

Dendritic morphology predicts pattern recognition performance in multi-compartmental model neurons with and without active conductances

Giseli de Sousa · Reinoud Maex · Rod Adams · Neil Davey · Volker Steuber

Received: 16 March 2014 / Revised: 6 September 2014 / Accepted: 20 October 2014 / Published online: 8 November 2014
© Springer Science+Business Media New York 2014

Abstract In this paper we examine how a neuron's dendritic morphology can affect its pattern recognition performance. We use two different algorithms to systematically explore the space of dendritic morphologies: an algorithm that generates all possible dendritic trees with 22 terminal points, and one that creates representative samples of trees with 128 terminal points. Based on these trees, we construct multi-compartmental models. To assess the performance of the resulting neuronal models, we quantify their ability to discriminate learnt and novel input patterns. We find that the dendritic morphology does have a considerable effect on pattern recognition performance and that the neuronal performance is inversely correlated with the mean depth of the dendritic tree. The results also reveal that the asymmetry index of the dendritic tree does not correlate with the performance for the full range of tree morphologies. The performance of neurons with dendritic tapering is best predicted by the mean and variance of the electrotonic distance of their synapses to the soma. All relationships found for

passive neuron models also hold, even in more accentuated form, for neurons with active membranes.

Keywords Dendritic tree · Associative memory · Synaptic plasticity · Asymmetry index · Electrotonic length

1 Introduction

This paper studies functional aspects of dendritic morphologies. The dendritic trees that are present in the brain exhibit a great variety of morphologies, and different types of neurons (such as pyramidal cells, Purkinje cells, Golgi cells) are characterised by their specific dendritic structures. It is unlikely that this is accidental, and several hypotheses have been posited to explain the existence of these variable dendritic dimensions and branching structures. For example, it has been suggested that the dendritic morphology of a neuron is optimised so that the cost of propagating signals from the synapses to the soma is minimal (Cuntz et al. 2007; Wen and Chklovskii 2008). It is also thought that the dendritic topology, how the dendritic segments are connected, could relate to the firing pattern of the neuron (Mainen and Sejnowski 1996; Fohlmeister and Miller 1997; Krichmar et al. 2002; van Ooyen et al. 2002; van Elburg and van Ooyen 2010).

Dendrites have an important role in the information processing that takes place in a neuron. They are involved in the generation, propagation and integration of synaptic potentials, the back propagation of action potentials and the induction of synaptic plasticity (Gulledge et al. 2005; London and Häusser 2005; Cuntz et al. 2007; Wen and Chklovskii 2008). The latter has been implicated in learning and therefore in the functioning of associative memory [for example Chen et al. (2011) and Steuber et al. (2007)].

Action Editor: Arnd Roth

G. de Sousa (✉)
Connectionism and Cognitive Science Laboratory,
Department of Informatics and Statistics,
Federal University of Santa Catarina,
Campus Universitário, Trindade, 88040-970,
Florianópolis, SC, Brazil
e-mail: giseli@inf.ufsc.br

G. de Sousa · R. Maex · R. Adams · N. Davey · V. Steuber
Biocomputation Group, Science and Technology
Research Institute, University of Hertfordshire,
Hatfield, Herts, UK

R. Maex
Group for Neural Theory, Department of Cognitive
Sciences, École Normale Supérieure, Paris, France

In the present study we take a model neuron, which may either be passive or contain active ion channels, and train it to perform a pattern recognition task. The synaptic strengths are set so that the neuron responds differently to patterns it has learnt, from purely random, novel patterns. Both the learnt and novel patterns are sparse, binary patterns that do not change over time. In order to see the effect of morphological variation in the dendrites we generated a variety of dendritic trees. In our first experiment, using a small neuron with only 22 terminal points and 43 synapses, we were able to generate every possible dendritic structure with binary bifurcations, and measure the performance of all of these model neurons. In the remaining experiments we used a bigger neuron with 128 terminal points and 255 synapses, in which the size of the morphological space entailed that we evaluated only a sample of all the possible morphologies. There are a variety of metrics that can be used to characterise a tree structure, such as its symmetry or mean depth. Using these measures we were able to examine how well these metrics were able to predict the performance of a neuron from its morphology.

2 Methods

We first present the neuron models and their biophysical parameters (Section 2.1), followed by four metrics used to quantify the morphological features of the dendrite (Section 2.2). Next, we describe the algorithms used for generating, in a systematic way, sample neurons that differed only in their dendritic topology (Section 2.3). Section 2.4 deals with the pattern recognition task (the patterns, their presentation and the metric used to assess neuronal performance), and the final Section 2.5 gives some implementation details. The simulation source code and neuronal model are freely available at <https://code.google.com/p/evol-patrec>.

2.1 The neuron model

The neuronal model used in this work is based on the model and dendritic morphologies described by van Ooyen et al. (2002). In their work, the authors built simple dendritic morphologies with the same electrophysiological properties but varying topological arrangements, shown to produce different firing patterns. This model was not based on an actual neuronal morphology, but rather used to represent members of an abstract morphology space. All morphologies were binary trees with a simplified structure, including all dendritic segments having the same length. Because of the simplicity of the morphologies produced by this model, they were chosen to be the basis of our search for optimal morphologies for pattern recognition.

2.1.1 Passive and active membrane properties

The experiments performed in our present work used two types of neuron models: passive models that did not contain any active conductances in the soma and dendrites and that were therefore unable to generate action potentials, and active models with voltage-gated ion channels in both soma and dendrites. Based on van Ooyen et al. (2002), the following values were used for the passive parameters membrane capacitance, membrane resistance and axial resistivity, respectively: $C_m = 0.75 \mu\text{F}/\text{cm}^2$, $R_m = 30 \text{ k}\Omega\text{cm}^2$ and $R_a = 150 \Omega\text{cm}$. For the active models, Hodgkin-Huxley-type kinetic descriptions of ion channels were also taken from van Ooyen et al. (2002), based on Mainen and Sejnowski's two-compartmental model (Mainen and Sejnowski 1996). Their conductance densities and reversal potentials are presented in Table 1. Note that a different E_{leak} value was used for the passive model (-65 mV), in agreement with a study on pattern recognition in a model hippocampal pyramidal cell by Graham (2001).

Table 1 Ion channel conductances from van Ooyen's model. These conductances are expressed in $\text{pS}/\mu\text{m}^2$, reversal potentials in mV

Conductance	Symbol	Soma	Dendrite	Reversal potential
Fast Na^+	\bar{g}_{Na}	3000	15	60
Fast non-inactivating K^+	\bar{g}_{Kv}	150	–	–90
Slow voltage-dependent non-inactivating K^+	\bar{g}_{Km}	–	0.1	–90
Slow Ca^{2+} -activated K^+	\bar{g}_{KCa}	–	3	–90
High voltage-activated Ca^{2+}	\bar{g}_{Ca}	–	0.3	140
Leak	\bar{g}_{Leak}	–	0.33	–70

2.1.2 Compartmentalisation and synapses

Each edge of the binary tree representing the dendritic morphology was implemented as an isopotential compartment receiving one synapse. As a result (see below), the neurons simulated in the exhaustive search, having $m = 22$ terminal branches, counted 43 ($= 2m - 1$) dendritic compartments. Those sampled from the space of trees with 128 terminal branches had 255 compartments. The lengths and diameters of the soma and dendritic compartments were also based on the values used in van Ooyen et al. (2002). The soma was a cylinder of 20 μm length and diameter. Each dendritic compartment had a diameter of 2.5 μm . The passive models were simulated with dendritic compartmental lengths of either 10 μm (all simulations apart from Fig. 8) or 5 μm (Fig. 8, which contains a direct comparison between active and passive models in the same panels). Only 5 μm long compartments were used in the active models, which required some initial tuning to make them appropriate for the pattern recognition task (see also Section 2.4). To analyse the effect of tapering on neuronal performance, we also introduced a new parameter called the *tapering factor*. The tapering factor is the ratio between the diameter of the child branch and its parent ($\textit{tapering factor} = \frac{\textit{diam}_{\textit{child}}}{\textit{diam}_{\textit{parent}}}$). Hence, $\textit{tapering factor} = 1$ means no tapering, and $\textit{tapering factor} = 0.8$ means that the diameter of each child branch will be 20% smaller than that of its parent branch. To prevent the generation of unrealistically thin dendrites through tapering, a minimum dendritic diameter of 0.1 μm was reinforced.

The models were excited through activation of synapses of the AMPA-receptor type, one on each compartment. These synapses were modelled as time-varying conductances with a dual-exponential time-course and implemented as *Exp2Syn* objects in the NEURON simulator (Carnevale and Hines 2006), with parameter values of 0.2 and 2 ms for the rise and decay time-constant, respectively, and a reversal potential of 0 mV. The peak conductance amplitude of a naive synapse (before learning) was set to 1 nS for passive models and 1.5 nS for active models. After learning, these conductances were scaled by multiplying them with the resulting synaptic weights (see Section 2.4).

2.2 Morphological tree metrics

To distinguish between different dendritic tree topologies, we used four morphological metrics: asymmetry index, mean depth, and mean and variance of electrotonic path length. The asymmetry index is simply the mean of the partition asymmetries of all vertices (bifurcation points) in the tree (van Pelt et al. 1992). So, the index of asymmetry A_t for a given tree α^n , with n terminal segments and $n - 1$

bifurcation points, is defined as:

$$A_t(\alpha^n) = \frac{1}{n - 1} \sum_{j=1}^{n-1} A_p(r_j, s_j) \tag{1}$$

The partition asymmetry A_p at a given vertex j is defined as :

$$A_p(r_j, s_j) = \frac{|r_j - s_j|}{r_j + s_j - 2} \tag{2}$$

where r_j and s_j are the number of terminal segments in the two subtrees of the vertex j and $A_p(1, 1)$ is equal to zero. Given this equation, we find that the asymmetry index is zero for the most symmetric tree and close to one for the most asymmetric one.

The second metric used, mean depth, is calculated as the mean number of steps between the soma and the dendritic synapses. Thus, for a given tree α^n with n terminal segments, the mean depth P_t is defined as:

$$P_t(\alpha^n) = \frac{1}{2n - 1} \sum_{i=1}^{2n-1} P_i \tag{3}$$

where P_i is the total number of edges on the path from the i th segment to the soma. Notice that the mean depth is calculated over all dendritic segments instead of just the terminal ones as previously done in related work (van Ooyen et al. 2002). This is required as we need to consider the location of all synapses, which are uniformly distributed over all dendritic segments.

The third metric, mean electrotonic path length, also used by van Elburg and van Ooyen (2010), is also calculated using the path from each dendritic segment to soma. To calculate the electrotonic path length, each dendritic segment i has its length ℓ_i normalised by an electrotonic length constant λ_i , which is defined as:

$$\lambda_i = \sqrt{\frac{d_i R_m}{4R_a}} \tag{4}$$

where d_i is the diameter of the dendritic segment i . So, the normalised electrotonic length Λ_i is given as:

$$\Lambda_i = \frac{\ell_i}{\lambda_i} \tag{5}$$

To calculate the mean electrotonic path length (MEP) for the dendritic tree with n terminal segments, the following equation is used:

$$MEP(\alpha^n) = \frac{1}{2n - 1} \sum_{i=1}^{2n-1} \Pi_i \tag{6}$$

where Π_i is the sum of the electrotonic lengths Λ_j of all the dendritic segments on the path from dendritic segment i to soma. It is important to notice that when all compartments have the same length and diameter (no tapering), the mean electrotonic path length is proportional to the mean depth

metric. For this reason, mean electrotonic path length is only reported in the final section where tapering is examined.

The last metric, variance of electrotonic path length, calculates the variance of Π_i across all synapses. This metric was introduced as it may better correlate with the signal-to-noise metric used to quantify neuronal performance, which also involves the calculation of variances (as later explained in Section 2.4).

2.3 Systematic tree generation

2.3.1 Representation of dendritic trees

To represent and generate dendritic trees, the partition notation from van Pelt and Verwer (1985) was used. A partition at a bifurcation point in a binary tree is defined by a pair of numbers which denote the degree of each subtree. Each partition represents a bifurcation point, where the nodes of its subtree are split into those on its left and those on its right branch. The topology of the whole tree can therefore be characterised by the set of partitions at its bifurcation points. So, for example, the most asymmetric tree with 5 terminal points can be described by the partitions 5(1 4(1 3(1 2(1 1))))).

In general, a binary tree T_n with n terminal points can be described using the following rule:

$$T_n = n(T_a T_b) \quad (7)$$

where $a + b = n$; $a, b > 0$ and $T_1 = 1$.

2.3.2 Trees exhaustively generated

To compare a large set of neuronal morphologies, the initial idea was to cover the whole search space by generating all possible binary trees for a given number of terminal points. To do this, we implemented an algorithm of which the pseudocode is presented in Algorithm 1 in the Appendix. Note that the tree-space scales exponentially with the number of terminal branches (Harding 1971). For practical reasons, we chose a tree order of 22 terminal points, which, using the code above, generated a total number of 1,514,661 trees

(a tree of 24 terminal branches would have 8,197,377 different morphologies!). Samples of the trees generated are presented in Fig. 1.

2.3.3 Trees selectively generated

As it was not possible to simulate the whole range of neuronal morphologies for the desired dendritic tree order (128 terminal points), a second method was used, comparing randomly generated morphologies. To achieve this, we implemented an algorithm to produce samples of dendritic trees with a given number of terminal points (see pseudocode given in Algorithm 2 in the Appendix). This algorithm differs from the one to generate trees exhaustively mainly at the splitting function. Instead of generating the whole possible range of partitions for each pair of branches a and b , the algorithm depends on a bias value, which controls the partition. In summary, low bias is more likely to generate trees with extreme values, which means more symmetric or asymmetric trees. On the other hand, if the bias is equal to 0.5 the algorithm generates completely random trees.

A sample of trees generated with 128 terminal points using this algorithm is presented in Fig. 2, where the trees are ordered by their degree of symmetry.

2.4 The pattern recognition task

The neuronal model was trained to discriminate between stored and novel spatial input patterns. A pattern was a random vector of binary numbers with one number for each compartment, a positive bit meaning that the associated synapse was to be activated. The patterns were sparse (only about 10% of the synapses were activated per pattern). The selectively generated neurons with 128 terminal branches (255 dendritic compartments) had 255-bit input patterns, with 25 positive bits. In the exhaustive search of neurons with 22 terminal branches, 43-bit patterns were used with 4 positive bits. Throughout all neurons and trials, each presented pattern was newly generated. Potential effects of randomness were avoided by averaging, for selected neuron samples, over 100 trials (Figs. 7a, 8, 10 and 11).

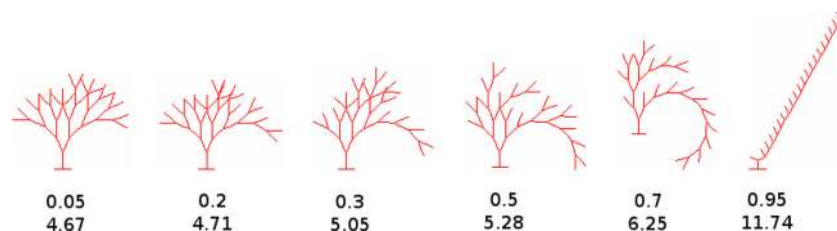


Fig. 1 Samples of tree morphologies with 22 terminal points generated by an exhaustive tree-generation algorithm. The values shown indicate asymmetry index (*top*) and mean depth (*bottom*). For these neurons, all compartments were $10 \mu\text{m}$ long

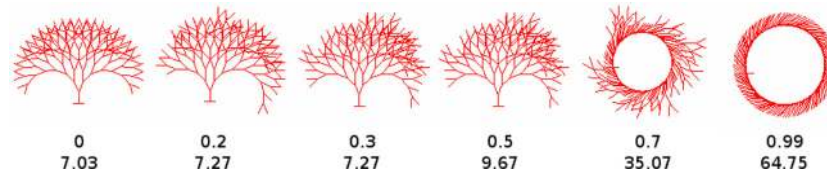


Fig. 2 Six sample trees with 128 terminal points generated by the selective tree-generation algorithm. The values indicate the tree asymmetry index (*top*) and the mean depth (*bottom*). The trees are visualised using the NEURON simulator tool (Hines and Carnevale 1997), which displays the long dendritic shafts of the most asymmetric trees

To present the patterns to the neuronal model, each bit of the input pattern was mapped to a specific synapse. To do this, each synapse was numbered by the location of its dendritic compartment in the tree. The compartments were indexed from the left side of the tree to the right. An example is given in Fig. 3 where the same pattern was mapped to the dendritic trees from both the most symmetric and the most asymmetric morphologies.

The learning rule was simple one-shot Hebbian learning: when the N patterns to be learnt were x^μ (μ is $1..N$), the (dimensionless) weight at synapse i was given by $w_i = \sum_{\mu} x_i^\mu$. In the recall phase, the performance was measured by comparing the neuronal responses after the presentation of learnt and novel input patterns. In the passive models, the comparison was done by using the somatic EPSP amplitudes as shown in Fig. 4; in the active models, the number of evoked action potentials in a 100 ms time window was used as shown in Fig. 5. It should

be noted that the accuracy of this performance evaluation is determined by the number of trials, the greater the number of trials the more accurate is the performance measurement.

The discrimination between stored and novel patterns was evaluated by calculating a signal-to-noise ratio (s/n), which is given as (Dayan and Willshaw 1991):

$$s/n = \frac{(\mu_s - \mu_n)^2}{0.5(\sigma_s^2 + \sigma_n^2)} \tag{8}$$

where μ_s and μ_n represent the mean values and σ_s^2 and σ_n^2 the variances of the responses to stored and novel patterns, respectively. From the histograms presented in Figs. 4 and 5, it is possible to find a clear discrimination between stored (blue bars) and novel patterns (red bars), which result in high s/n ratios in both figures (23.76 and 16.20 respectively). Note that all signal-to-noise ratios mentioned in Section 3 are the averages of at least five complete trials of this learning and testing procedure.

Fig. 3 Mapping input pattern to trees. The diagram shows how the same input pattern is mapped to each synapse in the most symmetric and the most asymmetric morphologies

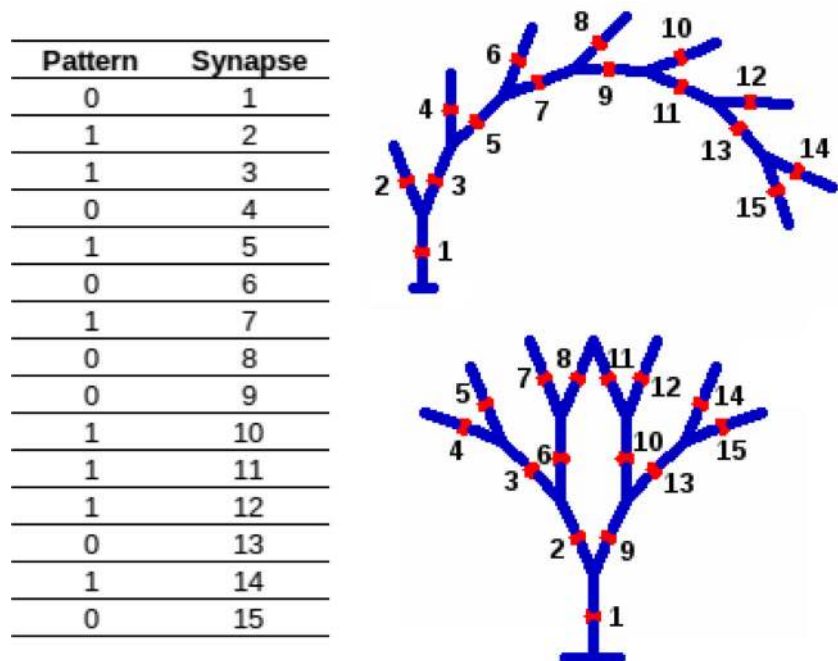
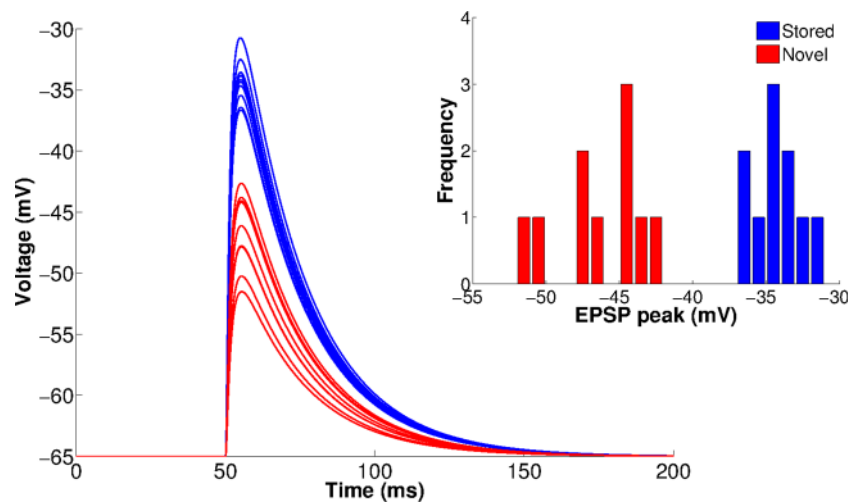


Fig. 4 Pattern recognition in the passive neuron model. The voltage traces (*left*) show the EPSP responses at the soma to 10 stored patterns (*blue traces*) and 10 novel patterns (*red traces*). The histogram shows the frequency of the EPSP peak responses for both stored and novel patterns (bin-width 1 mV). The resulting signal-to-noise ratio is 23.76



Given that the measured responses were different in passive and active neurons (EPSP amplitude versus number of spikes), a slightly different strategy of pattern presentation was used in the active models in order to obtain signal-to-noise ratios of similar magnitude. More particularly, whereas in passive neurons all positive pattern bits activated their associated synapses simultaneously and only once, in active neurons these synapses were each activated by a train of 5 spikes with 3 ms interspike intervals. Moreover, as the outcome of the active neurons was discrete (their number of spikes), their range of responses was limited and it was not uncommon for some neuronal topologies to generate identical spike numbers to all patterns, rendering the variances zero and hence the s/n ratio non-existent. To avoid this problem noise was added to the synapses of the active models. This noise comprised both a jitter on the timing of

the spikes in the afferent train, and a random background 1 Hz activation of each synapse in the tree with a strength of 0.5 nS. As a control, we applied in Fig. 8 the same afferent trains and background noise to both active and passive models.

2.5 Implementation details

All trees were generated by LISP programs, stored using their partition representation, and read into the NEURON simulator (Hines and Carnevale 1997) by custom-made routines implemented in C++. Simulating a neuronal morphology with 128 terminal points took approximately 7 seconds for passive models and 24 seconds for active models, on an 2 Quad-Core Intel Xeon 2.8-GHz processor with 8Gb physical memory under MacOS X 10.6. The most intense

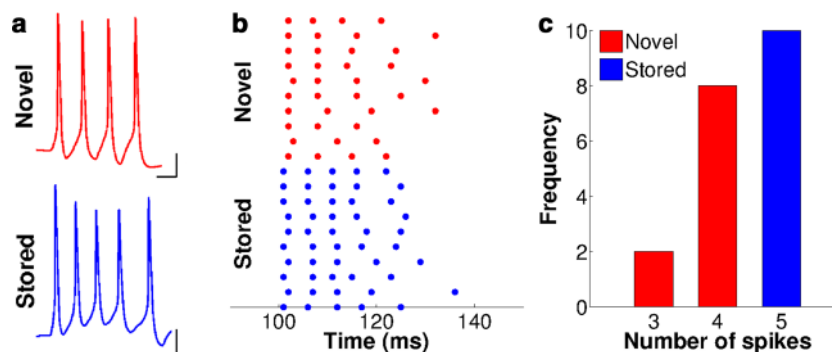


Fig. 5 Pattern recognition in the active model. The neuronal response in active neuronal models was determined by counting the number of spikes after pattern presentation. Two examples of neuronal response are shown in (a), for novel and stored patterns. The raster plot in (b)

represents the responses to 10 stored patterns (*blue dots*) and 10 novel patterns (*red dots*). The histogram in (c) shows the frequency of the number of spikes produced for stored and novel patterns. The resulting signal-to-noise ratio is 16.20. Scale bars: 5 ms, 20 mV

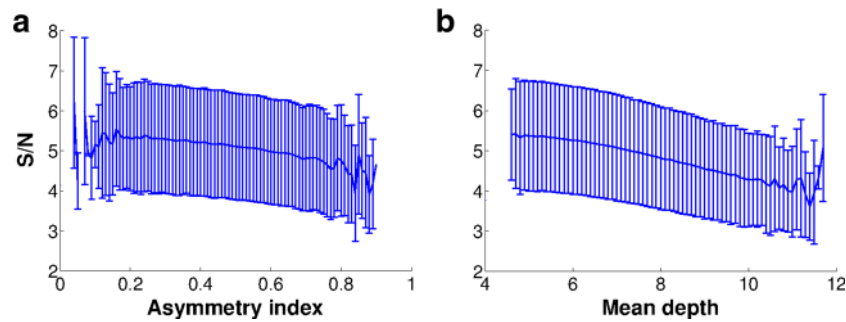


Fig. 6 The performance of every binary tree with 22 terminal points is plotted against asymmetry index (a) and mean depth (b). For each metric, a different bin width is used to result in a similar resolution of data (0.01 for asymmetry index and 0.1 for mean depth respectively).

The pattern recognition performance was calculated by averaging over the s/n ratio in response to five different sets of patterns presented to each neuronal morphology. *Error bars* represent the standard deviation calculated across all trees within the bin

simulations ran 155,000 passive morphologies for about 10 days distributed over five dual Quad-core computers. The data were analysed using MATLAB (MathWorks).

3 Results

In our exploration of the relationship between pattern recognition performance and dendritic morphology, and our search for the best metric to describe this relationship, we first compared all possible trees with 22 terminal points using a passive model neuron. Next we compared, using both passive and active neuronal models, an extensive sample of trees with 128 terminal points, for which the space of all morphologies is too large to be explored exhaustively. Finally, we studied the effects of tapering of the diameter of the dendritic compartments.

3.1 Comparing exhaustively generated trees

Each of all 1,514,661 possible trees with 22 terminal points was implemented as a passive model neuron, and its performance assessed as the mean signal-to-noise ratio of five trials of a 20-pattern recognition task (10 stored versus 10 novel patterns). To present this large amount of data, we partitioned the tree space according to the two morphometrics studied here, asymmetry index and mean depth. Figure 6 shows the mean and standard deviation of the signal-to-noise ratio for these partitioned data sets, using bin-widths of 0.01 for the asymmetry index in (a), and 0.1 for the mean depth in (b).

These results demonstrate that for trees with 22 terminal points, the trees with lower values of the two metrics, which represent the more symmetric morphologies, are the ones

with better pattern recognition performance. The trends presented in Fig. 6 show a decrease of performance when the morphologies become more asymmetric. The fluctuations found in the last bins of each metric as well as the initial bins for the asymmetry index metric result from the low number of trees that are contained in these bins.

3.2 Comparing selectively generated trees

Exhaustively generating a complete set of neurons with larger dendritic trees is not feasible. We therefore wrote a LISP program that drew 155,000 sample trees from the space of trees with 128 terminal branches. This program had a parameter (the *bias*) that could be set to ensure that trees with extreme morphologies (very symmetric and asymmetric trees) were sampled as well. The scatter plot of Fig. 7a shows, for each sampled tree implemented as a passive neuron, the signal-to-noise ratio averaged over 5 trials of 20 patterns. Notice that the LISP program sampled the entire range of asymmetry indices, though not uniformly, which generated the set of data found in the blue scatter plot of Fig. 7a. Initially, the result seemed unpromising. However, when a more detailed examination was undertaken a clearer pattern emerged. As for each neuron 20 new patterns (10 acting as stored patterns and 10 acting as novel patterns) were randomly generated at each trial, and as each tree was assessed over only 5 trials, the difficulty of the pattern recognition task was expected to depend on the particular set of 5 times 20 patterns that was used for each neuron (for example, it should be easier for neurons to distinguish orthogonal or non-overlapping patterns). Thus, as the accuracy of the performance measure increased with the number of trials, we randomly selected one neuron from each bin and averaged its performance over 100 trials of the same pattern recognition task. The results are shown as red data points in Fig. 7a, which allowed to verify the overall trend

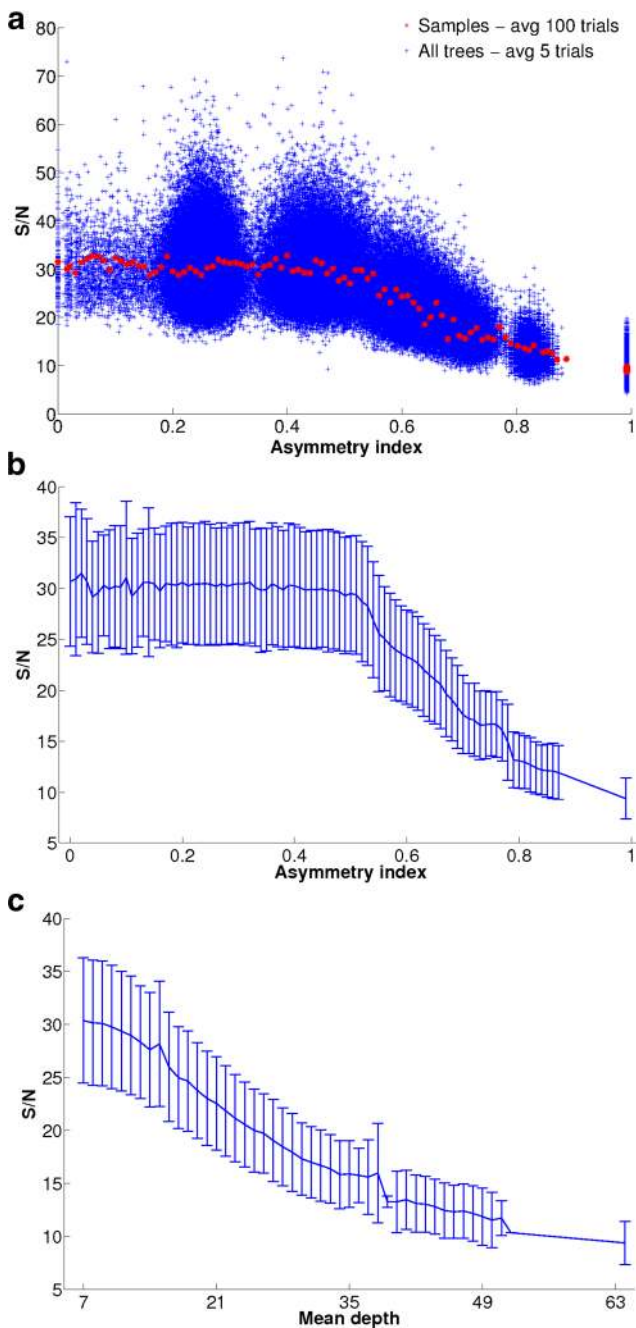


Fig. 7 Pattern recognition performance of passive neurons having dendrites selectively generated from the space of trees with 128 terminal points. The scatter plot in (a) shows, for each of 155,000 trees, the signal-to-noise ratios assessed over five trials of a 20-pattern recognition task (blue data points). For the construction of the red curve, one tree was randomly selected from each bin, and its performance recalculated over 100 trials of the pattern recognition task. (b) plots the mean (calculated from data points in (a)) and standard deviation of the signal-to-noise ratio over all trees within the same bin (bin-width 0.01) against the asymmetry index. (c) shows the signal-to-noise ratio of all trees within the same bin (bin-width 1) against the mean depth of these trees

visible in the data. In the next step, we partitioned the tree space into bins of 0.01 width along the asymmetry index axis as above. The blue error bars in (b) plot the mean and standard deviation in each bin, and the smooth character of the curve arises from the large number of samples averaged for each bin. For (c) a similar plot was produced where the mean and standard deviation of the signal-to-noise ratio for each bin is plotted against the mean depth of each tree. This metric shows a better correlation with the pattern recognition performance, which can be explained by the way this metric is calculated based on the distance of each synapse from the soma.

To study the effect of active conductances on the relationship between dendritic morphology and pattern recognition performance, and to compare this relationship for active and passive model neurons with selectively generated trees, a new experiment was designed where both models used the same set of trees, as given in the previous experiment, and the same set of parameters as originally determined for active models (compartments of $5\mu\text{m}$ length with synapses of 1.5 nS peak conductance, 5-spike input trains and background synaptic noise). The results are plotted in Fig. 8 against two metrics, asymmetry index (a) and mean depth (b). Each data point and error bar indicate the average and standard deviation over five randomly selected neurons in each bin. The results show that the negative correlation between pattern recognition performance and mean depth persists across the whole depth range for both active and passive models (b). In contrast, as shown in (a), the asymmetry index does not correlate with performance over its entire range for either the active or the passive models. This results from the fact that all of the trees with asymmetry indices between 0 and 0.4 correspond to a range of trees with very similar low mean depth (close to 7), as shown in Fig. 9. Interestingly, all of these trees with varying asymmetry index and therefore varying morphology but similar mean depth show an almost identical pattern recognition performance. This lack of effect of dendritic morphology on pattern recognition performance for very symmetric trees explains why mean depth but not asymmetry index correlates well with pattern recognition performance as shown in Fig. 8. Thus, the presence of active conductances does not affect the shape of the relationships between the two measures of tree morphology and the pattern recognition performance.

3.3 Robustness of the results

In order to investigate the robustness of our results, we varied the different parameters of our experiment described in Section 3.2. In particular we varied the amount of background noise, the loading (the number of stored patterns) and the sparsity (the number of active synapses) of the

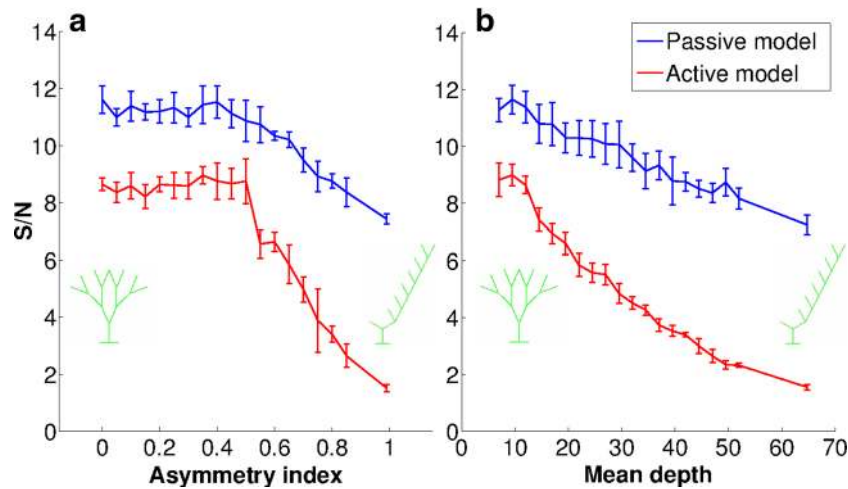


Fig. 8 Pattern recognition performance of active (red) and passive (blue) neuronal models with selectively generated dendritic morphologies drawn from the space of trees with 128 terminal points. Results were obtained by generating a population of 155,000 trees spanning the full range of each metric, from which five trees were randomly selected in each bin, using bin-widths of 0.05 for the asymmetry index

(a) and 2.5 for the mean depth (b). Each data point and error bar plots the average and standard deviation over the five selected neurons, each tested over 100 trials in a 20-pattern recognition task. Note that in these simulations, as explained in Section 2.4, the passive neurons received the same afferent spike trains and background noise as used for the active neurons

patterns. We also checked the effect of adding NMDA receptors. These results are shown in Fig. 10. Panel (a) shows that the addition of background noise led to a small decrease in pattern recognition performance, without affecting the shape of the relationship between the mean depth of the dendritic trees and the signal-to-noise ratio. As expected, similar results were obtained when the loading and the sparsity were varied. Panel (b) shows that the signal-to-noise ratio decreased as the loading increased, but performance

was still inversely correlated with mean depth. Panel (c) shows that as anticipated performance decreased as sparsity increased, but the anti-correlation between mean depth and performance was maintained at all three sparsities tested. Finally panel (d) shows that changing the ratio of NMDA and AMPA receptors affected the performance of the model and interestingly, when the conductance ratio was 0.5 the neuron became less sensitive to variations of its morphology. The slow time-course of the NMDA receptor conductances made the responses to the input patterns less sensitive to low-pass filtering by the dendrite, and hence pattern recognition less sensitive to the precise location of the synapses and the morphology of the tree. For NMDA/AMPA receptor conductance ratios of 0.5 or more, this improved the pattern recognition performance of neurons with asymmetric dendritic trees (in mean \pm standard deviation, NMDA/AMPA ratio 0.5: $s/n = 32.59 \pm 11.27$ for fully symmetric trees, $s/n = 24.59 \pm 9.01$ for fully asymmetric trees; NMDA/AMPA ratio 1: $s/n = 29.31 \pm 12.62$ for fully symmetric trees, $s/n = 23.18 \pm 10.52$ for fully asymmetric trees; compare Fig. 10d).

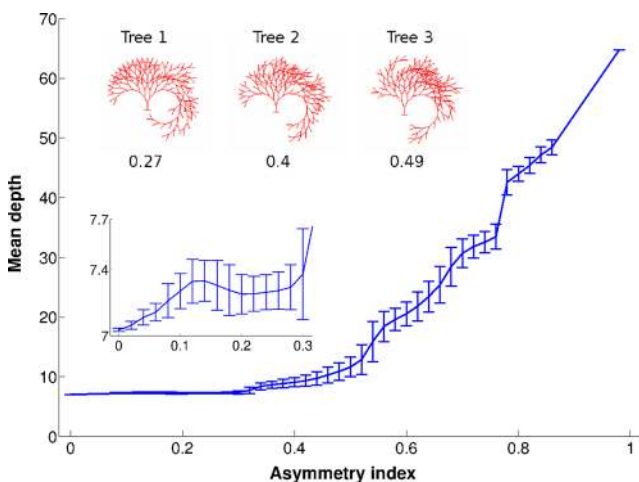


Fig. 9 Asymmetry index against mean depth for selectively generated trees with 128 terminal points. The plot shows that all trees with asymmetry index up to 0.4 have a similar low mean depth. The inset in the top left corner shows example trees with the same mean depth (8.15) but different asymmetry indices (presented below each tree). The bottom left inset highlights the trees with asymmetry index between 0 and 0.3, which all have a mean depth around 7.2. The bin width used for the asymmetry index is 0.02

3.4 The effect of dendritic tapering

The results presented so far concerned trees with branches of uniform thickness; all compartments had the same diameter. In the remaining set of simulations, we investigated the effect of dendritic tapering on neuronal performance. In these experiments, the tapering factor (explained in Section 2.1.2) was varied from 1 down to 0.7 and applied to all dendritic branches, from the soma to the terminal

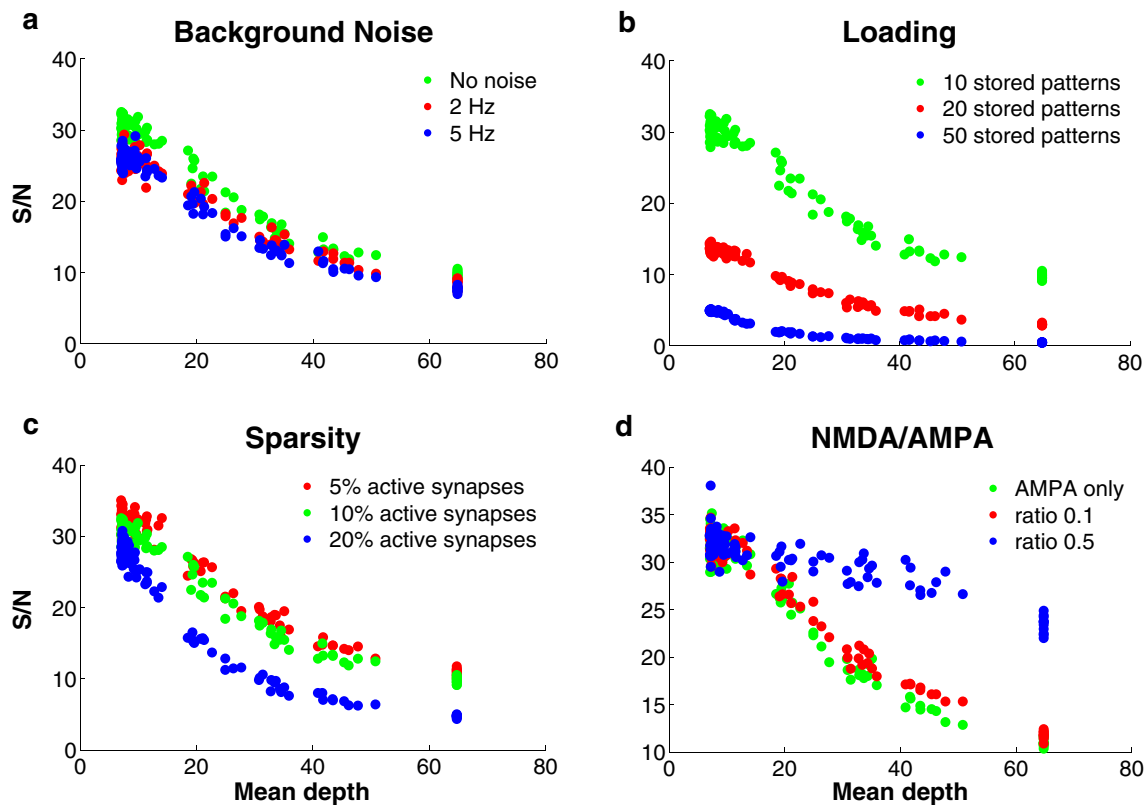


Fig. 10 Robustness of the pattern recognition performance in passive neurons with 128 terminal points. Results were obtained by averaging the pattern recognition task over 100 trials for each of the 100 samples of morphologies used in Figure 7a. The green data points in each panel represent the control simulation, where the parameters used were the same as in Figure 7. Panel (a) shows that the pattern recognition task is robust to varying amounts of background noise. In (b), we demonstrate that the pattern recognition performance was affected by increasing

the number of stored patterns presented to the model, but the overall pattern of performance against mean depth persists. Panel (c) shows that the results are robust when the sparsity varied. Panel (d) shows that the ratio of NMDA and AMPA receptors affects pattern recognition performance, whilst preserving the inverse correlation between performance and mean depth. The NMDA receptors were modelled as described in Graham (2001)

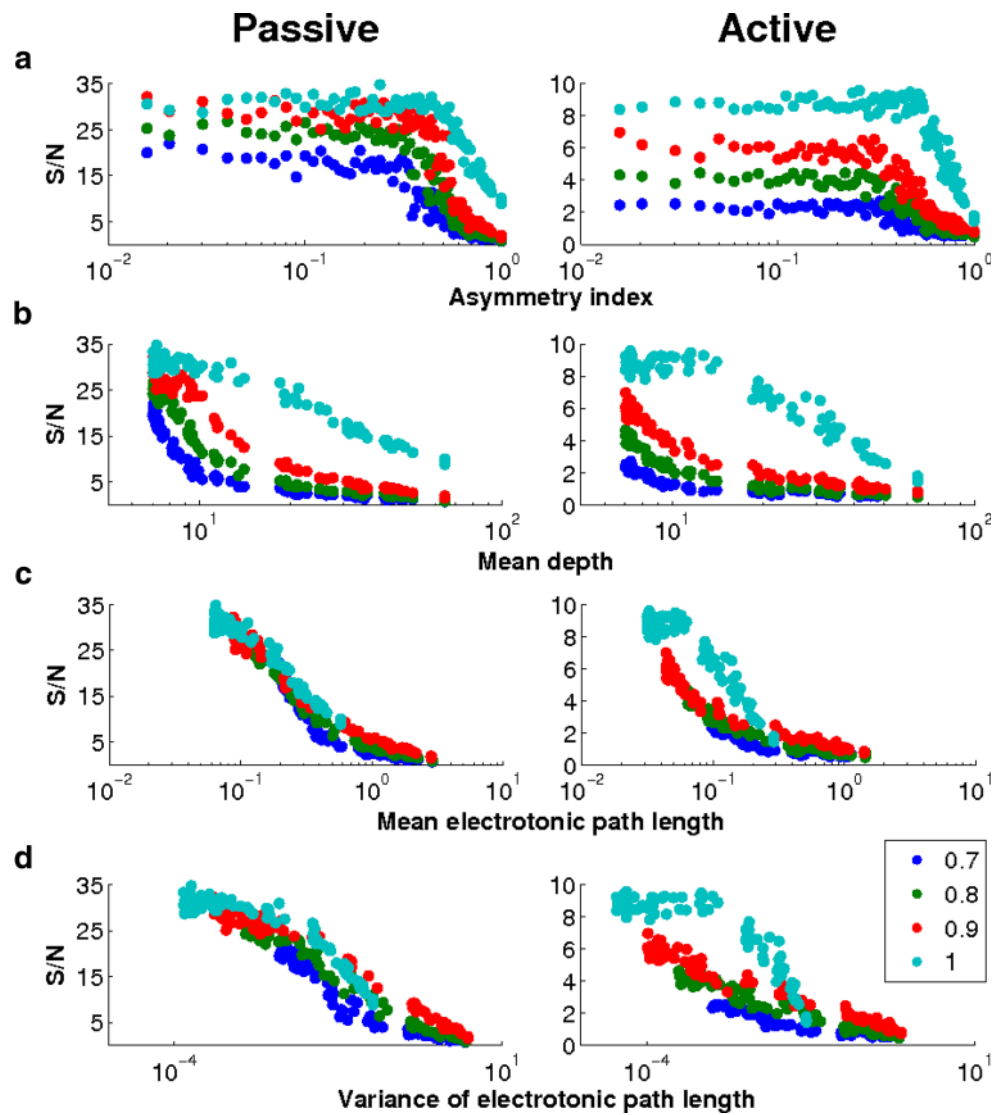
points, until a minimum allowed diameter of $0.1 \mu\text{m}$ was reached. To analyse the results for dendritic trees in the presence of tapering, we calculated two new metrics, the mean and variance of the electrotonic path length, which take into account the dendritic compartmental diameter (see Eqs. (4) to (6)). Comparing the results for these metrics with both metrics used in the previous experiments, asymmetry index and mean depth, we found that the mean and variance of the electrotonic path length correlated better with neuronal performance in both passive and active models (see Fig. 11). From this figure, we can see that the mean and variance of the electrotonic path length are robust predictors of the pattern recognition performance of passive neuronal models even when trees with different degrees of tapering are compared (Fig. 11c and d). In neuronal models with active conductances, the mean and variance of the electrotonic path length correlate well

with performance for tapering factors between 0.7 and 0.9 (Fig. 11c and d). The other two metrics, in contrast, show a much poorer relationship, which moreover strongly depends on the tapering factor used (Fig. 11a and b).

4 Discussion

The main result of this paper is that the dendritic morphology of a neuron has a major effect on its pattern recognition performance. To study how dendritic morphology affects pattern recognition performance, we generated all possible dendritic trees with 22 terminal points, and compared simulations of these smaller trees to a representative selection of larger trees with 128 terminal points. In both cases, the fully symmetric morphologies showed a better performance when compared to the fully asymmetric ones. However, the

Fig. 11 Pattern recognition performance of model neurons in the presence of dendritic tapering. The tapering factor was varied from 0.7 to 1. The morphological metrics used, asymmetry index (a), mean depth (b), and mean and variance of electrotonic path length (c, d) are plotted on a logarithmic scale for visualisation purposes. The signal-to-noise ratio was calculated by averaging over 100 trials of 20 patterns

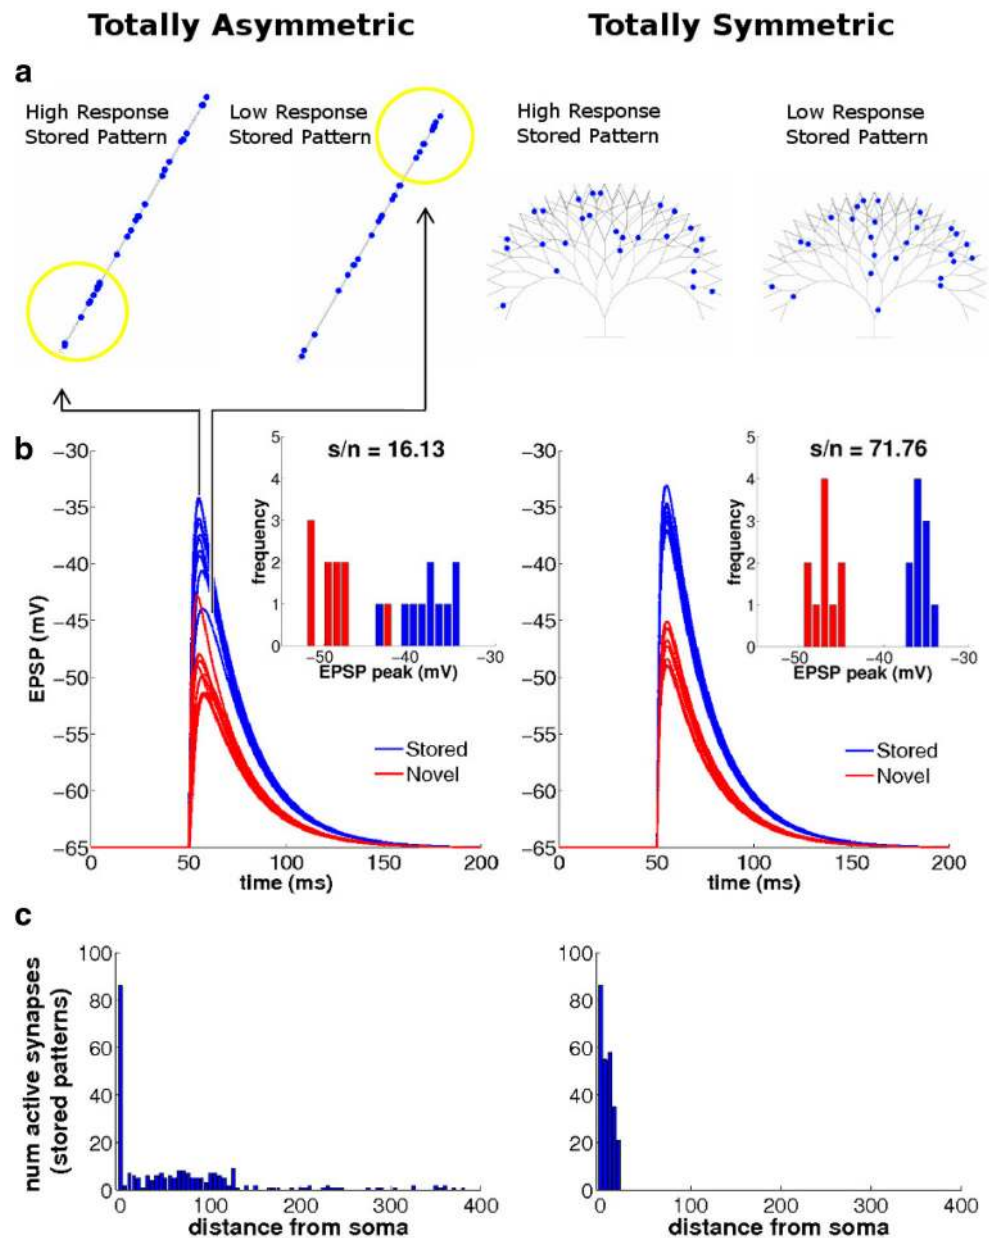


results for the selectively generated trees with 128 terminal points showed that the dendritic morphologies with an asymmetry index up to 0.4 performed as well as the most symmetric ones. The inability of morphology to affect performance for very symmetric trees can be explained by the fact that all trees with an asymmetry index up to 0.4 correspond to the same set of trees with the lowest mean depth, which results in a poor overall correlation between asymmetry index and neuronal performance.

We also found that the mean depth of the dendritic tree correlates with neuronal performance, when tapering is not present, for both active and passive models (Fig. 8b). However, when dendritic tapering was introduced, the mean depth correlated less well with performance (Fig. 11b). The same experiment showed that the mean and variance of the electrotonic path length were the best predictors of pattern recognition performance in both active and passive

models (Figs. 11c and d). The reason for the better pattern recognition performance of the neurons with dendritic trees with smaller mean electrotonic path lengths (or, in the absence of tapering and for constant compartment lengths, equivalently, smaller mean depths or mean path lengths) is illustrated in Fig. 12. The dendritic trees with the smallest possible mean electrotonic path length are the most symmetric ones (shown on the right in Fig. 12). In these fully symmetric dendritic trees, the variance of somatic responses to dendritic synaptic input is minimised, maximising the signal-to-noise ratio Eq. (8) and pattern recognition performance. In contrast, asymmetric neurons with large mean electrotonic path lengths (illustrated on the left of Fig. 12) are more likely to receive input patterns with active synapses located predominantly near the distal or proximal end of the dendrite (highlighted by the yellow circles in Fig. 12). These input patterns with distal or proximal activation biases

Fig. 12 Comparison of pattern recognition in neurons with the most symmetric and the most asymmetric dendritic morphologies. The *s/n* ratios, shown on the top of each histogram, are calculated using the EPSP peaks resulting from the presentation of sets of stored (*blue*) and novel (*red*) patterns. On top of each of the EPSP traces the neuronal morphologies used in this experiment are shown, with *blue dots* that represent the location of each active synapse for the lowest and highest response obtained for stored patterns. The *yellow circles* shown for the asymmetric morphology indicate the location of clusters of active synapses that are located predominantly at the distal or proximal end of the dendrite, leading to the lowest or highest response, respectively. The spatial distributions of active synapses for the stored patterns (*bottom graphs, x-axis in μm*) also explain why the performance is better for the most symmetric morphology (*right*), as this morphology has a larger number of active synapses closer to the soma and consequently, a smaller variance of synaptic distances and somatic voltage responses when compared to the most asymmetric morphology (*left*)



will result in particularly small or large somatic potentials, respectively (or, in the active model, small or large numbers of spikes), which increases the range of responses to input patterns, leads to a larger response variance and a smaller signal-to-noise ratio.

The present study is an extension of previous work by us (Steuber et al. 2007) and others (Graham 2001) that has used input patterns with synapses that are activated synchronously by single pulses. The nature of the patterns of neuronal activity that are stored and recalled in real neuronal systems is not known, although the sparse activity that has been recorded in many neuronal systems suggests that some neurons have to decode patterns of single spikes or bursts of spikes (Chadderton et al. 2004). Other studies (Poirazi

et al. 2003b) have used input patterns where synapses were activated by high-frequency spike trains. However, a companion paper by the same authors (Poirazi et al. 2003a) has shown that although the type of input pattern (single pulses vs spike train) affects the exact shape of the neuronal input-output relation, the type of arithmetic operations performed by the neurons is the same for both types of input patterns. Whilst we have used an evolutionary algorithm to optimise the number of spikes that each active synapse receives in another study (de Sousa et al. 2012), in the current study we have therefore used simple types of input patterns with one spike or a short burst of spikes for each synapse. Although the type of input pattern may affect the value of the signal-to-noise ratio and hence the

pattern recognition performance, we never found it to affect the shape of the relationships between dendritic morphology and pattern recognition described in the present paper.

Although our conclusions are based on simulations of binary trees trained to recognise random (and hence uncorrelated) input patterns, we think they can be generalised to more anatomically constrained input configurations, like those which neurons receive for instance in layered structures such as neocortex. Indeed, in the present pattern recognition task, the neurons had to summate the weights of the synapses that were activated by the pattern, and at least in the active models, compare this sum to a reference, their spike threshold (Willshaw et al. 1969). As the input patterns caused transient responses, the neurons also had to act as coincidence detectors, to which symmetric neurons arguably are better suited. In a study of model neurons for interaural coincidence detection, Agmon-Snir et al. (1998) reported another advantage of neurons with symmetric trees, which holds even in the absence of learning: the threshold intensity needed to fire them is lower for spatially balanced than for unbalanced stimuli. However, this is not to say that asymmetric trees would always have a negative effect on pattern recognition. If neurons were to recognise asynchronous input patterns, asymmetric trees would offer an advantage by slowing down the propagation of the earliest EPSPs so as to synchronise their arrival at the soma with the EPSPs of later activated synapses (Rall 1964). Moreover, the optimal shape of a dendritic tree will be affected by other factors, such as the desire to maximise the number of possible connectivity patterns between dendrites and neighbouring axons (Wen et al. 2009).

From the present results, and more particularly from the inverse relationship between electrotonic distance and pattern recognition performance (Fig. 11), one might be inclined to conclude that smaller neurons would always be better pattern recognisers than big neurons. In the limit, even a single-compartment neuron, though physically implausible, would be most cost beneficial. However, this conclusion is mistaken, as it is based on the comparison of trees built of compartments of a given fixed length. Indeed, when in another study we used genetic algorithms to optimise dendritic shape (de Sousa et al. 2012), treating compartmental length as a free parameter, we found no clear indication that neurons would minimise the length of their compartments, which obviously would further minimise both electrotonic distance and its variance. The reason is that individual synapses must be sufficiently isolated, or compartmentalised, to prevent sublinearities in the generation and summation of their EPSPs, which inevitably arise due to shunting of the current at the synapse's reversal potential (Rall 1964). One could of course linearise the interaction between synapses

by reducing their weights, but in actual neurons membrane noise may put a limit on this miniaturisation, as it does for axons (Faisal et al. 2005). Hence the trade-off between minimising the synapses' distance from the soma, and preventing sublinear interference by maximising the distance between them may be best satisfied by symmetric multi-compartmental trees. Another strategy for neurons, not covered by the present study, may be to enhance their computational capacity by taking advantage of dendritic nonlinearities and expanding them through localised, branch-specific interactions (Legenstein and Maass 2011; Poirazi and Mel 2001; Poirazi et al. 2003a, b; Cazé et al. 2013).

Conflict of interests The authors declare that they have no conflict of interest.

Appendix

Algorithm 1 Exhaustive tree generation algorithm

1. For a given number of terminal points N , create every possible partition $(a\ b)$ such that: $a \leq b$ and $a + b = N$
 2. For each new branch a and b , call the partition method recursively (step 1)
 3. Repeat steps 1-2 until $N=1$ (when a leaf has been reached)
-

Algorithm 2 Selective tree generation algorithm

1. For a given number of terminal points N , create a partition by randomly splitting N in two branches $(a\ b)$, for a and b such that: $a \leq b$ and $a + b = N$
 2. Use parameter *asym* to define whether to generate trees that are more symmetric (*asym*=0) or more asymmetric (*asym*=1)
 3. Use parameter *bias* to control the range of possible tree asymmetries, where *bias* varies from 0.01 to 0.5. To generate random partitions use *bias* equal to 0.5; for more symmetric or asymmetric partitions use *bias* equal to 0.01. For example for $m=100$ *bias*=0.1: *asym*=0 => (46 54) to (49 51); *asym*=1 => (5 95) to (10 90)
 4. For each new branch a and b , call the partition method recursively (step 1)
 5. Repeat steps 1-4 until $N=1$ (when a leaf has been reached)
 6. Repeat steps 1-5 until number of desired trees is achieved
-

References

- Agmon-Snir, H., Carr, C.E., Rinzel, J. (1998). The role of dendrites in auditory coincidence detection. *Nature*, *393*, 268–272.
- Carnevale, N., & Hines, M. (2006). *The NEURON Book*: Cambridge University Press.
- Cazé, R.D., Humphries, M., Gutkin, B. (2013). Passive dendrites enable single neurons to compute linearly non-separable functions. *PLoS Comput Biology*, *9*, e1002867.
- Chadderton, P., Margrie, T.W., Häusser, M. (2004). Integration of quanta in cerebellar granule cells during sensory processing. *Nature*, *428*, 856–860.
- Chen, J.L., Lin, W.C., Cha, J.W., So, P.T., Kubota, Y., Nedivi, E. (2011). Structural basis for the role of inhibition in facilitating adult brain plasticity. *Nat Neurosci*, *14*, 587–594.
- Cuntz, H., Borst, A., Segev, I. (2007). Optimization principles of dendritic structure. *Theoretical Biology and Medical Modelling*, *4*, 21.
- Dayan, P., & Willshaw, D.J. (1991). Optimising synaptic learning rules in linear associative memories. *Biol Cybern*, *65*, 253–265.
- de Sousa, G., Maex, R., Adams, R., Davey, N., Steuber, V. (2012). Evolving dendritic morphology and parameters in biologically realistic model neurons for pattern recognition. In Villa, A., Duch, W., Érdi, P., Masulli, F., Palm, G. (Eds.) *Artificial Neural Networks and Machine Learning - ICANN 2012 Vol. 7552 of Lecture Notes in Computer Science*, (pp. 355–362). Heidelberg: Springer Berlin.
- Faisal, A.A., White, J.A., Laughlin, S.B. (2005). Ion-channel noise places limits on the miniaturization of the brain's wiring. *Curr Biol*, *15*, 1143–1149.
- Fohlmeister, J.F., & Miller, R.F. (1997). Mechanisms by which cell geometry controls repetitive firing. *Journal of Neurophysiology*, *78*, 1948–1964.
- Graham, B.P. (2001). Pattern recognition in a compartmental model of a CA1 pyramidal neuron. *Network: Computation in Neural Systems*, *12*, 473–492.
- Gulledge, A.T., Kampa, B.M., Stuart, G.J. (2005). Synaptic integration in dendritic trees. *Journal of Neurobiology*, *64*, 75–90.
- Harding, E.F. (1971). The probabilities of rooted tree-shapes generated by random bifurcation. *Advances in Applied Probability*, *3*, 44–77.
- Hines, M.L., & Carnevale, N.T. (1997). The NEURON simulation environment. *Neural Computation*, *9*, 1179–1209.
- Krichmar, J.L., Nasuto, S.J., Scorcioni, R., Washington, S.D., Ascoli, G.A. (2002). Effects of dendritic morphology on CA3 pyramidal cell electrophysiology: a simulation study. *Brain Research*, *941*, 11–28.
- Legenstein, R., & Maass, W. (2011). Branch-specific plasticity enables self-organization of nonlinear computation in single neurons. *The Journal of Neuroscience*, *31*, 10787–10802.
- London, M., & Häusser, M. (2005). Dendritic computation. *Annual Review of Neuroscience*, *28*, 503–532.
- Mainen, Z.F., & Sejnowski, T.J. (1996). Influence of dendritic structure on firing pattern in model neocortical neurons. *Nature*, *382*, 363–366.
- Poirazi, P., Brannon, T., Mel, B.W. (2003a). Arithmetic of subthreshold synaptic summation in a model CA1 pyramidal cell. *Neuron*, *37*, 977–987.
- Poirazi, P., Brannon, T., Mel, B.W. (2003b). Pyramidal neuron as two-layer neural network. *Neuron*, *37*, 989–999.
- Poirazi, P., & Mel, B.W. (2001). Impact of active dendrites and structural plasticity on the memory capacity of neural tissue. *Neuron*, *29*, 779–796.
- Rall, W. (1964). Theoretical significance of dendritic trees for neuronal input-output relations. In Reiss, R. (Ed.) *Neural theory and modeling*, (pp. 73–97). Stanford: Stanford University Press.
- Steuber, V., Mittmann, W., Hoebeek, F.E., Silver, R.A., De Zeeuw, C.I., Häusser, M., De Schutter, E. (2007). Cerebellar LTD and pattern recognition by Purkinje cells. *Neuron*, *54*, 121–136.
- van Elburg, R.A.J., & van Ooyen, A. (2010). Impact of dendritic size and dendritic topology on burst firing in pyramidal cells. *PLoS Comput Biol*, *6*, e1000781.
- van Ooyen, A., Duijnhouwer, J., Remme, M., van Pelt, J. (2002). The effect of dendritic topology on firing patterns in model neurons. *Network: Computation in Neural Systems*, *13*, 311–325.
- van Pelt, J., & Verwer, R. (1985). Growth models (including terminal and segmental branching) for topological binary trees. *Bulletin of Mathematical Biology*, *47*, 323–336.
- van Pelt, J., Uylings, H.B.M., Verwer, R.W.H., Pentney, R.J., Woldenberg, M.J. (1992). Tree asymmetry—a sensitive and practical measure for binary topological trees. *Bulletin of Mathematical Biology*, *54*, 759–784.
- Wen, Q., & Chklovskii, D.B. (2008). A cost-benefit analysis of neuronal morphology. *Journal of Neurophysiology*, *99*, 2320–2328.
- Wen, Q., Stepanyants, A., Elston, G.N., Grosberg, A.Y., Chklovskii, D.B. (2009). Maximization of the connectivity repertoire as a statistical principle governing the shapes of dendritic arbors. *Proceedings of the National Academy of Sciences of the USA*, *106*, 12536–12541.
- Willshaw, D.J., Buneman, O.P., Longuet-Higgins, H.C. (1969). Non-holographic associative memory. *Nature*, *222*, 960–962.