Marr's Theory of the Hippocampus
Part II: Effect of Recurrent Collaterals

Computational Models of Neural Systems

Lecture 3.4

David S. Touretzky

October, 2017

Two Layer Model

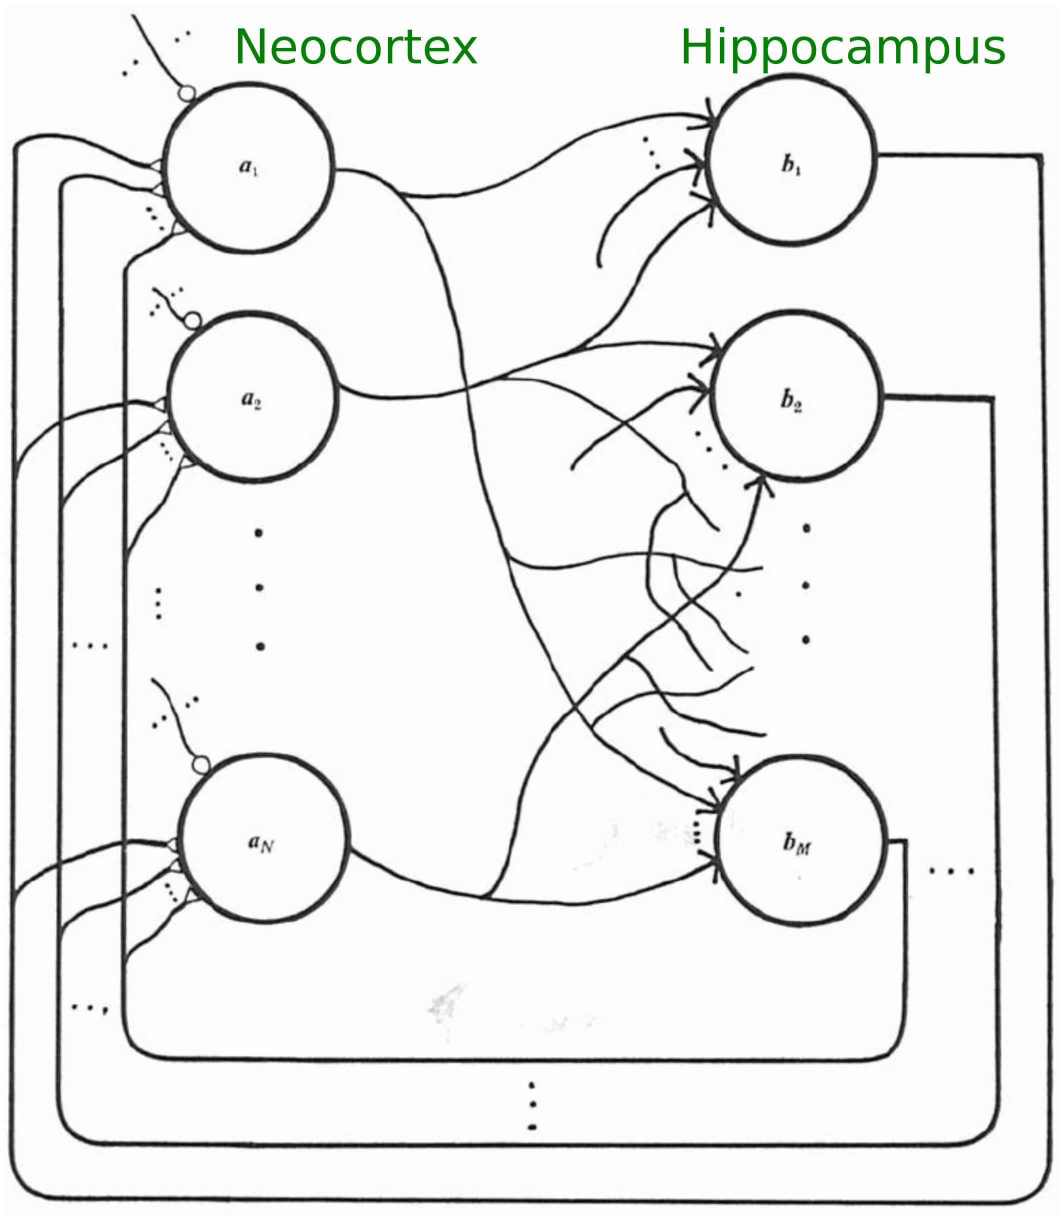

Two Layer Model Insufficient?

- Marr claimed the two layer model could not satisfy all the constraints he had established concerning:
  - number of stored memories *n*
  - number of cells
  - sparse activity: $n \, \alpha_i \, \alpha_{i-1} \leq 1$
  - but patterns not too sparse for effective retrieval
  - number of synapses per cell: $S_i \, \alpha_i \, N_i \geq 20 \, N_{i-1}$
- But this was really because he assumed the number of output cells was just 10^4.
- He switched to a three layer model, with neocortical cells, evidence cells (codons), and output cells.
- The output cells had recurrent collaterals.

The Three-Layer Model

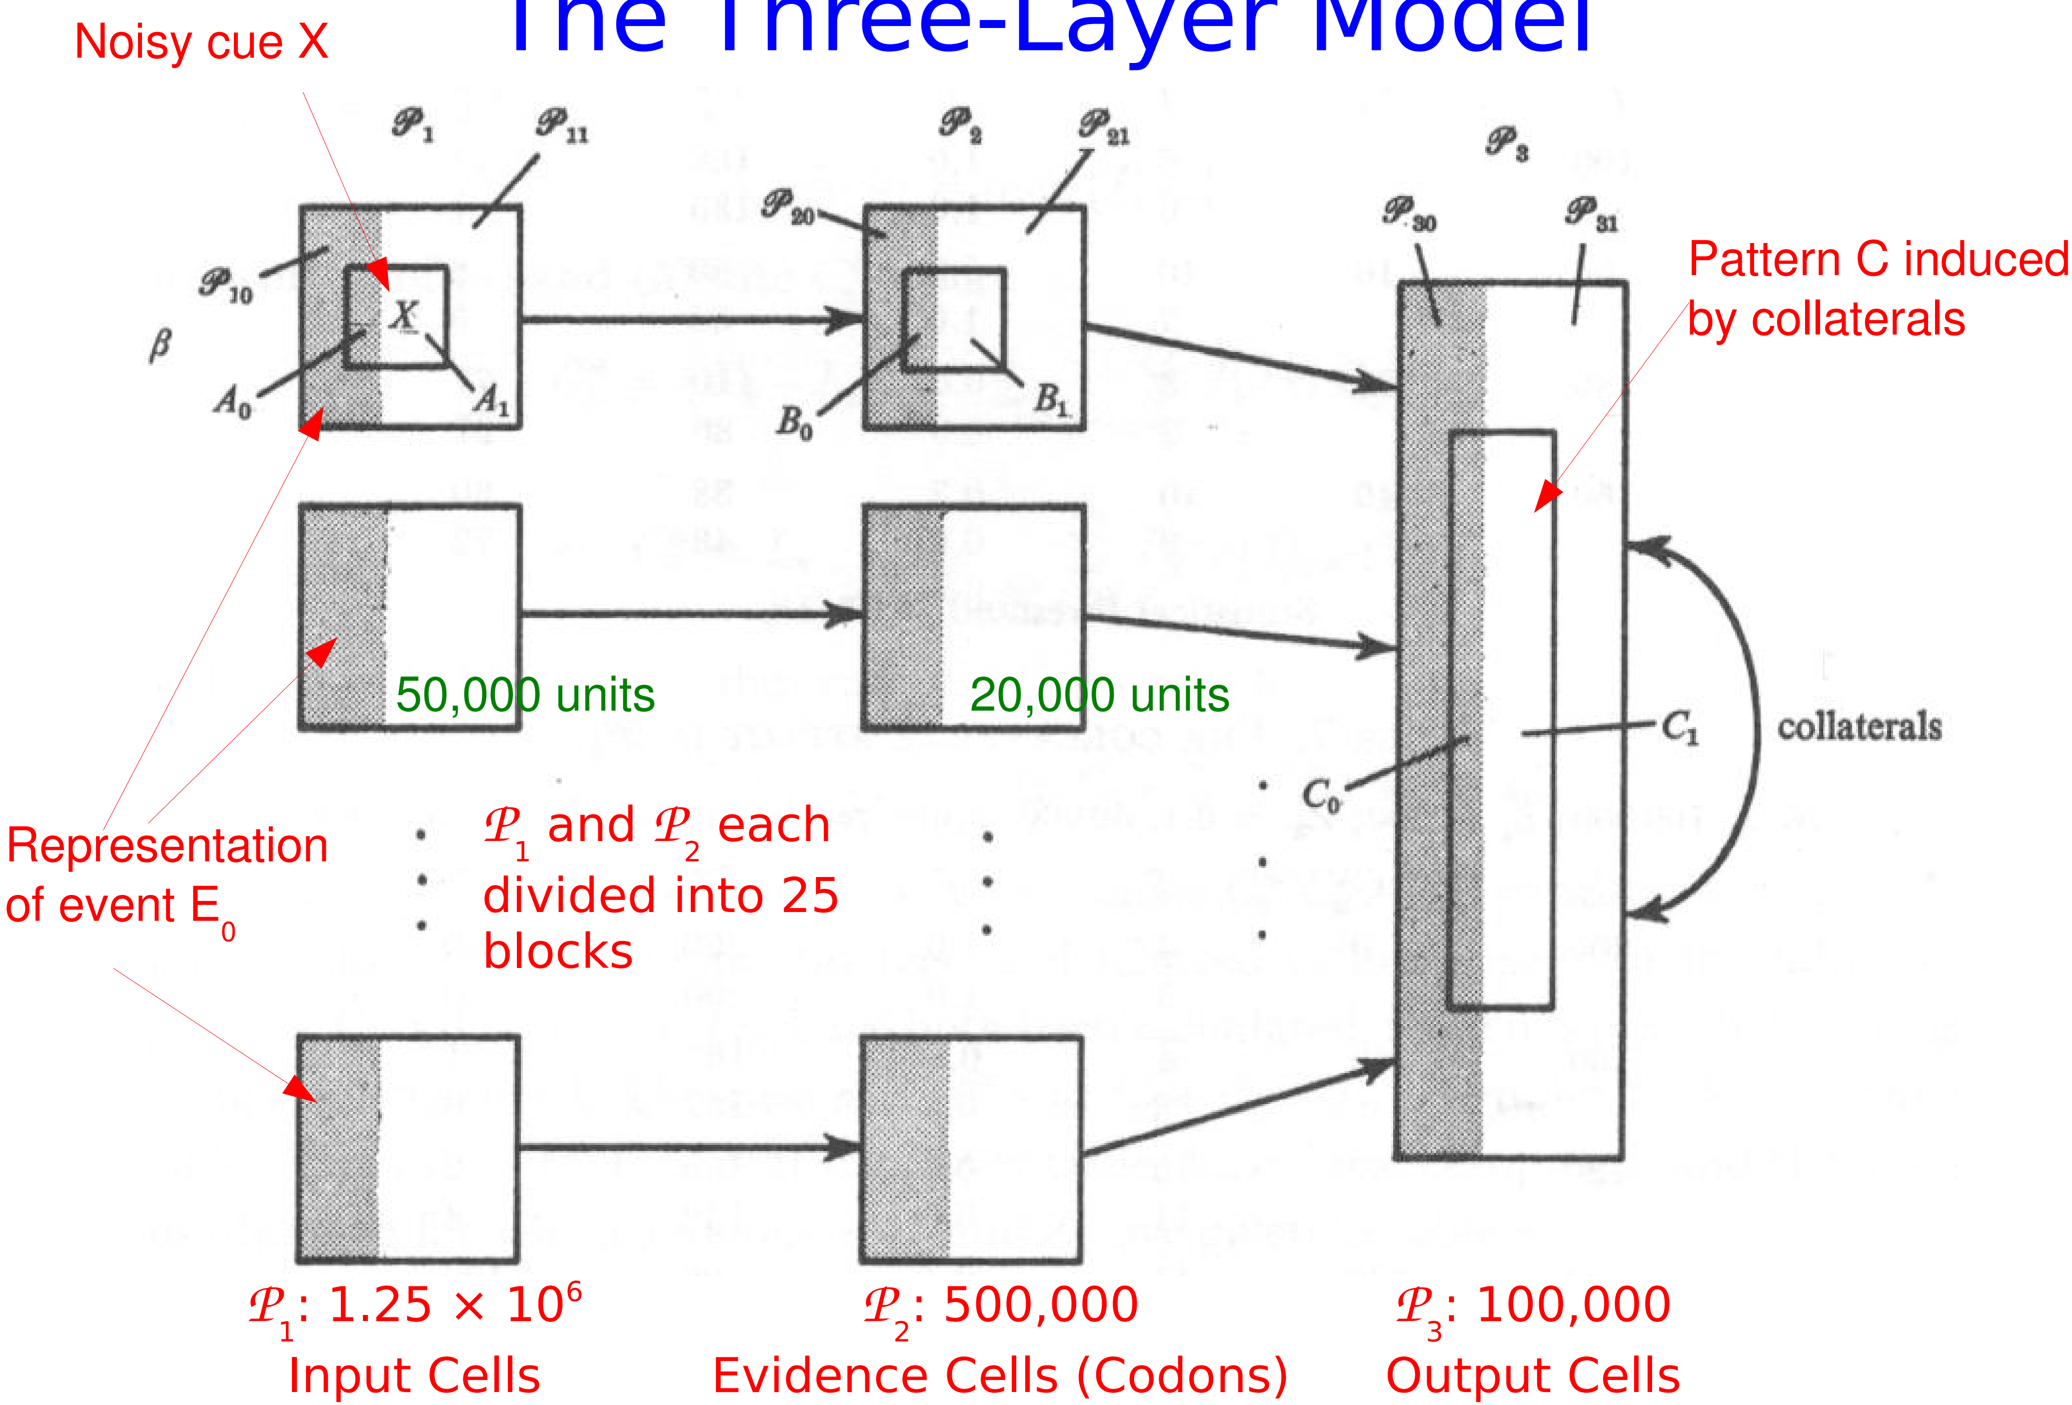

The Collateral Effect

- Let $\mathcal{P}_i$ be a population of cells forming a simple representation.
- Each cell can learn about 100 input events.
- Population as a whole learns $n = 10^5$ events.
- Hence α_i must be around 10^{-3}.
- We require $n\, \alpha_i\, \alpha_{i-1}$ to be at most 1.
  Estimated value based on the above is 0.1.
- Hence we can let $\mathcal{P}_{i-1} = \mathcal{P}_i$ and use recurrent collaterals to help clean up the simple representation.
- Result: external input to $\mathcal{P}_i$ need not be sufficient by itself to reproduce the entire simple representation.

Parameters of the Three-Layer Model

- $\mathcal{P}_1$ has 1.25 × 10^6 cells divided into 25 blocks of 50,000.
- $\mathcal{P}_2$ has 500,000 cells divided into 25 blocks of 20,000.
- $\mathcal{P}_3$ has a single block of 100,000 cells.
- Let number of synapses/cell S_3 = 50,000.
- Let x_i be number of active synapses on a cell, i.e., the number used to store one event.
- $n\alpha_i$ is the expected number of events a cell encodes.
- Probability of a synapse being potentiated is:

$$\Pi_i = 1 - (1 - x_i / S_i)^{n\alpha_i}$$

Parameters of the Three-Layer Model

$$\Pi_i = 1 - (1 - x_i / S_i)^{n\alpha_i}$$

$$x_i = \sum_{r \geq R_i} P_i(r) \cdot r$$

- $P_l(r)$ is the probability that a cell in layer *i* has exactly r active afferent synapses.
- From the above, we have $L_3 = \alpha_3 N_3 = 217$, and $\alpha_3 = 0.002$.
- If we want useful collateral synapses in $\mathcal{P}_3$, must have $n(\alpha_3)^2 \leq 1$.
- So with $n = 10^5$ events, we have α_3 = at most 0.003.

Retrieval With Partial/Noisy Cues

- Let P_{30} be the simple representation of E_0 in $\mathcal{P}_3$.
- Let P_{31} be the remaining cells in $\mathcal{P}_3$.
- Let C_0 be the active cells in P_{30} representing subevent X.
- Let C_1 be the active cells in P_{31} (noise).
- Note that C_0+C_1 = pattern size L_3.

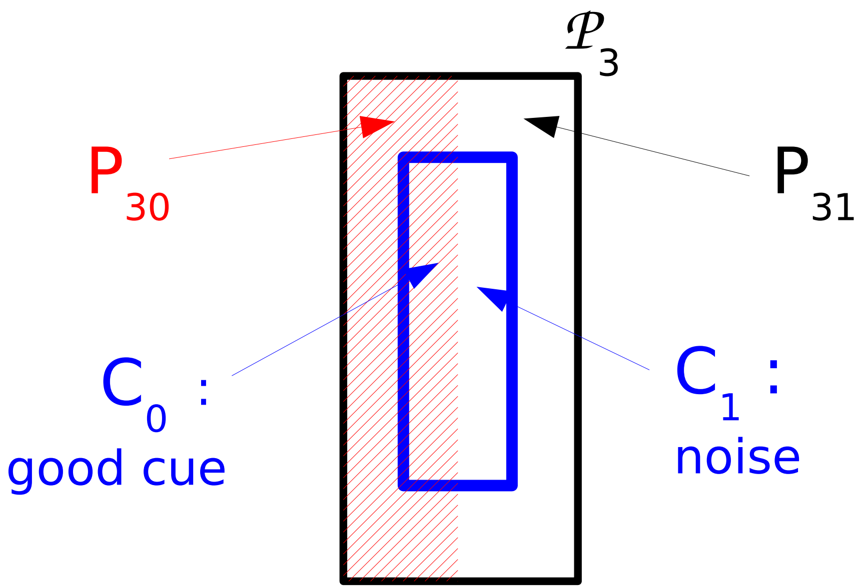

Collateral Connections

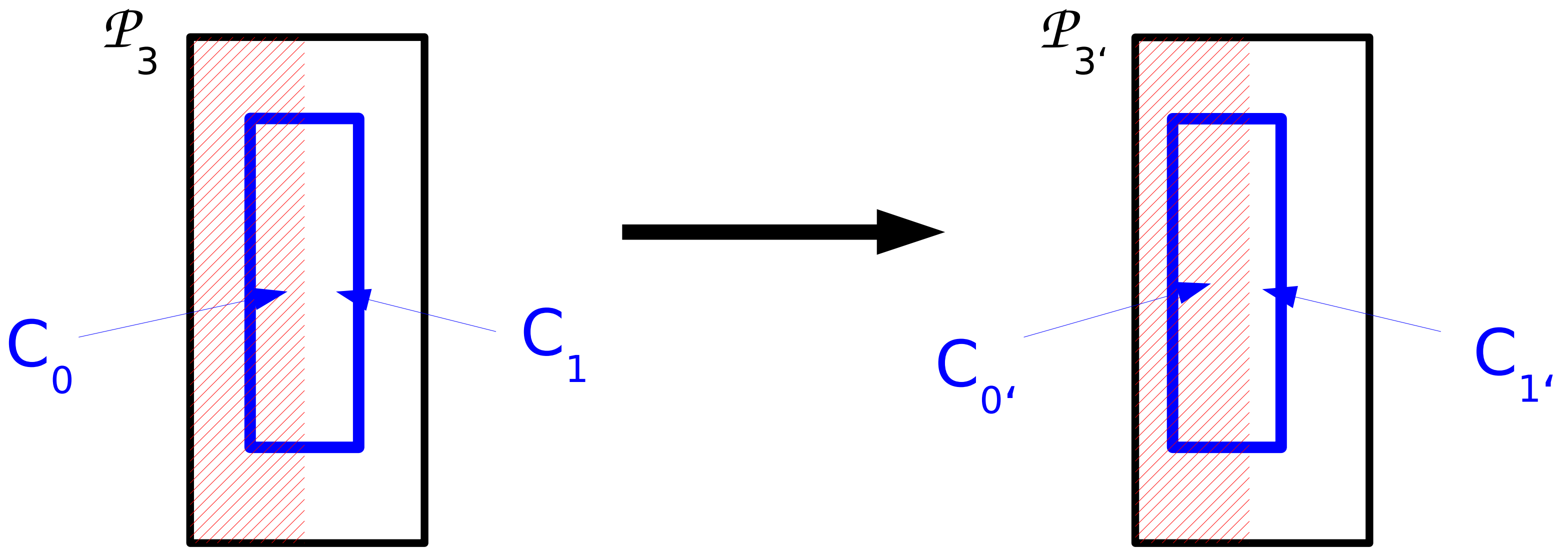

- The <u>statistical threshold</u> is the ratio $C_0{:}C_1$ such that the effect of collaterals is zero: $C_0{:}C_1 = C_{0'}{:}C_{1'}$
- Collaterals help when statistical threshold is exceeded.
- Calculating $C_{0'}{:}C_{1'}$ is a bit tricky because there is both a subtractive and a divisive threshold; see Marr §3.1.2.

Collateral Effect in $P_{3'}$

- Let **b** be an arbitrary cell in $P_{3'}$.
- $Z_{3'}$ is probability of a recurrent synapse onto **b**.
- Number of <u>active</u> recurrent synapses onto **b** is distributed as Binomial(L_3; $Z_{3'}$) with expected value $L_3 Z_{3'}$.
- Probability that **b** has exactly *x* active synapses onto it:

$$P_{3'}(x) = \binom{L_3}{x} \cdot Z_3^x \cdot (1 - Z_3)^{L_3 - x}$$

- **b** is either in P_{30} or not. We'll consider each case:

- Suppose **b** is in P_{31}, so not in P_{30}.
- Of the *x* active synapses onto **b**, the number of facilitated synapses *r* is distributed as Binomial(x; $\Pi_{3'}$).
- Probability that exactly *r* of the *x* active synapses onto **b** have been modified when **b** is in P_{31} is:

$$Q_{3'1}(r) = \binom{x}{r} \cdot \Pi_{3'}^{r} \cdot \left(1 - \Pi_{3'}\right)^{x-r}$$

- Suppose **b** is in P_{30}.
- All afferent synapses from other cells in P_{30} onto **b** will have been modified.
- Active synapses onto **b** are drawn from two distributions:
  - Binomial(C_0; $Z_{3'}$) for cells in P_{30} – modified with probability 1
  - Binomial(C_1; $Z_{3'}$) for cells in P_{31} – modified with probability $\Pi_{3'}$
- Approximate this mixture with a single distribution for the number of modified active synapses:
  - Binomial(x; $(C_0+C_1\Pi_{3'})/(C_0+C_1)$)

- Let C be the expected fraction of synapses onto **b** in the subevent X that have been modified:

$$C = \frac{C_0 + C_1 \Pi_{3'}}{C_0 + C_1}$$

- Probability that *r* of *x* active synapses have been modified when **b** is in P_{30} is:

$$Q_{3'0}(r) = \binom{x}{r} \cdot C^r \cdot (1-C)^{x-r}$$

- Note: this differs from Marr's formula 3.3.

- If all cells in $P_{3'}$ have threshold R, then:

Size of the simple representation P_{30}

$$C_{0'} = L_3 \cdot \sum_{r \geq R} \sum_{x=r}^{L_3} P_{3'}(x) Q_{3'0}(r)$$

Prob. that a cell in P_{30} has enough active modified synapses to be above threshold

Number of potential P_{31} noise cells

$$C_{1'} = (N_3 - L_3) \cdot \sum_{r \geq R} \sum_{x=r}^{L_3} P_{3'}(x) Q_{3'1}(r)$$

Prob. that a cell in P_{31} has enough active modified synapses to be above threshold

- Statistical threshold is the ratio where

$$C_0 : C_1 = C_{0'} : C_{1'}$$

subject to

$$C_0 + C_1 = C_{0'} + C_{1'} \approx L_3$$

Dealing With Variable Thresholds

- In reality, cells in $\mathcal{P}_3$ do not have fixed thresholds R. They have:
  - A subtractive threshold T
  - A divisive threshold *f*

- Combined threshold:

$$R(b) = \max(T, fx)$$

- Can calculate $C0^*$ and $C1^*$ using R(b) instead of R.

- Details are in Marr §3.1.2.

Results

- More synapses help: $Z_{3'} = 0.2$ gives a statistical threshold twice as good as $Z_{3'} = 0.1$.
- Good performance depends on adjusting T and *f*. (*f* should start out low and increase; T should decrease to compensate.)
- Collaterals can have a *big* effect.
- Recovery of E_0 is almost certain for inputs that are more than 0.1 L_3 above the statistical threshold.
- Example: Marr table 7: $L_3 = 200$, threshold is 60:140.
- In general: collaterals help whenever $n\alpha^2 \leq 1$.

  (Sparse patterns; not too many stored memories.)

Marr's Performance Estimate

- Input patterns: L_1 = 2500 units active out of 1.25 million (25 blocks of 50,000; 100 active units in each block)
- Output patterns: L_3 = 217 units out of 100,000.
- With $n = 10^5$ stored events, accurate retrieval from:
  - 30 active fibers in one block, all of which are in E_0
  - 100 active fibers in one block, of which 70 are in E_0 and 30 are noise
- With $n = 10^6$ stored events, accurate retrieval from:
  - 60 active fibers in one block, all of which are in E_0
  - 100 active fibers in one block, of which 90 are in E_0

Willshaw and Buckingham's Model

- Willshaw and Buckingham implemented a simplified 1/100 scale model of Marr's architecture
- Didn't bother partitioning $\mathcal{P}_1$ and $\mathcal{P}_2$ into blocks.
- $\mathcal{P}_1$ = 8000 cells, $\mathcal{P}_2$ = 4000 cells, and $\mathcal{P}_3$ = 1024 cells.
- For two-layer version, omit $\mathcal{P}_2$.
- Performance was similar for both architectures.
- Memory capacity was roughly 1000 events.
  - Partial cue of 8% gave perfect retrieval 66% of the time.
  - In two-layer net, 16% cue gave perfect retrieval 99% of the time.
  - In three-layer version, 25% cue gave 100% perfect retrieval.

Three-Layer Model Parameters

$\alpha_1 = 0.03$	$\alpha_2 = 0.03$	$\alpha_3 = 0.03$
$N_1 = 8000$	$N_2 = 4000$	$N_3 = 1024$
	$S_2 = 1333$	$S_3 = 2666$

calc.:

$L_1 = 240$	$L_2 = 120$	$L_3 = 30$
	$Z_2 = 0.17$	$Z_3 = 0.67$
	$\Pi_2 = 0.41$	$\Pi_3 = 0.41$

Two vs. Three Layers

- Dashed line is two layer; solid is three layer.
- Open circles: partial cue. Solid circles: noisy cue.
- Two and three layer models perform similarly.

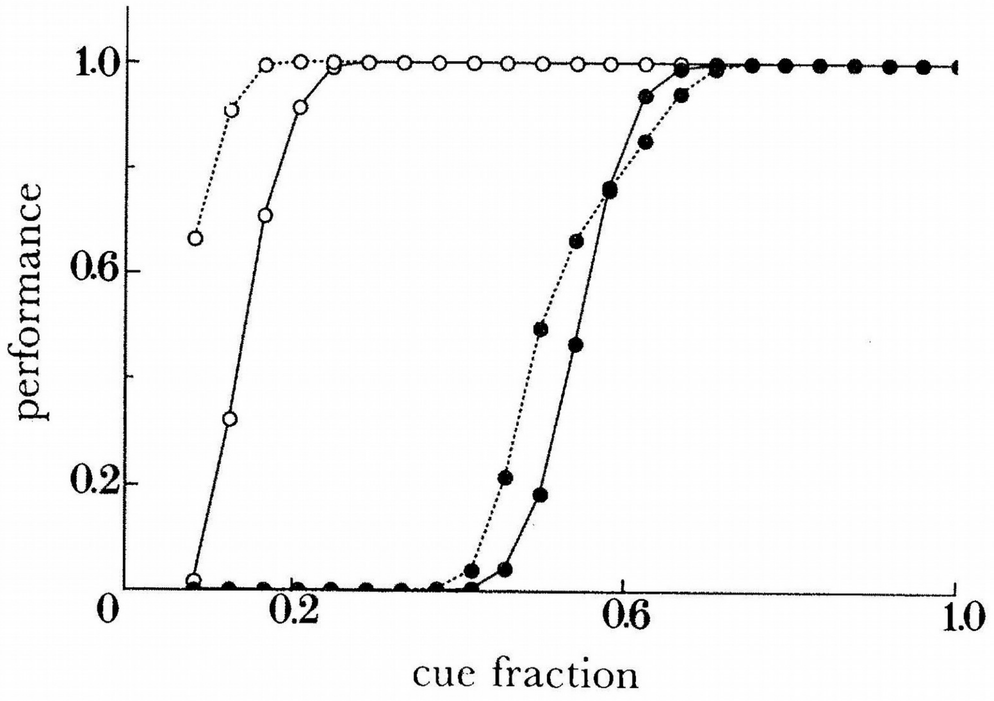

Effects of Memory Load

Two Layer

Three Layer

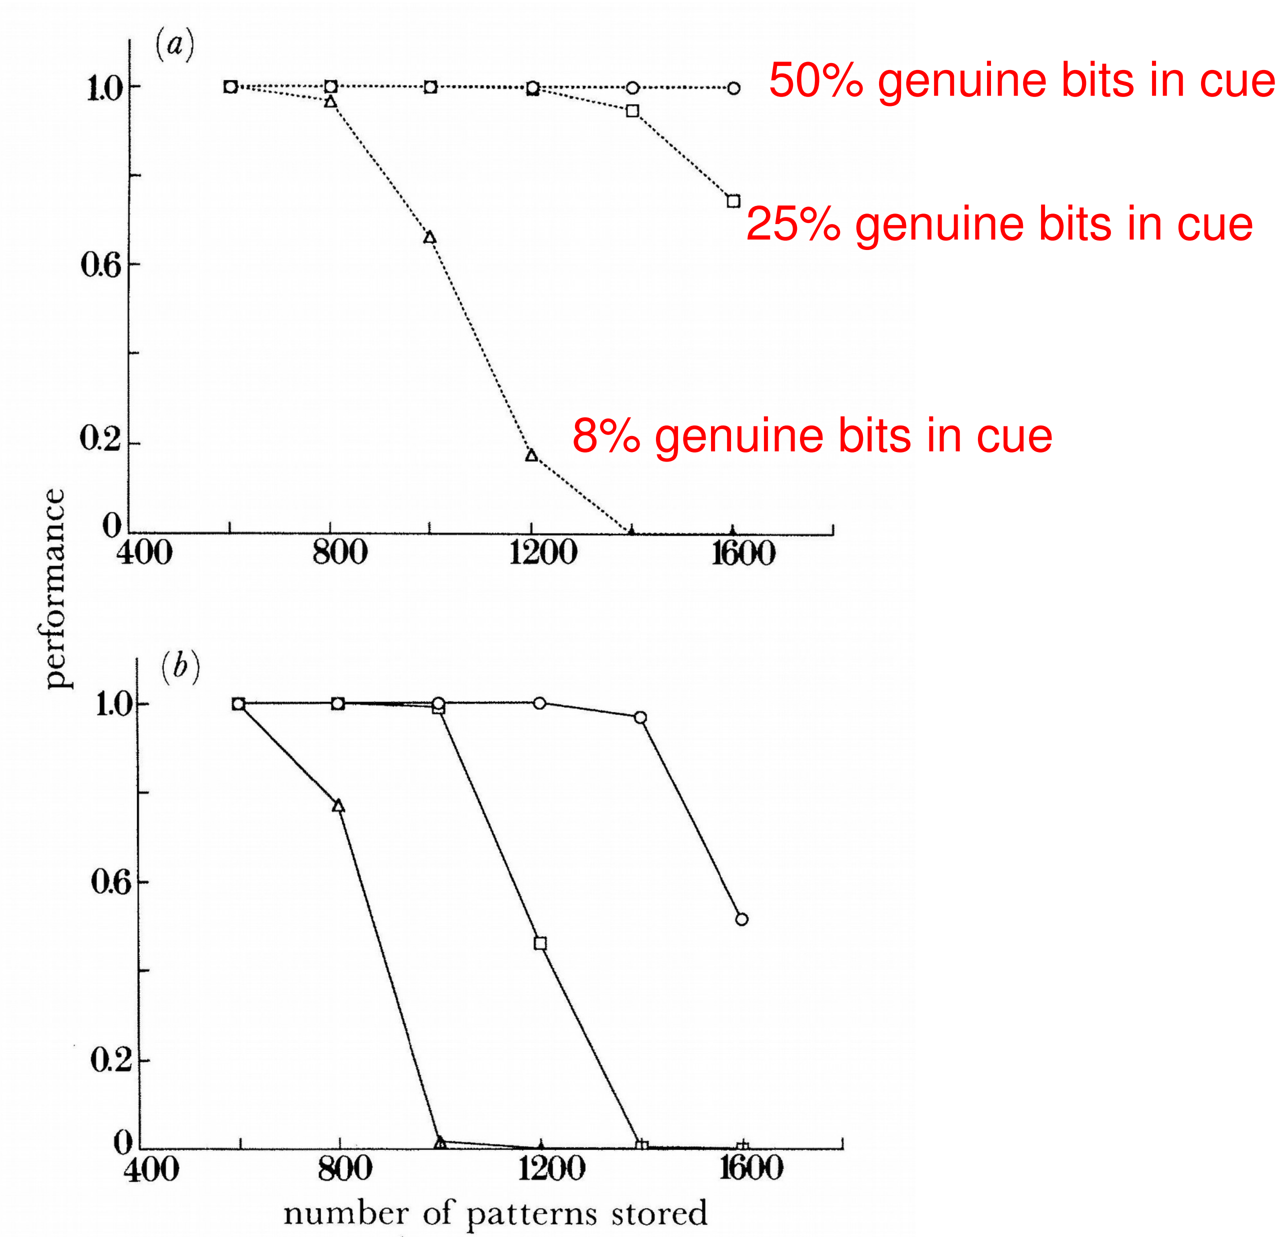

Division Threshold

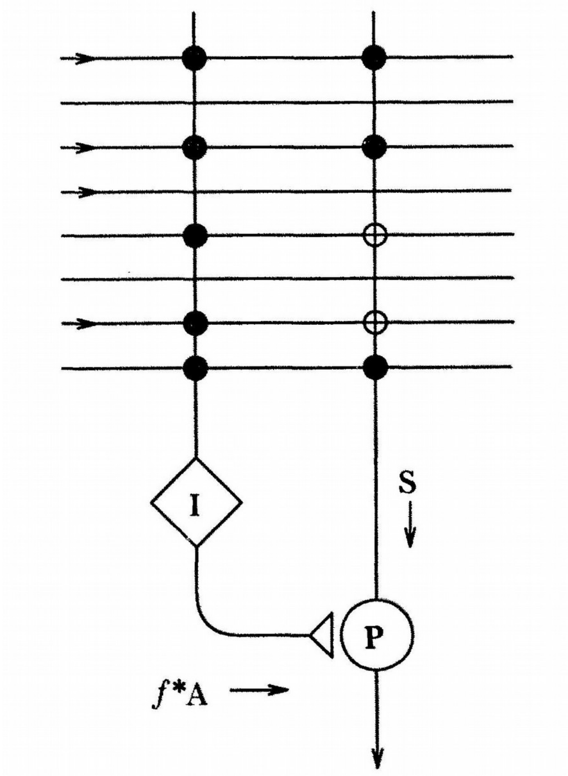

- I cell supplies divisive inhibition based on the number of active input lines that synapse onto the pyramidal cell, independent of whether they've been modified.
- P cell measures number of active synapses that have been modified, S. Has absolute threshold T (not shown).
- Cell should fire if S > *f* A and S > T.

How to Set the Thresholds?

- Maximal similarity strategy: choose T and f that cause the smallest number of cells to be in the wrong state. (May not be biologically realizable.)
- Staircase strategy: start with small f and high T. Lower T until enough cells become active. Then raise f slightly and lower T to restore the activity level. Repeat until can no longer maintain activity level or $f = 1$.
- Competitive strategy: set $f = 0$ and lower T until the required activity level is reached. This is a k-winner-take-all strategy.
- Measure performance as: # of perfectly recalled patterns divided by total # of patterns. Used 1000 patterns in most experiments.

Comparing Threshold Setting Methods

Two Layer

Three Layer

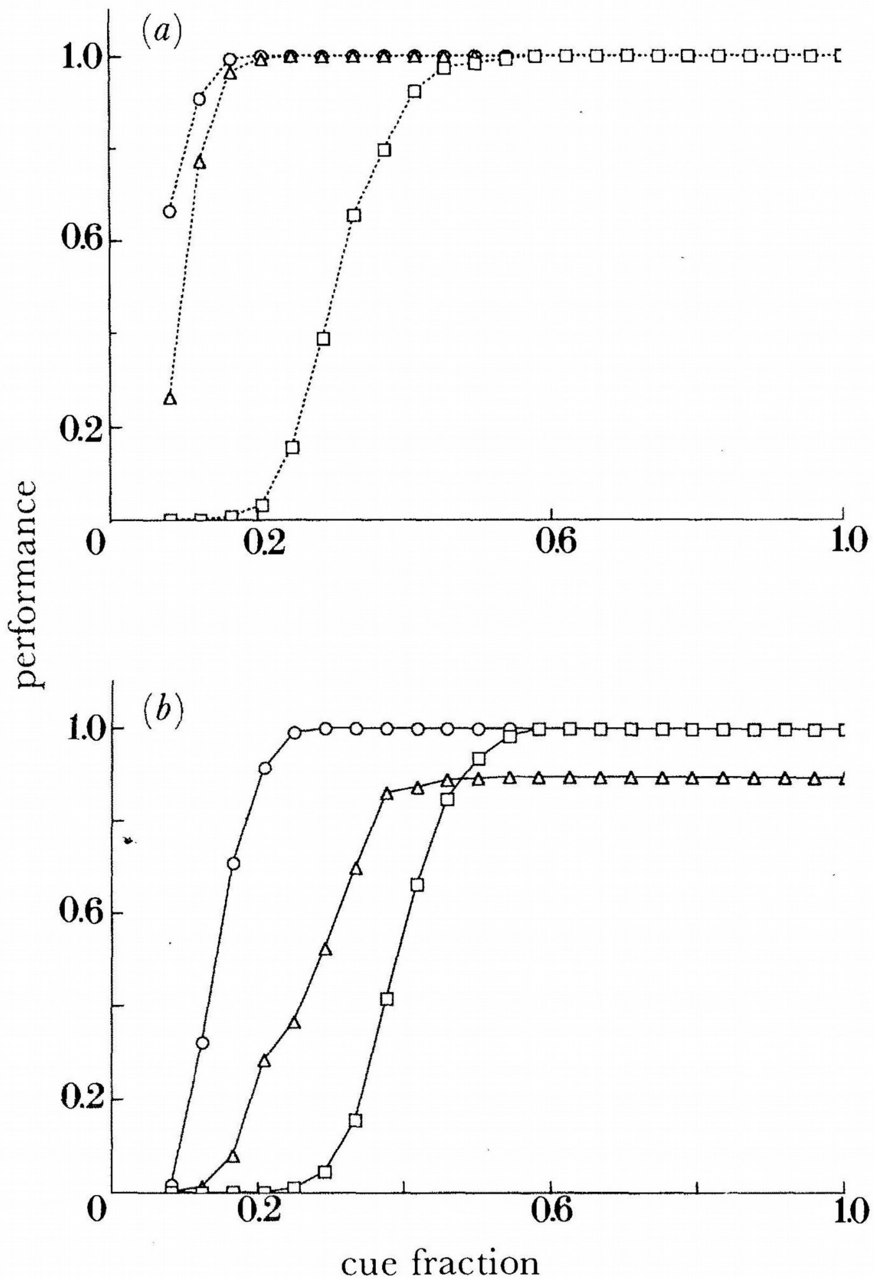

o - max similarity
Δ - staircase
❐ - competitive

Effect of Collaterals

- Marr estimates that the collaterals should have made their full contribution to recovering the event in about 3 cycles. Additional cycles would provide no benefit.
- McNaughton's commentary:
  - Oscillating cycle of excitation and inhibition in hippocampus, known as the theta rhythm: around 7 Hz (140 msec cycle).
  - Hippocampal cell output is phase-locked to the theta rhythm.
  - Assume pattern completion takes place in the ¼ cycle where excitation is increasing: 35 msec window.
  - Conduction delay and synaptic delay total 6–8 msec.
  - This leaves room for just 4–6 cycles in that 35 msec window: very close to Marr's prediction.

Assessment of Marr's Theory

- Strong points:
  - Sparse, topographic connectivity: more biologically realistic.
  - Multiple inhibitory mechanisms: subtraction and division.
  - Predicts when recurrent collaterals will help retrieval.
  - Anticipated many important findings: LTP, division operations, information transfer during sleep.
- Weak points:
  - Ignores the trisynaptic circuit (EC → DG → CA3 → CA1). It seems like $\mathcal{P}_1$ is neocortex, $\mathcal{P}_2$ is EC, and $\mathcal{P}_3$ is CA3.
  - Says nothing about DG or CA1. Ignores the direct perforant path input to CA1.
  - Claim that three layers of cells are necessary was unjustified.
  - Unanswered question: how are memories transferred from hippocampus to the neocortex?