Basics

•••

Outline

Introduction

- Synaptic transmission
- Plasticity and learning

Papers

- Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue
- A synaptic learning rule for exploiting nonlinear dendritic computation

I. Synaptic transmission

Axonal boutons

An axonal bouton, also known as a presynaptic terminal, is a specialized varicosity on an axon that can form a synapse with a dendritic spine. The axonal bouton contains synaptic vesicles that store neurotransmitters, which are released into the synaptic cleft when an action potential arrives at the bouton.

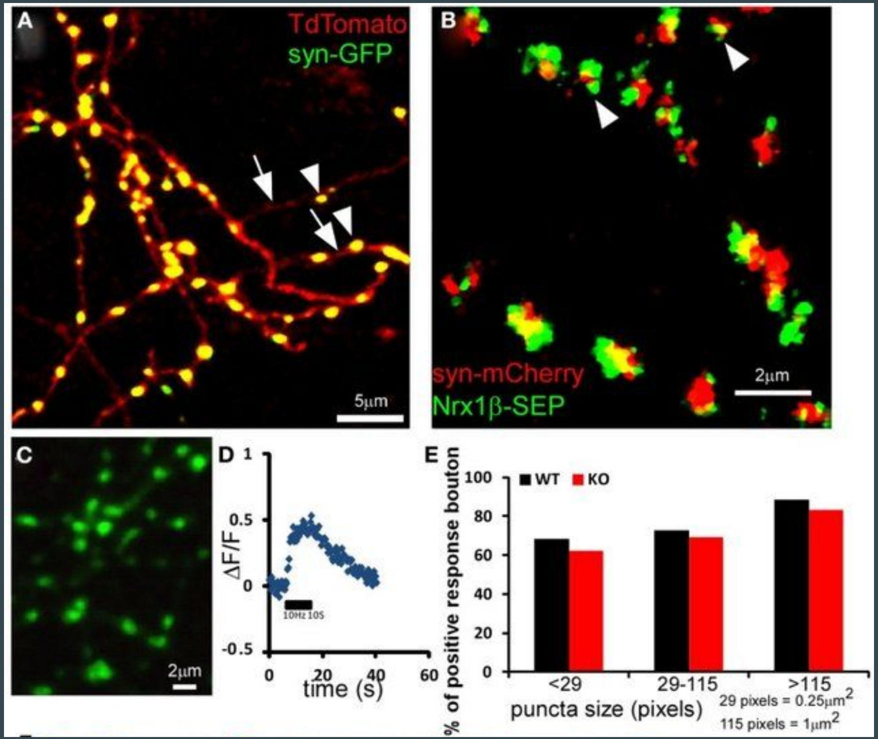

Yu Fu et al., 2012

Dendritic spines

A dendritic spine is a small protrusion on a dendrite that receives signals from other neurons. Dendritic spines are highly dynamic structures that can change shape and size in response to neuronal activity, and thus they play a crucial role in synaptic plasticity and learning.

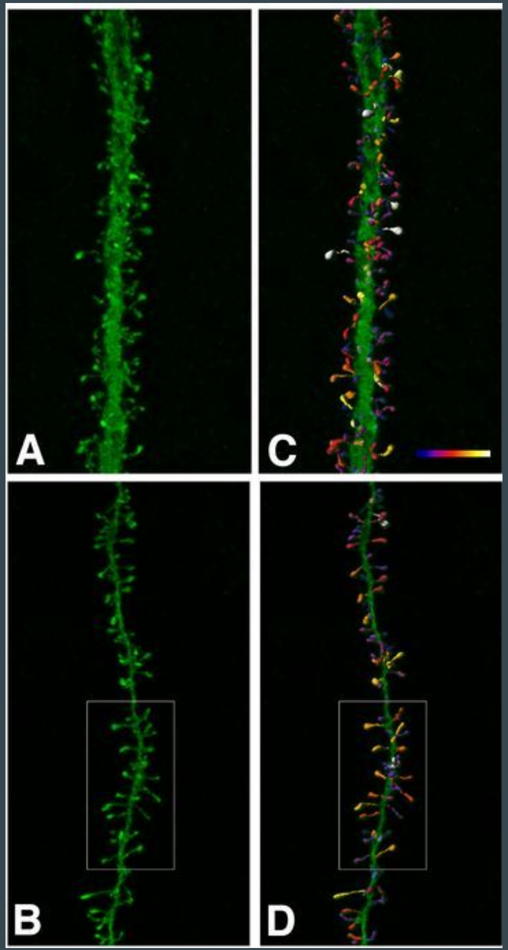

Ofer et al. 2022

Synaptic transmission

Synaptic transmission is the process by which information is transmitted from one neuron's axons (presynaptic) to another neuron's dendrites (postsynaptic) across a tiny gap called the synapse.

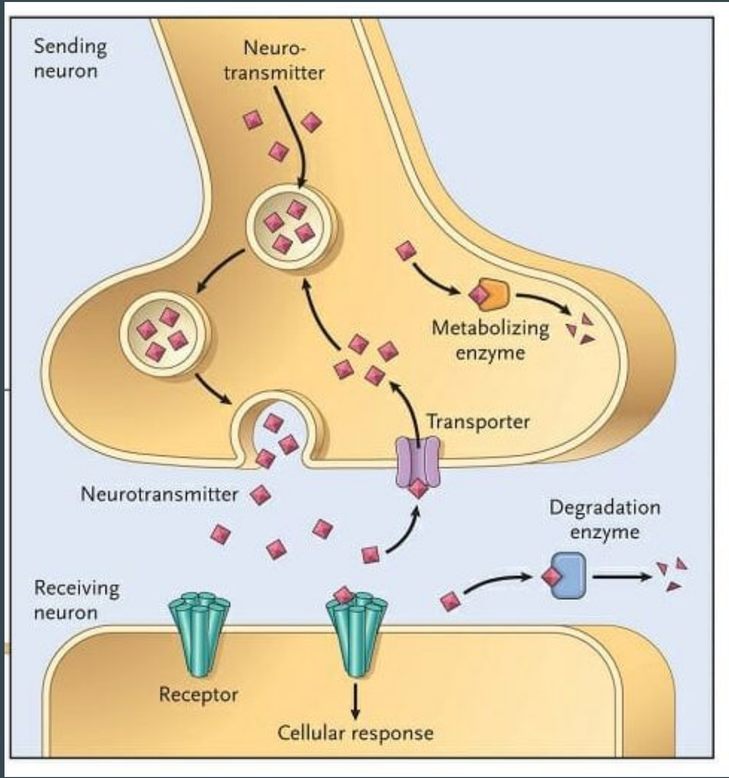

wikipedia

Synaptic transmission

The transmission of information across the synapse is accomplished through the release of neurotransmitters, which are chemical messengers that are synthesized and stored in the presynaptic neuron. When an action potential reaches the presynaptic neuron, it causes the release of neurotransmitters into the synaptic cleft, where they bind to specific receptors on the postsynaptic neuron or effector cell.

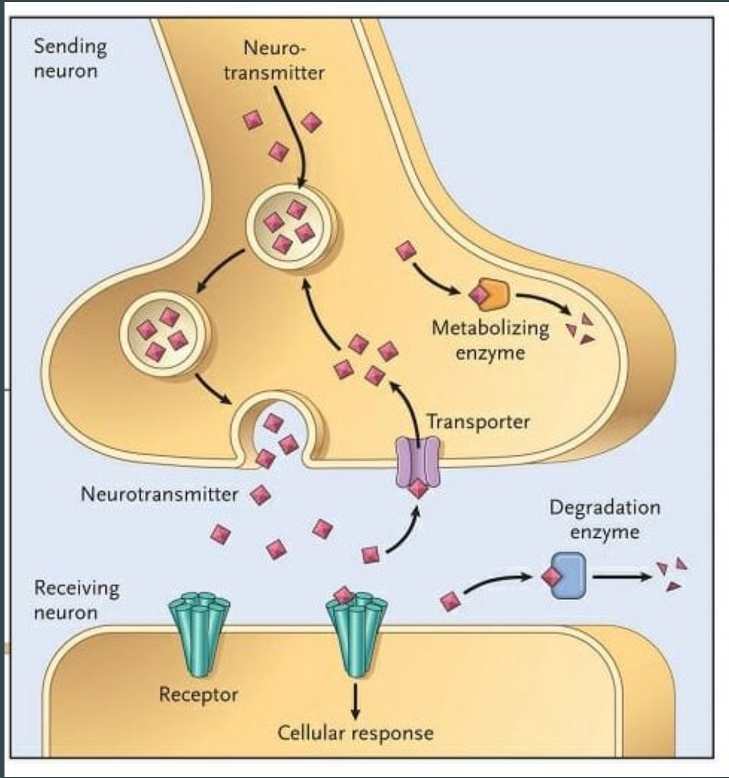

Synaptic transmission

The binding of neurotransmitters to receptors can either excite or inhibit the activity of the postsynaptic neuron or effector cell, depending on the type of receptor and neurotransmitter involved.

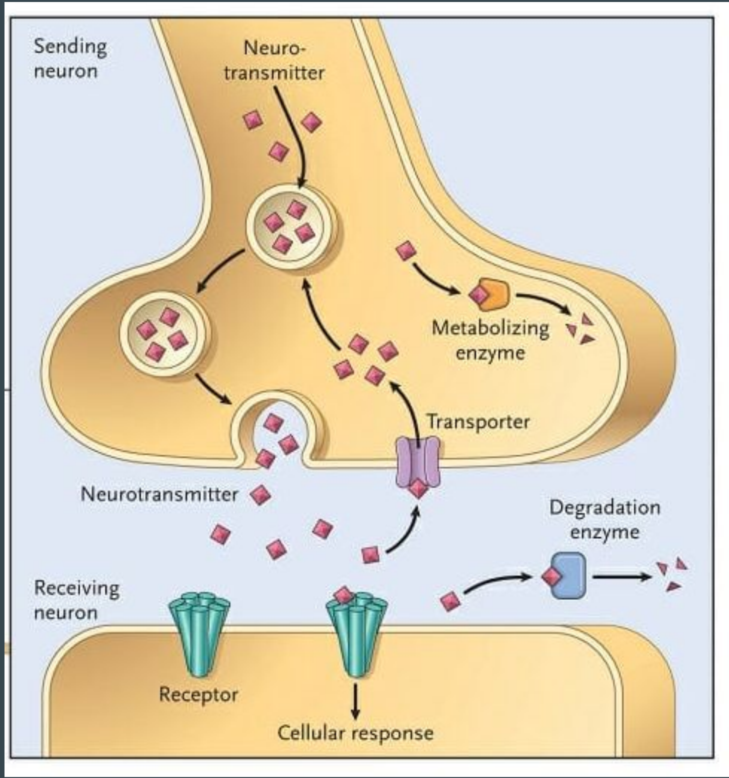

Synaptic transmission

The strength and duration of synaptic transmission can be modulated by a variety of factors, such as the availability of neurotransmitters, the number and sensitivity of receptors on the postsynaptic neuron and the activity of other neurons that make connections with the postsynaptic neuron.

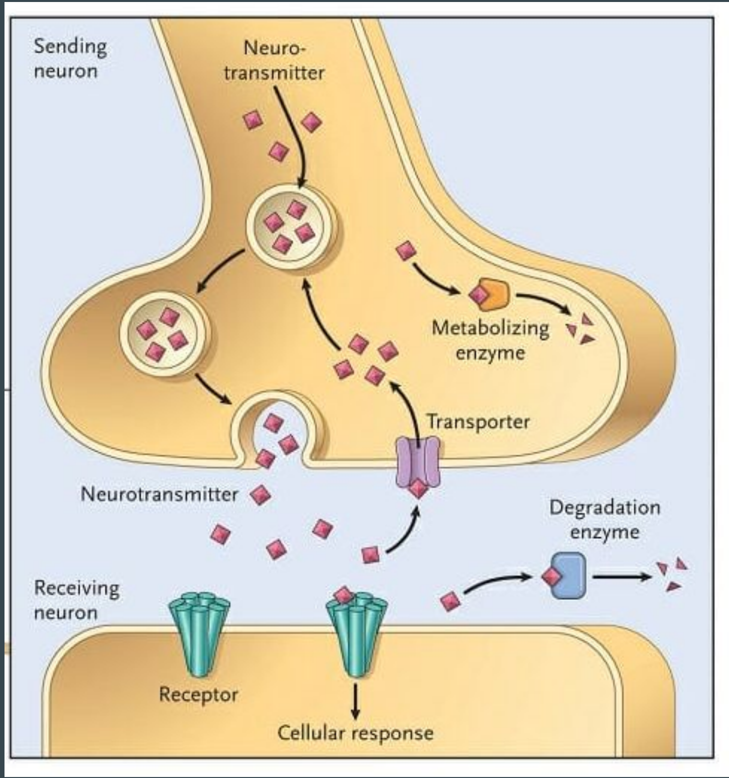

II. Learning / Plasticity

Synaptic Plasticity

Synaptic plasticity is the ability of synapses to become stronger or weaker over time, in response to increases or decreases in their activity, i.e. the action potentials that reach them. Synaptic plasticity is one of the most important neurochemical foundations of learning and memory.

Synaptic Plasticity

The strength of a synapse can be controlled by a variety of mechanisms that modulate the release of neurotransmitters from the *presynaptic* neuron or the sensitivity of the *postsynaptic* neuron to neurotransmitters.

presynaptic

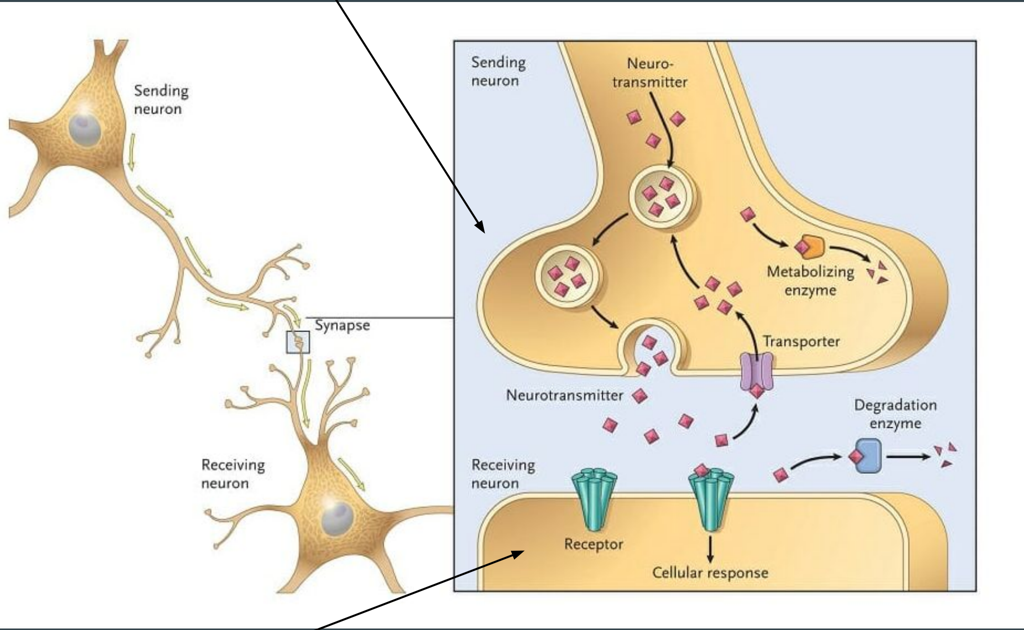

postsynaptic

Synaptic Plasticity

There are several ways in which synaptic strength can be controlled presynaptically:

a) Autoreceptors / reuptake: Many neurons have autoreceptors, i.e. a receptor for a neurotransmitter that is expressed on the same neuron that releases it. Once released into the synaptic cleft, neurotransmitters follow the rules of Brownian motion and they can be reabsorbed by the neuron. This can terminate the signal and inhibit or facilitate the release of further quantities of the neurotransmitter.

Synaptic Plasticity

b) Modulation of calcium influx: The influx of calcium ions into the presynaptic terminal is a key step in the release of neurotransmitter. The level of calcium influx can be regulated by changes in the extracellular calcium concentration, activation of voltage-gated calcium channels, and modulation of calcium-binding proteins. Changes in calcium influx can alter the amount of neurotransmitter released, and thus modulate synaptic strength.

Synaptic Plasticity

c) Modulation of extracellular matrix proteins: The extracellular matrix surrounding the presynaptic terminal can influence the release of neurotransmitter by modulating the diffusion and availability of the involved molecules. Changes in the composition or structure of the extracellular matrix can alter synaptic strength.

Synaptic Plasticity

There are several ways in which synaptic strength can be controlled postsynaptically, i.e. after neurotransmitters have been released from the presynaptic terminal:

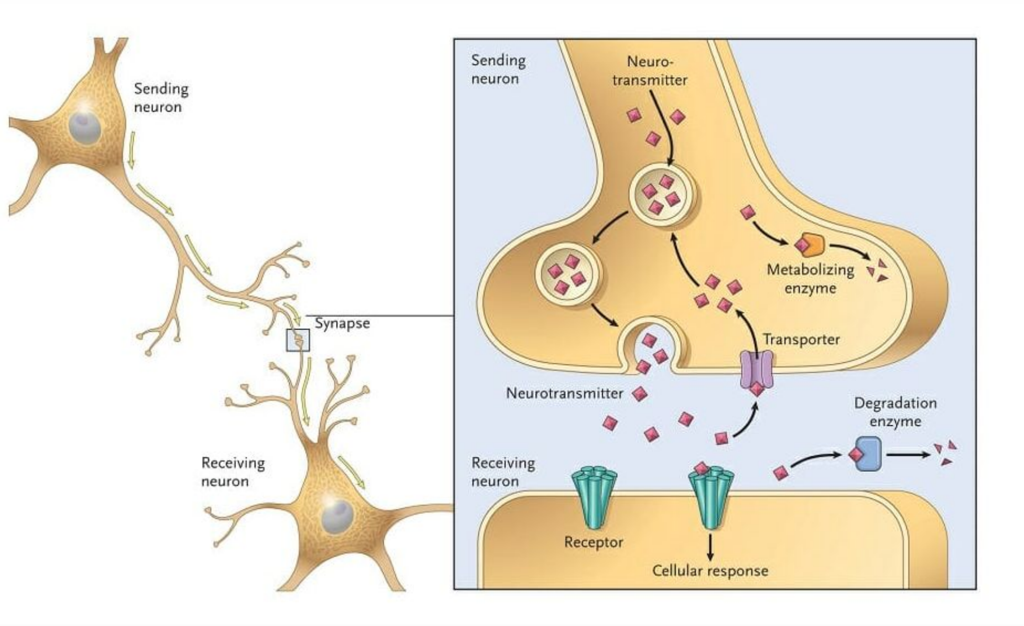

Synaptic Plasticity

There are several ways in which synaptic strength can be controlled postsynaptically:

a) Receptor desensitization: Receptors can become desensitized to neurotransmitters if they are exposed to high levels of it for an extended period of time. This can reduce the responsiveness of the postsynaptic neuron to subsequent neurotransmitter release, effectively decreasing the strength of the synapse.

Synaptic Plasticity

b) Number of receptors: The number and distribution of receptors on the postsynaptic membrane can be regulated by various mechanisms, including receptor insertion and removal from the membrane, as well as lateral movement of receptors within the membrane. Changes in receptor composition can alter the sensitivity of the postsynaptic neuron to neurotransmitters.

Synaptic Plasticity

c) Modulation of intracellular signaling pathways: The activation of neurotransmitter receptors can trigger intracellular signaling pathways that can lead to changes in gene expression, protein synthesis, and the morphology of dendritic spines. These changes can alter the structure and function of the synapse, resulting in changes of synaptic strength.

Long term potentiation (LTP)

Long-term potentiation (LTP) is a subcellular mechanism of synaptic plasticity that underlies learning and memory in the brain. It refers to the long term strengthening of synaptic connections between neurons that occurs in response to repeated neuronal activity. The process involves the activation of glutamate receptors on the *postsynaptic* membrane, which leads to an influx of calcium ions. This triggers a cascade of intracellular signaling pathways that ultimately result in the insertion of new AMPA receptors into the postsynaptic membrane or increase of their strength, increasing the sensitivity of the postsynaptic neuron to neurotransmitter release, effectively strengthening the synaptic connection.

Long term depression (LTD)

Long-term depression (LTD) has the opposite effect of LTP, i.e. the long term decrease in synaptic strength. LTD also occurs in response to repeated neuronal activity. LTD leads to a decrease in the number or the strength of AMPA receptors on the postsynaptic membrane, therefore decreasing the sensitivity of the postsynaptic neuron and weakening the synaptic transmission.

Spike-timing-dependent plasticity (STDP)

STDP is a form of synaptic plasticity that depends on the relative timing of pre- and postsynaptic action potentials. The strength of a synaptic connection between two neurons is modified based on the temporal relationship between the action potentials generated by the two neurons. If the presynaptic neuron fires before the postsynaptic neuron, the strength of the synaptic connection is potentiated, whereas if the postsynaptic neuron fires before the presynaptic neuron, the strength of the synaptic connection is depressed.

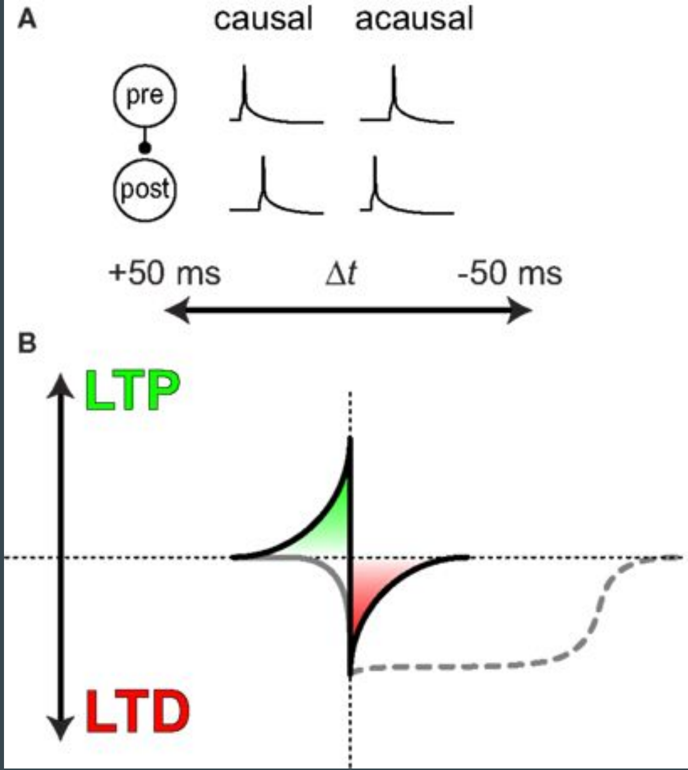

Markram et al. 2011

Dendrites: bug or feature?

Hausser and Mel

Dendrites: bug or feature?

The integrative properties of dendrites are determined by a complex mixture of factors, including their morphology, the spatio-temporal patterning of synaptic inputs, the balance of excitation and inhibition, and neuromodulatory influences, all of which interact with the many voltage-gated conductances present in the dendritic membrane. Recent efforts to grapple with this complexity have focused on identifying functional compartments in the dendritic tree, the number and size of which depend on the aspect of dendritic function being considered. We discuss how dendritic compartments and the interactions between them help to enhance the computational power of the neuron and define the rules for the induction of synaptic plasticity.

Dendrites: bug or feature?

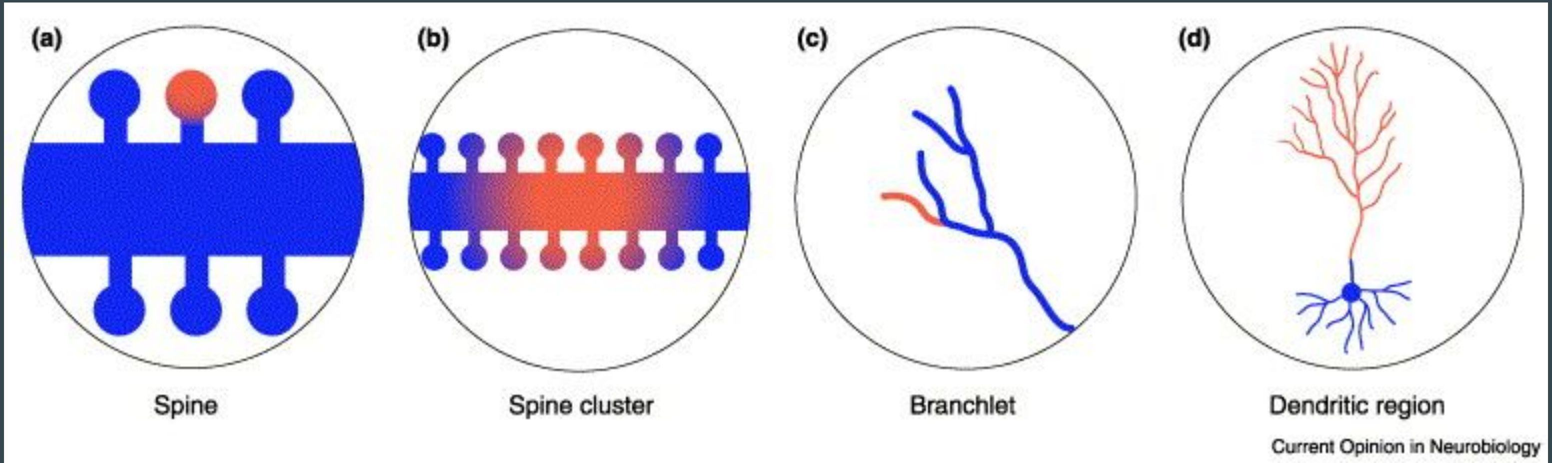

What are the functional compartments in neurons? A schematic representation of different levels of granularity in neuronal processing. (a) Calcium signalling restricted to single spines. (b) Signalling restricted to a small cluster of spines. (c) Signalling restricted to a single terminal branchlet. (d) Signalling extending across the entire apical dendritic tree.

Dendrites: bug or feature?

Two extreme cases:

- Point neuron hypothesis: the neuron functions as a simple one-compartment (usually linear) summing unit, where, all synapses have an equal opportunity to influence neuronal output through the axon.
- Spatio-temporal interactions among synaptic inputs and the local responses they trigger may suggest the importance of dendritic space and time for various aspects of neuronal information processing:
  - back- and forward-propagating action potentials (APs)
  - synaptic inputs to spatially defined dendritic compartments
  - synaptic plasticity

Dendrites: bug or feature?

Point neuron (classic approach)

The rule for combining the effects of many synapses under this model is generally assumed to be linear, and can thus be expressed as a weighted sum of all excitatory and inhibitory synaptic inputs.

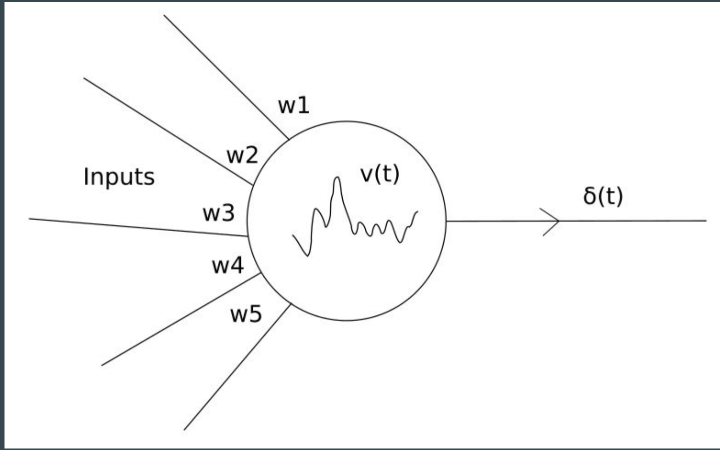

Dendrites: bug or feature?

Modern point neuron

The work of Wilfred Rall provided the first demonstration that from an electrical point of view, dendrites can be treated as spatially extended, branched coaxial cables subject to the laws of cable theory.

First, scaling of synaptic conductances depending on electrotonic distance from the soma could function to equalise the effects of synapses regardless of location.

Second, dendritic voltage-dependent Na+, Ca2+ and NMDA channels can boost the effectiveness of weak distal synaptic inputs.

Third, a dendritic normalisation, whose function is to counteract the classical synaptic saturation non-linearity, could result from a patch of voltage-dependent Ca2+ channels in the apical tree.

Dendrites: bug or feature?

Dendritic spikes have a clear voltage and stimulus intensity threshold and can occur without triggering axonal action potentials. Similarly, action potentials initiated in the axon do not propagate fully into the distal dendrites of neurons.

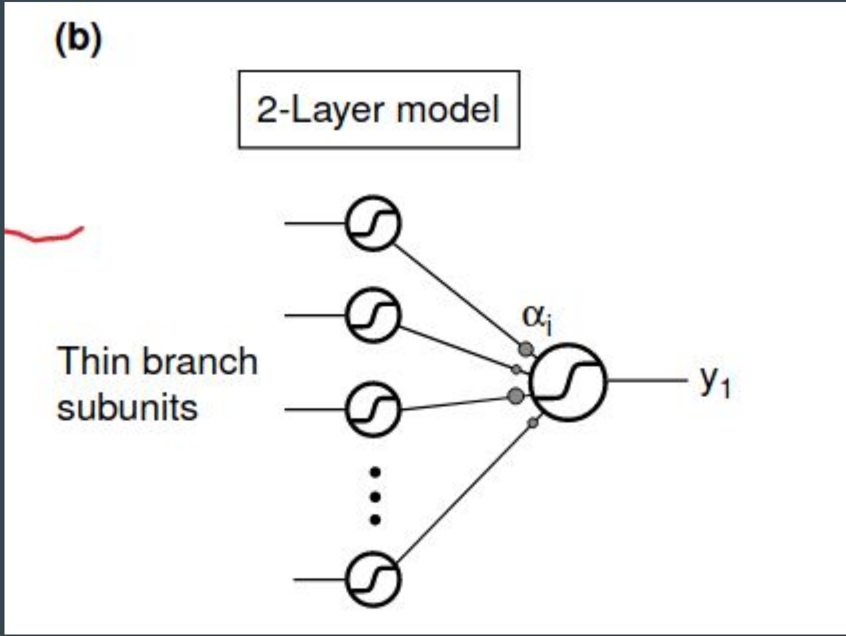

Dendrites: bug or feature?

2 layer model

one proximal compartment, including the soma, basal dendrites and axon, in which classical Na+ action potentials are generated

one distal compartment, representing the distal apical tree, in which fast Na+ and slow Ca2+ spikes are initiated

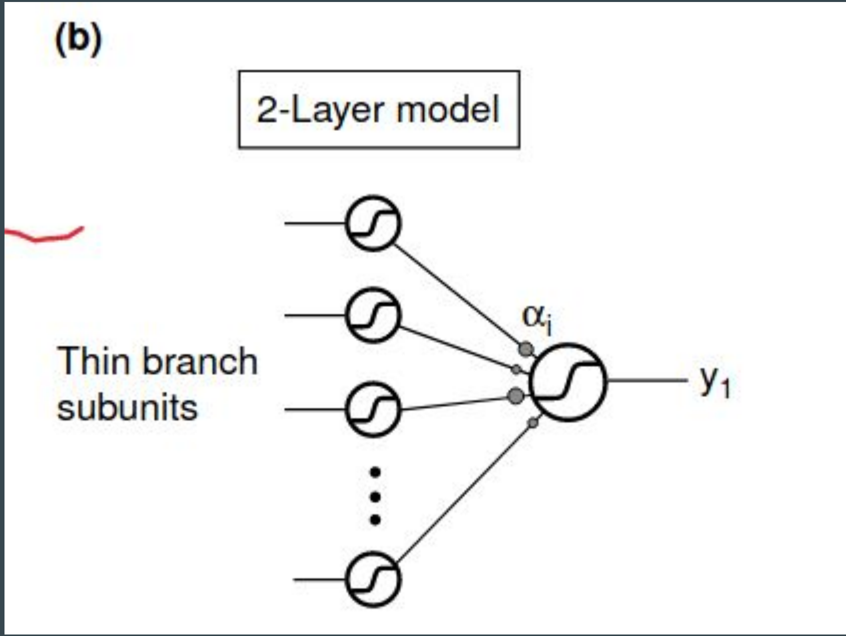

Dendrites: bug or feature?

The two-layer sum-of-sigmoids model is attractive from a computational point of view, and could have broad implications for information processing & learning.

Only with steady state input and output variables (i.e. spike rates) does not address the question of spike timing, which can be important in dendrites.

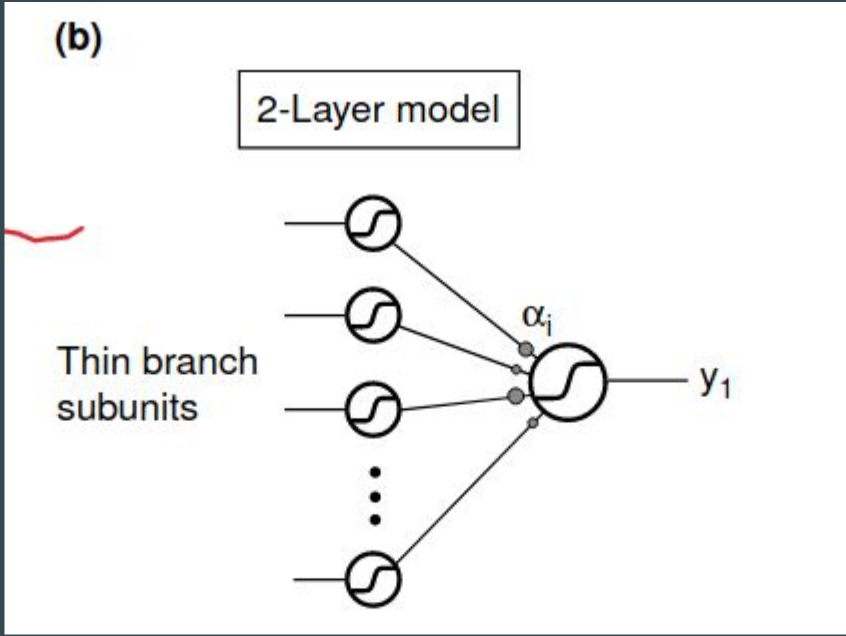

Dendrites: bug or feature?

Is the 2 layer model sufficient?

Schiller et al (2000) used focal laser-activated release of caged glutamate, to stimulate clusters of excitatory synapses (within an approximately 10 micron radius) on fine basal dendrites of neocortical pyramidal cells. They found highly localised all-or-nothing spike-like responses that were initially triggered by AMPA receptors, and followed by co-activation of voltage-dependent Na+, Ca2+ and NMDA channels -> NMDA spikes.

These constitute non linear highly localized dendritic processing

Dendrites: bug or feature?

Is the 2 layer model sufficient?

Highly localised dendritic processing can also be found in the retina, in which (Euler et al 2002) using calcium imaging techniques demonstrated that individual dendritic branches of retinal starburst amacrine cells show directionally selective calcium signals, whereas the somatic voltage response shows no such selectivity. Individual dendritic branches of amacrine cells can act as independent integrative units with branch-specific outputs.

Dendrites: bug or feature?

3-layer model

A next generation single neuron model could include a multiplicative interaction between proximal and distal integrative regions of the cell.

Overall output of such a three-layer model might be expressed using the form y1+αy2.

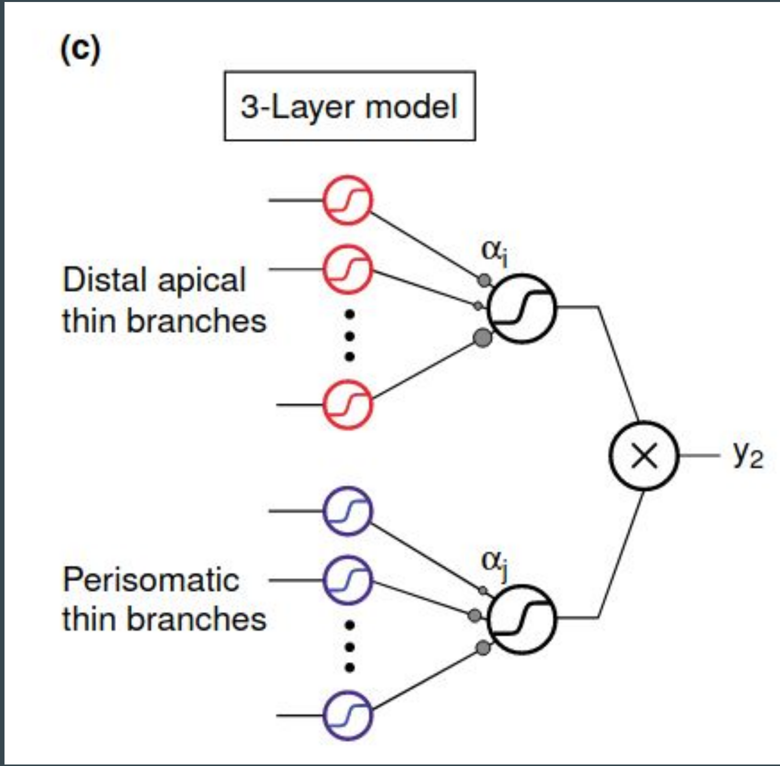

Dendrites: bug or feature?

Two independent strong synaptic inputs is very different for inputs near the soma and for those in the distal dendrites.

As the AP conductance is concentrated in the axon, distal inputs are 'protected'. Distal dendrites thus represent a separate functional compartment in which processing can continue relatively uninterrupted by somatic AP firing.

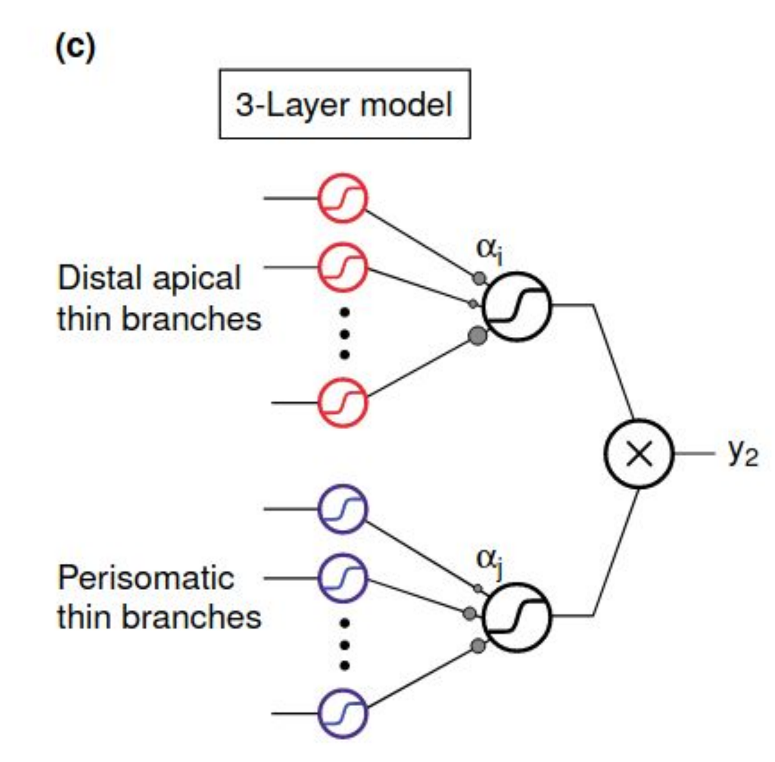

Dendrites: bug or feature?

Proximal–distal interactions could play a role in several modulatory effects in cortical sensory neurophysiology:

- contour completion
- attentional modulation
- multiplicative 'gain fields'

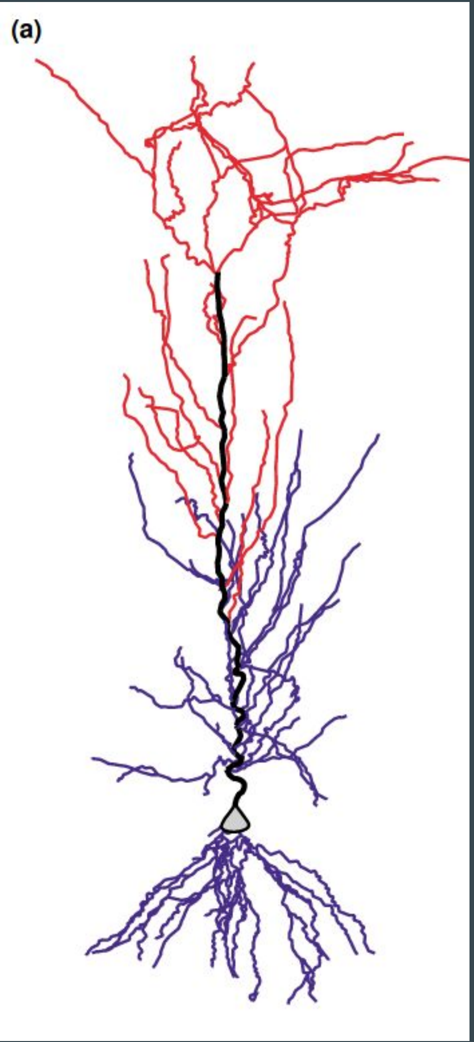

Dendrites: bug or feature?

Results / Conclusion

- Ultimately, whether particular dendritic properties represent bugs or features must be determined in the context of the intact brain.
- To link these and other aspects of dendritic phenomenology to behaviour, it is essential to develop techniques that make this possible in the awake animal.
- New approaches can help determine when and how dendrites, and their compartments, contribute to the brain’s remarkable capacities for perception, action and memory.

Are Dendrites Conceptually Useful?

Matthew E. Larkum

Are Dendrites Conceptually Useful?

This article presents the argument that, while understanding the brain will require a multi-level approach, there is nevertheless something fundamental about understanding the components of the brain. The author argues that the standard description of neurons is not merely too simplistic, but also misses the true nature of how they operate at the computational level. In particular, the humble point neuron, devoid of dendrites with their powerful computational properties, prevents conceptual progress at higher levels of understanding.

Are Dendrites Conceptually Useful?

Distinguish conceptual necessity from conceptual usefulness.

It might be possible to show formally that any neural network with dendrites can be replaced, without loss of function, with some other network using point neurons, and therefore argue that they are not conceptually *necessary*.

Neural networks (with or without dendrites) are also not conceptually necessary for computation. For example, they could be in principle, replaced by 0 - 1 s on a ticker tape with the appropriate finite state machine, i.e., a Turing machine.

Are Dendrites Conceptually Useful?

The essence of computation. (A) The Turing machine, a ticker tape and a finite state machine. (B) The key components of neural network design are architectures, learning rules and objective functions according to Richards et al. (2019).

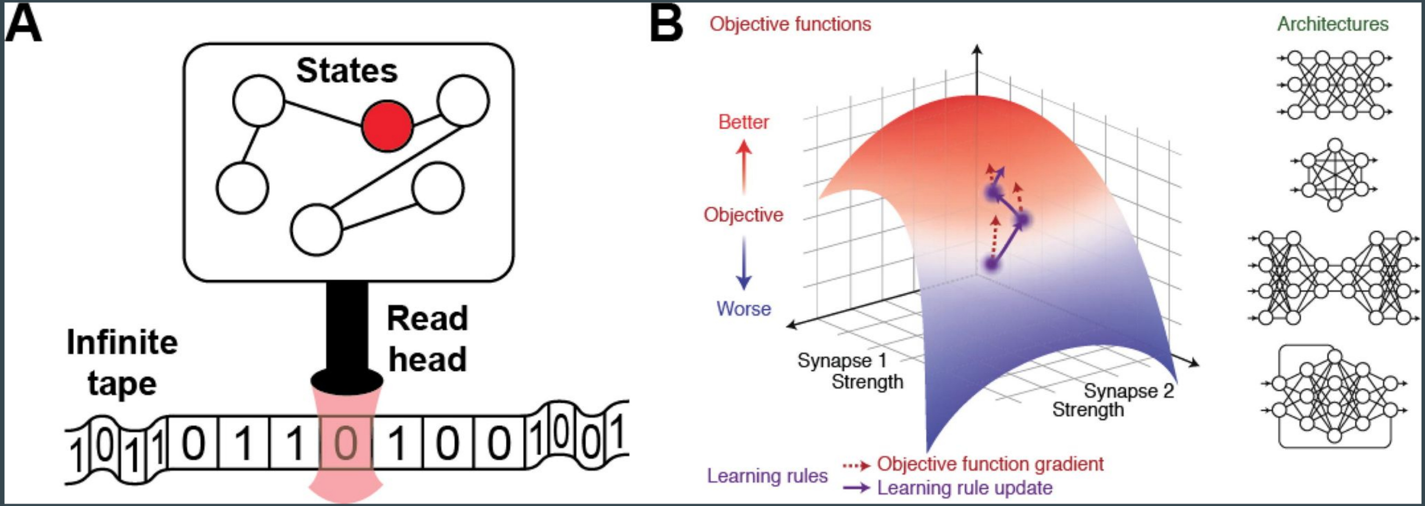

Are Dendrites Conceptually Useful?

Distinguish conceptual necessity from conceptual usefulness.

It is now unarguable, however, that the idea of neural networks (Hopfield, 1982) introduces extremely useful concepts for solving computationally challenging problems: it introduces higher level concepts such as network architecture, learning rules and cost functions.

So the question here is what are the concepts that dendrites introduce that are useful for different levels of computations and cannot be sufficiently explained by point neuron models?

Are Dendrites Conceptually Useful?

Computers can be understood at high level functions because we know they work on binary systems (as we build them).

However, that's not known for neurons. If neurons *don't* actually behave like point neurons, many of the models in neuroscience are being built on false assumptions, and they will impact our understanding of how the brain actually works.

Are Dendrites Conceptually Useful?

The operation of a point neuron is usually described as “integrate-and-fire”:

- linearly summated weighted synaptic inputs
- determine, via an activation function
- if the neuron should be ‘on’ or ‘off’.

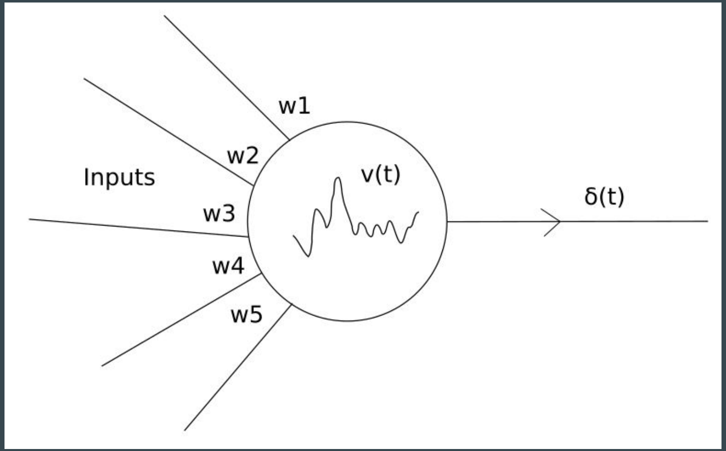

Are Dendrites Conceptually Useful?

Neural networks

The artificial neural network comes historically from an attempt to encapsulate the essential features of brains as parallel distributed processors.

It might be possible to show formally that any neural network with dendrites can be replaced, without loss of function, with some other network using point neurons, and therefore argue that they are not conceptually necessary.

Neural networks (with or without dendrites) are also not conceptually necessary for computation, per se. Any standard neural network can, in fact, run on a digital computer

Are Dendrites Conceptually Useful?

A common phenomenon involving barrages of synaptic input are so-called up/down-states, slow (1 Hz) oscillations between two membrane potential values, usually observed under anesthesia. Interestingly, up-state has a relatively fixed and stable amplitude all over the dendritic tree.

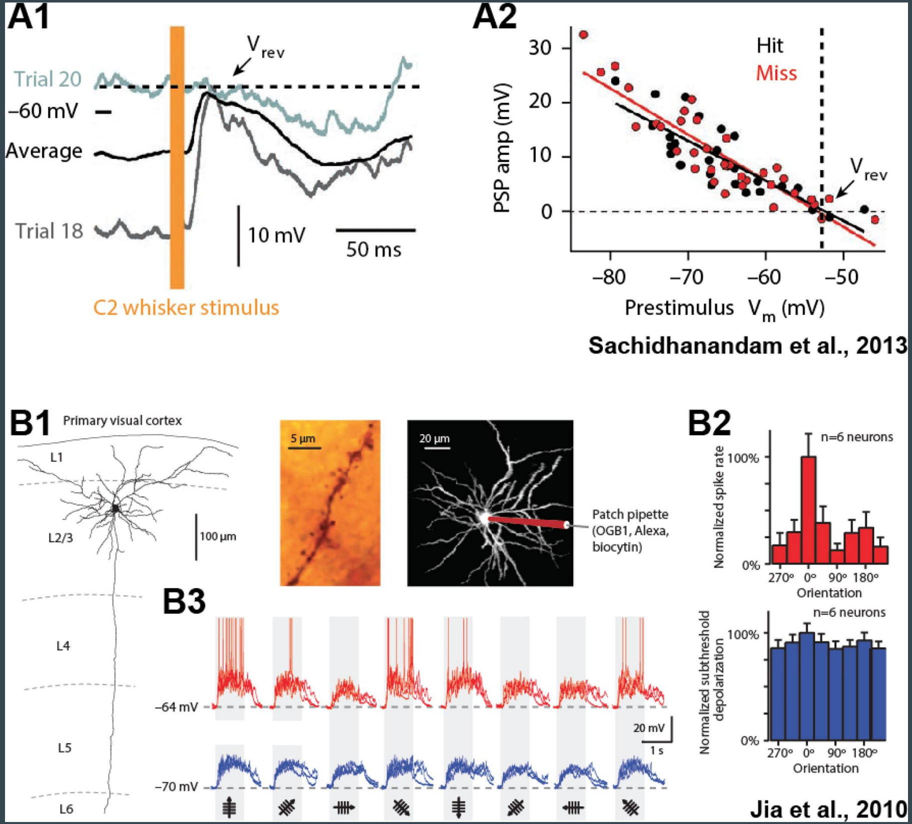

Are Dendrites Conceptually Useful?

During visual stimuli with orientated driftings that cause broadly similar depolarization of the neuron for each stimulus presentation, the preferred orientations robustly cause much greater action potential firing than others (Jia et al. 2010).

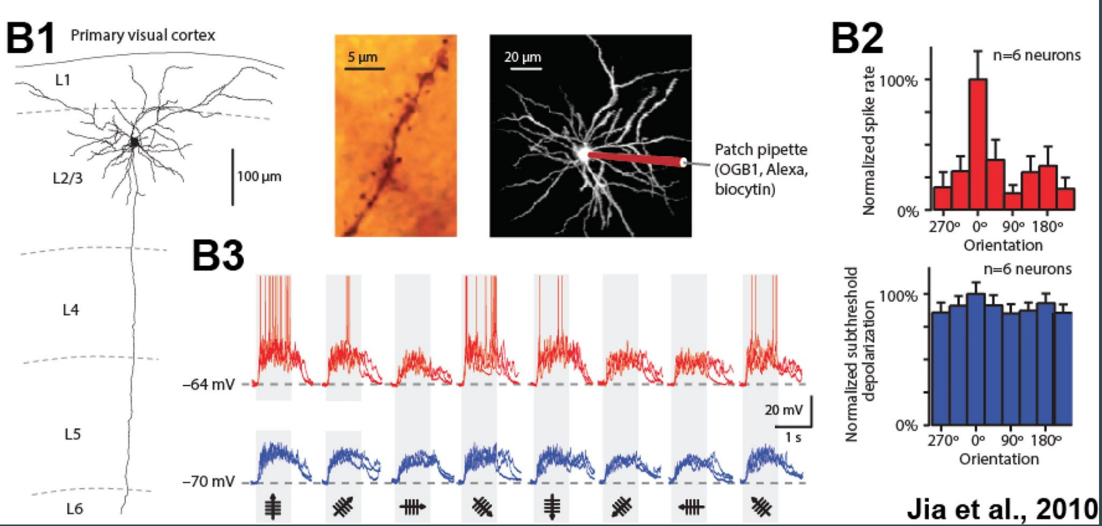

Are Dendrites Conceptually Useful?

Three experiments (Chen et al., 2013; Smith et al., 2013; Palmer et al., 2014) were carried out under similar conditions: in vivo intracellular (patch-clamp) recordings from layer 2/3 pyramidal neurons in anesthetized rodents.

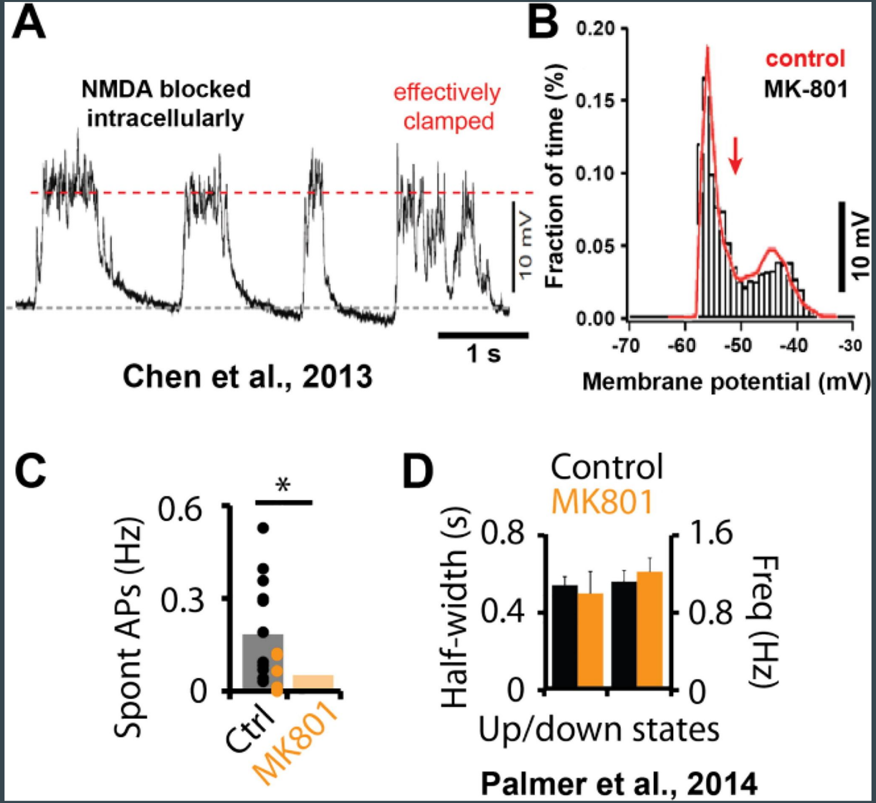

Are Dendrites Conceptually Useful?

Firstly, they showed that the amplitude of the depolarization caused by the up-state was unchanged by NMDA receptor channel block. This implies that AMPA receptors (the other main excitatory receptor type) are predominantly involved in subthreshold depolarization of the neuron.

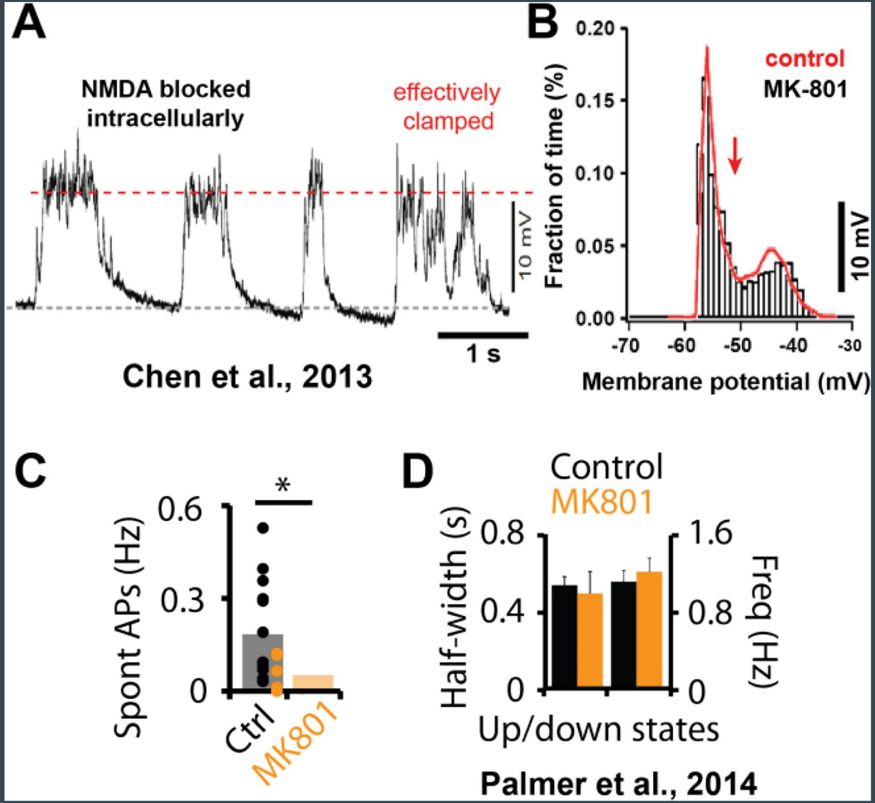

Are Dendrites Conceptually Useful?

More remarkably, AP generation was very much affected by NMDA receptor channel block (C) despite the fact that the amplitude of the up- states was unchanged.

In other words, it is the receptors that don't affect the amplitude of up-states (NMDA) that cause action potentials, not the ones that do (AMPA).

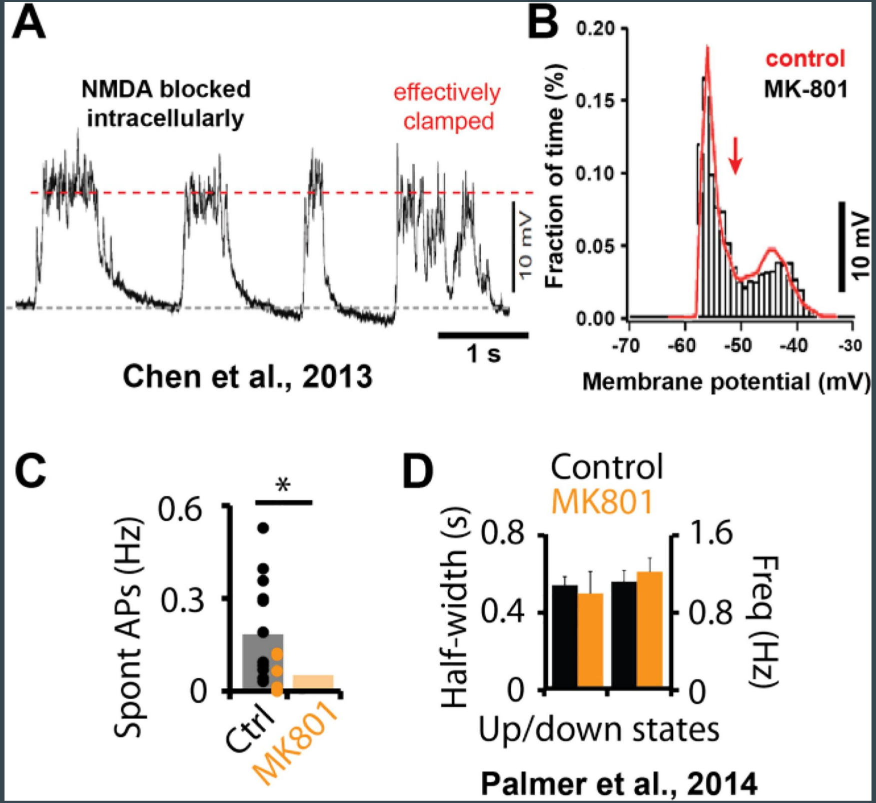

Are Dendrites Conceptually Useful?

Hypothesis for spike generation

The fixed up-state value results from a relatively uniform balance of excitatory (AMPA) and inhibitory (GABA) receptors such that the effective reversal potential is always the same. Regardless of the exact level of network activity, the balance of E / I needs to be constant as frequently observed

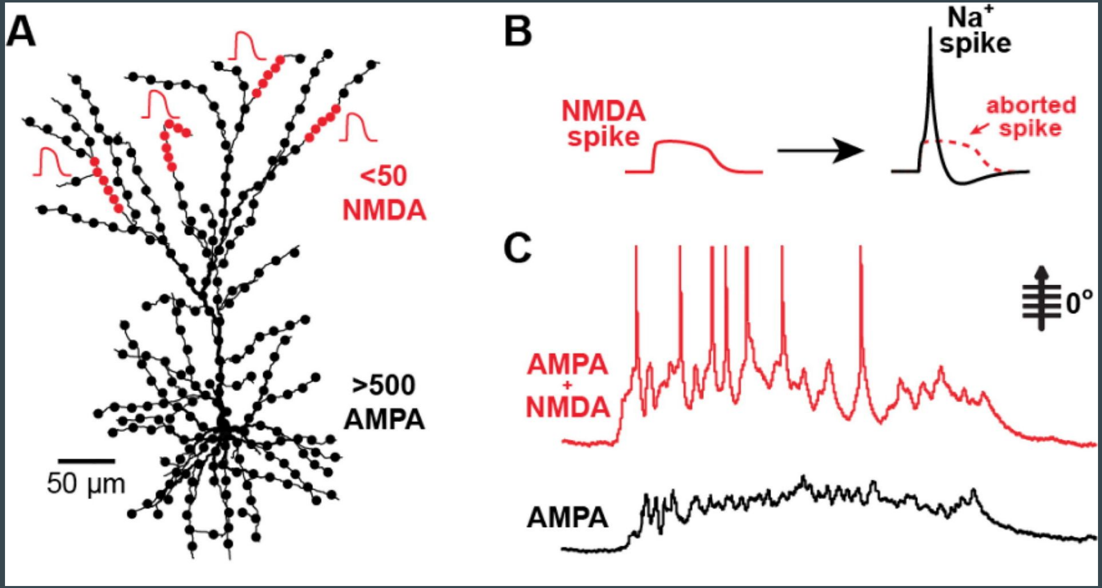

Are Dendrites Conceptually Useful?

Hypothesis for spike generation

Balanced E/I input effectively clamps the neuron at a fixed depolarization, reducing the impact of random fluctuations that might otherwise generate AP, and is consistent with the fact that intracellular blockade of NMDA receptors has little or no effect.

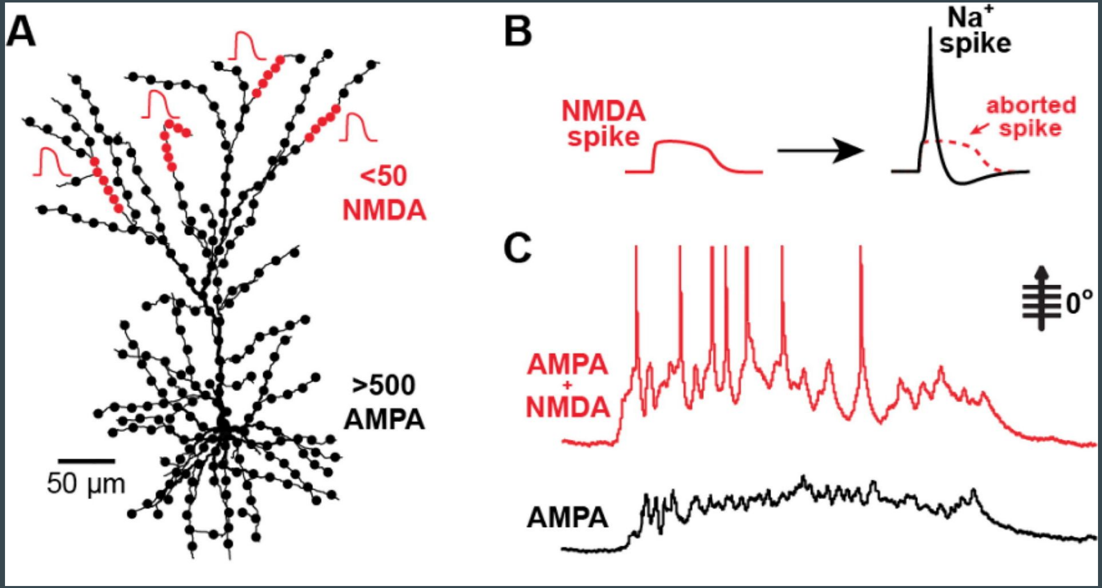

Are Dendrites Conceptually Useful?

Hypothesis for spike generation

Few (10) clustered inputs in thin dendrites can induce NMDA receptors to cooperatively open, generating local NMDA spikes (Larkum and Nevian, 2008; Larkum et al., 2009) causing ~10 mV depolarization at the cell body that can exceed threshold when starting from the up-state depolarization level.

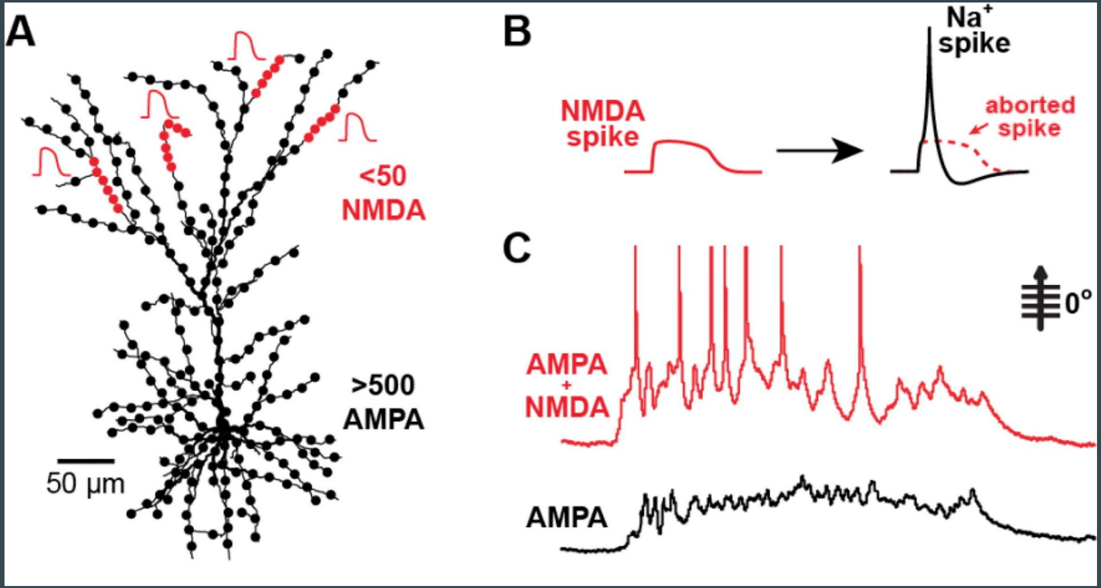

Are Dendrites Conceptually Useful?

Hypothesis for spike generation

Thousands of inputs combine to clamp the neuron at a particular subthreshold value, while a handful of synapses several orders of magnitude less in number but with explosive impact dictate the firing of the neuron

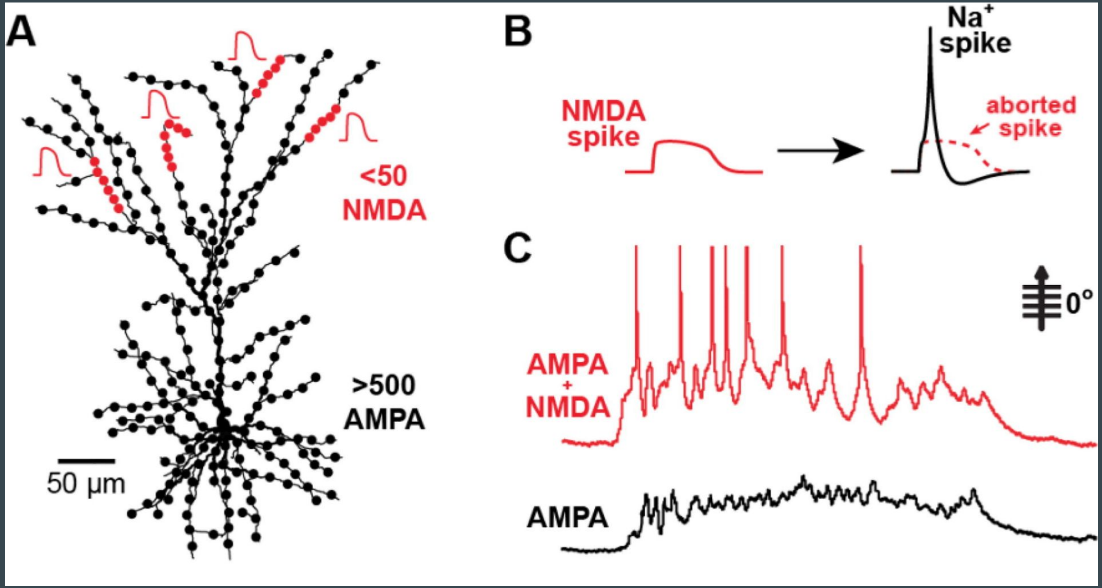

Are Dendrites Conceptually Useful?

Orientation-selective neurons in visual cortex are a case in point.

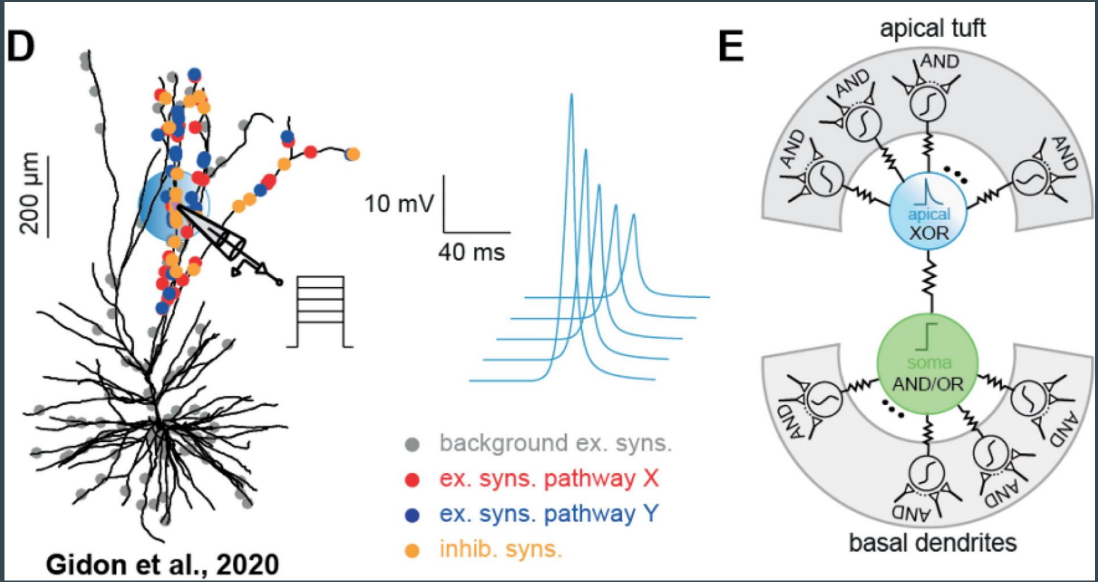

Are Dendrites Conceptually Useful?

Highlights

- A learning rule derived from cable theory is used in biophysical simulations
- Pyramidal cell I/O functions can be optimized for computation by synaptic plasticity
- Active and passive dendritic mechanisms enhance input pattern discrimination
- Single neurons can learn network-level computations simply by tuning synaptic weights

Combinations

A combination is a selection of items from a set of distinct objects, such as the order of selection does not matter. When the order of selection is important then we are referring to permutations instead.

Combinations

For example a combination of two balls from a set of 5 balls with different colors

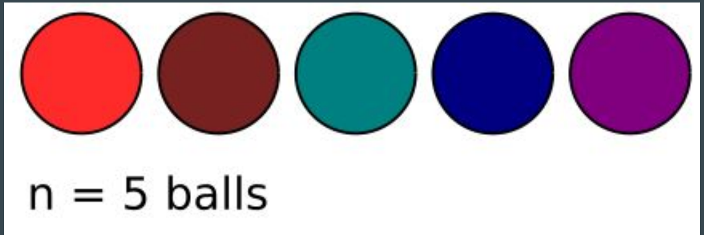

Combinations

For example a combination of two balls drawn from a set of 5 balls with different colors, can be any of the following:

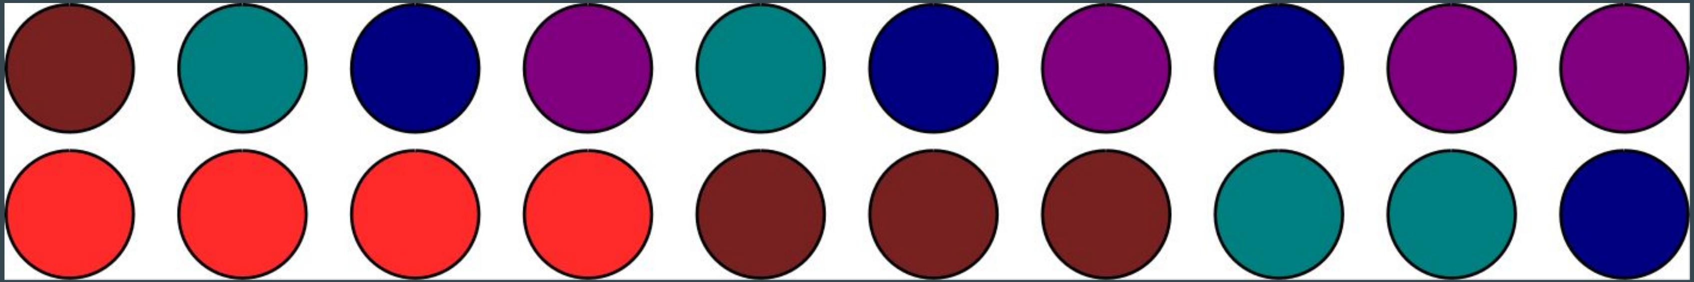

Combinations

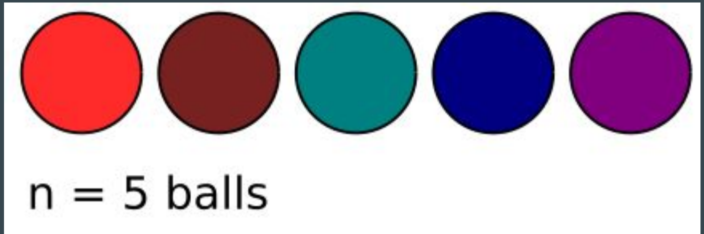

In order to compute the number of all the possible combination of k=2 balls from a set of n=5 balls, we can use the formula:

$$\binom{n}{k} = \frac{n!}{(n-k)!k!}$$

Combinations

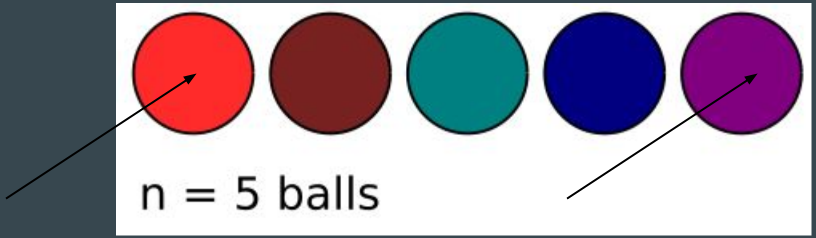

In order to compute the number of all the possible combination of k=2 balls from a set of n=5 balls, we can use the formula:

$$\binom{5}{2} = \frac{5!}{3!2!}$$

What if we want to compute combinations?

$$\binom{5}{3}$$

Combinations

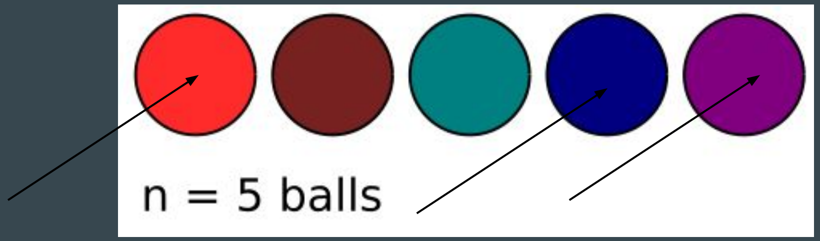

In order to compute the number of all the possible combination of k=3 balls from a set of n=5 balls, we can use the formula:

$$\binom{5}{3} = \frac{5!}{2!3!} = \frac{5!}{3!2!} = \binom{5}{2}$$

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

Panayiota Poirazi and Bartlett W. Mel

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

We consider the combined effects of active dendrites and structural plasticity on the storage capacity of neural tissue. We compare capacity for two different modes of dendritic integration: (1) linear, where synaptic inputs are summed across the entire dendritic arbor, and (2) nonlinear, where each dendritic compartment functions as a separately thresholded neuron-like summing unit. We calculate much larger storage capacities for cells with nonlinear subunits and show that this capacity is accessible to a structural learning rule that combines random synapse formation with activity-dependent stabilization/elimination. In a departure from the common view that memories are encoded in the overall connection strengths between neurons, our results suggest that long-term information storage in neural tissue could reside primarily in the selective addressing of synaptic contacts onto dendritic subunits.

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

Cell is modeled as a set of m identical branches connected to a soma, where each branch contains k excitatory synaptic contacts.

Each synapse is driven by one of d input lines and is given a small integer-valued weight.

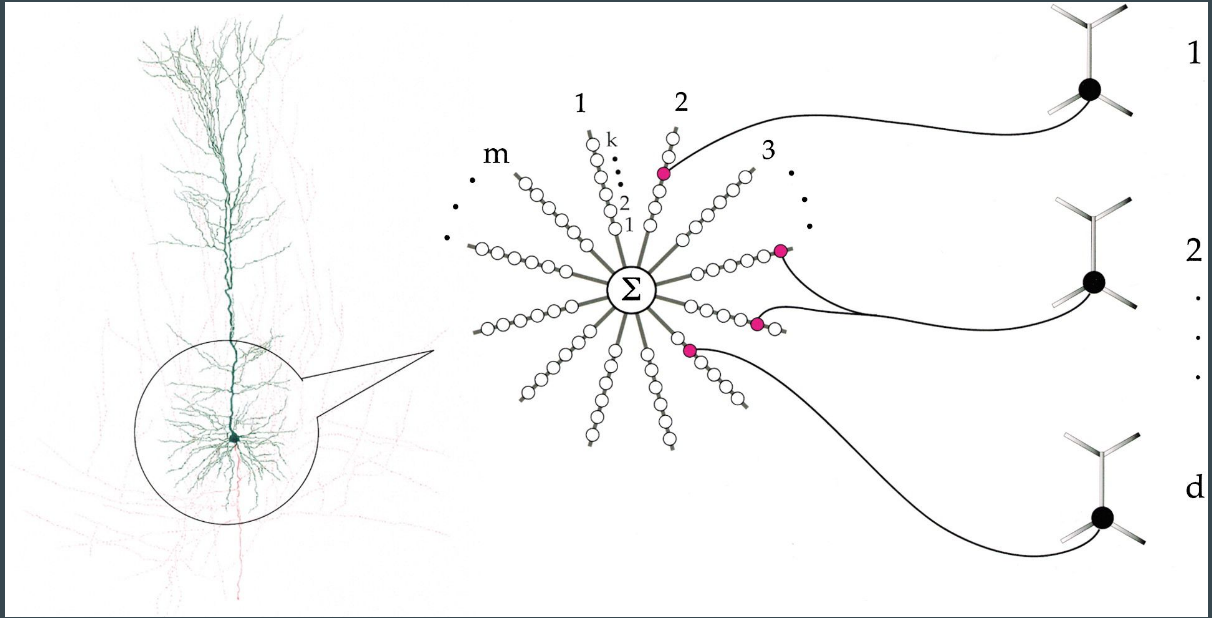

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

Two alternative models:

i) linear activation: sum of all inputs

ii) non-linear activation: integration within branches

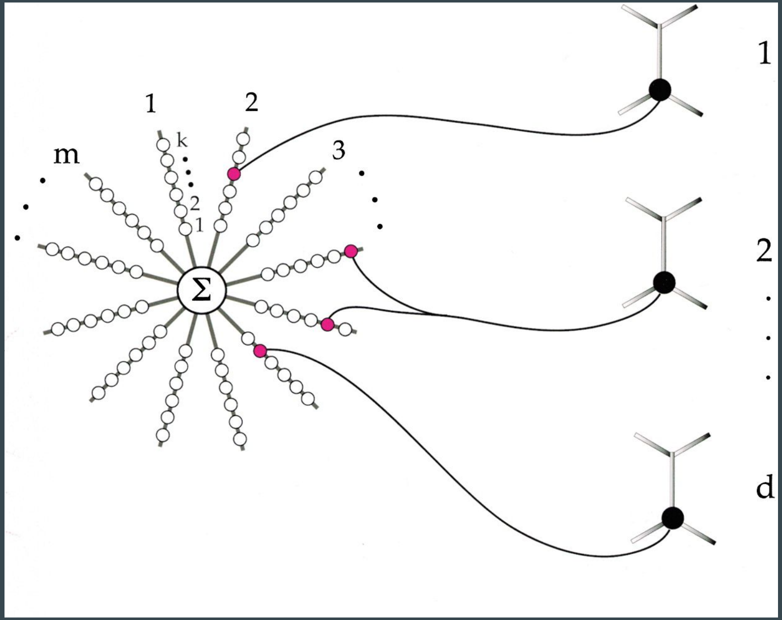

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

i) linear activation

linear integration (Equation 1), where the cell's activation level $a_L(x)$ prior to output thresholding is given by a weighted sum of inputs from across the entire cell

$$a_L(\mathbf{x}) = \sum_{j=1}^{m} \sum_{i=1}^{k} w_{ij} x_{ij}$$

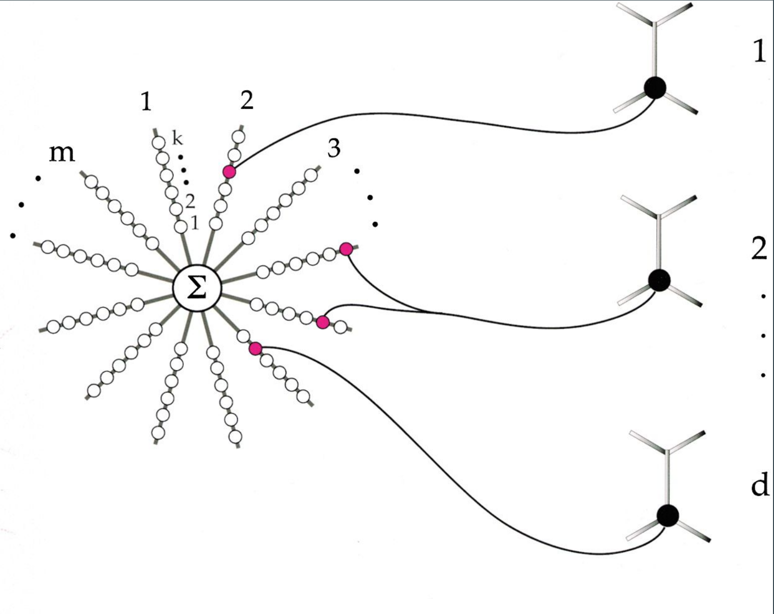

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

ii) non-linear activation

(Equation 2), where (1) the k inputs to each branch are combined in a weighted sum, (2) a static branch nonlinearity b, such as a sigmoid function, is applied to each of the m branchs subtotals, and (3) the nonlinear branch responses are summed to produce the cell's overall activation level $a_N(x)$:

$$a_N(\mathbf{x}) = \sum_{j=1}^{m} b\left(\sum_{i=1}^{k} w_{ij} x_{ij}\right).$$

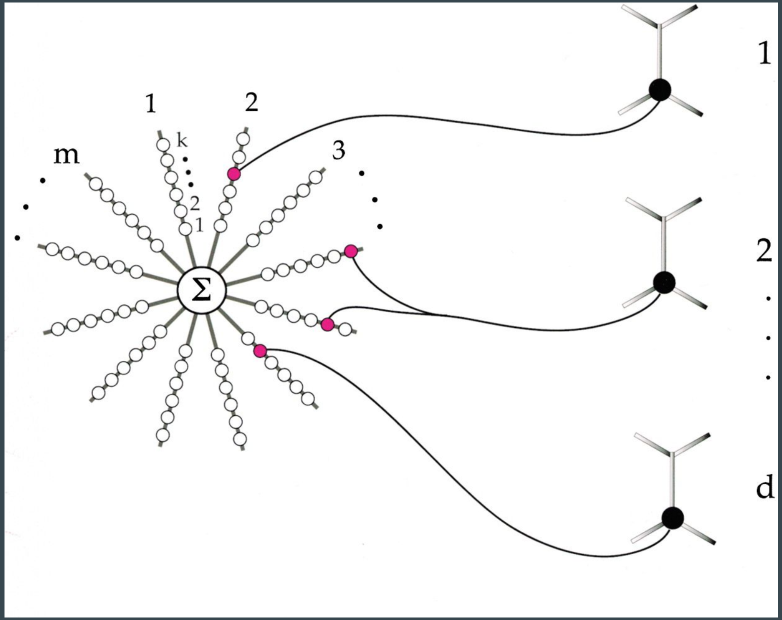

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

m: # branches, k: # synapses per branch

s = m · k: total number of synapses

d: dendritic inputs

$$a_L(\mathbf{x}) = \sum_{j=1}^{m} \sum_{i=1}^{k} w_{ij} x_{ij}$$

$$a_N(\mathbf{x}) = \sum_{j=1}^{m} b\left(\sum_{i=1}^{k} w_{ij} x_{ij}\right).$$

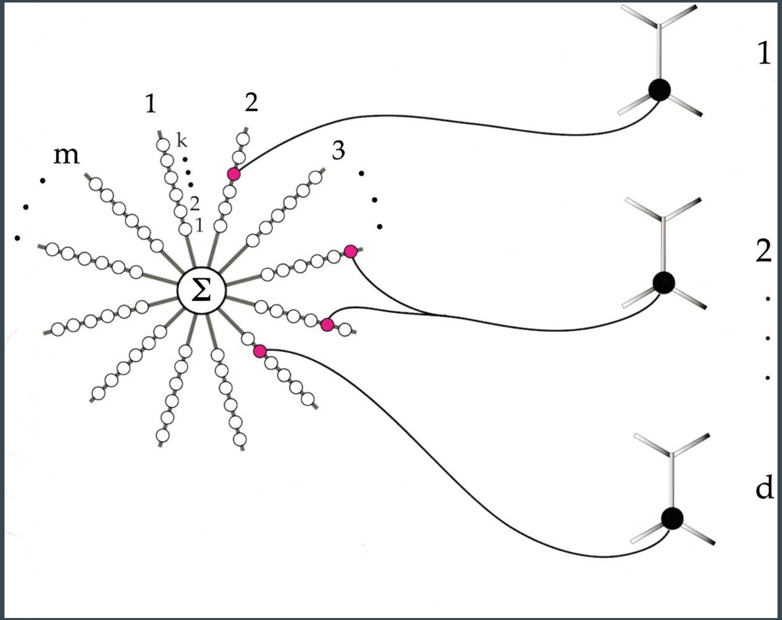

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

An upper bound can be computed based on the number of distinct memory fields expressible by the cell drawing $s = m \cdot k$ synaptic contacts, from d distinct input lines. Effectively that is the number of possible combinations that the cell receives from d inputs.

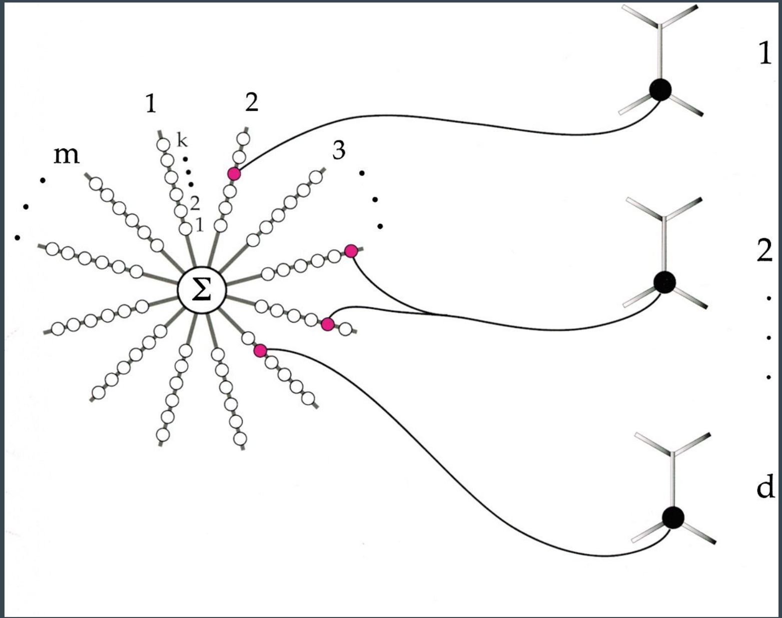

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

Linear integration:

$$B_L = 2\log_2 \binom{s+d-1}{s}$$

The combinatorial term in B_L gives the number of ways of assigning s synaptic sites to the d afferents, where only the number of contacts formed by each afferent is counted regardless of location on the cell.

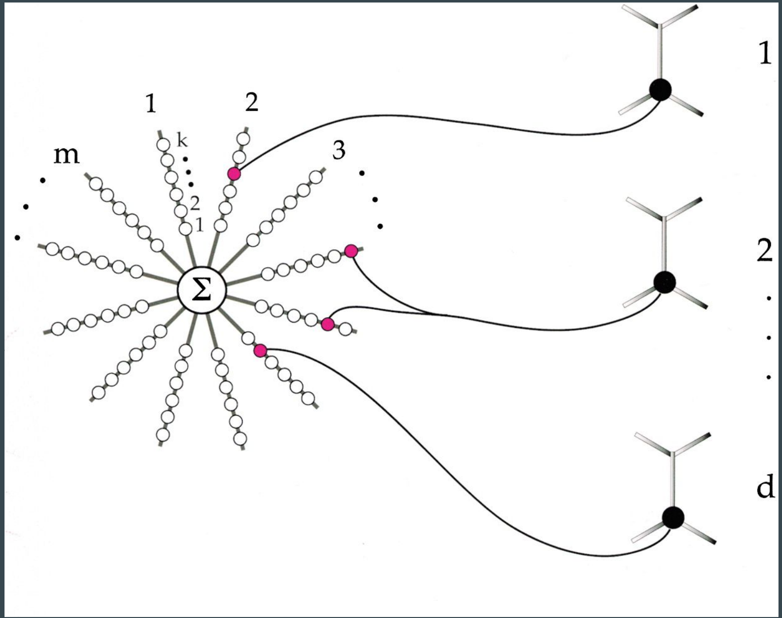

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

Non Linear integration:

$$B_N = 2\log_2 \begin{pmatrix} \begin{pmatrix} k+d-1 \\ k \end{pmatrix} + m - 1 \\ m \end{pmatrix}$$

The expression for BN was derived by applying the combinatorial expression in BL in two stages: (1) to calculate the number of distinct branch functions f expressible by drawing k synapses from d input lines with replacement, then (2) to calculate the number of distinct cells expressible by drawing m branches from f possible branch functions.

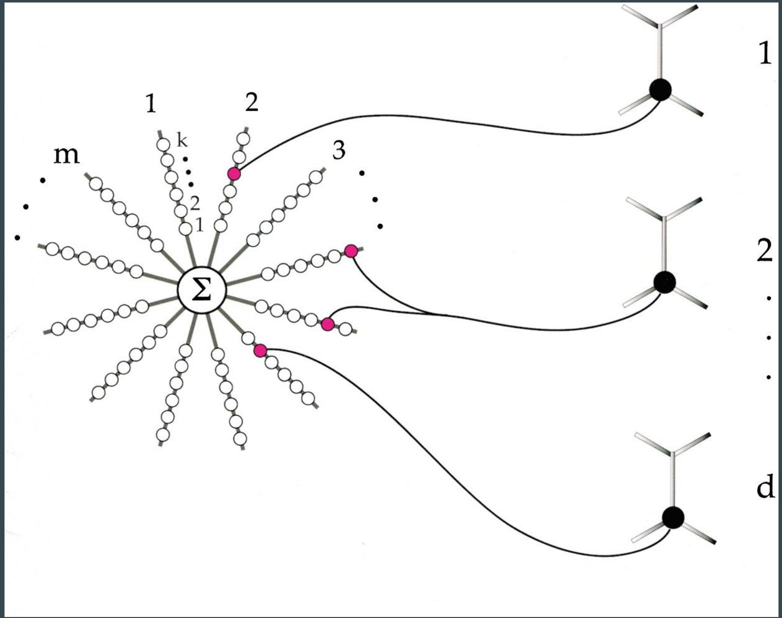

In which case would $B_L = B_N$

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

BL = BN in the special case of

a) one long branch (k = s) or
b) the cell has many branches containing only one synapse (m = s)

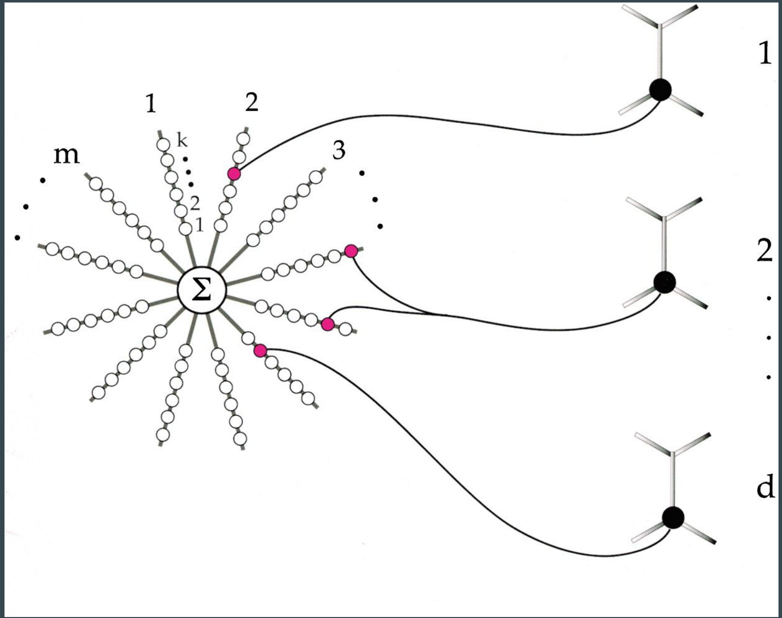

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

Simple example:

m: # branches, k: # synapses per branch

s = m · k: total number of synapses

d: dendritic inputs

For d=m=k=3: B_L = 110 for linear integration, while B_N = 220 (doubles) for non - linear integration

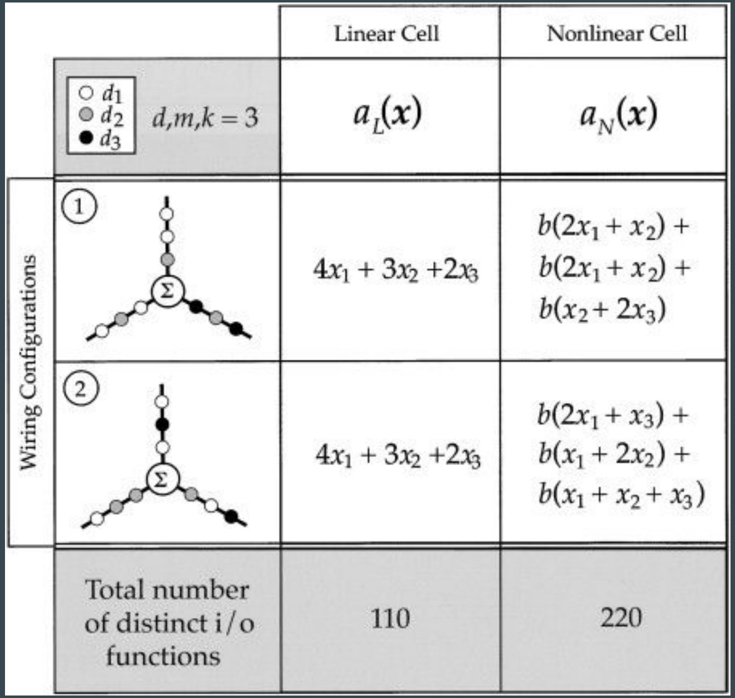

		Linear Cell	Nonlinear Cell
	d_1 d_2 d_3 $d,m,k = 3$	$a_L(\boldsymbol{x})$	$a_N(\boldsymbol{x})$
Wiring Configurations	1	$4x_1 + 3x_2 + 2x_3$	$b(2x_1 + x_2) +$ $b(2x_1 + x_2) +$ $b(x_2 + 2x_3)$
	2	$4x_1 + 3x_2 + 2x_3$	$b(2x_1 + x_3) +$ $b(x_1 + 2x_2) +$ $b(x_1 + x_2 + x_3)$
	Total number of distinct i/o functions	110	220

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

Capacity in bits for linear (lower curve) and several nonlinear cells (upper curves) with branches of different length; branch count increases from left to right as indicated schematically beneath x axis. Capacity of nonlinear model grows approximately linearly with the number of dendritic subunits

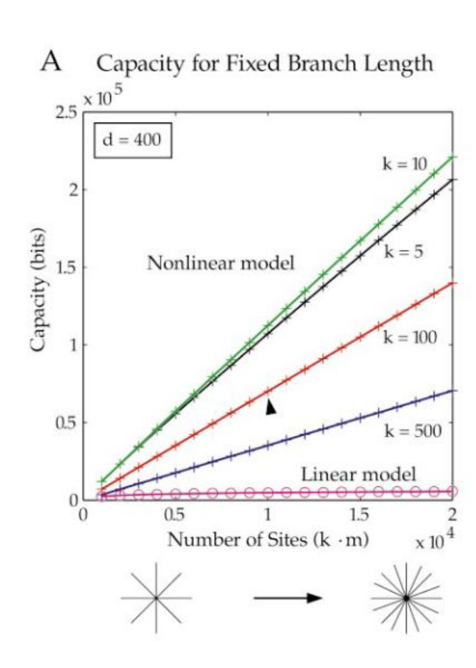

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

For fixed branch count, capacity increases monotonically as branches are lengthened. Each curve is indexed by branch count m; saturation is evident as branches become relatively few and relatively long.

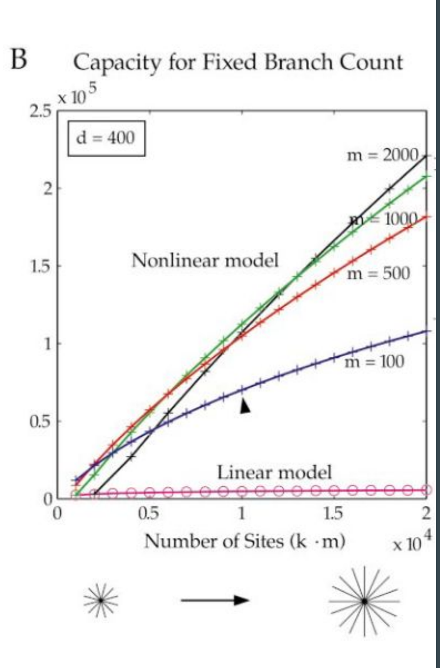

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

Capacity of a nonlinear cell with 10,000 sites for different values of d. Branch count m grows and branch size k correspondingly shrinks moving along x axis. Cells at both ends of x axis have capacity equivalent to that of linear model. Capacity of the nonlinear model is maximal for cells with 1250 branches containing eight synapses each. Asterisks indicate half-maximum capacity.

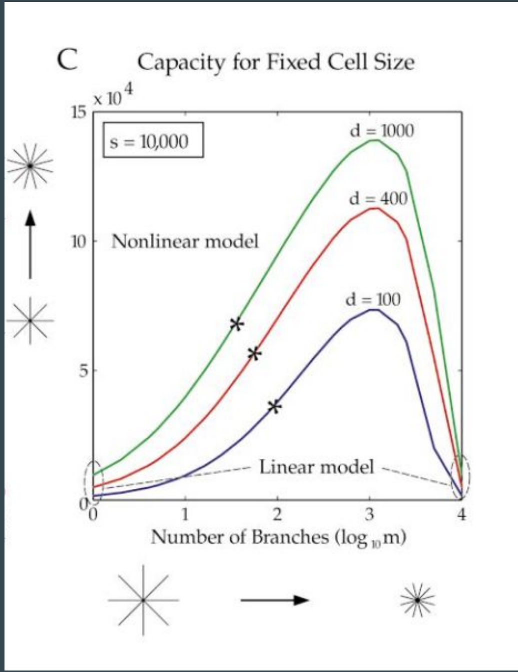

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

Dashed lines show analytical curves for linear and nonlinear cells. Solid curves show capacity measured empirically at 2% error criterion, using a subunit nonlinearity. Analytical curves were scaled down together by a factor of 3.5, to align peak analytical and empirical capacity values for the nonlinear model. Analytical and empirical curves were similar in form.

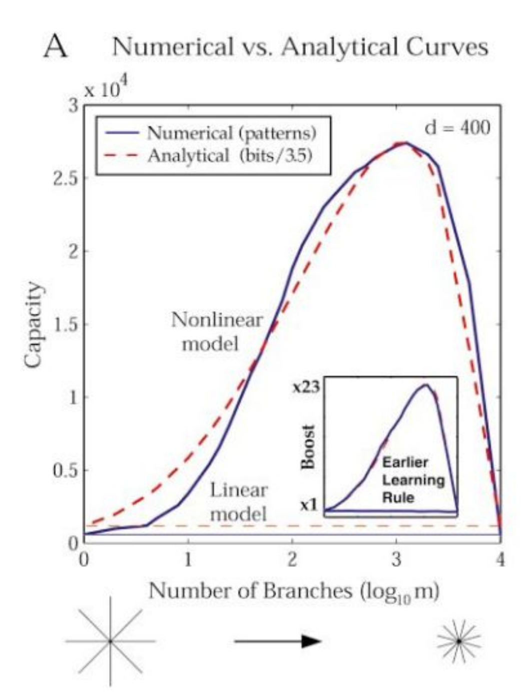

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

Cells within a population were trained independently with random initial conditions. Output of population was computed using a simple voting scheme. Positive and negative training examples were in this case drawn from two non-Gaussian distributions. A total of 30,000 training examples were drawn, half positive, half negative. Error rates are plotted for populations ranging from 1 to 50 cells. All cells were trained with m = 400, k = 25.

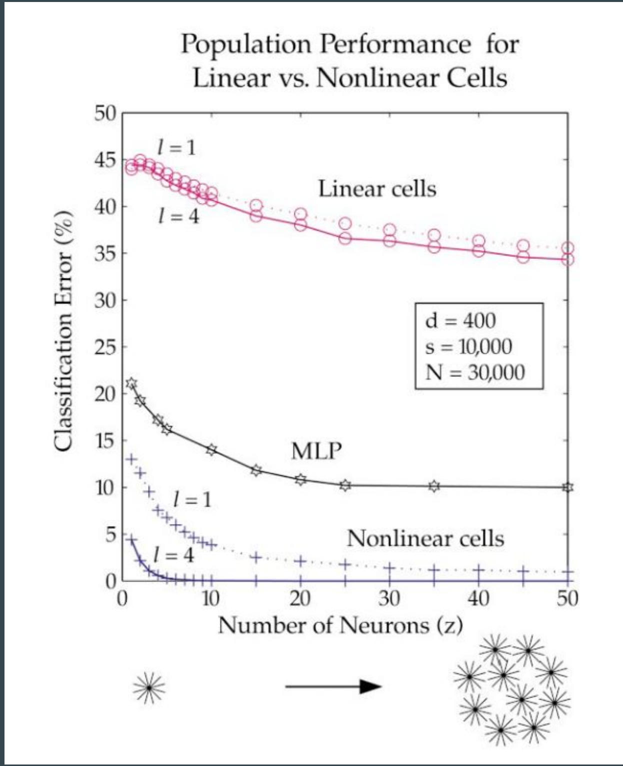

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

(1) cells with nonlinear subunits learn substantially more than cells without

(2) peak capacity occurs for subunits that are neither too small nor too large—with near-maximum capacity over a wide range of subunit sizes

(3) when subunits are of optimal size, memory capacity increases in direct proportion to the number of dendritic subunits available.

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

The main biophysical assumption underlying capacity calculations is that a neuron's integrative behavior can be captured by a simple form, which says that the neuron's output can be expressed as a sum of independent nonlinear subunit responses.

Surprisingly, however, the particular form of the subunit nonlinearity b, whether a power function, exponential, sigmoid, or other nonlinear relation, has no bearing on the function counts for nonlinear cells, since the sole role of the branch nonlinearity from the perspective of the combinatorial expression is to break the symmetry among otherwise identical branches.

Impact of Active Dendrites and Structural Plasticity on the Memory Capacity of Neural Tissue

For a cell of realistic size, storage capacity is maximized when the cell contains a large number of small subunits. For example, a cell with 10,000 synaptic contacts learns the most when it is broken into roughly 1000 independent subunits containing ten synapses each. This is not a realistic result for biological morphologies. However, the authors indicate that:

1) Around the optimal neuronal geometry, the dependence of capacity on subunit size is relatively weak, so different combinations have small impact on overall capacity

2) It is likely that their results underestimate the optimal subunit size for an individual cell because cells are described by binary weights, (all-or-none) values for the strength of a synaptic contact. In addition, cells learn at a population level not in isolation, but this effect has not been taken into account.

Questions?