

Dendritic Computation in Multi-Compartment Neurons

A Technical Strategy Brief for the
Neuro-fuzzy Soft Computing Program

R. Wells

Introduction

In accordance with our strategic plan, this upcoming year for our project might be subtitled “the year of the networks.” On the hardware side of things, I think our neuron designs are in pretty good shape, and there has been good progress in the EC research. Therefore I see our next major phase as one that involves the development of the methods of analysis, design tools, and functional definitions for designing pulse-coded neural networks.

Toward this end, I think it is important for us to have an understanding of what kinds of neuronal computation schemes are available to us from the perspective of neurobiology. Up to this point, I think it is fair to say that we have all been pretty much focused on what are generally termed “single-compartment” neuron schemata (e.g. the venerable integrate-and-fire neuron model). Biological neurons are not so simple as this, and the computational flexibility they exhibit may prove useful to us in designing neurocomputers. The purpose of this brief is to discuss multi-compartment neuron models operating non-adaptively. Structure-modifying adaptations will be addressed in a different brief.

There is a large body of experimental evidence that dendrite structures in neurons should be viewed as simple computing elements in their own right. We will refer to these structures here as “dendritic compartments” or DCs. While there are many types of small interneurons where the single-compartment model seems adequate and appropriate, there are also a great many other types of neurons that have rather extended dendritic “arbors” (trees). Dendritic branches in many (not all) neurons also tend to sprout small protrusions – the dendritic spines – and the larger fraction of all synaptic connections to a neuron occur out on these spines. One spine can support multiple synaptic connections, although it is more typical for a simple spine to support only one synapse. A dendritic shaft typically will have a great many spines. Figure 1 (below) is a photograph of a neuron showing the spines in its dendritic arbor.

Dendrite compartments that are physically near the cell body (soma) are called “proximal dendrites” or, in our terminology, PDCs. Dendrite compartments located farther away from the soma are called “distal dendrites” or DDCs. The physical location of a synapse within the dendritic arbor is an important factor in the signal processing carried out in biological neurons. Figure 2 is a cartoon sketch of multi-compartmental arrangement in a neuron’s dendritic arbor.

Compartment Organization and Responses

Although every “rule” in neuroscience has numerous exceptions, within the cerebral cortex about 79% of all excitatory synapses occur on dendritic spines. The rest occur on the dendrite’s shaft. In addition, about 31% of all inhibitory synapses occur on the spines (with the rest occurring on the cell body). A spine never has only inhibitory synapses; if a spine has an inhibitory synapse then it also has excitatory synapses. About 15% of all spines carry both excitatory and inhibitory synapses; the remaining 85% have only excitatory synapses. As a general rule, inhibitory synapses tend to occur on proximal dendrites, while excitatory synapses tend to occur in the distal dendrites.

Excitatory post synaptic potentials (EPSPs) in spines are typically small in amplitude (less than 20 mV of depolarizing change in the spine's membrane potential). Owing to the small size of a spine, the time constants accompanying these EPSPs are very fast, typically only a fraction of the duration of a neuron's action potential (AP). Multiple excitatory inputs in a DC “integrate” to produce the total EPSP, but this integration is nonlinear. If e_1 and e_2 represent the EPSPs that would result individually from two synaptic inputs, the “sum total” of the response to both inputs is of the form

$$e_T = e_1 \circ e_2 < e_1 + e_2 \quad (1).$$

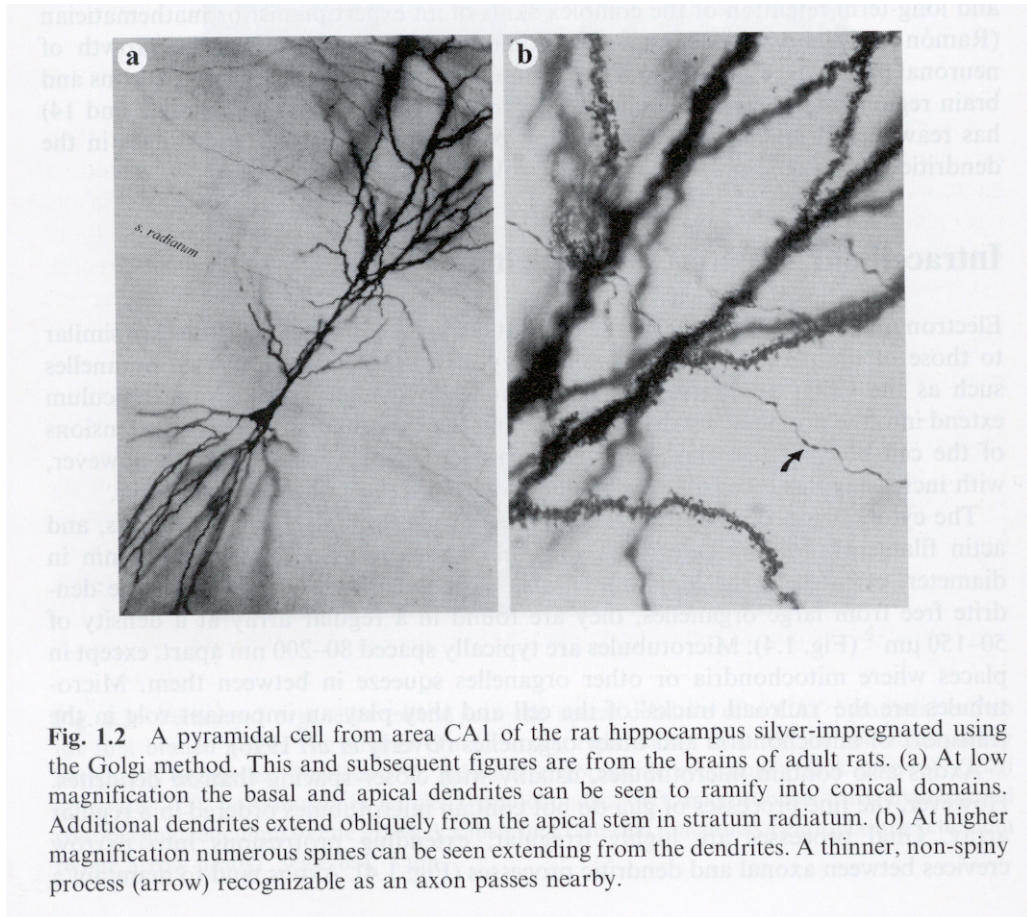


Figure 1 (taken from *Dendrites* by Stuart et al.). The spines are the little knobby things that look like fuzz on the dendrites in figure b.

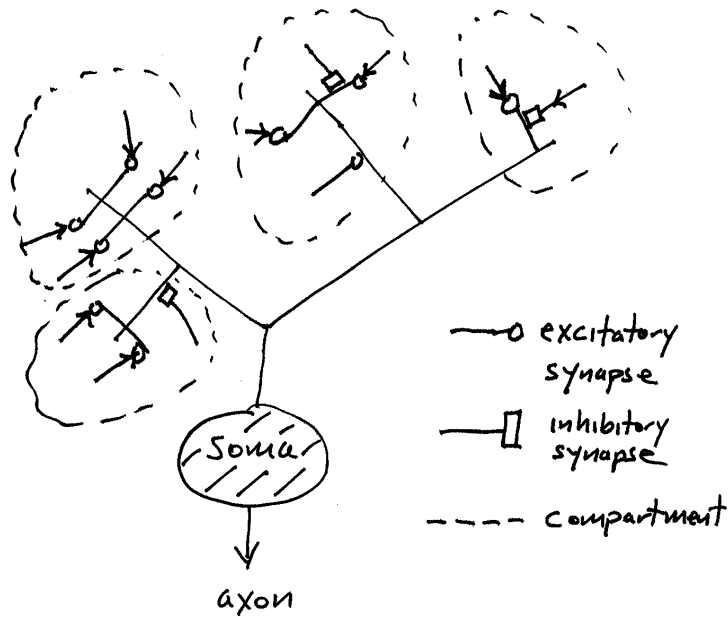


Figure 2: Sketch of a compartmental structure in the dendritic arbor

For some DCs a single AP at one synapse suffices to produce a response at the neuron's cell body at the trigger zone (TZ) of the postsynaptic neuron; for other DCs, multiple synaptic APs are needed to provoke a response at the TZ. In the first case, we can regard the spine as performing a logic OR function on its synaptic inputs; in the second case, the spine can be viewed as performing an AND function (if all synapses must be active to provoke a response) or a McCulloch-Pitts-like threshold function

$$a = \begin{cases} 0, & k < \theta \\ 1, & k \geq \theta \end{cases} \quad (2)$$

where a is the DC's output bit, k is the number of synapses receiving excitatory APs, and θ is the McCulloch-Pitts activation threshold.

Inhibitory DC inputs exert a "veto" function on "upstream" excitatory synapses (synapses farther away from the soma than the inhibitory synapse) but *have no effect* on "downstream" synaptic inputs (inputs closer to the soma than the inhibitory input). If a is an "upstream" activity given by (2) and b is a downstream inhibitory synaptic input, this relationship can be regarded as forming the logic function

$$c = a \wedge \bar{b} \quad (3).$$

We can, of course, also extend this binary type of response to an "analog logic" response in the a and b signals by appropriate definition of nonlinear "summing" and "blocking" operators. For example, if a in (3) is replaced by e_T in (1) then we might express the "analog" form of (3) as

$$c = (e_1 \circ e_2) \otimes \bar{b} \quad (4)$$

where the blocking operator \otimes is such that an increasing value of b leads to a decreasing value of c . The specific form of the mathematical operators we invent, such as \circ and \otimes , can in principle be defined in whatever way seems appropriate for our purposes so long as these operations are not at open odds with biological phenomena. For example, in (4) the “veto” exercised by the inhibitory signal b is such that, owing to the nonlinearity of the “summing” operator, we probably should not assume that \otimes distributes over \circ . In other cases, involving three or more terms in the “summing” operator, it seems likely to be the case that we should not assume associativity necessarily applies to the \circ operator.

Dendritic Integration in the Dendritic Arbor: Simple Model

Other than in a few exceptional cases¹, the dendrites of most neurons are incapable of propagating action potentials. The attenuation of EPSPs in the normal neuron is a strong function of distance. Complicated “cable models” of passive dendritic signal propagation have been proposed, but basically the attenuation over distance d goes roughly as

$$\exp\left[-(d/\lambda)^2\right]$$

where λ is the “space constant” for the dendrite. This means that proximal dendritic synapses (synapses close to the soma) have a much greater effect on the soma membrane voltage than do synapses on distal dendrites (synapses far away from the soma). Figure 3 illustrates this attenuation effect.

It is generally agreed that the severe attenuation of signals from distal dendrites is basically a “voltage divider” effect of the series resistance of the dendritic shaft. The attenuation is accompanied by significant time dispersion due to the distributed capacitance of the dendrite. It is also generally agreed that most of the excitatory current generated at the distal synapse does invade the soma. However, this current is going to be a small current for the same reason that the voltage attenuation is so severe. Nonetheless, everyone also agrees that synaptic signaling at DDCs does have some kind of important effect on the neuron’s ability to fire action potentials. The question is: how could it given the small amplitude and low current it produces?

¹ Exceptions include the pyramidal CA1 cells in the hippocampus, mitral cells in the olfactory bulb, some pyramidal neurons in the neocortex, some spinal motor neurons, and dopaminergic and GABAergic neurons in the substantia nigra (part of the basal ganglia). These cells do backpropagate APs from the TZ into the dendrites. The amplitude of the backpropagating AP in these cells diminish with distance from the soma but still remain well above the level expected for passive spread.

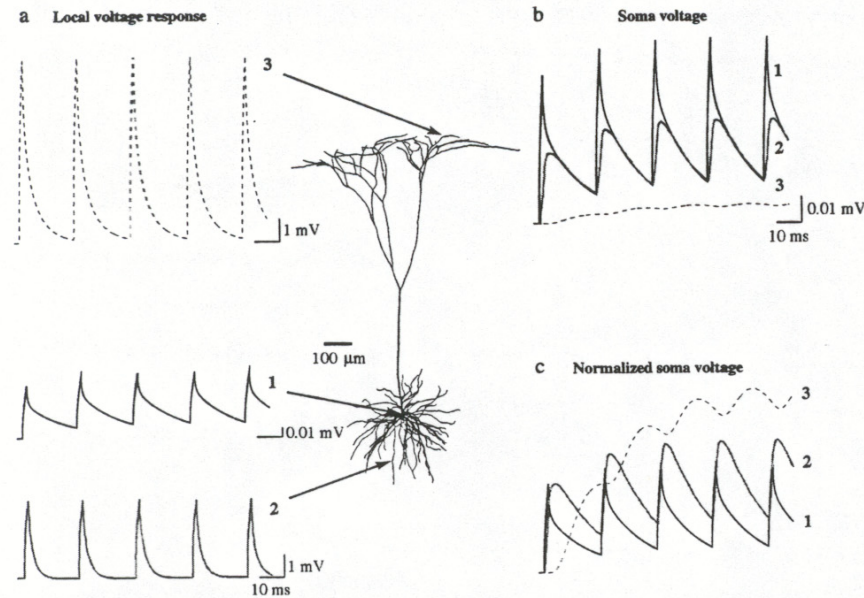


Fig. 9.2 Temporal summation in the dendritic tree depends on input location. Five consecutive current pulses (at 50 Hz) were injected at locations in the dendritic tree of a layer V pyramidal cell depicted at the middle (reconstructed by J. C. Anderson). Point 3 is in the distal apical tuft, point 2 is in a basal arbor and point 1 is at the soma. The local voltage response at the respective input sites is shown in (a). Note that at the two dendritic input locations (2 and 3) there is almost no temporal summation of potential, owing to the small 'effective local time constant' which results from the large 'current sink' imposed on these sites by the dendritic tree. The amplitude of the local voltage response is, however, relatively large because of the large input impedance in these regions. The voltage recorded at the soma in response to these same current injections are shown in (b). Note the very large voltage attenuation of the signal from the dendrites to the soma (compare trace 3 in (a) with that in (b)). In (c), the traces in (b) are normalized to the peak of the first EPSP. This normalization enables one to compare the amounts of temporal summations at the soma as a function of the input location. The summation is larger for the distal apical input. This is because the EPSP arriving from this site becomes significantly broader because of the capacitance that it encounters when spreading towards the soma and consequently, successive EPSPs can summate better with each other. (Model parameters: $R_i = 200 \Omega \text{ cm}$, $R_m = 20,000 \Omega \text{ cm}^2$ and $C_m = 1 \mu\text{F cm}^{-2}$. Current pulse amplitude is 0.01 nA and 2 ms duration).

Figure 3. Model of voltage attenuation for proximal and distal dendritic synapses. This figure is taken from *Dendrites* by Stuart et al.

I have seen no satisfactory answer to this question, although admittedly my literature search on this question falls rather short of being exhaustive. Nonetheless, I think there are likely to be at least two possible mechanisms that could support a significant signal processing role for distal dendritic synapses. The first is ionotropic, and this is the one we'll discuss in this brief. The other is metabotropic, and I plan to discuss this one in a later brief.

In the ionotropic case, it is known that many DCs do contain Na^+ or Ca^{2+} voltage-gated calcium channels (which open when the membrane sufficiently depolarizes and which then conduct excitatory currents into the spine). The density of these channels is not large enough to support AP generation (except in the cases of a few types of neurons; even in these cases the dendrite is typically not able to support propagation of the dendritic AP to the soma). However, the opening of voltage-gated channels (VGCs) in the spine, if it occurs, does increase the EPSP and leads to greater excitation of the cell's soma.




Figure 4: Illustration of effect of EPSP of a DDC on nearby DCs. APs arriving at DDC A produce an EPSP. The slight depolarization of the membrane potential at A spreads to nearby DDC B, thus raising the membrane voltage there. This increase in membrane potential makes it easier for APs arriving at B to open excitatory VGCs in DDC B.

Now, although the EPSP of a DDC is not large enough to significantly affect the soma, it may well be large enough to slightly depolarize the membrane potential of nearby DCs. In this case, if those nearby DCs are also receiving excitatory AP inputs, this incremental increase in membrane potential added to the EPSP that these APs would produce might be enough to open the second DDC's VGCs. The end result is that DDC B (figure 4) would be more excitable owing to synaptic activity at DDC A (again, figure 4), and this increased excitability could then lead to a greater synaptic response at B than would be the case if spine A had not received any excitatory input signals.

We can regard this hypothesis as the basis of a “domino effect” along the dendrite. As distal synapses are excited, they increase the excitability of their downstream neighbors. If these neighbors are also receiving excitatory inputs, they produce even more excitability in DCs further “downstream”, etc. This action is basically the equivalent to DDC B having its “gain” (synaptic weight) modulated by DDC A.

This hypothesis is consistent from what we can expect from the Hodgkin-Huxley dynamics of a multi-compartment neuron model. While I am not aware of any technical papers that support this hypothesis, I am also not aware of any that refute it. Consequently, I think we can regard this hypothesis as a legitimate working premise for our neural network designs. Hardware implementation of this model is relatively straightforward and is low-cost in terms of chip real estate. There are, as before, both analog and digital implementations for mimicking this signal processing scheme.

Let A_1 and A_2 be DDCs (A_2 downstream from A_1). Let B be a DC proximal to the soma. Let us further assume that the effect of the distal DCs is to increase the synaptic weight of the cell's response to B. In the BAN design, synaptic weight is set by a current source whose value is determined by a current mirror arrangement. Figure 5 illustrates this scheme. The PDC, B, controls a switch in the BAN that switches weight current I_{syn} (cf. figure 5) into the BAN's excitatory summing resistor (ESR). For a single DDC, A, we can model the increase in the effect of PDC B on the BAN's soma by sending a binary-valued output from A to the WSR depicted in Figure 5. This arrangement is illustrated in Figure 6.

Extension of this idea to a series of DDCs is straightforward. Figure 7 illustrates one possible approach for the case of two distal “A” DDCs converging on one proximal “B” DC (not shown in the figure). In this simple scheme, the overall strength of the weight enhancement due to DDCs is represented by the number of bits transmitted and the bit position (to allow for different WSR weightings of different DC signals).

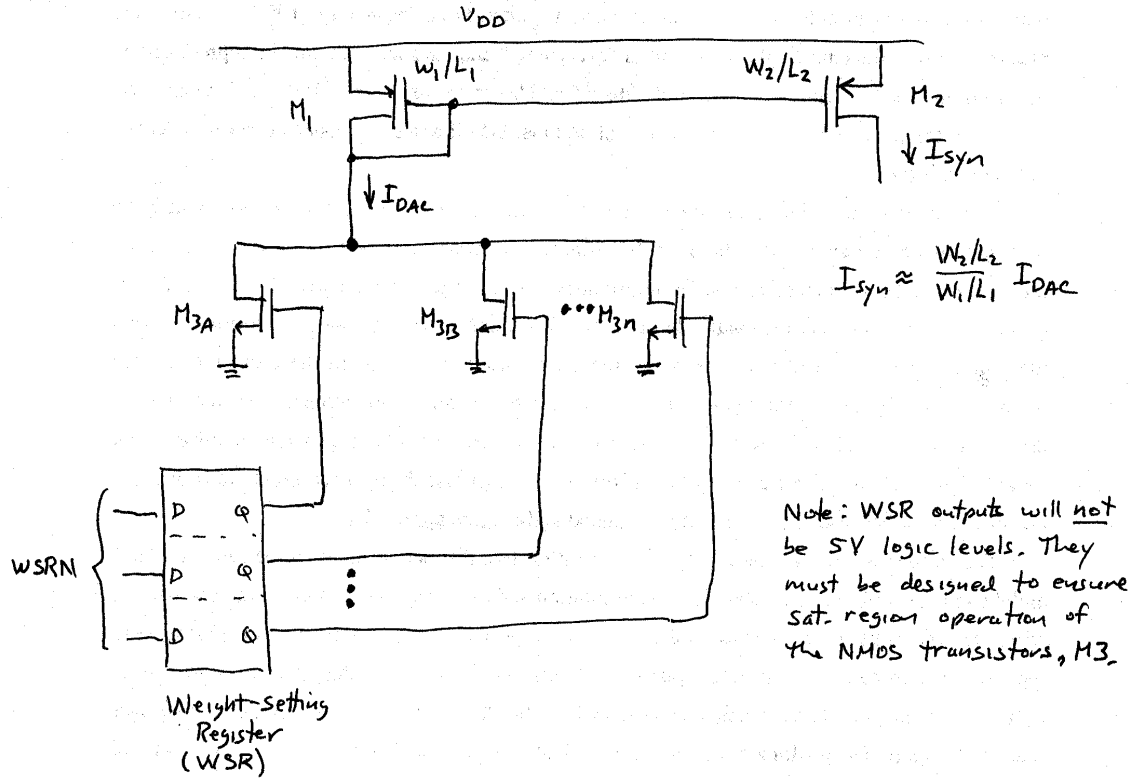


Figure 5: Synaptic weight setting via a current mirror. Transistor M2 is the current source for the weight. M1 is the mirror transistor. The current in M1 is determined by transistors M3; the more M3 transistors that are turned ON via the weight-setting register, the greater the weight I_{syn} .

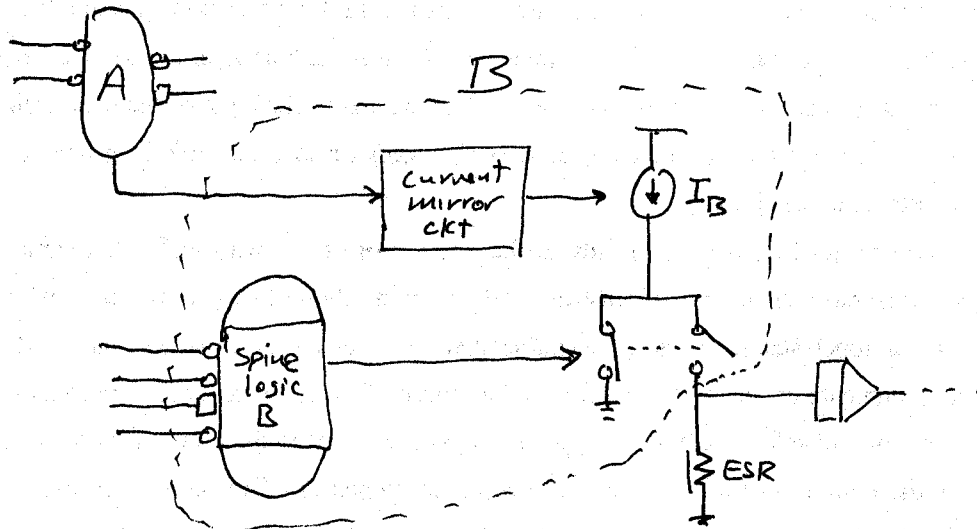


Figure 6: Enhancement of the effect of proximal spine B by distal spine A. The current mirror input shown here corresponds to one of the bits of the WSR in figure 5. Spine logic B carries out its own local “logic” preprocessing on its synaptic inputs and produces a McCulloch-Pitts-type response that controls the synaptic switch in the BAN cell body, diverting the weight current I_B into the ESR.





Figure 7: Cascading successive distal spine signals. In this figure DC A_1 is more distal to the soma than DC A_2 . The effect of A_1 propagates to proximal DC B only if DDC A_2 is active at the same time. (This is represented by the AND gate in the figure). A_2 is able to propagate its effect regardless of the activity of the more distal A_1 . A variety of dendritic “integration” effects can be mimicked by appropriate design of the cascade logic.



$$I_T = I_{A_1} \circ I_{A_2}$$

Figure 8: Analog summation of gain effects of DDCs. DDC outputs are binary-valued for both DDCs. A_1 is the more distal DC. Transistor M_1 is a transmission gate and is enabled by the output of A_2 . If A_1 is LOW and A_2 is HIGH, current flow from A_1 to summing resistor M_3 is blocked by a diode (not shown). Sizing of M_1 and diode M_2 allows independent sizing of the current contributions to I_T . The voltage across M_3 determines the amount of incremental mirror current I_M produced by M_4 . M_5 sets the minimum mirror current in current mirror device M_6 .

The basic scheme can also be extended to analog signaling. This is represented in Figure 8 above for the case of two distal DCs. In this arrangement, conversion of the digital signals from the distal DCs into independently-valued weight control currents is achieved through the design of the gate transistor M_1 , the diode-connected transistor M_2 , the current summing resistor M_3 , and the current source M_4 . Again, this scheme can be extended to multiple distal DC signals.

Dendritic Integration in the Dendritic Arbor: McCulloch-Pitts Model

The dendritic integration scheme discussed above is a very simple approximation to what one could expect from Hodgkin-Huxley dynamics out in the arbor. It may indeed be a bit too simple. The capabilities of our neural networks can be potentially be enhanced, and the number of degrees of freedom for the EC algorithms increased, if we adopt a somewhat more complex integration scheme.

In the original McCulloch-Pitts model with $k \leq K$ active excitatory inputs and $l \leq L$ inhibitory inputs the output was given by the expression

$$a = \begin{cases} 0, & k - l < \theta \\ 1, & k - l \geq \theta \end{cases} \tag{5}$$

This expression is a generalization of (2). However, in the arrangement of dendritic compartments any downstream inhibitory input vetoes all upstream excitation but has no effect on anything happening downstream in the arbor. This implies that we can specialize the McCulloch-Pitts model using the form shown below in Figure 9.

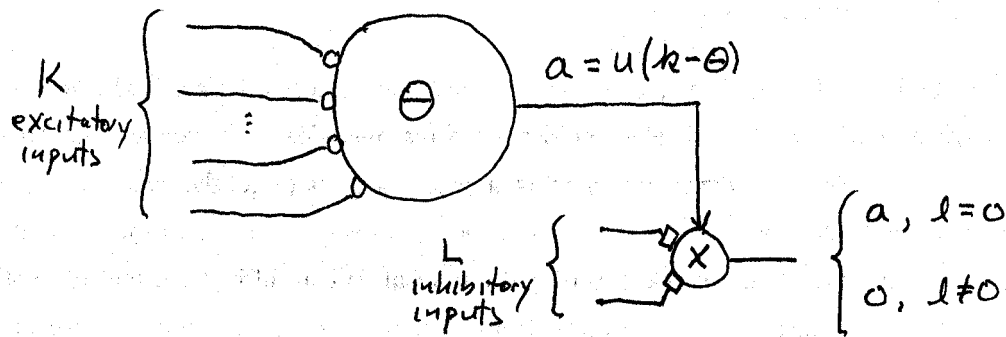


Figure 9: McCulloch-Pitts model of a dendritic compartment

Here the activity a is given by (2). The blocking operator \otimes is defined such that any active inhibitory input cancels activity a and produces a zero output.

The a calculation can be implemented using one of James' sorting networks. The scheme is illustrated in Figure 10.




Figure 10: Sorting network implementation of McCulloch-Pitts a function.

This basic compartment core shown in Figure 10 has a fixed threshold θ . In order to realize the integration scheme discussed in the previous section, this threshold must be variable. Increased DC excitability implies reduction of θ by one or more discrete steps, subject to the constraint that θ can never be made less than one. This is easy to do with a minor change to the circuit in Figure 10. The idea is illustrated in Figure 11 for the case of a unit decrement in θ . Generalization to any number of decrements $\Delta < \theta$ is accomplished by the addition of more OR functions in the output paths of the sorting network. (I assume here that the sorting network sorts all “1” inputs to the “top” of the sorting block).

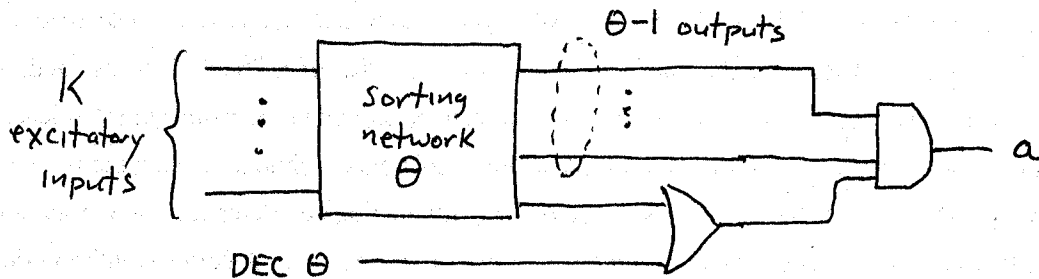


Figure 11: Variable threshold McCulloch-Pitts a function. DEC θ is a signal coming in from an upstream DC and indicates an increase in the excitability of this dendritic compartment.

In order to incorporate any inhibitory inputs this DC may have, signal a is combined with the blocking operator as illustrated in Figure 12.




Figure 12: McCulloch-Pitts dendritic compartment. The DC’s output signal is sent to the next DDC (or DDCs) downstream or is combined with a PDC unit to change the input weight of the BAN at the soma.

I propose that we standardize on a schematic symbol for the DC structure in Figure 12. My suggestion is shown below in Figure 13.




Figure 13: Wells' proposed schematic symbol for a dendritic compartment.

Time-Varying Dendritic Structure

In this tech brief I have confined my discussion to the case of a fixed (non-adaptive) dendritic arbor. In real neurons, the dendritic arbor is the site of a number of different modulations and longer-term changes. While the BAN soma design incorporates (or will incorporate) a number of short-term signal modulations, dendritic synapses are subject to changes in their efficacy to input signals from short-term facilitation, short-term potentiation, short-term depression, and long-term potentiation/long-term depression.

I am still sorting through my research materials to come up with a good general description of each of these important effects and their putative causes. What is perhaps most important for now is simply to make it known to all team members that in our design of networks *the properties of the DCs and of the soma (BAN circuit) need not (and probably should not) be regarded as having to assume one fixed set of parameters* or even a fixed number of synaptic inputs. I hope to come up with some “adaptation” and some “learning” rules for us to use in the near future. Until then, however, I think that if the EC algorithms designing our pulse-coded neural networks come up with adaptations “on their own” that’s going to be okay. Putting it another way, the EC designs need not confine themselves to producing merely “fixed” DCs. It is desirable for us to also have “mapping rules” for adaptively altering the structures *in response to signal activity*. While this is, I know, sort of vaguely worded here, I do hope to be able to come up with a brief in the near future that can provide some ideas or a general sense of direction for what these mapping/ modulation functions could look like.