ACTIVE DENDRITES ENABLE ROBUST SPIKING COMPUTATIONS DESPITE TIMING JITTER

Thomas SJ Burger[1,*], **Michael E Rule**[1], **and Timothy O'Leary**[1,2,*]

[1]Department of Engineering, University of Cambridge, UK
[2]Theoretical Sciences Visiting Program, Okinawa Institute of Science and Technology Graduate University, Onna, 904-0495, Japan
[*]Correspondence: tsjb2@cam.ac.uk, timothy.oleary@eng.cam.ac.uk

ABSTRACT

Dendritic action potentials exhibit long plateaus of many tens of milliseconds, outliving axonal spikes by an order of magnitude. The computational role of these slow events seems at odds with any need to rapidly integrate and relay information throughout large nervous systems. We propose that the timescale of dendritic potentials allows reliable integration of asynchronous inputs. We develop a physiologically grounded model in which the extended duration of dendritic spikes equips each dendrite with a resettable memory of incoming signals. This provides a tractable model for capturing dendritic nonlinearities observed in experiments and in more complex, detailed models. Using this model, we show that long-lived, nonlinear dendritic plateau potentials allow reliable integration of asynchronous spikes. We demonstrate this model supports non-trivial computations in a network solving an arbitrary association/discrimination task using sparse spiking that is subject to timing jitter. This demonstrates a computational role for the specific timecourse of dendritic potentials in situations where decisions occur quickly, reliably, and with a low number of spikes. Our results provide empirically testable hypotheses for the role of dendritic action potentials in cortical function as well as a potential bio-inspired means of realising neuromorphic spiking computations in analog hardware.

***K*eywords** neurophysiology, spiking neural networks, dendrites, neuromorphic computing

Introduction

Across species, many types of neurons possess active dendrites that produce strongly nonlinear responses to synaptic input [31, 43, 20, 8]. The computational role of these nonlinearities is diverse and will depend on function of the wider neural circuit they inhabit. Some of the most intensely studied examples of dendritic excitability are found in cortical excitatory neurons, which produce regenerative action currents in response to excitatory synaptic drive [43, 55, 38].

Cortical excitatory dendritic action currents last for many tens of milliseconds [55, 54, 53, 44, 40, 38]. This feature is conspicuous because it is an order of magnitude longer than unitary synaptic inputs and axonal spikes. Reconciling the slow timecourse of dendritic potentials with rapid signalling and computation therefore poses a challenge, particularly when such computations may involve relaying information over multiple brain areas in a short time interval [62]. Furthermore, the duration of dendritic events incurs heavy energetic costs, because dendritic currents contribute significantly to the ATP budget of the brain [4]. What computational benefit might couterbalance these signalling and metabolic costs?

We propose that the duration and threshold-like properties of dendritic currents support robust computation in the face of spike timing jitter. This is especially relevant to integration of inputs during high conductance states that are prevalent in-vivo. In these states the effective time constant of the neuronal membrane is extremely short and varies substantially depending on synaptic drive [13, 34, 49]. As a consequence, computations that rely on passive summation of multiple inputs place punishing constraints on spike timing precision. Dendritic action potentials, by contrast, have a consistently long duration that is ensured by the kinetic properties of voltage gated ion channels and NMDA receptors [54, 47, 10, 3]. These properties are largely determined by the amino acid sequence of receptor

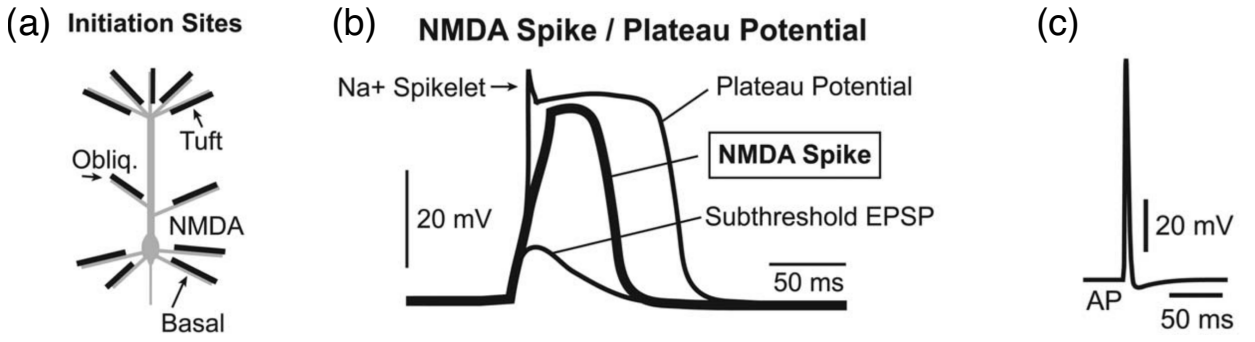

Figure 1: Dendritic NMDA-dependent action currents. **(a)** Long lived calcium and voltage transients initiated within distal dendritic branches. **(b)** If the input to a dendritic branch with NMDA receptors is sufficiently strong, it can cross a threshold to produce an NMDA spike (bold trace). The NMDA response to inputs is super-linear, until it saturates in a plateau potential (top trace). **(c)** A somatic spike mediated by voltage-gated sodium channels. Note the order of magnitude difference in timescales with (b). Reproduced from [3].

and channel proteins that are specifically expressed in dendrites [45, 44, 40]. This suggests dendritic properties are specifically tuned to produce localised, suprathreshold events that outlive rapid membrane fluctuations.

We extract these core features of the biophysics of dendritic integration to construct and analyse a simplified model, showing that rapid computation remains possible and is in fact facilitated by dendritic transients that exceed the integration timeconstant of single neurons. We focus on computations that take place on the most rapid timescale, because short integration windows are necessarily more sensitive to timing jitter. An interesting side product of this analysis is the interpretation of rapid cortical computations as operating in a binary regime, with each neuron possessing an integration window that can accommodate at most one spike from each input. A number of studies find empirical evidence for such an operating regime in different parts of the nervous system [14, 25, 61, 62]. We show how dendritic potentials in this regime allow non-trivial, robust and rapid spiking computations at the network level.

Numerous studies point out that nonlinear summation in dendrites can make neurons computationally equivalent to entire networks of simplified point models, or 'units' in a traditional neural network [9, 21, 38, 40, 45, 48, 50, 51]. Other work has shown that the dynamic properties of dendritic action potentials enrich computational and signal processing capacity by providing additional long timescales over which input-driven membrane potential dynamics evolve [22, 51, 7, 37, 50, 23]. These ideas and the specific examples that support them are complimentary to what we are proposing here. With the dendritic potential as a backbone, our work adds to the computational repertoire by allowing neurons to tune sensitivity to spike timing so as to achieve robust computation on rapid timescales. Our work therefore suggests that long-lived dendritic potentials can paradoxically assist in the most rapid computations possible in a spiking network.

1 Results

1.1 Abstract model

Key features of NMDA action currents are their long duration and their super-linear integration of inputs [43]. Figure 2A and B show a recording of an NMDA spike from a cortical neuron in a rat, from [18]. [18] triggered two NMDA spikes by glutamate uncaging at the indicated (red, blue) sites. The voltage response (Figure 2B) reveals the orders-of-magnitude difference in timescale of an NMDA spike (left) compared to a sodium spike in the soma (right). We took their extensive biophysical computational model (85 compartments, 439 segments—for details see [18] and subsection 3.1) and simulated glutamate releases 50 ms apart in the three dendritic sites indicated in Figure 2C, thereby triggering three NMDA spikes at those sites. Despite these dendritic spikes being initiated at different times, they still sum in the soma, leading to a sodium spike there (Figure 2E).

We developed an abstract model of these NMDA spikes to capture the essence of their role in circuit computations and for computational expediency. The model consists of a somatic compartment coupled passively to multiple dendritic compartments, each of which corresponds to a single branch on the dendritic tree (Figure 2E; for details see subsection 3.2). We model NMDA spikes by thresholding the voltage of a leaky dendritic compartment. When the dendritic voltage exceeds threshold it remains depolarized for some time before returning to rest (Figure 2F). The voltage dynamics are thus parametrised by the threshold and duration of the NMDA spike. We refer to this behavior

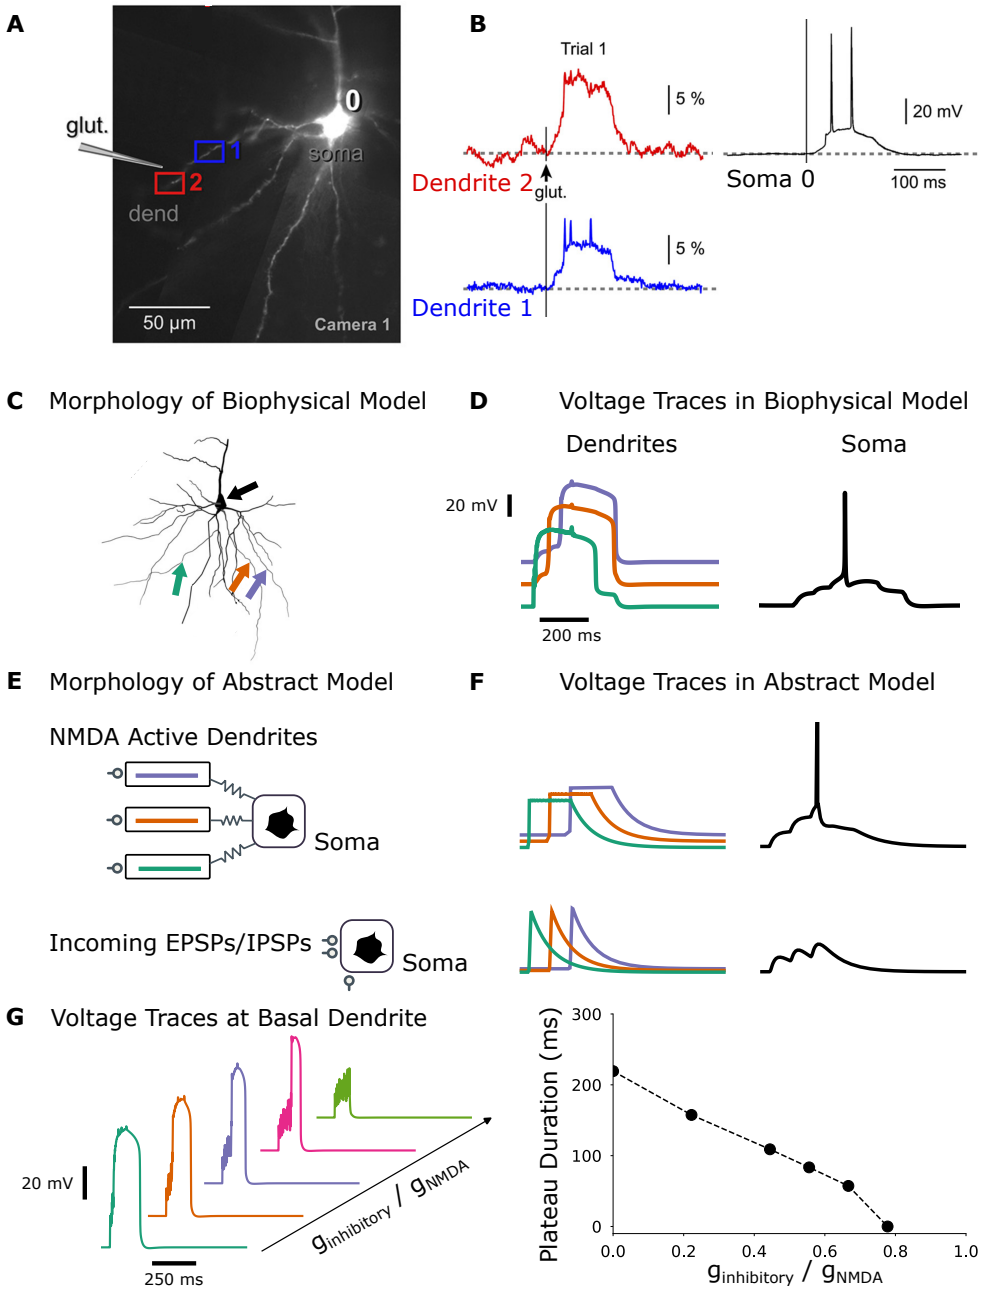

Figure 2: **A** Fluorescent dye fill image of a layer V pyramidal neuron recorded intracellularly in a rat cortical slice [18]. Two dendrites are marked where an NMDA spike is triggered through release of glutamate. **B** Voltage traces of the NMDA spikes of the dendritic patches marked in A. Note the long duration of the NMDA spike compared to the somatic sodium spikes evident in the top-left panel. **C** The morphology of the biophysical model used for simulating detailed NMDA plateau potentials. The arrows mark the positions where glutamate release was simulated. **D** The three traces of the NMDA spikes triggered at the sites marked in C, and the resulting somatic spike. **E** The morphology of the abstract model, with and without active NMDA dendrites. **F** The voltage traces of the abstract model, with and without plateaus. Because of the extended time duration of the plateau potentials, they sum accurately to produce a somatic spike. In the case where the plateau potentials are absent, they do not sum due to the short membrane time-constant of the soma. **G** Voltage traces of a basal dendrite with an NMDA spike, in the biophysical model, with an increasingly strong inhibitory current added (left). The plateau duration decreases linearly for a linear increase in the inhibitory conductance (right).

as "Leaky Integrate-and-Hold" (LIH). It captures the salient features of the NMDA spikes, namely the threshold plus saturation of the super-linear integration, and the long-lived plateau of the dendritic voltage.

Due to passive coupling between compartments in the model, excitation in the dendrites depolarises the soma membrane potential, potentially leading to "axonal" output spikes. We do not model an axonal compartment, we instead implement standard leaky integrate-and-fire (LIF) dynamics to the somatic compartment. A detailed description of the model, along with links to code, is provided in methods.

We compared the behaviour of our simplified model with that of the full, detailed biophysical model. The plateau potentials in the abstract model have a qualitatively similar effect on somatic membrane potential as the NMDA spikes in the biophysical model: Figure 2F shows that spikes arriving at different times are summed in an integrate and hold-like manner.

We compared this to a situation where all inputs arrive at a soma with standard LIF dynamics and a 10 ms membrane time constant. This time constant is consistent with the high-conductance state of pyramidal neurons in the cortex [6]: Inputs decay after 2–3 ms, and fail to sum to spike threshold (Figure 2F, lower).

To partially account for effects of inhibition, we assessed the robustness of dendritic plateaus to tonic inhibitory conductance. As can be seen in Figure 2G, dendritic plateaus survive inhibitory conductance up to values where the total conductance is roughly equal. Thus, to a crude approximation, dendritic potentials provide an integrate-and-hold mechanism that could function e.g. in the balanced regimen observed in cortical circuits. In the present study we did not attempt to account for temporal variation in inhibition, which will likely play a role in providing further spike synchrony, among other things. This is an important issue we intend to return to in future work. In the scope of what remains here we want to ask if integrate-and-hold is minimally feasible, and if so, whether it can easily and plausibly facilitate network computations with spikes.

1.2 Single Neurons Struggle to Integrate Asynchronous Spikes

The simplified model captures a key feature of the detailed biophysics of pyramidal neuron dendrites: the ability to integrate and hold inputs for a duration exceeding the membrane timeconstant. We hypothesized that this feature would be useful in situations where neurons need to integrate asynchronous input and reliably threshold it despite fluctuations in arrival times of the input.

In effect, each dendrite is performing a binary classification on its inputs. If input spikes arrive in a narrow time window, reliably integrating them is trivial (Figure 3A,B left). However, millisecond-scale synchrony is unlikely in a large network that is subject to uncertainty and noise. Empirically, spike timing jitter is commonly observed at the population level.

To illustrate the severity of this problem, we modelled a single neuron using our abstract model, and fed it input spikes. We drew the times of these input spikes from a normal distribution, and varied the degree of input synchrony by changing the standard deviation of this normal distribution. We took the standard deviation to be a function of the membrane time constant τ, which defines the timescale of the neuron dynamics.

Spikes arriving even slightly out of sync with each other introduces noise in the membrane potential of the receiving neuron (Figure 3B,D), which can lead to the neuron failing to spike when it should have, or vice-versa. Asynchrony reduces the effective drive of the inputs (Figure 3C), which means that a failure to spike will occur more often than an errant spike. This loss of drive could be compensated by lowering the postsynaptic cell's threshold, but variability due to jitter remains. This is shown in (Figure 3D), where we used the coefficient of variation of the peak membrane potential (standard deviation divided by the mean) to summarize the membrane-voltage uncertainty. This grows with increasing input-timing jitter.

Having extended NMDA spikes remedies these issues (Figure 3A). Because of the extension of the time-constant of the spikes, the uncertainty is filtered out, and the spikes are integrated as if they had arrived synchronously.

It is worth noting that these problems can also be addressed by increasing the number of inputs, thereby reducing the uncertainty through averaging. However, many inputs are required to keep the uncertainty in the membrane potential low; Figure 3E shows that there is a linear relationship between the spike timing jitter in the inputs, and the number of input spikes necessary to have low uncertainty. Even for a relatively low amount of jitter in the input spike timings of 10 ms, the number of inputs required is in the hundreds. Furthermore, it is only possible to average-away timing jitter if timing variations are uncorrelated. This need not be the case, especially if timing jitter arises from variable conduction delays from common sources.

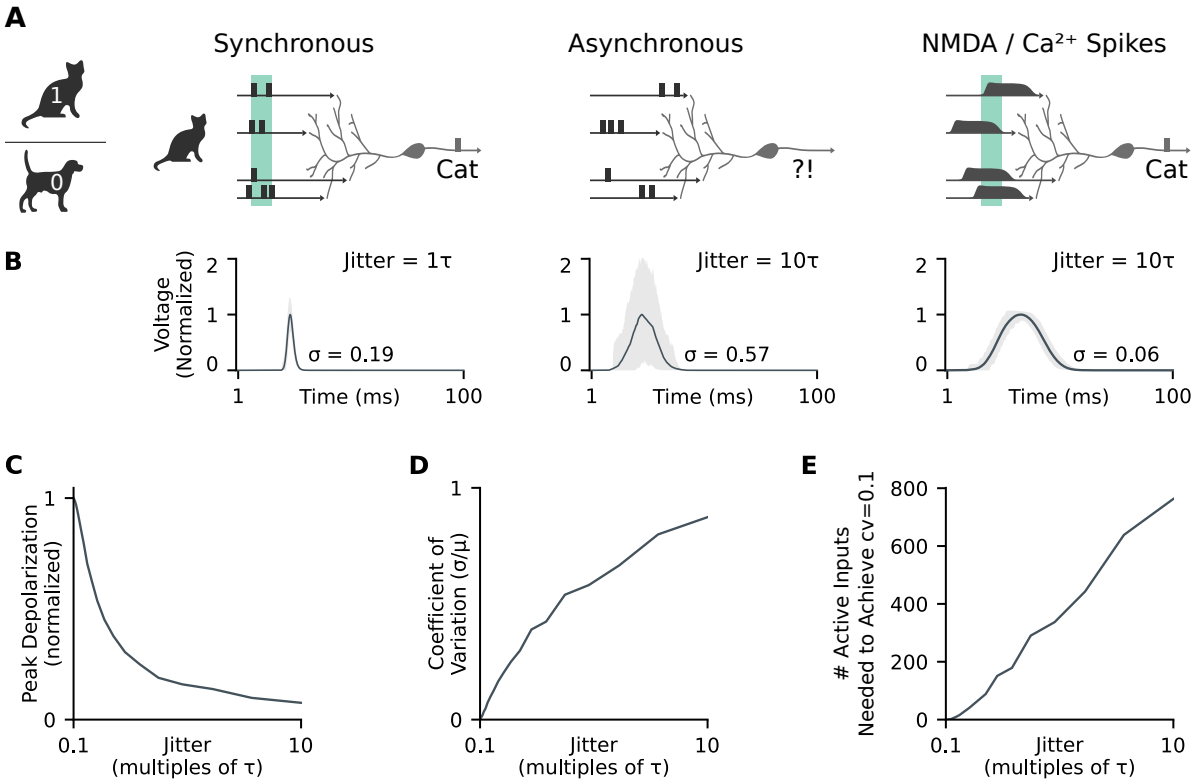

Figure 3: **(A)** Neurons integrate inputs and compare the result to a firing threshold, which is comparable to performing a binary classification. When inputs are synchronous this can be done with a low number of spikes (left). But when spike timing is unpredictable (middle), this falls apart. Extended depolarizing potentials within dendritic compartments acts as a hold mechanism, allowing asynchronous events from different compartments to summate (right). **(B)** Simulation of summed voltage for 10 dendritic compartments, for small amounts of input-event timing jitter (on the order of one EPSP duration $\tau \sim 1$ ms; left), and larger amounts jitter (10τ, middle). Increased jitter increases the variability of the net depolarization. Extended depolarizing potentials on the duration of ~ 20 ms reduce the variability in net depolarization. **(C)** Increased jitter in the timing of input events reduces the net summed depolarization. **(D)** Increased jitter in the timing of input events increases the variability in membrane voltage depolarization. **(E)** Variability can be restored by increasing the number of inputs, but this is not cost-effective.

1.3 Active Dendrites confer Robustness in Spiking Computations

So far we have shown how a biophysical mechanism extracted from detailed biophysics naturally extends the integration window for excitatory inputs. While this might be useful in principle, it remains for us to show how such dynamics can permit non-trivial computations in a network.

In our model, the crucial difference between summation at the soma and summation at the dendrite is that each dendrite can sum subthreshold inputs passively and independently, while the soma is summing sustained plateau potentials from all dendrites that happen to be active. We have made fairly simple assumptions that the dendrite has linear properties beneath the threshold for a dendritic spike, and a relatively short timeconstant. If neither of these assumptions hold, then dendrites might have even more robust integration properties than our model assumes. In this sense our claims and results in this section are rather conservative.

We assumed that inputs to a network arrive at the dendrites within some time window, and their combined depolarisations are either sufficient to either elicit a dendritic spike or not, as shown in Figure 3. We consider *fast computations*, that is, a regime where the window in which spikes arrive is small, but not so small as to be equivalent to assuming perfect synchrony.

In this regime, each dendrite integrates over a time window and either reaches threshold or does not. Because we assume spike timing jitter in all inputs, each dendrite might reach threshold at different times for different presentations of the same nominal input. However, because dendritic spikes are sustained, jitter in the onset of these events across and entire cell has relatively little effect on whether the soma reaches threshold or not. This effect should confer spike timing robustness at the network level, which is the main claim we will test.

Before describing the implementation of the model and the results, we introduce an interpretation of the operating regime of the network that will be very useful. Computations occurring on short timescales can be interpreted as a binary computation, where incoming connections can be represented with a 1 (a spike arrives) or a 0 (absence of a spike), and the dendrite in turn produces a 1 (it fires) or a 0 (it does not). Connections between dendrites and soma are interpreted analogously: the dendrites produce 1s or 0s, and the soma sums these and compares the result to a firing threshold, thereby computing a 1 or 0. Interestingly, neurons and dendrites operating in this regime have been observed empirically, see e.g. [14, 63].

We used this binary interpretation in both a philosophical and practical sense in our spiking model. Philosophically, the binary analogy provides a clean intuition for how fast computations operate, with each unit only having sufficient time to spike once during a single "pass" or "cycle" of a computation. Such situations do appear in biology; see e.g. [61] for an example where neurons at each synaptic stage have about 10ms to process presynaptic spikes and fire during an animal's reaction time, leaving room for the receiving and firing of about one action potential per unit.

On a practical level the binary interpretation gives us a means for finding synaptic weights that allow us to train a network to perform a non-trivial computation, then test its robustness to timing jitter. We remind the reader that the focus of our investigation is not on the training or learning procedure, so the fact that we can train binary networks and use the weights in a dynamic, spiking network is not strictly relevant to the biology. However, it may give a powerful practical means for optimising hardware implementations of dynamic spiking networks. It may also hint that biological learning rules can operate in a somewhat equivalent manner.

We now outline the implementation. We built a Spiking Neural Network (SNN) where the individual neurons consist of our abstract neuron model. We constructed a separate Binary Neural Network (BNN) with the same number of equivalent units, and trained it using standard optimisation methods to perform a classification task (see: subsection 3.4 for details). The BNN is a static network where each unit's state is either 0 or 1. BNNs can thus be regarded as the saturated limit of regular sigmoidal networks, i.e. with weights of large absolute value [42]. As an aside we point out that BNNs are not restrictive computationally: any computable function can be approximated and implemented with a BNN [41, 58].

The task we train the BNN for is shown in Figure 4A. The 2D input points were first projected onto a binary feature space, to obtain 13D binary vectors. The dimensionality of 13 was chosen because this was the lowest dimensionality in which the binary network could still cope with the loss of information due to the binarization of the continuous coordinates. If the ith input of the binary vector was a 1, a randomly generated timepoint t_i was added to produce an input spike (i, t_i), meaning that input neuron i was fed an impulse so that it would spike at time t_i. If the ith element of the binary vector was 0, it meant that neuron i would not fire for that input vector.

The network architecture is set up so that each dendrite is connected to both a unique upstream neuron and a unique downstream soma (see Figure 4C for a sketch). The assumption that each neuron connects to one dendrite of an upstream neuron is actually grounded in physiology, although it may appear like a strong assumption at a first glance: related inputs arrive at local clusters of spines synchronously [60]. We have modelled these dendritic patches that synchronously get excited by correlated inputs as one dendritic compartment. When a sufficiently large number of dendritic compartments have been excited, the soma will spike. We have not explicitly accounted for inhibition in this model. Because our focus is to account for how transient signals can be summed and thresholded robustly, we are assuming that inhibition is implicitly accounted for in the lumped abstraction.

We transplanted the weight matrices of the BNN onto the spiking network, thereby obtaining a spiking network that can do the classification. When the input neurons all spike exactly simultaneously, the spiking network mimics the BNN exactly, i.e. the same units are active in both. But when asychrony is introduced, a discrepancy can arise. We introduce asynchrony in the network by moving the timing of the input spikes. Two examples of the spiking network receiving jittered input spikes are shown in Figure 4D and E. In Figure 4D the dendrites are furnished with plateau potentials, (as in Figure 2F, top). For these three input vectors, the network gives the correct answer despite the input spike jitter. In fact, the identity of the neurons spiking are still the same units that emit a 1 in its BNN counterpart.

This stands in contrast with the network performance without dendritic plateaus (Figure 4E), where the dendrites had no plateau potentials, (as in Figure 2F, bottom). In Figure 4E it can be seen that the network now fails to process two of the three input vectors correctly. The duration of the dendritic spikes is too short so that the dendritic spikes are separated in time, and the soma fails to sum them all. To test how quickly this leads to a degradation of performance, we tested the accuracy (as percentage of inputs classified correctly) of the network as the asynchrony increased. The network with active dendrites coped well, but the performance of the network without dendrites with plateau potentials degraded rapidly. This is quantified in Figure 4F, where we see that classification accuracy drops precipitously if spike timing jitter exceeds the membrane timeconstant significantly. In contrast, dendritic plateaus maintain performance even when spike jitter exceeds membrane timeconstant by an order of magnitude.

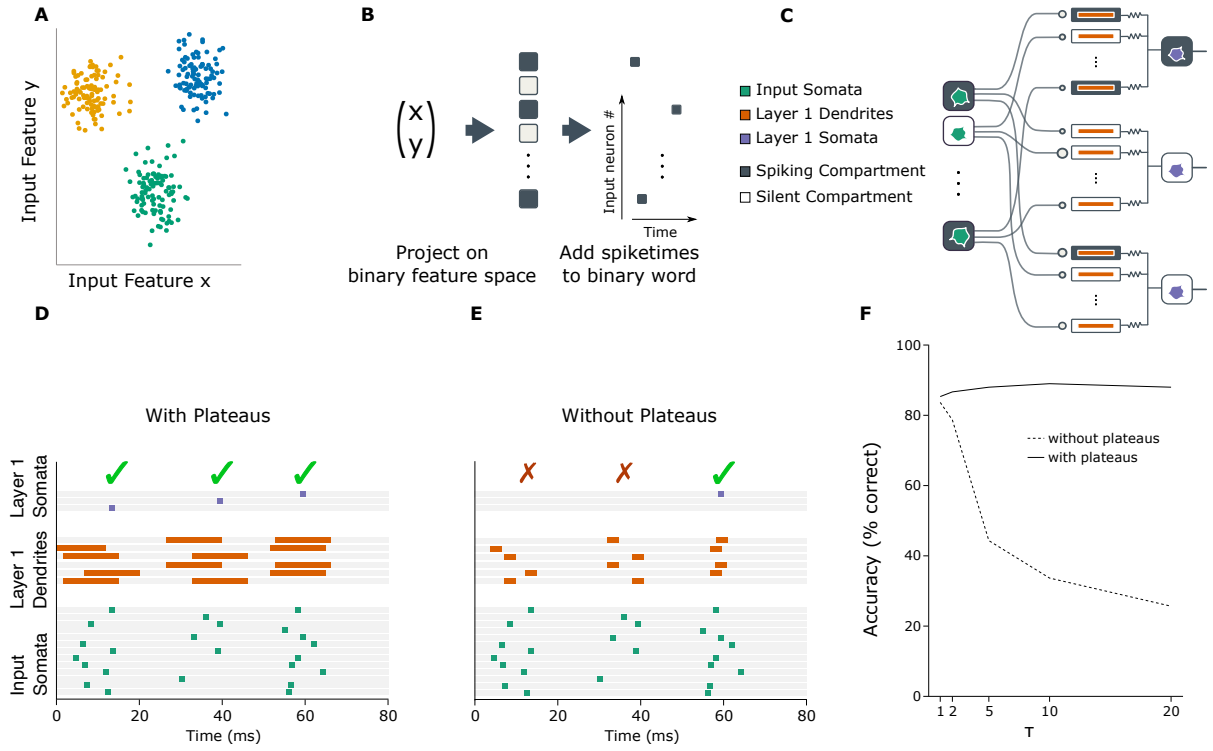

Figure 4: A simple conductance based model displays the same qualitative behavior as a detailed biophysical model. **(A)** The classification task performed by the spiking network of figures D, E, F. Each point is a 2D input vector x, the colors represent the different classes. **(B)** Procedure of transforming continuous 2D inputs x into input spikes for the network. First, x is projected onto a binary feature space to obtain a binary vector in a higher dimensional space. Then, spiketimes are added to this binary vector to produce the series of input spikes. For details, see subsection 3.4. **(C)** Schematic of the network architecture. The input somas spike according to the spiketimes obtained from the binary input vector. Each soma in the next layer has one dendrite per upstream soma, and each dendrite is connected to both one downstream and one upstream soma only. The dendrite-soma coupling is a bidirectional passive resistive coupling, whereas the upstream somas have a one-directional synaptic coupling onto the dendrites. **(D)** Example of the spiking network equipped with plateaus in the dendrites receiving asynchronous input spikes. It classifies the three inputs correctly in spite of the asynchrony. **(E)** Example of the spiking network equipped with dendrites without plateaus receiving asynchronous input spikes. The first two points are classified incorrectly, the network gets the third answer correct. **(F)** Summary of how well the network with plateaus and the network without plateaus deal with asynchrony τ, with the performance measured as the percentage of points of the classification task classified correctly. Without plateaus the performance drops off quickly, whereas the network with plateaus does not suffer from performance degradation for this range of τ.

Together, these results show in principle how a cellular mechanism that captures the essential abstract features of dendritic spikes can serve to enhance robustness of non-trivial spiking computations in a network. Furthermore, they provide an abstract interpretation of rapid spiking computations as binary neural network computations.

2 Discussion

An animal's survival often depends on its ability to make rapid decisions. Consequently, there will be evolutionary pressure for neural circuits to function at the most rapid timescale possible in some situations. For example, studies of primate visual reaction time estimate each neuron in the synaptic pathway has approximately 10 ms to make a firing decision, a time window allowing 1-2 spikes in each unit on average [62, 61]. This places the excitatory units in these pathways in an effectively binary regime: either a neuron fires once during the entire computation or it does not. We asked how cortical neurons might exploit dendritic nonlinearities to make such rapid computations feasible in the biologically realistic situation of spike timing jitter and signalling noise.

Traditionally thought to be passive cables, we now know that dendrites possess a zoo of voltage-gated ion channels that generate nonlinear membrane potential dynamics [59, 53, 38, 55, 32, 43]. As with axonal conduction, dendritic

excitability provides a means for signals to overcome spatial attenuation so it is perhaps not surprising to find regenerative currents in dendrites, particularly the long, thin dendrites of cortical neurons.

It is far less obvious why dendritic action currents are so much slower than their axonal counterparts. Their temporal dynamics, along with their nonlinear amplitude dependence opens numerous ways for neurons to process time-varying signals. For instance, the dendrites of pyramidal neurons can perform complex tasks such as the discrimination of temporal signals [7] or the detection of coincident inputs [36]. In parallel with providing rich signal processing capabilities, dendritic currents also shape activity-dependent synaptic plasticity dynamics, and may thus allow neural circuits to learn temporal patterns [29, 28, 23].

We considered a complementary role for dendritic action currents that is not in conflict with any of these ideas, yet it addresses an outlying problem we believe is essential: making rapid cortical computation robust. Conduction delays and noise make asynchrony unavoidable in communication between circuits in the brain [15]. This poses a fundamental problem for the integration of related inputs: neurons with short membrane time constants can only integrate coincident inputs that arrive simultaneously within $\sim$1 ms of one another. Here, we have shown that slow time constants, which are provided by NMDA depolarization events within dendritic branches, can remedy the situation by widening the integration time window of neurons.

Our hypothesis is consistent with several known experimental findings. It has been shown that the blocking of NMDA receptors impairs precise spike-timing asynchrony and inter-area communication [67]. This hints at an important role for NMDA in facilitating reliable synchronous communication between neuronal circuits. Recently it was shown that NMDA spikes are a major determinant of neuronal output in vivo, and that these dendritic spikes can be triggered by a handful of synaptic inputs [21, 44, 56, 9]. This is in line with the image we have sketched here, where NMDA spikes allow the network to perform computations with sparse spiking patterns.

We have been careful to respect the essence of basic physiological facts while trying to build an abstraction of how elementary spiking computations might occur. One conspicuous omission is to account for temporal variation in inhibition, which plays an important role in determining when and if spikes can fire in a network. We have two motivations for leaving this issue to one side in the present work: First, we wanted to isolate a mechanistic 'kernel' for dealing with spiking jitter in excitatory input by assuming that inhibition is present, and, at minimum, not making matters worse. In this setting one may interpret the excitatory inputs to the abstract model in the network (Figure 4) as a *net* drive in the presence of both inhibition and excitation. We feel this is reasonable because inhibitory signals in many local circuits reflect local population activity and often reliably track excitatory input [24]. Secondly, overwhelming evidence shows that inhibition itself plays an important role in enhancing synchrony in neural populations [12, 64]. We want to return to the question of integrating these features of the physiology in future work, our hypothesis being that integrate-and-hold can serve to improve computational robustness, as we have shown, and furthermore permit information to be preserved throughout the phase of prominent network level oscillations in the brain that are largely orchestrated by inhibition.

An alternative approach to building spiking computations that uses sparse spiking is for post-synaptic targets to become sensitive to specific, predictable, patterns of asynchronous spikes. Several computational studies have shown this is possible in principle using surrogate gradient learning rules that allow networks to perform computations based on relative spike-timings [46, 66]. However, these solutions are by design sensitive—rather than invariant—to the precise timing and order of inputs. It is therefore not clear that such solutions would work when networks are required to operate robustly on the fastest possible timescale.

Synchrony could potentially be maintained in networks that are organized as feed-forward "synfire chains", with relatively homogeneous transmission delays between nodes in each "rung" [1, 26]. [33] emphasize a role for refractoriness in maintaining synchrony, nothing that post-spike inhibition "clips" late inputs, thereby maintaining a localized packet in time. [30] explore further the importance of dendritic nonlinearities in stabilizing packet synchrony.

The significance of our work is to show that sparse yet reliable spiking computations may not require precisely synchronized inputs, and may aid in making computations robust when inputs are only partially synchronised. This perceived necessity has sometimes been raised as an objection to the possibility of spiking computations based on spike times [39]. Much effort has been devoted to finding out the extent to which neuronal noise can deteriorate the function of neural networks [5, 17, 68]. Sometimes out of practical necessity, these studies may assume a point model of a neuron [19]. Dendritic arbors and their dynamics may obviate some of the apparent fragility of spiking computations.

Other, simpler solutions to the asynchrony problem have been proposed. On possibility is that neural circuits exploit population averaging to overcome spike timing jitter [16, 57, 52, 35, 11]. However, this this is energetically costly [4] and would amount to scaling up the number of cells in a network to perform computations that could, in principle, be performed by more robust single units.

In contrast, we have argued that NMDA currents in distal dendrites can achieve robust, reliable threshold-based computation using relatively few resources. The longer duration of these potentials confers robustness to input timing without additional learning, reducing the number of neurons that must spike to achieve reliable signal transmission. Our key prediction is that neurons in some circuits use these dendritic potentials to make coincidence detection robust, allowing them to fire reliably despite input-timing variability larger than the membrane time constant. We would expect to find these mechanisms in circuits that detect or distinguish specific inputs rapidly, such as in the early stages of perception, or in circuits that coordinate long-range communication across multiple areas of the brain.

Fast action potentials allow rapid, massively parallel communication in the brain. In principle this gives spiking networks the ability to perform complex computations efficiently. However, decades of research has identified obstacles to implementing spiking computations under biologically realistic conditions and in hardware. Our work offers a possible means for rapid spiking computations to function robustly by providing a resettable temporal buffer in the input to each spiking unit. We would be excited to see experimental tests of whether dendrites do indeed operate in this way in the nervous system, and whether this simple principle offers a bio-inspired means to scale up reliable spiking computations in artificial neural networks.

Table 1: Standard values of parameters used

C^d	C^m	τ^r	τ^s	g_l	g_i	Θ^d	Θ^s	I^r_{max}	V_E
1 nF	1 nF	1 ms	1 ms	1 μS	1 μS	1 mV	1 mV	5 nA	100mV

3 Methods

3.1 Biophysical model

The biophysical model used to test the plausibility of the abstract model was taken from [18], and was implemented in NEURON (version 7.5) [27]. The model can be found on ModelDB (no. 249705). The model represented a layer 5 pyramidal neuron [2], and consisted of 85 compartments and 439 segments; 36 compartments for basal dendrites, 45 for apical dendrites, 3 for the soma, and one for the axon. A range of voltage-gated ion channels is included: sodium, A-type potassium, both high-voltage and low-voltage gated calcium, HCN, calcium-activated potassium, and K_v type channels are present.

The glutamate release is simulated at basal dendrites with the indices 14, 15, and 34. These indices were chosen for no particular reason; any collection of numbers would have worked. The glutamate stimuli were all given a fraction of 0.9 away from the soma (where 1 is the total length of the dendrite), and were separated in time from each other with a 50ms delay.

For Figure 2G, a passive inhibitory current was added to basal dendrite 34 in which an NMDA spike was triggered. The reversal potential of this current was that of $GABA_A$, i.e. -80 mV. The maximal conductance of this current was simulated as $g_{\text{inhib}} = \bar{g}_{\text{inhib}} w_{\text{inhib}}$, with $\bar{g}_{\text{inhib}}$ being the maximal conductance of 0.001 mS / cm^2, and w_{inhib} the weight that was varied during the simulation, such that $w = \{0.0, 1.0, 2.5, 2.0, 3.0, 3.5, 4.0\}$. In Figure 2G we plotted against the dimensionless quantity $g_{\text{inhib}}/g_{\text{NMDA}}$, where g_{NMDA} is 0.005 mS / cm^2.

3.2 Abstract model

The simplified model in Figure 2E,F is described by two differential equations for each dendritic branch, and two for the soma. The dynamics of the dendritic membrane potential V^d and somatic potential V^s are given by

$$C^d \dot{V}_i^{\text{dendrite}}(t) = -g_l V - (I_i^{\text{input}} + I_i^{sd}) \tag{1}$$

$$I_i^{sd} = g_i (V_i^{\text{dendrite}} - V^{\text{soma}}) \tag{2}$$

$$C^m \dot{V}^{\text{soma}}(t) = -g_l V^{\text{soma}} - I^r + \textstyle\sum_i I_i^{sd} \tag{3}$$

$$\tau^r \dot{I}_r(t) = -I^r \tag{4}$$

$$\tau^s \dot{s}(t) = -s \tag{5}$$

$$, \tag{6}$$

where g_l is the leak conductance, I_i^{input} is the current triggered by a spike arriving at dendrite i, I_i^{sd} is the current flowing between dendrite i and the soma, with g_i the conductance between the two, τ^x is the time constant of variable x, I^r is a refractory current, and s is the postsynaptic conductance of the neuron. When the dendrite reaches threshold Θ^{dendrite} the dendrite remains at threshold for P ms:

$$\begin{aligned} &\text{if } V^{\text{dendrite}} > \Theta^{\text{dendrite}} : \\ &\qquad \text{set } V^{\text{dendrite}} = \Theta^{\text{dendrite}} \text{ for } P \text{ ms.} \end{aligned} \tag{7}$$

When the soma reaches its threshold Θ^{soma} a spike is triggered:

$$\text{if } V^{\text{soma}} > \Theta^{\text{soma}} : \tag{8}$$

$$\text{set } V = 0 \tag{9}$$

$$\text{set } I^r = I^r_{\text{max}} \tag{10}$$

$$\text{set } s = 1, \tag{11}$$

i.e. the membrane potential is reset, a refractory current is activated, and the postsynaptic conductance increases.

Unless mentioned otherwise, the parameter values that were used in simulations are given in Table 1.

For Figure 2E,F, for the case with active dendrites, we used this abstract model, furnished with three dendritic compartments. The dendrites were given boxcar functions, with an amplitude of 50 and a width of 1ms, as input pulses:

$$I_i^{\text{input}}(t) = 50 \left(H(t - t_{\text{input},i}) - H\left(t - (t_{\text{input},i} + 1)\right) \right), \tag{12}$$

and $i \in \{1, 2, 3\}$. These input pulses were spaced 10 ms apart; $t_{\text{input},i} = 10i$ ms. Each pulse is strong enough by itself to trigger the plateau in each dendrite, thereby extending its duration. Taken together these three plateaus are strong enough to trigger a spike in the soma. For Figure 2E,F, for the case without dendrites, these pulses were fed to the soma directly. Due to the lack of plateaus they do not sum and fail to trigger a spike.

3.3 Neuron with asynchronous inputs

For Figure 3 the abstract model was used. The spike times of the incoming spikes were drawn from a normal distribution. The jitter of the spike times was defined as the standard deviation of this distribution. This jitter should always be compared with the membrane time constant τ of the neuron, for which we used 1 ms.

We used jitter values of 1τ for the synchronous case and 10τ for the asynchronous case respectively, i.e. $t_{\text{input}} \sim \mathcal{N}(0, \alpha\tau), \alpha \in \{1, 10\}$.

For all the computed values 500 simulations with a single neuron and randomly initialized spiketimes were performed. During each simulation, the neuron received 10 input spikes. For Figure 3B, the solid line is the mean voltage over the ensemble of simulations, and the shaded region is the inner 90% of the voltage distribution.

In Figure 3C,D,E all variables were computed inside the window of spikes arriving. The peak depolarization plotted is the mean of the ensemble of peak depolarizations for each simulation. Similarly, the coefficient of variation C_v plotted is the C_v for each jitter value, where that particular C_v is computed over all 500 simulations for that jitter value.

For Figure 3E the minimum number of input spikes was computed that would push C_v down to 0.1 for a particular value of the jitter value.

3.4 Binary network

The data consisted of 2D points which were assigned to different classes. The three classes were Gaussian clusters, with means $\mu_i = \{(4, 5), (-12, 2), (10, -7)\}$ respectively, and isotropic covariance of $\sigma = 1.5$. 100 points x_i per class i were generated from $x_i \sim \mathcal{N}(\mu_i, \sigma)$.

The task of the binary neural network (BNN) was to classify the points correctly. Because we wanted to interpret the continuous 2D points as input spikes to our network, we binarized the data first. To this end the input vectors $x \in \mathbb{R}^2$ were transformed by mapping them onto $\{0, 1\}^{13}$, i.e. every point was mapped onto a 13 dimensional vector of ones and zeros. This was achieved by randomly generating $W \sim \mathcal{N}_{13}(\mathbf{0}, I)$ and $b \sim \mathcal{N}(0, 1)$. Then we applied

$$x \mapsto \frac{1}{2}\left(\operatorname{sgn}(Wx + b) + 1\right).$$

To train our binary network we made use of surrogate gradients [46]. In short, we defined a function ϕ that for the forward pass (i.e. the network unit outputs) acted like a step function

$$\phi(x) = H(x).$$

But for the purposes of backpropagation the derivative of ϕ is defined to be

$$\phi'(x) = \frac{1}{4}\left(\left|\frac{x}{2}\right| + 1\right)^{-2},$$

which is the superspike derivative [65], equivalent to the derivative of a fast sigmoid. Using this derivative allows us to train the network, in spite of the step function having a derivative that is zero almost everywhere.

The network weights were initialized with the distribution:

$$W \sim U(0, x),$$

where $x = \sqrt{6/(\text{no. inputs} + \text{no. outputs})}$.

The biases of the BNN units were set to -1, to enforce a positive firing threshold for the neurons.

Each dendrite receives input from one presynaptic neuron. Let N_d be the number of dendrites per neuron for a layer, and N_o the number of somas. First an $N_d \times N_d$ diagonal matrix is constructed for each soma in the next layer. Then the total weight matrix will be the vertical concatenation of all these diagonal matrices.

Similarly, each dendrite is coupled with only one soma. This is achieved by constructing an $N_o \times N_d$ matrix for output unit i where only the ith row is non-zero. The overall weight matrix is the horizontal concatenation of these matrices.

The weight matrices used by the network alternate between the matrices with 'dendrites constraints' (a matrix that projects onto a layer of dendrites), and matrices with 'soma constraints' (a matrix projecting onto a layer of somas).

For both weight matrices elements that are initially zero, remain zero during the training procedure, and the elements of the weight matrices were constrained to be nonnegative, by having the activation of a unit in the BNN be

$$\phi\left(|W|x + b\right).$$

We implemented this constraint because we modelled pyramidal neurons, which are excitatory neurons. Therefore, the weights from the connections of these neurons should be positive.

The BNN was trained using stochastic gradient descent and surrogate gradient methods. First the network outputs $\hat{y}(x)$ where put through the softmax function

$$s(\hat{\mathbf{y}})_i = \frac{e^{\hat{y}_i}}{\sum_{j=1}^{K} e^{\hat{y}_j}} \tag{13}$$

which was then parsed through the cross-entropy function

$$H(\hat{\mathbf{y}}(\mathbf{x}), \mathbf{y}) = -\frac{1}{N}\sum_{i=1}^{B}\sum_{j=1}^{K} \log(s_j(\hat{\mathbf{y}}_i))\, y_{i,j}, \tag{14}$$

where $\hat{\mathbf{y}}_i$ is the output in response to input $\mathbf{x}_i$ belonging to the batch of size B and $y_{i,j}$ the jth element of the corresponding target.

The stochastic gradient descent was performed by randomly shuffling the inputs and labels and choosing a batch size (30 in our case). Then in each epoch all batches were iterated over, and the parameters were updated using the Adam algorithm. Training was terminated when the network gave the correct answer for more than 90% of all input points.

3.5 Spiking network

After having a set of weights that gave good performance with the BNN those weights were transplanted to a spiking network with the abstract models from subsection 3.2. This was done through setting the input currents to

$$I_i^{\text{input}} = w_{ij}^{\text{dendrite}} s_j\big(V_E - V_i^{\text{dendrite}}\big),$$

where $s_j(t)$ is the postsynaptic conductance of neuron j, which projects on dendrites i with weight w_{ij}^d. Here, w_{ij}^d is the nonzero element of the ith row of the weight matrix of this layer, trained with the 'dendrite constraints' (see subsection 3.4). This element is unique due to the constraints imposed on the weight matrices during training.

This dendrite is coupled to its soma with weight w_i^{sd} to give the current

$$w_i^{sd}\big(V_i^{\text{dendrite}} - V^{\text{soma}}\big),$$

with w_i^{sd} the (again unique) nonzero element of the ith row of the next weight matrix, trained with the 'soma constraints', and so on.

The 13 dimensional spikes were fed into the network by means of an input layer. This layer consisted of 13 neurons, one for each element in the binary input vectors. If the corresponding element was a 1, this neuron would be made to fire by giving them the same boxcar-shaped input current as described in subsection 3.2.

The degree of synchrony in the input spikes could then be controlled by varying the timing at which these impulses were given to the input neurons. Every input spike packet was centered around an input time t_{input}^0. Then the degree of asynchrony τ was varied by letting each spiketime, for each input neuron i, be

$$t_{\text{input}}^i \sim U(t_{\text{input}}^0 - \tau/2, t_{\text{input}}^0 + \tau/2).$$

To compare the performance of a spiking network without the dendritic plateaus to networks with plateaus, the 'hold' on the dynamics of V_i^d was removed. Therefore, the dendrite potential would decay immediately back to rest after reaching threshold. Otherwise, the architecture remained unchanged. For Figure 4D,E the input volleys were centered at 15ms, 35ms, and 55ms, with an asynchrony measure of $\tau = 10$ms.

An answer given by the spiking network was considered correct if neuron i in the output layer produced one spike when the input point belonged to class i, and the other neurons remained silent.

To compare the accuracies with and without plateaus in Figure 4F, the spiketimes of the input neurons were produced independently for each data point and each value of τ, with the spike packet being centered at 10 ms. Then the same spiketimes were used for both network versions, and the accuracy of the network was measured as the percentage of points classified correctly by the network. In each simulation, each network saw one input vector only, to prevent interference from multiple overlapping inputs for high jitter values.

4 Acknowledgements

This work was supported by ERC grant 716643 FLEXNEURO (TO), an Engineering and Physical Sciences Research Council DTP studentship (TJSB) and a Leverhulme and Isaac Newton Trust Fellowship (MER).

References

[1] ABELES, M. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, Feb. 1991.

[2] ACKER, C. D., AND ANTIC, S. D. Quantitative Assessment of the Distributions of Membrane Conductances Involved in Action Potential Backpropagation Along Basal Dendrites. *Journal of Neurophysiology 101*, 3 (Mar. 2009), 1524–1541.

[3] ANTIC, S. D., ZHOU, W.-L., MOORE, A. R., SHORT, S. M., AND IKONOMU, K. D. The decade of the dendritic NMDA spike. *Journal of Neuroscience Research 88*, 14 (2010), 2991–3001.

[4] ATTWELL, D., AND LAUGHLIN, S. B. An Energy Budget for Signaling in the Grey Matter of the Brain. *Journal of Cerebral Blood Flow & Metabolism 21*, 10 (Oct. 2001), 1133–1145.

[5] BARTOLO, R., SAUNDERS, R. C., MITZ, A. R., AND AVERBECK, B. B. Information-Limiting Correlations in Large Neural Populations. *Journal of Neuroscience 40*, 8 (Feb. 2020), 1668–1678.

[6] BERNANDER, O., DOUGLAS, R. J., MARTIN, K. A., AND KOCH, C. Synaptic background activity influences spatiotemporal integration in single pyramidal cells. *Proceedings of the National Academy of Sciences 88*, 24 (Dec. 1991), 11569–11573.

[7] BRANCO, T., CLARK, B. A., AND HÄUSSER, M. Dendritic Discrimination of Temporal Input Sequences in Cortical Neurons. *Science 329*, 5999 (Sept. 2010), 1671–1675.

[8] BRANCO, T., AND HÄUSSER, M. The single dendritic branch as a fundamental functional unit in the nervous system. *Current opinion in neurobiology 20*, 4 (2010), 494–502.

[9] BRANCO, T., AND HÄUSSER, M. Synaptic Integration Gradients in Single Cortical Pyramidal Cell Dendrites. *Neuron 69*, 5 (Mar. 2011), 885–892.

[10] BRODIN, L., TRAVEN, H. G., LANSNER, A., WALLEN, P., EKEBERG, O., AND GRILLNER, S. Computer simulations of N-methyl-D-aspartate receptor-induced membrane properties in a neuron model. *Journal of Neurophysiology 66*, 2 (Aug. 1991), 473–484.

[11] BUTTS, D. A., AND GOLDMAN, M. S. Tuning Curves, Neuronal Variability, and Sensory Coding. *PLOS Biology 4*, 4 (Mar. 2006), e92.

[12] BUZSAKI, G., AND DRAGUHN, A. Neuronal oscillations in cortical networks. *science 304*, 5679 (2004), 1926–1929.

[13] DESTEXHE, A., RUDOLPH, M., AND PARÉ, D. The high-conductance state of neocortical neurons in vivo. *Nature reviews neuroscience 4*, 9 (2003), 739–751.

[14] DEWEESE, M. R., WEHR, M., AND ZADOR, A. M. Binary Spiking in Auditory Cortex. *Journal of Neuroscience 23*, 21 (Aug. 2003), 7940–7949.

[15] FAISAL, A. A., SELEN, L. P. J., AND WOLPERT, D. M. Noise in the nervous system. *Nature Reviews Neuroscience 9*, 4 (Apr. 2008), 292–303.

[16] FÖLDIÁK, P. The 'Ideal Homunculus': Statistical Inference from Neural Population Responses. In *Computation and Neural Systems*, F. H. Eeckman and J. M. Bower, Eds. Springer US, Boston, MA, 1993, pp. 55–60.

[17] FRANKE, F., FISCELLA, M., SEVELEV, M., ROSKA, B., HIERLEMANN, A., AND AZEREDO DA SILVEIRA, R. Structures of Neural Correlation and How They Favor Coding. *Neuron 89*, 2 (Jan. 2016), 409–422.

[18] GAO, P. P., GRAHAM, J. W., ZHOU, W.-L., JANG, J., ANGULO, S., DURA-BERNAL, S., HINES, M., LYTTON, W. W., AND ANTIC, S. D. Local glutamate-mediated dendritic plateau potentials change the state of the cortical pyramidal neuron. *Journal of Neurophysiology 125*, 1 (Jan. 2021), 23–42.

[19] GERSTNER, W., AND KISTLER, W. M. *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press, Aug. 2002.

[20] GOAILLARD, J.-M., MOUBARAK, E., TAPIA, M., AND TELL, F. Diversity of axonal and dendritic contributions to neuronal output. *Frontiers in cellular neuroscience 13* (2020), 570.

[21] GOETZ, L., ROTH, A., AND HÄUSSER, M. Active dendrites enable strong but sparse inputs to determine orientation selectivity. *Proceedings of the National Academy of Sciences 118*, 30 (July 2021), e2017339118.

[22] GÓMEZ GONZÁLEZ, J. F., MEL, B., AND POIRAZI, P. Distinguishing Linear vs. Non-Linear Integration in CA1 Radial Oblique Dendrites: It's about Time. *Frontiers in Computational Neuroscience 5* (2011).

[23] GÜTIG, R., AND SOMPOLINSKY, H. The tempotron: A neuron that learns spike timing–based decisions. *Nature Neuroscience 9*, 3 (Mar. 2006), 420–428.

[24] HAIDER, B., DUQUE, A., HASENSTAUB, A. R., AND MCCORMICK, D. A. Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. *Journal of Neuroscience 26*, 17 (2006), 4535–4545.

[25] HEMBERGER, M., SHEIN-IDELSON, M., PAMMER, L., AND LAURENT, G. Reliable sequential activation of neural assemblies by single pyramidal cells in a three-layered cortex. *Neuron 104*, 2 (2019), 353–369.

[26] HERRMANN, M., HERTZ, J. A., AND PRÜGEL-BENNETT, A. Analysis of synfire chains. *Network: Computation in Neural Systems 6*, 3 (Jan. 1995), 403–414.

[27] HINES, M. L., AND CARNEVALE, N. T. The NEURON Simulation Environment. *Neural Computation 9*, 6 (Aug. 1997), 1179–1209.

[28] HISTED, M. H., AND MAUNSELL, J. H. Cortical neural populations can guide behavior by integrating inputs linearly, independent of synchrony. *Proceedings of the National Academy of Sciences 111*, 1 (2014), E178–E187.

[29] JACOBS, A. L., FRIDMAN, G., DOUGLAS, R. M., ALAM, N. M., LATHAM, P. E., PRUSKY, G. T., AND NIRENBERG, S. Ruling out and ruling in neural codes. *Proceedings of the National Academy of Sciences 106*, 14 (2009), 5936–5941.

[30] JAHNKE, S., MEMMESHEIMER, R.-M., AND TIMME, M. Propagating synchrony in feed-forward networks. *Frontiers in Computational Neuroscience 7* (2013), 153.

[31] JOHNSTON, D., MAGEE, J. C., COLBERT, C. M., AND CHRISTIE, B. R. Active properties of neuronal dendrites. *Annual review of neuroscience 19*, 1 (1996), 165–186.

[32] KIM, H. G., AND CONNORS, B. W. Apical dendrites of the neocortex: Correlation between sodium- and calcium-dependent spiking and pyramidal cell morphology. *Journal of Neuroscience 13*, 12 (Dec. 1993), 5301–5311.

[33] KISTLER, W. M., AND GERSTNER, W. Stable Propagation of Activity Pulses in Populations of Spiking Neurons. *Neural Computation 14*, 5 (May 2002), 987–997.

[34] KOCH, C., RAPP, M., AND SEGEV, I. A brief history of time (constants). *Cerebral cortex 6*, 2 (1996), 93–101.

[35] KOHN, A., COEN-CAGLI, R., KANITSCHEIDER, I., AND POUGET, A. Correlations and Neuronal Population Information. *Annual review of neuroscience 39* (July 2016), 237–256.

[36] LARKUM, M. E., ZHU, J. J., AND SAKMANN, B. A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature 398*, 6725 (Mar. 1999), 338–341.

[37] LOEWENSTEIN, Y., AND SOMPOLINSKY, H. Temporal integration by calcium dynamics in a model neuron. *Nature neuroscience 6*, 9 (2003), 961–967.

[38] LONDON, M., AND HÄUSSER, M. Dendritic Computation. *Annual Review of Neuroscience 28*, 1 (2005), 503–532.

[39] LONDON, M., ROTH, A., BEEREN, L., HÄUSSER, M., AND LATHAM, P. E. Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature 466*, 7302 (July 2010), 123–127.

[40] LOSONCZY, A., AND MAGEE, J. C. Integrative properties of radial oblique dendrites in hippocampal ca1 pyramidal neurons. *Neuron 50*, 2 (2006), 291–307.

[41] LUPANOV, O. B. Circuits using threshold elements. *Doklady Akademii Nauk 202*, 6 (1972), 1288–1291.

[42] MAASS, W., SCHNITGER, G., AND SONTAG, E. On the computational power of sigmoid versus Boolean threshold circuits. In *[1991] Proceedings 32nd Annual Symposium of Foundations of Computer Science* (San Juan, Puerto Rico, 1991), IEEE Comput. Soc. Press, pp. 767–776.

[43] MAJOR, G., LARKUM, M. E., AND SCHILLER, J. Active Properties of Neocortical Pyramidal Neuron Dendrites. *Annual Review of Neuroscience 36*, 1 (July 2013), 1–24.

[44] MAJOR, G., POLSKY, A., DENK, W., SCHILLER, J., AND TANK, D. W. Spatiotemporally Graded NMDA Spike/Plateau Potentials in Basal Dendrites of Neocortical Pyramidal Neurons. *Journal of Neurophysiology 99*, 5 (May 2008), 2584–2601.

[45] MAKARA, J. K., AND MAGEE, J. C. Variable dendritic integration in hippocampal ca3 pyramidal neurons. *Neuron 80*, 6 (2013), 1438–1450.

[46] NEFTCI, E. O., MOSTAFA, H., AND ZENKE, F. Surrogate Gradient Learning in Spiking Neural Networks. *arXiv:1901.09948 [cs, q-bio]* (May 2019).

[47] O'LEARY, T., AND WYLLIE, D. J. Single-channel properties of n-methyl-d-aspartate receptors containing chimaeric glun2a/glun2d subunits, 2009.

[48] POIRAZI, P., BRANNON, T., AND MEL, B. W. Pyramidal neuron as two-layer neural network. *Neuron 37*, 6 (2003), 989–999.

[49] POLACK, P.-O., FRIEDMAN, J., AND GOLSHANI, P. Cellular mechanisms of brain state–dependent gain modulation in visual cortex. *Nature neuroscience 16*, 9 (2013), 1331–1339.

[50] POLEG-POLSKY, A. Dendritic Spikes Expand the Range of Well Tolerated Population Noise Structures. *The Journal of Neuroscience 39*, 46 (Nov. 2019), 9173–9184.

[51] POLSKY, A., MEL, B. W., AND SCHILLER, J. Computational subunits in thin dendrites of pyramidal cells. *Nature neuroscience 7*, 6 (2004), 621–627.

[52] SANGER, T. D. Probability density estimation for the interpretation of neural population codes. *Journal of Neurophysiology 76*, 4 (Oct. 1996), 2790–2793.

[53] SCHILLER, J., MAJOR, G., KOESTER, H. J., AND SCHILLER, Y. NMDA spikes in basal dendrites of cortical pyramidal neurons. *Nature 404*, 6775 (Mar. 2000), 285–289.

[54] SCHILLER, J., AND SCHILLER, Y. NMDA receptor-mediated dendritic spikes and coincident signal amplification. *Current Opinion in Neurobiology 11*, 3 (June 2001), 343–348.

[55] SCHILLER, J., SCHILLER, Y., STUART, G., AND SAKMANN, B. Calcium action potentials restricted to distal apical dendrites of rat neocortical pyramidal neurons. *The Journal of Physiology 505*, 3 (1997), 605–616.

[56] SCHMIDT-HIEBER, C., TOLEIKYTE, G., AITCHISON, L., ROTH, A., CLARK, B. A., BRANCO, T., AND HÄUSSER, M. Active dendritic integration as a mechanism for robust and precise grid cell firing. *Nature Neuroscience 20*, 8 (Aug. 2017), 1114–1121.

[57] SEUNG, H. S., AND SOMPOLINSKY, H. Simple models for reading neuronal population codes. *Proceedings of the National Academy of Sciences 90*, 22 (Nov. 1993), 10749–10753.

[58] ŠÍMA, J., AND ORPONEN, P. General-Purpose Computation with Neural Networks: A Survey of Complexity Theoretic Results. *Neural Computation 15*, 12 (Dec. 2003), 2727–2778.

[59] SPENCER, W. A., AND KANDEL, E. R. Electrophysiology of hippocampal neurons: Iv. fast prepotentials. *Journal of Neurophysiology 24*, 3 (May 1961), 272–285.

[60] TAKAHASHI, N., KITAMURA, K., MATSUO, N., MAYFORD, M., KANO, M., MATSUKI, N., AND IKEGAYA, Y. Locally Synchronized Synaptic Inputs. *Science 335*, 6066 (Jan. 2012), 353–356.

[61] THORPE, S. J., AND IMBERT, M. Biological constraints on connectionist modelling. In *Connectionism in Perspective* (1989), Elsevier, pp. 63–92.

[62] VANRULLEN, R., AND THORPE, S. J. Is it a Bird? Is it a Plane? Ultra-Rapid Visual Categorisation of Natural and Artifactual Objects. *Perception 30*, 6 (June 2001), 655–668.

[63] WEI, D.-S., MEI, Y.-A., BAGAL, A., KAO, J. P. Y., THOMPSON, S. M., AND TANG, C.-M. Compartmentalized and Binary Behavior of Terminal Dendrites in Hippocampal Pyramidal Neurons. *Science 293*, 5538 (Sept. 2001), 2272–2275.

[64] WHITE, J. A., CHOW, C. C., RIT, J., SOTO-TREVIÑO, C., AND KOPELL, N. Synchronization and oscillatory dynamics in heterogeneous, mutually inhibited neurons. *Journal of computational neuroscience 5* (1998), 5–16.

[65] ZENKE, F., AND GANGULI, S. SuperSpike: Supervised Learning in Multilayer Spiking Neural Networks. *Neural Computation 30*, 6 (June 2018), 1514–1541.

[66] ZENKE, F., AND VOGELS, T. P. The remarkable robustness of surrogate gradient learning for instilling complex function in spiking neural networks. *bioRxiv* (June 2020), 2020.06.29.176925.

[67] ZICK, J. L., CROWE, D. A., BLACKMAN, R. K., SCHULTZ, K., BERGSTRAND, D. W., DENICOLA, A. L., CARTER, R. E., EBNER, T. J., LANIER, L. M., NETOFF, T. I., AND CHAFEE, M. V. Disparate insults relevant to schizophrenia converge on impaired spike synchrony and weaker synaptic interactions in prefrontal local circuits. *Current Biology 32*, 1 (Jan. 2022), 14–25.e4.

[68] ZYLBERBERG, J., POUGET, A., LATHAM, P. E., AND SHEA-BROWN, E. Robust information propagation through noisy neural circuits. *PLOS Computational Biology 13*, 4 (Apr. 2017), e1005497.