EVENT-BASED PATTERN DETECTION IN ACTIVE DENDRITES

PREPRINT

Johannes Leugering[*]
Osnabrück University, Germany
jleugeri@uni-osnabrueck.de

Pascal Nieters[*]
Osnabrück University, Germany
pnieters@uni-osnabrueck.de

Gordon Pipa
Osnabrück University, Germany
gpipa@uni-osnabrueck.de

August 7, 2020

ABSTRACT

Many behavioural tasks require an animal to integrate information on a slow timescale that can exceed hundreds of milliseconds. How this is realized by neurons with membrane time constants on the order of tens of milliseconds or less remains an open question. We show, how the interaction of two kinds of events within the dendritic tree, *excitatory postsynaptic potentials* and locally generated *dendritic plateau potentials*, can allow a single neuron to detect specific sequences of spiking input on such slow timescales. Our conceptual model reveals, how the morphology of a neuron's dendritic tree determines its computational function, which can range from a simple logic gate to the gradual integration of evidence to the detection of complex spatio-temporal spike-sequences on long time-scales. As an example, we illustrate in a simulated navigation task how this mechanism can even allow individual neurons to reliably detect specific movement trajectories with high tolerance for timing variability. We relate our results to conclusive findings in neurobiology and discuss implications for both experimental and theoretical neuroscience.

Introduction

The ability to detect long-lasting sequences of neural activity is crucial for complex behavior, but poses a serious challenge for most established neuron models. Consider a rodent navigating through an environment in search for food. Receptive fields of place and grid cells tile a spatial map of the environment and encode the current position by their respective population activities [1, 2]. But in order to find its way back, the animal needs to know not only its present location, but also which path it took to get there. Decoding this path from the sequential activation of place and grid cells requires the integration of information on behavioural timescales that can span hundreds of milliseconds or more [3, 4]. Relevant patterns on such long timescales may prove to be a ubiquitous phenomenon, and have already been documented for a wide range of sensory processing tasks, such as olfaction [5, 6] or cortical auditory processing [7].

This raises the puzzling question, how such long sequences of neural activity can be processed by volatile neurons with membrane time constants on the timescale of tens of milliseconds or less [8]. While this problem is typically addressed on a network level, e.g. by relying on effects of fast-acting synaptic plasticity [9] or slow emergent dynamics due to recurrent connections [10], we argue that it can be solved on the level of individual neurons by active processes within their dendritic trees. These localized processes endow neurons with internal memory traces on the timescale of hundreds of milliseconds, and can be captured in a simple, conceptual model that adheres to recent biological evidence not accounted for in integrate-and-fire neuron models.

By investigating the computational properties of neurons with active dendrites, we conclude that:

1. Active dendritic processes can implement complex spatio-temporal receptive fields for ordered sequences of synaptic inputs
2. Active dendritic processes enable the robust integration of weak signals over long time-scales.
3. When analyzed from a rate-coding perspective, active dendritic processes can produce sophisticated nonlinear computations determined by the neuron's dendritic morphology

[*]Both authors contributed equally.

We demonstrate these propositions in a general computational framework for event-based, active dendritic sequence processing (ADSP), which offers an elegant solution to the problem of detecting highly variable, long lasting patterns in a neuron's input.

Neural dynamics is driven by active dendritic processes

We derive our abstract model of dendritic computation from a few basic biological observations: Most of a cortical pyramidal neuron's excitatory synaptic inputs terminate on dendritic spines [11], where post-synaptic ion channels are activated via the stochastic, pre-synaptic release of glutamate-carrying vesicles [12, 13]. The activated channels, primarily controlled by α-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid receptors (AMPARs) [14], become conductive to a mixture of ions, which leads to a brief depolarization in the corresponding spine, referred to as the *excitatory post-synaptic potential* (EPSP) [15]. These voltage changes in nearby spines induce a modest depolarization in the local dendritic membrane potential [16], which passively propagates along the dendrite as described by neural cable theory (**Fig. 1c**). For very specific branching patterns, the passive propagation of activity along a neuron's dendrite can be simplified to an equivalent model of a cylinder, in which the contribution of individual synaptic inputs sum (sub-)linearly [17]. Since propagation along the cylinder is very fast, abstract point-neuron models such as leaky integrate-and-fire neurons ignore the spatial dimension of the dendritic tree entirely and model the neuron as if it were a single electric compartment [18]. However, in this purely passive model of dendritic integration, the attenuation of signals along the dendritic cable is so strong, that synaptic input onto thin apical dendrites should have little, if any, measurable effect on the membrane potential at the soma far away [19, 20]. This apparent problem could be resolved by a synaptic plasticity mechanism that proportional up-scales synaptic efficacies to compensate for the distance-dependent attenuation. This phenomenon, aptly termed "dendritic democracy"[21], has been shown in hippocampal pyramidal neurons [22], where it results in a similar contribution of synaptic inputs onto the somatic membrane potential — regardless of the synapse's position along the dendrite. We instead look at a different mechanism to boost weak synaptic inputs, which relies on localized depolarizations that are actively generated and maintained within the dendritic tree.

Such active dendritic processes are ubiquitous [23, 24] and largely rely on N-methyl-D-asparate receptor (NMDAR) gated ion-channels [14] (see **Fig. 1c** for a schematic representation of this mechanism). NMDAR gated channels, like their AMPAR gated counterparts, are activated in the presence of glutamate, but do not become conductive unless a channel-blocking Mg^+ ion is first displaced by a sufficiently strong depolarization [25, 26]. This depolarization can be achieved by the coactivation of multiple AMPAR channels on nearby spines within a short time-window. Experimental as well as simulation studies report that this requires a volley of 4-20 or even up to 50 spikes within 1-4ms, depending on the location along the dendritic tree [16, 27, 28, 29]. The opening of NMDAR channels triggers a massive influx of different ionic currents that lead to a complete depolarization of a small segment of the dendritic arbor. While the isolated NMDAR response itself is reported to last on the order of at least 25ms [30], in vivo recordings reveal that voltage-gated channels in the dendritic membrane [20] prolong this effect, resulting in a depolarization that can last from tens to hundreds of milliseconds [31]. We focus on these longer lasting events, which we collectively refer to as *dendritic plateau potentials*, and argue, that they provide useful memory traces within the dendritic tree that can last hundreds of milliseconds.

The much larger depolarization during a plateau potential propagates further along the dendrite than the weaker effect of individual EPSPs and thus extends the range at which they can contribute to somatic action potential generation. This may even be required for generating or spiking [32] or bursting [33] output. Just like EPSPs, however, plateau potentials are still subject to considerable attenuation along the dendritic cable and thus have a strong effect only in their direct neighbourhood[2]. This leads to a division of complex dendritic arbors into functional subunits [34, 35, 36], which we here refer to as *dendritic segments*. How local plateau potentials in these segments interact within a dendritic tree depends on its morphology. In particular, the depolarizing effect on other directly connected dendritic segments is effectively raising their resting potential for the whole duration of the plateau potential, thus lowering the amount of coinciding spikes required to initiate a plateau potential there [37]. As [38] demonstrates, this local nonlinear interaction of dendritic segments due to NMDAR-gated channels can allow neural dendrites to become selective to specific sequences of synaptic inputs. While their work uses a biophysical, spatially extended neuron model to explain this behaviour, we instead derive a much simplified model composed of discrete dendritic segments. This helps explain how local interactions between connected segments lead to cascades of plateau potentials, which in turn allow the detection of specific long-lasting sequences within the dendritic tree.

Each segment of a dendritic tree tends to receive strongly correlated volleys of spikes on clustered synaptic inputs from some subpopulation of neurons [39, 40]. We propose, that such incoming spike volleys constitute elementary events that convey relevant information. The morphology of the dendritic tree then determines how this information is processed and retained in memory, and thereby endows the ADSP neuron with an intricate computational function.

[2]Unlike EPSPs, this attenuation cannot be circumvented by synaptic scaling as for dendritic democracy.

The interaction of active dendritic processes realizes event-based computation.

Based on the biological observations in the previous section, we can derive an abstract mathematical model of active dendritic sequence processing: Conceptually, the complex dynamics of dendritic membrane potentials can be reduced to the interactions of two kinds of events, EPSPs and actively generated plateau potentials, in a tree structure of dendritic segments. Since both of these events result in localized stereotypical effects on the dendritic membrane potential, we abstractly model them as rectangular pulses of unit magnitude and fixed duration τ^{syn} and τ^{den}, respectively. The qualitative behaviour of the dendritic arbor can thus be explained purely in terms of the locations and times at which EPSPs and plateau potential are initiated in its dendritic segments.

Only those incoming spikes that are successfully transmitted by the probabilistic synapses induce EPSPs in the postsynaptic segment, which sum up and constitute the total *synaptic input* into the segment. This input is particularly strong when a *volley* of multiple spikes occurs in a time-window short enough for their EPSPs to overlap. In addition to synaptic input, the electric coupling between directly connected dendritic segments provides another source of *dendritic input*.

When both the *synaptic* and *dendritic input* into a segment exceed critical thresholds, the segment enters a prolonged *plateau* state. For the whole duration of the plateau, all other directly connected segments receive depolarizing dendritic input. Segments of the dendritic tree therefore act as coincidence detectors that respond to highly synchronized volleys of spikes with plateau potentials. The precise thresholds for synaptic and dendritic input depend on the segment's location within the dendritic tree. While a large volley of spikes alone suffices to trigger a plateau in the outermost segments of the dendritic tree, internal segments require the additional dendritic input due to plateau potentials in connected segments. For segments that lie at branching points in the dendritic tree, more than one of their neighbours may have to be in a plateau state concurrently to have a sufficient effect. If the soma, which lies at the root of the dendritic tree, receives sufficient synaptic and dendritic input, a somatic action potential, rather than a plateau potential, is generated.

Since the small effects of EPSPs remain confined to the postsynaptic dendrite segment, they can only affect the neuron's behaviour indirectly by contributing to the generation of local plateau potentials. It is the plateau potentials and their interaction across neighbouring segments that drives the dendritic membrane potential, and therefore implements an event-based framework of dendritic computation on two distinct timescales orders of magnitude apart. On a fast timescale, the combined effect of a volley of coincident spikes can initiate a localized plateau potential. On a much slower timescale, the interaction of these plateaus provides an ephemeral memory of the recent history. The computation we have described here can be fully formalized in terms of synaptic spikes and plateau events as provided in the Methods section.

In **Fig. 1** we describe an exemplary ADSP neuron that receives input from five populations of neurons on five segments (**Fig. 1a**). Each segment, if sufficiently excited, responds to a spike volley in its respective input populations by emitting a plateau event at the time of the volley (**Fig. 1b**). The morphology of the dendritic tree determines how these plateaus interact along the dendritic tree. For example, segment C will only activate if both segments A and B are already active once segment C receives a spike volley. We can formalize the relative timing requirement for these three segments by the expression $(A + B) \rightarrow_2 C$, which indicates that all two child branches A and B must be simultaneously active to enable the parent segment C, allowing it to emit a plateau in response to a spike-volley. We can read this as "A and B, and then C" (see also **Fig. 1d**). If the threshold was lowered, such that input from either segment A or B alone would suffice, the expression would correspondingly become $(A + B) \rightarrow_1 C$, which translates to "A or B, and then C" (see also **Fig. 1 e**). Generally, the expression $(X_1 + X_2 + \ldots + X_n) \rightarrow_m Y$ translates to "At least m out of the n segments $X_1, X_2, \ldots, X_n$ must be simultaneously active to enable segment Y". By chaining multiple segments together, these timing relations and nonlinear combinations can be arbitrarily nested, as for example in **Fig. 1f** that shows a neuron implementing $A \rightarrow_1 B \rightarrow_1 C$, which can be read as "A, and then B, and then C". Using this formal notation, we can express the complex ADSP neuron example in **Fig. 1a** as $(((A + B) \rightarrow_2 C) + D) \rightarrow_1 E$, a computation on spike volleys originating from the input populations associated with segments $A, \ldots, E$.

The interaction between connected dendritic segments facilitates cascades of plateau potentials along the dendritic tree, as illustrated in **Fig. 1b**. Starting in a distal segment, a leaf-node in our diagrams, a spike volley can initiate a plateau, which then provides dendritic input for the parent segment. Next, that segment responds to an incoming spike volley with a plateau of its own, in turn providing dendritic input to yet another segment. Whenever such a continuous chain of plateau potentials proceeds all the way to the soma, it culminates in a somatic action potential.

This signals to other neurons, that a specific sequence of spike volleys has been detected – on a timescale that may be as long as the number of segments times the plateau duration, i.e. hundreds of milliseconds. The precise timing between spike volleys is not prescribed exactly, as long as the distance between two successive volleys does not exceed the duration of one plateau potential.

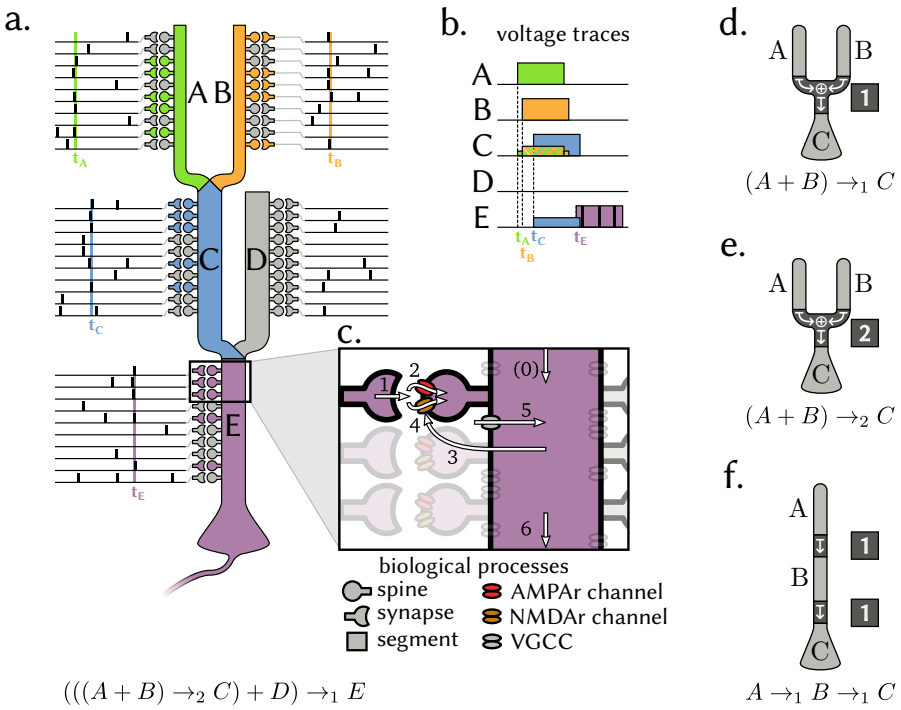

Figure 1: Schematic representation of a complex dendritic tree and its function. **a** A neuron receives on each of its 5 dendritic segments 10 synaptic connections from a corresponding neural population. Sufficiently many coincident spikes (here ≥ 6 out of 10) from population A can lead the corresponding dendritic segment to generate a plateau potential (t_A). Similarly, coincident spikes from population B can induce a plateau in a parallel branch (t_B). A third segment requires simultaneous input from both of these segments in addition to coincident synaptic input from population C, in order to fire a plateau of its own (t_C). On another branch, a fourth segment receives its input from population D but does not trigger a plateau. A somatic spike is triggered when coincident synaptic input from population E arrives (t_E) during dendritic input from either of its two upstream segments (in this case C). **b** Local membrane potentials show a cascade of plateau potentials. **c** The steps involved in the generation of a plateau: The membrane potential is already elevated due to a plateau potential in a neighbouring segment (0). Presynaptic input arrives at a synapse (1), which leads to a postsynaptic EPSP via AMPAr mediated ion channels (2). Once the local membrane potential is sufficiently depolarized due to coincident EPSPs and prior depolarization, voltage gated, NMDAr mediated ion channels open, causing additional depolarization (4) which can be further facilitated by the opening of voltage gated calcium channels (5). This strong depolarization initiates a longer lasting plateau potential in the dendritic segment, which has a modest depolarizing effect on other neighbouring segments (6). Different dendritic morphologies correspond to different computed functions, indicated in the respective formula under each schematic illustration. **d** If activating one of two dendritic branches with input from either population A or B, followed by a somatic spike initiated by input from population C, is sufficient to produce a spike, the neuron implements the operation $(A + B) \rightarrow_1 C$, which constitutes an "or"-operation between population A and B. **e** If simultaneous input from A and B is required, the neuron calculates an "and"-operation between inputs A and B. **f** A simple neuron that requires sequential activation of first A "and then" B before C.

The branching morphology of a dendritic tree therefore determines the computation performed by the neuron, which allows even single neurons to detect complex compositions of sequential patterns. This event-based computation is what we call active dendritic sequence processing (ADSP).

Results

Dendritic processing allows the rapid detection of long, time-invariant patterns

To demonstrate the practical implications of such neuronal sequence detection, we return to the example of a rat navigating an environment. We assume that the rat has an internal representation of its environment, tiled by the receptive fields of distinct populations of place cells. While the animal resides within such a receptive field, the corresponding population emits spike volleys with a magnitude that is largest when the animal is close to the center of the receptive field. Different paths lead the animal through some of these receptive fields in different order, and result in different sequences of spike volleys.

Each individual spike volley consists of several coincident spikes, the EPSPs of which have to be integrated and thresholded on a millisecond time-scale to detect sufficiently significant events in the presence of noise. To detect whether the animal has taken a specific path through the environment, only specific sequences of such significant spike volleys must be detected on a much slower behavioural time-scale. These two distinct time-scales pose a challenge for conventional spiking neuron models, which is further exacerbated by the fact, that the precise timing of the spike-volleys can vary substantially, depending e.g. on the speed with which the animal traverses its environment. While a solution to this problem may be found on a population level, we illustrate in **Fig. 2** how a single neuron can implement a solution very elegantly with just three active dendritic segments.

To simulate the rat's behaviour, we generate random movement trajectories through the environment by a stochastic process (see Methods section). Each place-cell population fires spike-volleys with a magnitude determined by the population's tuning-curve, a two-dimensional Gaussian function centered at the population's preferred location on a hexagonal grid. In this example, we are interested in paths that traverse three specific receptive fields, respectively color-coded in blue, orange and purple, and hence look at a neuron that consists of a chain of three dendritic segments, each receiving input from just one of these place-cell populations (**Fig. 2b**). The only trajectories that effectively drive the neuron to spike are those that sequentially traverse the three receptive fields in the correct order Blue $\rightarrow_1$ Orange $\rightarrow_1$ Purple (**Fig. 2a**).

During the example path shown in solid black, the three place cell populations are activated in the correct order over the course of 200ms and emit sufficiently large spike volleys to trigger a cascade of plateau potentials that lead the neuron to emit a somatic spike **Fig. 2b**. To illustrate how reliable of a detector an individual neuron can be — even when its synaptic inputs are stochastic with a transmission probability of 0.5 —, we systematically evaluate the probability of the neuron to fire in response to different paths with varying directions and lateral offsets. For an ideal straight 200ms long path through the center of all three place cell populations, the firing probability of the neuron is around 75%. When the orientation of the path is varied, this probability sharply decreases to 0%, indicating that the neuron is both highly sensitive and highly specific for paths with this orientation (**Fig. 2c**). Similarly, when the path is shifted orthogonally to the movement direction, the response probability falls quickly, confirming that the neuron is sensitive to the absolute location of the path as well as its direction (**Fig. 2d**).

A remarkable feature of this mechanism is, that it is invariant to changes in the precise timing of the individual volleys as long as two consecutive segments are activated within one plateau duration τ of each other. The ADSP Neuron can therefore detect paths of any duration from 0ms to $N\tau^{\text{den}}$ms, where $N = 3$ is the number of consecutive segments. We believe this source of timing-invariance to be a highly beneficial feature for generalization that might help explain phenomena, where the same sequence of events must be detected across multiple time-scales.

Plateaus integrate evidence on long time-scales

In the previous example, specific paths are recognized by memorizing the sequential activation of different neural populations on a slow behavioural time-scale. A seemingly different, yet in fact closely related problem is the integration of individually unreliable bits of evidence over time. Consider, for example, a population of neurons that extract some relevant feature of a stimulus, such as the local movement direction in a visual moving dots stimulus. If we assume a retinotopic mapping, neighbouring neurons are highly correlated, and whenever the local movement direction is apparent, we expect a couple of neighbouring neurons coding for that direction to produce a volley of spikes. However, these events are unlikely to occur at the exact same point in time throughout the entire input space. The decision, whether or not the visual flow is in a certain direction, therefore requires that a neuron can integrate many such pieces of evidence, each indicated by a spike volley event, over a longer time-scale. Despite the all-or-none response of dendritic plateaus, a neuron with sufficiently many dendritic segments can in fact approximate such a smooth integration of evidence on time-scales of hundreds of milliseconds!

We give an example of evidence integration using dendritic plateau potentials in a simplified experiment, in which a neuron with 1000 dendritic compartments receives input from a population of 1000 input neurons through a total of

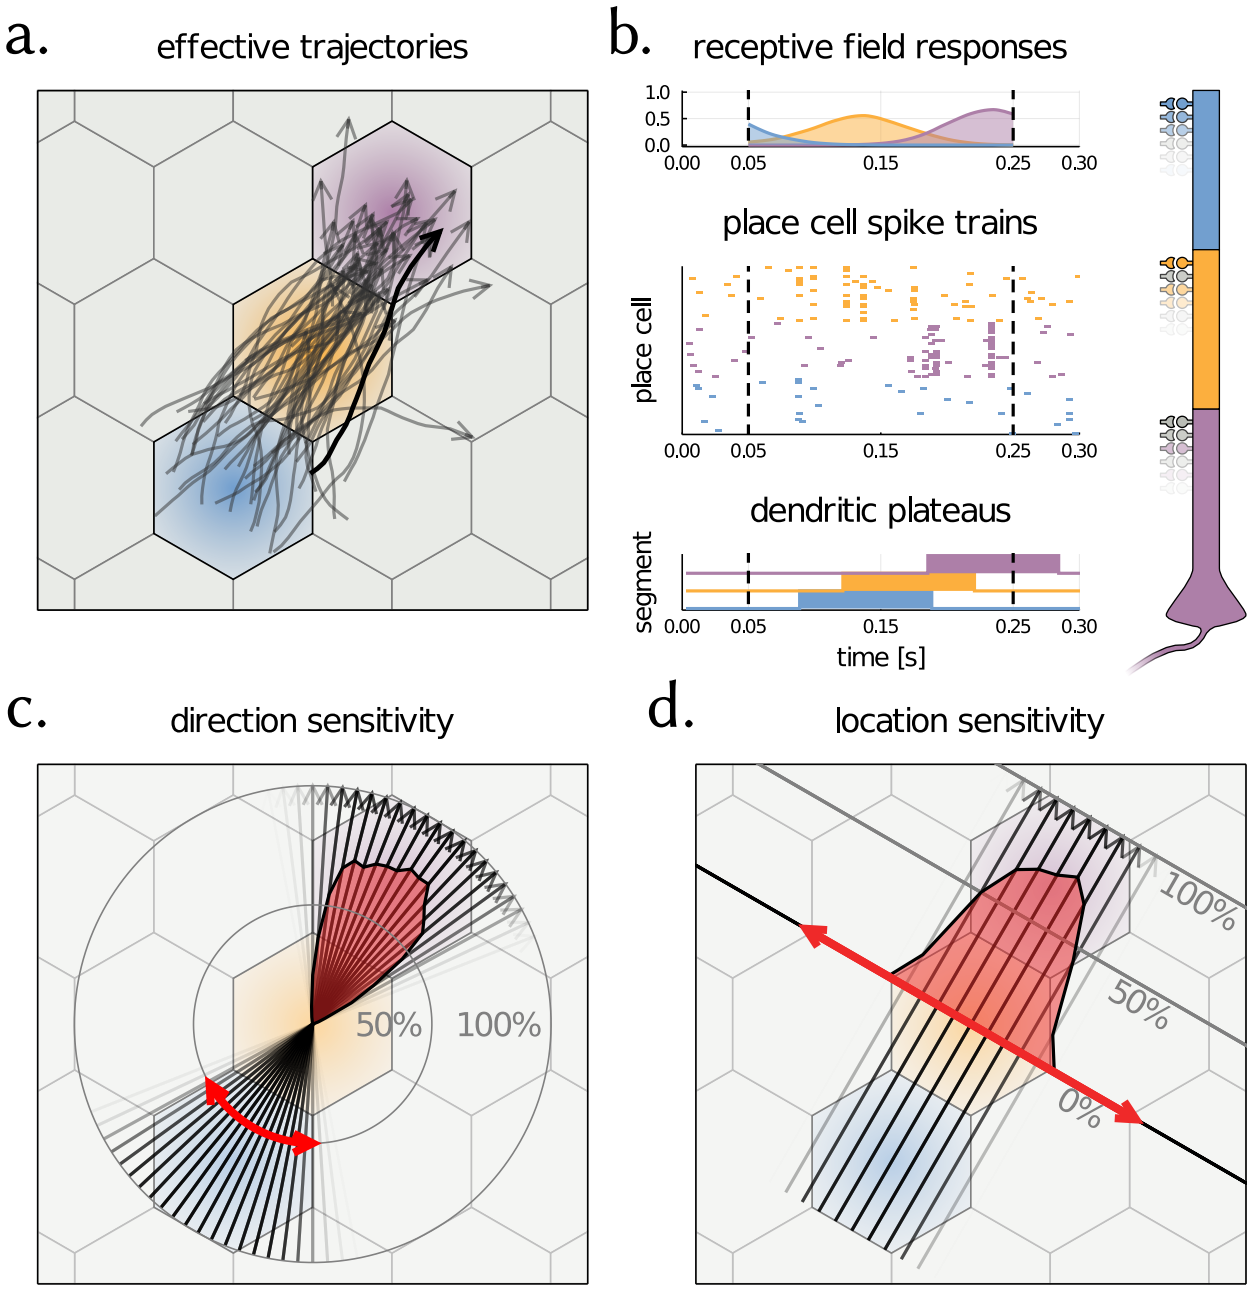

Figure 2: A simple neuron with three dendritic segments arranged as shown to the right of panel b can detect directed paths on a timescale of $300ms$. **a.** The receptive fields of place cell populations tile the environment through which the animal moves in a hexagonal grid. Random trajectories are generated through a stochastic process with randomized initial positions, velocities and angular heading to simulate the animal's movements. **b.** While the animal follows the black trajectory through space, the response of the place cell populations' tuning curves show the sequential activation of the populations over time (top panel). The generated spikes (middle panel) lead to a temporal sequence of dendritic plateaus (bottom panel) that results in a somatic spike. **c. and d.** The neuron responds with high probability to exactly those paths that traverse the desired receptive fields in the correct direction and with little lateral offset. The empirical firing probabilities to the black paths are shown by the superimposed density plot in red in polar and Cartesian coordinates, respectively.

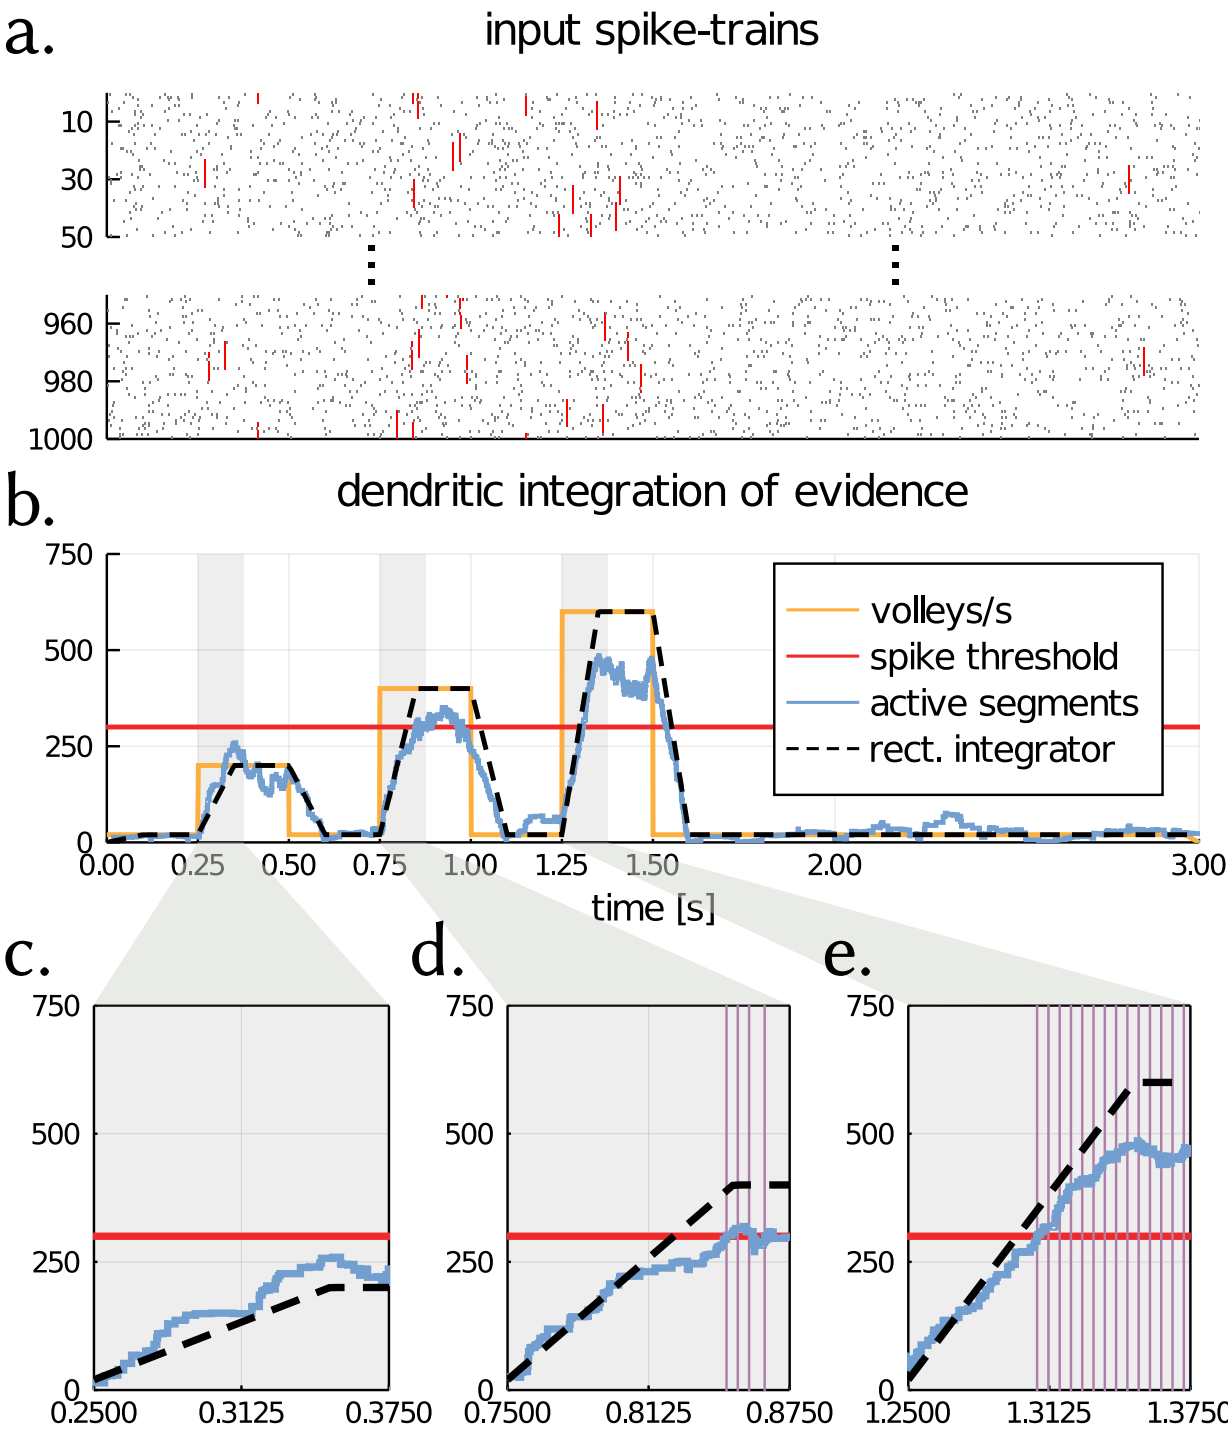

Figure 3: Dendritic plateaus can be used to gradually integrate evidence over long time periods. **a** A neuron with 1000 dendritic segments is driven by 1000 incoming spike-trains. Embedded in these spike-trains are spike volleys of 10 coincident spikes each, spread across 10 neighbouring neurons (shown in red). **b** The rate of spike volleys is determined by an input signal (organge line). Each segment receives input from 20 consecutive neurons through stochastic synapses with transmission probability $p = 0.5$, and requires 5 coincident spikes to trigger a plateau potential. The total number of co-activated dendritic segments (blue line) follows the convolution of the stimulus signal with a rectangular filter of length 100ms (black dashed line). **c-e** For increasing levels of stimulation, the number of co-activated segments rises faster and saturates at a higher level, crossing the threshold required for spike initiation (horizontal red line) at an earlier point in time or not at all, resulting in a sequence of spikes (vertical purple lines).

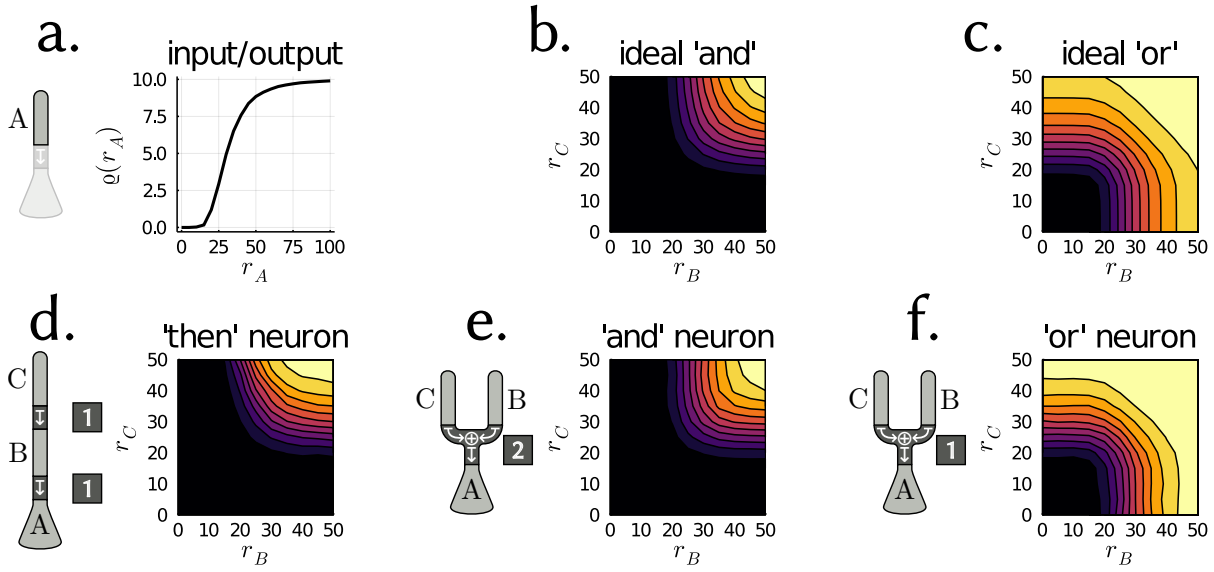

Figure 4: A rate-based analysis reveals well-known computational primitives. **a** A single dendritic compartment that receives independent Poisson-spike trains at a fixed rate r_A from a population of 25 neurons responds with plateaus at a rate that can be expressed as a non-linear sigmoidal function $\varrho(r_A)$. For multiple dendritic segments, each of which receives input from an identical but independent population A,B or C, the neuron's computation depends on the dendritic morphology. **d and e** If both segments C and B are required to enable a somatic spike, the neuron's firing rate is proportional an **b** idealized "and" operation between the two inputs. **f** If either of the two segments suffices, **c** the firing rate instead resembles an idealized "or" operation.

$20,000$ stochastic synapses (**Fig. 3**). The weak signal to be integrated by the ADSP neuron is encoded into spike volleys of 10 simultaneous spikes from adjacent neurons of the input population. Each dendritic segment of the ADSP neuron is connected to a different set of 20 adjacent neurons in the input population, and a total of 300 dendritic segments are required be in simultaneous plateau states for the neuron to emit a somatic spike.

Because each spike volley is likely to activate a different dendritic segment, we expect the number of simultaneously active dendritic compartments to reflect the average rate of incoming spike volleys during a time-interval of one plateau duration. This corresponds to a filtering of the time-varying rate by a rectangular filter, and, for a brief interval after stimulus onset, represents an ideal integrator. We observe this exact behavior by driving the rate, at which spike volleys are generated by the input population, to three different levels for brief time-intervals (**Fig. 3b**, orange line). The number of co-activated dendritic segments (blue line) closely follows the theoretical prediction of an ideal rectangular filter (black dashed line) until saturation. In particular, during the rising flanks right after stimulus onset (**Fig. 3c, d and e.**), we see the number of co-active segments rise with a slope proportional to the intensity of the stimulus until it saturates after 100ms. The neuron begins firing spikes once sufficiently many segments are active (red line). This is exactly the behavior expected for evidence integration: The ADSP neuron will fire sooner if the amount of evidence encoded in the stimulus is stronger, and will not fire at all if it remains sub-critical.

Interestingly, the stochasticity of synaptic transmission helps to further decorrelate the partially overlapping input to different dendritic segments, and can regulate the total amount of evidence required to reach the neuron's physiologically fixed spiking threshold. Also, while the example here makes use of just a single "layer" of dendritic segments directly driving the soma, this idea can be extended to deeper chains of multiple segments, such as in the previous example, to allow for the integration of evidence and non-linear combination thereof on time-scales even longer than one plateau duration.

Dendritic morphology determines computational function

In the two previous examples, we assume that each dendritic segment is driven by well-timed volleys of coincident spikes, the magnitudes of which represent the magnitude of an underlying signal. But in theoretical neuroscience, the function of a neuron is often analyzed in a rate-based framework, which relates only the average firing rate of a neuron to the average firing rates of its spiking inputs.

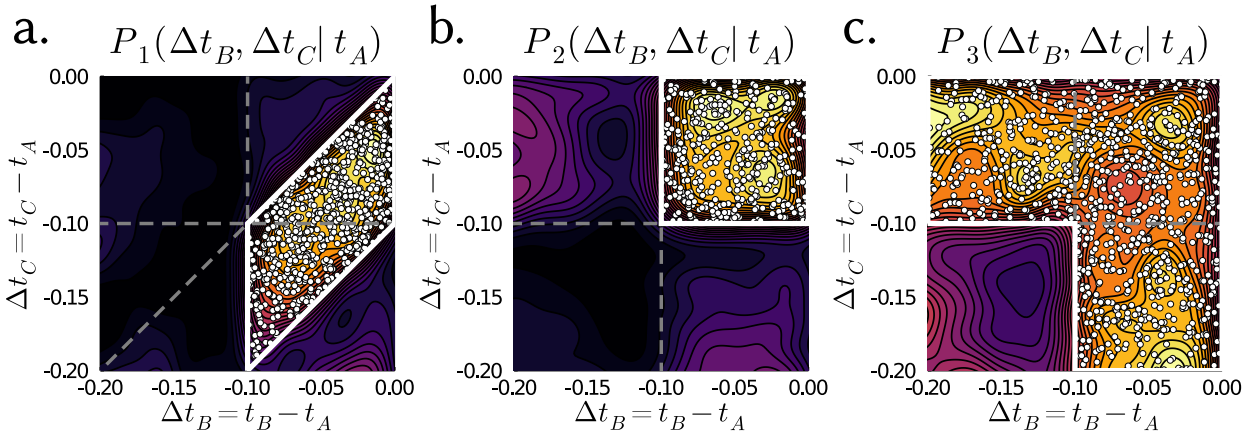

Figure 5: Dendritic morphology imposes timing constraints not revealed by rate-based analysis. For the neurons shown in figure 4, the joint probability distribution of relative timings Δt_B,Δt_C of dendritic plateaus directly preceding a somatic spike at t_A show a distinct temporal structure (contour-plots). **a** For the "then" neuron, a plateau in segment C must precede a plateau in segment B by at most 100ms, which in turn must occur at most 100ms before a somatic spike can be triggered. This is evident by the fact that all unambiguous cases, where exactly one plateau in each segment C and B was observed before a somatic spike, fall into the corresponding parallelogram-shaped domain (white dots). **b** This is in contrast to the "and" neuron, which despite showing a similar rate-response requires both inputs to occur within 100ms before the somatic spike. **c** The "or" neuron only requires either of the populations B or C to trigger a plateau within 100ms before a somatic spike.

Applying this sort of analysis to our proposed neuron model reveals, how different morphologies of dendritic arbors give rise to different non-linear computations. A dendritic segment driven by independent Poisson spike-trains originating from some population A of 25 neurons respond by triggering plateau potentials at a rate $\varrho(r_A)$ that continuously depend on the fixed firing-rate r_A of the populations' neurons. Here, 8 coincident spikes are required to trigger a plateau. As each plateau lasts for 100ms, ϱ saturates at a rate of 10 plateaus per second for large inputs (**Fig. 4a**). In more complex neurons composed of three dendritic segments, each of which is driven by an identical but independent population of neurons, we analyze the relative contributions of the populations B and C in the same way. In these experiments, we hold the firing rate $r_A = 25$ constant. For a neuron $C \to_1 B \to_1 A$, whose segments are sequentially chained together, a spike is generated if and only if both C and B are activated, and in the correct order. The resulting contour-plot, which shows how the output firing rate of this neuron scales with both r_C and r_B, illustrates that both a high firing rate of population C and B are required to result in a high firing rate of the neuron (**Fig. 4d**). This is similar to the neuron with two parallel segments $(C + B) \to_2 A$ (**Fig. 4e**), only that simultaneous activation of both segments, not sequential activation, is required. The shape of this function closely matches an idealized "and" operation (**Fig. 4b**), the firing rate of which can be derived as just the product of the rates at which plateaus are triggered in all dendritic segments:

$$\varrho(A, B, C) \propto \tau^{\text{den}} \varrho(r_A) f_{\text{and}}(B, C) \quad \text{where} \quad f_{\text{and}}(B, C) = \tau^{\text{den}\,2} \varrho(r_C) \varrho(r_B)$$

Here, $\varrho(A, B, C)$ is the firing rate of the neuron, and $f_{\text{and}}(B, C)$ is the factor due to the segments B and C.

For a different dendritic morphology $(C + B) \to_1 A$, where a plateau in either segment C or B is sufficient (**Fig. 4f**), we see a response that closely resembles an idealized "or" operation (**Fig. 4c**):[3]

$$f_{\text{or}}(B, C) \propto \tau^{\text{den}} \varrho(C) + \tau^{\text{den}} \varrho(B) - f_{\text{and}}(B, C)$$

For a derivation of f_{and} and f_{or} see the Methods section. This rate-based functional description offers a very useful abstraction of the neurons' behaviours, but it necessarily neglects questions of timing. As we saw in the previous sections, depending on the morphology, a dendritic arbor can impose stringent requirements on the order in which different segments can be activated. For example, while both neurons $C \to_1 B \to_1 A$ and $(C + B) \to_2 A$ require strong input from both input population B and C and hence show the same "and"-like response in the rate-coding

[3] As the last equation shows, referring to this operation as an "or" is justified in the sense that the resulting rate is proportional to the addition of the segments' individual plateau-firing-rates minus the "and" operation applied to both, which generalizes the Boolean operation to real values.

paradigm, the former imposes the constraint that the input from population C must arrive *before* that from population B while the latter does not. Rather than an "and"-like operation, neuron $C \rightarrow_1 B \rightarrow_1 A$ in fact implemented an "and then" operation. This is apparent when looking at the joint probability density of the relative timing of dendritic plateaus in the respective segments directly preceding a somatic spike (**Fig. 5**). In particular, if we only consider the unambiguous cases of one dendritic plateau each occurring in each segment within a brief window before a somatic spike (shown by the white dots (**Fig. 5**)), we observe that for neuron $C \rightarrow_1 B \rightarrow_1 A$, a dendritic plateau in segment B can occur at most 100ms before the somatic spike and is preceded by a dendritic plateau in segment C by at most another 100ms for a maximum total delay of 200ms. In contrast for neuron $(C + B) \rightarrow_2 A$, both segments must trigger a plateau within 100ms to elicit a somatic spike. For neuron $(C + B) \rightarrow_1 A$, a plateau in either segment within a 100ms window suffices to trigger a somatic spike.

Discussion

In this theoretical study we showed how a well-known biological phenomenon, dendritic plateau potentials, can drastically improve the computational capabilities of spiking neurons, turning them into powerful spatio-temporal pattern detectors. Due to the long-lasting memory provided by these plateau potentials, it becomes possible for individual neurons to integrate evidence or distinguish specific sequences of input on a timescale of hundreds of milliseconds – an order of magnitude larger than commonly observed membrane time constants [41]. In our model, the morphology of a neural dendrite determines its computational function and, when viewed in a conventional rate-coding paradigm, allows an individual neuron to implement a wide range of nonlinear behaviours in a modular and intuitive way.

This is in line with the two-layer neuron model proposed in [42], which used a detailed biophysical simulation of a pyramidal neuron to investigate the nonlinear effect on the neuron's firing rate due to synaptic input at different dendritic branches. Using a diverse array of stimuli, they showed that a two-layer network of sigmoidal subunits provides a substantially better approximation of the neuron's firing rate than a linear point-neuron. They speculated, however, that the prediction could be improved further, if the nonlinear interactions between the branches were considered, which we did here. We also investigated the use of dendritic plateau potentials as long-lasting memory traces, which our results revealed to be particularly important for evidence integration and the detection of temporal sequences. Remarkably, our drastically simplified and inherently event-based model could qualitatively reproduce properties of the model in [42], such as the sigmoidal input-output firing rate response of each dendritic segment and the linear-nonlinear combination thereof at the soma (see methods section).

But on the fast time-scale of individual spikes, our model differs substantially from this and other rate-based point-neuron models, since it relies on the detection of volleys of coincident spikes on a millisecond time-scale as the basic units of information, which are then integrated on the slower time-scale of dendritic plateau potentials. Our model is more closely related to recent work by [43], which proposed the use of active coincidence detection in dendritic segments to model prolonged effects of basal dendrites on the soma. A similar line of reasoning can also be found in [44], which presented a very elegant two-compartment neuron model and corresponding learning rule with one somatic and one dendritic compartment. Both models assign a specific functional role to the (basal) dendrite segments, namely to predict subsequent activation at the soma from their local synaptic inputs, which allows individual neurons to learn to predict state-transitions ("prospective coding"). Longer sequences are then detected by networks of such laterally connected neurons, endowing the networks with a form of temporal sequence-memory ("hierarchical temporal memory"). In our work, we have focused on a more mechanistic model that heavily relies on biological phenomena observed in single neurons. This allowed us to describe a neuron's computational capability concretely as that of a sophisticated pattern detector with long-lasting memory, and to illustrate how these mechanisms at play would appear under a rate based analysis. We believe our results offer a very appealing explanation of spike-based computation that has wider implications in neuroscience and raises several important questions, which we briefly discuss in the following:

What is the role of inhibition for dendritic computation?

Our model only takes into account excitatory synapses, but has clear implications for the role of inhibition. The all-or-none response of dendritic plateau potentials in our model implies that the only significant effect an inhibitory synapse can have on the far-away soma is by either reducing the likelihood of plateaus, preventing the generation of plateaus altogether, or by disrupting already ongoing plateau potentials. In the first two cases, an inhibitory synapse's post-synaptic potential must be either well-timed to coincide with the volley of excitatory spikes or exhibit a longer time-scale. Experiments suggest that inhibition can affect the ability of dendrites to generate active plateaus and prevent them [45]. The disruption of ongoing plateaus has also been reported and analyzed [46] and requires no such precise timing a-priori, as long as the spike occurs within the plateau's duration. Inhibition may, however, exhibit different effects depending on when during the plateau processes it is received. In all cases, the likely effect is shunting, rather

than substractive, inhibition.
Shunting inhibition can provide an efficient mechanism to improve the computational capabilities of the neurons described above, for example as it would allow individual neurons to exclusively respond to a sequence $a \rightarrow c$ but not to the sequence $a \rightarrow b \rightarrow c$, which is impossible for a neuron with purely excitatory synapses. Inhibition may therefore play an important and distinct role in ADSP neuron that warrants further investigation.

What are the implications of this model for plasticity?

We discussed a fundamental mechanism of dendritic computation and its capabilities, but did not cover the important topic of learning and plasticity. Nevertheless, the model presented here imposes constraints on potential plasticity mechanisms. Due to the long-lasting plateau potentials, a synaptic input can have a relevant causal effect for a somatic spike at a much later time. This makes the temporal assignment of credit for spiking outputs to synaptic inputs fundamentally difficult. The timing-invariance shown by our model and the dependency on the complex nonlinear dynamics within a dendritic tree further exacerbate this problem.
The most prominent example of synaptic learning is spike-time dependent plasticity [47], which tunes synaptic efficacy based on the relative timing of pre- and post-synaptic activity. Since the active dendritic processes discussed here both dominate the post-synaptic membrane potential as well as local Ca^{2+} concentration, they have a major effect on Hebbian plasticity [48, 49].
This is at odds with the common assumption, that backpropagating action potentials (bAPs) from the soma into the dendrite act as the primary post synaptic signal driving synaptic plasticity [50]. Since dendritic plateau potentials strongly depolarize dendrite segments for an extended period of time and should similarly "backpropagate" throughout the dendritic tree, it seems unlikely to us that bAPs are the primary factor for synaptic plasticity in neurons with active dendritic processes. Resolving this inconsistency is an important, but open research question.
Additionally, our model is based on binary stochastic synapses, and which segment the synapse terminates on plays a more important role than its efficacy. We therefore believe that structural plasticity mechanisms are particularly relevant for this kind of model. Furthermore, homeostatic plasticity mechanisms, e.g. scaling synaptic transmission probabilities[51], could be important here to ensure that only sufficiently large spike-volleys, but not randomly correlated inputs, can reliably trigger plateau potentials.

Is neuronal computation based on plateau processes?

Dendritic processes are thought to implement solutions to a number of specific computational problems in neurons [52], often distributed across many functional dendritic compartments [53, 54]. Based on convincing biological evidence for the mechanism of plateau generation and the interaction of such plateaus, we have argued that they are indeed the primary building block for the implementation of behaviorally highly relevant computations. How can this claim be experimentally verified or falsified?
Direct experimental verification, that computation in single neurons is well described by our proposed ADSP neuron model requires simultaneous measurement of synaptic inputs and local membrane potentials along a single neuron's dendrite on a fine temporal and spatial resolution over a long-time span.
As a first step, since our model is driven by incoming spike volleys from multiple intact neuron populations, *in vivo* measurements could verify the existence of patterns of spike-volleys over different time-scales using newly developed statistical techniques [55, 56].
Secondly, a key part of the model, the detection and integration of information across two timescales, one on the order of a few milliseconds, the other on the order of a hundred milliseconds or more, can be refuted for any type of neuron that achieves this without reliance on active dendritic processes. This may be the case either for neurons incapable of generating plateaus in the first place, or if plateau-generating processes have been pharmacologically disabled.
Thirdly, we predict single neurons that use active dendritic sequence processing to have spatio-temporal receptive fields on long temporal time-scales, but with high tolerance to variations in the precise timing of individual plateaus, qualitatively described in **Fig. 2**. Because of this invariance, we propose to go beyond linear analysis such as spike-triggered averages and instead measure both somatic response, as well as the timing of plateaus across the dendritic tree to find structures in the joint distributions as demonstrated in **Fig. 5**. Experimentally, spatio-temporal receptive fields of this kind could also be found by systematically varying stimuli, and should disappear when plateau-generating processes are disrupted.
While we have based our analysis on NMDAr-mediated plateaus in pyramidal cells [20], the same computational principle may be found in other neuron types, as well. For example, Purkinje cells in the cerebellum also generate localized Ca^{2+} events in response to coincident input on individual dendritic segments [57, 58], and thalamo-cortical neurons respond to strong synaptic input by localized plateaus in distal dendritic branches [59]. This indicates that the underlying ADSP mechanism, possibly implemented through diverse means in a case of convergent evolution, may be very general and ubiquitous in the brain.

Methods

Formal description of the event-based framework for computation in active dendrites

Mathematically, we approximate both EPSPs and plateau potentials by rectangular pulses with fixed duration τ^{syn} and τ^{den}, respectively. Here, we chose $\tau^{\text{syn}} = 5\text{ms}$ and $\tau^{\text{den}} = 100\text{ms}$ for all experiments if not stated otherwise. The dynamics of each dendritic segment can then be fully described in terms of the arrival times of incoming spikes as well as the times at which plateau potentials are initiated within the segment itself or in other directly connected segments. For some segment i, the synaptic input X_i and the dendritic input Y_i take the form of equations (1) and (2), respectively:

$$X_i(t) = \sum_{j \in S_i} \sum_k \chi_{i,j,k} \cdot \mathbf{1}_{[s_k^j, s_k^j + \tau^{\text{syn}}]}(t) \quad \text{where } \chi_{i,j,k} \sim \text{Bernoulli}(\omega_{i,j}) \tag{1}$$

$$Y_i(t) = \sum_{j \in D_i} \sum_k \mathbf{1}_{[t_k^j, t_k^j + \tau^{\text{den}}]}(t) \tag{2}$$

$$t_{m+1}^i = \min \left\{ t \in \mathbb{R} \;\middle|\; t \geq t_m^i + \tau^{\text{den}}, X_i(t) \geq \theta_i^{\text{syn}} \text{ and } Y_i(t) \geq \theta_i^{\text{den}} \right\}, \tag{3}$$

where $\mathbf{1}_{[a,b]}$ represents a unit pulse during the time interval $[a, b]$, and s_k^j and t_j^i are the times of spikes arriving from some presynaptic neuron j and the plateau onset times on segment i, respectively. The random variable $\chi_{i,j,k}$ represents the independent probabilistic transmission of every spike k from source j via a synapse to dendritic segment i, where the transmission occurs with the synapse specific probability $\omega_{i,j}$. The sets S_i and D_i respectively identify the segment's synaptic connections to other neurons and which other dendritic segments it is directly coupled to, and therefore reflect the morphology of the neuron's dendritic tree. Equation (3) states that, if the segment is not in a plateau state already, a new plateau is initiated as soon as both synaptic and dendritic inputs exceed their respective thresholds θ_i^{syn} and θ_i^{den}.

Implementation of the navigation experiments

To simulate the stochastic movements of a rat, random paths are generated with time-varying location $l(t) = (X(t), Y(t)) \in \mathbb{R}^2$ as solutions of the following system of stochastic differential equations:

$$\begin{aligned} dX &= \cos(2\pi A) V dt \\ dY &= \sin(2\pi A) V dt \\ dA &= 0.25 dW_A \\ dV &= 10.0(0.25 - V)dt + 0.1 dW_V \end{aligned}$$

A represents the angular heading of the animal, V represents its velocity in $\frac{m}{s}$ and W_A, W_V represent independent standard Brownian motion processes. Each path is generated with a randomized initial position within a rectangular domain of $10cm \times 9.5cm$, a random angular heading and a random velocity according to the marginal stationary distribution of V in the equation above, and is simulated for a fixed duration of $200ms$. Three populations of place cells, each 20 neurons strong, are centered on a hexagonal grid with center-to-center distance of $r \approx 2.9cm$. Each population randomly emits spike volleys following a homogeneous Poisson process with rate $\lambda = 50Hz$. The magnitude of each spike volley is determined by the population's mean activity at the time, which depends on the animal's location within the environment through a receptive field tuning curve. The tuning curves model the probability of each individual neuron within the population to participate in a given spike volley by the bell-curves $f_i(x) = \exp(-\frac{x-\mu_i}{2\sigma^2})$ with coefficient $\sigma = 9.7mm$, centered on the tiles of the hexagonal grid. The total number of spikes emitted during a volley from population i at time t is therefore a random variable distributed according to a Binomial distribution with population size $n = 20$ and probability $p = f_i(l(t))$. Additionally, each neuron in the population emits random spikes at a rate of $5Hz$ to emulate background activity. Each spike is transmitted through stochastic synapses independently with probability 0.5.

Each of the simulated neuron's dendritic segments receives spiking input from the 20 neurons of one population and requires at least 5 coincident spikes to trigger a plateau potential. The three segments are connected in a chain that requires sequential activation by spike volleys from the input populations in correct order to fire a spike. A random path is considered to be accepted by the neuron, if the neuron responds with a spike at any point in time during the corresponding simulation run.

To evaluate the rotation and location sensitivity of the neuron, we also generate straight paths with constant movement speed $v = \frac{3r}{200\text{ms}} \approx 43^{\text{cm}}/_{s}$ that are either rotated around the center of the environment by an angle α or offset from the center by a distance Δx orthogonal to the optimal movement direction. For each angle or offset, respectively, the empirical firing probability of the neuron in response to that path is estimated by simulating the path and the neuron's responses for 500 times each.

Implementation of the evidence-integration experiments

The input to the evidence-integrating neuron is generated by superimposing spike volleys onto 1000 independent Poisson processes with a constant firing rate of 10Hz. The volleys times are generated by a Poisson process with a time-varying rate $\lambda(t)$ representing the incoming "evidence". Here, $\lambda(t) = 200\text{Hz} \cdot (\mathbf{1}_{[0.25,0.5]}(t) + 2 \cdot \mathbf{1}_{[0.75,1.0]}(t) + 3 \cdot \mathbf{1}_{[1.25,1.5]}(t)) + 20\text{Hz}$. Each volley consists of simultaneous spikes from a randomly chosen set of ten input neurons with consecutive indices (wrapping around from 1000 to 1). Since each EPSP is assumed to last for a duration of 5ms, volleys and individual spikes are discarded if they occur less than 5ms after a preceding volley or spike. Each of the neuron's 1000 dendritic segments receives synaptic input via stochastic synapses with transmission probability 0.5 from 20 consecutive input neurons. As the number of input neurons and dendritic segments matches in this example, there is exactly one dendritic segment for every group of 20 consecutive input neurons, and each input neuron projects to exactly 20 dendritic segments. The total number of the neuron's synapses in this example is therefore 20000. Over time, the number of simultaneously active dendritic compartments as well as the times of generated somatic spikes is recorded. As a reference, the convolution $(\lambda \star \Pi)(t)$ of the time-varying rate-function λ with a rectangular filter Π of length 100ms and unit-integral is calculated.

Implementation of the rate-based analysis

For the rate-based analysis, four different neurons are constructed. First, a neuron consisting of a single dendritic compartment is driven by a total of 25 independent Poisson spike-trains with constant firing rate r_A. As in all other experiments, the duration of each spike is set to $\tau^{syn} = 5\text{ms}$, the duration of a plateau potential is set to $\tau^{den} = 100\text{ms}$. By systematically varying r_A and, for each choice, recording the number of plateau potentials generated during a simulation time-interval of 250s we can estimate the smooth function $\varrho(r_A)$, which relates the firing rate of the input population A to the resulting rate at which plateau potentials are generated.

For each of the three morphologies representing the $C \to_1 B \to_1 A$ neuron, the $(C + B) \to_2 A$ neuron and the $(C + B) \to_1 A$ neuron, we systematically vary the input firing rates of both populations B and C independently while keeping the firing rate of population A fixed at a constant 25Hz. For each combination, we again record the number of somatic spikes generated over a time-interval of 250s. As a reference for these two-dimensional functions, we use an idealized "and" and "or" function defined as:

$$f_{\text{and}}(B, C) = {\tau^{\text{den}}}^2 \varrho(r_C)\varrho(r_B) \tag{4}$$

$$f_{\text{or}}(B, C) = \tau^{\text{den}}\varrho(C) + \tau^{\text{den}}\varrho(B) - f_{\text{and}}(B, C) \tag{5}$$

$$= 1 - (1 - \tau^{\text{den}}\varrho(C))(1 - \tau^{\text{den}}\varrho(B)) \tag{6}$$

At a firing rate r_X, a segment driven by population X is in a plateau state at a given point in time with probability $\tau^{\text{den}}\varrho(r_X)$, therefore the probability that a segment driven by population C is active at the time that an input from population B arrives, which could in turn activate the next segment, is $\tau^{\text{den}}\varrho(r_C)$. The probability that this second segment is still active, when yet another volley from population A arrives to possibly trigger a somatic spike is also $\tau^{\text{den}}\varrho(r_B)$. Therefore the neuron's firing rate is proportional to ${\tau^{\text{den}}}^2\varrho(r_C)\varrho(r_B)$. Similarly, the probability that two parallel upstream segments driven by populations C and B are simultaneously active at a given point in time is ${\tau^{\text{den}}}^2\varrho(r_C)\varrho(r_B)$. In contrast, the probability that either upstream segment is active at a given point in time is just the probability that not both are simultaneously inactive, i.e. $1 - (1 - \tau^{\text{den}}\varrho(C))(1 - \tau^{\text{den}}\varrho(B))$. This expression has the nice alternative form $c + b - cb$, where $c = \tau^{\text{den}}\varrho(C)$, $b = \tau^{\text{den}}\varrho(B)$ and $cb = f_{\text{and}}(B, C)$, which generalizes the Boolean "or" operation to real-valued firing rates. When identifying $true$ with 1 and $false$ with 0, the truth-table of this expressions matches that of the logic expression "c or b".

To evaluate timing requirements for each of these three neuron morphologies, we run another simulation at constant input rates $r_A = r_B = r_C = 25\text{Hz}$ for a duration of 1h of simulated time. We record the time of each plateau-initiation-event in both upstream segments driven by population C and B for a time-interval of 200ms preceding each somatic spike. If there is exactly one plateau-event from each segment in such a time-interval, we record this as an *unambiguous* pair of plateau events. If there is more than one plateau-event on either of the dendritic segments, we record all pairs of

plateau-events in that time-interval composed of one plateau event for each segment. We refer to these latter pairs as *ambiguous*. Using these ambiguous pairs, we estimate the joint probability distribution $P_i(\Delta t_B, \Delta t_C | t_A)$ over relative times Δt_B and Δt_C between a plateau triggered by population B or C and a somatic spike triggered at time t_A by population A. For a more reliable estimate of the timing constraints, we consider only the unambiguous pairs, which evidently fall into distinct domains of these joint probability distributions that uniquely characterize the precise timing requirements of the respective neuron morphologies. This can be seen in figure 5. E.g. for the $C \to_1 B \to_1 A$ neuron, all plateaus triggered by population C must precede those triggered by B, but cannot precede them by more than one plateau duration of 100ms, therefore they fall into a parallelogram below the diagonal. For the $(C + B) \to_2 A$ neuron, on the other hand, both plateau events must independently occur within 100ms before a somatic spike, and hence fall into the upper quadrant of the joint density.

Code availability

All simulations are implemented in a custom developed package in the Julia programming language [60], publicly available via the code repository hosted at https://github.com/jleugeri/ADSP.jl. Further documentation of the simulator and implementation details can be found there.

References

[1] J O'Keefe and J Dostrovsky. The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain Res.*, 34(1):171–175, November 1971.

[2] Torkel Hafting, Marianne Fyhn, Sturla Molden, May-Britt Moser, and Edvard I Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806, August 2005.

[3] Martin Stemmler, Alexander Mathis, and Andreas V M Herz. Connecting multiple spatial scales to decode the population activity of grid cells. *Sci Adv*, 1(11):e1500816, December 2015.

[4] Howard Eichenbaum. On the integration of space, time, and memory. *Neuron*, 95(5):1007–1018, August 2017.

[5] Brice Bathellier, Derek L Buhl, Riccardo Accolla, and Alan Carleton. Dynamic ensemble odor coding in the mammalian olfactory bulb: sensory information at different timescales. *Neuron*, 57(4):586–598, February 2008.

[6] Bede M Broome, Vivek Jayaraman, and Gilles Laurent. Encoding and decoding of overlapping odor sequences. *Neuron*, 51(4):467–482, August 2006.

[7] Huan Luo and David Poeppel. Phase patterns of neuronal responses reliably discriminate speech in human auditory cortex. *Neuron*, 54(6):1001–1010, June 2007.

[8] Ofer Melamed, Wulfram Gerstner, Wolfgang Maass, Misha Tsodyks, and Henry Markram. Coding and learning of behavioral sequences. *Trends Neurosci.*, 27(1):11–4; discussion 14–5, January 2004.

[9] Gianluigi Mongillo, Omri Barak, and Misha Tsodyks. Synaptic theory of working memory. *Science*, 319(5869):1543–1546, March 2008.

[10] Yulia Sandamirskaya and Gregor Schöner. An embodied account of serial order: how instabilities drive sequence generation. *Neural Netw.*, 23(10):1164–1179, December 2010.

[11] C Beaulieu and M Colonnier. A laminar analysis of the number of round-asymmetrical and flat-symmetrical synapses on spines, dendritic trunks, and cell bodies in area 17 of the cat. *J. Comp. Neurol.*, 231(2):180–189, January 1985.

[12] Katz. *The Release of Neural Transmitter Substances (The Sherrington Lectures)*. Liverpool University Press, December 1969.

[13] C F Stevens. Quantal release of neurotransmitter and long-term potentiation. *Cell*, 72 Suppl:55–63, January 1993.

[14] Michael Hollmann and Stephen Heinemann. Cloned glutamate receptors. *Annual review of neuroscience*, November 2003.

[15] JC Watkins and RH Evans. Excitatory amino acid transmitters. *Annual review of pharmacology and toxicology*, 21(1):165–204, 1981.

[16] Attila Losonczy and Jeffrey C Magee. Integrative properties of radial oblique dendrites in hippocampal CA1 pyramidal neurons. *Neuron*, 50(2):291–307, April 2006.

[17] W Rall. Electrophysiology of a dendritic neuron model. *Biophys. J.*, 2(2 Pt 2):145–167, March 1962.

[18] A N Burkitt. A review of the integrate-and-fire neuron model: I. homogeneous synaptic input. *Biol. Cybern.*, 95(1):1–19, July 2006.

[19] Greg Stuart and Nelson Spruston. Determinants of voltage attenuation in neocortical pyramidal neuron dendrites. *Journal of Neuroscience*, 18(10):3501–3510, 1998.

[20] Nelson Spruston. Pyramidal neurons: dendritic structure and synaptic integration. *Nat. Rev. Neurosci.*, 9(3):206–221, March 2008.

[21] Michael Häusser. Synaptic function: dendritic democracy. *Current Biology*, 11(1):R10–R12, 2001.

[22] Jeffrey C Magee and Erik P Cook. Somatic epsp amplitude is independent of synapse location in hippocampal pyramidal neurons. *Nature neuroscience*, 3(9):895–903, 2000.

[23] Srdjan D Antic, Wen-Liang Zhou, Anna R Moore, Shaina M Short, and Katerina D Ikonomu. The decade of the dendritic NMDA spike. *J. Neurosci. Res.*, 88(14):2991–3001, November 2010.

[24] Katerina D Oikonomou, Mandakini B Singh, Enas V Sterjanaj, and Srdjan D Antic. Spiny neurons of amygdala, striatum, and cortex use dendritic plateau potentials to detect network UP states. *Front. Cell. Neurosci.*, 8:292, September 2014.

[25] H Monyer, N Burnashev, D J Laurie, B Sakmann, and P H Seeburg. Developmental and regional expression in the rat brain and functional properties of four NMDA receptors. *Neuron*, 12(3):529–540, March 1994.

[26] T Götz, U Kraushaar, J Geiger, J Lübke, T Berger, and P Jonas. Functional properties of AMPA and NMDA receptors expressed in identified types of basal ganglia neurons. *J. Neurosci.*, 17(1):204–215, January 1997.

[27] Sonia Gasparini, Michele Migliore, and Jeffrey C Magee. On the initiation and propagation of dendritic spikes in CA1 pyramidal neurons. *J. Neurosci.*, 24(49):11046–11056, December 2004.

[28] Sonia Gasparini and Jeffrey C Magee. State-dependent dendritic computation in hippocampal CA1 pyramidal neurons. *J. Neurosci.*, 26(7):2088–2100, February 2006.

[29] Jacopo Bono and Claudia Clopath. Modeling somatic and dendritic spike mediated plasticity at the single neuron and network level. *Nat. Commun.*, 8(1):706, September 2017.

[30] Paul Rhodes. The properties and implications of NMDA spikes in neocortical pyramidal cells. *J. Neurosci.*, 26(25):6704–6715, June 2006.

[31] Guy Major, Matthew E Larkum, and Jackie Schiller. Active properties of neocortical pyramidal neuron dendrites. *Annu. Rev. Neurosci.*, 36:1–24, July 2013.

[32] Tim Jarsky, Alex Roxin, William L Kath, and Nelson Spruston. Conditional dendritic spike propagation following distal synaptic activation of hippocampal CA1 pyramidal neurons. *Nat. Neurosci.*, 8(12):1667–1676, December 2005.

[33] Christine Grienberger, Xiaowei Chen, and Arthur Konnerth. Nmda receptor-dependent multidendrite ca2+ spikes required for hippocampal burst firing in vivo. *Neuron*, 81(6):1274–1281, 2014.

[34] C Koch, T Poggio, and V Torre. Retinal ganglion cells: a functional interpretation of dendritic morphology. *Philos. Trans. R. Soc. Lond. B Biol. Sci.*, 298(1090):227–263, July 1982.

[35] Alon Polsky, Bartlett W Mel, and Jackie Schiller. Computational subunits in thin dendrites of pyramidal cells. *Nat. Neurosci.*, 7(6):621–627, June 2004.

[36] Tiago Branco and Michael Häusser. The single dendritic branch as a fundamental functional unit in the nervous system. *Curr. Opin. Neurobiol.*, 20(4):494–502, August 2010.

[37] Guy Major, Alon Polsky, Winfried Denk, Jackie Schiller, and David W Tank. Spatiotemporally graded NMDA spike/plateau potentials in basal dendrites of neocortical pyramidal neurons. *J. Neurophysiol.*, 99(5):2584–2601, May 2008.

[38] Tiago Branco, Beverley A Clark, and Michael Häusser. Dendritic discrimination of temporal input sequences in cortical neurons. *Science*, 329(5999):1671–1675, 2010.

[39] Matthew E Larkum and Thomas Nevian. Synaptic clustering by dendritic signalling mechanisms. *Curr. Opin. Neurobiol.*, 18(3):321–331, June 2008.

[40] Naoya Takahashi, Kazuo Kitamura, Naoki Matsuo, Mark Mayford, Masanobu Kano, Norio Matsuki, and Yuji Ikegaya. Locally synchronized synaptic inputs. *Science*, 335(6066):353–356, January 2012.

[41] Debora Ledergerber and Matthew Evan Larkum. Properties of layer 6 pyramidal neuron apical dendrites. *Journal of Neuroscience*, 30(39):13031–13044, 2010.

[42] Panayiota Poirazi, Terrence Brannon, and Bartlett W Mel. Pyramidal neuron as two-layer neural network. *Neuron*, 37(6):989–999, 2003.

[43] Jeff Hawkins and Subutai Ahmad. Why neurons have thousands of synapses, a theory of sequence memory in neocortex. *Frontiers in neural circuits*, 10:23, 2016.

[44] Johanni Brea, Alexisz Tamás Gaál, Robert Urbanczik, and Walter Senn. Prospective coding by spiking neurons. *PLOS Computational Biology*, 12(6):1–25, 06 2016.

[45] Monika Jadi, Alon Polsky, Jackie Schiller, and Bartlett W Mel. Location-dependent effects of inhibition on local spiking in pyramidal neuron dendrites. *PLoS Comput Biol*, 8(6):e1002550, 2012.

[46] Michael Doron, Giuseppe Chindemi, Eilif Muller, Henry Markram, and Idan Segev. Timed synaptic inhibition shapes nmda spikes, influencing local dendritic processing and global i/o properties of cortical neurons. *Cell reports*, 21(6):1550–1561, 2017.

[47] Sen Song, Kenneth D Miller, and Larry F Abbott. Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nature neuroscience*, 3(9):919–926, 2000.

[48] John Lisman and Nelson Spruston. Postsynaptic depolarization requirements for LTP and LTD: a critique of spike timing-dependent plasticity. *Nat. Neurosci.*, 8(7):839–841, July 2005.

[49] Jason Hardie and Nelson Spruston. Synaptic depolarization is more effective than back-propagating action potentials during induction of associative long-term potentiation in hippocampal pyramidal neurons. *J. Neurosci.*, 29(10):3233–3241, March 2009.

[50] Nicola Kuczewski, Cristophe Porcher, Volkmar Lessmann, Igor Medina, and Jean-Luc Gaiarsa. Back-propagating action potential. *Communicative & Integrative Biology*, 1(2):153–155, 2008. PMID: 19704877.

[51] Gina Turrigiano. Homeostatic synaptic plasticity: local and global mechanisms for stabilizing neuronal function. *Cold Spring Harbor perspectives in biology*, 4(1):a005736, 2012.

[52] Michael London and Michael Häusser. Dendritic computation. *Annu. Rev. Neurosci.*, 28:503–532, 2005.

[53] Aaron Kerlin, Mohar Boaz, Daniel Flickinger, Bryan J MacLennan, Matthew B Dean, Courtney Davis, Nelson Spruston, and Karel Svoboda. Functional clustering of dendritic activity during decision-making. *Elife*, 8:e46966, 2019.

[54] Hongbo Jia, Nathalie L Rochefort, Xiaowei Chen, and Arthur Konnerth. Dendritic organization of sensory input to cortical neurons in vivo. *Nature*, 464(7293):1307–1312, April 2010.

[55] Rory G Townsend and Pulin Gong. Detection and analysis of spatiotemporal patterns in brain activity. *PLoS Comput. Biol.*, 14(12):e1006643, December 2018.

[56] Min Song, Minseok Kang, Hyeonsu Lee, Yong Jeong, and Se-Bum Paik. Classification of spatiotemporal neural activity patterns in brain imaging data. *Sci. Rep.*, 8(1):8231, May 2018.

[57] Yunliang Zang, Stéphane Dieudonné, and Erik De Schutter. Voltage- and Branch-Specific climbing fiber responses in purkinje cells. *Cell Rep.*, 24(6):1536–1549, August 2018.

[58] CF Ekerot and O Oscarsson. Prolonged depolarization elicited in purkinje cell dendrites by climbing fibre impulses in the cat. *The Journal of physiology*, 318(1):207–221, 1981.

[59] Sigita Augustinaite, Bernd Kuhn, Paul Johannes Helm, and Paul Heggelund. NMDA spike/plateau potentials in dendrites of thalamocortical neurons. *J. Neurosci.*, 34(33):10892–10905, August 2014.

[60] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017.