

Impact of dendritic non-linearities on the computational capabilities of neurons

Carlo Baldassi

Department of Computing Sciences, Bocconi University
20136 Milano, Italy

Nicolas Brunel

Department of Computing Sciences, Bocconi University
20136 Milano, Italy

Clarissa Lauditi

Department of Applied Math, John A. Paulson School of Engineering and Applied Sciences, Harvard University
02138 Cambridge, MA, USA

Enrico M. Malatesta

Department of Computing Sciences, Bocconi University
20136 Milano, Italy

Fabrizio Pittorino

Department of Electronics Information and Bioengineering, Politecnico di Milano
20133 Milano, Italy

Riccardo Zecchina

Department of Computing Sciences, Bocconi University
20136 Milano, Italy

Abstract

Recent experiments in neurophysiology, primarily in pyramidal cells, have shown that dendrites contribute to neuronal computational capabilities with non-linear synaptic input integration. In this work we model a single neuron as a two-layer network with non-overlapping synaptic weights and a biologically plausible form of dendritic non-linearity, which is analytically tractable with statistical physics methods. Analytical and numerical analysis of the model reveals key computational benefits of non-linear dendritic integration over traditional linear neuron models. We find that the dendritic non-linearity concurrently enhances the number of possible learned input-output associations and the learning speed. At variance with previously studied linear neuron models, we find that the experimentally observed synaptic weight sparsity naturally emerges as a consequence of non-linear dendritic integration, while the experimental synaptic weight distribution is consistently reproduced. Non-linearly induced sparsity comes with a second advantage for information processing, i.e. input and synaptic noise robustness. By testing our model on standard real-world benchmark datasets inspired by deep learning practice, we empirically observe that the non-linearity provides an enhancement in generalization performance - a desirable property of neurons for non-trivial information processing.

Introduction

Understanding the computational capabilities of single neurons is among the most fundamental open problems in neuroscience. A long-standing question concerns the role of dendrites in shaping neuronal information processing. Traditionally viewed as simple linear summing devices akin to perceptrons, whose learning capabilities have been studied extensively (Minsky & Papert, 1988; Gardner, 1988), this perspective has been challenged by the recognition of non-linear integration of synaptic inputs (London & Häusser, 2005; Larkum, 2013; Major, Larkum, & Schiller, 2013), due to active currents in dendrites, influenced by various voltage-gated ionic currents. Recent studies show that dendrites significantly impact biological learning and memory, performing linear or non-linear input processing depending on their spatial arrangement (Polsky, Mel, & Schiller, 2004). This insight has led to comparisons of neurons to complex computational models, from multi-layer perceptrons (Poirazi, Brannon, & Mel, 2003) to cascades of linear-nonlinear processes (Brunel, Hakim, & Richardson, 2014; Ujfalussy, Makara, Lengyel, & Branco, 2018). Overall, a potential connection between dendritic mechanisms and machine learning has been proposed (Pagkalos, Makarov, & Poirazi, 2024). Here, we set out to study the computational capabilities of a single neuron model with dendrites implementing experimentally observed non-linear integration and with sign-constrained positive synapses modelling excitatory connectivity. In that two-layer sign-constrained neuron model with a biologically plausible saturating non-linearity, neuron computational abilities are significantly enhanced. This includes improvements in the

capacity for input-output associations, training speed, noise robustness, and generalization to new inputs. Additionally, this model replicates the experimentally observed synaptic weight distribution and sparsity without needing ad hoc robustness parameters, as required in linear neuron models (Brunel, Hakim, Isope, Nadal, & Barbour, 2004; Brunel, 2016).

Single neuron model

We consider a single neuron model that transforms N binary synaptic inputs $\xi_i = \{0, 1\}^N$ into a binary output $\hat{\sigma} = \{0, 1\}$. In the standard linear perceptron model, the neuronal output is $\hat{\sigma} = \Theta(\sum_{i=1}^N W_i \xi_i - T)$, where Θ is the Heaviside function, W is a vector of synaptic weights, typically optimized by a learning process, and T is a threshold. Here, motivated by experiments that have revealed significant non-linearities in the summation of inputs within single dendritic branches, but not across branches (Polsky et al., 2004; Larkum, 2013), we consider a generalization of the perceptron model with K dendritic non-overlapping branches, and non-linear summation of inputs within each dendritic branch. In this model, the neuronal output is $\hat{\sigma} = \Theta(\Delta)$, where $\Delta = \frac{1}{\sqrt{K}} \sum_{l=1}^K g(\lambda_l) - \sqrt{K} \theta_s$ is the total input to the soma, proportional to the sum of the outputs of all dendritic branches; g is a non-linear function describing the dendritic non-linearity; $\lambda_l = \sqrt{\frac{K}{N}} \sum_{i=1}^{N/K} W_{li} \xi_{li} - \sqrt{\frac{N}{K}} \theta_d$ is the total input to dendritic branch l , which is a linear sum of inputs to this branch $\xi_{li} \in \{0, 1\}^{N/K}$, weighted by synaptic efficacies W_{li} . Usually, only the synaptic weights of excitatory (glutamatergic) pre-synaptic neurons are explicitly modeled and subject to plasticity, while inhibitory (GABAergic) synapses are considered static. In our model the latter are grouped under the two somatic θ_s and dendritic θ_d thresholds, so that each synaptic weight $W_{li} \geq 0$.

Dendritic non-linearity Experiments in neocortical pyramidal cells have indicated that the dendritic output is roughly linear at low stimulation intensities, and that it then increases in a strongly non-linear fashion beyond a threshold, before saturating (Polsky et al., 2004). To reproduce the shape of the experimentally-recorded function, we consider the dendritic non-linear activation

$$g_{\text{polsky}}(x) = \begin{cases} \max(0, x) & x < x_{\min} \\ \frac{2(1-x_{\min})}{1+e^{-\gamma(x-x_{\min})}} - 1 + 2x_{\min} & x \geq x_{\min} \end{cases} \quad (1)$$

where x_{\min} is the non-linearity threshold, and γ describes the strength of the non-linearity. We refer to this non-linearity as the Polsky transfer function, which interpolates between ReLU (when $x_{\min} \rightarrow \infty$) and the step function for $x_{\min} = 0$ when $\gamma \rightarrow \infty$. Note also that in the experiments of ref. (Polsky et al., 2004), only excitatory inputs are considered, and consequently only the positive side of the dendritic non-linearity is probed. For consistency, we set g to zero on the negative side.

Learning tasks We consider a standard classification task with the objective of learning a dataset $\mathcal{D} = \{\xi^\mu, \sigma^\mu\}_{\mu=1}^P$ com-

posed of P binary random input patterns ξ_{li}^μ and labels σ^μ that are both i.i.d. Bernoulli variables with $P(\xi_{li}^\mu = 1) = f_{\text{in}}$ (input coding level) and $P(\sigma^\mu = 1) = f_{\text{out}}$ (output coding level) respectively. The task of the neuron is to correctly classify all input patterns, i.e. produce the correct output $\hat{\sigma} = \sigma^\mu$ when input ξ^μ is presented. This classification task (often called "storage problem" in the literature), whose goal is to learn the associations by progressively modifying the synaptic weights, has been studied extensively for both perceptron and committee architectures (Gardner & Derrida, 1988; Baldassi, Malatesta, & Zecchina, 2019; Zavatone-Veth & Pehlevan, 2021), either by optimizing directly the number of errors or some surrogate loss functions. On the numerical side only, we also study classical benchmark classification tasks in machine learning, providing realistic correlated datasets, such as MNIST, Fashion-MNIST and CIFAR-10.

Computational capabilities

Critical and algorithmical capacities

To investigate the memory-related properties of our single neuron model in the storage setting, one leverage on asymptotic methods from statistical physics (Engel & Van den Broeck, 2001). Given a density of patterns $\alpha = \frac{P}{N}$, and in the asymptotic limit where the total number of dendrites N and the number of branches K are $N, K \rightarrow \infty$ with $\frac{K}{N} \rightarrow 0$, it is possible to analytically study what is the typical maximum threshold α_c for which the complexity of the model is sufficient to classify the activity patterns correctly. This value of α_c is a function of the external parameters $\theta_d, \theta_s, f_{\text{in/out}}$ of the model and of the specific choice of the activation g . With a linear activation $g(x) \equiv x$, our model behaves like a traditional one-layer neuron model, reproducing well-known results about critical capacity as explored in prior studies (Brunel et al., 2004; Brunel, 2016). Instead, by taking advantage of a non-linear integration of synaptic inputs as in (1), one can make the storage capacity of the model α_c to arbitrarily increase with respect to the linear case. By playing with the dendritic inhibitory threshold θ_d when $g(x) \equiv \text{Polsky}(x)$, one can *strongly* alters the expressivity of the model and switch from the linear regime when $\theta_d \rightarrow 0$ and $\alpha_c^{\text{perc}}(f_{\text{in/out}} = 0.5) \sim 1$ to high values of α_c when $\theta_d \gg 1$. On the algorithmic side, using the SGD and LAL (Barkai, Hansel, & Sompolinsky, 1992) algorithms (modified for positive weights) with optimally tuned hyper-parameters, we demonstrate that non-linear models can achieve higher algorithmic capacities than linear models, despite the non-convexity of the optimization challenge. Plus, the non-linear model requires fewer training epochs to perfectly learn a dataset compared to its linear counterpart, indicating enhanced learning speed thanks to dendritic non-linearities.

Distribution of synaptic weights and sparsity

The distribution of synaptic weights, considered to be informative of the brain learning processes, began to be observed in the early 2000s (Isope & Barbour, 2002). Electrophysiological




Figure 1: (left) Experimental distribution of synaptic weights compared with the analytical prediction. (right) Fraction of silent synapses as a function of the dendritic inhibition for Polsky ($x_{\text{min}} = 0.33, \gamma = 15$) and ReLU non-linearities.

studies on various types of neurons have consistently found a large fraction of silent synapses. In neocortex pyramidal cells, the authors of (Song, Sjöström, Reigl, Nelson, & Chklovskii, 2005) report a value for the fraction of silent synapses of about 90%. In Fig. 1 we successfully align the theoretical synaptic weight distribution of our non-linear neuron model with experimental observations (Song et al., 2005) of pyramidal cells. We set input and output coding levels at $f_{\text{in/out}} = 0.05$, adjusting the dendritic threshold to match the experimental average synaptic weight, and tuning the somatic threshold to achieve a 90% fraction of silent synapses. The critical capacity $\alpha_c \simeq 9.26$ corresponding to the analytical curve, can give us an estimate of the constraint density at which pyramidal neurons operate. To achieve the desired 90% fraction of silent synapses in the Polsky case, we increased the value of θ_d , while keeping other model parameters fixed (Fig. 1). Unlike the perceptron case, where achieving similar levels of sparsity requires adding a predefined robustness adjustments during learning (Brunel, 2016), the non-linearity in dendrites naturally induces sparsity in our setting.

Noise robustness and generalization

From an algorithmic standpoint, our numerical simulations (for both SGD and LAL) aligns with the theoretical findings. First, we assess robustness to input noise by flipping training set pattern entries with a probability ρ . Instead, for synaptic noise robustness, we introduce multiplicative Gaussian noise to the weights and compute the change in training errors. In both cases, the non-linear model keeps a robust performance against input and synaptic noises, enhancing both stability and performance under noisy conditions. To evaluate the generalization capabilities of our neuron model, we conducted binary classification tasks using standard machine learning datasets: MNIST, Fashion-MNIST, and CIFAR-10. In all cases, the update model of neuron outperforms its linear counterpart.

Conclusions

We explored a biologically inspired two-layer neural network model featuring non-linear dendritic integration as observed in neocortical pyramidal cells. Our findings highlight the advantages of dendritic non-linearity, notably enhancing neuron

expressivity across different metrics compared to the linear model: storage capacity, sparsity, robustness, and generalization capabilities.

References

- Baldassi, C., Malatesta, E. M., & Zecchina, R. (2019, Oct). Properties of the geometry of solutions and capacity of multilayer neural networks with rectified linear unit activations. *Physical Review Letters*, *123*(17).
- Barkai, E., Hansel, D., & Sompolinsky, H. (1992, Mar). Broken symmetries in multilayered perceptrons. *Phys. Rev. A*, *45*, 4146–4161.
- Brunel, N. (2016, 04). Is cortical connectivity optimized for storing information? *Nature neuroscience*, *19*.
- Brunel, N., Hakim, V., Isope, P., Nadal, J.-P., & Barbour, B. (2004). Optimal information storage and the distribution of synaptic weights: Perceptron versus purkinje cell. *Neuron*, *43*(5), 745-757.
- Brunel, N., Hakim, V., & Richardson, M. J. (2014). Single neuron dynamics and computation. *Current Opinion in Neurobiology*, *25*, 149-155. (Theoretical and computational neuroscience)
- Engel, A., & Van den Broeck, C. (2001). *Statistical mechanics of learning*. Cambridge University Press.
- Gardner, E. (1988, jan). The space of interactions in neural network models. *Journal of Physics A: Mathematical and General*, *21*(1), 257–270.
- Gardner, E., & Derrida, B. (1988, jan). Optimal storage properties of neural network models. *Journal of Physics A: Mathematical and General*, *21*(1), 271–284.
- Isope, P., & Barbour, B. (2002). Properties of unitary granule cell: Purkinje cell synapses in adult rat cerebellar slices. *Journal of Neuroscience*, *22*(22), 9668–9678.
- Larkum, M. (2013). A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. *Trends Neurosci.*, *36*, 141–151.
- London, M., & Häusser, M. (2005). Dendritic computation. *Annual Review of Neuroscience*, *28*(1), 503-532. (PMID: 16033324)
- Major, G., Larkum, M. E., & Schiller, J. (2013). Active properties of neocortical pyramidal neuron dendrites. *Annu. Rev. Neurosci.*, *36*, 1–24.
- Minsky, M., & Papert, S. A. (1988). *Perceptrons - expanded edition*. MIT Press.
- Pagkalos, M., Makarov, R., & Poirazi, P. (2024). Leveraging dendritic properties to advance machine learning and neuro-inspired computing. *Current Opinion in Neurobiology*, *85*, 102853.
- Poirazi, P., Brannon, T., & Mel, B. W. (2003). Pyramidal neuron as two-layer neural network. *Neuron*, *37*(6), 989-999.
- Polsky, A., Mel, B. W., & Schiller, J. (2004). Computational subunits in thin dendrites of pyramidal cells. *Nature neuroscience*, *7*(6), 621–627.
- Song, S., Sjöström, P. J., Reigl, M., Nelson, S., & Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS biology*, *3*(3), e68.
- Ujfalussy, B. B., Makara, J. K., Lengyel, M., & Branco, T. (2018). Global and Multiplexed Dendritic Computations under In-Vivo-like Conditions. *Neuron*, *100*, 579–592.
- Zavatone-Veth, J. A., & Pehlevan, C. (2021, Feb). Activation function dependence of the storage capacity of tree-like neural networks. *Phys. Rev. E*, *103*, L020301. doi: 10.1103/PhysRevE.103.L020301