NEUROSCIENCE

PERSPECTIVE

M. E. Larkum / Neuroscience 489 (2022) 4–14

Are Dendrites Conceptually Useful?

Matthew E. Larkum [a,b]*

[a] *Institute for Biology, Humboldt University of Berlin, D-10117 Berlin, Germany*

[b] *NeuroCure Cluster of Excellence, Charité – Universitätsmedizin Berlin, D-10117 Berlin, Germany*

Abstract—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. I argue here 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.

This article is part of a Special Issue entitled: SI: Dendritic contributions to biological and artificial computations.

Key words: levels of understanding, point neurons, integrate-and-fire, networks, NMDA spikes, intrinsic excitability.

This article discusses the *conceptual usefulness* of dendrites. That is, the extent to which they are necessary or useful for an abstract description of brain function and neural networks in general. A commonly used abstraction in neurobiology and computational neuroscience is the point neuron that has no dendrites. The operation of a point neuron is usually described as ''integrate-and-fire", where linearly summated weighted synaptic inputs determine, via an activation function, if the neuron should be 'on' or 'off'. It could be argued that this conception of a neuron collapses the role of dendrites in to the strength of synaptic inputs thus relegating their function to a mere after-thought, and therefore not conceptually useful. In fact, there is a long history of disregarding dendrites (De Schutter, 2008) and the relatively brief renaissance of interest in dendrites (Antic et al., 2010) has more recently been challenged by the incredible success of deep neural networks with point neurons. Indeed, neuroscience now turns to deep neural networks for insights into brain function (Richards et al., 2019), suggesting that neural networks might already use the optimal abstraction of a neuron. The burden of proof has once again shifted to the biologists to demonstrate that the point-neuron abstraction is conceptually lacking.

The artificial neural network comes historically from an attempt to encapsulate the essential features of brains as parallel distributed processors (Rumelhart and McClelland, 1986). The choice to use a minimal description of a neuron was acknowledged to be an oversimplification from the beginning (Durbin and Rumelhart, 1989). Since that time, a great deal has been discovered about dendrites and single-neuron computation (Koch and Segev, 2000; Poirazi and Papoutsi, 2020), and it not clear which of the new discoveries are conceptually necessary in order to capture the essential features of biological neurons in artificial networks. However, it is important to distinguish conceptual necessity from conceptual usefulness. For instance, 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. However, this is not the benchmark I wish to set here. 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 (but see Goldental et al., 2014). Therefore they can, in principle, be instantiated in 0 s and 1 s on a ticker tape with the appropriate finite state machine, i.e., a Turing machine (Fig. 1A). It is now unarguable, however, that the idea of neural networks (Hopfield, 1982) introduces extremely useful concepts for solving computationally challenging problems. On the other hand, as a conceptual tool for designing software, a Turing machine is nearly useless. Borrowing from Marr's three levels of understanding (Marr, 1982), the neural network can be viewed as an implementation-level conceptual improvement over a Turing machine, that allows us to focus on higher level concepts such as network architecture, learning rules and cost functions (Richards et al., 2019) (Fig. 1B). To use the neural network abstraction, we simply imagine that the computer is actually implemented with a set of connected neurons. Similarly, with respect to the conceptual usefulness of

*Address: Institute for Biology, Humboldt University of Berlin, D-10117 Berlin, Germany.
E-mail address: matthew.larkum@hu-berlin.de

https://doi.org/10.1016/j.neuroscience.2022.03.008
0306-4522/

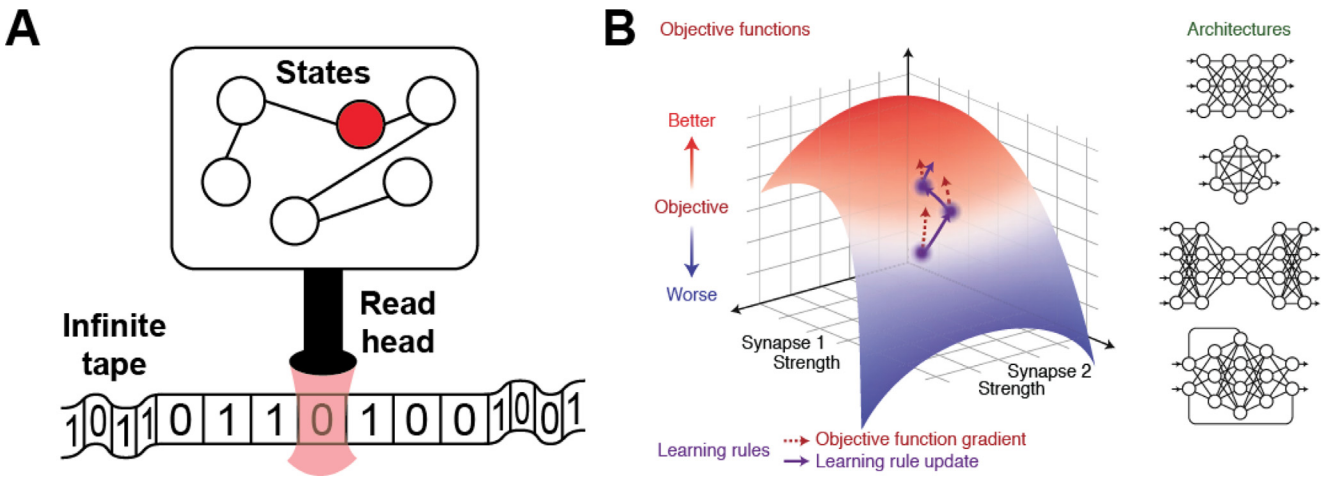

Fig. 1. Capturing the essence of computation. **(A)** The Turing machine, a ticker tape and a finite state machine. **(B)** The key components of neural network design are architectures, learning rules and objective functions according to Richards et al. (2019).

dendrites, the question is really whether any aspect of dendritic computation aids higher level concepts such as network architecture, learning rules and cost functions, or any other as yet undiscovered principles of computation. If so, this would mean that disregarding the conceptual improvement that dendritic computation offers, would hold back progress in understanding both the biological and artificial networks.

It is commonly argued that tinkering with the implementation level by including dendrites in the description of neurons will not impact the more fundamental high-level issues that should now be addressed at the network and behavioral levels (Krakauer et al., 2017). It has been pointed out, for instance, that implementation-level approaches used in neuroscience involving recordings from single neurons would be inadequate for determining what a neural network does (Hopfield and Brody, 2000). Indeed, the equivalent of single-cell recordings would also likely be inadequate for determining the program running on a digital computer (Jonas and Körding, 2016). We nevertheless believe we understand how computers work because we built them and therefore *know* how the fundamental components operate. We can go about using high-level concepts and always be sure that there is a throughline from the higher levels to the component parts. If, on the other hand, we investigated computing devices from an alien world, we couldn't assume that they use binary components like our digital computers. Recording from the brain is analogous to this problem. If neurons don't actually behave like point neurons, many of the models in neuroscience are being built on false assumptions. The correct abstraction for biological neurons is therefore fundamentally important for neuroscience. At this point in the history of neuroscience it might appear that, having 'looked inside the box', neuroscientists have apparently reached a consensus. The neurons of nearly all models of the brain and human-designed neural networks linearly count inputs and fire in an all-or-none manner. Can this nearly universal approach really be substantially wrong? The next sections give examples that suggest this is exactly where neuroscience as whole currently stands. The arguments are inevitably more technical and a reader seeking the upshot might want to skip to "Concluding remarks – Dendrites and what constitutes understanding").

METASTABLE STATES OF ACTIVITY INVOLVING EXCITATION AND INHIBITION

At first glance, it is not entirely clear whether dendrites are a bug or a feature (Häusser and Mel, 2003). Neurons typically receive huge amounts of synaptic inputs, many of which are so far from the cell body that their impact is difficult to detect (Williams and Stuart, 2002). It is generally assumed that this fact explains why neurons must receive a huge and constant barrage of inputs, rendering single inputs almost useless (London et al., 2002). However, this would seem to be a highly inefficient strategy out of place in the biological world. Other more recent explanations, involve spatial patterning that combined with active dendritic properties, leading to vast improvements in computational efficiency (Hawkins and Ahmad, 2016; Cui et al., 2016).

A common phenomenon involving barrages of synaptic input are so-called up/down-states, which has been intensely studied (Sanchez-Vives and McCormick, 2000; Hahn et al., 2006; Destexhe, 2009; McCormick, 2015; Tukker and Sachdev, 2020; Torao-Angosto et al., 2021). These are slow (~1 Hz) oscillations between two membrane potential values, usually observed under certain kinds of anesthesia, and similar to slow oscillations observed during sleep states (Wilson, 2008), that can also be induced in slice preparations (Sanchez-Vives and McCormick, 2000) and resemble high conductance states *in vivo* (Destexhe et al., 2003). One of the curiosities of the up-state is that it has a relatively fixed and stable amplitude all over the dendritic tree (Waters and Helmchen, 2004). The actual value is mostly likely determined by the precise balance of excitation and inhibition (E/I) that establishes an effective reversal potential for the combined synaptic input across the dendritic tree

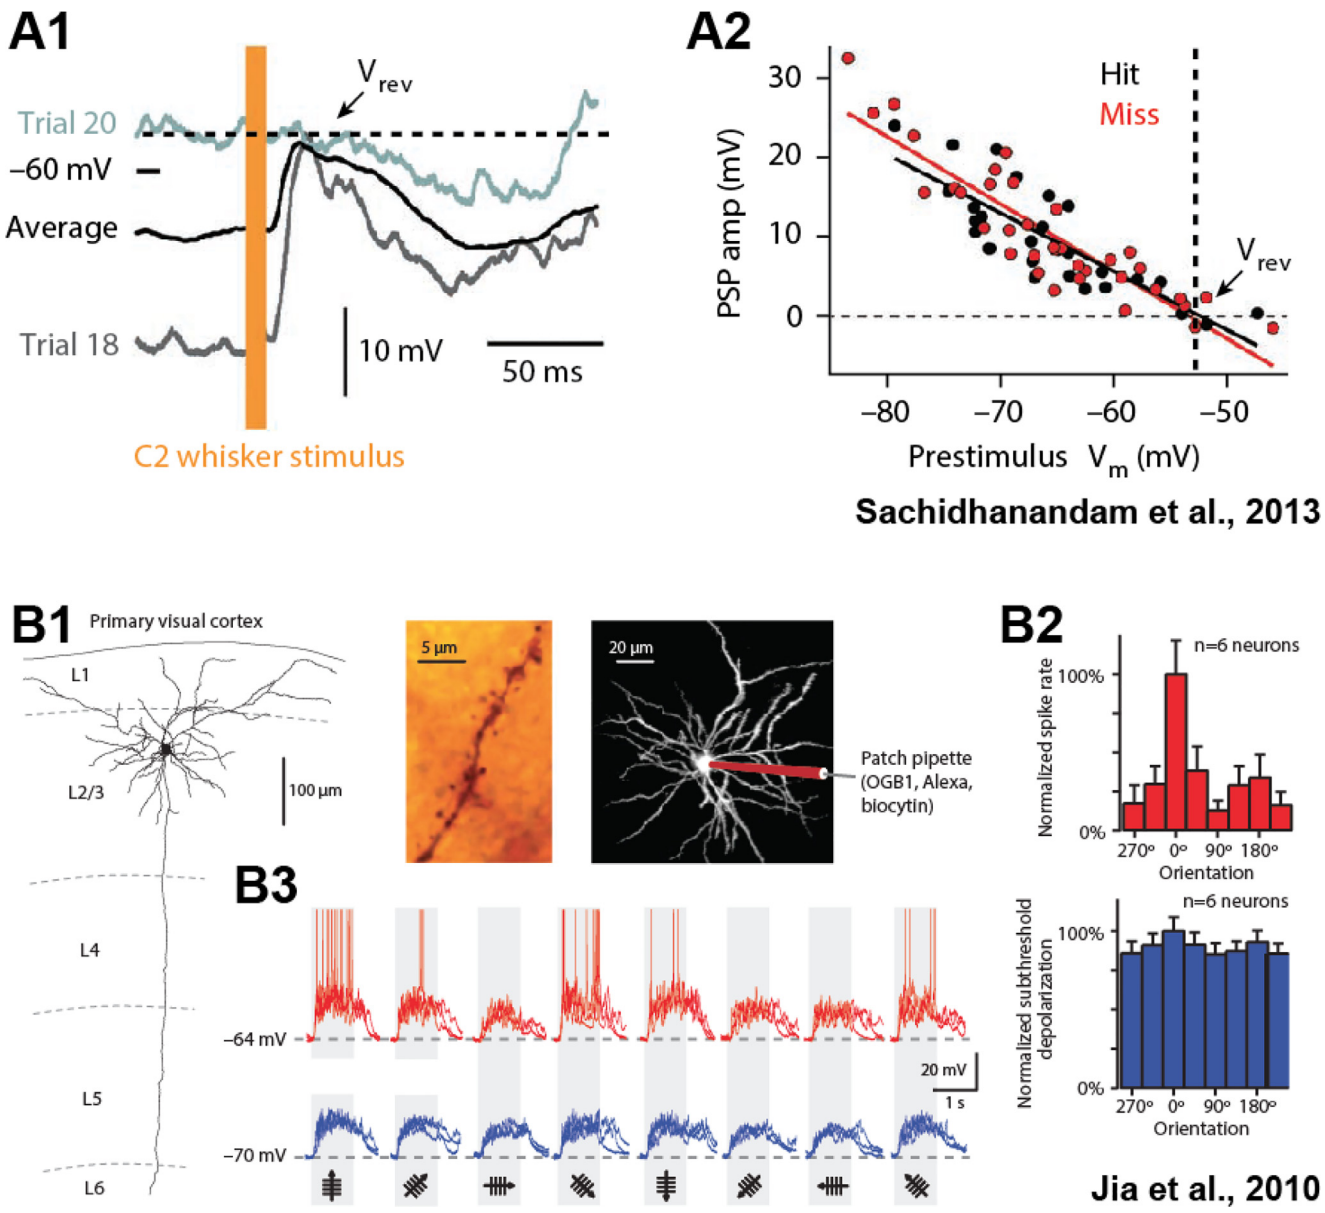

Fig. 2. The effective reversal potential of balanced E/I. **(A1)** The amplitude of a whisker stimulus in a L2/3 pyramidal neuron in mouse sensory cortex is a decreasing function of resting membrane potential (V_m). **(A2)** The "effective reversal potential" (V_{rev}) for whisker-evoked potentials is ∼−50 mV. **(B1)** Reconstructed L2/3 pyramidal neuron in the primary visual cortex of a mouse (insets showing a magnified dendritic branch with dendritic spines and a fluorescent neuron with patch clamp recording). **(B2)** Responses of the L2/3 neuron to oriented drifting gratings of different orientations at resting V_m (red) and hyperpolarized V_m (blue). **(B3)** Spikes generated at each orientation (top) versus subthreshold depolarization (bottom). Adapted from: A (Sachidhanandam et al. 2013), B (Jia et al. 2010).

(Shadlen and Newsome, 1998; Okun and Lampl, 2008; Okun et al., 2010). Balanced E/I is found under many conditions including in the awake state and can, for instance, be readily observed with stereotypical sensory stimuli. For instance, single-whisker movements in a mouse have a reversal potential at ∼−50 mV (Fig. 2A), which is frequently still below the threshold for action potentials (Sachidhanandam et al., 2013). A similar phenomenon is seen during visual stimuli with orientated driftings that cause broadly similar depolarization of the neuron for each stimulus presentation despite the fact that the preferred orientations robustly cause much greater action potential firing than others (Fig. 2B) (Jia et al., 2010).

It had been suggested that the up-state might be dependent on the activation of NMDA receptors (Antic et al., 2010), i.e. activity intrinsic to the cells, however in a series of publications from different laboratories, the situation was shown to be even more surprising. These experiments were by chance carried out under similar conditions: in vivo intracellular (patch-clamp) recordings from layer 2/3 pyramidal neurons in anesthetized rodents (Chen et al., 2013; Smith et al., 2013; Palmer et al., 2014).

None of these exciting studies examined phenomena that were directly related to up-states, per se, which perhaps explains why the following observations have gone 'under the radar'. In the course of experiments in all three studies, controls were carried out to intracellularly block NMDA receptors using MK-801 in the patch pipettes used for the recordings. This detail is important because it allowed the researchers to observe the post-synaptic contribution of NMDA receptors to the up-state without affecting network activity at large.

The findings were remarkable and consistent across all the laboratories, yet hard to interpret using an integrate-and-fire concept of synaptic integration. Firstly, all the studies observed that the amplitude of the depolarization caused by the up-state was unchanged by NMDA receptor channel block (Fig. 3, MK-801). This implies that AMPA receptors (the other main excitatory receptor type) are predominantly involved in subthreshold depolarization of the neuron. But even more remarkably, action potential generation was very much affected by NMDA receptor channel block (Fig. 3C) 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 ones that do (AMPA). This fact is hard to reconcile with a simple integrate-and-fire theory of spike generation, and argues for a more complex description of the neuron than the traditional models imply.

A NEW HYPOTHESIS FOR SPIKE GENERATION IN LAYER 2/3 PYRAMIDAL NEURONS

Two of the studies showed that NMDA-dependent dendritic spikes were generated during up-states (Smith et al., 2013; Palmer et al., 2014), which lays the groundwork for a new conceptual hypothesis for how these, and possibly other neurons, operate. The consensus view is that the binary nature of up/down-states, with two fixed and stable potentials, is best explained by a slow oscillation between high and low network activity (Sanchez-Vives and McCormick, 2000). The fixed up-state value, it is here hypothesized, results from a relatively uniform balance of excitatory (AMPA) and inhibitory (GABA) receptors such that the effective reversal potential is always the same. This would lead to both the stable and stereotypic amplitude over long durations, and yet be consistent with the fact that intracellular blockade of NMDA receptors has little or no effect. It would also imply that regardless of the exact level of network activity, the only important parameter that needs to be constant is the balance of excitation versus inhibition as is frequently observed (Okun and Lampl, 2008).

But why does the blockade of NMDA receptors have such a dramatic effect on the generation of action potentials? The necessity for high network activity for up-states entails the bombardment of the dendritic tree with synaptic inputs. Because the amplitude is set by the effective reversal potential, this means that, no matter how high the activity, balanced excitatory and inhibitory input effectively clamps the neuron at a fixed depolarization. Indeed, the higher the activity, the *more* the neuron is clamped to this potential, counterintuitively reducing the impact of random fluctuations that might otherwise generate action potentials. It has been shown, for instance, that sensory stimuli are less likely to cause action potentials during an up-state in layer 2/3 pyramidal neurons *in vivo* (Petersen et al., 2003).

At the same time, as predicted by the work of Mel, Schiller and others (see Box 1), the increase in synaptic bombardment also implies an increase in the density

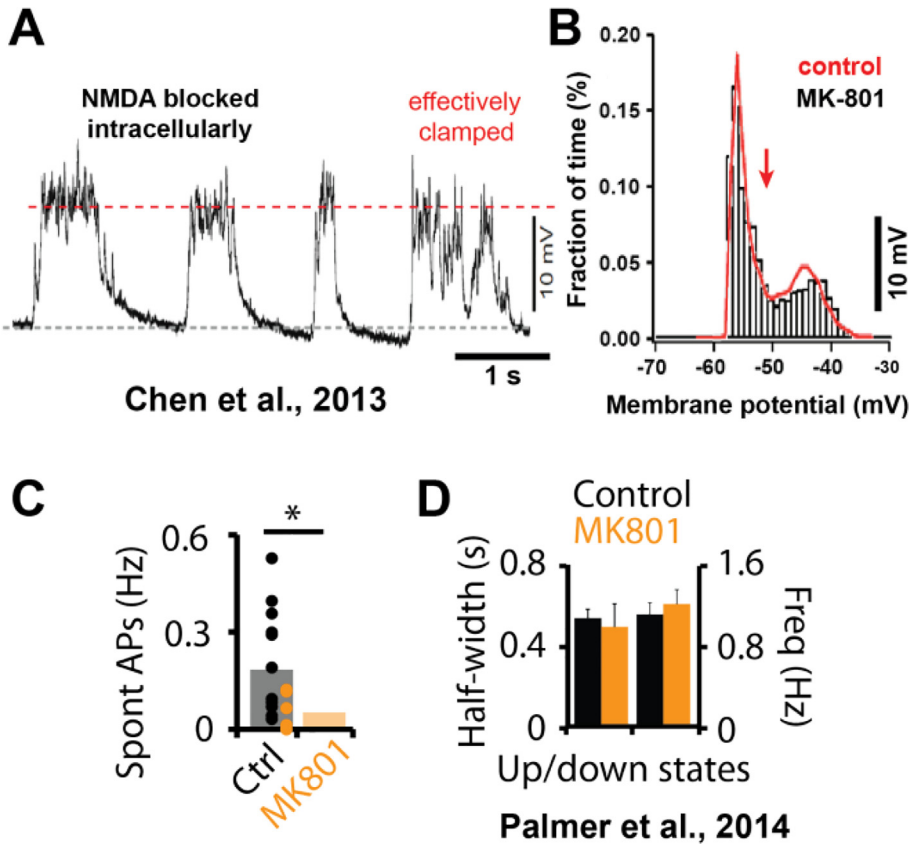

Fig. 3. NMDA channels during high activity. **(A)** Up/down states without NMDA receptors blocked with MK-801. **(B)** The same two peaks (the up-state V_m/down-state V_m) with and without NMDA receptors blocked. **(C)** Spontaneous action potentials (APs) during up-states with (black) and without (orange) NMDA receptors blocked. **(D)** Half-width and frequency of up/down states unaffected by block of NMDA receptors. A, B, adapted from (Chen et al., 2013), C, D, adapted (Palmer et al., 2014).

of glutamate-bound NMDA receptors and the increased chance of local clusters of synaptic inputs on thin dendritic branches. As few as ~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 s of mV depolarization at the cell body (Palmer et al., 2014) that can easily exceed threshold when starting from the up-state depolarization level. Unlike the up-state that is made up of balanced excitation and inhibition spread uniformly over the dendritic tree (Liu, 2004; Waters and Helmchen, 2004), spatial clustering of small groups of excitatory inputs on random dendritic branches are to be expected at high densities of inputs (Mel and Schiller, 2004) without the equal and opposite effect of clustered inhibitory inputs. These local islands of excitatory input are therefore expected to be largely unaffected by the net balance of excitatory and inhibitory inputs over the dendritic tree as a whole and can therefore operate according to their specific reversal potential of near 0 mV (Moriyoshi et al., 1991). The emerging picture is one where 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 (Fig. 4A–C).

The situation for the computational properties of layer 2/3 pyramidal neurons of rodent sensory cortex might seem on the face of it quite bizarre and is certainly quite unlike the integrate-and-fire picture (Fig. 4A–C). In this revised picture, real neurons ignore most synaptic inputs during high synaptic barrages. There may be, however, a perfectly simple explanation: to be optimally feature-specific, cortical neurons need to respond to a given feature in a manner that doesn't depend on the stimulus strength or other factors not related to saliency. Orientation-selective neurons in visual cortex are a case in point. If stimulus strength dominated the response of these neurons, non-preferred orientations at high stimulus strengths would evoke more output than preferred orientations at low stimulus strengths. To be optimally feature selective, a neuron should respond most vigorously to the pattern of inputs that relate to that particular feature and invariant to stimulus strength. This could in principle be achieved by having very few and specific synaptic connections, however cortical neurons typically have in the order of 10,000 synaptic inputs (Larkman, 1991), something quite unexpected in the world of machine learning that seems to demand an explanation (Hawkins and Ahmad, 2016). Furthermore, learning should involve becoming sensitive to new contexts or conditions. To do this, the cell needs to have access to information that is not (yet) relevant but nevertheless be able to ignore it until needed. Interestingly, it has long been observed that despite the huge differences in action potential firing in response to differently-oriented drifting gratings the subthreshold activity of neurons in primary visual cortex is very similar regardless of orientation (Carandini and Ferster, 2000). This implies that orientation-selective neurons get a substantial amount of synaptic inputs conveying information about all orientations (Jia et al., 2010), while nevertheless maintaining their orientation selectivity in terms of their output (Fig. 2B). In this way, neurons can be both feature-selective while also caring about the activity state of the network, perhaps indicating the general presence of a stimulus or brain state.

The sensitivity of the neuron to clustered input giving rise to local dendritic spikes that cause action potentials not only best explains the data but also makes computational sense. Using this approach, the neuron can be hyper-sensitive to a small number of inputs that represent precisely defined collections of synchronously firing presynaptic neurons, i.e. specific and preferred input patterns. The hundreds to thousands of inputs that contribute to the up-state serve the dual purpose of setting the neuron in the active state and increasing the overall density of synaptic inputs, thereby increasing the probability of spatial clusters of inputs in the first place. This interpretation would imply that, in terms of defining the precise conditions for firing action potentials moment to moment, the thousands of AMPA-dependent inputs are largely ignored. For the purposes of the current essay, it shows, at least in principle, that a radically different way of describing the function of a neuron that takes into account the active spatio-temporal properties of the dendritic tree, might lead to a more parsimonious and conceptually useful description of a bistable network.

ACTIVATION FUNCTIONS – CHALLENGING THE LINEAR SUMMATION MODEL

If a neuron could be compared to a logic gate, the integrate-and-fire model would be a kind of 'OR' gate, an inclusive-OR gate (Shepherd and Brayton, 1987). The binary version of this has two inputs either or both of which being 'ON' results in the output being 'ON' (or 'TRUE'). It is very natural and easy to implement with a counting device and a threshold function (an integrate-and-fire unit), and it also extrapolates conveniently to an N-gate with N inputs. However, inclusive-OR gates are not that useful computationally, partly because they don't do much actual work - they simply decide if there is or isn't enough input. Many real-world and computational issues involve decisions which are better characterized as either/or (i.e. exclusive-OR), e.g. should I choose vanilla or chocolate? It was recently shown that human layer 2/3 pyramidal neurons have a very interesting ion channel in their dendrites that allow them to compute the equivalent of an exclusive-OR, 'XOR', function (Fig. 4D, E) (Gidon et al., 2020). More broadly, this allows dendritic compartments to calculate the anti-coincidence of classes of inputs which adds to the range of functions that single biological neurons can compute (Segev, 1992; Herz et al., 2006). This aspect of dendritic computation is likely just the tip of the iceberg for what can be achieved with local dendritic spike mechanisms and interesting activation functions.

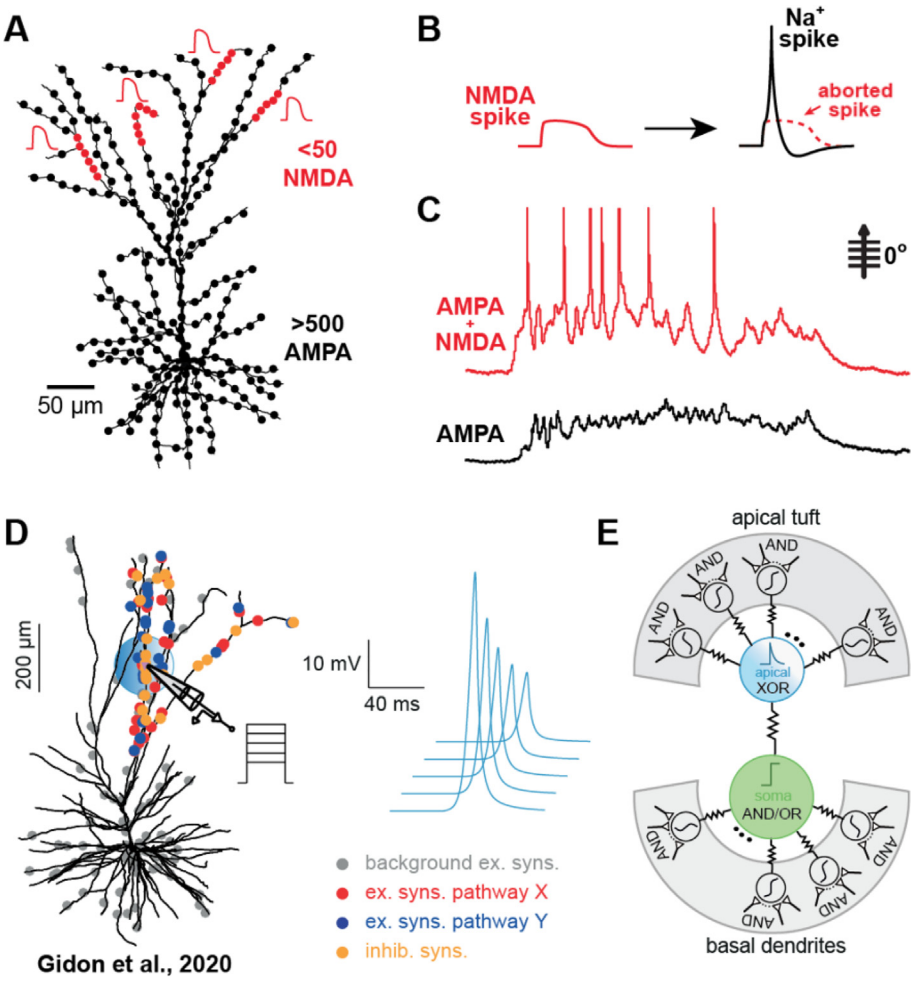

Fig. 4. Re-examining how neurons compute. **(A)** A L2/3 neuron receiving asynchronous distributed inputs that activate AMPA receptors at >500 (Chen et al. 2013) synapses (black dots) during an up-state and a few (<50) clustered synaptic inputs (red dots) that activate NMDA channels causing NMDA spikes. **(B)** Dendritic NMDA spikes (red, left) summate and cause up to 15 mV depolarization at the cell body (Palmer et al., 2014). This depolarization in addition to the up-state is likely to cause a Na^+ spike at the cell body which repolarizes due to the opening of K^+ channels aborting the longer NMDA spike. **(C)** It is hypothesized that, at the preferred orientation (here 0°), >500 AMPA broadly distributed EPSPs with balanced inhibition cause a steady up-state that with the addition of the combined NMDA spike current due to a handful (<50) clustered inputs, leads to robust AP spiking (red), whereas no spiking occurs without the <50 clustered inputs (black). This data is taken from (Jia et al. 2010). **(D)** A L2/3 human pyramidal neuron (left) with dendritic calcium action potentials (dCaAPs) that inactivate at currents above threshold (left, blue traces), can perform the XOR operation on two different excitatory pre-synaptic pathways that nevertheless distribute synapses over a similar region of the dendritic tree (red and blue dots). **(E)** A model of the human L2/3 pyramidal neuron with two distinct activation functions for spatially segregated compartments. D, E from (Gidon et al., 2020).

In a larger context, active dendrites serve not only to boost signals arriving at the dendrite (Stuart and Sakmann, 1995; Branco and Häusser, 2011), but also allow for multiple sites of spike initiation (Llinás et al., 1968; Calabrese and Kennedy, 1974; Mel, 1993; Larkum et al., 2009). Further complicating this situation is the fact that signals can travel actively in all directions throughout the dendritic tree. For instance, back-propagating sodium action potentials (Stuart and Sakmann, 1994) travel actively from the axonal initiation zone (Kole et al., 2008) allowing the output of the neuron to reverberate with inputs significantly complicating the input/output relationship (Waters et al., 2005). This aspect of biological neurons with dendritic trees is particularly hard to capture by using small networks of point neurons to represent dendrites (Poirazi et al., 2003) because neurons have privileged access in both directions to their dendrites compared to synaptic connections which only allow one-way interaction. The ability of neurons to signal freely throughout the dendritic tree means that the activation of the neuron can be a complex function over time and space which can only be captured with differential equations. It also means that some synaptic inputs or neuromodulation can serve to alter the coupling between dendritic compartments (Larkum et al., 2001; Schaefer et al., 2003; Suzuki and Larkum, 2020). All of these features speak to the conceptual efficiency of units with multiple compartments and non-linear properties.

DENDRITES FOR SEGREGATING INPUTS

The suggestion that dendrites might provide conceptual insights is not new and was championed by Wilfrid Rall as far back as the 1950s (reviewed in Jack et al., 1988; Segev, 1992). The core conceptual advantage of using dendrites is that they allow for the segregation of synaptic inputs, which in turn enables multiple parallel operations to be performed on the inputs, e.g. in retinal ganglion neurons (Koch et al., 1983). Point neurons cannot process different synaptic inputs separately. However, information can sometimes be categorically separated into different classes of information, such as in binaural processing (Agmon-Snir et al., 1998). Perhaps the most important partition for an information system is the dichotomy between external versus internal information which translates roughly to feed-forward and feedback information. In the cerebral cortex, these information streams follow broadly separable architectural

principles (Felleman and Van Essen, 1991; Harris et al., 2019). The class of neurons that is most abundant and ubiquitous in the cortex, the pyramidal neuron, is built and positioned to optimally separate and then integrate these two information streams allowing for the cortex to perform operations that have large-scale, network-wide consequences (Larkum, 2013; Larkum et al., 2018; Aru et al., 2020).

It would make sense that the elements of the cortex are built to be able to parse these information streams separately and could also significantly improve learning rules (Körding and König, 2000; Nielsen, 2003; Urbanczik and Senn, 2014; Kastellakis et al., 2015; Brea et al., 2016; Hawkins and Ahmad, 2016; Schiess et al., 2016; Sardi et al., 2018; Richards et al., 2019) and has been suggested as a way to solve the credit assignment problem in biological networks (Guerguiev et al., 2017). There is also growing empirical evidence for dendrite-specific learning (Letzkus et al., 2006; Sjöström and Häusser, 2006; Losonczy et al., 2008; Makara et al., 2009; Gambino et al., 2014; Cichon and Gan, 2015; Brandalise et al., 2016; Bittner et al., 2017; Sheffield et al., 2017; Abs et al., 2018; Roelfsema and Holtmaat, 2018; Doron et al., 2020; Shin et al., 2021). These data and models demonstrate that dendrites can function to segregate synaptic inputs so that particular rules can be applied to subsets of dendritic inputs. Various categories of information may be useful conceptually to segregate at the input stage. For example, information related to learning, like error signals, rewards, as well as states of arousal related to environmental conditions or cognitive states. The cerebellum is a prime example of the use of spatio-temporal principles efficiently implemented with parallel fibers coursing through rows of complex Purkinje cell dendritic trees. With point neurons, the segregation and re-integration of information would need to be performed by separate neurons making their integration more difficult both conceptually and practically. In summary, there is ample evidence that neurons of the brain do not behave like simple point neurons and that including dendrites in their description might be useful to efficiently capture their true function.

CONCLUDING REMARKS – DENDRITES AND WHAT CONSTITUTES UNDERSTANDING

Consider a scenario in which aliens correctly ascertain that the digital computers of the humans are binary devices. Could they immediately understand how Donkey Kong is programmed by reverse engineering? Probably not – at the very least, simply recording the transistors as they turn on and off would probably be unenlightening (Jonas and Körding, 2016). Whereas the original programmer could claim they, ‘‘don’t understand how the computer works”, for instance, and nevertheless write functional programs, we (humans) collectively do understand how digital computers work in the sense that we can describe how high-level descriptions are implemented. On the other hand, if the aliens incorrectly identified the binary nature of a digital computer, there would no way to implement high-level ideas and no conceptual through-line that could be described as understanding digital computers.

The key insight in terms of understanding is that the problem is not symmetric; it *is* necessary that the implementation level is understood (by someone at least) in order for anything to operate, but there is no *highest limit* that would constitute final understanding. The lower limit is defined by the level at which the components capture the principles of implementation. Here, replacing semiconductors with tubes would not change the important concepts at any level of understanding. And so it seems that the often cited claim of Richard Feynman that, ‘‘What I cannot create, I do not understand” (Way, 2017), is not invertible. That is, it is not necessarily true that ‘‘what I *can* create, I understand completely”. Understanding can be reached on many levels, and it is not necessarily true that understanding on one level entails understanding on another. For instance, do we understand artificial neural networks (Lillicrap and Kording, 2019)? Much effort is currently being invested in trying to understand how deep neural networks actually recognize faces or play ‘Go’ even though they can already perform better than humans.

What has this all to do with dendrites? Neurons with dendrites are clearly not linear counting devices. This is a category error that, it is argued here, will have consequences similar to that of aliens misinterpreting the function of transistors. What we are learning from looking closely at neurons is that real neurons are not only quite unlike point neurons, but under closer inspection they don’t even seem to operate according to the same principles. In other words, we are quite possibly trying to derive high-level theories of the brain while completely misunderstanding the substrate upon which they need to operate. But the problem may go deeper than this; there may well be aspects of the implementation that suggest higher-level principles. In my own work, for instance, I have argued that a useful description of a cortical pyramidal neuron is as a three-compartment neuron (Larkum et al., 2001) that compares external, feature-specific information with internally-generated predictions arriving at separate compartments (Larkum et al., 1999; Larkum, 2013) and that their privileged access and coupling under certain circumstances may be useful, for instance, for understanding of how anesthetics lead to loss of consciousness (Aru et al., 2020; Suzuki and Larkum, 2020) and what transpires at the moment of perception (Takahashi et al., 2016). I claim this should be described as implementation-level information that nevertheless suggests new ways to understand the brain at high levels of description. Furthermore, this understanding is fundamentally rooted in the conceptual importance of the dendrites.

ABANDONING THE SOMATO-CENTRIC PERSPECTIVE

This dendritically-enriched perspective of neurons offers a radical update of the integrate-and-fire model, which

simply counts inputs regardless of their dendritic location. This approach implies that the vast majority of current neural models, whether 'biologically-inspired' or not, are operating under misguided assumptions that may in effect render the models biologically inaccurate. How could this situation have gone on so long? I would argue this comes from the somato-centric perspective the neuroscience community has inherited (Larkum et al., 2018). Almost all recordings are either made from, or influenced by, the cell bodies of neurons. Extracellular methods that detect spikes in the electric field are much more sensitive to a big ball of excitable membrane (the soma) than to dendrites. Intracellular and imaging methods are also orders-of-magnitude easier and more common from cell bodies than intra-dendritic recordings. Furthermore, the point-neuron dogma is seemingly trivial to demonstrate: a student on their first day in an electrophysiology laboratory can easily show that the firing frequency of a neuron has a simple threshold/linear relationship to current injection at the cell body over most of its firing range. But this is really because the axon happens to be near the cell body, which means that the interesting aspects of single-cell computation are bypassed with recordings at the soma. Ironically, even the work done by the dendritic spikes is likely to be obscured at the cell body by the axo-somatic APs that repolarize the membrane (Fig. 4B). The experiments that need to be done to understand the input–output relationship of a neuron are much harder, however, and would require control of the ~10,000 synaptic inputs to the neuron. Better conceptual descriptions of neurons are not lacking. Perhaps, the simplest version of a pattern-dependent neuron is the 'clusteron' (Mel, 1994), however there are small network conceptual descriptions (Poirazi et al., 2003) and even full-scale, neural-network-in-a-neuron solutions being offered (Beniaguev et al., 2019; Jones and Kording, 2021). At the moment, empirical data is still lacking from neuroscience.

Since the brain arose from evolutionary processes, there is no one we can ask what the crucial base unit of computation in the brain is, nor how complex should be its description. At stake, is not only the amount of processing power that can be invested at the single-neuron level, but what conceptual efficiencies can be gained at higher levels of description, with better component-level descriptions. At a minimum, interactions between synapses and post-synaptic intrinsic excitability that are a function of space and time are almost certainly crucial for determining the input/output function of neurons and these facts appear to be lost in modern-day descriptions of neuronal function without dendrites. In any case, the best evidence we have from real neurons suggests that neuroscientists may have completely mischaracterized the conceptual operation of single neurons (Koch and Segev, 2000). There is every reason to suspect that better descriptions of sophisticated single-cell computation will lead to better descriptions at the network level blurring the distinction between Marr's algorithmic and implementation levels.

Box 1

Input clustering, NMDA spikes and local dendritic computation Another advantage of dendritic trees is that they allow for subsets of synaptic inputs to operate cooperatively. This was elegantly pointed out in the visionary work of Bartlett Mel and colleagues in the early 1990s (Mel, 1992, 1994; Poirazi and Mel, 2001), in which he speculated that NMDA receptors can in principle operate like other voltage-dependent ion channels making the neuron sensitive to clustered input. This led to the ground-breaking discovery of the so-called 'NMDA spike' by Jackie Schiller and colleagues (Schiller et al., 2000). Input clustering is quintessentially a spatial feature (Mel and Schiller, 2004), with possible computational advantages like sequence learning (Hawkins and Ahmad, 2016) and binocular rivalry (Archie and Mel, 2000). There is by now a lot of evidence for synaptic clustering in the brain (Larkum and Nevian, 2008; Kleindienst et al., 2011; Takahashi et al., 2012; Winnubst and Lohmann, 2012; Kastellakis et al., 2015; Weber et al., 2016; Wilson et al., 2016; Kerlin et al., 2019) but the impact of this on biological and artificial network models has so far been minimal (Kastellakis and Poirazi, 2019).

The location dependence of NMDA spikes is due not only to the necessity for clustered synaptic input but also to the fact that the NMDA spikes can only actively propagate along stretches of dendrites where NMDA receptors are activated, i.e. with bound glutamate (Major et al., 2013). That is, this form of spiking is intrinsically location-dependent with the properties aecting their spatial extent constantly fluctuating. A model incorporating the NMDA spike (Rao et al., 2021) is therefore the conceptual antithesis of the dendrite-less point neuron used ubiquitously in computational and machine learning models. The influence of the NMDA spike can also be orders of magnitude larger than the depolarization caused by EPSPs such that the underlying synapses can be transformed into having a massively disproportionate influence on the generation of somatic action potentials (Polsky et al., 2009). The concept of local cluster-driven NMDA spikes and other kinds of dendritic spikes stands to revolutionize our understanding of the way neurons compute and thus our understanding of the brain and the way we build artificial networks.

ACKNOWLEDGEMENTS

I thank Albert Gidon, Jaan Aru and James M. Shine for their comments on the manuscript. The research was supported by German Research Foundation (LA 3442/2-1, EXC 257 NeuroCure, and project number 327654276 – SFB 1315) and European Union Horizon 2020 Research and Innovation Programme (under grant agreement numbers HBP SGA3/945539 and 670118/ERC ActiveCortex).

REFERENCES

Abs E, Poorthuis RB, Apelblat D, Muhammad K, Pardi MB, Enke L, Kushinsky D, Pu D-L, Eizinger MF, Conzelmann K-K, Spiegel I,

Letzkus JJ (2018) Learning-related plasticity in dendrite-targeting layer 1 interneurons. Neuron 100:684–699.e6.
Agmon-Snir H, Carr CE, Rinzel J (1998) The role of dendrites in auditory coincidence detection. Nature 393:268–272.
Antic SD, Zhou W-L, Moore AR, Short SM, Ikonomu KD (2010) The decade of the dendritic NMDA spike. J Neurosci Res 88:2991–3001.
Archie KA, Mel BW (2000) A model for intradendritic computation of binocular disparity. Nat Neurosci 3:54–63.
Aru J, Suzuki M, Larkum ME (2020) Cellular mechanisms of conscious processing. Trends Cogn Sci:S1364661320301753.
Beniaguev D, Segev I, London M (2019) Single cortical neurons as deep artificial neural networks. Neuroscience. Available at: http://biorxiv.org/lookup/doi/10.1101/613141 [Accessed September 8, 2020].
Bittner KC, Milstein AD, Grienberger C, Romani S, Magee JC (2017) Behavioral time scale synaptic plasticity underlies CA1 place fields. Science 357:1033–1036.
Branco T, Häusser M (2011) Synaptic integration gradients in single cortical pyramidal cell dendrites. Neuron 69:885–892.
Brandalise F, Carta S, Helmchen F, Lisman J, Gerber U (2016) Dendritic NMDA spikes are necessary for timing-dependent associative LTP in CA3 pyramidal cells. Nat Commun 7:13480.
Brea J, Gaál AT, Urbanczik R, Senn W (2016) Prospective coding by spiking neurons Latham PE, ed.. PLoS Comput Biol 12 e1005003.
Calabrese RL, Kennedy D (1974) Multiple sites of spike initiation in a single dendritic system. Brain Research 82:316–321.
Carandini M, Ferster D (2000) Membrane potential and firing rate in cat primary visual cortex. J Neurosci 20:470–484.
Chen X, Rochefort NL, Sakmann B, Konnerth A (2013) Reactivation of the same synapses during spontaneous up states and sensory stimuli. Cell Reports 4:31–39.
Cichon J, Gan W-B (2015) Branch-specific dendritic Ca2+ spikes cause persistent synaptic plasticity. Nature 520:180–185.
Cui Yuwei, Ahmad Subutai, Hawkins Jeff (2016) Continuous online sequence learning with an unsupervised neural network model. Neural Computation 28:2474–2504. https://doi.org/10.1162/NECO_a_00893.
De Schutter E (2008) Why are computational neuroscience and systems biology so separate? Friston KJ, ed.. PLoS Comput Biol 4 e1000078.
Destexhe A (2009) Self-sustained asynchronous irregular states and Up–Down states in thalamic, cortical and thalamocortical networks of nonlinear integrate-and-fire neurons. J Comput Neurosci 27:493–506.
Destexhe A, Rudolph M, Paré D (2003) The high-conductance state of neocortical neurons in vivo. Nat Rev Neurosci 4:739–751.
Doron G, Shin JN, Takahashi N, Drüke M, Bocklisch C, Skenderi S, de Mont L, Toumazou M, Ledderose J, Brecht M, Naud R, Larkum ME (2020) Perirhinal input to neocortical layer 1 controls learning. Science 370.
Durbin R, Rumelhart DE (1989) Product units: a computationally powerful and biologically plausible extension to backpropagation networks. Neural Comput 1:133–142.
Felleman DJ, Van Essen DC (1991) Distributed hierarchical processing in the primate cerebral cortex. Cereb Cortex 1:1–47.
Gambino F, Pagès S, Kehayas V, Baptista D, Tatti R, Carleton A, Holtmaat A (2014) Sensory-evoked LTP driven by dendritic plateau potentials in vivo. Nature 515:116–119.
Gidon A, Zolnik TA, Fidzinski P, Bolduan F, Papoutsi A, Poirazi P, Holtkamp M, Vida I, Larkum ME (2020) Dendritic action potentials and computation in human layer 2/3 cortical neurons. Science 367:83–87.
Guerguiev J, Lillicrap TP, Richards BA (2017) Towards deep learning with segregated dendrites. eLife 6:e22901.
Goldental Amir, Guberman Shoshana, Vardi Roni, Kanter Ido (2014) A computational paradigm for dynamic logic-gates in neuronal activity. Frontiers in Computational Neuroscience 8, 52. https://doi.org/10.3389/fncom.2014.00052.
Hahn TTG, Sakmann B, Mehta MR (2006) Phase-locking of hippocampal interneurons' membrane potential to neocortical up-down states. Nat Neurosci 9:1359–1361.
Harris JA et al (2019) Hierarchical organization of cortical and thalamic connectivity. Nature 575:195–202.
Häusser M, Mel B (2003) Dendrites: bug or feature? Curr Opin Neurobiol 13:372–383.
Hawkins J, Ahmad S (2016) Why neurons have thousands of synapses, a theory of sequence memory in neocortex. Front Neural Circuits 10 Available at: http://journal.frontiersin.org/Article/10.3389/fncir.2016.00023/abstract [Accessed January 8, 2021].
Herz AVM, Gollisch T, Machens CK, Jaeger D (2006) Modeling single-neuron dynamics and computations: a balance of detail and abstraction. Science 314:80–85.
Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. Proc Natl Acad Sci 79:2554–2558.
Hopfield JJ, Brody CD (2000) What is a moment? "Cortical" sensory integration over a brief interval. Proc Acad Sci 97:13919–13924.
Jack JJB, Noble D, Tsien RW (1988) Electric current flow in excitable cells, Reprint. Oxford: Clarendon Press.
Jia H, Rochefort NL, Chen X, Konnerth A (2010) Dendritic organization of sensory input to cortical neurons in vivo. Nature 464:1307–1312.
Jonas E, Körding KP (2016) Could a neuroscientist understand a microprocessor? Neuroscience. Available at: http://biorxiv.org/lookup/doi/10.1101/055624 [Accessed February 3, 2021].
Jones IS, Kording KP (2021) Might a Single Neuron Solve Interesting Machine Learning Problems Through Successive Computations on Its Dendritic Tree? Neural Computation 33:1554–1571. https://doi.org/10.1162/neco_a_01390.
Kastellakis G, Cai DJ, Mednick SC, Silva AJ, Poirazi P (2015) Synaptic clustering within dendrites: An emerging theory of memory formation. Prog Neurobiol 126:19–35.
Kastellakis G, Poirazi P (2019) Synaptic clustering and memory formation. Front Mol Neurosci 12:300.
Kerlin A, Mohar B, Flickinger D, MacLennan BJ, Dean MB, Davis C, Spruston N, Svoboda K (2019) Functional clustering of dendritic activity during decision-making. eLife 8:e46966.
Kleindienst T, Winnubst J, Roth-Alpermann C, Bonhoeffer T, Lohmann C (2011) Activity-dependent clustering of functional synaptic inputs on developing hippocampal dendrites. Neuron 72:1012–1024.
Koch C, Poggio T, Torre V (1983) Nonlinear interactions in a dendritic tree: localization, timing, and role in information processing. Proc Natl Acad Sci 80:2799–2802.
Koch C, Segev I (2000) The role of single neurons in information processing. Nat Neurosci 3:1171–1177.
Kole MHP, Ilschner SU, Kampa BM, Williams SR, Ruben PC, Stuart GJ (2008) Action potential generation requires a high sodium channel density in the axon initial segment. Nat Neurosci 11:178–186.
Körding KP, König P (2000) Learning with two sites of synaptic integration. Network 11:25–39.
Krakauer JW, Ghazanfar AA, Gomez-Marin A, MacIver MA, Poeppel D (2017) Neuroscience needs behavior: correcting a reductionist bias. Neuron 93:480–490.
Larkman AU (1991) Dendritic morphology of pyramidal neurones of the visual cortex of the rat: III. Spine distributions. J Comp Neurol 306:332–343.
Larkum M (2013) A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. Trends Neurosci 36:141–151.
Larkum ME, Nevian T (2008) Synaptic clustering by dendritic signalling mechanisms. Curr Opin Neurobiol 18:321–331.
Larkum ME, Nevian T, Sandler M, Polsky A, Schiller J (2009) Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle. Science 325:756–760.

Larkum ME, Petro LS, Sachdev RNS, Muckli L (2018) A perspective on cortical layering and layer-spanning neuronal elements. Front Neuroanat 12:56.
Larkum ME, Zhu JJ, Sakmann B (2001) Dendritic mechanisms underlying the coupling of the dendritic with the axonal action potential initiation zone of adult rat layer 5 pyramidal neurons. J Physiol 533:447–466.
Larkum Matthew E, Zhu J Julius, Sakmann Bert (1999) A new cellular mechanismfor coupling inputs arriving at different cortical layers. Nature 398:338–341. https://doi.org/10.1038/18686.
Letzkus JJ, Kampa BM, Stuart GJ (2006) Learning rules for spike timing-dependent plasticity depend on dendritic synapse location. J Neurosci 26:10420–10429.
Lillicrap TP, Kording KP (2019) What does it mean to understand a neural network? arXiv:190706374 [cs, q-bio, stat] Available at: http://arxiv.org/abs/1907.06374 [Accessed March 3, 2021].
Liu G (2004) Local structural balance and functional interaction of excitatory and inhibitory synapses in hippocampal dendrites. Nat Neurosci 7:373–379.
Llinás R, Nicholson C, Freeman JA, Hillman DE (1968) Dendritic spikes and their inhibition in alligator purkinje cells. Science 160:1132–1135.
London M, Schreibman A, Häusser M, Larkum ME, Segev I (2002) The information efficacy of a synapse. Nat Neurosci 5:332–340.
Losonczy A, Makara JK, Magee JC (2008) Compartmentalized dendritic plasticity and input feature storage in neurons. Nature 452:436–441.
Major G, Larkum ME, Schiller J (2013) Active properties of neocortical pyramidal neuron dendrites. Annu Rev Neurosci 36:1–24.
Makara JK, Losonczy A, Wen Q, Magee JC (2009) Experience-dependent compartmentalized dendritic plasticity in rat hippocampal CA1 pyramidal neurons. Nat Neurosci 12:1485–1487.
McCormick DA (2015) Brain state dependent activity in the cortex and thalamus. Curr Opin Neurobiol:8.
Marr D (1982). Vision: A Computational Approach. MIT Press.
Mel BW (1992) NMDA-based pattern discrimination in a modeled cortical neuron. Neural Computation 4:502–517.
Mel BW (1993) Synaptic integration in an excitable dendritic tree. Journal of Neurophysiology 70:1086–1101.
Mel BW (1994) Information processing in dendritic trees. Neural Comput 6:1031–1085.
Mel BW, Schiller J (2004) On the fight between excitation and inhibition: location is everything. Sci Signal 2004:pe44–pe44.
Moriyoshi K, Masu M, Ishii T, Shigemoto R, Mizuno N, Nakanishi S (1991) Molecular cloning and characterization of the rat NMDA receptor. Nature 354:31–37.
Nielsen BG (2003) Sequence learning in differentially activated dendrites. Network: Comput Neural Syst 14:189–209.
Okun M, Lampl I (2008) Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. Nat Neurosci 11:535–537.
Okun M, Naim A, Lampl I (2010) The subthreshold relation between cortical local field potential and neuronal firing unveiled by intracellular recordings in awake rats. J Neurosci 30:4440–4448.
Palmer LM, Shai AS, Reeve JE, Anderson HL, Paulsen O, Larkum ME (2014) NMDA spikes enhance action potential generation during sensory input. Nat Neurosci 17:383–390.
Petersen CCH, Hahn TTG, Mehta M, Grinvald A, Sakmann B (2003) Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex. Proc Natl Acad Sci 100:13638–13643.
Poirazi P, Brannon T, Mel BW (2003) Pyramidal neuron as two-layer neural network. Neuron 37:989–999.
Poirazi P, Mel BW (2001) Impact of active dendrites and structural plasticity on the memory capacity of neural tissue. Neuron 29:779–796.
Poirazi P, Papoutsi A (2020) Illuminating dendritic function with computational models. Nat Rev Neurosci Available at: http://www.nature.com/articles/s41583-020-0301-7 [Accessed May 12, 2020].
Polsky A, Mel B, Schiller J (2009) Encoding and decoding bursts by NMDA spikes in basal dendrites of layer 5 pyramidal neurons. J Neurosci 29:11891–11903.
Rao A, Legenstein R, Subramoney A, Maass W (2021) Synaptic plasticity in distal dendrites can optimize probabilistic top-down predictions in pyramidal cells. Neuroscience. Available at: http://biorxiv.org/lookup/doi/10.1101/2021.03.04.433822 [Accessed March 10, 2021].
Richards BA et al (2019) A deep learning framework for neuroscience. Nat Neurosci 22:1761–1770.
Roelfsema PR, Holtmaat A (2018) Control of synaptic plasticity in deep cortical networks. Nat Rev Neurosci 19:166–180.
Rumelhart DE, McClelland JL (1986) Parallel distributed processing: explorations in the microstructure of cognition. Cambridge, Mass: MIT Press.
Sachidhanandam S, Sreenivasan V, Kyriakatos A, Kremer Y, Petersen CCH (2013) Membrane potential correlates of sensory perception in mouse barrel cortex. Nat Neurosci 16:1671–1677.
Sanchez-Vives MV, McCormick DA (2000) Cellular and network mechanisms of rhythmic recurrent activity in neocortex. Nat Neurosci 3:1027–1034.
Sardi S, Vardi R, Goldental A, Tugendhaft Y, Uzan H, Kanter I (2018) Dendritic learning as a paradigm shift in brain learning. ACS Chem Neurosci 9:1230–1232.
Schaefer AT, Larkum ME, Sakmann B, Roth A (2003) Coincidence detection in pyramidal neurons is tuned by their dendritic branching pattern. J Neurophysiol 89:3143–3154.
Schiess M, Urbanczik R, Senn W (2016) Somato-dendritic synaptic plasticity and error-backpropagation in active dendrites Gutkin BS, ed. PLoS Comput Biol 12 e1004638.
Schiller J, Major G, Koester HJ, Schiller Y (2000) NMDA spikes in basal dendrites of cortical pyramidal neurons. Nature 404:285–289.
Segev I (1992) Single neurone models: oversimple, complex and reduced. Trends Neurosci 15:414–421.
Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. J Neurosci 18:3870–3896.
Sheffield MEJ, Adoff MD, Dombeck DA (2017) Increased Prevalence of Calcium Transients across the Dendritic Arbor during Place Field Formation. Neuron 96:490–504.e5.
Shepherd GM, Brayton RK (1987) Logic operations are properties of computer-simulated interactions between excitable dendritic spines. Neuroscience 21:151–165.
Shin JN, Doron G, Larkum ME (2021) Memories off the top of your head. Science 374:538–539.
Sjöström PJ, Häusser M (2006) A cooperative switch determines the sign of synaptic plasticity in distal dendrites of neocortical pyramidal neurons. Neuron 51:227–238.
Smith SL, Smith IT, Branco T, Häusser M (2013) Dendritic spikes enhance stimulus selectivity in cortical neurons in vivo. Nature 503:115–120.
Stuart G, Sakmann B (1995) Amplification of EPSPs by axosomatic sodium channels in neocortical pyramidal neurons. Neuron 15:1065–1076.
Stuart GJ, Sakmann B (1994) Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. Nature 367:69–72.
Suzuki M, Larkum ME (2020) General anesthesia decouples cortical pyramidal neurons. Cell 180:666–676.e13.
Takahashi N, Kitamura K, Matsuo N, Mayford M, Kano M, Matsuki N, Ikegaya Y (2012) Locally synchronized synaptic inputs. Science 335:353–356.
Takahashi Naoya, Oertner Thomas, Hegemann Peter, Larkum Matthew E (2016) Active cortical dendrites modulate perception. Science 354(6319):1587–1590. https://doi.org/10.1126/science.aah6066.
Torao-Angosto M, Manasanch A, Mattia M, Sanchez-Vives MV (2021) Up and down states during slow oscillations in slow-wave sleep and different levels of anesthesia. Front Syst Neurosci 15 609645.

Tukker JJ, Sachdev RNS (2020) Up and down states and memory consolidation across somatosensory, entorhinal, and hippocampal cortices. Front Syst Neurosci 14:17.
Urbanczik R, Senn W (2014) Learning by the dendritic prediction of somatic spiking. Neuron 81:521–528.
Waters J, Helmchen F (2004) Boosting of action potential backpropagation by neocortical network activity in vivo. J Neurosci 24:11127–11136.
Waters J, Schaefer A, Sakmann B (2005) Backpropagating action potentials in neurones: measurement, mechanisms and potential functions. Prog Biophys Mol Biol 87:145–170.
Way M (2017) What I cannot create, I do not understand. J Cell Sci 130:2941–2942.
Weber JP, Andrásfalvy BK, Polito M, Magó Á, Ujfalussy BB, Makara JK (2016) Location-dependent synaptic plasticity rules by dendritic spine cooperativity. Nat Commun 7:11380.
Williams SR, Stuart GJ (2002) Dependence of EPSP efficacy on synapse location in neocortical pyramidal neurons. Science 295:1907–1910.
Wilson C (2008) Up and down states. Scholarpedia J 3:1410.
Wilson DE, Whitney DE, Scholl B, Fitzpatrick D (2016) Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex. Nat Neurosci 19:1003–1009.
Winnubst J, Lohmann C (2012) Synaptic clustering during development and learning: the why, when, and how. Front Mol Neurosci 5 Available at: http://journal.frontiersin.org/article/10.3389/fnmol.2012.00070/abstract [Accessed March 10, 2021].

(Received 19 March 2021, Accepted 5 March 2022)
(Available online 11 March 2022)