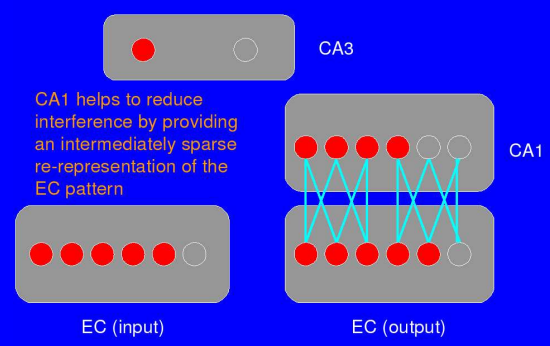

77

## Role of CA1

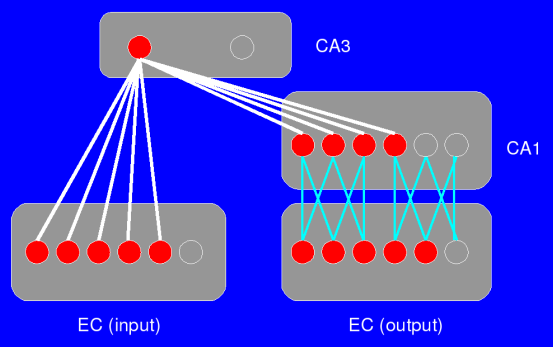


78

## Role of CA1

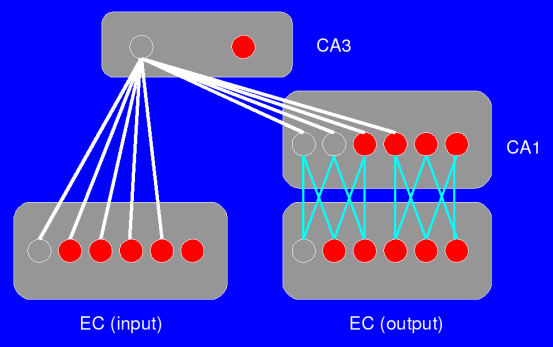


79

## Role of CA1

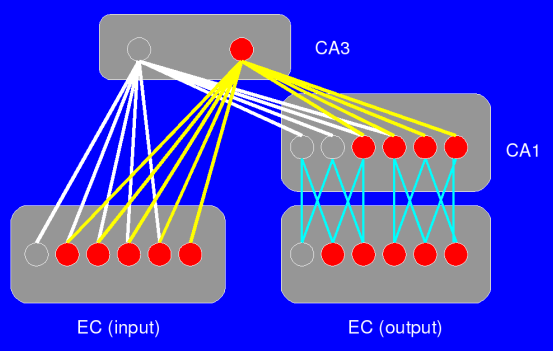


80

## Role of CA1

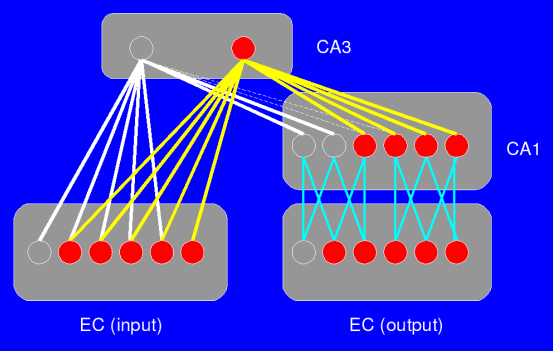


81

## Role of CA1

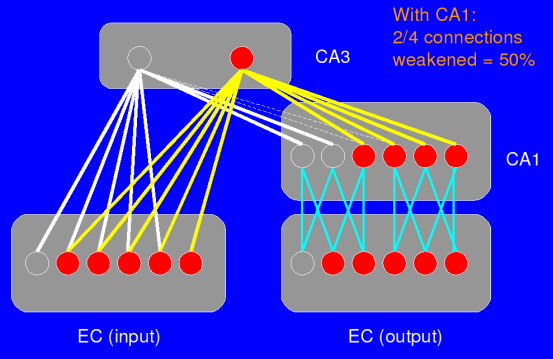


82

## Role of CA1

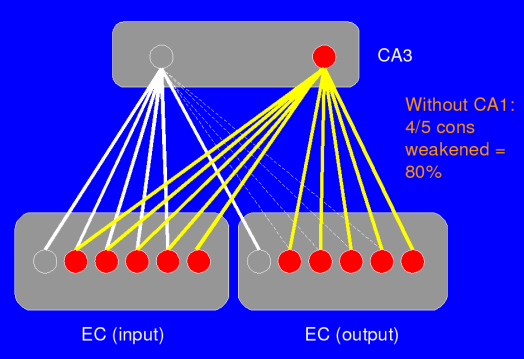

## 83 Hippocampus: Summary

- CA3 stores sparse, pattern-separated representations of cortical input patterns
- Recurrent self projections in CA3 facilitate recall (pattern completion)
- Dentate Gyrus (DG) acts as a *removable pattern separation turbocharger*
- *CA1 helps "translate" sparse, non-overlapping CA3 representations back into overlapping EC reps, by providing an* intermediately *sparse representation*

## 84 AB-AC Learning in the Hippo Model

[hip.proj]

- *Unlike cortical model, Hippocampus can rapidly and sequentially learn arbitrary information (AB-AC lists) without huge amounts of interference.*
- Cortex still critical for slow learning of overlapping, distributed representations, supporting *generalized* knowledge, semantic information, and similarity.
- Later: How learning/memory capacity can be enhanced with theta waves (Ken Norman)

## 85 Memory

Memory is not unitary.

1. Weights (long-lasting, requires re-activation) versus activations (short-term, already active, can influence processing).
2. Specialized neural systems: computational tradeoffs. Cortex shows priming, but suffers catastrophic interference. Abandon neural network models? *No, hippocampus can learn rapidly without interference using sparse, pattern-separated representations.*
3. Next time: Activation-based memory and activation-weight-based interactions.

## 86 Memory

Memory is not unitary.

1. Weights (long-lasting, requires re-activation) versus activations (short-term, already active, can influence processing).
2. Weight-based: Cortex shows priming, but suffers catastrophic interference. Hippocampus can learn rapidly without interference using sparse, pattern-separated representations.
3. Activation-based: Cortex shows priming, but can't do working memory.
4. Activation- and weight-based interactions.

## 87 Hippo and spatial topography: what about "grid cells"?

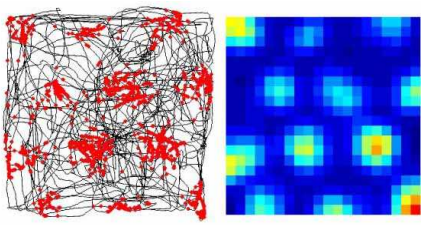

- Grid cells are in medial entorhinal cortex (Hafting et al, 2005), not hippo proper
- Hippo might integrate location with speed and direction ("head direction cells") to perform *path integration*

- This can be recast as just another example of conjunctive, pattern-separate representations

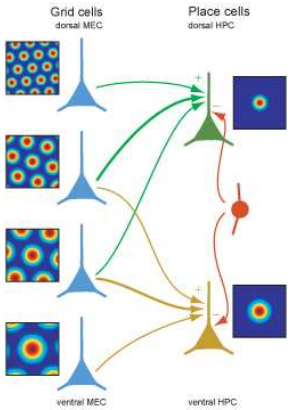


Solstad et al, 2006

## 88 But first: Why does this happen?

A-not-B

Kai A-not-B

Max Card Sort

## 89 Remember Weight-Based Priming?

| trial | Event | sum_se | Outp_dist | ev_nm | sm_nm | both_err |
|---|---|---|---|---|---|---|
| 0 | 0_a | 5.22935 | 0 | 0_b | 1 | 0 |
| 1 | 1_a | 6.48608 | 0 | 1_b | 1 | 0 |
| 2 | 2_a | 7.77501 | 0.273233 | 2_b | 1 | 0 |
| 3 | 3_a | 7.64788 | 0 | 3_b | 1 | 0 |
| 4 | 4_a | 5.41569 | 0.551383 | 4_b | 1 | 0 |
| 5 | 5_a | 0 | 0 | 5_a | 0 | 0 |
| 6 | 6_a | 10.2454 | 0 | 6_b | 1 | 0 |
| 7 | 7_a | 8.33851 | 0 | 7_b | 1 | 0 |
| 8 | 8_a | 5.64973 | 2.61438 | 8_b | 1 | 0 |
| 9 | 9_a | 10.2408 | 0 | 9_b | 1 | 0 |
| 10 | 10_a | 3.21385 | 1.06278 | 10_b | 1 | 0 |
| 11 | 11_a | 2.82117 | 2.42077 | 11_b | 1 | 0 |
| 12 | 12_a | 4.69916 | 0.253711 | 12_b | 1 | 0 |
| 13 | 0_b | 6.68981 | 0 | 0_a | 1 | 0 |
| 14 | 1_b | 5.40769 | 0.330821 | 1_a | 1 | 0 |
| 15 | 2_b | 7.51547 | 0 | 2_a | 1 | 0 |
| 16 | 3_b | 7.73557 | 0 | 3_a | 1 | 0 |
| 17 | 4_b | 1.94789 | 1.94789 | 4_b | 0 | 0 |
| 18 | 5_b | 0.414954 | 0.414954 | 5_b | 0 | 0 |
| 19 | 6_b | 10.5514 | 0 | 6_a | 1 | 0 |
| 20 | 7_b | 8.79166 | 0 | 7_a | 1 | 0 |
| 21 | 8_b | 9.64561 | 0 | 8_a | 1 | 0 |
| 22 | 9_b | 10.2245 | 0 | 9_a | 1 | 0 |
| 23 | 10_b | 3.53423 | 0.766472 | 10_a | 1 | 0 |
| 24 | 11_b | 7.46935 | 0 | 11_a | 1 | 0 |
| 25 | 12_b | 5.72054 | 0 | 12_a | 1 | 0 |

## 90 Activation-Based Priming

Residual activation can also result in priming: *act_priming.proj*

No learning (wt changes), to see effects of activation alone.

## 91 Activation-based Priming: Residual Activation

| trial | Event | sum_se | Outp_dist | ev_nm | sm_nm | both_err |
|---|---|---|---|---|---|---|
| 0 | 0_a | 0 | 0 | 0_a | 0 | 0 |
| 1 | 0_b | 1.7529 | 1.7529 | 0_b | 0 | 0 |
| 2 | 1_a | 0 | 0 | 1_a | 0 | 0 |
| 3 | 1_b | 2.18947 | 2.06997 | 1_a | 1 | 0 |
| 4 | 2_a | 0 | 0 | 2_a | 0 | 0 |
| 5 | 2_b | 5.43822 | 0.467382 | 2_a | 1 | 0 |
| 6 | 3_a | 0 | 0 | 3_a | 0 | 0 |
| 7 | 3_b | 1.05335 | 1.05335 | 3_b | 0 | 0 |
| 8 | 4_a | 0 | 0 | 4_a | 0 | 0 |
| 9 | 4_b | 6.26163 | 0.663053 | 4_a | 1 | 0 |
| 10 | 5_a | 0 | 0 | 5_a | 0 | 0 |
| 11 | 5_b | 4.02698 | 2.36882 | 5_a | 1 | 0 |
| 12 | 6_a | 0 | 0 | 6_a | 0 | 0 |
| 13 | 6_b | 5.74102 | 2.00435 | 6_a | 1 | 0 |
| 14 | 7_a | 0 | 0 | 7_a | 0 | 0 |
| 15 | 7_b | 8.85609 | 0 | 7_a | 1 | 0 |
| 16 | 8_a | 0 | 0 | 8_a | 0 | 0 |
| 17 | 8_b | 9.4205 | 0.444151 | 8_a | 1 | 0 |
| 18 | 9_a | 0 | 0 | 9_a | 0 | 0 |
| 19 | 9_b | 7.888 | 1.64196 | 9_a | 1 | 0 |
| 20 | 10_a | 0 | 0 | 10_a | 0 | 0 |
| 21 | 10_b | 5.20613 | 0.337607 | 10_a | 1 | 0 |
| 22 | 11_a | 0 | 0 | 11_a | 0 | 0 |
| 23 | 11_b | 6.4702 | 1.40431 | 11_a | 1 | 0 |
| 24 | 12_a | 0 | 0 | 12_a | 0 | 0 |
| 25 | 12_b | 5.32969 | 0.33391 | 12_a | 1 | 0 |

But what about when need to maintain over longer delays (working memory)??

## 92 Prefrontal Cortex: Delay-related activity

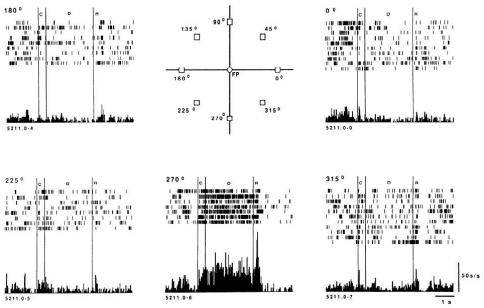


Spatial delayed-response task; Funahashi et al, 1989

## 93 Active Maintenance

Maintaining information in active form over longer time periods.

Can be used for *working memory* (e.g., in mental arithmetic).

Attractor = stable activation state:

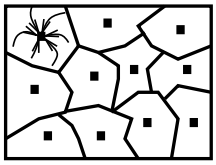

(don't want activity to spread)

## 94 Prefrontal vs. Posterior Cortex

Posterior cortex: interactive reps w/spreading activation

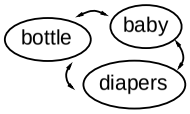


| *Advantages* | *Disadvantages* |
|---|---|
| Semantic associations | Memory confusion |
| Inference (diapers → baby) | |
| Schema (parenting) | |

Prefrontal: isolated reps, maintenance w/out activation spread

## 95 Attractors: Summary

- To get robustness from noise, you need *isolated* representations with strong recurrent connections
- This prevents activity from spreading
- Tradeoff #1: Preventing spreading activation (active maintenance) vs. allowing spreading activation (inference)
- Solution: Posterior cortex uses interconnected representations → spreading activation; prefrontal cortex (PFC) uses isolated reps → prevents spreading activation
- Evidence for isolated *stripes* in PFC (Levitt et al, 93; Pucak et al, 96)

## 96 Exploration: [act_maint.proj]

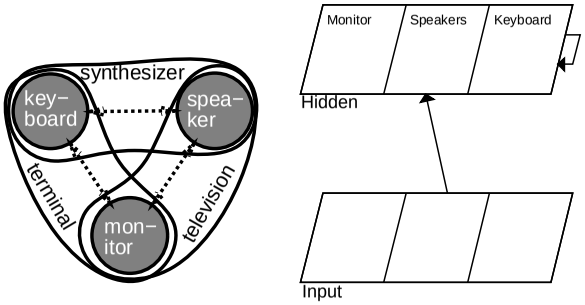


## 97 Weight- and Activation-Based Memory Interactions

A-not-B task

- Perseverative searching at A – also seen in patients with PFC damage
- Better peformance in gaze/expectation
- Inhibition problem?
- Model demonstrates maintenance problem.
- Same model accounts for various effects in different versions of A-not-B task not explained by any other unified theory (Munakata, 1998).

## 98 A-not-B Model

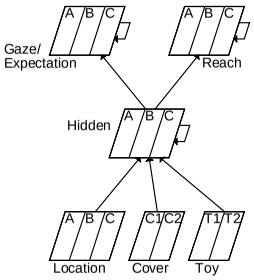

## 99 Knowledge-action dissociations in card-sort task

- Kids can tell you where trucks go in the shape game, even after sorting according to color!
- But if you ask "where do red trucks go in the shape game" they still fail! (Morton & Munakata, 2002)
- Explained by different levels of conflict experienced when faced with multiple stimuli-response associations..

## 100 Attractors: Summary

- Tradeoff #2: Within PFC, need for robust maintenance vs. need to update PFC activation when appropriate
- Strong recurrents (weak inputs) = robust maintenance
- Weak recurrents (strong inputs) = rapid updating
- We need a mechanism for switching PFC between the two modes
- Also, how to *learn* when to update?

## 101 Remember the SRN? (chap 6)

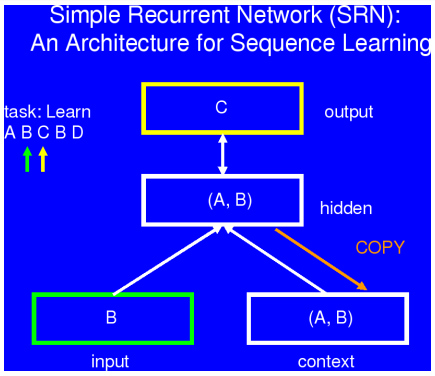


this is a *gating* network: context only updated at discrete timepoints

## 102 Simple SRN story is not flawless

- How is hidden→ "copy" function implemented biologically?
- During settling, context must be *actively maintained* (ongoing hidden activity has no effect on context).
- Assumes all context is relevant: What if distracting information presented in middle of sequence? Want to only hold on to *relevant* context.
- What if want to hold on to more than one piece of information in WM at a time?? Or to selectively update one part of WM while continuing to robustly maintain others?
- And what if the decision of whether or not to update information *depends* on currently internal WM state?

## 103 Working Memory Demands: Updating & Maintenance

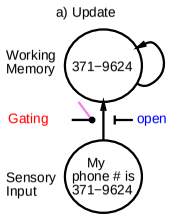


- Working memory: robust maintenance of information, but must also have ability to be rapidly updated — requires *gating*.
- You've got to know when to hold 'em, know when to fold 'em.

## 104 Working Memory Demands: Updating & Maintenance

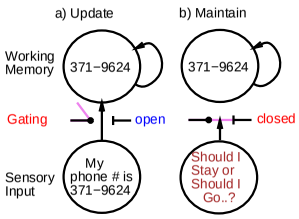


- Working memory: robust maintenance of information, but must also have ability to be rapidly updated — requires *gating*.
- You've got to know when to hold 'em, know when to fold 'em.

## 105 Working Memory Demands: Updating & Maintenance

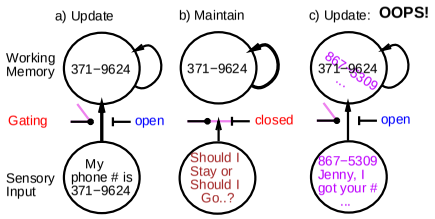


- Working memory: robust maintenance of information, but must also have ability to be rapidly updated — requires *gating*.
- You've got to know when to hold 'em, know when to fold 'em.

## 106 But who controls the controller??

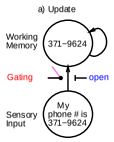


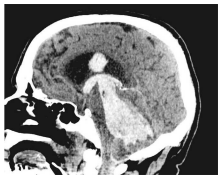

## 107 But who controls the controller??

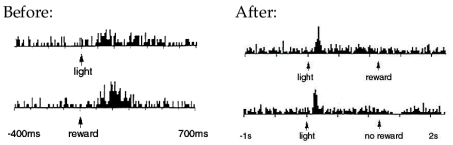


[First pass story (Braver & Cohen, '00 and text):]
*Dopamine* provides dynamic *gating* mechanism:

- Positive TD $\delta$ (reward) = DA burst = update PFC.
- No TD $\delta$ = constant DA = maintain PFC.
- Negative TD $\delta$ (error) = DA dip = clear PFC.

The *same* DA signal that learns to predict reward can be used to drive updating of PFC states!

## 108 DA solves part of the problem

- Learning signal for gating.
- But DA is very global signal projecting to all of PFC – sufficent for updating and maintaining one item at a time.
- How to *selectively* update some aspects of WM but not others?
- Also prev DA-PFC model had awkward catch-22 problem: the stimulus is only predictive of reward if it is maintained (ie in PFC). But then stim needs to be gated into PFC in the first place to generate DA!
- Solution: separate learning from gating... and link to now well established role of basal ganglia-thalamus in gating.

## 109 Dynamic Gating: Current Story

- DA signals are important for learning/knowing when to gate
- But actual gating signals are implemented via more complex circuit interactions with the Basal Ganglia Go/NoGo system
- DA used to train Go/NoGo system exactly like in the motor and simple decision making domains...
- BG-gating solves multiple computational and biological plausibility issues that are problematic with pure-DA based gating
- Goto BG_PFC_WM1.pdf slides for more info and evidence

## 110 A Simple WM Task

| Trial | Input | Maint | Output |
|---|---|---|---|
| 1 | STORE-A | A | A |
| 2 | IGNORE-B | A | B |
| 3 | IGNORE-C | A | C |
| 4 | IGNORE-D | A | D |
| 5 | RECALL | A | A |

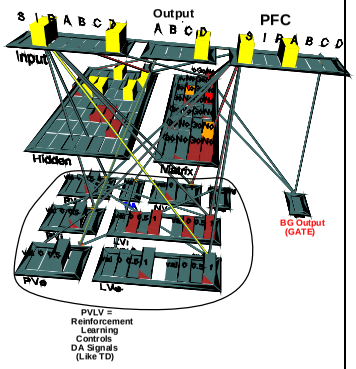

## 111 PFC/BG Model: pfc_maint_updt.proj

## 112 Reinforcement learning and WM gating

- Network learns to associate stimuli with rewards via DA system
- DA gets information not only from outside world, but also PFC state
- *Desired outcome:* Network learns that having the STORE pattern in PFC leads to rewards, but having the IGNORE pattern does not

## 113 Reinforcement learning and WM gating

- Bursts and dips of DA train the basal ganglia Go/NoGo gating system: DA burst→Go; DA dip→NoGo
- If BG system gates an input into PFC **and** that PFC pattern had been associated with reward → DA burst (DA system recognizes this new PFC pattern as rewarding)
- This DA burst reinforces *Go* activity in the BG units that caused the gating in the first place, making it even more likely that the BG will gate this pattern into PFC on *future trials*. (ie phasic DA does *not* directly drive updating itself, but is learning signal)
- *Desired outcome:* Networks learns "Go" to gate STORE into PFC, but learns "NoGo" to IGNORE

## 114 Sketch of how the network learns

- Begins with trial-and-error learning (both at response output and in BG gating system)
- Explore different gating "policies" and reinforce ones that work. (some amount noise helpful!)
- If correct response happens to occur when STORE is active in PFC (initially due to guessing) ⇒ Reward
- Resulting DA burst trains PVLV (or TD) system to learn that having "STORE" in PFC is a good thing

## 115 Sketch of how the network learns

- Next time STORE is represented in PFC, PVLV system triggers a DA burst, based on its learned PFC-reward association (without needing external reward)
- This DA burst drives Go learning in BG so that good stimuli are more likely to be gated
- In turn, stored information is more likely to be present in PFC during RECALL trial.
- At this point, Hidden layer simply has to learn to map PFC representation of stored stimulus to the Output response.
- This leads to increased rewards, further training gating system, and leading to stable state.

## 116 Four "Stripe" PFC/BG Model, Learns with DA

(O'Reilly & Frank, 2006)