From dendritic computation to symbolic operation.

Pascal Nieters
Institute of Cognitive Science, Osnabrück University

Summary. Neuromorphic computing aims to provide new technologies and substrates for artificial neural networks to implement cognitive computations otherwise implemented by brains. However, it is not clear how symbolic manipulations our minds carry out routinely are implemented in subsymbolic neural computation. Therefore, we ask what neuromorphic theories are missing. Our hypothesis is that neural dendrites are substrates for logical operations and that single neurons embedded in larger networks are already capable structured representations of knowledge. We discuss how new neuromorphic computers can benefit from implementing local, structured computers with memory as part of larger networks to unlock a new category of computations.

Traditional von-Neumann computers are well-formalized and understood systems that do however become more and more difficult to scale. Newly developed approaches to computation, for example neuromorphic computing inspired by brain circuits, promise to break through this apparent glass-ceiling but do not offer a clear-cut formalisation of computation and often end up as black-box approaches to computation. This tension between understanding computation and scaling new computational paradigms is also reflected in the debate between connectionists and proponents of traditional cognitive architectures in neural network research [1, 2]. Proposed solutions usually entail implementing reasoning on top of a connectionist implementation. For example, conceptors have recently been introduced as a method to assign transient dynamics characterized by regions in the phase space of recurrent networks to symbols that can be combined with logical connectives [3]. In all cases, models fundamentally agree on the basic notion of what a neural network is: A network of single, atomic nodes following the point-neuron model established early pioneers of neural computation.

However, we now know that the picture is incomplete. Cellular recordings and observational studies during animal behavior have now convincingly shown that the dendritic trees of neurons are heavily involved in neural processing [4, 5]. In response to significant coincident spiking input, segments of the dendritic tree actively respond by actively generating a local plateau potential significantly longer than the effect individual spikes have on the membrane potential. The significant difference in time scale and the all-or-non nature of the plateau response imply that the picture of individual neurons implementing a linear non-linear mapping of the rate of a spike train onto the rate of its own spike train cannot hold for neurons in which plateau potentials play an important role for computation. Instead, the distributed representations underlying the spiking input to a neuron turns into a localized representation of a coincidence detection which is kept in memory within a dendrite segment and accessible to its neighbours for computational purposes.

We have shown in previous work how this simple mechanism is ideally suited to recognizing sequences with significant timing invariance, a necessary feature for any circuit processing inputs from a real, dynamic environment [6]. Here, we go beyond simple sequences and describe computation in branching dendritic trees as structured, logical sentences over symbols rank-ordered in time.

The underlying model is simple enough (Fig. 1 **b.**): Dendrite segments are organized in a tree, each segments output is input to its parent. For each segment to generate a plateau, a certain number of dendritic inputs from child segments are required in addition to excitatory synaptic input that crosses a threshold and implements coincidence detection. Inhibitory inputs interrupt the plateau process and either prevent it entirely if they occur just before excitatory inputs or cancel the plateau outright if they occur during plateau generation.

Each detection represents a coinciding combination of input spikes as predicate symbol and inhibition as the negation of it. Because inhibition only interacts only with plateau processes, this is not a true negation but a veto – similar to McCulloch and Pitts circuits [7]. Dendritic input from child segments must always occur before the detection of a symbol at a parent segment, so the dendritic tree naturally operates on symbolic expressions that are rank-ordered in time. Depending on the dendritic threshold of the parent segment, different logical connections between sibling segments can be implemented. Starting at the root of the tree, the recursive definition of the computation of a single neuron with segmented dendritic tree is (see Fig. 1 **c.** for an example):

$$\mathbf{S}_t = \underbrace{s_t}_{\text{excitation}} \wedge \underbrace{\neg s_t}_{\text{inhibition}} \wedge \underbrace{\sum_{\pi \in \kappa} \prod_{j \in \pi} \mathbf{S}[j]_{t-1}}_{\text{dendritic gating}}$$

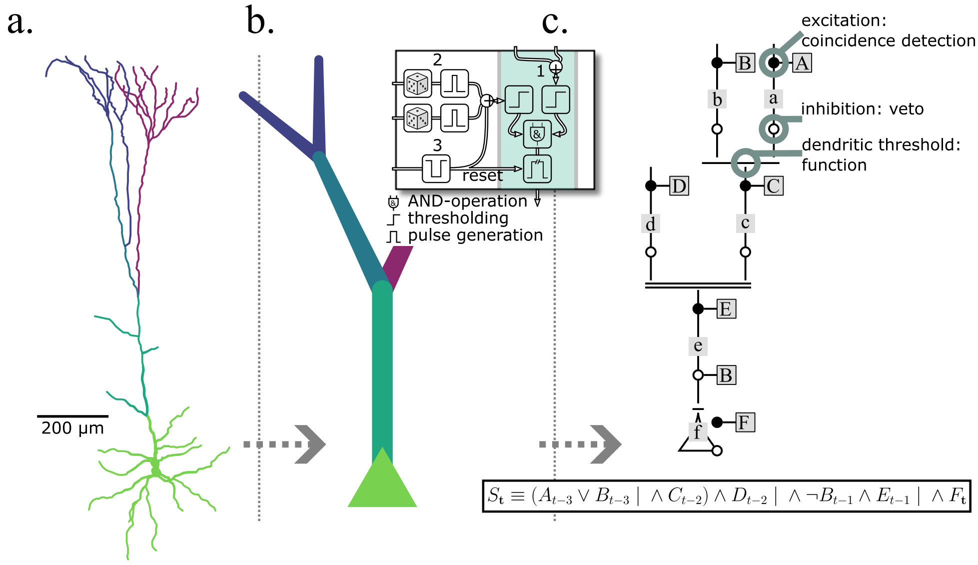

Figure 1: **Neuron to logical circuit.** (a.) Segments of a pyramidal neuron map onto a (b.) segmented dendritic tree driven by dendritic inputs on top of synaptic excitation and inhibition. (c.) The computational function of these neurons can be formalized as predicate sentences over symbols rank-ordered in time when connected to input populations.

We do not argue that symbolic thought is a higher level process implemented on top of a subsymbolic substrate. Instead, we argue, structured and symbolic representations are inherently integrated in neural networks in the "subatomic" structures of dendritic trees of single neurons. Representations are distributed among many different neurons on the network level, but local to individual dendrite segments and part of constituent structures on the neuron level.

Our work provides a fundamentally new approach to computation in neural network with direct implications for neuromorphic hardware: Dendritic structures and longer time scales are decidedly important for computational function in our model. But all other new computational architectures that follow the parallel and distributed computation paradigm exemplified by neural networks can also learn from the biological phenomenon of plateau potentials in dendritic trees. Generally, local memory traces in tiny, symbolic computers embedded in larger, dynamic networks may are a promising path to implement cognitive functions in new computing systems.

References

[1] Fodor, J. A., and Z. W. Pylyshyn, "Connectionism and Cognitive Architecture: A Critical Analysis." Cognition 28 (1-2): 3–71, 1988

[2] Smolensky, Paul, "On the Proper Treatment of Connectionism." The Behavioral and Brain Sciences 11 (1): 1–23, 1988

[3] Jaeger, Herbert,"Controlling Recurrent Neural Networks by Conceptors." arXiv [cs.NE]. arXiv. http://arxiv.org/abs/1403.3369, 2014

[4] Antic, Srdjan D. et al, "The Decade of the Dendritic NMDA Spike." Journal of Neuroscience Research 88 (14): 2991–3001, 2010

[5] Kerlin, Aaron, et al, "Functional Clustering of Dendritic Activity during Decision-Making." eLife 8, 2019

[6] Leugering, Johannes, Pascal Nieters, and Gordon Pipa. "Dendritic Plateau Potentials Can Process Spike Sequences across Multiple Time-Scales." Frontiers in Cognition 2. 2023

[7] W. S. McCulloch and W. Pitts, "A logical calculus of the ideas immanent in nervous activity," Bull. Math. Biophys., 1943.