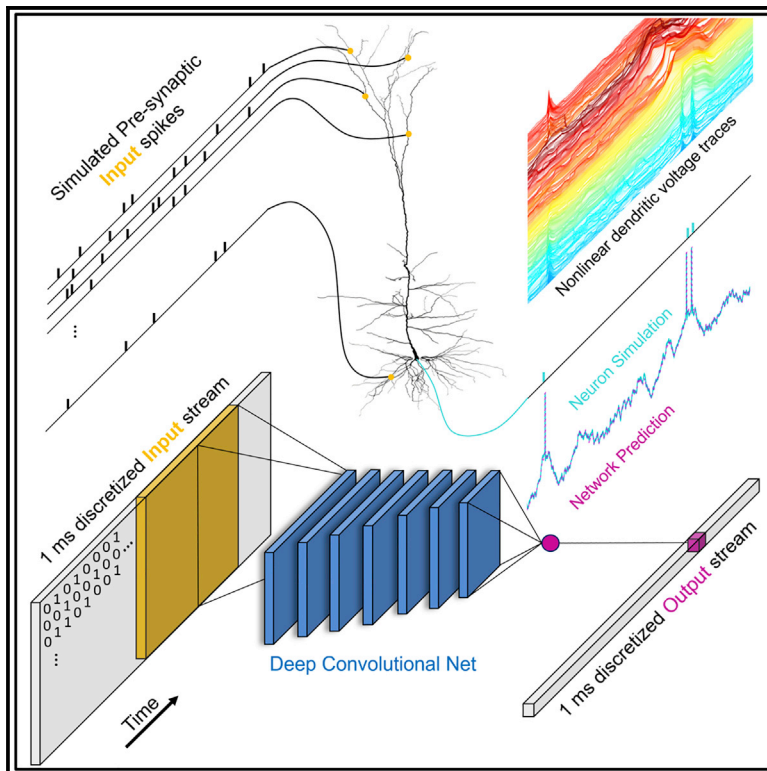


Neuron

Single cortical neurons as deep artificial neural networks

Graphical abstract



Authors

David Beniaguev, Idan Segev,
Michael London

Correspondence

david.beniaguev@gmail.com

In brief

Using a modern machine learning approach, we show that the I/O characteristics of cortical pyramidal neurons can be approximated, at the millisecond resolution (single spike precision), by a temporally convolutional neural network with five to eight layers. This computational complexity stems mainly from the interplay between NMDA receptors and dendritic morphology.

Highlights

- Cortical neurons are well approximated by a deep neural network (DNN) with 5–8 layers
- DNN's depth arises from the interaction between NMDA receptors and dendritic morphology
- Dendritic branches can be conceptualized as a set of spatiotemporal pattern detectors
- We provide a unified method to assess the computational complexity of any neuron type

Article

Single cortical neurons as deep artificial neural networks

David Beniaguev,^{1,3,*} Idan Segev,^{1,2} and Michael London^{1,2}

¹Edmond and Lily Safra Center for Brain Sciences (ELSC), The Hebrew University of Jerusalem, Jerusalem 91904, Israel

²Department of Neurobiology, The Hebrew University of Jerusalem, Jerusalem 91904, Israel

³Lead contact

*Correspondence: david.beniaguev@gmail.com

<https://doi.org/10.1016/j.neuron.2021.07.002>

SUMMARY

Utilizing recent advances in machine learning, we introduce a systematic approach to characterize neurons' input/output (I/O) mapping complexity. Deep neural networks (DNNs) were trained to faithfully replicate the I/O function of various biophysical models of cortical neurons at millisecond (spiking) resolution. A temporally convolutional DNN with five to eight layers was required to capture the I/O mapping of a realistic model of a layer 5 cortical pyramidal cell (L5PC). This DNN generalized well when presented with inputs widely outside the training distribution. When NMDA receptors were removed, a much simpler network (fully connected neural network with one hidden layer) was sufficient to fit the model. Analysis of the DNNs' weight matrices revealed that synaptic integration in dendritic branches could be conceptualized as pattern matching from a set of spatiotemporal templates. This study provides a unified characterization of the computational complexity of single neurons and suggests that cortical networks therefore have a unique architecture, potentially supporting their computational power.

INTRODUCTION

Neurons are the computational building blocks of the brain. Understanding their input/output (I/O) transformation has therefore been a major quest in neuroscience since Ramon y Cajal's "neuron doctrine." With the recent development of sophisticated genetic, optical, and electrical techniques, it has become clear that many key neuronal types (e.g., cortical and hippocampal pyramidal neurons, cerebellar Purkinje cells) are highly complicated I/O information processing devices. They receive a barrage of thousands of synaptic inputs via their elaborated dendritic branches; these inputs interact with a plethora of local nonlinear regenerative processes, including the back-propagating (Na^+ -dependent) action potential (Stuart and Sakmann, 1994), the multiple local dendritic NMDA-dependent spikes (Schiller et al., 2000; Polsky et al., 2004; Branco et al., 2010; Kastellakis et al., 2015), and the large and prolonged Ca^{2+} spike at the apical dendrite of layer 5 (L5) cortical pyramidal neurons (Schiller et al., 1997; Larkum et al., 1999). The input synapses interact with these local nonlinear dendritic properties to eventually generate a train of output spikes in the neuron's axon, carrying information that is communicated, via synapses, to thousands of other (postsynaptic) neurons. Indeed, as a consequence of their inherent nonlinear mechanisms, neurons can implement highly complicated I/O functions (Bar-Ilan et al., 2013; Behabadi and Mel, 2014; Cazé et al., 2013; Doron et al., 2017; Häusser and Mel, 2003; Hawkins and Ahmad, 2016; Katz et al., 2009; Koch and Segev, 2014; Koch et al., 1982; London and Häusser, 2005; Mel, 1992; Moldwin and Segev, 2018;

Poirazi et al., 2003b, 2003a; Shepherd et al., 1985; Tzivilaki et al., 2019; Zador et al., 1991; see recent work on nonlinear dendritic computations in human cortical neurons in Gidon et al., 2020).

A classical approach to study the I/O relationship of neurons is to construct a simplified model that omits many of their detailed biological mechanisms. These models present a highly reduced phenomenological abstraction of the neuron's I/O characteristics (Lapicque, 1907; McCulloch and Pitts, 1943). One such abstraction is the "perceptron" (Rosenblatt, 1958), which lies at the heart of some of the most advanced pattern recognition techniques to date (LeCun et al., 2015). However, the perceptron's basic function, a linear summation of its inputs and thresholding for output generation, ignores the nonlinear synaptic integration processes and the temporal characteristics of the output, which take place in real neurons. Some more recent modeling studies have addressed this gap (Gütig and Sompolinsky, 2006; Poirazi et al., 2003a; Polsky et al., 2004; Ujfalussy et al., 2018) but have either not considered fully diverse synaptic inputs distributed over the full nonlinear dendritic tree or did not aim to capture the I/O transformation of neurons at a millisecond temporal precision of output spikes. Attempts to predict the spiking activity of neurons in response to somatic input current/conductance, rather than dendritic synaptic input, can be found in Jolivet et al. (2008) and Naud et al. (2014), and attempts to predict the spiking activity of neurons in response to natural images can be found in Cadena et al. (2019) and Keat et al. (2001).

Another common approach to study the I/O characteristics of neurons is to simulate, via a set of partial differential equations,

the fine electrical and anatomical details of the neurons using the cable and compartmental modeling methods introduced by Rall (Rall, 1959, 1964; Segev and Rall, 1988). Using these models, it is possible to account for nearly all of the above experimental phenomena and explore conditions that are not accessible with current experimental techniques. While this is the only method to date to account for the full I/O transformation in a neuron, this success comes at a price. Compartmental and cable models are composed of a high-dimensional system of coupled nonlinear differential equations, which is notoriously challenging to understand (Strogatz, 2001). Specifically, it is a daunting task to extract general principles that govern the transformation of thousands of synaptic inputs to a train of spike output at millisecond precision from such detailed simulations (but see Amsalem et al., 2020; Larkum et al., 2009; Magee and Johnston, 1995; Rapp et al., 1992; Schiller et al., 2000; Spruston et al., 1995; Stuart and Sakmann, 1994; Stuart et al., 1997; Wybo et al., 2021).

Here, we propose a novel approach to study the neuron as a sophisticated I/O information processing unit by utilizing recent advances in the field of machine learning. Specifically, we exploited the capability of deep neural networks (DNNs) to learn very complex I/O mappings (Holden et al., 2019; Kasim et al., 2020; Senior et al., 2020), in this case, that of neurons. Toward this end, we trained DNNs with rich spatial and temporal synaptic input patterns to mimic the I/O behavior of a L5 cortical pyramidal neuron model with its full complexity, including its elaborated dendritic morphology, the highly nonlinear local dendritic membrane properties, and a large number of excitatory and inhibitory inputs that bombard the neuron. Consequently, we obtained a highly computationally efficient DNN model that faithfully predicted this neuron's output at millisecond (spiking) temporal resolution. We then analyzed the weight matrices of the DNN to gain insights into the mechanisms that shape the I/O function of cortical neurons. By systematically varying the DNN size, this approach allowed us to characterize the functional complexity of a single biological neuron, pin down the ion-channel-based and morphologically based origin of this complexity, and examine the generality of the resultant DNN to synaptic inputs that were outside of the training set distribution. We demonstrated that cortical pyramidal neurons, and the networks they form, are potentially computationally much more powerful and "deeper" than previously assumed.

RESULTS

Analogous DNN for integrate and fire (I&F) neuron model: Method overview and filters interpretation

Our goal is to fit the I/O relationship of a detailed biophysical neuron model by an analogous DNN. This DNN receives, as a training set, the identical synaptic input and the respective axonal output of the biophysical model. By changing the connection strengths of the DNN using a backpropagation learning algorithm, the DNN should replicate the I/O transformation of the detailed model. To accommodate the neuron's temporal aspect, we employed temporal-convolutional networks (TCNs) throughout the study.

Figure 1 illustrates this paradigm's feasibility and usefulness as a first demonstrative step, starting with the I/O transformation

of a well-understood neuron model: the I&F neuron (Burkitt, 2006; Lapicque, 1907). This neuron receives a train of random synaptic inputs and produces a subthreshold voltage response as well as a spiking output (see STAR Methods). While this I/O transformation appears to be simple, it is unclear whether it can be learned from data by an artificial neural network using the backpropagation algorithm at millisecond temporal resolution with a compact architecture (this indeed has not been previously demonstrated). If successful in achieving a high degree of accuracy with a simple DNN for the I&F model, it will demonstrate that our approach, consisting of the specific way we represent the neuronal I/O data and subsequent fitting of an DNN on these data, is able to represent the functional relationship of the I&F neuron model compactly.

What is the simplest DNN that faithfully captures the I/O properties of this most basic single-neuron model? To answer this question, we constructed a minimal "DNN" consisting of one hidden layer with a single hidden unit (Figure 1A) and verified that it does indeed capture the complexity of this simple neuron model (Figure 1F). The time axis was divided into 1 ms bins in which only a single spike can occur in the I&F neuron model. The objective of the network is to predict the binary spike output of the I&F model at time t_0 , based on the preceding input spike trains (the time-window history) up to t_0 . This input is represented using a binary matrix of size $N_{syn} \times T$, where N_{syn} is the number of input synapses and T is the number of preceding time bins considered (Figure 1B). We used $N_{syn} = 100$, and trained a DNN with a single hidden unit on 7,200 s of simulated data. When using $T = 80$ ms, we achieved a very good fit, namely, a simple DNN with a single hidden unit that accurately predicted both the subthreshold voltage dynamics as well as the spike output of the respective I&F neuron model at millisecond precision (Figure 1C).

Figure 1D depicts the weights ("filters") of the single hidden unit of the respective DNN as a heatmap. It shows that the learning process automatically produced two classes of weights (filters), one positive and one negative, corresponding to the excitatory and inhibitory inputs impinging on the I&F model. In agreement with our understanding of the I&F model, the excitatory inputs contribute positively to output spike prediction (red), whereas the inhibitory inputs contribute negatively to it (blue). Earlier inputs, either inhibitory or excitatory, contribute less to this prediction (teal). Figure 1E depicts the temporal cross-section of those filters and reveals an exponential profile that reflects the temporal decay of postsynaptic potentials in the I&F model (in the reverse time direction). From these filters, one can recover the precise membrane time constant of the I&F model. These two temporal filters (excitatory and inhibitory) are easily interpretable, as they agree with our previous understanding of the temporal behavior of synaptic inputs that give rise to an output spike in the I&F model.

Figure 1F quantifies the model performance in terms of spike prediction using the receiver operating characteristic (ROC) curve (Figure 1F, left; see STAR Methods) and the area under it (area under the curve [AUC]). The AUC for the I&F case is 0.9973, indicating a very good fit. Figure 1F (right) shows an additional quantification of spike temporal precision using the DNN prediction by plotting the cross-correlation between the

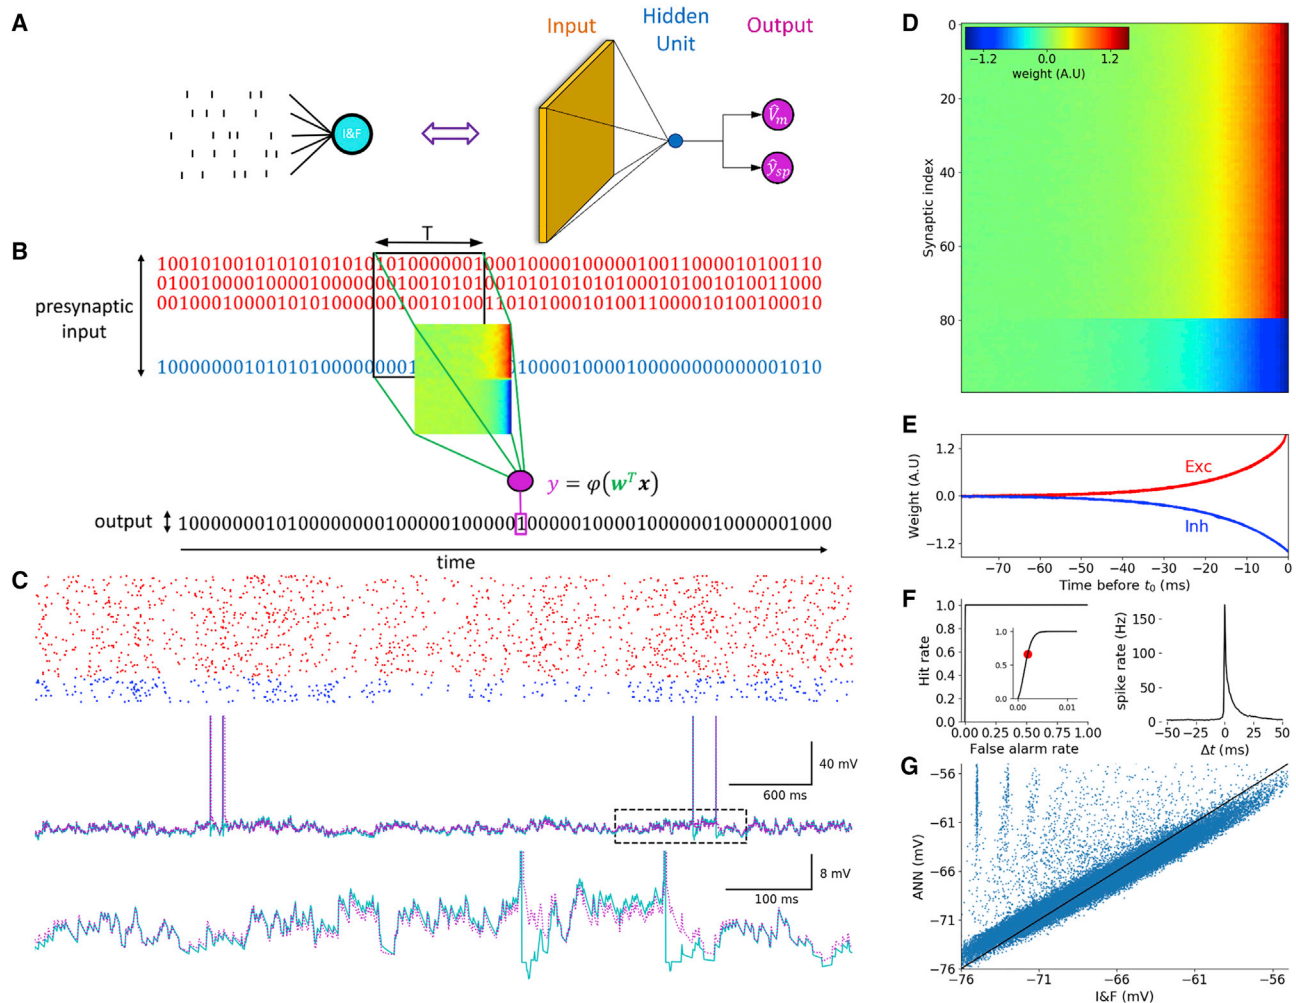


Figure 1. The integrate and fire (I&F) neuron model is faithfully captured by a NN with one hidden layer consisting of a single hidden unit
 (A) Illustration of an I&F neuron model receiving a barrage of random synaptic inputs and generating voltage and spiking output (left) and its analogous DNN (right). Orange, blue, and magenta represent the input layer, the hidden layer, and the DNN output, respectively.
 (B) Schematic overview of our prediction approach. The objective of the DNN is to predict the spike output of the respective I&F model based on its synaptic input. The binary matrix, denoted by x , represents the input spikes in a time window T (black rectangle) preceding t_0 . x is multiplied by the synaptic weight matrix, w (represented by the heatmap image), and summed up to produce the activation value of the output unit, y . This value is used to predict the output (magenta rectangle) at $t = t_0$. Excitatory inputs are denoted in red and inhibitory in blue. Note that unlike the I&F, the DNN has no *a priori* information about the type of the synaptic inputs (excitation or inhibition).
 (C) Top: example inputs (red, excitatory; blue, inhibitory) presented to the I&F neuron model. Middle: response of the I&F model (cyan) and the analogous DNN (magenta). Bottom: zoom in on the dashed-rectangle region in the top trace. Note the great similarity between the two traces.
 (D) Learned weights of the DNN modeled synapses. The top 80 rows are excitatory synapses to the I&F model; the bottom 20 rows are its inhibitory synapses. Columns correspond to different time points relative to t_0 (rightmost time point). The prediction probability for having a spike at t_0 increases if the number of active excitatory synapses increases just before t_0 (red) and the number of active inhibitory synapses decreases (blue).
 (E) Temporal cross-section of the learned weights in (D).
 (F) Left: receiver operator characteristic (ROC) curve of spike prediction. The area under the curve (AUC) is 0.997, indicating high prediction accuracy at 1 ms precision. Inset: zoom in on up to 1% false alarm rate. The red circle denotes the threshold selected for the DNN model shown in (C). Right: cross-correlation between the I&F spike train (ground truth) and the predicted spike train of the respective DNN, when the prediction threshold was set to 0.2% false-positive (FP) rate (red circle in left plot).
 (G) Scatterplot of the predicted DNN subthreshold voltage versus ground-truth voltage produced by the I&F model.

predicted spike train and the target I&F simulated spike train (the “ground truth”). The cross-correlation shows a sharp peak at 0 ms and has a short (~ 10 ms) half-width, suggesting high temporal accuracy of the DNN. We also quantified the DNN performance in predicting the subthreshold membrane potential by

using standard regression metrics, and, in Figure 1G, depict the scatterplot of the predicted voltage versus the ground-truth simulated output voltage. The root mean square error (RMSE) is 1.73 mV (79.8% variance explained), indicating a good fit between the I&F and the respective DNN. Note that high accuracy

on both binary spike prediction and continuous somatic voltage is a dual prediction attempt achieved with only a single hidden unit that is enforcing a strict bottleneck. This is possible only due to the strong relationship between outputs spikes and somatic voltages in the I&F model case.

In conclusion, as a proof of concept, we have demonstrated that a very simple DNN can learn the I/O transformation of an I&F model with a high degree of temporal accuracy. Importantly, the resulting weight matrix (the filter) obtained by the learning process is interpretable, as it represents known features of the I&F model, including the existence of two classes of inputs (excitatory and inhibitory), the convolution of the synaptic inputs with the exponential decay representing the passive membrane properties (resistance, capacitance), and the transformation from subthreshold membrane potential to spike output.

Analogous DNN for the full complexity of the L5 cortical pyramidal neuron model

We next applied our paradigm to a morphologically and electrically complex detailed biophysical compartmental model of a 3D-reconstructed L5 cortical pyramidal cell (L5PC) from rat somatosensory cortex (Figure 2A). The model is equipped with complex nonlinear membrane properties, a somatic spike generation mechanism, and an excitable apical nexus capable of generating calcium spikes (Hay et al., 2011; Larkum et al., 1999; Schiller et al., 1997). The excitatory synaptic inputs are mediated through both voltage-independent AMPA-based conductance and voltage-dependent NMDA-type conductance (Jahr and Stevens, 1993); the inhibitory inputs are mediated through conductance-based GABA_A-type synapses. Both excitatory and inhibitory synapses are uniformly distributed across the dendritic tree of the model neuron (see STAR Methods and Figures S1 and S4 for more details). The training data consisted of a combined total of more than ~200 h of simulated time, with excitatory and inhibitory inputs randomly activated in time according to a Poisson distribution with a firing rate consisting of a piecewise constant temporal trajectory.

A thorough search of configurations of deep and wide fully connected neural network (FCN) architectures have failed to provide a good fit to the I/O characteristics of the L5PC model. These failures suggest a substantial increase in the complexity of I/O transformation compared to that of I&F model. Indeed, only a TCN architecture with seven layers (depth), 128 channels per layer (width), and $T = 153$ ms (history), provided a high precision fit (Figures 2B, 2C, and S2). The example in Figure 2C shows that this TCN can predict the somatic subthreshold voltage and spikes of a highly complex neuron with high precision when provided with a previously unseen input pattern from the test set. Although this was the first configuration of a network that met our criterion for a fit, we consequently managed to find other DNNs that provided comparable results. An extended analysis of the DNN depth, width, and time-window history required to replicate the I/O of this L5PC model faithfully is presented below.

It is important to note that the model's accuracy was relatively insensitive to the temporal kernel sizes of the different DNN layers when keeping the total temporal extent of the entire network fixed. Therefore, the first layer's temporal extent was selected to be larger than the subsequent layers, mainly for visu-

alization purposes (see Figure 2G–2I). Figure 2H shows a filter from a unit in the first layer of the DNN. This filter is somewhat similar to the filter in Figure 1D but integrates only basal and oblique subtrees and ignores the apical tree's inputs. Moreover, the filters have different shapes, representing the differential contribution of inputs arriving at different distances from the soma as predicted by cable theory for dendrites (Rall 1967). Figure 2I, however, shows a filter of another unit that, in contrast to the filter in Figure 2H, has negligible weights assigned for basal and oblique dendrites but a very strong apical tuft dependency. By examining additional first layer filters (not shown), we found a wide variety of different activation patterns that the TCN utilized as an intermediate representation, including many temporally directionally selective filters (similar to those shown in Figure 5D below). Figures 2D–2F show the quantitative performance evaluation of this DNN model. For binary spike prediction (Figure 2D), the AUC is 0.9911. For somatic voltage prediction (Figure 2E), the RMSE is 0.71 mV, and 94.6% of the variance is explained by this model. Note that despite its seemingly large size, the resulting TCN represents a substantial decrease in computational resources relative to a full simulation of a detailed biophysical model (involving numerical integration of thousands of nonlinear differential equations), as indicated by a speedup of simulation time by several orders of magnitude.

NMDA synapses are major contributors to the I/O complexity (“depth”) of L5PCs

Now that we have obtained a DNN model that can replicate the I/O relationship of a detailed biophysical/compartmental model of a real neuron very accurately, can we learn from it what the essential features that contribute to neuron complexity are? Detailed studies of synaptic integration in dendrites of cortical pyramidal neurons suggested the primary role of the voltage-dependent current through synaptic NMDA receptors, including at the subthreshold and suprathreshold (NMDA spike) regimes (Polsky, Mel, and Schiller 2004; Branco, Clark, and Häusser 2010). As NMDA receptors (NMDARs) depend nonlinearly on the voltage, they are highly sensitive not only to the activity of the synapse in which they are located but also to the activity of (and the voltage generated by) neighboring synapses and their dendritic location. Moreover, the NMDA current has slow dynamics, promoting integration over a time window of tens of milliseconds (Doron et al., 2017; Jahr and Stevens, 1993; Major et al., 2013). Consequently, we hypothesized that removing NMDA-dependent synaptic currents from our L5PC model will significantly decrease the size of the respective DNN to achieve similar levels of accuracy, implying a reduction in the complexity of the I/O transformation.

In Figure 3, we present the results of a new set of simulations where the NMDA voltage-dependent conductances were removed, such that the excitatory input relies only on AMPA-mediated conductances. We compensated for the significant reduction of excitatory current resulting from the NMDA removal by adjusting the input firing rate of the AMPA-type synapses to maintain the same average output firing rate of the L5PC neuron model as in Figure 2 (see precise details in Figure S4). The figure shows that for this model, we have managed to achieve a similar quality fit as in Figure 2 with a much smaller (and shallower)

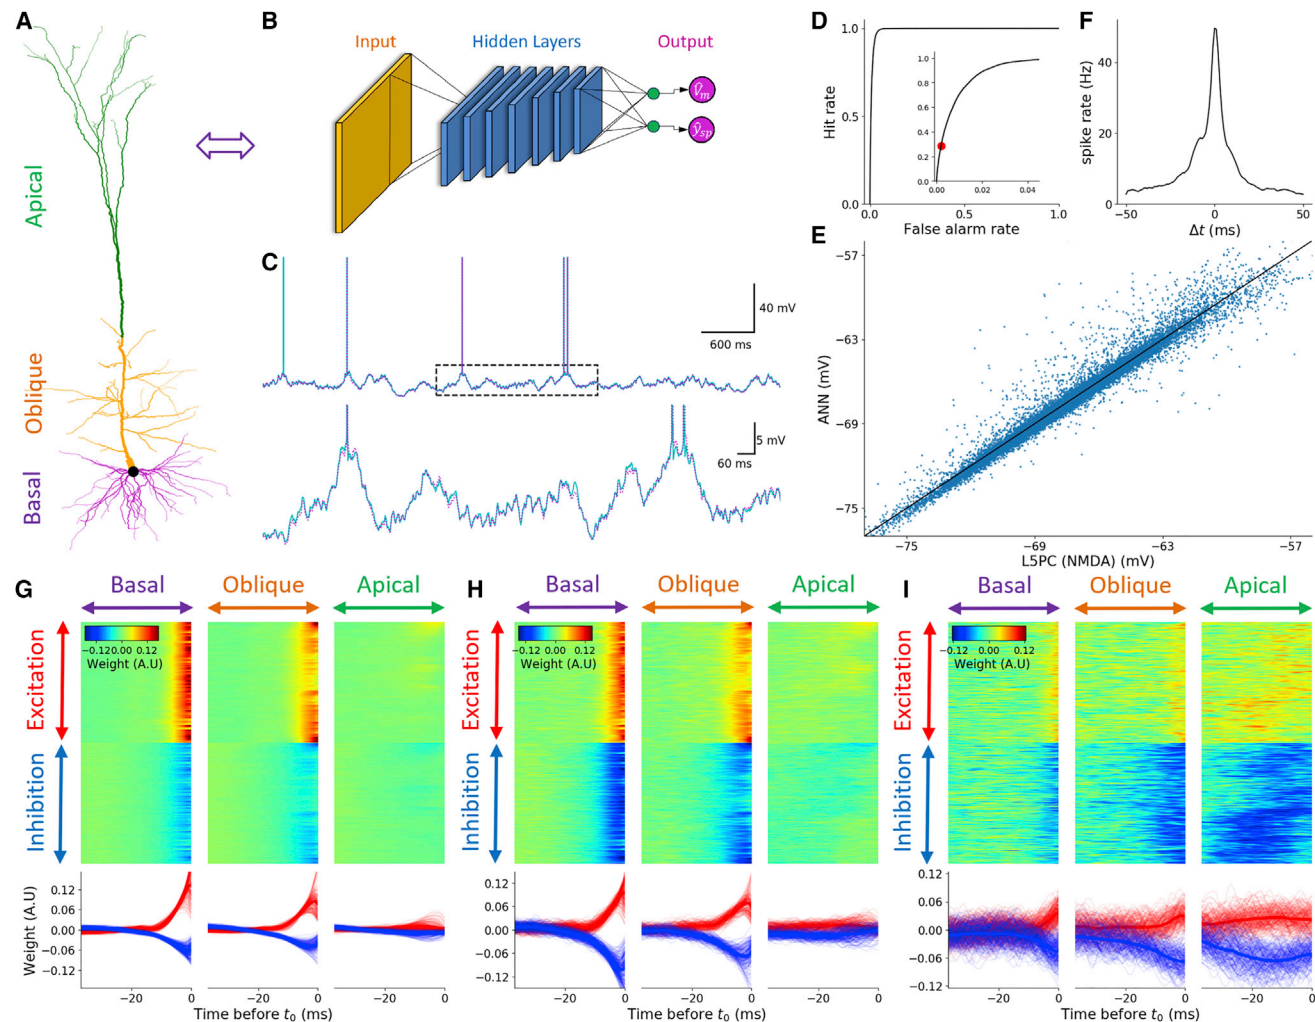


Figure 2. A detailed model of an L5 cortical pyramidal neuron with AMPA and NMDA synapses is faithfully captured by a TCN with seven hidden layers consisting of 128 feature maps per layer and a history of 153 ms

For a Figure360 author presentation of this figure, see <https://doi.org/10.1016/j.neuron.2021.07.002>.

(A) The modeled L5PC. Basal, oblique, and apical dendrites are marked in purple, orange, and green, respectively.

(B) Analogous DNN with seven hidden layers. Orange, blue, and magenta circles represent the input layer, the hidden layer, and the DNN output, respectively. Green units represent linear activation units (see STAR Methods).

(C) Top: example voltage response of the L5PC model with AMPA and NMDA synapses (cyan) and the analogous DNN (magenta) to random synaptic input. Bottom: zoom in on the dashed-rectangle region in the top trace.

(D) ROC curve of spike prediction; the AUC is 0.9911, indicating high prediction accuracy at 1 ms precision. A zoom in on up to 4% false alarm rates is shown in the inset. The red circle denotes the threshold selected for the DNN model shown in (B).

(E) Scatterplot of the predicted DNN subthreshold voltage versus ground-truth voltage.

(F) Cross-correlation plot between the ground truth (L5PC model with AMPA and NMDA synapses) spike train, and the predicted spike train of the respective DNN, when the prediction threshold was set according to the red circle in (D).

(G) Learned weights of a selected unit in the first layer of the DNN. The top left, center, and right panels show inputs located on the basal dendrites, oblique dendrites, and the apical tuft, respectively. For each case, excitatory synapses are shown in the top half of the rows, whereas inhibitory synapses are shown in the bottom half. Different columns correspond to different time points relative to t_0 (rightmost time point). Bottom: temporal cross section of the learned weights above.

(H) Similar to (G), first layer weights for a different unit in the first layer but with a different spatiotemporal pattern.

(I) An additional unit that is weakly selective to whatever happens in the basal dendrites, weakly sensitive to oblique dendrites, but very sensitive to apical tuft dendrites. The output of this hidden unit is increased when there are apical excitation and a lack of apical inhibition in a time window of 40 ms before t_0 . Note the asymmetry between the amplitudes of the temporal profiles of excitatory and inhibitory synapses, indicating that inhibition decreases the activity of this unit more than excitation increases it.

See also Figure S1 for more details about simulation.

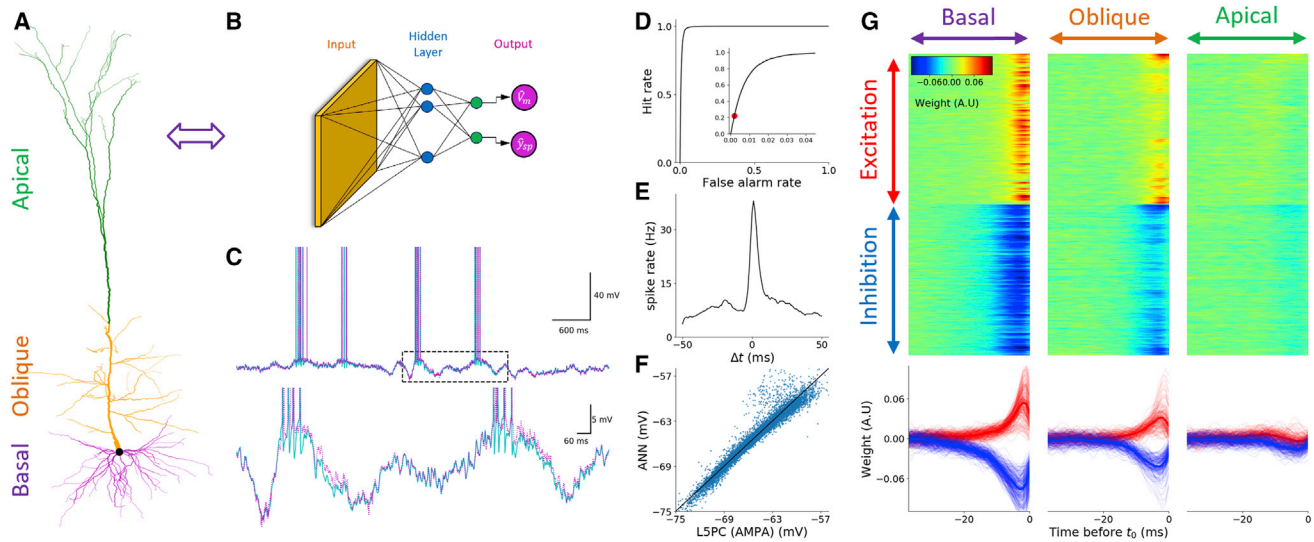


Figure 3. A detailed model of an L5PC neuron with AMPA synapses is faithfully captured by a DNN with one hidden layer consisting of 128 hidden units

(A) Illustration of the L5PC model. Excitation in this model is mediated only by AMPA conductances.
 (B) Analogous DNN. As in Figure 2, orange, blue, and magenta circles represent the input layer, the hidden layer, and the DNN output, respectively. Green units represent linear activation units.
 (C) Top: response of the L5PC model (cyan) and the analogous DNN (magenta) to random AMPA-based excitatory and GABA_A-based inhibitory synaptic input (see STAR Methods). Bottom: zoom in on the dashed-rectangle region in the top trace. Note the great similarity between the two traces.
 (D) ROC curve of spike prediction; the AUC is 0.9913, indicating high prediction accuracy at 1 ms precision. Inset: zoom in on up to 4% false alarm rates. The red circle denotes the threshold selected for the DNN model shown in (B).
 (E) Cross-correlation plot between the ground truth (L5PC model response) and the predicted spike train of the respective DNN for prediction threshold indicated by the red circle in the left plot.
 (F) Scatterplot of the predicted DNN subthreshold voltage versus ground-truth voltage.
 (G) Learned weights of selected units in the DNN, separated by their morphological (basal, oblique, and apical) location. Like in Figure 2, in each case, excitatory synapses are shown in the top half of the rows, and the bottom half shows inhibitory synapses. As in Figure 1D, different columns correspond to different time points relative to t_0 (rightmost time point). Note that just before t_0 , the output of this hidden unit increases if the number of active excitatory synapses increases at the basal and oblique dendrites (red), whereas the number of active inhibitory synapses decreases (blue) at these locations. However, this unit is nonselective to activity at the apical tuft, indicating the lack of influence of the tuft synapses on the neuron's output by this unit.
 See also Figures S3 and S4 for more details about simulation and fitting results with a larger TCN.

network. The network consists of a fully connected DNN (FCN) with 128 hidden units and only a single hidden layer (Figure 3B) and $T = 43$ ms (history). This significant reduction in complexity is due to the ablation of NMDA channels. Also, in our DNN training attempts, we have failed to achieve a good fit when using the smaller architecture that was successful for the I&F model neuron shown in Figure 1. This indicates that whereas the DNN model for L5PC is greatly simplified in the absence of NMDA conductance, additional neuronal mechanisms still contribute to the richness of its I/O transformation of the L5PC as compared to that of the I&F neuron model.

Figure 3C shows an exemplar test trace for the DNN illustrated in Figure 3B, whereas Figure 3H depicts a representative exemplar of the weight matrix for one of the hidden units of the DNN. By examining the filters of the hidden layer of the DNN, we observed that the weights representing inputs to the oblique and basal dendrites had profiles that resemble postsynaptic potentials (PSPs; mirrored in time). Interestingly, the weights associated with synapses of the apical tuft are essentially zero. This pattern remains consistent for all first layer filters of the network, implying that, for this model, the apical dendritic synapses had

negligible information regarding predictions of the output spikes of the neuron, even in the presence of calcium spikes occasionally occurring in the nexus. Contrasting this filter with the one presented in Figure 2I suggests that the NMDA nonlinearity greatly assists in the activation of apical tuft dendrites. Figures 3D–3F show the quantitative performance evaluation. For binary spike prediction (Figure 3E), the AUC is 0.9913. For somatic voltage prediction (Figure 3F), the RMSE is 0.58 mV, and 95.0% of the variance is explained by this model.

In order to systematically compare the complexity of DNNs for the AMPA-only case (Figure 3) versus the AMPA and NMDA case (Figure 2), we selected a minimal approximation threshold (AUC = 0.9910) as a “good enough” performance threshold for spiking accuracy of the DNN as compared to the spiking of the respective biophysical model. DNNs that performed better than this threshold were considered to be “a good approximation.” We then asked what is the minimal sized network that satisfies this threshold for AMPA-only case versus the AMPA and NMDA case. Toward this end, we trained a total of 137 DNNs (62 for AMPA only and 75 for AMPA and NMDA) while varying three main hyperparameters: depth, width [number of channels per

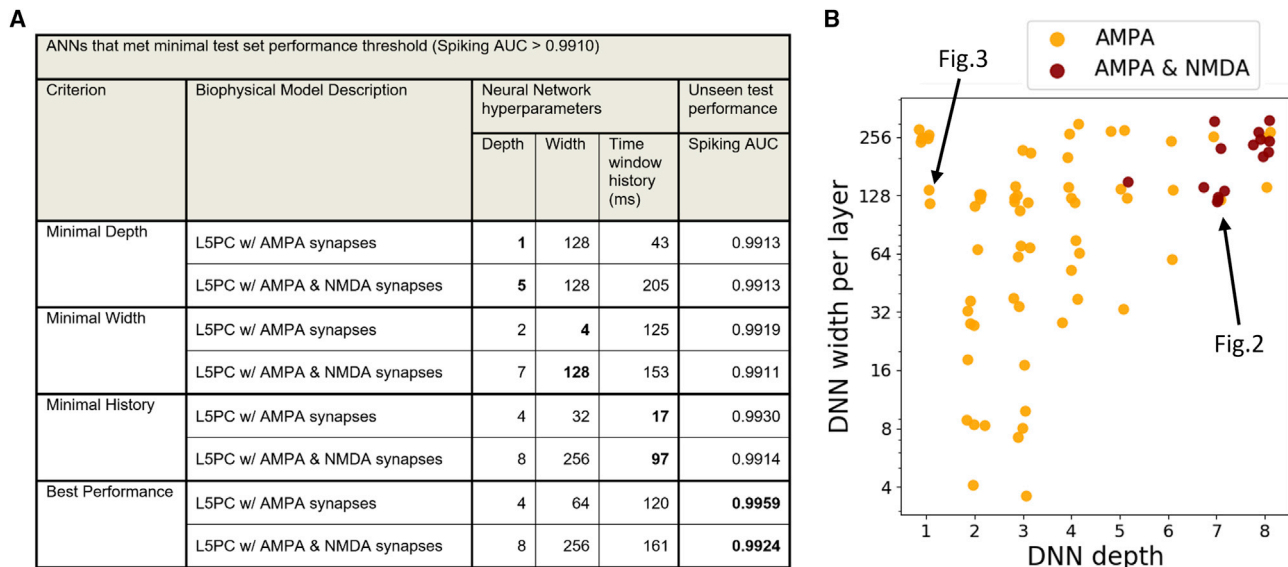


Figure 4. The minimal DNN size required to achieve a good fit is larger for AMPA and NMDA synapses compared to AMPA-only synapses across all tested hyperparameters

(A) Deep networks of various depths, widths, and maximal temporal history dependencies were fitted on two L5PC models, one with AMPA and NMDA synapses and the other with only AMPA synapses. The table summarizes the results by showing the minimal depth, width, and temporal dependence extent DNNs for each of the two biophysical model variations that successfully passed a minimal “good approximation” threshold of AUC = 0.9910. A noticeably clear trend arises that, with NMDA-based synapses, DNNs of depth 5–8 are required to achieve this threshold criterion, whereas shallower DNNs are required for the case of AMPA-only synapses.

(B) 2D scatterplot jointly depicting depth and width of DNNs that meet minimal performance criteria as in (A) for the two cases among all attempted networks for AMPA-only synapses (orange) and AMPA and NMDA synapses (dark red). A small Gaussian noise was added to depth and width axes for visualization purposes to avoid overlap of the respective points. Arrows indicate the networks utilized in Figures 2 and 3. See also Figure S2 for more details.

layer], and temporal extent of the input history used to make the prediction. Figure 4A depicts a table that summarizes the results. It is clear that in order to reach similar levels of accuracy, much smaller DNNs are required for the AMPA case for all three hyperparameters considered. Figure 4B depicts a 2D scatterplot of both depth and width of DNNs that meet the minimum performance threshold among all attempted hyperparameter configurations. It is evident that only networks with large depths and widths can achieve this minimal level of accuracy for the AMPA and NMDA case (dark red dots), whereas much smaller (shallower) networks are sufficient for the AMPA-only case (orange dots). Note that although the precise details of the results will vary when selecting a different AUC threshold value, the overall trend stays the same. For an extended comparison, see Figure S2.

To investigate further the relationship between NMDARs and their interaction with dendritic morphology, as well as the influence of NMDAR density on I/O transformation complexity, we conducted a set of additional analyses. As shown in Figures S5A–S5C, we examined the I/O complexity when the AMPA and NMDA synapses were placed only on part of the L5PC morphology. We examined four such cases whereby the synapses impinge on different regions of the dendritic tree; these cases are strictly encompassed within each other so that we have a clear axis of complexity (full morphology, excluding tuft synapses, synapses on the basal tree only, and synapses

placed only proximal compartments of the basal tree). Figure S5G shows a summary of these results, clearly indicating a reduction in I/O complexity as morphology containing synapses is progressively more restricted. In Figures S5D–S5F, we examine the interaction between the case where NMDA synapses are placed on proximal and distal parts of the basal tree as compared to the respective AMPA-only case. Figure S5H shows a summary of these results, which indicate that there exists an additive interaction between synapse type and morphology, specifically that segregated dendritic compartments with synapses are harder to model also in the AMPA-only case; the nonlinearities of NMDAR synapses in this case add an additional complexity to the modeling effort. In Figure S6, we examine the dependence of the I/O relationship on the maximal NMDA conductance. It is evident (Figure S6D) that even a small increase in NMDA conductance has a significant impact on I/O complexity and that the I/O complexity increases as the NMDAR density increases further, albeit with “diminishing returns.”

DNN analysis of a single dendritic branch provides new insights for the contribution of NMDA conductance to the computational complexity of neurons

To deepen the understanding of the contribution of NMDA synapses to the computational complexity of neurons, we next studied a simplified case. Here, the NMDA synapses were activated

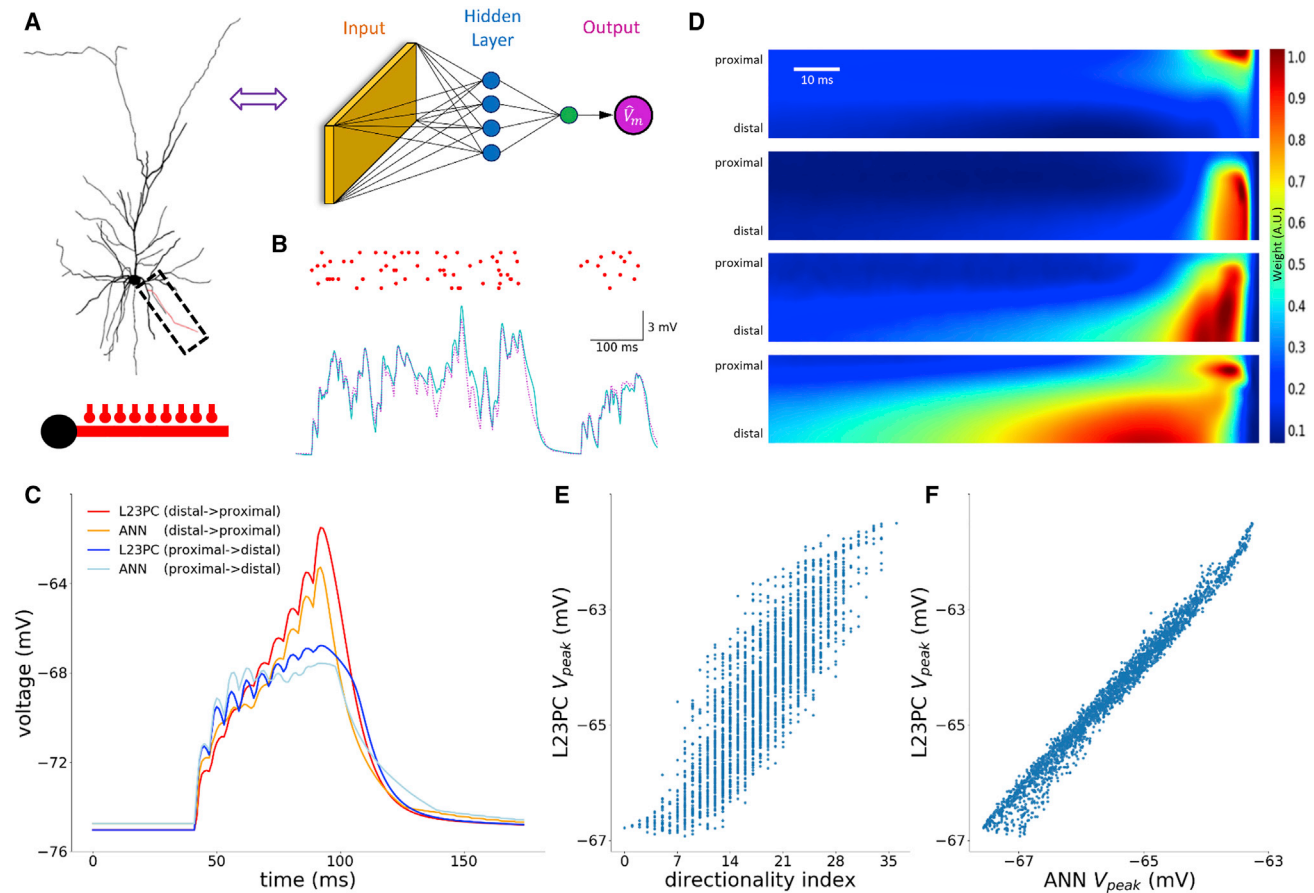


Figure 5. Analysis of a DNN that successfully replicates the I/O of a single L23PC dendritic branch receiving NMDA synapses reveals spatio-temporal pattern matching with four distinct convolutional filters

(A) Left: a layer 2/3 pyramidal neuron was used in the simulations with a zoom in on one selected basal branch (dashed rectangle). This modeled dendritic branch receiving nine excitatory synapses, depicted schematically by the “ball and stick” model at the bottom, was also used in a previous study (Branco et al., 2010). Right: illustration of the analogous DNN that was trained on random synaptic inputs impinging on this basal dendrite. Colors as in Figure 2A.

(B) Example of the somatic voltage response (cyan) and DNN predicted output (magenta) to a random input spike pattern impinging on the basal dendritic branch (red dots above).

(C) Example of somatic response to two spatiotemporal sequences of synaptic activation patterns (red, distal-to-proximal direction; blue, proximal-to-distal direction) and the DNN predicted output for these same sequences (orange and light blue traces, respectively).

(D) Learned weights of the four hidden units consisting the respective DNN model. Heatmaps are spatiotemporal filters, as shown in Figures 1D and 2C. Note the direction-selective shapes and long temporal extent of influence by distal synaptic activations.

(E) Scatterplot that shows the discrimination ability between different temporal orders of synaptic activations on the modeled basal branch. The vertical axis is the ground-truth maximum voltage at the soma during a specific synaptic order of activation. The horizontal axis is the directionality index proposed in Branco et al. (2010). The correlation coefficient is 0.86.

(F) Same as (E), but showing the DNN estimation of the maximum voltage of the respective order of activation. The correlation coefficient is 0.99.

See also Figure S7.

along a single basal branch of a layer 2/3 pyramidal cell (L23PC) from the mouse visual cortex, as in the experimental and modeling study of Branco et al. (2010). This modeled dendrite received a random activation of nine excitatory synapses uniformly distributed across the dendritic length. The temporal activation of each synapse followed a Poisson distribution, and the instantaneous firing rate was identical for all synapses (Figure 5A). We found that the output of this single dendritic branch is faithfully captured by a single layer of a fully connected DNN with four hidden units (Figures 5A and 5B). Examining the four filters of the first layer reveals interesting shapes that make intuitive sense, as first

explored by the pioneering theoretical studies of Rall (1964). The topmost filter in Figure 5D appears to be summing only very recent and proximal dendritic activation. The second-from-top hidden unit sums up recent distal dendritic synaptic inputs. The third filter clearly shows a direction-selective hidden unit, preferring patterns in which synaptic activation is temporally activated sequentially from distal to proximal, whereas the last hidden unit responds to a prolonged distal dendrite summation of activity combined with precisely timed proximal input activation. Thus, the rich and complex integration of inputs on this dendritic branch, which involved the contribution of the nonlinear NMDA

current, can be conceptualized as pattern matching of a set of four specific spatiotemporal templates.

To further examine the generalization capability of the DNNs, [Figure 5C](#) examines the special cases studied by [Branco et al. \(2010\)](#) of a sequential temporal activation of the nine synapses once in the distal-to-proximal direction and, conversely, in the proximal-to-distal direction. Importantly, our DNN network was trained on random synaptic activation patterns and was not exposed to these highly organized spatiotemporal input patterns during training. However, the DNN successfully replicated the response of the L23PC modeled branch to these spatiotemporal sequences. [Figure 5E](#) shows our reconstruction of the results of [Branco et al. \(2010\)](#), whereby a directionality index was suggested as a predictor for the peak somatic voltage for a random activation sequence of the nine input synapses. [Figure 5F](#) shows the much-improved prediction of the respective DNN for the same activation sequences as in [Figure 5E](#). It is important to note that the special case of nine synaptic activations equally spaced in time is highly unlikely to occur during the random input stimulation regime that was used to train the DNN. Nevertheless, as shown in [Figure 5](#), the network can generalize even to this new input regime with high precision.

These results illuminate the interpretability power of our approach. By examining the four kernels (filters), we provide an intuitive ([Figure 5D](#)), yet powerful ([Figure 5F](#)), interpretation for the complex process of nonlinear spatiotemporal synaptic integration in a single dendrite with NMDA synapses. It also further demonstrates the ability of our DNN models to generalize to previously unseen input patterns (out-of-distribution generalization), as will be discussed below.

Generalization of single-neuron analogous DNNs to spatiotemporal structured inputs

The analogous DNN for the L5PC with NMDA synapses shown in [Figure 2](#) was trained on a large set of synaptic inputs that were uniformly distributed across the dendritic trees and randomly activated in time (see [STAR Methods](#)). However, how well does this DNN capture the case of spatially clustered and temporally synchronized inputs that may give rise to highly nonlinear dendritic phenomena (e.g., NMDA spike)? [Figure 6](#) shows that this DNN generalizes very well to a wide range of spatiotemporally structured stimulation protocols without retraining. [Figure 6A](#) depicts the case whereby excitatory and inhibitory synapses impinge on restricted subtrees of the modeled cell (purple dendritic regions); in this case, the temporal patterns of the instantaneous input rates are similar to those during the training of the DNN shown in [Figure 2](#), and only the spatial pattern is altered. The voltage response traces of the L5PC model (in cyan) and the analogous DNN (magenta) are shown at right. Comparisons of the output of the L5PC model and that of the respective DNN for additional spatiotemporal structured synaptic input cases are depicted in [Figures 6B–6D](#), including situations with synchronous temporal input patterns ([Figures 6C](#) and [6D](#)) and without inhibition ([Figures 6B](#) and [6D](#)). In all cases, a close similarity between the subthreshold voltage responses and spiking activity was found (see [Figures S8B](#) and [S8C](#) for aggregate details for all the above-mentioned conditions and combined simulation test time of ~230 min). We conclude

that, albeit trained on a set of input synapses that were activated randomly in time and uniformly distributed over the L5PC model dendrites, this training input set was sufficiently rich so that the respective DNN for this modeled cell successfully captures (generalizes well) the I/O of the L5PC model for a wide range of spatially and temporally structured inputs. We note that the ability to generalize to different input statistics, as depicted in [Figure 6](#), is greatly dependent on the DNN size. Indeed, we have found that the deeper the analogous DNN is, the better it generalizes (see [Figures S8E](#) and [S8F](#) and [Table S1](#) for a full comparative analysis).

DISCUSSION

Recent advances in the field of DNNs provide, for the first time, a powerful general-purpose tool that can learn complex mappings from examples. In this study, we used these tools to study the I/O mappings of single complex nonlinear neurons at millisecond temporal resolution. We constructed a large dataset of pairs of (synaptic) input and (axonal) output examples by simulating a neuron model of L5PC receiving a rich repertoire of synaptic inputs over its dendritic surface and recorded its spike output at millisecond temporal resolution, as well as its somatic sub-threshold membrane potential. We then trained networks of various configurations on these I/O pairs until we obtained an analogous “deep” network with close performance to that of the neuron’s detailed simulation. We applied this framework to a series of neuron models with various levels of morpho-electrical complexity and obtained new insights regarding the computational complexity of cortical neurons.

For simple I&F neuron models, our framework provides simple analogous DNNs with one hidden layer consisting of a single hidden unit that captures the full I/O relationship of the model. The respective DNN filters provided key biophysical insights that are consistent with our understanding of the parameters shaping the I/O relationship of I&F models ([Figure 1](#)). In the case where only a single basal dendritic branch of a cortical neuron, consisting of NMDA synapses, was modeled, a shallow (one hidden layer) DNN with only a few units was required to capture a different aspect of the spatiotemporal integration of synaptic inputs ([Figure 5](#)). Surprisingly, even a model of L5 cortical pyramidal neuron with the full complexity of its dendritic trees and a host of dendritic voltage-dependent currents and AMPA-based synapses is well captured by a relatively simple network with a single hidden layer ([Figure 3](#)). However, in a full model of an L5 pyramidal neuron consisting of NMDA-based synapses, the complexity of the analogous DNN is significantly increased; we found a good fit to the I/O of this modeled cell when using a TCN that has five to eight hidden layers ([Figures 2](#) and [4](#); see also [Figures S2](#) and [S3](#)). Furthermore, a seven-layer analogous DNN for a L5PC that was trained on random inputs successfully generalizes to new out-of-distribution set of clustered and synchronous inputs ([Figure 6](#)).

These results suggest that the single cortical neuron with its nonlinear synaptic inputs is already, on its own, a sophisticated computational unit. Consequently, cortical networks built from such units are deeper and computationally more powerful than they seem to be just based on their anatomical (pre- to post-synaptic) connections. Importantly, the implementation of the I/O

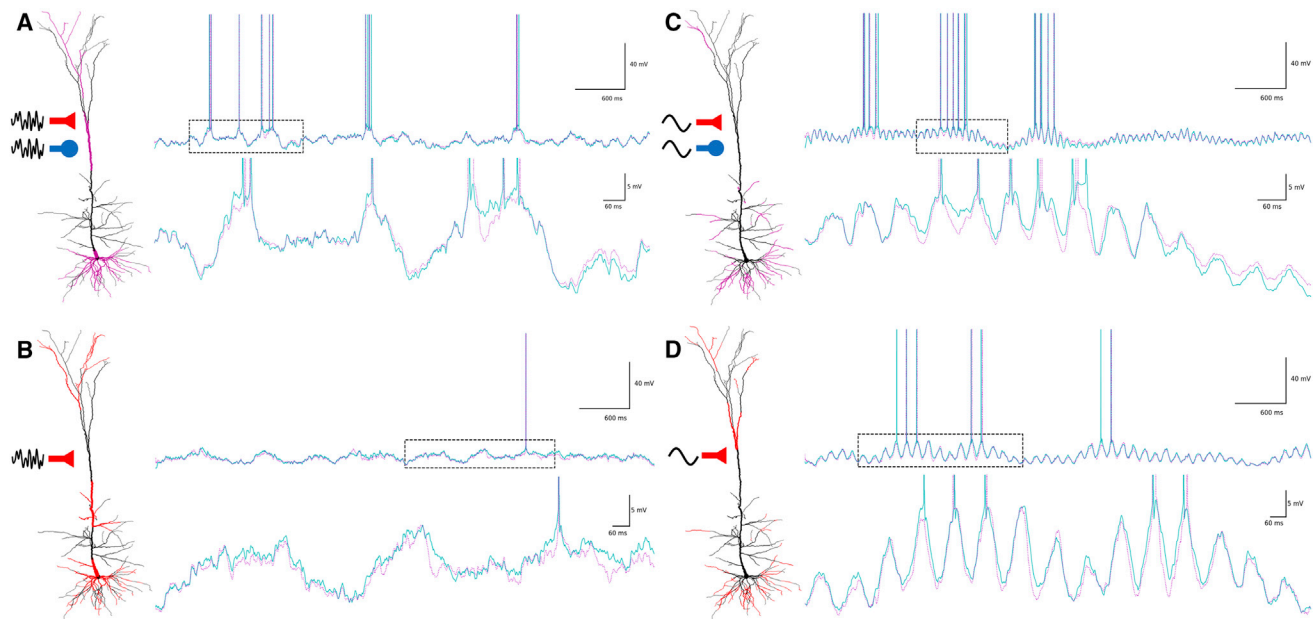


Figure 6. The seven-layer analogous DNN for L5PC that was trained on random inputs successfully generalizes to new out-of-distribution clustered and synchronous inputs

(A) Left: modeled L5PC with AMPA and NMDA synapses as in Figure 2 with excitatory (red) and inhibitory (blue) synapses impinging only on a specific subtree (purple dendritic regions). Instantaneous input rates for the excitatory (E) & inhibitory (I) synapses are similar to that during the training of the analogous DNN. Top right: voltage response traces of the modeled cell (cyan) and the analogous DNN (magenta). Bottom right: zoom in on the dashed-rectangle region in the top trace. (B) As in (A), but this time with excitatory-only synapses impinging on subparts of the dendritic tree (red regions). The temporal profile of instantaneous input rates is like that given during training.

(C) As in (A), but with E&I synapses impinging on different subtrees (purple regions) of the modeled L5PC. Instantaneous rates of excitatory and inhibitory synapses are driven by a synchronous sinusoidal pattern that was not seen during training.

(D) As in (A), but here, only excitatory synapses are impinging on the dendritic tree (red regions) and are activated with synchronous sinusoidal instantaneous input rates.

See also Figure S8.

function of single neurons using a DNN also provides practical advantages. It is computationally much more efficient than the traditional compartmental model, which required the solution of many thousands of partial differential equations (PDEs) per neuron. Indeed, for the full model of L5PCs, we obtained a speedup of $\sim 2,000\times$ when using the DNN instead of its compartmental-model counterpart. However, we did not perform a comprehensive empirical or theoretical comparison of the computational advantage of our DNN method for different neuron models, different input conditions, or different hardware or software versions, as this lies beyond the scope of the present study. In general, the computational speedup of our approach results from (1) compression/simplification (via the computational nodes and layers of the DNN) of all dendritic computations performed in the cell, keeping only the computations that are relevant for output spike generation; and (2) utilization of specialized graphics processing unit (GPU) hardware and a corresponding software ecosystem, which are continuously developed for the increasingly growing needs of the deep-learning community. Together, our tool can potentially be useful for simulating large-scale realistic neuronal networks (Egger et al., 2014; Markram et al., 2015). Furthermore, the size of the respective DNN for a given neuron could be used (under certain assumptions [see below]) as an index for its computational power; the

larger it is, the more sophisticated computations this neuron could perform. Such an index will enable a systematic comparison between different neuron types (e.g., CA1 pyramidal cell, cortical pyramidal cell, and Purkinje cell) or the same type of cell in different species (e.g., mouse versus human cortical pyramidal cells).

It is important to emphasize that for optimization reasons, the complexity of the analogous DNN described above is an upper bound of the true computational complexity of the I/O of the respective single neuron, i.e., it is possible that there exists a smaller DNN that mimics the biophysical neuron with a similar degree of accuracy but the training process we used did not find it. Additionally, we note that we have limited our architecture search space only to FCN and TCN neural network architectures. It is likely that additional architectural search could yield simpler and more compact models for any desired degree of prediction accuracy. Nevertheless, we stress that this upper bound is in fact several orders of magnitude computationally less intensive when compared to the detailed biophysical simulations. In addition, the results presented in Figure 1 clearly indicate that our proposed method would have been able to uncover simple I/O relationships if the I/O transformation of cortical neurons was in fact rather simple. This suggests that the DNN size that we have found might be a relatively tight upper bound. In order to facilitate

this search in the scientific community, we have released our large readymade dataset of simulated inputs and outputs of a fully complex single L5 cortical neuron in an *in-vivo*-like regime so that the community can focus on modeling various aspects of this endeavor and avoid re-running the simulations themselves.

The analysis of DNNs is a challenging task and a rapidly growing field (Mahendran and Vedaldi, 2015; Mordvintsev et al., 2015; Olah et al., 2017). Nevertheless, observing the weight matrix of units (filters) in the first layer of the respective DNN is straightforward and can provide ample insights regarding the I/O transformation of the neuron. The full network can be interpreted as consisting of a basic set of filters that span the space of possible spatiotemporal patterns of synaptic inputs that will drive the original neuron to spike. The first layer defines this space, and the rest of the network mixes and matches within that space. For example, as shown in Figure 3, in the case of a pyramidal neuron without NMDA synapses, most filters have significant weights only for basal and apical oblique inputs, and the weight given for apical tuft synapses is negligible (despite the existence of voltage-dependent Ca^{2+} and other nonlinear currents in this model, which results in occasional Ca^{2+} spikes in the distal dendritic tuft). The picture is fundamentally different when NMDA synapses were included in the model. In this case, the weights assigned to apical dendrite synapses are significant (Figure 2; see also Figure S6). Moreover, the filters devoted to these apical inputs tend to have a temporal structure that is significantly broader (in time) than for that of the proximal synapses, suggesting that the temporal precision of input to the apical synapses is less important for spike generation in the soma. These are basic insights that could be drawn by observing first layer filters of the resulting analogous DNN.

This work opens multiple additional avenues for future research. One important direction is isolating the contribution of specific mechanisms to the computational power of the neuron in a similar way to that performed here for the NMDA-based receptors (see also Figure S2). Fitting a DNN while manipulating specific voltage-dependent dendritic currents (e.g., voltage-gated calcium channels [VGCCs], voltage-gated potassium [Kv] channels or hyperpolarization-activated cyclic nucleotide-gated [HCN] channels), will provide a deeper understanding of their contribution to the overall synaptic integration process and to the complexity of the respective DNN. An additional interesting direction is to utilize this work to explore how real neurons can use their rich biophysical repertoire in order to perform specific computations from the class computed by the equivalent DNNs. By taking advantage of gradient-descent optimization and specialized GPU hardware acceleration, one can efficiently train the DNN representing the neuron to compute an interesting, meaningful function (e.g., training it to classify images of handwritten digits or to classify sequences of auditory sounds). Then it might be possible to map this DNN back to the original biophysical neuron model. One can then both directly validate the hypothesis that single neurons could perform complex and useful computational tasks and investigate how these neurons and specific spatiotemporal distribution of synapses can actually implement such tasks.

If indeed one cortical neuron is equivalent to a multilayered DNN, then what are the implications for the cortical microcircuit?

Is that circuit merely a deeper classical DNN composed of simple “point neurons”? A key difference between the classical DNN and a cortical circuit composed of deep neurons is that, in the latter case, synaptic plasticity can take place mainly in the synaptic (input) layer of the analogous DNN for a single cortical neuron, whereas the weights of its hidden layers are fixed (and dedicated to represent the I/O function of that single cortical neuron). It is important to note, however, that there are other forms of (non-synaptic) plasticity in neurons, such as branch-specific plasticity or intrinsic plasticity (Losonczy et al., 2008) that can, perhaps, shape the weights of deep layers of the analogous DNN in addition to synaptic efficacies. Taken together with the myriad of recurrent connections and network motifs between cortical neurons of different types (Markram et al., 2015), we hereby propose a concrete, biologically inspired network architecture for cortical networks that seamlessly incorporates single-neuron complexity. Focusing on architecture as a key element of research was recently advocated by (Richards et al., 2019). Indeed, the search for the appropriate architecture of artificial neural networks is one of the most rewarding avenues of machine learning today (He et al., 2015; Lin et al., 2014; Vaswani et al., 2017), and studying the specific architecture suggested in the present study may unravel some of the inductive bias hidden within the cortical microcircuit and harness it for future artificial intelligence (AI) applications.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- **KEY RESOURCES TABLE**
- **RESOURCE AVAILABILITY**
 - Lead contact
 - Materials availability
 - Data and code availability
- **METHOD DETAILS**
 - I&F simulation
 - L5PC simulations
 - L23PC simulations
 - DNN fitting
 - Out-of-Distribution (OOD) simulations
- **QUANTIFICATION AND STATISTICAL ANALYSIS**
 - Model evaluation
 - Comparison of AMPA versus NMDA model complexity
 - Assessing the complexity resulting from the interaction between dendritic morphology and NMDA receptor density

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.neuron.2021.07.002>.

ACKNOWLEDGMENTS

We thank Oren Amsalem, Guy Eyal, Michael Doron, Toviah Moldwin, Yair Deitcher, Eyal Gal, and all lab members of the Segev and London Labs for many fruitful discussions and valuable feedback regarding this work. This work was

supported by ONR grant N00014-19-1-2036, Israeli Science Foundation grant 1024/17 (to M.L.), and a grant from the Gatsby Charitable Foundation.

AUTHOR CONTRIBUTIONS

D.B., conceptualization, methodology, investigation, visualization, software, validation, data curation, writing – original draft; I.S. and M.L., conceptualization, methodology, writing – review & editing, supervision, resources, funding acquisition.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: October 7, 2020

Revised: March 4, 2021

Accepted: June 30, 2021

Published: August 10, 2021

REFERENCES

Amsalem, O., Eyal, G., Rogozinski, N., Gevaert, M., Kumbhar, P., Schürmann, F., and Segev, I. (2020). An efficient analytical reduction of detailed nonlinear neuron models. *Nat. Commun.* *11*, 288.

Bai, S., Zico Kolter, J., and Koltun, V. (2018). An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. *arXiv*, 1803.01271. <https://arxiv.org/abs/1803.01271>.

Bar-Ilan, L., Gidon, A., and Segev, I. (2013). The role of dendritic inhibition in shaping the plasticity of excitatory synapses. *Front. Neural Circuits* *6*, 118.

Behabadi, B.F., and Mel, B.W. (2014). Mechanisms underlying subunit independence in pyramidal neuron dendrites. *Proc. Natl. Acad. Sci. USA* *111*, 498–503.

Branco, T., Clark, B.A., and Häusser, M. (2010). Dendritic discrimination of temporal input sequences in cortical neurons. *Science* *329*, 1671–1675.

Burkitt, A.N. (2006). A Review of the Integrate-and-fire Neuron Model: I. Homogeneous Synaptic Input. *Biol. Cybern.* *95*, 1–19.

Cadena, S.A., Denfield, G.H., Walker, E.Y., Gatys, L.A., Tolia, A.S., Bethge, M., and Ecker, A.S. (2019). Deep convolutional models improve predictions of macaque V1 responses to natural images. *PLoS Comput. Biol.* *15*, e1006897.

Cazé, R.D., Humphries, M., and Gutkin, B. (2013). Passive dendrites enable single neurons to compute linearly non-separable functions. *PLoS Comput. Biol.* *9*, e1002867.

Doron, M., Chindemi, G., Muller, E., Markram, H., and Segev, I. (2017). Timed Synaptic Inhibition Shapes NMDA Spikes, Influencing Local Dendritic Processing and Global I/O Properties of Cortical Neurons. *Cell Rep.* *21*, 1550–1561.

Egger, R., Dercksen, V.J., Udvardy, D., Hege, H.-C., and Oberlaender, M. (2014). Generation of dense statistical connectomes from sparse morphological data. *Front. Neuroanat.* *8*, 129.

Gidon, A., Zolnick, T.A., Fidzinski, P., Bolduan, F., Papoutsi, A., Poirazi, P., Holtkamp, M., Vida, I., and Larkum, M.E. (2020). Dendritic action potentials and computation in human layer 2/3 cortical neurons. *Science* *367*, 83–87.

Gütig, R., and Sompolinsky, H. (2006). The tempotron: a neuron that learns spike timing-based decisions. *Nat. Neurosci.* *9*, 420–428.

Häusser, M., and Mel, B. (2003). Dendrites: bug or feature? *Curr. Opin. Neurobiol.* *13*, 372–383.

Hawkins, J., and Ahmad, S. (2016). Why Neurons Have Thousands of Synapses, a Theory of Sequence Memory in Neocortex. *Front. Neural Circuits* *10*, 23.

Hay, E., Hill, S., Schürmann, F., Markram, H., and Segev, I. (2011). Models of neocortical layer 5b pyramidal cells capturing a wide range of dendritic and perisomatic active properties. *PLoS Comput. Biol.* *7*, e1002107.

He, K., Zhang, X., Ren, S., and Sun, J. (2015). Deep Residual Learning for Image Recognition. *arXiv*, 1512.03385v1 <https://arxiv.org/abs/1512.03385>.

Holden, D., Duong, B.C., Datta, S., and Nowrouzezahrai, D. (2019). Subspace neural physics: Fast data-driven interactive simulation. In *Proceedings of the SCA 2019: ACM SIGGRAPH / Eurographics Symposium on Computer Animation (Association for Computing Machinery)*, pp. 1–12.

Jahr, C.E., and Stevens, C.F. (1993). Calcium permeability of the N-methyl-D-aspartate receptor channel in hippocampal neurons in culture. *Proc. Natl. Acad. Sci. USA* *90*, 11573–11577.

Jolivet, R., Schürmann, F., Berger, T.K., Naud, R., Gerstner, W., and Roth, A. (2008). The quantitative single-neuron modeling competition. *Biol. Cybern.* *99*, 417–426.

Kasim, M.F., Watson-Parris, D., Deaconu, L., Oliver, S., Hatfield, P., Froula, D.H., Gregori, G., Jarvis, M., Khatiwala, S., Korenaga, J., et al. (2020). Up to two billion times acceleration of scientific simulations with deep neural architecture search. *arXiv*, 2001.08055 <https://arxiv.org/abs/2001.08055>.

Kastellakis, G., Cai, D.J., Mednick, S.C., Silva, A.J., and Poirazi, P. (2015). Synaptic clustering within dendrites: an emerging theory of memory formation. *Prog. Neurobiol.* *126*, 19–35.

Katz, Y., Menon, V., Nicholson, D.A., Geinisman, Y., Kath, W.L., and Spruston, N. (2009). Synapse distribution suggests a two-stage model of dendritic integration in CA1 pyramidal neurons. *Neuron* *63*, 171–177.

Keat, J., Reinagel, P., Reid, R.C., and Meister, M. (2001). Predicting every spike: a model for the responses of visual neurons. *Neuron* *30*, 803–817.

Koch, C., and Segev, I. (2014). The role of single neurons in information processing. *Nat. Neurosci.* *3*, 1171–1177.

Koch, C., Poggio, T., and Torres, V. (1982). Retinal Ganglion Cells: A Functional Interpretation of Dendritic Morphology. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* *298*, 227–263.

Lapicque, L. (1907). Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarization. *J. Physiol. Pathol. Gen.* *9*, 620–635.

Larkum, M.E., Zhu, J.J., and Sakmann, B. (1999). A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* *398*, 338–341.

Larkum, M.E., Nevian, T., Sandler, M., Polsky, A., and Schiller, J. (2009). Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle. *Science* *325*, 756–760.

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* *521*, 436–444.

Lin, M., Chen, Q., and Yan, S. (2014). Network in network. In *2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings (International Conference on Learning Representations)*.

London, M., and Häusser, M. (2005). Dendritic computation. *Annu. Rev. Neurosci.* *28*, 503–532.

Losonczy, A., Makara, J.K., and Magee, J.C. (2008). Compartmentalized dendritic plasticity and input feature storage in neurons. *Nature* *452*, 436–441.

Magee, J.C., and Johnston, D. (1995). Characterization of single voltage-gated Na⁺ and Ca²⁺ channels in apical dendrites of rat CA1 pyramidal neurons. *J. Physiol.* *487*, 67–90.

Mahendran, A., and Vedaldi, A. (2015). Understanding deep image representations by inverting them. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (IEEE)*, pp. 5188–5196.

Major, G., Larkum, M.E., and Schiller, J. (2013). Active properties of neocortical pyramidal neuron dendrites. *Annu. Rev. Neurosci.* *36*, 1–24.

Markram, H., Muller, E., Ramaswamy, S., Reimann, M.W., Abdellah, M., Sanchez, C.A., Ailamaki, A., Alonso-Nanclares, L., Antille, N., Arsever, S., et al. (2015). Reconstruction and Simulation of Neocortical Microcircuitry. *Cell* *163*, 456–492.

McCulloch, W.S., and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *Bull. Math. Biophys.* *5*, 115–133.

Mel, B.W. (1992). NMDA-Based Pattern Discrimination in a Modeled Cortical Neuron. *Neural Comput.* *4*, 502–517.

- Moldwin, T., and Segev, I. (2018). Perceptron learning and classification in a modeled cortical pyramidal cell. *BioRxiv*, 464826.
- Mordvintsev, A., Olah, C., and Tyka, M. (2015). Inceptionism: Going deeper into neural networks (Google Res. Blog).
- Naud, R., Bathellier, B., and Gerstner, W. (2014). Spike-timing prediction in cortical neurons with active dendrites. *Front. Comput. Neurosci.* 8, 90.
- Olah, C., Mordvintsev, A., and Schubert, L. (2017). Feature Visualization. *Distill* 2, e7.
- Poirazi, P., Brannon, T., and Mel, B.W. (2003a). Pyramidal neuron as two-layer neural network. *Neuron* 37, 989–999.
- Poirazi, P., Brannon, T., and Mel, B.W. (2003b). Arithmetic of subthreshold synaptic summation in a model CA1 pyramidal cell. *Neuron* 37, 977–987.
- Polsky, A., Mel, B.W., and Schiller, J. (2004). Computational subunits in thin dendrites of pyramidal cells. *Nat. Neurosci.* 7, 621–627.
- Rall, W. (1959). Branching dendritic trees and motoneuron membrane resistivity. *Exp. Neurol.* 1, 491–527.
- Rall, W. (1964). Theoretical significance of dendritic trees for neuronal input-output relations. *Neural Theory Model.* 73–97.
- Rall, W. (1967). Distinguishing theoretical synaptic potentials computed for different soma-dendritic distributions of synaptic input. *J. Neurophysiol.* 30, 1138–1168.
- Rapp, M., Yarom, Y., and Segev, I. (1992). The Impact of Parallel Fiber Background Activity on the Cable Properties of Cerebellar Purkinje Cells. *Neural Comput.* 4, 518–533.
- Richards, B.A., Lillicrap, T.P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., Clopath, C., Costa, R.P., de Berker, A., Ganguli, S., et al. (2019). A deep learning framework for neuroscience. *Nat. Neurosci.* 22, 1761–1770.
- Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol. Rev.* 65, 386–408.
- Schiller, J., Schiller, Y., Stuart, G., and Sakmann, B. (1997). Calcium action potentials restricted to distal apical dendrites of rat neocortical pyramidal neurons. *J. Physiol.* 505, 605–616.
- Schiller, J., Major, G., Koester, H.J., and Schiller, Y. (2000). NMDA spikes in basal dendrites of cortical pyramidal neurons. *Nature* 404, 285–289.
- Segev, I., and Rall, W. (1988). Computational study of an excitable dendritic spine. *J. Neurophysiol.* 60, 499–523.
- Senior, A.W., Evans, R., Jumper, J., Kirkpatrick, J., Sifre, L., Green, T., Qin, C., Židek, A., Nelson, A.W.R., Bridgland, A., et al. (2020). Improved protein structure prediction using potentials from deep learning. *Nature* 577, 706–710.
- Shepherd, G.M., Brayton, R.K., Miller, J.P., Segev, I., Rinzel, J., and Rall, W. (1985). Signal enhancement in distal cortical dendrites by means of interactions between active dendritic spines. *Proc. Natl. Acad. Sci. USA* 82, 2192–2195.
- Spruston, N., Schiller, Y., Stuart, G., and Sakmann, B. (1995). Activity-dependent action potential invasion and calcium influx into hippocampal CA1 dendrites. *Science* 268, 297–300.
- Strogatz, S. (2001). *Nonlinear Dynamics And Chaos: With Applications To Physics, Biology, Chemistry, And Engineering* (CRC Press).
- Stuart, G., Spruston, N., Sakmann, B., and Häusser, M. (1997). Action potential initiation and backpropagation in neurons of the mammalian CNS. *Trends Neurosci.* 20, 125–131.
- Stuart, G.J., and Sakmann, B. (1994). Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature* 367, 69–72.
- Tzilivaki, A., Kastellakis, G., and Poirazi, P. (2019). Challenging the point neuron dogma: FS basket cells as 2-stage nonlinear integrators. *Nat. Commun.* 10, 3664.
- Ujfalussy, B.B., Makara, J.K., Lengyel, M., and Branco, T. (2018). Global and Multiplexed Dendritic Computations under In Vivo-like Conditions. *Neuron* 100, 579–592.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., and Polosukhin, I. (2017). Attention Is All You Need. *arXiv*, 1706.03762v5 <https://arxiv.org/abs/1706.03762>.
- Wybo, W.A.M., Jordan, J., Ellenberger, B., Marti Mengual, U., Nevian, T., and Senn, W. (2021). Data-driven reduction of dendritic morphologies with preserved dendro-somatic responses. *eLife* 10, 1–26.
- Zador, A.M., Claiborne, B.J., and Brown, T.H. (1991). Nonlinear Pattern Separation in Single Hippocampal Neurons with Active Dendritic Membrane. In *Proceedings of the 4th International Conference on Neural Information Processing Systems (NIPS)*, pp. 51–58.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
Simulation data, pretrained models, and key results	This paper	https://doi.org/10.34740/kaggle/ds/417817
Software and algorithms		
NEURON 7.6.2	NEURON	https://github.com/neuronsimulator/nrn
Python version 3.6	Python	https://www.python.org
Custom Code for simulation, fitting and analysis	This paper	https://github.com/SelfishGene/neuron_as_deep_net

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, David Beniaguev (david.beniaguev@gmail.com)

Materials availability

This study did not generate new unique reagents.

Data and code availability

All data and pre-trained networks that were used in this work are available on Kaggle datasets platform (<https://doi.org/10.34740/kaggle/ds/417817>) at the following link:

<https://www.kaggle.com/selfishgene/single-neurons-as-deep-nets-nmda-test-data>

Additionally, the dataset was deposited to Mendeley Data (<https://doi.org/10.17632/xjvsp3dhzf.2>) at the link:

<https://data.mendeley.com/datasets/xjvsp3dhzf/2>

A github repository of all simulation, fitting and evaluation code can be found in the following link:

https://github.com/SelfishGene/neuron_as_deep_net.

Additionally, we provide a python script that loads a pretrained artificial network and makes a prediction on the entire NMDA test set that replicates the main result of the paper (Figure 2):

<https://www.kaggle.com/selfishgene/single-neuron-as-deep-net-replicating-key-result>.

Also, a python script that loads the data and explores the dataset (Figure S1) can be found in the following link: <https://www.kaggle.com/selfishgene/exploring-a-single-cortical-neuron>.

METHOD DETAILS

I&F simulation

For Figure 1 simulations, membrane voltage was modeled using a leaky I&F simulation $V(t) = \sum_{j=1}^{N_{syn}} w_j \sum_{t_i} K(t - t_i)$, where w_j denotes synaptic efficacy for each synapse, t_i denotes presynaptic spike times, and $K(t - t_i)$ denotes the temporal kernel of each postsynaptic potential (PSP). We used a temporal kernel with exponential decay $K(t - t_i) = e^{-\frac{t-t_i}{\tau}} \cdot u(t - t_i)$ where $u(t)$ is the Heaviside function $u(t) = \begin{cases} 0, & |t| < 0 \\ 1, & |t| \geq 0 \end{cases}$ and $\tau = 20\text{ms}$ is the membrane time constant. When the threshold was reached, an output spike was recorded, and the voltage was reset to $V_{rest} = -77\text{mV}$. As input to the simulated I&F neuron, $N_{exc} = 80$ excitatory synapses and $N_{inh} = 20$ inhibitory synapses were used. Synaptic efficacies of $w_{exc} = 2\text{mV}$ were used for excitatory synapses and $w_{inh} = -2\text{mV}$ for inhibitory synapses. Each presynaptic spike train was taken from a Poisson process with a constant instantaneous firing rate. Values used $f_{exc} = 1.4\text{Hz}$ for excitatory synapses and $f_{inh} = 1.3\text{Hz}$ for inhibitory synapses. The resulting output average firing rate for these simulation values was 0.9 Hz.

L5PC simulations

For [Figure 2](#) and [Figure 3](#) simulations, we used a detailed compartmental biophysical model of cortical L5PC **as is**, modeled by [Hay et al., 2011](#). For a full description of the model please see [STAR Methods](#) in the original paper. Briefly, this model contains in total 12 ion channels for each dendritic compartment. Some of the channels are unevenly distributed over the dendritic arbor. In [Figure 3](#) double exponential conductances based AMPA synapses were used in simulations with $\tau_{rise} = 0.3\text{ms}$, $\tau_{decay} = 3\text{ms}$ and $g_{max} = 0.4\text{nS}$. For [Figure 2](#) and [Figure 4](#), in related simulations we used the standard NMDA model ([Jahr and Stevens, 1993](#)), with $\tau_{rise} = 2\text{ms}$, $\tau_{decay} = 70\text{ms}$, $\gamma = 0.08\text{mV}^{-1}$ and $g_{max} = 0.4\text{nS}$. For both [Figure 2](#) and [Figure 3](#), we also used double exponential GABA_A synapses with $\tau_{rise} = 2\text{ms}$, $\tau_{decay} = 8\text{ms}$ and $g_{max} = 1\text{nS}$ on each independent dendritic segment, we placed a single AMPA (for [Figure 3](#)) or AMPA + NMDA (for [Figure 2](#)) synapse as well as a single GABA_A synapse. In order to mimic uniform coverage of excitatory and inhibitory synapses on the entire dendritic tree, we stimulated each compartment with a firing rate proportional to the segment's length. Each presynaptic spike train was taken from a Poisson process with a smoothed piecewise constant instantaneous firing rate. The Gaussian smoothing sigma, as well as the time window of constant rate before smoothing were independently resampled for each 6 s simulation from the range 10ms to 1000 ms ([Figure S1D](#)). This was chosen, as opposed to a constant firing rate, to create additional temporal variety in the data in order to increase the applicability of the results to all possible situations. For [Figure 2](#) simulations with NMDA synapses, the total amount of excitatory and inhibitory presynaptic spikes per 100 ms range between 0 and 800 spikes ([Figure S1](#)). This is equivalent to 8000 excitatory synapses with an average rate of 1 Hz and 2000 inhibitory synapses with an average rate of 4 Hz. The average output firing rate of the simulated cell across all simulations was 1.0 Hz. For [Figure 3](#) simulations, with AMPA only synapses, the total amount of excitatory and inhibitory presynaptic spikes per 100 ms range were increased in order to account for the smaller amount of total current injected due to lack of NMDA current, with the purpose to achieve similar output firing rates of 1.0 Hz. See [Figure S4](#) for detailed comparison.

L23PC simulations

For [Figure 4](#) simulations, we used a detailed compartmental biophysical model of cortical L23PC **as is**, modeled by [Branco et al., 2010](#). In these experiments we stimulated a single branch with 9 dendritic segments with an NMDA synapse on each compartment, with parameters as in the simulation for [Figure 4](#). The branch was selected as in [Branco et al., 2010](#) in order to perform a comparison with the original paper. Similarly, to [Figure 2](#) and [Figure 3](#) simulations, each presynaptic spike train was taken from a Poisson process with a smoothed piecewise constant instantaneous firing rate. The number of presynaptic input spikes to the branch per 100 ms ranged between 0 and 15 in simulations used for training. In [Figures 5C](#), [5E](#), and [5F](#), we repeated input stimulation protocol suggested by Branco et. al, 2010, consisting of single presynaptic spike per synapse with constant time intervals of 5 ms between subsequent synaptic activations, only randomly permuting the order of activation between trials.

DNN fitting

In order to represent the input in a suitable manner for fitting with a DNN, we discretize time using 1 ms time bin Δt . Using this discretization, we can represent a spike train as a sequence of binary values $S[t]$, such that $S[t] \in \{0, 1\}$, since the length of a spike is approximately 1 ms there cannot be more than a single spike in such a time interval. We denote the spike trains the neuron receives as input as $X[s, t]$, $s \in \{1, 2, \dots, N_{syn}\}$, $t \in \{1, 2, \dots, T\}$, where s denotes the synapse index, and t denotes time. The spike trains a neuron emits as output we denote as $y_{spike}[t]$. The somatic voltage trace we denote as $y_{voltage}[t]$. For every point in time, we attempt to predict both somatic spiking $y_{spike}[t]$ and somatic voltage $y_{voltage}[t]$ based only a T_{input} sized window of presynaptic input spikes. i.e., define the vector $\vec{x}_t = [X[s, t]]$, $s \in \{1, 2, \dots, N_{syn}\}$, $t \in \{t, t-1, t-2, \dots, t-T_{input}\}$ and a neural network that maps \vec{x}_t to $\hat{y}_{spike}[t]$ and $\hat{y}_{voltage}[t]$. i.e., $\hat{y}_{voltage}[t], \hat{y}_{spike}[t] = DNN(\vec{x}_t, \theta)$. We treat spike prediction as a binary classification task and use standard log loss and treat voltage prediction as regression task and use standard MSE loss. We wish to find a model's parameters θ such that we minimize a combined loss $L(\theta) = L_{LogLoss}(y_{spike}, \hat{y}_{spike}) + w_{voltage} \cdot L_{MSE}(y_{voltage}, \hat{y}_{voltage})$, where $w_{voltage}$ is the relative importance of the spike prediction loss with respect to the somatic voltage prediction loss. For most of the experiments we set $w_{voltage}$ to be about half the size of the spike loss. The DNN architecture we used was a temporally convolutional network (TCN) ([Bai et al., 2018](#)) and we applied it in a fully convolutional manner on all possible time points. Note that when the temporal filter size after the first layer is 1 in a TCN applied as described, this is effectively a fully connected neural network. In most of our experiments we used fully connected neural networks, except for [Figure 2](#) in which we used a proper TCN with a hierarchical convolutional structure. After every convolutional layer, a batch normalization layer immediately follows. We employed a learning schedule regime in which we lowered the learning rate and increased batch size as we progressed through training. Full details of the learning schedule in each case are in the attached code repository. For the generation of [Figure S2](#) we trained many networks with different hyperparameters and trained each network for 2-14 days on a GPU cluster consisting of several V100, K80 and 2080Ti Nvidia GPUs. All results of the different hyperparameters and results can be found in the data link on the Kaggle platform. The total amount of single GPU years needed to fit all DNNs throughout the entire study was ~ 3.4 years.

Out-of-Distribution (OOD) simulations

To find spatial clusters in a data-dependent manner, we applied a K-means clustering algorithm on the cross-correlation matrix of all dendritic voltages. We use k to be 64 different spatial clusters. During the simulations used to create the data for [Figures 4](#) and [S8](#), we randomly selected up to 30% of all spatial clusters for each simulation and stimulated only these clusters during that simulation. In 50%

of the simulations, the temporal profile of instantaneous rates for all synapses was modulated by a sinusoidal function. The period of this temporal modulation was randomly selected to be between 15 ms and 300 ms for the duration of the simulation. Additionally, in 50% of the simulations we canceled all inhibition and kept only the excitation. The full OOD test set consisted of 2688 simulations, 6 s of simulation time each. The full details of this process can be found in the code repository on github and results of our spatial clustering can be found on the dataset on Kaggle. Together, these 3 manipulations (spatial clustering, temporal synchronization, and excitation only input) consist of large deviations from the statistics of the input during training and thus provide a robust OOD test.

QUANTIFICATION AND STATISTICAL ANALYSIS

Model evaluation

We divided our simulations to train, validation, and test datasets. We fitted all DNN models on the training dataset, and all reported results are on an unseen test dataset. A validation dataset was used for modeling decisions, hyperparameter tuning and snapshot selection during the training process (early stopping). We evaluated binary spike prediction results using the receiver operator characteristic (ROC) curve and calculated the area under the curve (AUC). We note that due to the relatively low firing rate of the neuron, the binary classification problem of the instantaneous spike prediction problem is highly unbalanced. For every second of simulation there was on average 1 positive sample (spike) for every 999 negative samples (non-spikes). Therefore, we used a very conservative threshold over the binary spike probability prediction output of the DNN in order to create the final spike train prediction and examine the cross-correlation plot. Note also that a prediction without a single True Positive on the 1 ms time binning binary spike prediction problem can still be in fact, a very good solution, e.g., if our model outputs, as its prediction, the exact same spike train as the original but offset by 1 ms in time. In this case, there will be no True positives and many False positives, but the predicted spike train is quite good nonetheless. To summarize: the temporal cross-correlation between the original and predicted spike trains is not directly related to binary prediction metrics used and therefore we display it as it's not redundant. For creating a binary prediction, we chose a threshold over the continuous model prediction that corresponds to 0.2% false-positive rate on the validation set after training (a different threshold was used for each model trained). In order to evaluate the temporal precision of the binary spike prediction we plotted the normalized cross-correlation between the predicted output spike train and the ground truth simulated spike train for a 50 ms temporal offset in either direction. To quantify the width of cross correlation plot, we fit a Gaussian to it and use the sigma parameter (see [Table S1](#)). In order to evaluate the voltage prediction, we calculated the RMSE and plot the scatterplot between predicted voltage and the ground truth simulated voltage.

Comparison of AMPA versus NMDA model complexity

In order to systematically compare the complexity of DNNs for AMPA case ([Figure 3](#)) versus NMDA case ([Figure 2](#)), we select a minimal approximation threshold (AUC = 0.9910) for spiking accuracy of the DNN as compared to the spiking of the respective biophysical model. We consider DNNs that perform better than this threshold to be a good approximation. We then ask what is the minimal sized network that satisfies this threshold for AMPA versus NMDA cases. For this, we trained a total of 137 DNNs (62 for AMPA, 75 for NMDA) while varying 3 main hyperparameters - depth, width (number of channels per layer), and temporal extent of the input history required to make a prediction. For an extended comparison of these data please see [Figures 4](#) and [S2](#). For full fitting results with many additional various metrics for all 137 DNN fits, see released accompanying data. Although the quantitative results are dependent on the precise threshold used (AUC = 0.9910), they are mostly insensitive to small changes in this threshold.

Assessing the complexity resulting from the interaction between dendritic morphology and NMDA receptor density

To investigate further the relationship between NMDAR and its interaction with morphology as well as the influence of NMDA receptor density on I/O transformation complexity, we conducted a set of additional analyses ([Figures S5](#) and [S6](#)). When assessing the I/O complexity of NMDA synapses when placed only on part of the morphology, we perform simulations similar to those described previously, with the only difference being the part of the dendritic tree receiving synaptic input. In order to conduct a proper comparison, we tune the input firing rates such that the average output firing rate will be identical for all simulations (both test and train datasets). In this case we train DNNs of {1,2,3,7} layers for a fixed width (128) and fixed history dependence (100 ms). We limit training time to ~24 hours since our goal is to perform a quantitative comparison and not achieve maximally performance models like for the case shown in [Figures 2](#) and [3](#). To establish a "morphological complexity" axis we select different regions of the dendritic tree of the modeled L5PC that are strictly contained within each other. These regions are: full morphology, full morphology excluding tuft synapses, the basal tree only, and only proximal compartments of the basal tree. Results depicted in [Figures S5A–S5C](#), and summary results in [Figure S5G](#). summary results depict the spiking AUC accuracy measure for the case of DNNs with 3 layers for all 4 morphological complexities. Higher I/O complexity is interpreted as corresponding to lower fit accuracy. Similarly, In [Figures S5D–S5F](#) and [S5H](#) we examine the interaction between the existence of NMDARs on proximal and distal parts of the basal tree as compared to AMPA only case. The only difference from what was described above for the morphological case is that now we fitted DNNs with a smaller fixed width (32 filters/layer). In a similar way, In [Figure S6](#) we examine the dependence of the I/O relationship on varying the NMDAR max conductance level (which models receptor density). Here again we use 32 wide DNNs for this analysis.