Pattern Separation and Completion in the Hippocampus

Computational Models of Neural Systems
Lecture 3.5

David S. Touretzky
October, 2023

Overview

- Pattern separation
  - Pulling similar patterns apart reduces memory interference.
- Pattern Completion
  - Noisy or incomplete patterns should be mapped to more complete or correct versions.
- How can both functions be accomplished in the same architecture?
  - Use conjunction (codon units; DG) for pattern separation.
  - Learned weights plus thresholding gives pattern completion.
  - Recurrent connections (CA3) can help with completion as in Marr's model, but aren't used in the model described here.

Information Flow

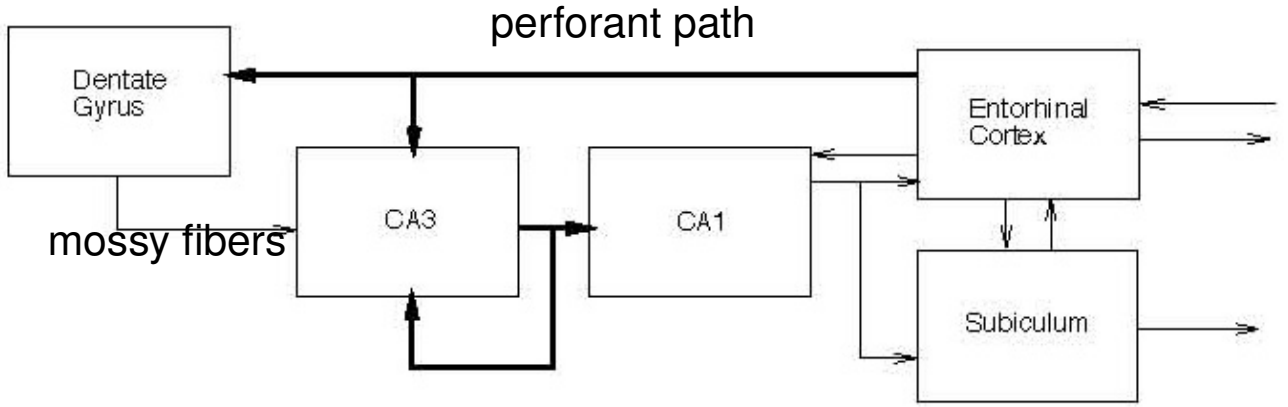

- Cortical projections from many areas form an EC representation of an event.
- EC layer II projects to CA3 (both directly and via DG), forming a new representation better suited to storage and retrieval.
- EC layer III projects to CA1, forming an invertible representation that can reconstitute the EC pattern.
- Learning occurs in all these connections.

Features of Hippocampal Organization

- Local inhibitory interneurons in each region.
  - May regulate overall activity levels, as in a kWTA network.
- CA3 and CA1 have less activity than EC and subiculum. DG has less activity than CA3/CA1.
  - Less activity means representation is more sparse, hence can be more highly orthogonal.

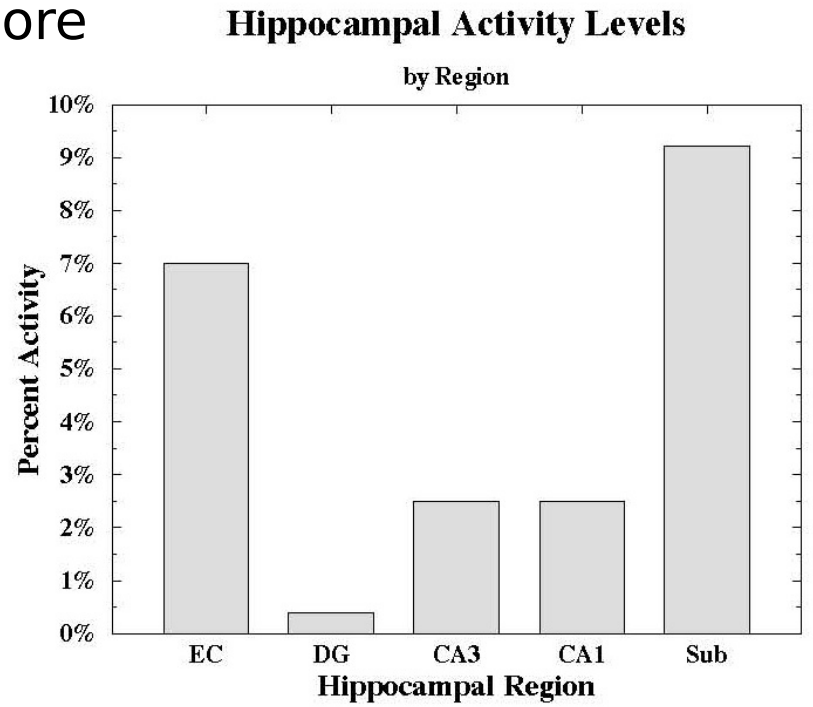

Connections in the Rat

- EC layer II (perf. path) projects diffusely to DG and CA3.
  - Each DG granule cell receives 5,000 inputs from EC.
  - Each CA3 pyramidal cell receives 3750-4500 inputs from EC. This is about 2% of the rat's 200,000 EC layer II neurons.
- DG has roughly 1 million granule cells.
  CA3 has 160,000 pyramidal cells; CA1 has 250,000.
- DG to CA3 projection (mossy fibers) is sparse and topographic. CA3 cells receive 52-87 mossy fiber synapses.
- NMDA-dependent LTP has been demonstrated in perforant path and Schaffer collaterals. LTP also demonstrated in mossy fiber pathway (non-NMDA).
- LTD may also be present in these pathways.

Model Parameters

- O'Reilly & McClelland investigated several models, starting with a simple two-layer k-WTA model (like Marr).
- N_i, N_o = # units in the layer
- k_i, k_o = # active inputs in one pattern
- α_i, α_o = fractional activity in the layer; $\alpha_o = k_o/N_o$
- F = fan-in of units in the output layer (must be < N_i)
- H_a = # of hits for pattern A

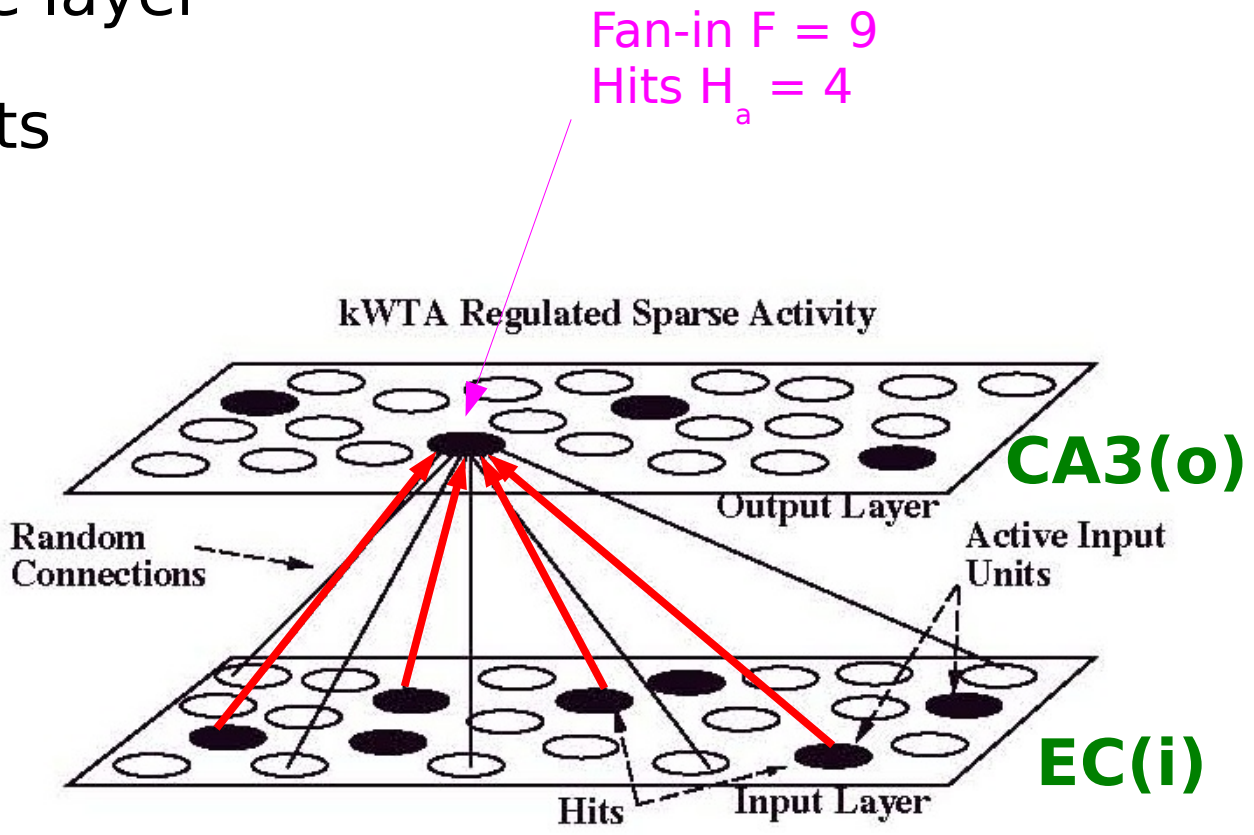

Measuring the Hits a Unit Receives

- How many input patterns?

$$\binom{N_i}{k_i}$$

- What is the expected number of hits H_a for an output unit?

$$\langle H_a \rangle = \frac{k_i}{N_i} F = \alpha_i F$$

- What is the distribution of hits, $P(H_a)$?

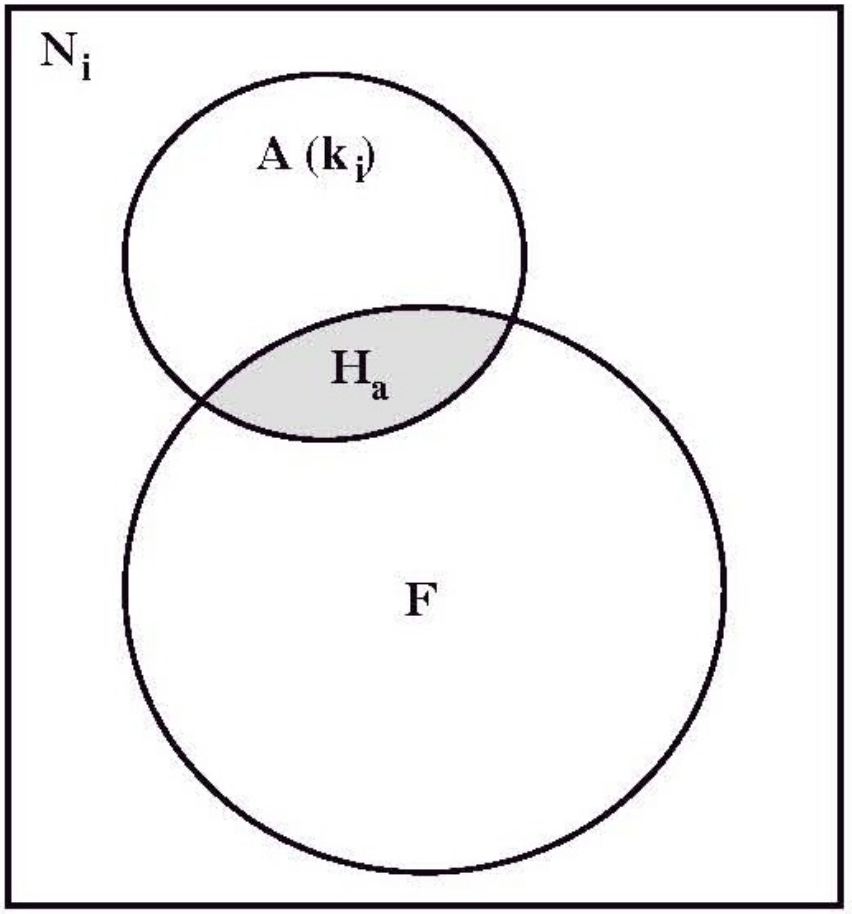

Hypergeomtric (not binomial; we're drawing without replacement)

Hypergeometric Distribution

- What is the probability of getting exactly H_a hits from an input pattern with k_i active units, given that the fan-in is F and the total input size is N_i?
  - $C(k_i, H_a)$ ways of choosing active units to be hits
  - $C(N_i\text{-}k_i, F\text{-}H_a)$ ways of choosing inactive units for the remaining ones sampled by the fan-in
  - $C(N_i, F)$ ways of sampling F inputs from a population of size N_i

$$P(H_a \mid k_i, N_i, F) = \frac{\binom{k_i}{H_a}\binom{N_i - k_i}{F - H_a}}{\binom{N_i}{F}}$$

# of ways to wire an output cell with H_a hits

# of ways to wire an output cell

Determining the kWTA Threshold

- Assume we want the output layer to have an expected activity level of α_o.
- Must set the threshold for output units to select the tail of the hit distribution. Call this H_a^t.
- Use the summation to choose H_a^t to produce the desired value of α_o.

$$\alpha_o = \sum_{H_a = H_a^t}^{\min(k_i, F)} P(H_a)$$

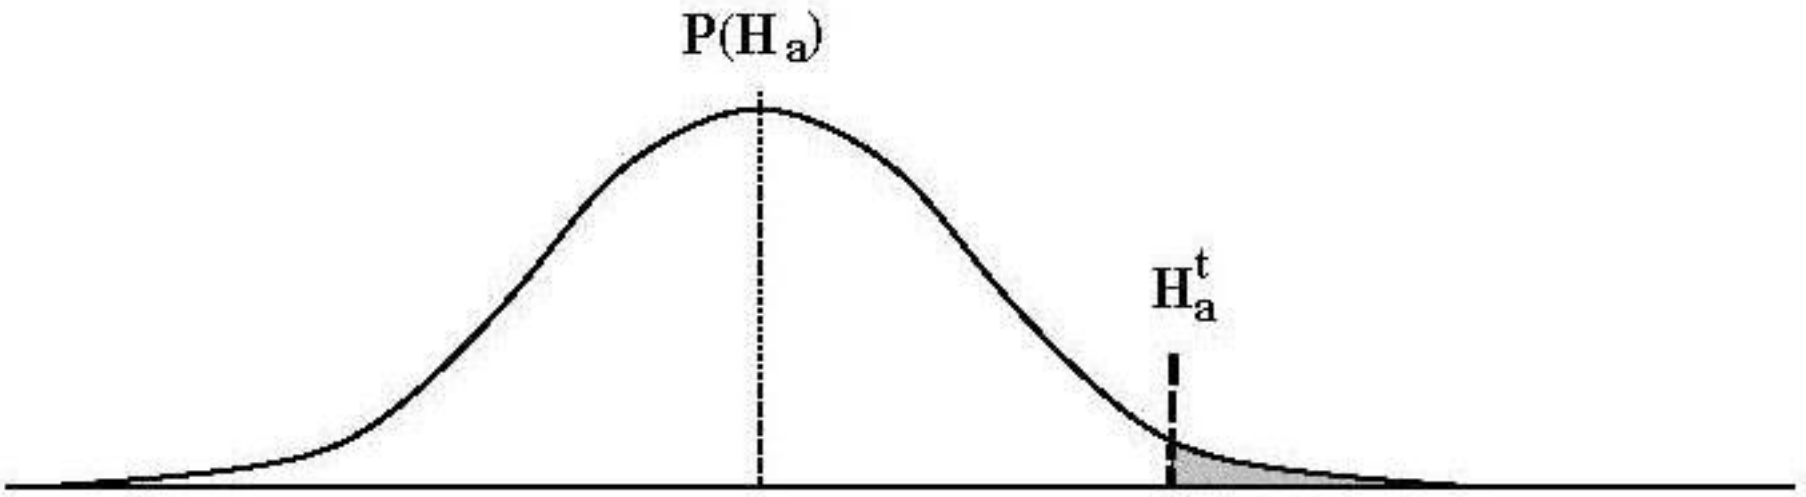

Pattern Overlap

- In order to measure pattern separation properties of the two-layer model, consider two patterns A and B.
  - Measure the input overlap Ω_i = number of units in common.
  - Compute the expected output overlap Ω_o as a function of Ω_i.
- If $\Omega_o < \Omega_i$ the model is doing pattern separation.

- To calculate output overlap we need to know H_{ab}, the number of hits an output unit receives for pattern B given that it is already known to be part of the representation for pattern A.

Distribution of H_{ab}

- For small input overlap, the patterns are virtually independent, and H_{ab} is distributed like H_a.
- As input overlap increases, H_{ab} moves rightward (more hits expected), and narrows: output overlap increases.
- But the relationship is nonlinear.

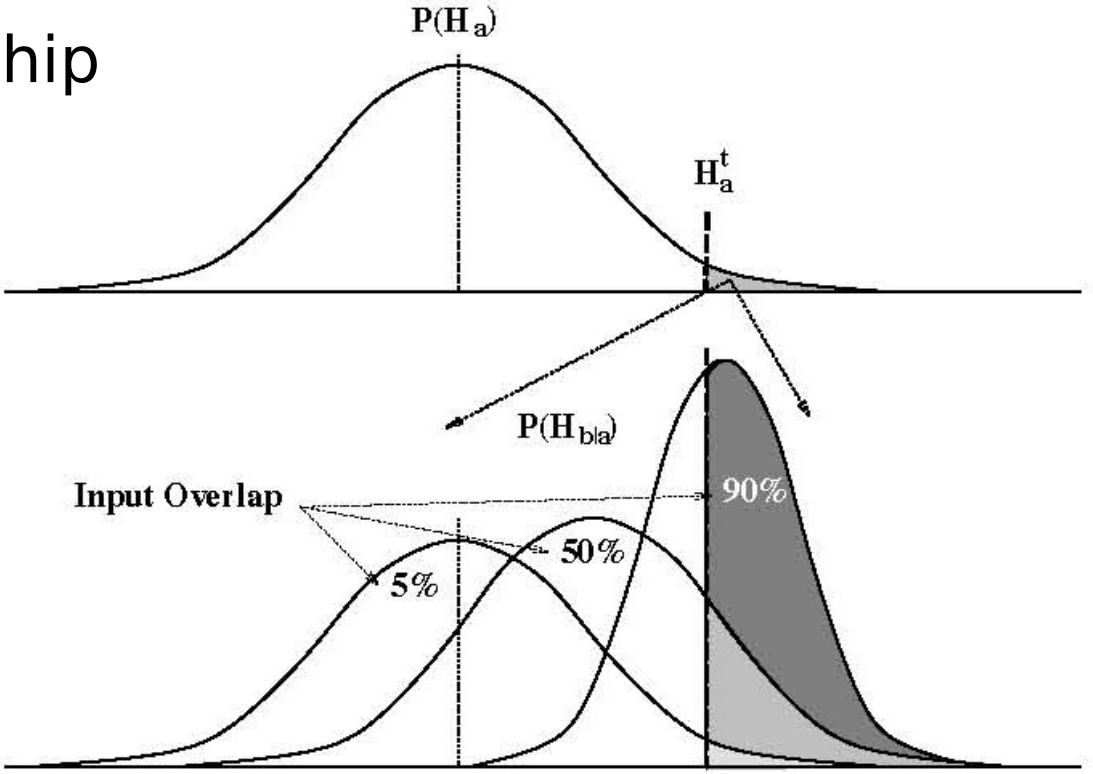

Visualizing the Overlap

a) Hits from pattern A.

b) H_{ab} = overlap of A&B hits

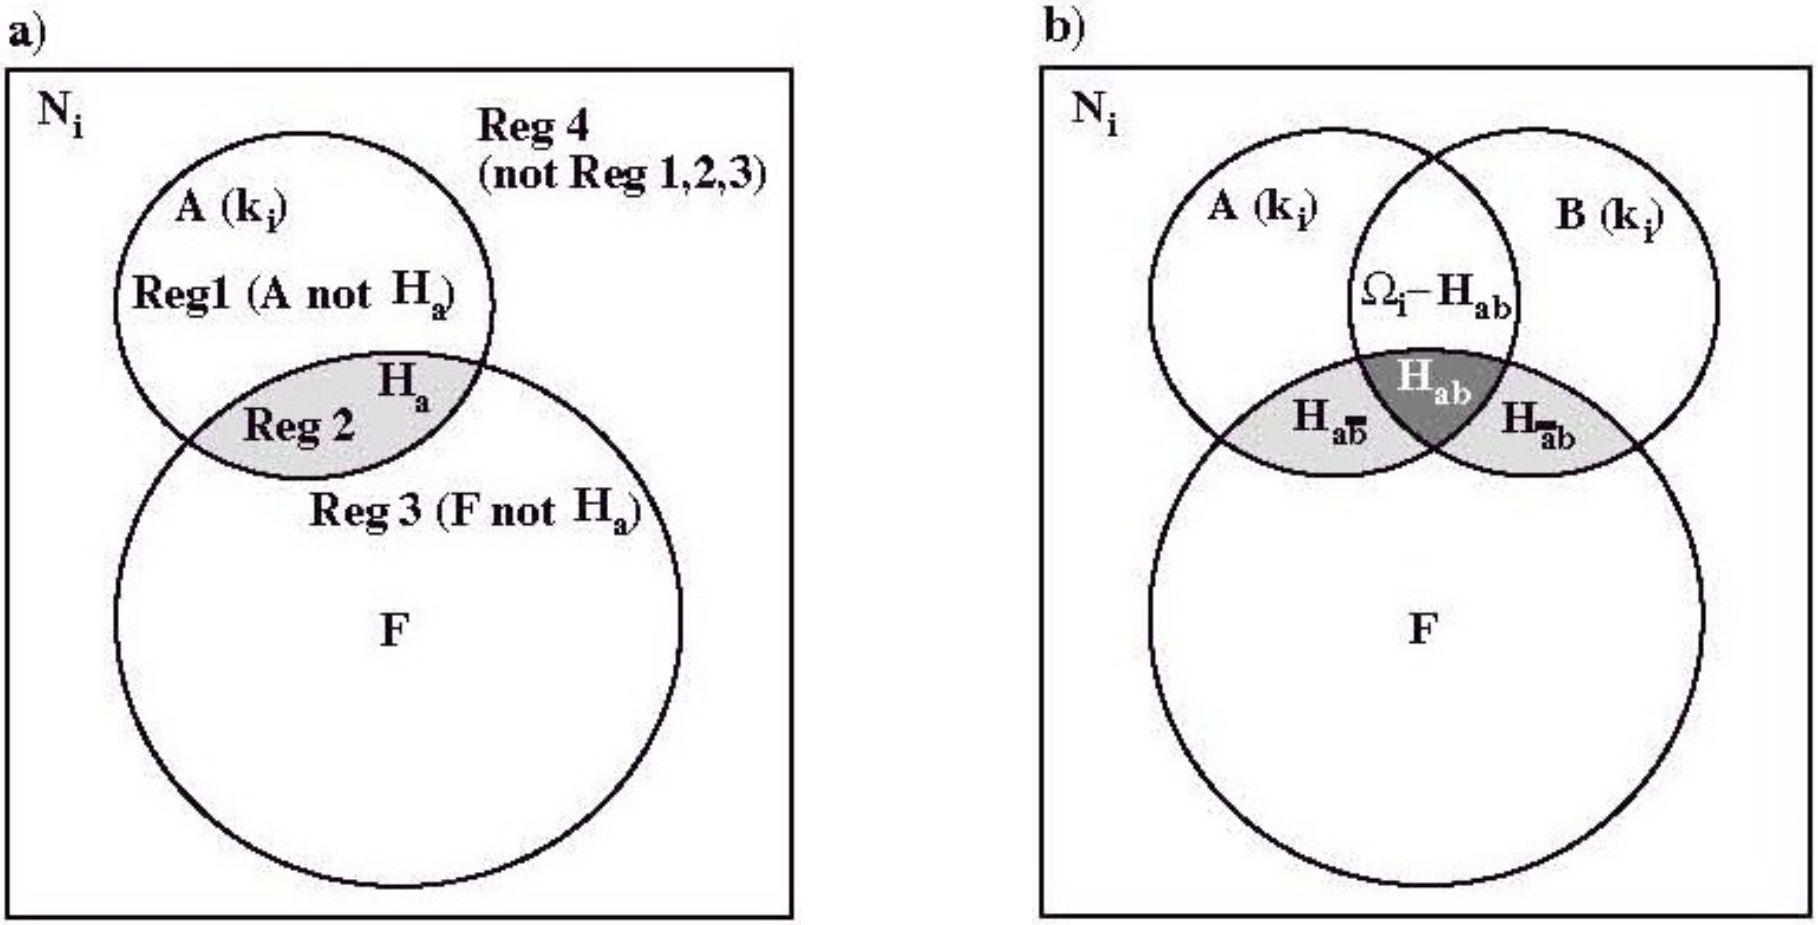

Prob. of b Hits And Specific Values for H_a, H_{ab}, $H_{\bar{a}b}$ Given Overlap Ω_i

$$P_b(H_a,\ \Omega_i,\ H_{ab},\ H_{\bar{a}b}) =$$

# of ways of achieving H_b hits given overlap Ω_i

$$\frac{\binom{H_a}{H_{ab}}\binom{F-H_a}{H_{\bar{a}b}}\binom{k_i-H_a}{\Omega_i-H_{ab}}\binom{N_i-k_i-F+H_a}{k_i-\Omega_i-H_{\bar{a}b}}}{\binom{k_i}{\Omega_i}\binom{N_i-k_i}{k_i-\Omega_i}}$$

1 2 3 4

# of ways of achieving $k_i - H_b$ non-hits given overlap Ω_i

# of ways of achieving overlap Ω_i

1,3 2,4

Note: $H_b = H_{ab} + H_{\bar{a}b}$

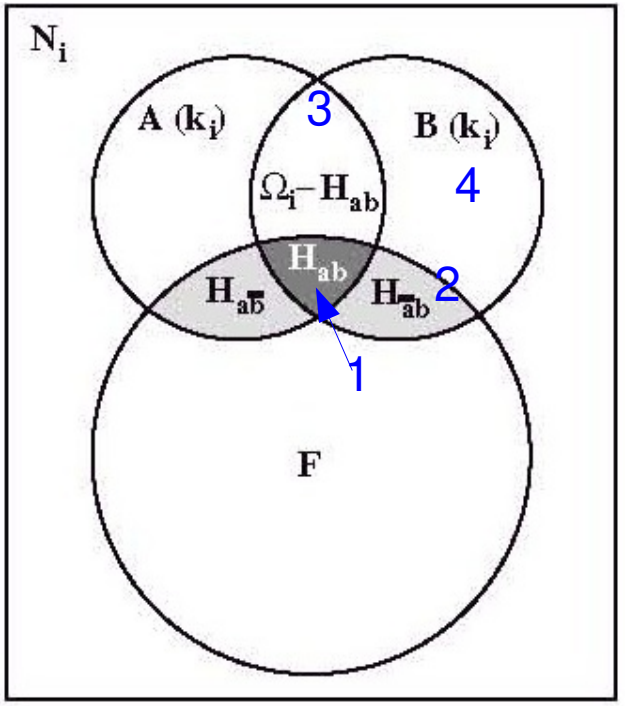

To calculate $P(H_b)$ we must sum P_b over all combinations of $H_a, H_{ab}, H_{\bar{a}b}$

Estimating Overlap for Rat Hippocampus

- We can use the formula for P_b to calculate expected output overlap as a function of input overlap.
- To do this for rodent hippocampus, O'Reilly & McClelland chose numbers close to the biology but tailored to avoid round-off problems in the overlap formula.

Area	α	N	F_{EC}	F_{DG}
EC	6.25%	200,000	—	—
DG	0.39%	850,000	4,006	—
CA3	2.42%	160,000	4,003	64

Estimated Pattern Separation in CA3

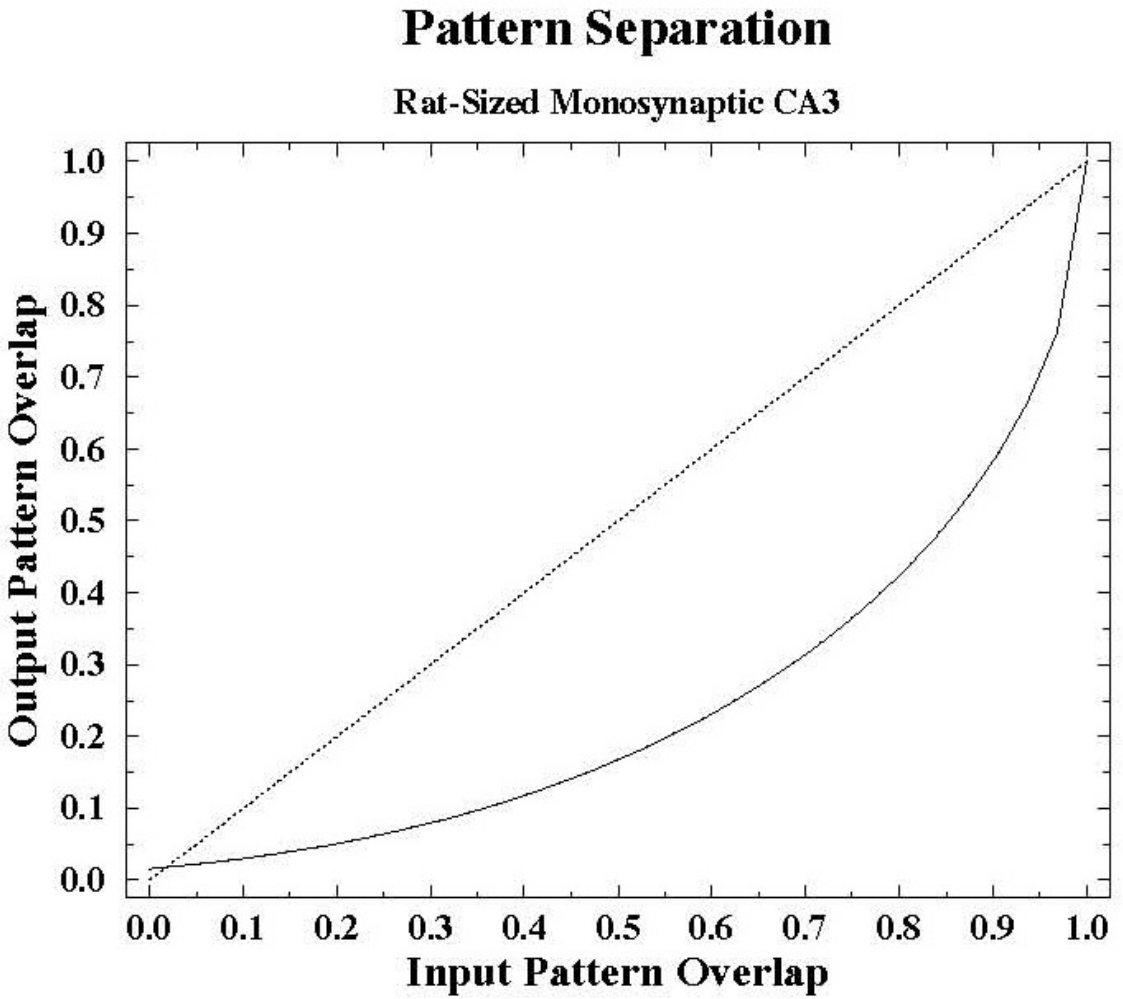

Sparsity Increases Pattern Separation

Pattern separation performance of a generic network with activity levels comparable to EC, CA3, or DG. Sparse patterns yield greater separation.

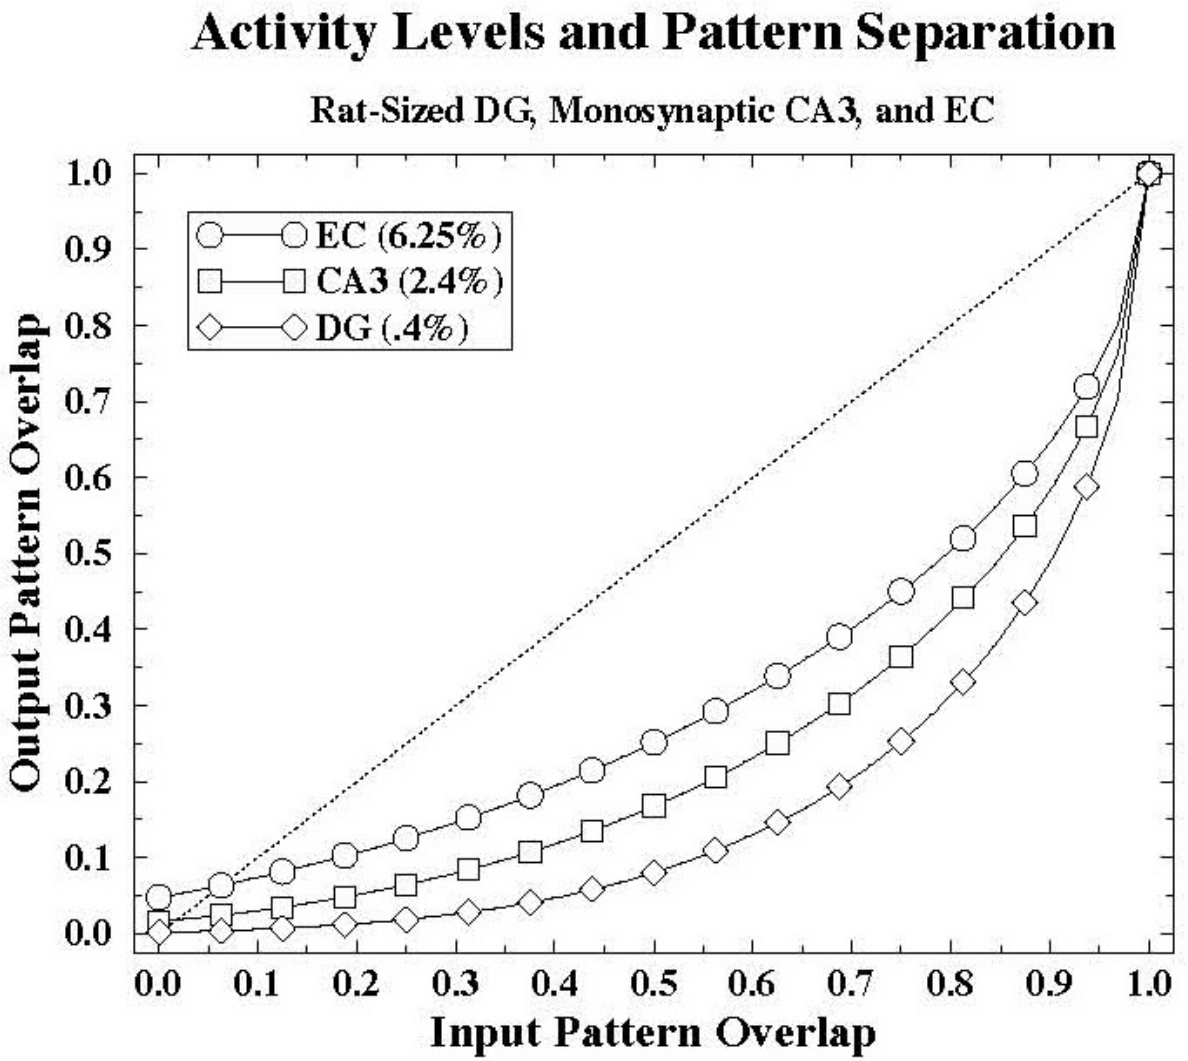

Fan-In Size Has Little Effect

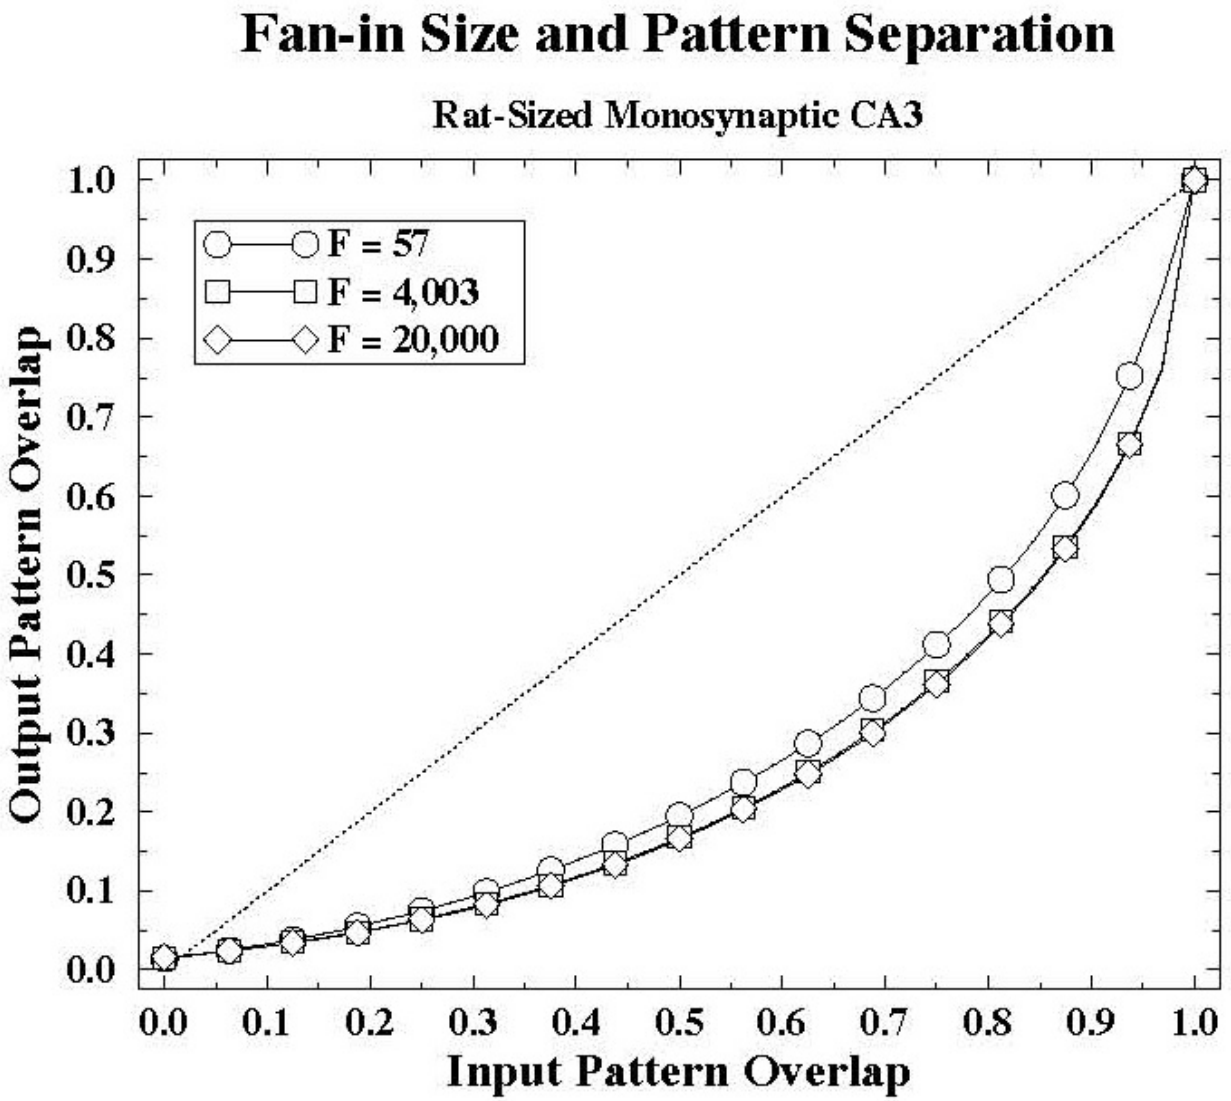

Adding Input from DG

- DG makes far fewer connections (64 vs. 4003), but they may have higher strength. Let M = mossy fiber strength.
- Separation in DG (*) is better than in CA3 without DG (○).
- DG connections help for M ≥ 15.
- With M=50, DG projection alone (+) is as good as DG+EC (▽).

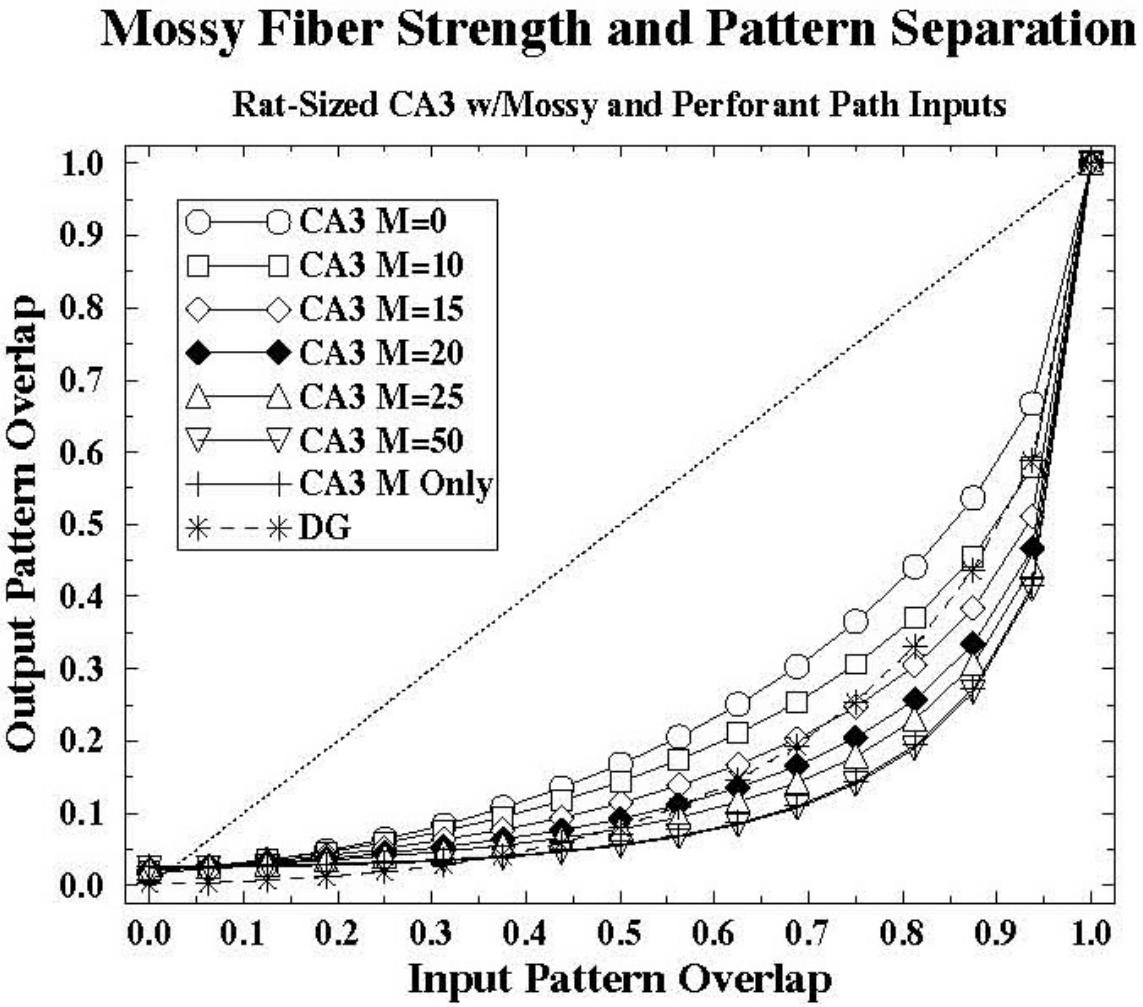

Combining Two Distributions

- CA3 has far fewer inputs from DG than from EC.
- But the DG input has greater variance in hit distribution.
- When combining two equally-weighted distributions, the one with the greater variance has the most effect on the tail.
- For 0.25 input overlap:
  - DG hit distribution has std. dev. of 0.76
  - EC hit distribution has std. dev. of 15.
  - Setting M=20 would balance the effects of the two projections.
- In the preceding plot, the M=20 line appears in between the M=0 line (EC only) and the “M only” line.

Without Learning, Partial Inputs Are Separated, Not Completed

Less separation between A and subset(A) than between patterns A and B, because there are no noise inputs.

But Ω_o is still less than Ω_i

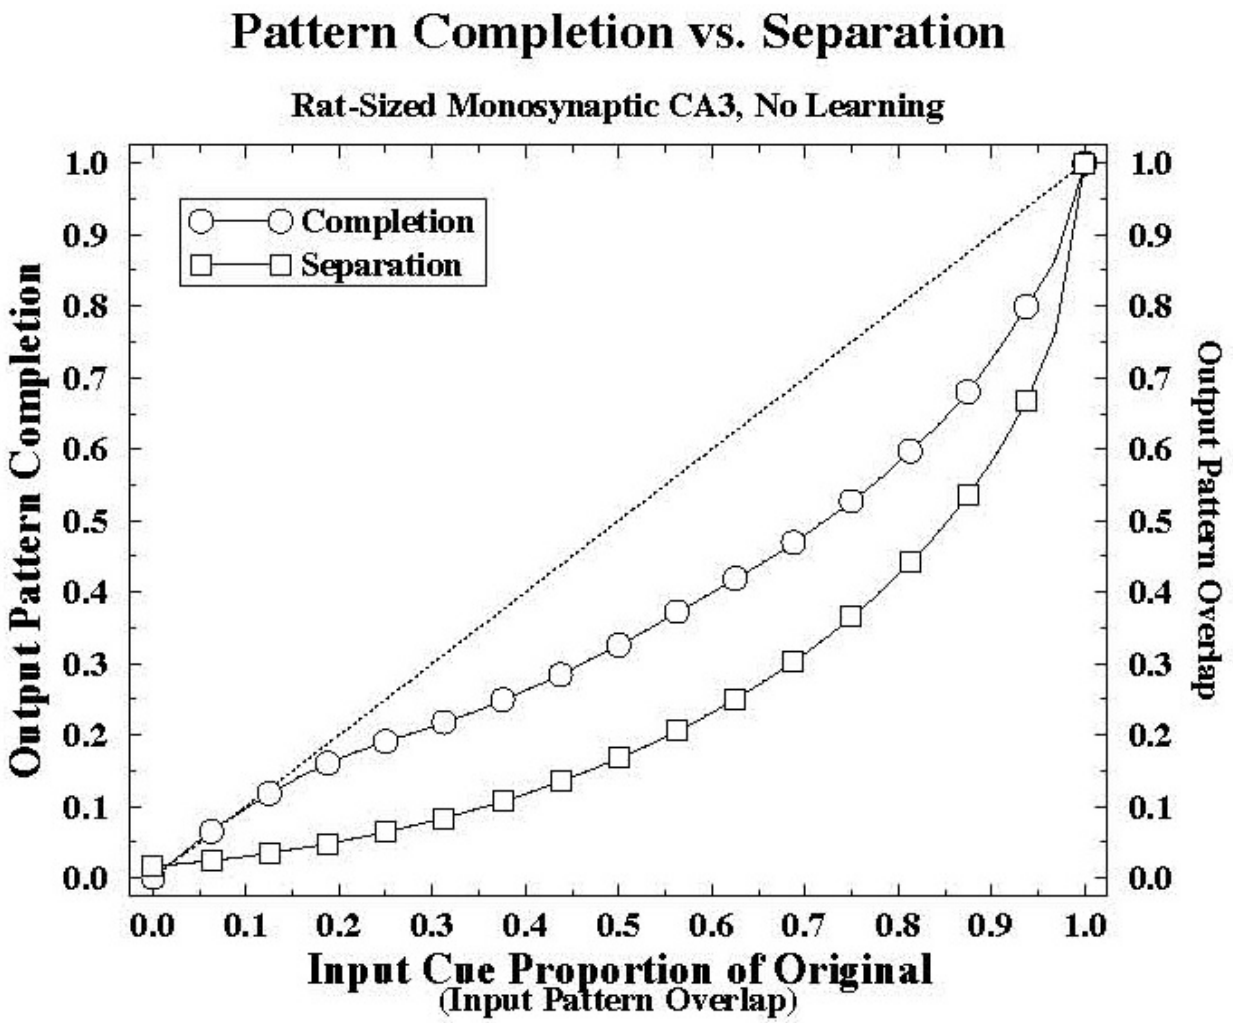

Pattern Completion

- Without learning, completion cannot happen.
- Two learning rules were tried:
  - WI: Weight Increase (like Marr)
  - WID: Weight Increase/Decrease
- WI learning multiplies weights in H_{ab} by $(1+L_{rate})$.
- WID learning increases weights as per WI, but also exponentially decreases weights to units in $F\text{-}H_a$ by multiplying by $(1-L_{rate})$.
- Result: WID learning improves both separation and completion.

WI Learning and Pattern Completion

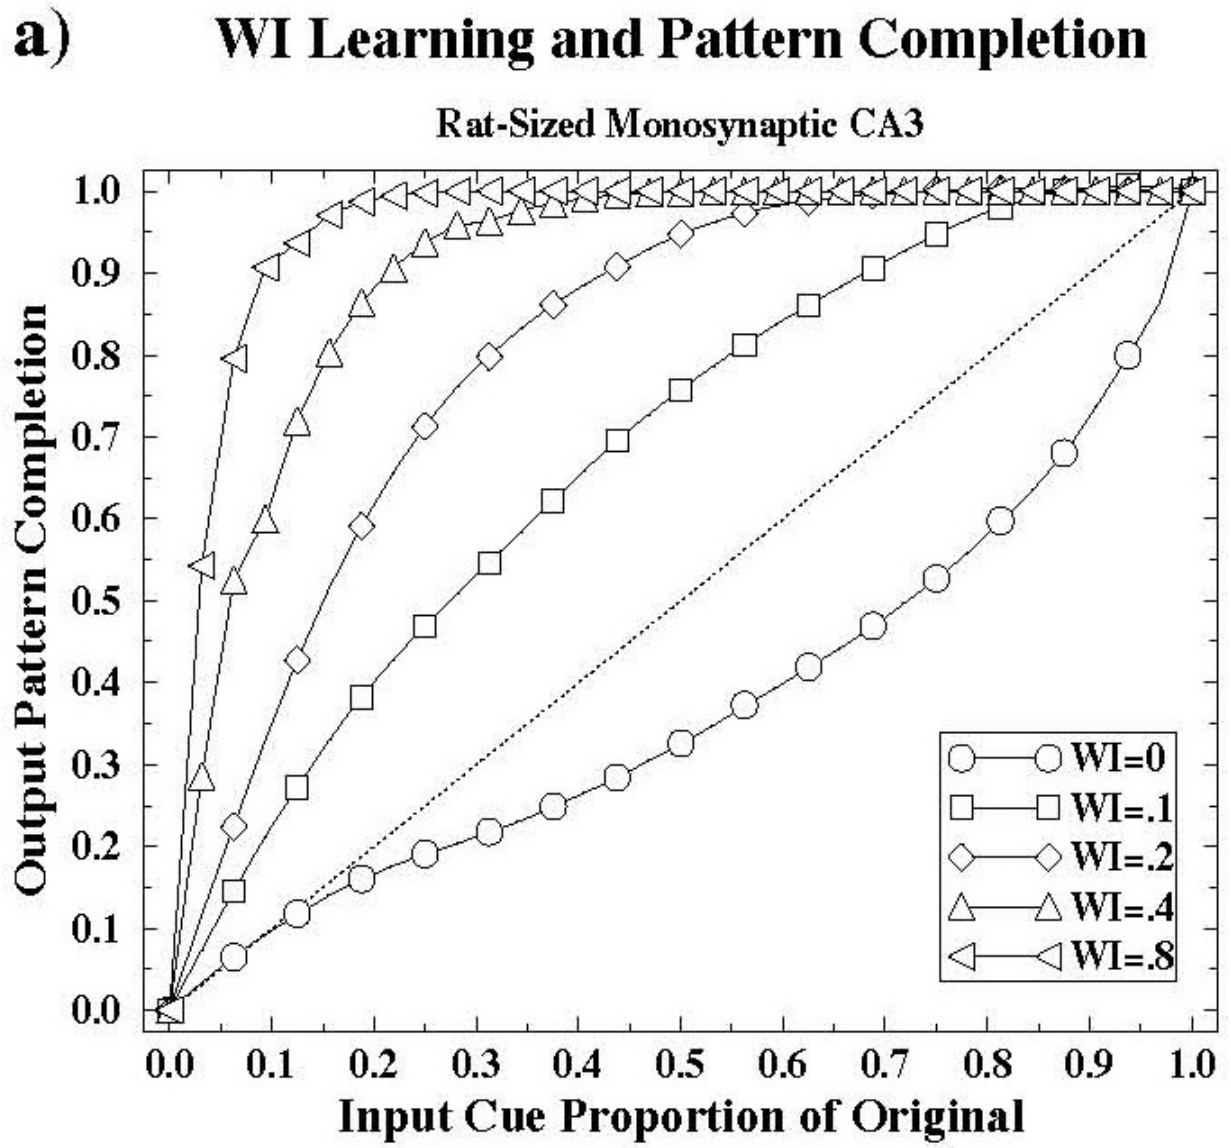

WI Learning Reduces Pattern Separation

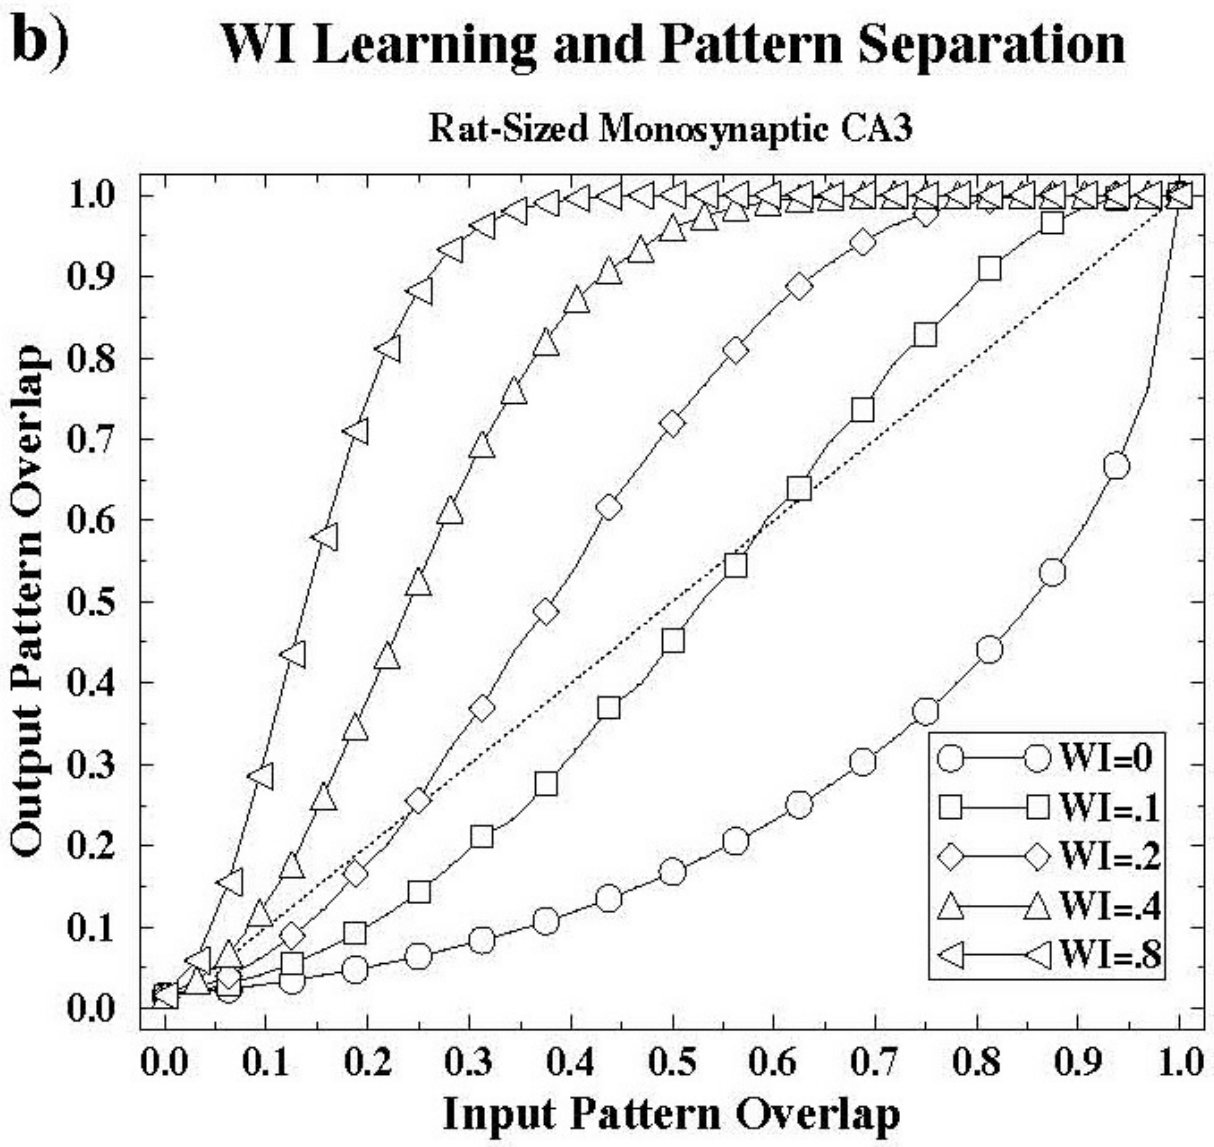

WI Learning Hurts Separation

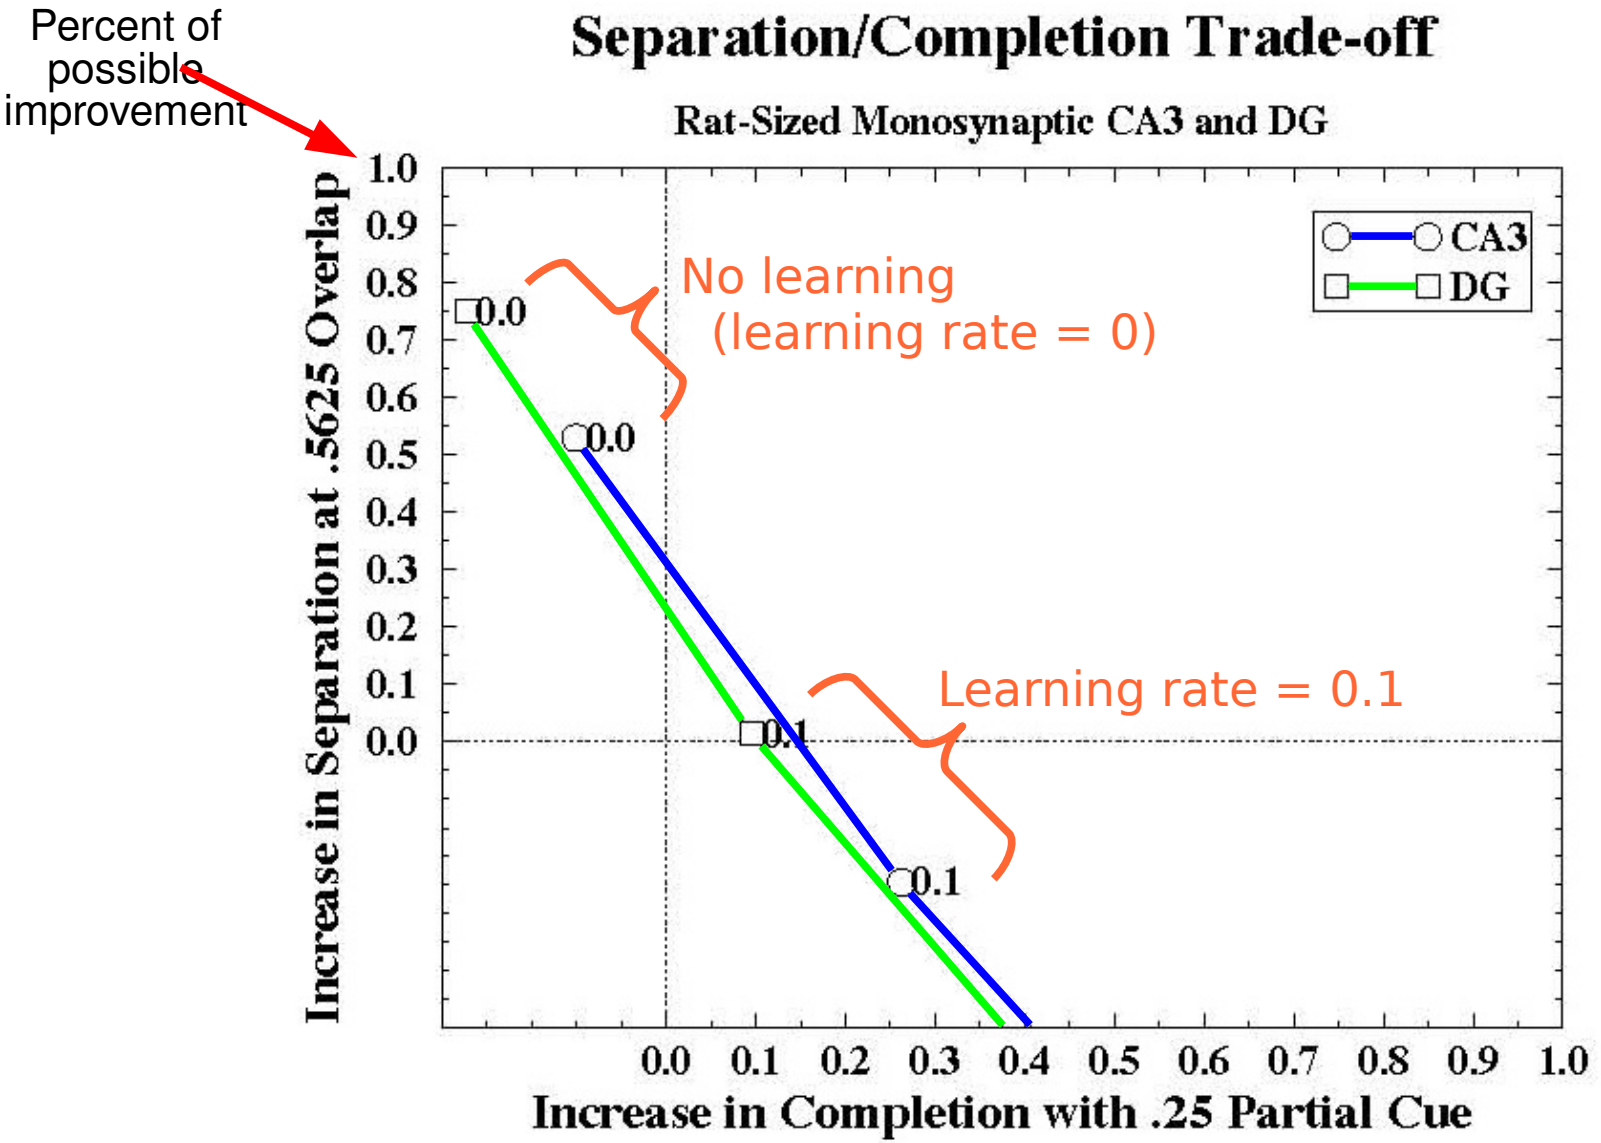

WID Learning Has A Good Tradeoff

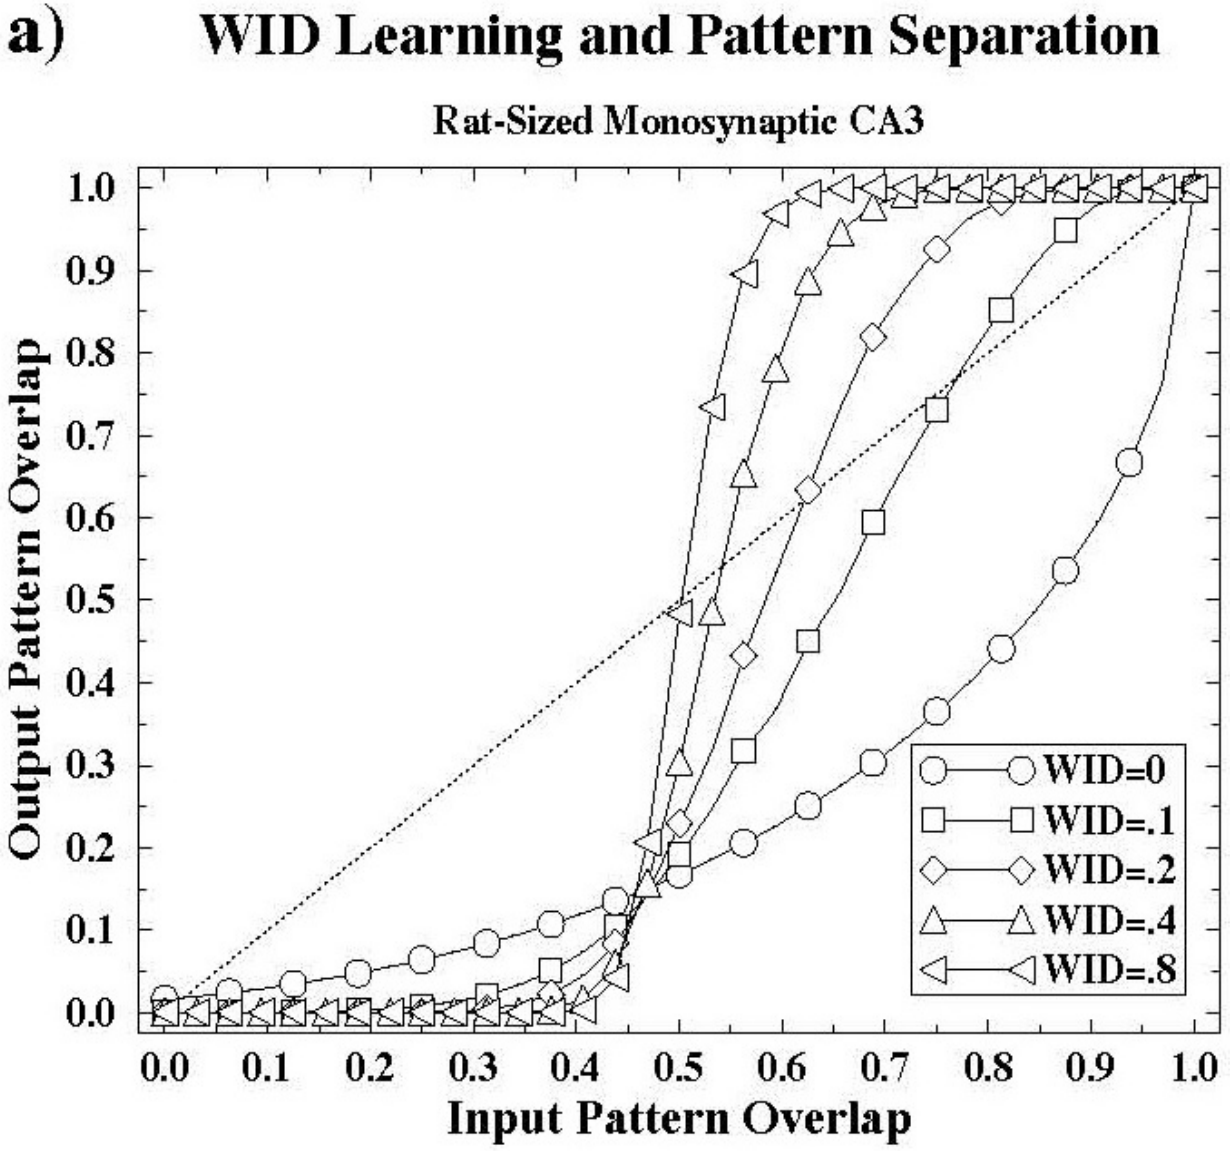

WI vs. WID Learning

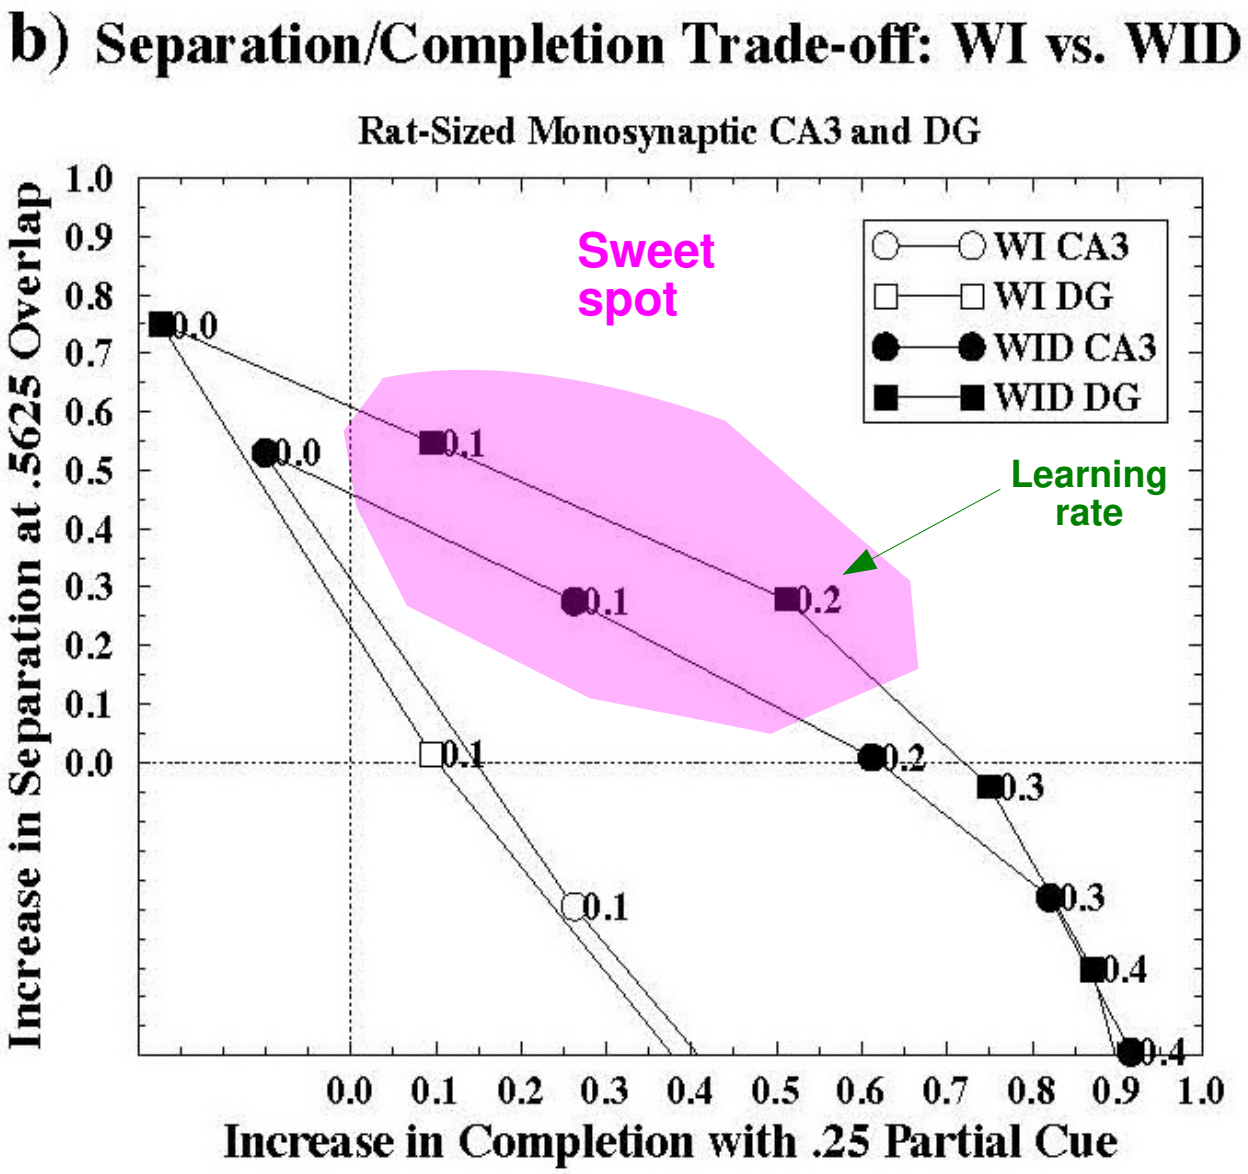

Hybrid Systems

- Multiple completion stages don't help (cf. Willshaw & Buckingham's comparison of Marr models.)
  - With noisy cues, completion produces a somewhat noisy result which would lead to further separation at the next stage.
- MSEPO — mossy fibers only for separation (learning).
  - Perhaps partial EC inputs aren't strong enough to drive DG.
- FM —fixed mossy system: no learning on these fibers.
  - Learning reduces pattern separation. Real mossy fibers undergo LTP, but it's not NMDA-dependent (so non-Hebbian).
- FMSEPO — combination of FM + SEPO.
  - Optimal tradeoff between separation and completion.

Performance of Hybrid Models

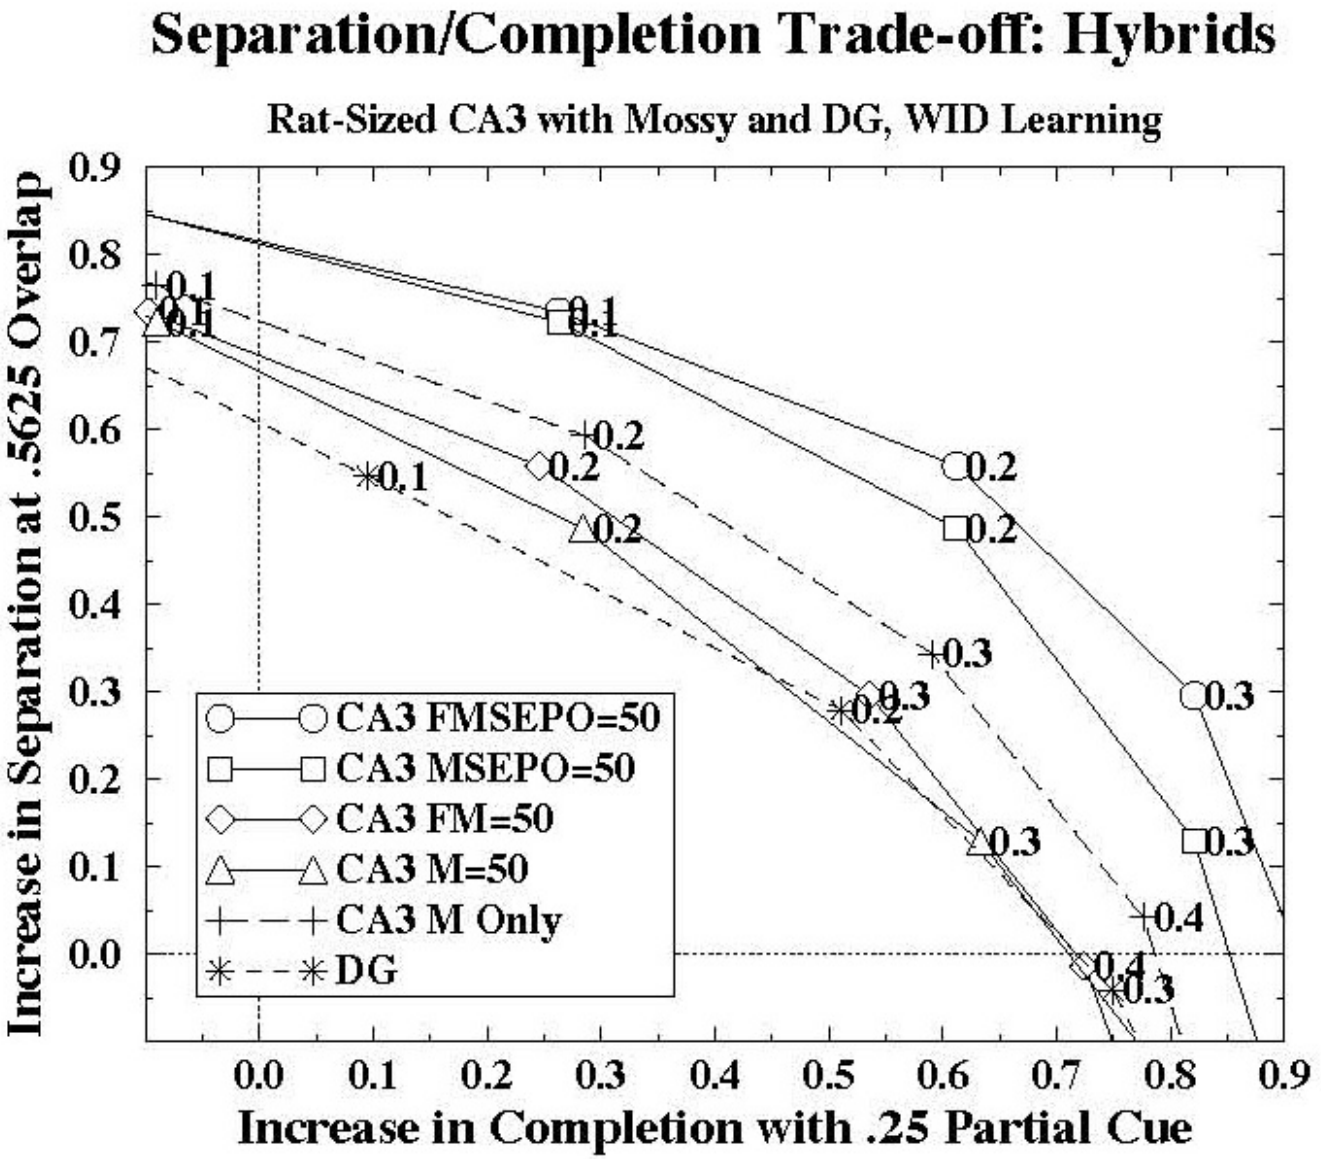

What Is the Mossy Fiber Pathway Doing?

- Adds a high variance signal to the CA3 input, which...
- Selects a random subset of CA3 cells that are already highly activated by EC input.
- This enhances separation when recruiting the representations of stored patterns.
- But it hurts retrieval with partial or noisy cues.
  - So don't use it. Use MSEPO or FMSEPO.

Conclusions

- The main contribution of this work is to show how separation and completion can be accomplished in the same architecture.
- The model uses realistic figures for numbers of units and connections.
- Fan-in size doesn't seem to matter.
- WID learning is necessary for a satisfactory tradeoff between separation and completion.
- DG contributes to separation but perhaps not to completion.

Limitations of the Model

- Simplified anatomy: the model only included EC→CA3 and EC→DG→CA3 connections.
- No CA3 recurrent connections.
- No CA1.
- Only a single pattern stored at a time:
  - Store A, measure overlap with B.
  - No attempt to measure memory capacity.
- A more realistic model would be too hard to analyze.

Possible Different Functions of CA3 and CA1

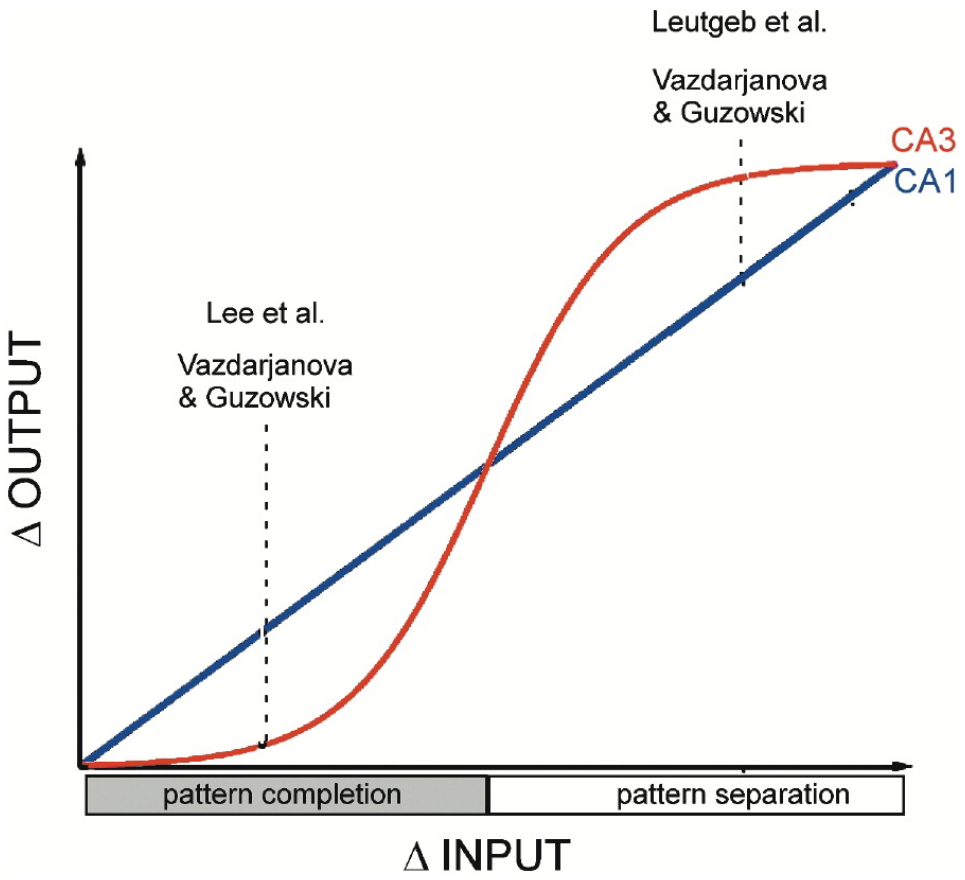

Measured by IEG (Immediate Early Genes): Arc/H1a catFISH method

Guzowski, Knierim, and Moser (2004)

Expose rats to two environments 30 minutes apart. Environments can be (i) identical , (ii) similar but with changes to local or distal cues, or (iii) completely different.

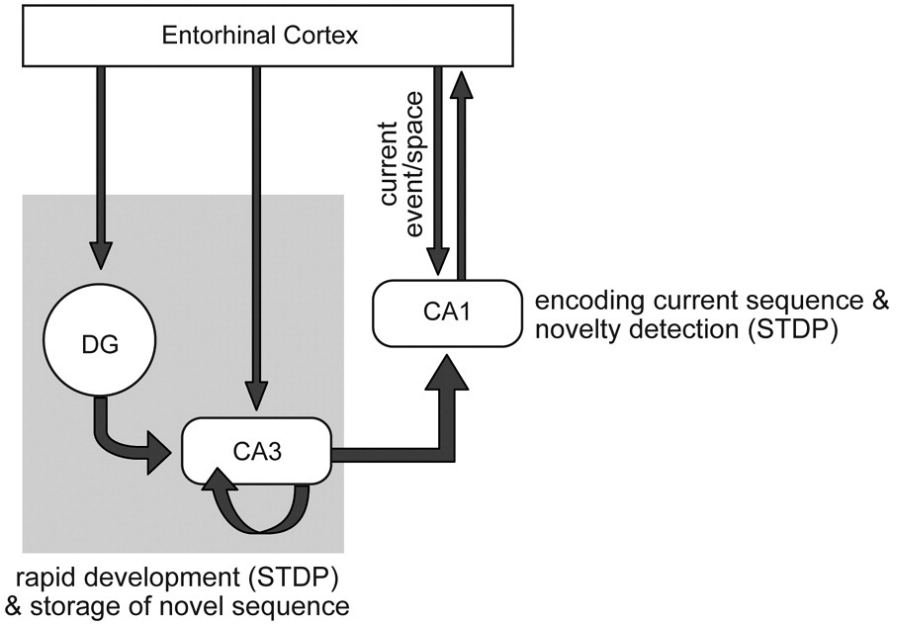

Hasselmo's Model: Novelty Detection

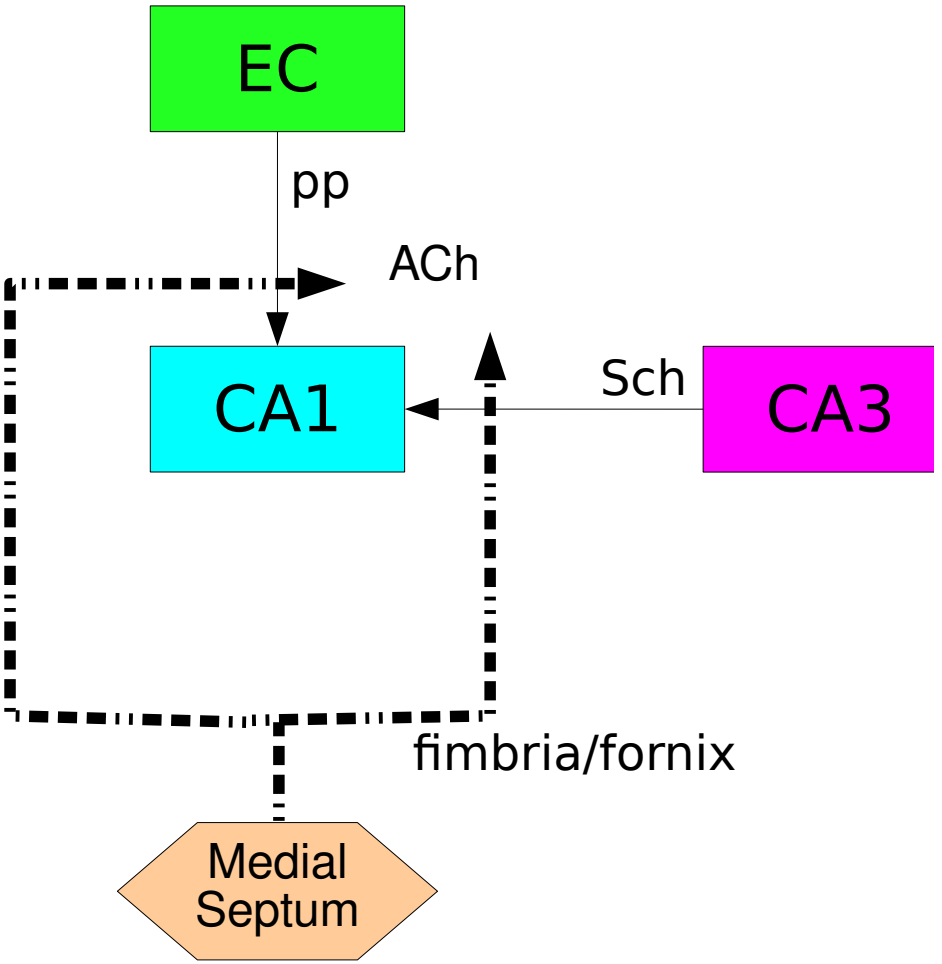

Acetycholine reduces synaptic efficacy (prevents CA3 from altering CA1 pattern) and enhances synaptic plasticity.

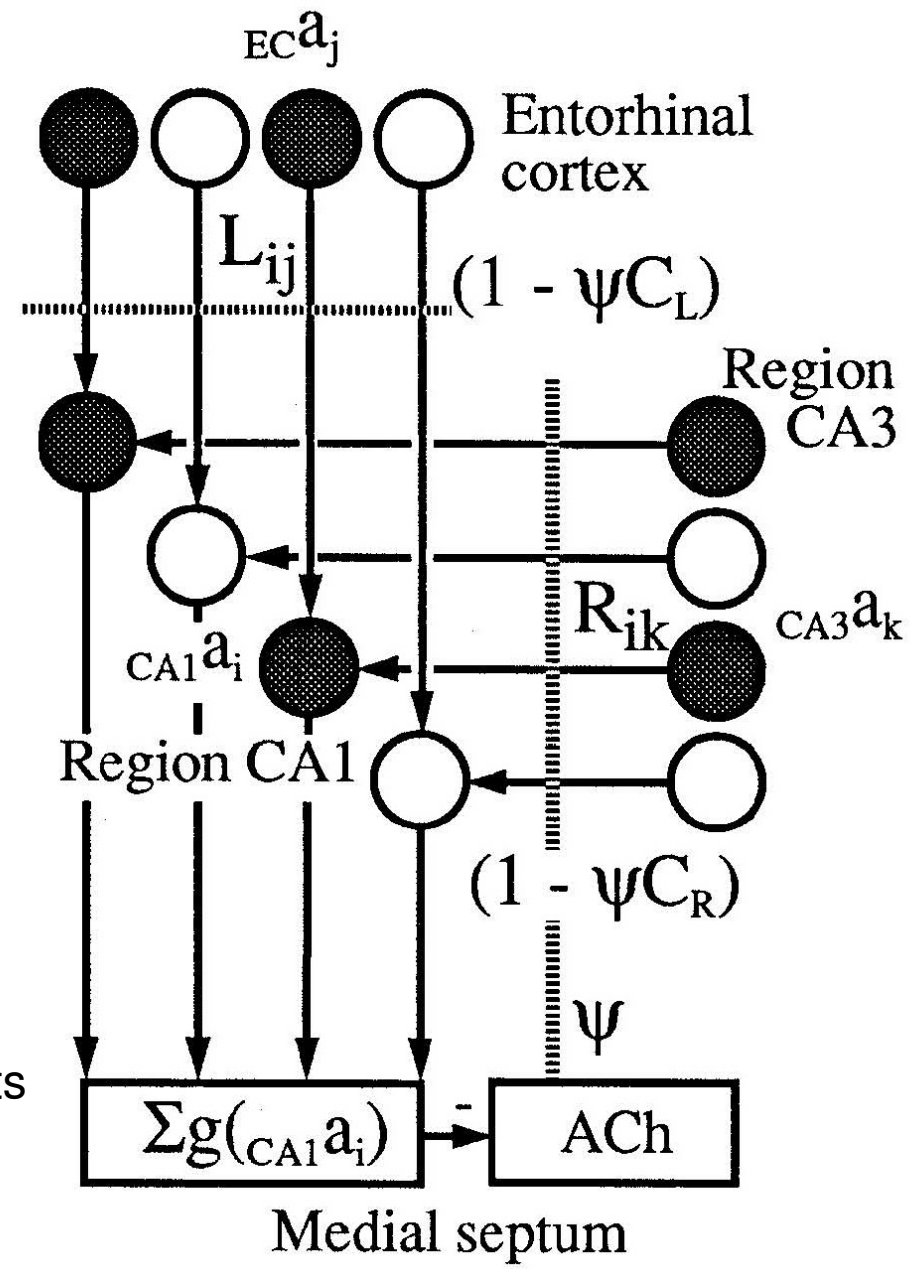

Pattern Separation in Human Hippocampus

- Bakker et al., *Science,* March 2008: fMRI study
- Subjects were shown 144 pairs of images that differed slightly, plus additional foils. Asked for an unrelated judgment about each image (indoor vs. outdoor object).

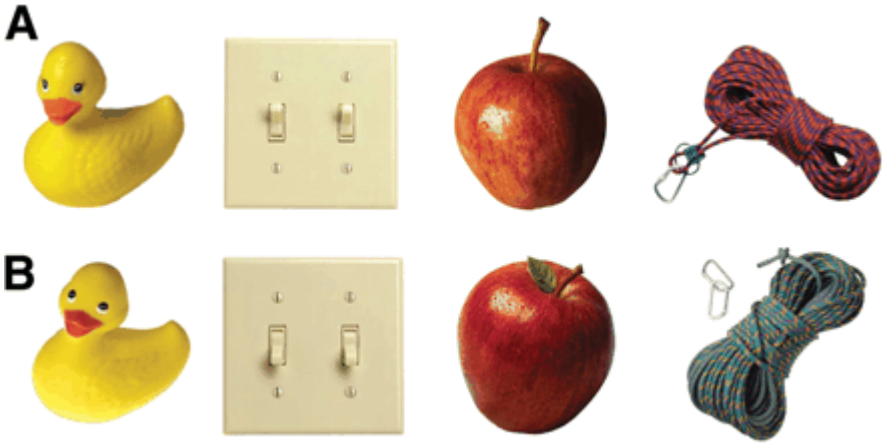

- Three types of trials: (i) new object, (ii) repetition of a previously seen object, (iii) slightly different version of a previously seen object: a *lure*.

Eight ROIs Found

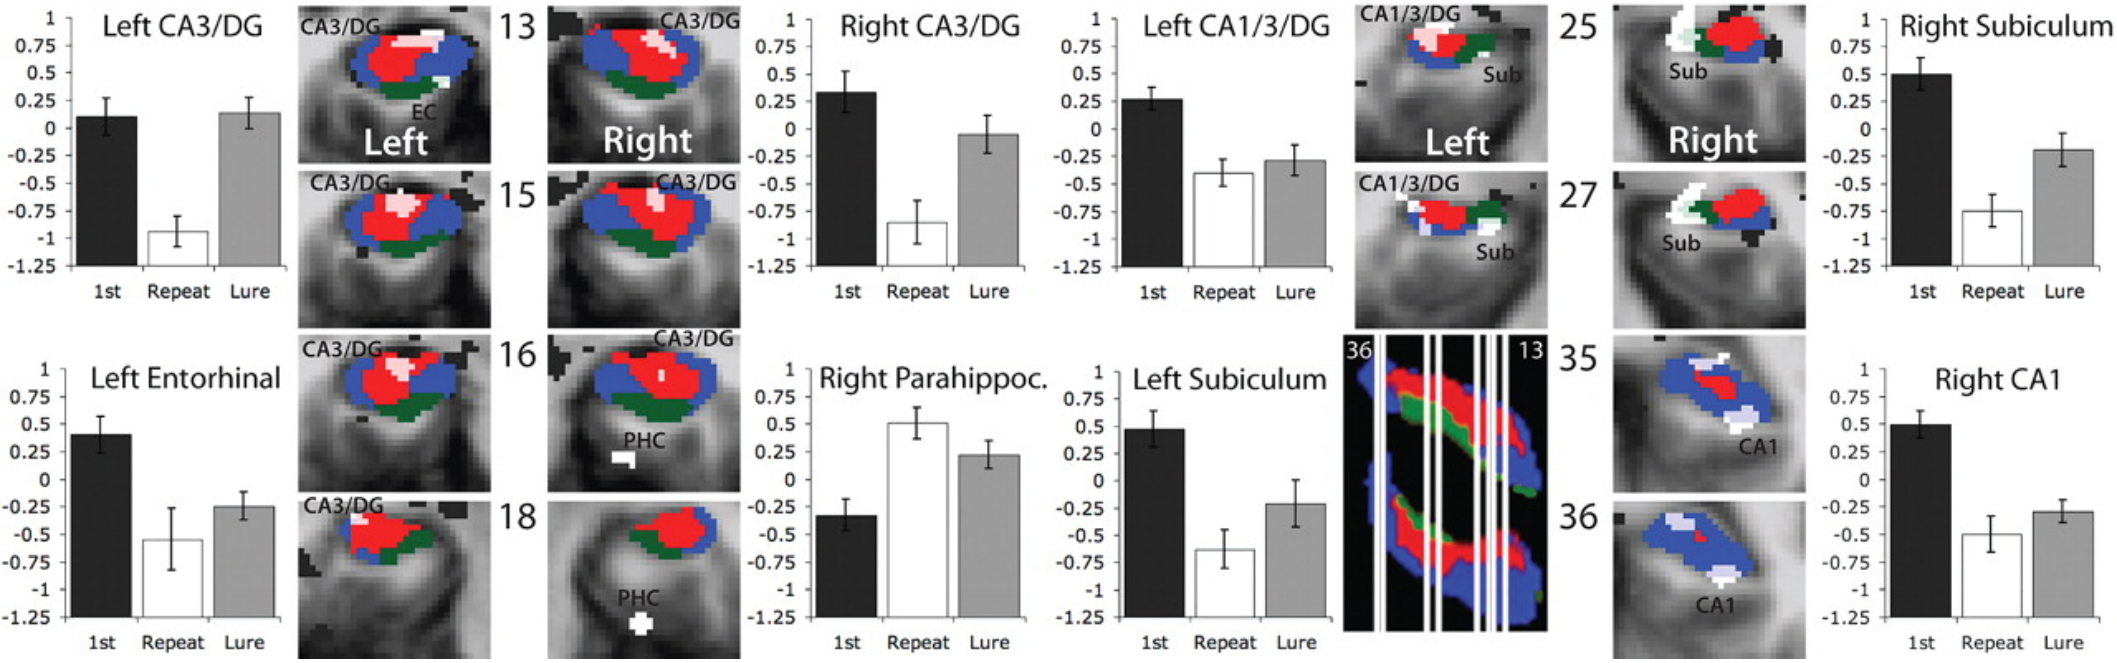

- Couldn't resolve DG vs. CA3 so treated as one region.
- Regions outlined above: CA3/DG CA1 Subiculum
- Areas of significant activity within MTL shown in white.
- New objects, repetitions, and lures were reliably discriminable. Generally, repetitions → lower activity.

Bias Scores for ROIs

- *bias* = (*first – lure) / (first – repetition)*
- Scores close to 1 → completion; 0 → separation.
- CA3/DG shows more pattern separation than other areas.