A MINIMAL MODEL OF NEURAL COMPUTATION WITH DENDRITIC PLATEAU POTENTIALS.

Johannes Leugering*
Fraunhofer Institute for Integrated Circuits
johannes.leugering@iis.fraunhofer.de

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

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

May 31, 2021

ABSTRACT

Over the last two decades, advances in neurobiology have established the essential role of active processes in neural dendrites for almost every aspect of cognition, but how these processes contribute to neural *computation* remains an open question. We show how two kinds of events within the dendrite, *synaptic spikes* and localized *dendritic plateau potentials*, interact on two distinct timescales to give rise to a powerful model of neural computation. In this theoretical model called *dendritic plateau computation*, a neuron's computational function is determined by the compartmentalization of its dendritic tree into functionally independent but mutually coupled segments. We demonstrate the versatility of this mechanism in a simulated navigation experiment, where it allows an individual neuron to reliably detect a specific movement trajectory over hundreds of milliseconds with a high tolerance for timing variability. We conclude by discussing the implications of this model for our understanding of neural computation.

1 Introduction

The vast majority of neural tissue is occupied by neural dendrites [1], the extensively branching tree structures on which almost all synapses terminate. Yet, most simplified neuron models have focused on the diverse dynamics underlying somatic spike generation [2]. In recent years, however, it has become increasingly evident that the computational function of a neuron is largely determined by properties and dynamics of its dendrite [3]. In particular, a steadily increasing number of studies find dendrites that generate active responses to spiking input. Further, these active dendritic processes appear to be fundamental to brain function [4, 5, 6, 7]. The observed effects range from short-lived Na^+ spikes [8] to the particularly striking dendritic plateau potential [9] — an intricate dynamic response mediated by NMDA receptors that maintains a strong depolarization of the local dendritic membrane potential for long periods of time. The complexity of these mechanisms and the diversity of neuron types makes determining the right level of abstraction for a single model of neural computation difficult [10].

For example, a recent study by Ujfalussy et al. [11] concluded that the somatic membrane potential in layer II/III pyramidal neurons can largely be explained by a linear statistical model. A simple non-linear model can improve the result, but the introduction of additional non-linearities only leads to minor improvements. Li et al. [12] also argue, that the properties of dendritic integration can be approximated well by a point-neuron model, if specific synaptic current effects are incorporated. On the other hand, seminal work by Poirazi et al. [13] suggests that in fact a 2-layer artificial neural network may be necessary to capture the input-output mapping of a single neuron, implying that the neuron's expressive power may be on a similar level. Whereas these earlier results analyzed the neuron in a static framework, Beniaguev et al. [14] incorporate temporal dynamics as well, and instead conclude that a temporally convolutional deep

*Both authors contributed equally.

neural network is better suited to model the single neuron's behavior. However, it is difficult to draw direct conclusions about the computational capabilities of a complex neuron from the complexity of a quantitative model of the neuron's membrane potential.

The most popular metaphor for neural computation to date is the linear-nonlinear (LN) point neuron that inspired the development of artificial neural networks [15]. LN neuron models make no use of dendritic complexity at all, and instead rely on a complex network of synaptic interconnections between individually simple neurons. This approach has worked exceedingly well for deep learning [16] and provides a compelling model for some forms of fast sensory processing e.g. in early visual areas [17], but it neglects a dimension critical for any interaction with the real world — time. Many behavioral or higher cognitive tasks require the ability to integrate, process and retain information across multiple, dynamically varying time scales. Consider, for example, a rodent navigating an environment in search of food. Receptive fields of place and grid cells tile a spatial map of this environment and encode the current position by their population activities [18, 19]. To navigate successfully, the animal needs to know not only its present location, but also the path it took through the environment. Decoding this path from sequential place and grid cell activity requires the integration of information on behavioral time scales that can span hundreds of milliseconds or more [20, 21]. Similar long sequential patterns can be found also in olfaction [22, 23] and cortical auditory processing [24], and they are likely involved in higher cognitive tasks such as language understanding, as well.

But how can such long temporal patterns of neural activity be processed by volatile neurons with membrane dynamics on the time scale of only tens of milliseconds or less [25]?

Our main idea is this: Neurons with active dendrites that generate dendritic plateau potentials have a form of working memory on a much longer time scale than that of individual spike responses. The interaction of these dendritic plateaus establishes a computation that enables single neurons to process information on long time scales and in a structured way. We derive this concept from a wealth of recent biological findings, which we categorize into four fundamental modelling assumptions (Section 2), and find that:

- A qualitative neuron model captures dendritic plateau computation in a tree structure of dendrite segments. The computational complexity of the model results from the interactions of these segments on the long timescale of plateau potentials (Section 3)
- The computational capabilities of dendritic trees can be characterized by a small set of elementary motifs (Section 4)
- Dendritic segments can robustly decode sequential activations of neuron populations. An example of path detection from place cell activity illustrates two key properties of dendritic plateau computation, timing invariance and a graded stochastic response, that lead to an intricate spatiotemporal receptive field (Section 5)
- Single neurons can implement structured computations over symbol-like inputs, motivating a new view on neural computation (Section 6)

2 Biological evidence for neural computation based on dendritic plateaus

Active generation of localized dendritic plateau potentials. Most of a cortical pyramidal neuron's excitatory synaptic inputs terminate on dendritic spines [26], where post-synaptic ion channels are activated via the stochastic, pre-synaptic release of glutamate-carrying vesicles [27, 28]. The activated channels, primarily controlled by α-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid receptors (AMPArs) [29], 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) [30]. In addition to AMPArs, the synaptic release of glutamate can also activate N-methyl-D-asparate receptor (NMDAr) gated ion-channels [29], but they do not become conductive unless a channel-blocking Mg^{+} ion is first displaced by a sufficiently strong depolarization [31, 32]. However, coincident EPSPs from multiple nearby spines can accumulate and thus induce this required depolarization of the local dendritic membrane potential [33]. Experimental as well as simulation studies report that this requires a volley of 4-20 or even up to 50 spikes within $1\,\mathrm{ms}$ to $4\,\mathrm{ms}$, depending on the location along the dendritic tree [33, 34, 35, 36]. The opening of NMDAr channels triggers a massive influx of different ionic currents that lead to a full depolarization of the local dendritic membrane potential. Although the isolated NMDAr response itself is reported to only last on the order of around $25\,\mathrm{ms}$ [37], in vivo recordings reveal that voltage-gated channels in the dendritic membrane [38] prolong this effect, resulting in an actively maintained depolarization that can last from tens to hundreds of milliseconds [39] (see **Fig. 1b** for an illustration of this mechanism). Such active long-lasting dendritic processes, *dendritic plateau potentials*, are ubiquitous [9, 40], and provide neurons with potentially useful memory traces that can last hundreds of milliseconds. Because NMDAr channels are gated by both depolarization and the presence of glutamate, plateau potentials remain localized, and do not actively propagate along the dendrite [41].

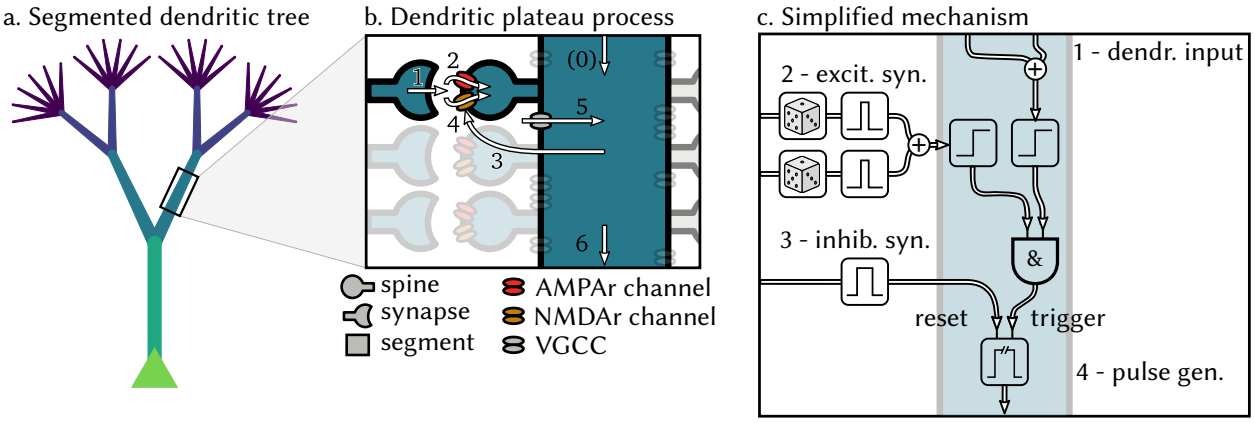

Figure 1: **How dendritic plateau potentials are generated, and how they interact. (a.)** A stylized neuron with dendritic arbor. **(b.)**Summary of the biological processes involved. A spike (1) releases glutamate, which opens AMPAr-gated ion-channels that depolarize the post-synaptic spine and cause an EPSP (2). If sufficiently many EPSPs coincide with up-stream dendritic input (0), the local membrane potentials rises (3) and NMDAr-gated ion-channels become de-inactivated, causing a further localized depolarization (4). Additional voltage-gated calcium channels can amplify and prolong this process (5) and cause a plateau potential, which can in turn moderately depolarize the parent segment (6 & 0). **(c.)** An algorithmic approximation of the biological mechanisms. If a dendrite segment is depolarized by sufficiently strong input from its child segments (1) and receives sufficiently strong excitatory input from its stochastic synapses (2), a local plateau potential is initiated. If the plateau is not interrupted by shunting inhibitory input (3), it depolarlizes the parent segment for an extended period of time.

Passive asymmetrical propagation of membrane potentials. The passive propagation of membrane potential through the dendrite is described by neural cable theory [42]. Only for very specific branching patterns, the complex dendritic tree can be reduced to an equivalent model of a cylinder, *Rall's ball-and-stick model*, in which the contribution of individual synaptic inputs sum (sub-)linearly and nearly instantaneously [43]. Due to its simplicity, this special case is often used to motivate abstract point-neuron models such as leaky integrate-and-fire neurons, which ignore the spatial dimension of the dendritic tree entirely and instead model the neuron as if it were a single electric compartment [44]. But in general, the passive spread of membrane potentials depends on the morphology and electrical impedance of the dendritic tree, specifically on the relationships of branch-diameters at the branching points [8]. For example, a back-propagating action potential moving in retrograde direction from soma to apical dendrite is only slightly attenuated if the dendritic branches become progressively thinner. However, the attenuation of singals in the anterograde direction is so strong, that synaptic input onto thin apical dendrites has little measurable effect on the membrane potential at the soma [45, 38]. One proposed solution to restoring "dendritic democracy" [46] and ensuring similar contribution of all synapses to the somatic membrane potential regardless of its position along the dendrite is an increased synaptic efficacy at distal synapses, which has been observed in hippocampal pyramidal neuron [47]. Another solution is the active amplification of distal synaptic input by active dendritic processes, which was shown to be required for somatic spiking [48]. Complex spike bursts in particular require the activation of NMDA receptors [49]. The resulting plateau potentials are also subject to anterograde attenuation along the dendritic cable, and thus only have a moderately depolarizing effect on their immediate neighborhood [50]. This effectively raises the local resting potential for the duration of the plateau potential, thus lowering the amount of coinciding spikes required to initiate a plateau potential in this neighbourhood [51].

Functional compartmentalization of dendrites. The structure of dendritic arbors has long been conjectured to play an important role for neural computation. Koch et al. [52] calculated, that due to impedance mismatch at branching points in the dendrites of various types of retinal ganglion cells, distinct electrically isolated *functional subunits* emerge, i.e. regions with a roughly equal local membrane potential throughout, which are only weakly coupled to their neighboring regions. For example, experiments in rats confirmed that thin dendrites in neorcortical pyramidal neurons can act as independent computational subunits that provide neurons with an additional non-linear integration site, increasing the potential computational power of the single neuron [53]. This behavior is not limited to pyramidal neurons, but rather appears to be a general principle that can be found in various forms across different cell types. For example, Purkinje cells in the cerebellum also generate localized Ca^{2+} events in response to coincident input on individual dendrite segments [54, 55], and thalamo-cortical neurons respond to strong synaptic input by localized plateaus in distal dendritic branches [56]. Branco and Häusser [57] identify such functional subunits with individual dendritic branches,

which they suggest constitute the "atomic unit" of computation in neural systems.
Rather than a single branch, one such functional compartment can also, stretch across multiple nearby branches, as long as synapses carrying correlated input signals cooperate to trigger local, regenerative events. Wybo et al. [58] present evidence for such compartmentalization and point out that structural plasticity may even allow the neuron to dynamically change the structure of its functional compartmentalization. We view dendrites as complex structures composed of functional subunits in this sense and will refer to them as *dendrite segments*.[2]. The segmentation implies, that only nearby synapses cooperate to trigger local regenerative events such as plateau potentials. To effectively drive a neuron, inputs therefore have to be clustered such that correlated spikes arrive at the same dendrite segment at the same time. This is observed in experiments [59] and suggests an alternative view of spike-based communication in which clustered groups of synaptic spines receive highly synchronized spiking inputs [60]. Since the resulting simultaneous EPSPs are required to trigger the NMDAr response, we therefore consider these highly synchronized *spike volleys* as the atomic unit of spike-based communication.

Stochastic excitation and shunting inhibition. An AMPAr or NMDAr response to an afferent neuron's input spike requires the prior release of neurotransmitter at the pre-synaptic terminal of a synapse. This process, however, is stochastic and best described by a "quantal" theory of neurotransmitter release [27, 28], according to which the successful transmission of a spike at a synapse is a random event with probability p_r. Branco et al. [61] found that in hippocampal synapses, p_r is distributed with a median of 0.22, which emphasizes the fundamentally stochastic nature of neural computation. The considerable variance of this distribution can be largely explained by the location of the synapse in the dendritic tree. The release probabilities of nearby synapses are much more homogeneous, which provides further evidence of the aforementioned functional segmentation of the dendrite.
AMPAr mediated EPSP responses as well as NMDAr mediated plateau responses can also be modified by inhibitory input at GABA_A or GABA_B synapses. The resulting shunting inhibition current can have both a subtractive and divisive effect on post-synaptic membrane potential [62]. The effect of shunting inhibition on active dendrites can be even more dramatic, outright stopping the generation of plateau potentials [63]. The interactions of plateaus and inhibition can be intricate [64], and inhibitory synapses tend to be placed critical positions within the dendrite to control dendritic excitability [65] or gate layer specific input from reaching the soma [66].

3 A computational model for dendritic plateau computation

From the biological observations outlined above, we derive a simple, qualitative model of active dendrites. At the core of this model lies the interaction of two types of events on distinct time-scales — short, spike-triggered EPSPs and long, actively generated dendritic plateau potentials — in a tree structure of dendrite segments. We define a *segment* as a minimal part of the dendritic tree, e.g. a single physical branch or stretching across multiple branches, that behaves as one functional, electrically isolated integration site. In other words, the synaptic inputs of a segment can cooperatively generate a plateau potential that stays confined to the segment. These dendrite segments form a tree structure with the soma at its root and thin dendrite branchlets as leaves.

Let's consider the function of one individual dendrite segment i in more detail (see figure 1c for a schematic). We distinguish excitatory and inhibitory synapses, which respectively produce excitatory (EPSPs) and inhibitory postsynaptic potentials (IPSPs). An excitatory synapse from neuron k to segment i only successfully transmits each spike with probability $p_{i,k}$. If the synapse transmits the spike, it induces an EPSP $\kappa_E(t)$ with duration τ_E and a magnitude $w_{i,k}$, which depends on the synaptic efficacy. Likewise for an inhibitory synapse, only that the duration τ_I of the IPSP is typically slightly longer. We model the shape of the post-synaptic potentials by rectangular pulses:

$$\kappa_E(t) = \begin{cases} 1 & \text{if } 0 \leq t \leq \tau_\text{E} \\ 0 & \text{otherwise} \end{cases}, \quad \kappa_I(t) = \begin{cases} 1 & \text{if } 0 \leq t \leq \tau_\text{I} \\ 0 & \text{otherwise} \end{cases}$$

We use exc_i and inh_i to represent the set of excitatory and inhibitory neurons targeting segment i, we denote the time of the m^th spike by neuron k with t_k^m, and introduce the i.i.d. random variables $\xi_{i,k}^m \sim \text{Bernoulli}(p_{i,k})$ to simplify notation. We can then define the combined effect of excitatory as well as inhibitory input for segment i[3]:

[2]We avoid the term "compartment" to prevent confusion with the concept of multi-compartment neuron models, which are commonly used as a spatially discretized solution to partial differential equation models of neurons.

[3]We assume that spike arrival times $t_{i,k}^m$ are at least τ_E apart.

$$\mathrm{EPSP}_i(t) = \sum_{k \in \mathrm{exc}_i} \sum_{m | t_k^m \leq t} \xi_{i,k}^m w_{i,k} \kappa_E(t - t_k^m) \tag{1}$$

$$\mathrm{IPSP}_i(t) = \sum_{k \in \mathrm{inh}_i} \sum_{m | t_k^m \leq t} \xi_{i,k}^m w_{i,k} \kappa_I(t - t_k^m) \tag{2}$$

$$\mathrm{PSP}_i(t) = \mathrm{EPSP}_i(t) - \mathrm{IPSP}_i(t) \tag{3}$$

One of the necessary preconditions for generating a dendritic plateau potential is a sufficiently strong net depolarization of the dendrite by synaptic input, i.e. larger than a segment-specific synaptic threshold TS_i, caused by the coincidence of multiple synchronous spikes. In thin dendrite branchlets, i.e. the leaf nodes of our tree structure, this is sufficient to trigger a plateau potential. But in the general case, additional depolarizing input from dendritic child branches is required. Here, we are only interested in the large depolarizing effects that actively generated plateau potentials have on directly adjacent segments, and we ignore the much weaker passive propagation of sub-threshold voltages along the dendrite.

We therefore introduce additional notation: child_i denotes the set of the direct children of segment i (if any), and $\mathrm{O}_k(t), k \in \mathrm{child}_i$ is the effect that the child segment k exerts on i at time t. Just like we did for the post-synaptic potentials, we can then define the total *dendritic input* $\mathrm{D}_i(t)$ into segment i:

$$\mathrm{D}_i(t) = \sum_{k \in \mathrm{child}_i} \mathrm{O}_k(t) \tag{4}$$

The segment-specific dendritic threshold TD_i determines, how much dendritic input is required in addition to synaptic input to trigger a plateau potential in segment i. For leaf nodes of the dendritic tree, i.e. segments without any children of their own, we set $\mathrm{TD}_i = 0$.

When both conditions become satisfied, i.e. there is sufficient synaptic and dendritic input, then a plateau potential is initiated. We use T_i^m to denote the starting-time of the m^{th} plateau potential in segment i:

$$T_i^m = \min t \geq T_i^{m-1} \quad \text{such that} \quad \mathrm{PSP}_i(t) \geq \mathrm{TS}_i \wedge \mathrm{D}_i(t) \geq \mathrm{TD}_i \tag{5}$$

The plateau plateau then typically ends at time $\tilde{T}_i^m = T_i^m + \tau_P$ after the fixed duration τ_P, unless it is interrupted by shunting inhibition, or it is prolonged by additional synaptic inputs.[4] We formalize these special cases as follows: The first inhibitory spike, if any, from neuron $k \in \mathrm{inh}_i$ at time $t_k^l \in [T_i^m; T_i^m + \tau_P]$ can end the plateau, i.e. in that case $\tilde{T}_i^m = t_k^l$. Otherwise, if another plateau is triggered at time $T_i^{m+1} \in [T_i^m; T_i^m + \tau_P]$ before the previous plateau has run its course, the first seamlessly flows into the second, i.e. $\tilde{T}_i^m = T_i^{m+1}$.

We can now define the output of segment i as a sequence of binary pulses, the plateau potentials:

$$\mathrm{O}_i(t) = \begin{cases} 1 & \text{if } \exists m : t \in [T_i^m; \tilde{T}_i^m] \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

This formalism can be iteratively applied to all segments of a neuron, including the soma, only that the segment produces a spike event followed by a brief refractory period τ_{refrac} instead of each long-lasting plateau potential. [5]

Conceptually, each dendrite segment acts first and foremost as a coincidence detector for a volley of synchronized spikes on the fast time scale of EPSPs. On the second, slower time scale of dendritic plateaus each segment is gated by its children in the dendritic tree. The computation of the neuron thus depends on a sequence of activations of its segments by spike volleys, which can be interrupted by shunting inhibition.

4 Motifs of dendritic plateau computation

The structure of the tree of dendrite segments determines which activation patterns lead to dendritic plateau potentials in all dendrite segments of a neuron, and therefore determines the computation implemented by the neuron. Each segment

[4] In engineering terms, this resembles a re-triggerable monoflop with reset.

[5] In addition to the forward-propagation of membrane potentials that we focused on so far (i.e. from child branches to the parent), the reverse direction typically has an even stronger effect — strong enough for the parent segment to depolarize its child segments by itself. To capture this effect, we recursively define that a neuron segment k's membrane potential $V_k(t) = O_k(t) \vee V_i(t), k \in \mathrm{child}_i$ is depolarized whenever either the segment itself or any of its ancestors produces a plateau potential. However, while this peculiarity may be relevant for learning, it cannot impact the forward model of dendritic plateau computation that we present here.

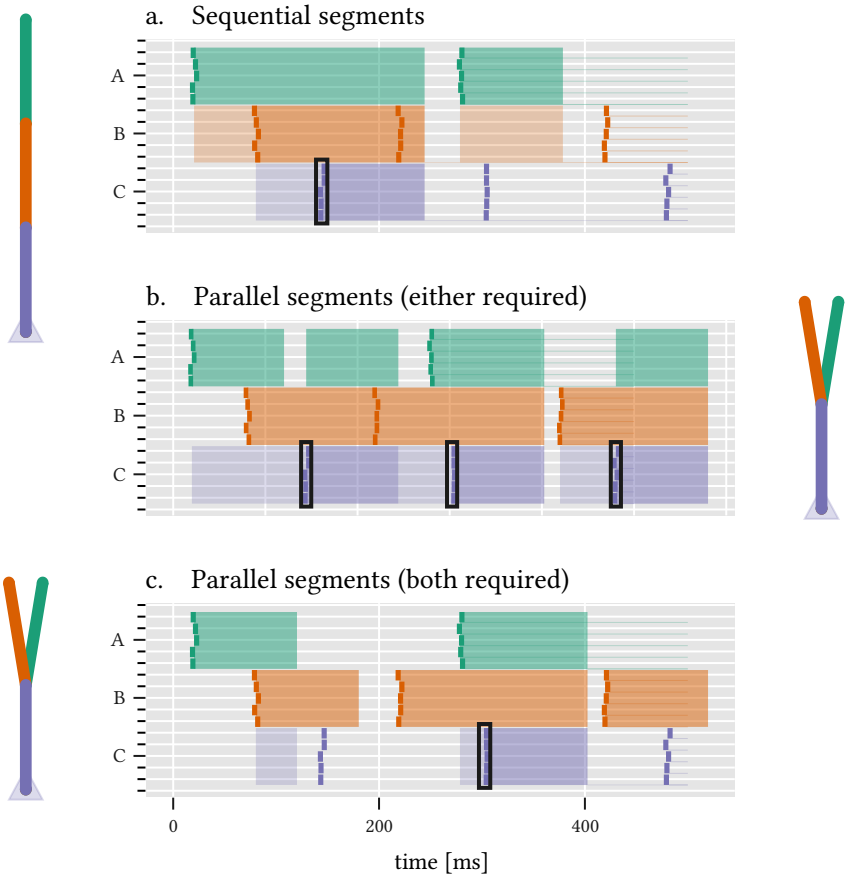

Figure 2: Various dendrite motives respond to different inputs. dendrite segments A, B and C receive spike volleys from corresponding neuron populations (color coded). We indicate for each segment, when it is enabled by its children (weakly shaded) or in a plateau state itself (shaded). **(a.)** If segments A, B and C form a chain, then C can only be activated while B is in a plateau state, whereas B can only be activated while A is in a plateau state. **(b.)** If both A and B are child-branches of C, either of which suffices to enable C, then C can be activated at any point where either A or B is in a plateau state. **(c.)** If both A and B are required, then C can only be activated while both A and B are in a plateau state.

is a coincidence detector for spike volleys, but additional input from a number of child segments may be necessary for a plateau potential to be triggered. The morphology of the dendritic tree defines these parent-child relations and the thresholds TD for required dendritic input. Changing these two variables changes the computation implemented, which we demonstrate in three prototypical motifs of dendritic plateau computation. Inhibition augments these motifs, for example to increase the specificity of pattern detection. Finally, stochastic synapses turn the otherwise deterministic neuron into a probabilistic pattern detector.

In the following examples, we look at neurons with several dendrite segments, each of which is connected to a small population of neurons that occasionally emits a volley of synchronized spikes. We are primarily interested in which patterns of spike volleys successfully trigger a somatic spike, and which do not. All experiments are simulated using open-source software (Section 8.2).

Dendrite structure determines computation

For example, dendritic segments can form a chain (**Fig. 2 a**), where each segment requires the previous one to be active ($\mathrm{TD}_i = 1$). A spike volley of at least five coincident spikes ($\mathrm{TD}_i = 1$) can therefore only trigger the most proximal segment, if a specific sequence of spike volleys activates each segment in the chain in the correct consecutive order.

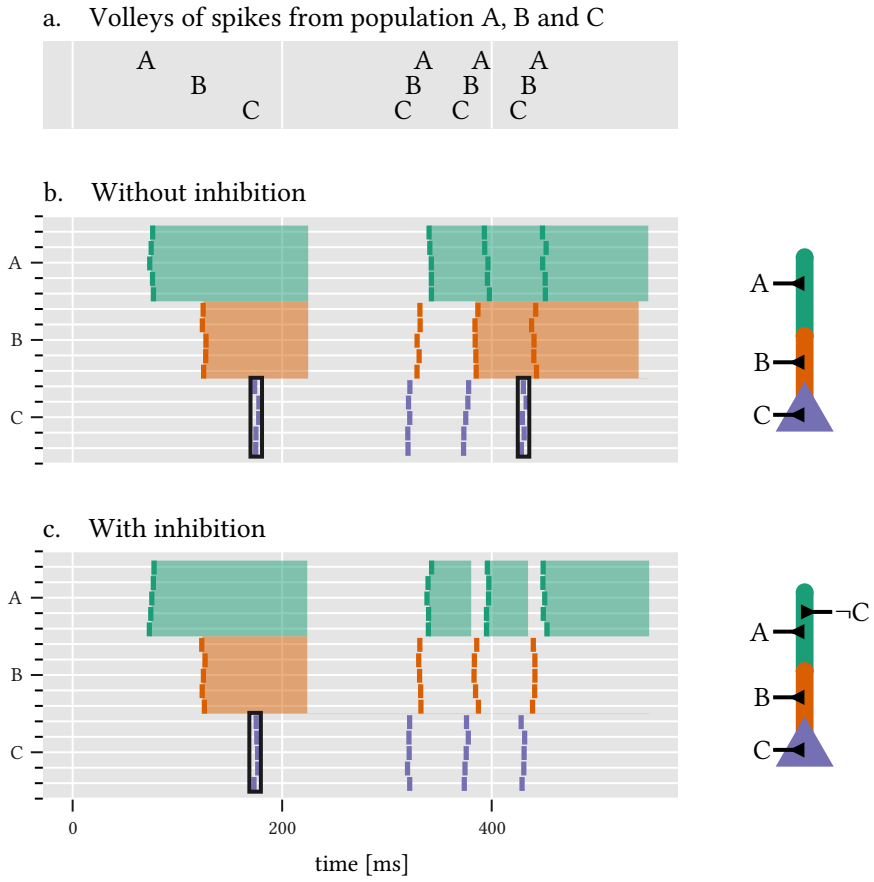

Figure 3: Shunting inhibition can prevent false detections. **(a.)** A neuron receives a sequence of spike volleys from three populations A, B and C. **(b.)** A neuron with a chain of dendrite segments A and B and soma C fires whenever they are activated in the correct order $A \to B \to C$, e.g. at time t_1. This also results in a false detection at t_2 if the desired sequence $A \to B \to C$ is contained in fast repetitions of the undesired sequence $C \to B \to A$. **(c.)** By adding shunting inhibition, the false detection at t_2 can be prevented.

However, because plateau potentials last for a comparatively long duration, the exact timing of the spike volleys within these time-intervals is not crucial, making the detection of spike-volley sequences largely timing invariant.

In the example **Fig. 2 a**, a spike volley from population A can trigger a plateau in the first segment, which in turn enables the second segment for the duration of the plateau $\tau_P = 100\,\mathrm{ms}$. If a volley from population B arrives at the second segment during that time-interval, it will trigger another plateau there, which in turn enables the third segment, and so on. If the first segment is not triggered first, or if its plateau has already ended, a volley from population B to the second segment remains ineffective. The chain of three segments shown here would therefore detect the sequence $A \to B \to C$ of spike volleys from populations A, B and C in that order, as long as the volleys come within $100\,\mathrm{ms}$ of each other or less.

What happens when one dendrite segment branches into two child segments? If either of the two child segments can provide enough dendritic input to enable the parent segment, i.e. $\mathrm{TD}_i = 1$ (see **Fig. 2 b**), then a spike volley from either population A or B can trigger a plateau that enables the third segment. In this motif, the third segment fires whenever a volley from C occurs within $100\,\mathrm{ms}$ after a volley from A *or* a volley from B. If instead the dendritic input from both child segments is required, i.e. $\mathrm{TD}_i = 2$ (see **Fig. 2 c**), then the third segment only fires if a volley from C occurs within $100\,\mathrm{ms}$ after both a volley from A *and* B.

Shunting inhibition prevents false positives

So far, we only looked at excitatory synaptic inputs and how they generate plateau potentials, but the shunting effect of inhibitory synapses plays an equally important role. To illustrate this, consider the example in **Fig. 3**. Our objective is to detect (within some timing constraints) any sequence $A \to B \to C$ of spike volleys from the populations A, B and C in that order. A chain of two dendrite segments and the soma, a motif we already saw above, will do just that. However, if we rely solely on excitatory synaptic input, any additional unnecessary inputs have no effect. This may be desirable in some cases but it may lead to false positives in others. For example, we might actually want to recognize the sequence $A \to A \to B \to B \to C \to C$ to contain the desired sub-sequence $A \to B \to C$. But the sequence $C \to B \to A \to C \to B \to A \to C \to B \to A$ shows three fast repetitions of the undesired sequence $C \to B \to A$, while still containing the desired sub-sequence $A \to B \to C$, and the neuron would fire all the same (see **Fig. 3** a).

To prevent the response to the anti-pattern $C \to B \to A$, we can add inhibitory input from population C to the two dendrite segments tasked with detecting A and B (see **Fig. 3** b). In that case, a volley from population C would terminate any ongoing plateau potentials in these two segments, thus preventing a response to the undesired sequence $C \to B \to A \to C \to B \to A \to C \to B \to A$ while leaving the response to the desired sequence $A \to B \to C$ unaffected. We can write this sequences with inhibition as $(A \wedge \neg C) \to (B \wedge \neg C) \to C$.

Shunting inhibition is therefore an important complementary mechanism for dendritic plateau computation, in particular if we consider that in our model, inhibiting the output of one segment at a branching point can effectively "veto" the entire computation of the corresponding subtree.

Stochastic synapses enable probabilistic computation

The various motifs shown above in combination with shunting inhibition can realize a wide range of operations via dendritic plateau computation. However, this mechanism responds to a rather long sequence of incoming spike volleys (potentially hundreds of milliseconds long) with an all-or-none response, i.e. a somatic spiking or nothing. Because the inputs to a neuron are also typically noisy, this might make an individual neuron's output too sparse and unreliable to base important decisions on it. We can overcome this problem, because the inherent stochasticity of synaptic transmission turn the probability that the neurons response into a graded response.

Let's consider an individual dendrite segment $i = 0$, that receives a spike volley from a population of $n = 10$ neurons (see **Fig. 4 a**). If each synapse independently transmits each spike it receives with the same probability P_0, then the number of actually transmitted spikes in a spike volley is a binomial random variable $\sim \text{Binomial}(P_0, n)$, and the probability that this number suffices to trigger a plateau potential depends on both P_0 and the segment's synaptic threshold TD_0. For a given threshold, the plateau probability $P_{\text{plateau}} = f(P_0)$ is hence a non-linear, sigmoidal function of both the volley size and the synaptic transmission probability. Despite the fact that the neuron has an all-or-none response for any individual spike volley, the expected value of its output, i.e. the probability to fire, is non-linear, graded response that reflects the size of the incoming volley.

If we extend this analysis to motifs of multiple dendrite segments, then the neuron's probability to fire is a non-linear function of the size of all incoming spike volleys. For example, to trigger a chain of two sequential segments with high probability, both segments have to be individually triggered with high probability, i.e. the neuron will only respond with high probability if both incoming spike volleys are large (see **Fig. 4 b**). The AND-like operation between plateaus that we observed in the deterministic case thus becomes a multiplication $P_{\text{chain}} = f(P_1) \cdot f(P_2)$ of plateau probabilities in the stochastic case. Similarly, if only one of two parallel segments needs to be activated (see **Fig. 4 c**), this happens with a probability $P_{\text{or}} = 1 - (1 - f(P_1)) \cdot (1 - f(P_2)) = f(P_1) + f(P_2) - f(P_1) \cdot f(P_2)$. The shown simulation results confirm this prediction.

This procedure can be applied inductively to more complex dendritic trees, as well. The neuron responds with a probability that depends on the size of all incoming spike volleys, and the expected value of the spike response thus encodes the "confidence" of the neuron in the result of a computation or detection. By combining multiple neurons with identical structure and synaptic input from the same source populations, we can construct an ensemble of neurons with a graded, probabilistic response. This ensemble can then respond to any potentially relevant sequence of incoming spike volleys with a volley of its own, such that the size of the emitted volley encodes the "confidence" of the ensemble in this detection.

5 Detecting movement trajectories from place cell activity

A good example to illustrate how dendritic plateau computation can function in a close-to-real-world example is the detection of sequential patterns in place cells. The location of an animal in its environment is represented by

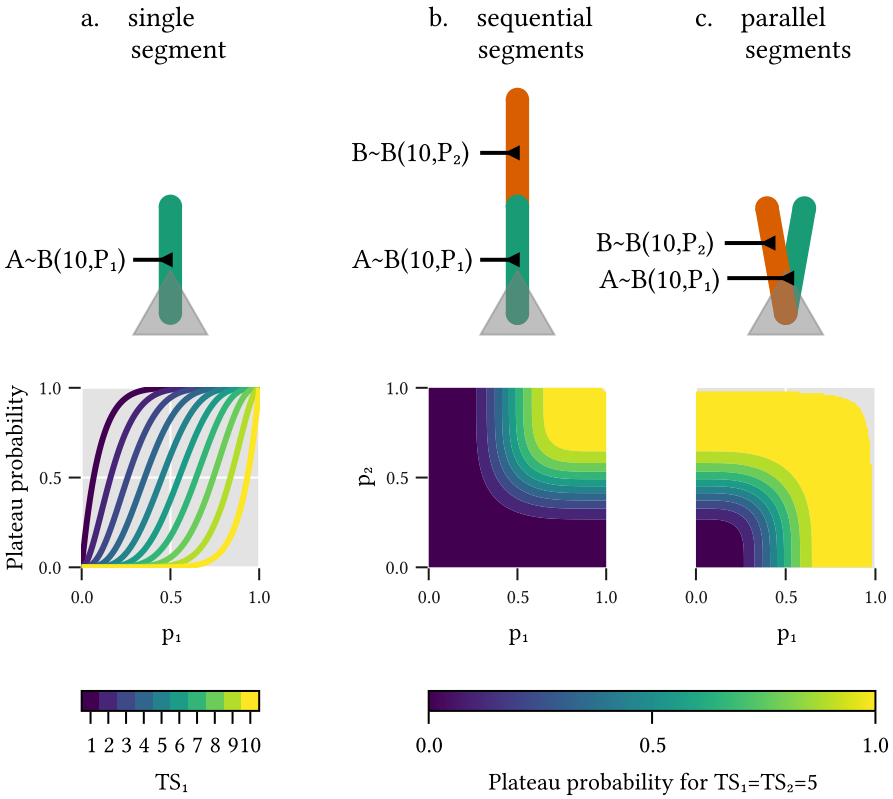

Figure 4: Stochastic synapses allow for graded responses. **(a.)** Out of a volley of ten spikes, the number of transmitted spikes is a Binomial distribution $B(10, P_1)$. A single dendrite segment, excited by such a volley, thus generates a plateau with a probability that depends on the synaptic transmission probability P_1 and the threshold TS_1 (color coded). **(b.)** Assuming two appropriately timed spike volleys activate two chained segments, then the probably that the second segment fires depends on both transmission probabilities P_1 and P_2 of the two synapse populations. The result resembles a probabilistic AND-gate. **(c.)** The probability of triggering at least one of two parallel segments resembles a probabilistic OR-gate.

place-cells [18, 19], each of which has a "receptive field" centered at a specific location. Navigation naturally produces sequential activation patterns as different locations are visited. The time scale of these patterns can be long and is variable because it is directly linked to the movement speed of the animal [21]. Further, active dendritic process have been shown to be selective for specific sequences of synaptic inputs [67] and dendritic spikes occur much more frequently in cortical pyramidal neurons of freely moving rats [68]. Applying our model to the problem of path decoding at varying movement shows how single neurons can solve this detection problem across multiple time scales.

We numerically simulate a rat moving through a small, 2-dimensional environment by generating stochastic paths at varying movement speed (more details in Section 8.1). The environment is tiled in a hexagonal grid by the receptive fields of place cell populations, each 20 neurons strong. These populations emit spike volleys with a varying magnitude that depends on the animal's distance to the center of the respective receptive field (**Figure 5 a and b**). Dendritic plateau computation allows a single neuron to detect specifically those paths, that traverse the receptive fields of three place cell populations in the correct order: from the bottom left (in green) through the center (in orange) to the top right (in purple). The neuron is composed of two sequentially chained dendrite segments and the soma, each of which receives synaptic input from exactly one of the place-cell populations and requires $\mathrm{TS_i} = 8$ coincident spikes to fire a plateau. In the presence of noise, this requires the fast detection of coincident spikes from each place cell population in order to distinguish legitimate spike volleys from background noise, as well as the interaction of long-lasting plateau potentials to detect the slow transition from one receptive field to the next on a behavioral time-scale. The problem thus has to two distinct time scales: fast estimation of the current location and slow integration of the traversed path.

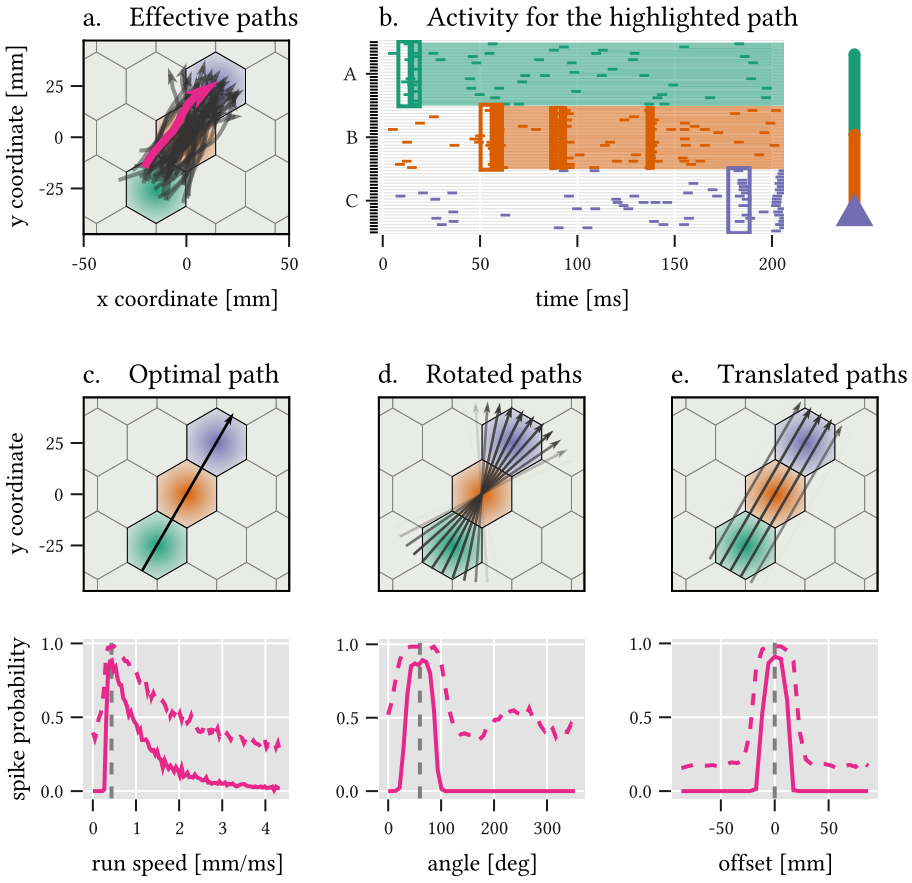

Figure 5: A simple neuron with three dendrite segments as shown to the right of panel b can detect directed paths on a timescale of 300 ms. **(a.)** The receptive fields of place cell populations tile the environment into a hexagonal grid. Random trajectories are generated through a stochastic process with randomized initial positions, velocities and angular heading to simulate the animal's movements. Only those trajectories are shown that elicit a spike response by the neuron. **(b.)** While the animal follows the highlighted trajectory (purple) through space, the place cell populations generate spikes, which in turn trigger plateau potentials in the corresponding dendrite segments (color coded). After initiation, plateau potentials can be extended by super-threshold inputs, as shown by the vertical lines. **(c.)** The neuron responds with highest probability to a path that traverses the center of the desired receptive fields at an optimal speed (top). Varying the movement speed (bottom) affects the probability of the neuron to detect the sequence (solid line). For a decreased threshold, the neuron's sensitivity is decreased, and the overall firing-probability is increased (dotted line). **(d.)** Changing the orientation or **(e.)** laterally shifting the path away from the optimal path rapidly decreases the probability of the neuron to fire, as well. Note that while the neuron is highly selective to orientation and offset, the firing probability only gradually changes as speed varies over an order of magnitude.

We can characterize the behavior of this path-detecting neuron by changing the speed, orientation and lateral offset of the path through then environment, and recording the neuron's probability to respond with a spike (**Fig. 5 c-e**). Firstly, the response probability for the optimal path is largest (almost $90\,\%$) with an optimal run-speed of around $0.5\,\mathrm{m\,s^{-1}}$, but even for a three times faster run speed, the neuron would still be able to detect it with roughly $30\,\%$ probability. This is due to the fact that the dendritic plateau computation is, within some limits, invariant to the specific timings of individual spike volleys. However, if the animal moves so slowly that the time-difference between spike volleys exceeds the plateau duration, or if it moves so quickly that a place-cell population fails to produce a spike volley, at all, then the probability to fire decreases. If a lower threshold $\mathrm{TS}_i = 4$ is chosen for each segment i, the sensitivity decreases, because even locations relatively far from the center of a receptive field can occasionally generate spike volleys of small magnitude (**Fig. 5 c dashed line**).

Secondly, our model neuron can be highly sensitive to the specific orientation and lateral offset of the desired path in space (**Fig. 5 d and e**). For a high threshold $\mathrm{TS}_i = 8$, only a narrow range of $\pm 30°$ around the optimal direction of $60°$ are reliably detected. By lowering the threshold to $\mathrm{TS}_i = 4$, each segment can be made less selective and the orientation-specificity of the neuron decreases substantially (**Fig. 5 d dashed line**), but as a side-effect the false detection probability also increases. When we shift the path orthogonally to the direction to the optimal path, only those pathes shifted by at most $\pm 10\,\mathrm{mm}$ from the center are reliably detected (**Fig. 5 e**). Just like for the rotated paths, decreasing the plateau firing threshold in each segment decreases sensitivity and increases the noise floor in the neuron's expected response.

6 Structured computation in single neurons

The issue of how working memory bridges the fast time scales of synaptic responses and the slower time scales of behavior is often addressed in recurrent networks of neurons, for example by slow emergent network dynamics [69], by fast synaptic plasticity [70], or as fading memory inherent in network dynamics [71]. However, if neurons generate dendritic plateaus in dendrite segments, then they already have access to internal memory. Moreover, this internal memory enables structured computation: the plateaus maintain a hidden state that can only be advanced if the correct input is seen at the correct segment in the correct time-frame. More formally, the single neuron is a (hidden-) state machine that accepts expressions of the form "A and B, then C" or similar, where the relative timing is restricted to the interval in which plateau potentials remain active. The specific form of the expression is determined by the dendritic tree itself, i.e. by the location of synapses and the strength of the coupling between the individual segments.
The "symbols" in these expressions are encoded into spike volleys and the timing of these spike volleys matters. Their order is particularly important, because it determines whether a sequence of spike volleys can activate a neuron or not. Therefore, neurons can process information in an event-based fashion where coherent spike events on a fast time scale trigger interacting plateau events on a slower time scale. This mode of computation allows neurons to respond much faster than a rate-code would permit, which is in line with empirical evidence on the level of single spikes in somatosensory cortex [72].

Structured, symbolic representations of information have been explored extensively in cognitive psychology, cognitive science, linguistics and artificial intelligence, but the most widely used analogy for neural computation are artificial neural networks, which are rooted in a connectionist view of cognition [73]: Simple and homogeneous individual units interact in a complex network to form distributed representations. Our perspective lies between these two extremes, and may help to resolve the apparent disconnect [74]: We conjecture that much of the richness of neural computation is derived from the structured internal memory processes of diverse and intricate neurons.

In our model, dendritic plateau computation allows neurons to detect specific, rare sequences of events and indicate this detection with as little as a single spike. This can result in an extremely sparse and hence metabolically efficient code, but it naturally comes at a price: if each pattern to be detected can last hundreds of milliseconds, then each neuron can only reliably detect one of these every couple of hundred milliseconds. The independent response of neurons in an ensemble can help in this regard. We suggest that the temporal coherence of spikes reflected in the magnitude of spike volleys is a mechanism through which a graded response can be encoded, for example the uncertainty about the occurrence of an input symbol A . As a result of the independent and stochastic synaptic transmission of spikes, dendrite segments can respond to this input symbol with a probability proportional to the volley magnitude. This in turn leads ensembles of neurons, which are sensitive to the same incoming patterns of spike volleys, to respond with spike volleys of their own. We therefore expect such ensembles that decode and encode information in spike timing and magnitude of spike volleys to be an integral building block in neural computation.

7 Discussion

Our qualitative model of dendritic plateau computation aims to explain and formalize the mounting biological evidence of neural computation in active dendrites. It asserts, that dendrites are segmented into functional units, each of which can generate and maintain a plateau potential when excited by a volley of spikes, and that the interaction of these plateau states allows a single neuron to detect remarkably complex temporal patterns. Our work is closely related to recent work by Hawkins and Ahmad [75], which proposes the use of active coincidence detection in dendrite segments to generate a long UP state at the soma. Similarly, Brea et al. [76] present an elegant two compartment model and a corresponding learning rule in which a basal dendrite segment learns to predict activation at the soma. In the "hierarchical temporal memory" model of neural computation, longer temporal sequences are detected by laterally connected neurons of this type [77].

We analyzed computation in single neurons, but what are the broader implications of this shift in perspective? Firstly, we have not addressed plasticity and learning. In fact, the reliance of our model on long-lasting plateau potentials and the ordering, rather than precise timing, of spike volleys poses a real challenge to most commonly used learning rules: On the one hand, this complicates temporal credit assignment for training paradigms that rely on instantaneous error signals, such as gradient backpropagation or related methods. On the other hand, the substantial, long-lasting depolarization of the membrane potential by localized plateau potentials within the dendrite with an accompanying high Ca^{2+} was shown to be the primary driver of synaptic plasticity [78, 79], which calls the role of backpropagating action potentials into question. In our model, the plateaus precede any potential somatic spiking. Besides the magnitude of synaptic weights, our model also makes extensive use of other properties of synapses, namely the transmission probability, the location of the synapse within the dendrite, and the delay, which can affect the synchronization of spike volleys. This increases the importance of structural plasticitity [80], homeostatic processes that adjust synaptic transmission probabilities [81] and recently proposed mechanisms for optimizing transmission delays through controlled (de-)myelinization [82] and opens the door for new learning rules insipred by these mechanisms.

Given a better understanding of plasticity and learning in dendritic plateau computation, we can approach the second challenge: constructing and optimizing large and useful networks from such complex neurons. The central theoretical questions are how the structured representation of information *inside* a neuron can be utilized in the context of a network, and in turn, how networks can reliably produce the synchronous spike volleys that encode relevant information in their magnitude, for example feature familiarity [83]. Recent advances in analysis techniques of brain data have shown promising results in this direction, which may help to establish direct empirical evidence of the spike-volley based representation of information we have proposed in this paper [84, 85]. On a more conceptual level, this event-based, symbolic view of neural computation could help to substantially reduce the gap between neural networks and cognitive architectures [74].

The prospect of energy efficient dendritic computation has also motivated research of potential implementations in neuromorphic hardware. For example, Intel's Loihi chip [86] and the DYNAPSE architecture [87] support some form of active non-linear processing in functionally isolated dendrite segments. Our model provides a new perspective on how these existing capabilities could be utilized for computation. But the simplicity of our proposed mechanism also suggests novel hardware implementations that use complex dendrite structures, rather than complex ion-channel dynamics or larger networks, to boost computational efficiency. We believe that this trade-off between structural and dynamic complexity of neurons will remain a critical topic for further research.

8 Materials and Methods

8.1 Implementation of the navigation experiment

To simulate the stochastic movements of an animal in a two-dimensional environment, 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:

$$\begin{aligned} \mathrm{d}X &= \cos(2\pi A)V\mathrm{d}t \\ \mathrm{d}Y &= \sin(2\pi A)V\mathrm{d}t \\ \mathrm{d}A &= 0.25\mathrm{d}W_A \\ \mathrm{d}V &= 10.0(0.25 - V)\mathrm{d}t + 0.1\mathrm{d}W_V \end{aligned} \tag{7}$$

A represents the angular heading of the animal, V represents its velocity in $\mathrm{m\,s^{-1}}$ 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 $10\,\mathrm{cm} \times 9.5\,\mathrm{cm}$, 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 $200\,\mathrm{ms}$. Three populations of place cells, each 20 neurons strong, are centered on a hexagonal grid with center-to-center distance of $r \approx 2.9\,\mathrm{cm}$. Each population randomly emits spike volleys following a homogeneous Poisson process with rate $\lambda = 50\,\mathrm{Hz}$. 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 $\sigma = 9.7\,\mathrm{mm}$, 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 $5\,\mathrm{Hz}$ to emulate background activity. Each spike is transmitted through stochastic synapses independently with probability 0.5.

Each of the simulated neuron's dendrite segments receives spiking input from the 20 neurons of one population and requires either 8 or 4 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 $v = \frac{3r}{200\,\mathrm{ms}} \approx 43\,\mathrm{cm\,s^{-1}}$ 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 500 times each.

8.2 Simulation framework for dendritic plateau computation

All simulations are implemented in a custom package developed in the Julia programming language [88], publicly available via the code repository hosted at https://github.com/jleugeri/DPC.jl. The simulator implements the neuron model outlined in this paper using a fast and extensible event-based formalism. All experiments and configuration files can be found in the `examples` subfolder of the repository.

Further documentation of the simulator, its interfaces ,and implementation details can be found there as well.

References

[1] Valentino Braitenberg and Almut Schüz. *Cortex: statistics and geometry of neuronal connectivity*. Springer Science & Business Media, 2013.

[2] Eugene M Izhikevich. Simple model of spiking neurons. *IEEE Transactions on neural networks*, 14(6):1569–1572, 2003.

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

[4] 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.

[5] Ning-Long Xu, Mark T Harnett, Stephen R Williams, Daniel Huber, Daniel H O'Connor, Karel Svoboda, and Jeffrey C Magee. Nonlinear dendritic integration of sensory and motor input during an active sensing task. *Nature*, 492(7428):247–251, December 2012.

[6] Naoya Takahashi, Thomas G Oertner, Peter Hegemann, and Matthew E Larkum. Active cortical dendrites modulate perception. *Science*, 354(6319):1587–1590, 2016.

[7] 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.

[8] Nelson Spruston, Greg Stuart, and Michael Häusser. Principles of dendritic integration. *Dendrites*, 351(597):1, 2016.

[9] 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.

[10] Andreas V M Herz, Tim Gollisch, Christian K Machens, and Dieter Jaeger. Modeling single-neuron dynamics and computations: a balance of detail and abstraction. *Science*, 314(5796):80–85, October 2006.

[11] Balázs B Ujfalussy, Judit K Makara, Máté Lengyel, and Tiago Branco. Global and multiplexed dendritic computations under in vivo-like conditions. *Neuron*, 100(3):579–592.e5, November 2018.

[12] Songting Li, Nan Liu, Xiaohui Zhang, David W McLaughlin, Douglas Zhou, and David Cai. Dendritic computations captured by an effective point neuron model. *Proc. Natl. Acad. Sci. U. S. A.*, 116(30):15244–15252, July 2019.

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

[14] David Beniaguev, Idan Segev, and Michael London. Single cortical neurons as deep artificial neural networks. *Cold Spring Harbor Laboratory*, page 613141, March 2020.

[15] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.

[16] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.

[17] Nikolaus Kriegeskorte. Deep neural networks: a new framework for modeling biological vision and brain information processing. *Annual review of vision science*, 1:417–446, 2015.

[18] 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.

[19] 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.

[20] 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.

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

[22] 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.

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

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

[25] 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.

[26] 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.

[27] J del Castillo and B Katz. Quantal components of the end-plate potential. *J. Physiol.*, 124(3):560–573, June 1954.

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

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

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

[31] 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.

[32] 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.

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

[34] 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.

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

[36] 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.

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

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

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

[40] 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.

[41] R Angus Silver, Andrew F MacAskill, and Mark Farrant. Neurotransmitter-gated ion channels in dendrites. *Dendrites, 3rd edn. Oxford University Press, New York*, pages 217–257, 2016.

[42] Steven S Goldstein and Wilfrid Rall. Changes of action potential shape and velocity for changing core conductor geometry. *Biophysical journal*, 14(10):731–757, 1974.

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

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

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

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

[47] 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.

[48] 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.

[49] 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.

[50] Matthew E Larkum, Thomas Nevian, Maya Sandler, Alon Polsky, and Jackie Schiller. Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle. *Science*, 325(5941):756–760, 2009.

[51] 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.

[52] 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.

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

[54] 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.

[55] 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.

[56] 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.

[57] 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.

[58] Willem A M Wybo, Benjamin Torben-Nielsen, Thomas Nevian, and Marc-Oliver Gewaltig. Electrical compartmentalization in neurons. *Cell Rep.*, 26(7):1759–1773.e7, February 2019.

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

[60] 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.

[61] Tiago Branco, Kevin Staras, Kevin J Darcy, and Yukiko Goda. Local dendritic activity sets release probability at hippocampal synapses. *Neuron*, 59(3):475–485, August 2008.

[62] Jiang Hao, Xu-Dong Wang, Yang Dan, Mu-Ming Poo, and Xiao-Hui Zhang. An arithmetic rule for spatial summation of excitatory and inhibitory inputs in pyramidal neurons. *Proc. Natl. Acad. Sci. U. S. A.*, 106(51): 21906–21911, December 2009.

[63] 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 Rep.*, 21(6):1550–1561, November 2017.

[64] K Du, Y W Wu, R Lindroos, Y Liu, and others. Cell-type–specific inhibition of the dendritic plateau potential in striatal spiny projection neurons. *Proceedings of the*, 2017.

[65] Albert Gidon and Idan Segev. Principles governing the operation of synaptic inhibition in dendrites. *Neuron*, 75 (2):330–341, July 2012.

[66] William Muñoz, Robin Tremblay, Daniel Levenstein, and Bernardo Rudy. Layer-specific modulation of neocortical dendritic inhibition during active wakefulness. *Science*, 355(6328):954–959, March 2017.

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

[68] Jason J Moore, Pascal M Ravassard, David Ho, Lavanya Acharya, Ashley L Kees, Cliff Vuong, and Mayank R Mehta. Dynamics of cortical dendritic membrane potential and spikes in freely behaving rats. *Science*, 355(6331), March 2017.

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

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

[71] Wolfgang Maass, Thomas Natschläger, and Henry Markram. Fading memory and kernel properties of generic cortical microcircuit models. *Journal of Physiology-Paris*, 98(4):315–330, July 2004.

[72] Rufin VanRullen, Rudy Guyonneau, and Simon J Thorpe. Spike times make sense. *Trends Neurosci.*, 28(1):1–4, January 2005.

[73] Paul Smolensky. On the proper treatment of connectionism. *Behav. Brain Sci.*, 11(1):1–23, March 1988.

[74] J A Fodor and Z W Pylyshyn. Connectionism and cognitive architecture: a critical analysis. *Cognition*, 28(1-2): 3–71, March 1988.

[75] 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.

[76] Johanni Brea, Alexisz Tamás Gaál, Robert Urbanczik, and Walter Senn. Prospective coding by spiking neurons. *PLOS Computational Biology*, 12(6):1–25, 06 2016. doi: 10.1371/journal.pcbi.1005003. URL `https://doi.org/10.1371/journal.pcbi.1005003`.

[77] Dileep George and Jeff Hawkins. Towards a mathematical theory of cortical micro-circuits. *PLoS Comput. Biol.*, 5(10):e1000532, October 2009.

[78] 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.

[79] 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.

[80] Andreas Knoblauch and Friedrich T Sommer. Structural plasticity, effectual connectivity, and memory in cortex. *Front. Neuroanat.*, 10:63, June 2016.

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

[82] R Douglas Fields. A new mechanism of nervous system plasticity: activity-dependent myelination. *Nat. Rev. Neurosci.*, 16(12):756–767, December 2015.

[83] Clemens Korndörfer, Ekkehard Ullner, Jordi García-Ojalvo, and Gordon Pipa. Cortical spike synchrony as a measure of input familiarity. *Neural computation*, 29(9):2491–2510, 2017.

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

[85] 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.

[86] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, Yuyun Liao, Chit-Kwan Lin, Andrew Lines, Ruokun Liu, Deepak Mathaikutty, Steven McCoy, Arnab Paul, Jonathan Tse, Guruguhanathan Venkataramanan, Yi-Hsin Weng, Andreas Wild, Yoonseok Yang, and Hong Wang. Loihi: A neuromorphic manycore processor with On-Chip learning. *IEEE Micro*, 38(1):82–99, January 2018.

[87] Saber Moradi, Ning Qiao, Fabio Stefanini, and Giacomo Indiveri. A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs). *IEEE Trans. Biomed. Circuits Syst.*, 12(1):106–122, February 2018.

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