A Tale of Two Trees: Modeling Apical and Basal Tree Contribution to L2/3 V1 Pyramidal Cell Orientation Selectivity

Petousakis Konstantinos-Evangelos

Master's Thesis, "Brain and Mind" Programme

University of Crete, School of Medicine, October 2018

Table of Contents

Τίτλος / Περίληψη 1

Title / Abstract 2

1. Introduction 3
    1.1. The Visual System 3
    1.2. The Hubel & Wiesel Theory of Orientation Selectivity 4
    1.3. Challenging the Hubel & Wiesel Theory 8
    1.4. Current View of V1 9
    1.5. Research Goals 11

2. Methods 12
    2.1. Model Description 12
    2.2. Model Manipulations 17
    2.3. Recording Information 18
    2.4. Simulation Protocols 18
    2.5. Causal Classification 22
    2.6. Data Acquisition and Analysis 25

3. Results 27
    3.1. Model Validation 27
    3.2. Evaluation of Dendritic Non-Linearities 28
    3.3. Dendritic Contribution to Orientation Selectivity 32
    3.4. Orientation Tuning Dominance 34
    3.5. Causal Interventions 35

4. Discussion 38

5. Acknowledgments 44

6. References 44

7. Supplementary Information 47

Μία Ιστορία Δύο Δέντρων: Μοντελοποίηση της Συνεισφοράς Κορυφαίων και Βασικών Δενδριτών στην Επιλεκτικότητα Διεύθυνσης Πυραμιδικών Νευρώνων των Στοιβάδων 2/3 του Πρωτοταγούς Οπτικού Φλοιού

Περίληψη

Οι πυραμιδικοί νευρώνες αποτελούν βασικό στοιχείο των φλοιικών περιοχών και δέχονται πληθώρα σημάτων από διάφορες περιοχές. Προσαγωγές συνάψεις άπτονται είτε του κορυφαίου είτε του βασικού δενδριτικού δένδρου, περιοχές με έντονη μορφολογική ποικιλομορφία. Και τα δύο δένδρα συνεισφέρουν διαφορικά στην απόκριση του σώματος του νευρώνα, αλλά οι ακριβείς ρόλοι τους παραμένουν ασαφείς. Ανατροφοδοτικά σήματα προς τους κορυφαίους δενδρίτες ολοκληρώνονται μαζικά στον κορυφαίο κορμό και μεταβαίνουν προς το σώμα. Οι βασικοί δενδρίτες, από την άλλη, εκφύονται άμεσα από το σώμα και δέχονται εμπροσθόδρομα σήματα που ολοκληρώνονται ημιανεξάρτητα. Άρα, τα δένδρα αυτά αποτελούν διακριτές ανατομικές και πιθανώς λειτουργικές υπομονάδες. Για να αξιολογήσουμε την ορθότητα του δευτέρου, μοντελοποιήσαμε το περίπλοκο μοτίβο απόκρισης ενός πυραμιδικού νευρώνα των στοιβάδων 2/3 του πρωτοταγούς οπτικού φλοιού σε χωρικά καλώς διατεταγμένα οπτικά ερεθίσματα. Ο στόχος μας ήταν η διαλεύκανση της συνεισφοράς του κάθε δένδρου στο μοτίβο απόκρισης του νευρώνα. Για την επίτευξη του στόχου αυτού, δημιουργήσαμε ένα μορφολογικά λεπτομερές μοντέλο ενός κυττάρου στο περιβάλλον προσομοιώσεων NEURON. Η ορθότητα του μοντέλου επαληθεύτηκε μέσω σύγκρισης της συμπεριφοράς του με δεδομένα ηλεκτροφυσιολογίας από *in vivo* και *in vitro* καταγραφές. Ερευνήσαμε το ρόλο της δενδριτικής ολοκλήρωσης στους βασικούς και κορυφαίους δενδρίτες, καθώς και τη συνεισφορά της στο σχηματισμό της νευρωνικής απόκρισης. Τα αποτελέσματα υποδεικνύουν ότι σωματικά

δυναμικά ενέργειας παράγονται μόνο όταν σήματα εισόδου συμπίπτουν αμφιπλεύρως, καθώς μονόπλευρα ερεθίσματα γενικά δεν είναι ικανά να παράξουν επαρκή απόκριση του σώματος. Επιπρόσθετα, δεδομένου ότι υπάρχει ισοκατανομή των συνάψεων, οι αποκρίσης του νευρώνα φαίνεται να εκκινούνται από το κορυφαίο δένδρο, καθώς η παραγωγή αιχμών δυναμικού του προηγείται χρονικά αντίστοιχης δραστηριότητας του σώματος. Τελικά, η δραστηριότητα του βασικού δένδρου, ως εκπόλωση ή παραγωγή αιχμών, είναι απαραίτητη για την παραγωγή σωματικής δραστηριότητας, παρά το γεγονός ότι οι περισσότερες αιχμές δυναμικού του σώματος εκκινούνται από το κορυφαίο δένδρο. Το παρόν μοντέλο παρέχει στοιχεία υπερ διακριτών υπολογισμών που λαμβάνουν χώρα στα βασικά και κορυφαία δενδριτικά πεδία, και τονίζει το ρόλο σημάτων πρόβλεψης και προσοχής.

A Tale of Two Trees: Modeling Apical and Basal Tree Contribution to L2/3 V1 Pyramidal Cell Orientation Selectivity

Abstract

Pyramidal neurons, a mainstay of cortical regions, receive a plethora of inputs from various areas. Afferent synapses are received by either the apical or basal dendritic trees, which are morphologically distinct. Both trees differentially contribute to the somatic response, although their exact functional roles remain unclear. Feedback inputs to apical dendrites are integrated en masse at the apical trunk and propagate to the soma. Basal dendrites, on the other hand, branch out from the soma, with feedforward inputs being integrated semi-independently. Thus, these trees define distinct anatomical and possibly functional sub-units. To assess the latter, we modeled the complex response pattern of the L2/3 V1 pyramidal neuron to spatially tuned

synaptic input. Our goal was to elucidate the contribution of each tree to the response pattern of the neuron. Towards this goal, we created a morphologically detailed computational model of a single cell in the NEURON simulation environment. The model was validated against electrophysiological data recorded *in vitro* and *in vivo*. We investigated the role of dendritic integration at the basal and apical trees, and its contribution in shaping cell responses. Results indicate that somatic action potentials are generated only when input coincides bilaterally, as unilateral stimuli are generally unable to evoke an adequate response at the soma. In addition, given equal synaptic drive, the responses of the neuron appear to be initiated by the apical tree, as its dendritic spiking activity temporally precedes somatic spike-like activity. Finally, basal tree activity, in the form of either depolarization or spiking, is essential for producing somatic activity, despite the fact that most somatic spikes are apically-driven. This model provides evidence for distinct computations taking place in the basal and apical trees of the neuron, and emphasizes the role of predictive and attentional feedback input.

1. Introduction

1.1. The Visual System

It is generally accepted that of all the senses humans possess, vision is the one that is most relied upon. As such, it is natural that in our attempts to understand sensory perception, the visual system is one of the most heavily examined. In broad terms, light enters our eyes through the pupil and is focused through the lens, which is warped into shape by the ciliary muscle. Reaching the retina, it passes through layers of ocular cells, reaching the photoreceptor layer. There,

104 photons excite the opsins that lie within the cones and rods, producing electrochemical impulses
105 that travel outward through connections with bipolar cells to the ganglion cell layer, and from
106 there, to the optic nerve. The two optic nerves travel through the cranium, with half of each
107 nerve crossing over to the contralateral side at the optic chiasm. Reaching the lateral geniculate
108 nucleus of the thalamus (LGN), the signal propagates through thalamocortical relay connections
109 mostly to layer 4 (L4) of the primary visual cortex (V1) (*Figure 1*). The hierarchy of visual
110 processing thereafter involves forwarding signals to L2/3 of V1, then L5, and afterwards to higher-
111 order cortices of the visual pathway (Salin & Bullier, 1995; Sun, Tan, Mensh, & Ji, 2016). The
112 output of visual cortex neurons is clearly influenced by visual stimuli, and as such, there are areas
113 of the visual space where the presence of a visual stimulus elicits a neuronal response. This space
114 is the *receptive field* of the responding neuron. The receptive field can also be described in terms
115 of retinal position; light stimulation of retinal areas will elicit a response from the neuron only if
116 the area being stimulated is part of the retinal receptive field of the neuron.
117

118 **1.2. The Hubel & Wiesel Theory of Orientation Selectivity**

119 In 1962, David Hubel and Torsten Wiesel introduced their theory on the architecture of the visual
120 cortex of the cat (Hubel & Wiesel, 1962). Two of the questions to be answered concerned the
121 organization of the receptive fields of V1 neurons, as well as their responses to varying visual
122 stimuli. To investigate this, the researchers anesthetized and prepared cats, inserting recording
123 electrodes into the top layers of the V1 area (3 – 4 mm), while presenting light stimuli (either
124 small dots or rod-shaped) to the contralateral eye of the animal.

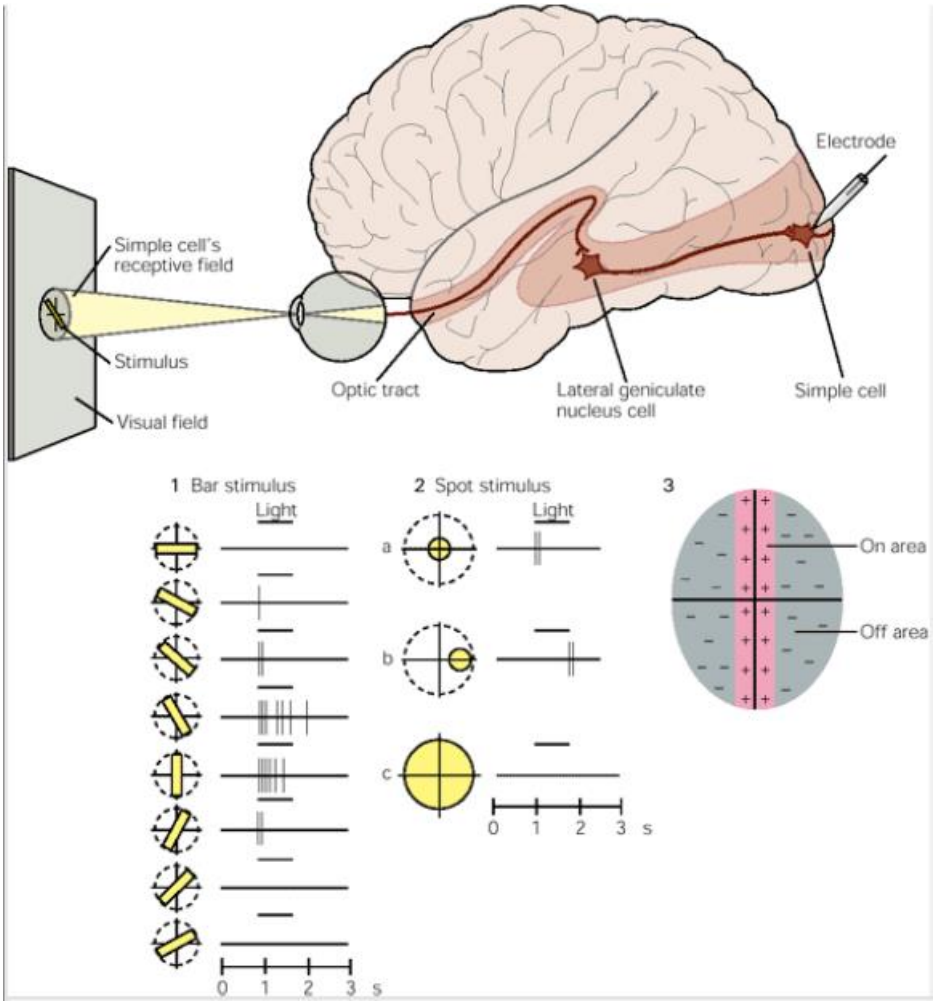

Figure 1. *Simple sketch of the visual pathway in a human brain. Presentation of bar (1) or spot (2) light stimuli generates a response in V1 simple cells commensurate with the orientation or receptive field position of the stimulus, a phenomenon known as orientation selectivity in the first case. These receptive fields are subdivided into on-areas and off-areas (3) that produce excitation or inhibition in the neuron, respectively, when excited by light. Reproduced from Principles of Neural Science, 4th Edition (Kandel et al., 2000).*

It was already known at the time that retinal ganglion cells exhibited a particular receptive field organization conforming to one of two types: *on-center* and *off-center* (Kuffler, 1953). This same organization was observed in the LGN cells that relayed signals from the retina to V1. In general, light stimulation of retinal areas was observed to produce either excitation or inhibition in upstream neurons. These regions are thus named ON and OFF, respectively, in relation to their

effects on a specific neuron. This nomenclature greatly facilitates the description of visual cortex receptive fields.

While mapping the receptive fields of specific V1 cells, it was observed that they tended to conform to one of two categories: *simple* and *complex*. Simple cell receptive fields consisted of an elongated central ON area flanked by one (or two) OFF area(s), or vice versa (*Figure 1*). These simple cells could thus be labeled as on-center or off-center, similarly to LGN relay cells. Complex cells, on the other hand, conformed to no such stereotypical receptive field structure. Instead, they "*responded to variously-shaped stationary or moving forms in a way that could not be predicted from maps made with small circular spots*" (Hubel & Wiesel, 1962). In addition, even when ON and OFF regions could be identified, they did not have the same configuration or properties as these of simple cells.

Another important observation was that the action potential generation rate (*firing rate*) of V1 neurons was modulated by the orientation of a rod-like light stimulus – a phenomenon now named *orientation selectivity*. These neurons exhibited a maximum firing rate for a particular, preferred orientation (*Figure 1*). At the same time, they exhibited a minimal response to stimuli of an orientation orthogonal to the preferred one. The preferred orientation of the neuron could be reliably predicted through the orientation of the ON region of its receptive field (Hubel & Wiesel, 1962). Given the responses of a neuron to stimuli of various orientations, we can generate the *tuning curve* of the neuron (see *Figure 9* for an example), which allows the visualization of the response pattern of the cell by plotting the firing frequency against the corresponding stimulus orientations.

Driven by this observation, Hubel and Wiesel proposed a theory to explain how simple

cell receptive fields arise. They hypothesized that simple cells receive input from LGN relay cells of the same type (on-center or off-center) whose retinal receptive fields lie along a well-defined orientation. As a result, the receptive field of the simple cell is a summation of the receptive fields of these LGN relay cells (*Figure 2-A*), and thus exhibits the structure that was experimentally observed: elongated ON/OFF areas flanked by one (or two) OFF/ON area(s). Similarly, the erratic receptive field structure of complex cells could be explained, Hubel and Wiesel argued, as consisting of the combined receptive fields of multiple simple cells, all of which project to this one complex cell (*Figure 2-B*).

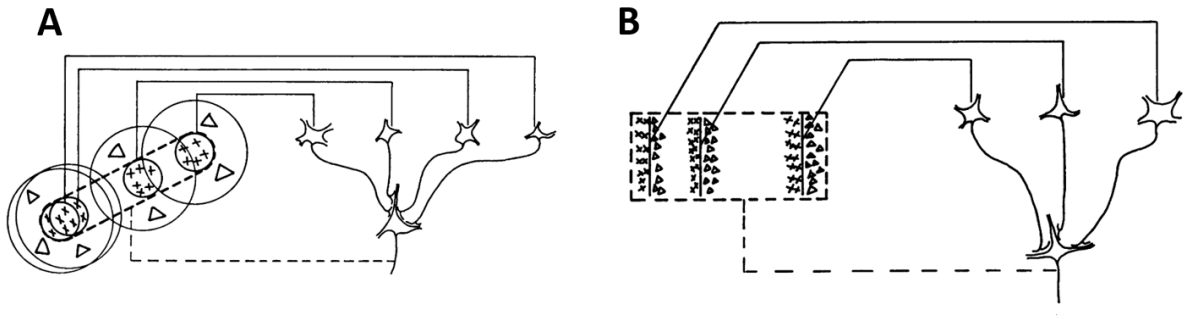

Figure 2. *Feedforward models of simple (A) and complex (B) cells. ON regions noted with plus signs. OFF regions noted with triangle signs. Reproduced from Hubel & Wiesel, 1962.*

When comparing the visual system across different model species, however, a number of differences emerge. In rodents, orientation-selective cells akin to V1 simple cells have been discovered in the thalamus itself (Scholl, Tan, Corey, & Priebe, 2013), and direct thalamocortical projections to L1 of V1 have also been observed (Roth et al., 2016). In addition, rodent V1 organization follows a dispersed salt-and-pepper motif, unlike the highly structured organization typical of higher mammals like primates and cats (Ohki & Reid, 2007). These differences are useful in understanding the function of the visual system when trying to infer general rules of function from information derived from different model animals.

1.3. Challenging the Hubel & Wiesel theory

The theory of Hubel and Wiesel describes a feedforward model, as it explains the properties of receptive fields of V1 neurons solely in terms of well-arranged visual (feedforward) input to the V1 neurons. However, it fails to explain certain properties of V1 neurons, such as cross-orientation suppression, contrast-invariant orientation tuning (Priebe, 2016) and more. Cross-orientation suppression refers to the reduction in neuronal output when stimuli of both the preferred and orthogonal orientation are presented. Contrast-invariant orientation tuning is the robustness of orientation selectivity despite large changes in stimulus contrast (Priebe & Ferster, 2012). Another overlooked factor is inhibition, which also plays a crucial role in shaping neuronal activity, both as a modulator of neuronal responses to stimulation (Haider, Häusser, & Carandini, 2013), as well as a regulator of dendritic excitability and plasticity (Gidon & Segev, 2012). Given the recurrent nature of connections in the cortex, inhibitory regulation of pyramidal cell activity feeds back into the interneurons themselves, largely defining the activity of the network (Palmer, Murayama, & Larkum, 2012). It is thus evident that the simple feedforward model cannot fully explain the behavior of V1 neurons.

More recently, a new type of model has been gaining traction; predictive coding. In this model, V1 neurons rely on feedforward as well as feedback input to produce their output. More

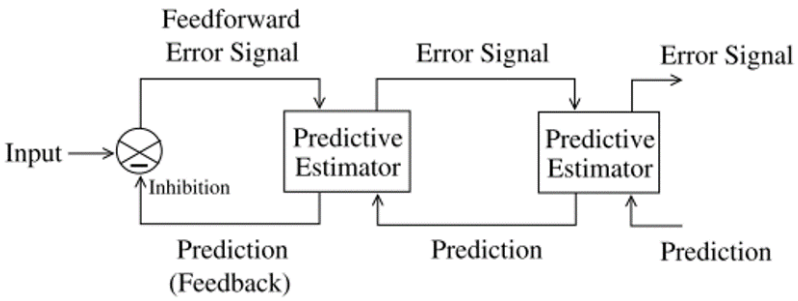

Figure 3. *Schematic representation of the predictive coding model. Reproduced from Rao & Ballard, 1999.*

specifically, feedback connections from higher-order cortical areas carry a prediction signal that attempts to predict the activity of the lower-level area, and feedforward connections from lower-order areas to higher-order areas carry a residual error signal, relaying the difference of the actual activity from the predicted activity to the neurons of the higher-order areas (Rao & Ballard, 1999) (*Figure 3*). In this model, predictive input relies on "*an efficient internal model of natural images*" (Rao & Ballard, 1999), which is to say, Bayesian prior probabilities derived from the environment. This type of model, when simulated computationally using multi-layer artificial neural networks (ANNs), manages to independently exhibit the phenomenon of endstopping (Rao & Ballard, 1999); the suppression of neuronal response when the stimulus extends beyond its receptive field (Gilbert, 1977). As this behavior cannot be explained with the simple feedforward model, it provides evidence in support of predictive coding being in use in the visual cortex. This idea has been receiving more exposure lately, having been implicated in such phenomena as saliency maps (Spratling, 2012), the interpolation of vision in the retinal blind spot (Raman & Sarkar, 2016), as well as motion perception (Edwards, Vetter, McGruer, Petro, & Muckli, 2017).

1.4. Current view of V1

Currently, we know that in V1, layer 2 and 3 (L2/3) pyramidal neurons conform to a stereotypical morphology characterized by separate apical and basal dendritic arbors. The apical tree consists of a thick apical trunk that extends into L1 and splits into an apical tuft. The basal tree consists of numerous dendrites, the majority of which sprout directly from the base of the soma and can thus directly influence somatic output (Spruston, 2008). This morphological

compartmentalization of V1 neurons is reflected in the wiring diagram describing each compartment: L4 and L2/3 pyramidal neurons synapse with basal dendrites of L2/3 pyramidal neurons, providing them with feedforward input. The apical dendrites of the L2/3 pyramidal neurons instead receive signals from hierarchically superior areas of the cortex, supplying the neurons with feedback input (Coogan & Burkhalter, 1990; M. Larkum, 2013), as well as orientation-tuned thalamocortical input from L1 (Chen et al., 2013; Cruz-Martín et al., 2014; Jia, Rochefort, Chen, & Konnerth, 2010; Roth et al., 2015). These two input streams are transformed via the different properties of the two dendritic trees and are integrated at the soma, producing neuronal output. The anatomical features of the apical and basal trees are also complemented by their distinct biophysical properties, documented as ion channels of different types and conductances along the respective trees (Cho et al., 2008). Further complicating input processing, in vitro studies have revealed multiple non-linear properties of synaptic integration, including active dendritic spikes (Häusser, Spruston, & Stuart, 2000; Spruston, 2008), the backpropagation of action potentials (G. J. Stuart & Sakmann, 1994; G. Stuart, Spruston, Sakmann, & Hausser, 1997), dendritic properties of coincidence detection (M. E. Larkum, Zhu, & Sakmann, 1999; Shai, Anastassiou, Larkum, & Koch, 2015) and others. These dendritic non-linearities greatly augment the repertoire of single neuron computations (Poirazi, Brannon, & Mel, 2003; Silver, 2010). Thus, in a realistic theory of visual processing, it is essential to scrutinize the input structure along the distinct dendritic arbors.

Few studies have undertaken the technically difficult task of measuring single spine and dendritic branch tuning properties in the visual cortex (Chen et al., 2013; Iacaruso, Gasler, & Hofer, 2017; Jia et al., 2010). Jia et al. (2010) showed that synaptic inputs of different orientation

preferences are distributed pseudorandomly throughout the dendritic tree. However, although dendrites receive functionally diverse inputs from extended regions of the visual space, spines with receptive fields similar to the soma are located more proximally on dendrites with low branch order. In contrast, spines with different receptive fields synapse on higher-order dendritic branches (Iacaruso et al., 2017). Most interestingly, Smith et al. (Smith, Smith, Branco, & Häusser, 2013) suggested that dendritic spikes in the apical tuft of L2/3 V1 pyramidal neurons enhance orientation selectivity, thereby contributing to a behaviorally relevant computation. These seminal studies (Chen et al., 2013; Iacaruso et al., 2017; Jia et al., 2010; Smith et al., 2013) began to decipher how dendritic processes shape neuronal properties.

Unfortunately, the numerous experimental and technical difficulties associated with the examination of dendritic activity (i.e. difficulties in long-term electrode placement, deconvolution of dendritic signal from somatic activity) has impeded rapid progress in the field. Thus, a viable alternative is turning to computational models, the domain of Computational Neuroscience. These models allow the researcher to rapidly search through the expansive parameter space of the model, locating conditions that produce interesting results. In computational experiments we can simulate nerve cells, or even networks of such, in an attempt to pinpoint promising hypotheses to be tested *in vivo*. In addition, this approach does not require any animals, special reagents or other materials, making it very inexpensive in comparison with *in vivo* and *in vitro* models.

1.5. Research Goals

In this work, we have created a computational model of a L2/3 V1 pyramidal cell as a biophysically

accurate and morphologically detailed reconstruction. After extensively validating the properties of this model, we have attempted to understand how somatic output is produced through the activity of the apical and basal dendritic trees. It is known that dendrites are capable of producing nonlinear input-output relations (*dendritic nonlinearities*) through the activity of their sodium, NMDA and calcium channels (Häusser et al., 2000; G. J. Stuart & Spruston, 2015), and that cooperativity between dendrites of the same tree facilitates action potential generation (Smith et al., 2013). Given that the pyramidal neurons of V1 respond to appropriate visual stimuli by producing action potentials, it is obvious that the combined activity of the apical and basal trees near the time of action potential generation must causally contribute to that event. Thus, our goal was to examine the activity of these neuronal compartments during a short time window before and after a spiking event in order to ascertain the causal triggers of neuronal activity in our L2/3 V1 pyramidal neuron model. To accomplish our goal, the model was subjected to extensive interventions aimed at modifying its input and biophysical properties in a carefully controlled manner. This allowed us to locate model configurations which, through their activity, betrayed the inner workings of the neuron. In doing so, we hoped to be able to identify the computations performed by these two distinct areas, and thus describe the *modus operandi* of the entire neuron.

2. Methods

2.1. Model Description

The model used in this work was derived from the L2/3 V1 pyramidal cell model of Papoutsi et

al., (Papoutsi, Park, Ash, Smirnakis & Poirazi, 2017) created in the NEURON simulation environment (Hines & Carnevale, 2001). It is a morphologically detailed reconstruction, featuring 43 apical dendrites (IDs A0 - A42), 7 basal dendrites (IDs B0 - B6) and an axon (*Figure 4*). As the model makes use of a variety of passive and active mechanisms (*Tables 1-4*), all electrophysiological properties were also validated against experimental data (Papoutsi et al., 2017).

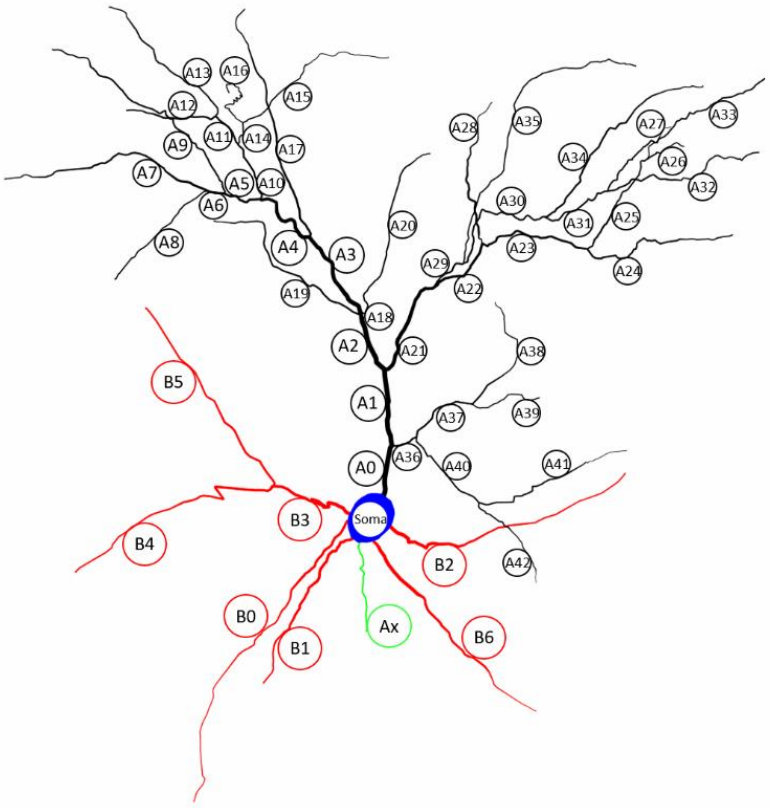

Figure 4*. Schematic of the L2/3 V1 pyramidal cell model. Compartments are annotated with their IDs. Blue: Soma. Red: Basal dendrites. Black: Apical dendrites. Green: Axon.*

Compartment Type	Passive/Active Mechanisms	Synaptic Mechanisms
Soma	Hodgkin/Huxley voltage-gated Na^+ channels Hodgkin/Huxley voltage-gated K^+ channels Muscarinic voltage-gated K^+ channels A-Type voltage-gated K^+ channels T-Type Ca^{++} channels High voltage activated (HVA) Ca^{++} channels Calcium-dependent K^+ channels Active ATP Ca^{++} pumps	$GABA_A$ (background-driven)
Basal Dendrites		AMPA (background-driven) NMDA (background-driven) $GABA_A$ (background-driven) AMPA (stimulus-driven) NMDA (stimulus-driven) $GABA_A$ (stimulus-driven)
Apical Dendrites		

281 ***Table 1****. Outline of passive, active and synaptic mechanisms present in the model neuron.*

282

Conductance (mS/cm^2)	Soma	Apical	Basal
g_{Na}	0.505	0.303	0.303
g_{Kdr}	0.05	$1.5*10^{-3}$	$1.5*10^{-3}$
g_{Km}	$2.8*10^{-3}$	$1.27*10^{-3}$	$1.27*10^{-3}$
g_A	5.4	Diameter ≤ 0.8 μm: 108 Diameter > 0.8 μm: 10.8	Diameter ≤ 0.8 μm: 108 Diameter > 0.8 μm: 10.8
g_T	0.03	x ≤ 260 μm: $0.029*\sin(0.009*x+0.88)$ x > 260μm: 0.012	$0.03+6*10^{-5}*x$
g_{HVA}	$0.05*10^{-3}$	x ≤ 260 μm: $0.049*10^{-3}*\sin(0.009*x+0.88)$ x > 260μm: $0.02*10^{-3}$	$0.05*10^{-3}+10^{-7}*x$
g_{KCa}	$2.1*10^{-3}$	$2.1*10^{-3}$	$2.1*10^{-3}$

283 ***Table 2****. Outline of membrane mechanism conductances (not synaptic). Reproduced from Papoutsi et al., 2017.*

284

	Model	Cho et al., 2010
RMP, mV	-79	-78.56 ± 1.34
IR, MΩ	123.6	125.2 ± 8.2
τ, ms	17.3	16 ± 0.7
AP amplitude, mV	66.1	67.8 ± 1.8
AP threshold, mV	-41.8	-37.7 ± 1.3
AHP, mV	17.9	13.3 ± 0.5
P-T time, ms	38.6	55.3 ± 2.7
AP adaptation	1.16	1.18 ± 0.02

285 ***Table 3****. Outline of model electrophysiological properties. RMP: resting membrane potential, IR: Input Resistance*
286 *measured at hyperpolarizing current (-0.04 nA), AP: action potential, AHP: after hyperpolarization measured at*
287 *depolarizing current (0.16 nA), P-T peak-trough. Reproduced from Papoutsi et al., 2017.*

288

	Conductance (nS)	τ_1, ms	τ_2, ms
NMDA	1.15	2	30
AMPA	0.84	0.1	2.5
$GABA_A$	1.25	0.2	1.4

Table 4. Outline of synaptic mechanism conductances and time constants. Reproduced from Papoutsi et al., 2017.

The model features both excitatory and inhibitory synaptic input. The first type is subdivided into background-driven (noise) and stimulus-driven, with the latter subdivision consisting of orientation-tuned synapses. Inhibitory synapses are modeled solely as background-driven. Synaptic input to the neuron is distributed uniformly on the target compartments using a density of 2 synapses per μm (DeFelipe & Fariñas, 1992; Schuz & Palm, 1989). Neurons with identical structural morphology but different patterns of synaptic distribution and/or input spike trains can be simulated by changing the neuron ID or the simulation ID respectively, as the random number generator seed used to produce these features is a function of the aforementioned IDs. Activation of synapses occurs via pseudo-randomly generated Poisson spike trains. Background-driven and stimulus-driven synapses feature different activation frequencies, with the frequency of the latter also varying as a function of synaptic orientation preference.

Orientation preference may be present on three levels: synaptic, dendritic and neuronal. The effect of single synapse orientation preference is explicitly modeled as an orientation weight vector biased towards a predefined mean orientation (μ_{pref}), which factors into the activation frequency of afferent stimulus-driven, orientation-tuned synapses. Thus, each synapse is assigned an orientation "tag", increasing the synaptic weight for stimuli of that orientation. Dendritic "tags" are then distributed to all synapses of each dendritic tree using a Gaussian probability density function wrapped around the unit circle (*Figure 5*). As a result, the dendritic

$$f(s) = \frac{1}{W\sqrt{2\pi}} \sum_{k=-1}^{1} e^{-\frac{(s-b+2\pi k)^2}{2W^2}}$$

Figure 5. *Wrapped Gaussian distribution, used to allocate synaptic tags to all synapses. W: Standard deviation (tuning width). s: Synaptic orientation preference (tag). b: Mean of the distribution (preferred orientation). k: Circularization term.*

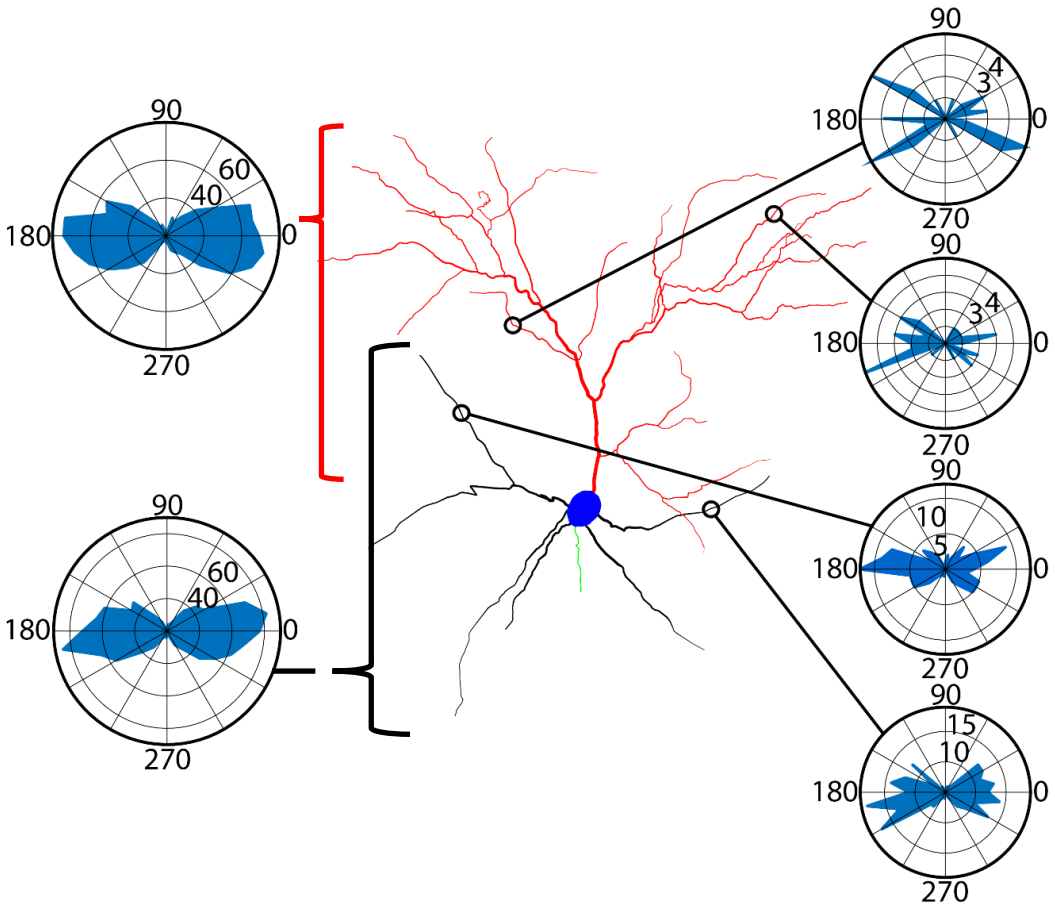

Figure 6*. Polar plots of synaptic tag distribution onto dendritic arbors. Note that each dendrite can have a wide array of synaptic tags.*

orientation preference distribution is biased towards the chosen mean orientation preference for each tree and features a nonzero width as a result of the standard deviation of the distribution (*Figure 6*). The apical and basal dendritic arbors can have different mean orientation preferences, with their difference being tuning disparity (Δ). Naturally, neuronal orientation preference

emerges through the combinatory effect of all synaptic and dendritic orientation preferences, without requiring explicit, *ad hoc* modeling.

The proportions of excitatory and inhibitory synapses to the apical and basal tree can follow one of two different configurations. Initial synapse count distribution was 40% apical to 60% basal, in adherence with experimental data (DeFelipe & Fariñas, 1992). A modified version of the model is also used, with a distribution of 50% to both trees, so as to investigate the effects of morphology on visual processing separately from those produced from input structure. The amount of inhibition to the soma was the same in both configurations, at 7% of total inhibitory input. The model can thus be used in either of two configurations of synaptic distribution – *even* (excitation: 50% apical, 50% basal; inhibition: 46.5% apical, 46.5% basal, 7% soma) and *biased* (excitation: 40% apical, 60% basal; inhibition: 33% apical, 60% basal, 7% soma). Unless noted otherwise, the even distribution is used in simulations.

2.2. Model Manipulations

A series of procedures were implemented to facilitate alteration of model parameters, allowing for the performance of simulation experiments under multiple different sets of conditions:

(1) Sodium (channel) blockage: Used to selectively nullify sodium channel conductance (g_{Na}) in a designated compartment. Eliminates action potentials when applied at the soma, allowing for dendritic voltage recordings free of back-propagating action potentials.

(2) Sodium channel weighting: Used to increase or decrease sodium channel conductance (g_{Na}) of all dendrites on the apical and/or basal tree. Achieved by multiplying the corresponding conductance values with a pre-defined weight factor for the apical and/or basal tree.

(3) Synaptic silencing: Used to selectively nullify synaptic mechanism conductances ($g_{AMPA}/g_{NMDA}/g_{GABAA}$) in a designated compartment.

(4) Input manipulation: Used to selectively de-activate any and all types of input: excitatory stimulus-driven, excitatory background-driven, inhibitory background-driven.

2.3. Recording Information

Unless noted otherwise, all recordings of model neuron output (voltage/current) were obtained at a sampling rate of 40 KHz (0.025 ms interval between data points). All data points are products of the recording. No interpolation was used to generate additional data points.

2.4. Simulation Protocols

The following simulation protocols were computationally implemented for validation and simulation experiment purposes:

(1) Paired-pulse protocol: Used for validation of dendritic non-linearities. Following complete synaptic silencing and sodium block of the entire model neuron excluding the dendritic compartment under investigation, a variable number of clustered (i.e. same attachment point on the target dendrite) excitatory synapses on the compartment are simultaneously activated twice, with a 20 ms interval (50 Hz activation frequency). Voltage at the midpoint was recorded for each dendrite examined.

(2) Iterative paired-pulse protocol: Used to evaluate dendritic non-linearities. This protocol is similar to the paired-pulse protocol described previously, with the exception that afferent

synapses increase in number from iteration to iteration, from 1 to 100 in steps of 1. We evaluated non-linear behavior caused by sodium and NMDA spikes separately. To achieve the latter, we blocked sodium channels on the selected dendrite during the protocol, allowing only AMPA and NMDA currents to act upon dendritic potential.

(3) Regular stimulation protocol: Using either the even or biased model configuration, the operation of the neuron is simulated for 2500 ms (2.5 s), with oriented stimulus onset at 500 ms (0.5 s). Voltage and/or current recordings can be obtained from any and all compartments, generally from their midpoint. No additional protocols are applied, with the exception of blocking somatic sodium in the cases where dendritic recordings are required.

(4) Orientation tuning validation protocol: Using the regular stimulation protocol, we independently simulate 10 neurons using 10 different simulation IDs for each one, resulting in 100 separate simulation categories. For each category, stimuli of 19 different orientations (0° to 180° in steps of 10°) are presented in separate simulations. Using the resulting data, we can derive the maximum and minimum firing rates of the model cell, comparing them to ones obtained from live cell recordings. By adjusting the input resistance of the neuron, as well as the frequency of synaptic activation, we constrain the minimum and maximum firing rates so as to agree with experimental data.

(5) Disparity protocol: Akin to the orientation tuning validation protocol, we independently simulate 10 neurons with 10 simulation IDs for each one, for 100 separate simulations. We also present stimuli of 4 different orientations (0° to 90° in steps of 30°). To evaluate the response of the neuron when there the orientation preference of the apical and basal trees is disparate, we keep the mean apical tree orientation preference fixed at 0°, and set the mean basal tree

orientation preference to any one of 10 different values (0° to 90° in steps of 10°). Thus, we introduce a degree of orientation tuning disparity in the model, ranging from 0° to 90°. Then, we derive the mean orientation preference of the neuron for each degree of disparity. Should the orientation preference of the neuron be closer to 0° (apical mean orientation preference) than to the mean orientation preference of the basal tree, then the apical tree dominates. Otherwise, the basal tree dominates. To easily visualize this fact, we plot neuronal responses per degree of disparity. The main diagonal of this plot represents the threshold for apical and basal dominance. Responses above the diagonal represent basal dominance, as the neuron favors orientations closer to the basal orientation preference (which will most often be greater than 0°, and thus higher on the y-axis). On the other hand, responses below the diagonal represent apical dominance, as the neuron tends to favor orientations closer to 0° regardless of basal orientation preference.

(6) Causal intervention protocol: Using fixed neuron and simulation IDs, prior knowledge of action potential occurrence is obtained via the regular stimulation protocol, and the exact timing of all action potentials is recorded. Afterwards, the protocol is similar to the regular stimulation protocol with somatic sodium blockage, until a time point t_i before a somatic action potential is to occur. At that time, sodium is blocked on either the apical or basal tree dendrites in two separate simulations. The time point t_i is defined as the time of somatic spike occurrence, offset towards zero by 1 ms plus the temporal distance of the earliest dendritic spike (of the tree to be blocked) to the somatic spike timing, limited to a time window of 3 ms prior to the somatic spike

(*Figure 7*). This protocol is repeated for all pre-recorded action potentials. Voltage recordings are obtained from all 51 compartments of the neuron. The ensuing traces are analyzed to ascertain whether the somatic action potential under investigation was rendered extinct or survived the manipulation. Depending on the result of this protocol, all action potentials can be classified in

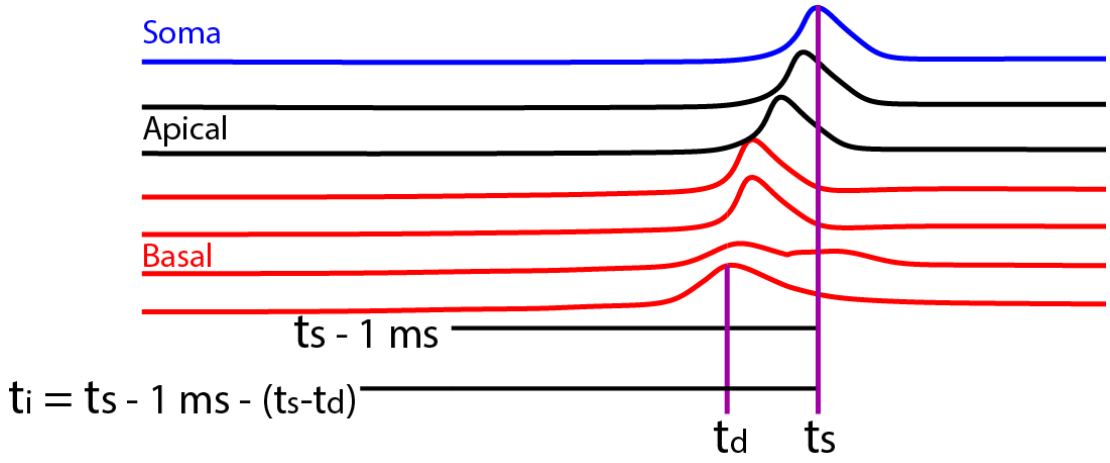

***Figure 7**. Schematic description of t_i. The timepoint t_s is the time of somatic spike occurrence. Timepoint t_d represents the occurrence of the earlies dendritic spike preceding the somatic spike, but limited to within 3 ms of somatic spike occurrence.*

one of four different categories, based on the most likely causal instigator (*Figure 8*): apically-driven, basally-driven, bistable and unstable. *Apically-driven* spikes are somatic spikes produced by apical tree dendritic spiking activity. Similarly, *basally-driven* spikes are somatic spikes produced by basal tree dendritic spiking. *Bistable* spikes are somatic spikes that cannot be rendered extinct by removing either one of the two dendritic components. *Unstable* spikes, finally, are somatic spikes that are lost when either of the two dendritic arbors is silenced.

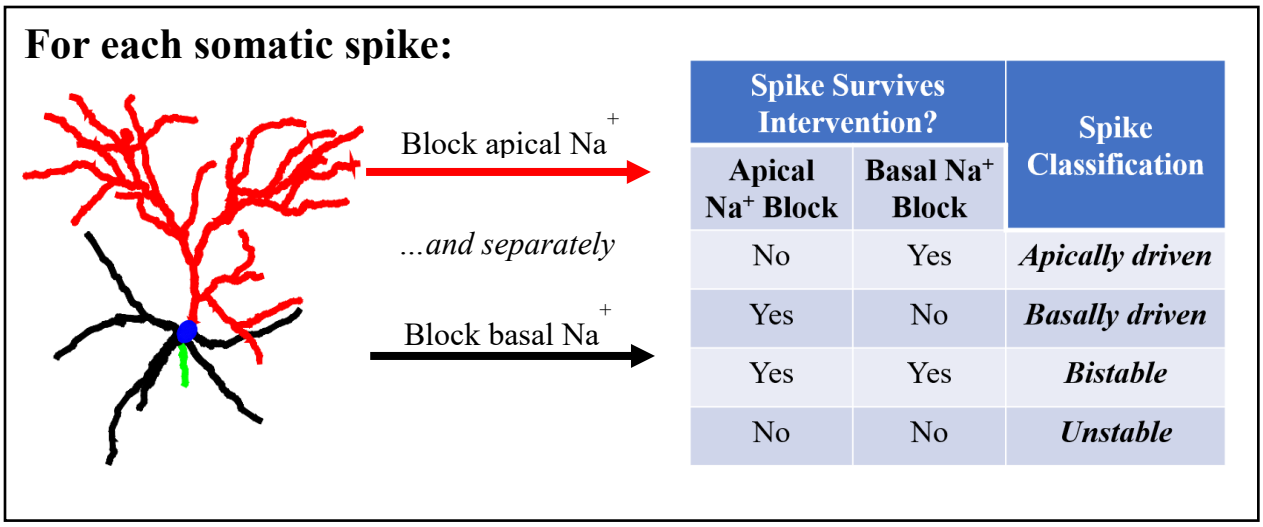

Spike Survives Intervention?		Spike Classification
Apical Na^+ Block	Basal Na^+ Block	
No	Yes	***Apically driven***
Yes	No	***Basally driven***
Yes	Yes	***Bistable***
No	No	***Unstable***

Figure 8. Causal Intervention protocol description, alongside the types of spikes that can be thus inferred.

408

409 **2.5. Causal Classification**

410 We executed regular stimulation experiments with somatic sodium blockage, using both the even
411 and biased distribution models while also applying a weight factor on the sodium channel
412 conductance of the basal tree dendrites. For the even model, the weight factor corresponded to
413 an effective increase of basal sodium conductance by 0 to 20%, in steps of 1%. For the biased
414 model, the weight factor had the opposite effect, reducing basal sodium conductance by 0 to
415 20% in steps of 1%. We then used two approaches in an attempt to classify causal triggers of
416 neuronal output:

417 (1) Simple Classification Algorithm: This algorithm was created to coarsely label the data obtained
418 by the experiments outlined above, in an attempt to later use the resulting dataset to train a
419 Machine Learning classifier that would then accurately classify the causal triggers of neuronal
420 output. It takes as input 51 voltage traces generated by the simulated neuron. Afterwards, it
421 evaluates the percentage of apical and basal dendrites that exhibited a dendritic sodium spike

shortly before the soma (3 ms) and compares the two percentages. Whichever dendritic tree had a greater percentage would be labeled as the causal trigger. The calculated apical and basal percentages corresponded to the posterior probability of any given apical or basal dendrite exhibiting a dendritic spike before the corresponding somatic spike. Using this reasoning, we designed an evaluation function that took as input the output of the simple classifier. It would then calculate the logarithm of the quotient of the returned percentages – henceforth "log-likelihood". Comparing it to a significance level α, it determined whether the absolute value of the log-likelihood of each sample in the data was sufficient to warrant labeling it with high confidence. In our case, in order to completely exclude ambiguous data (i.e. somatic spikes that were temporally preceded by a mix of apical and basal dendritic spikes), we chose to exclude all samples for which the absolute value of the log-likelihood was finite. This excluded all ambiguous data while still leaving a sufficiently large number of samples for training. In the end, the evaluation function returned two subsets of data: trivially classifiable data, labeled with high confidence, and trivially unclassifiable data, labeled with low confidence. The former consists of somatic spikes that were preceded by dendritic spikes from only one tree (*single tree spiking data*), while the latter category contains somatic spikes that were preceded by a mix of dendritic spikes from both trees (*coincident spiking data*).

(2) Causal Classification Algorithm: Machine Learning algorithms were used in conjunction with the high-confidence data labeled using the Simple Classification Algorithm and verified via causal intervention experiments (n = 1729) as well as with the low-confidence data labeled solely using the results of causal intervention experiments (n = 659). To simplify classification by rendering it binary, samples of unstable and bistable spikes (n = 69 and n = 121, respectively) were excluded.

In order to ensure selection of the best model possible, we used a 10-times Repeated, Stratified, Nested 10-Fold Cross-Validation (RS-NCV) protocol (Tsamardinos, Rakhshani, & Lagani, 2015) that allows us to test the stability of the chosen model, as well as account for non-uniform class priors. Feature extraction and transformation was required to reduce data dimensionality, as the original dimensions of the dataset used were more than $5 * 10^6$. For every recorded spike, we extracted the timing difference of each dendritic spike to the somatic spike (50 features), the absolute timing of each dendritic spike (50 features), the max depolarization amplitude (51 features) and the total area under the compartment voltage trace (51 features). The soma was not included in the first two features, as the corresponding values bore no information (zero variance). We scaled the extracted features using min-max normalization applied on all members of each feature category on a per-sample basis, to avoid carrying information across samples through normalization. Classification accuracy and area under the Receiver Operating Characteristic curve (auROC) were selected as model performance metrics. To ensure an adequate level of performance, the values of these metrics for the selected model were compared to those of a trivial classifier, which always classifies all samples to the most populated class (apically-driven). The protocol evaluated and selected among the following classifiers: Random Forests (RF), Naïve Bayes (NB), K-Nearest Neighbors (KNN), and Support Vector Machines (SVM). A narrow set of hyperparameters were provided for each classifier in an attempt to further optimize performance (Table 5). NB prior class probabilities were not selected for optimization, as attempts to use other probability distributions (i.e. uniform) drastically reduced the performance of the classifier. We avoided Artificial Neural Networks (ANN) because of the difficulties involved in using them in a Cross-Validation protocol (Tsamardinos et al., 2015).

RS-NCV was preferred over the Tibshirani and Tibshirani method (TT) (Tibshirani & Tibshirani, 2009) because of its tendency to underestimate the true performance of the model, giving it a more conservative nature (Tsamardinos et al., 2015), which is desirable in this type of analysis, where class labels are relatively uncertain.

Classifier	Hyperparameters	Possible Values
RF	Number of trees, Minimum leaf node size	51, [1,2,3,4,5]
NB	Prior class probabilities	Empirically calculated
KNN	Number of neighbors (K)	[1,2,3,4,5,6,7,8,9,10]
SVM	Kernel function	Linear, Polynomial, Gaussian

Table 5. Classifier models and hyperparameter sets used in the RS-NCV protocol.

2.6. Data Acquisition and Analysis

All simulations were performed on the High-Performance Computational Cluster at IMBB-FORTH, featuring 312 high-performance CPU cores and 1,150 GB of RAM, through the NEURON simulation environment (Hines & Carnevale, 2001). Data analysis was performed on MATLAB R2017a (*Mathworks Inc.*), using publicly available libraries as well as custom-made scripts and functions. These include:

(1) Dendritic Spike Detection: Presence of sodium spikes in data obtained from paired-pulse protocol experiments was verified via a simple spike detection algorithm that identified short-lived (around 1 ms) depolarizations exceeding a -20 mV threshold. Identification of NMDA spikes was implemented via an algorithm that located inflection points on the voltage trace. We exploited the characteristic shape of the NMDA spike and calculated the number of inflection points immediately after the second pulse. Excitatory Post-Synaptic Potentials (EPSPs) and sodium spikes only have one inflection point, as the exponential decay in membrane voltage continues up to the resting potential. As the NMDA spike does not conform to this description,

and instead exhibits a voltage plateau, 2 or more inflection points indicate an NMDA spike. Inflection points were discovered by taking the points where the second derivative of the voltage trace is zero. As our voltage measurements are not continuous, we are unlikely to encounter a measurement point for which the second derivative is exactly zero. Thus, we assume an inflection point exists at some point P_n if and only if $(P_{n-1} \cdot P_{n+1}) < 0$, where P_{n-1} and P_{n+1} refer to existing points in the second derivative of the voltage trace, immediately preceding and anteceding the theorized inflection point.

(2) Quantification of Dendritic Nonlinearities: Using data from iterative paired-pulse protocol experiments, the neuronal output signal is extracted. The selected signal is usually the maximum amplitude of depolarization, although the width of the excitatory post-synaptic potential (EPSP) at half amplitude is used in cases of sodium blockage. A linear input-output curve is then generated, extrapolating from the response of the dendrite to a single synaptic input. This curve is compared to the actual response of the dendrite to increasing input in an "Expected vs. Actual" plot, thus characterizing the dendrite as sub-linear or supra-linear. The non-linear behavior of the dendrites was quantified using the Nonlinearity Relative to Linear Extrapolation (NRLE) metric (Behabadi, Polsky, Jadi, Schiller, & Mel, 2012), which is defined as the maximum ratio of actual to expected neuronal output signal. An NRLE value of less than 1 denotes a sub-linear dendrite. NRLE of exactly 1 indicates a linear dendrite. NRLE values over 1 characterize a dendrite as supra-linear. The non-linearity threshold of all dendrites was also calculated, measured as the minimum number of synapses required to elicit the corresponding electrogenic event (sodium or NMDA dendritic spike).

3. Results

3.1. Model Validation

Constraining our model is a crucial step that ensures validity of results. This involves the replication of experimentally-derived response values using permissible alterations in free model parameters. In our case, we needed to ensure our model cell responded to stimuli of a preferred orientation with a firing rate of approximately 1.50 Hz, and to orthogonal stimuli with a firing rate of approximately 0.26 Hz, on average (Adesnik, Bruns, Taniguchi, Huang, & Scanziani, 2012). We used an orientation tuning validation protocol (see "*Methods*") to minimize model bias and ensure robustness of results. The ensuing orientation tuning curve resembles a Normal (Gaussian) distribution with a mean of 0° and a standard deviation of 30° (*Figure 9*).

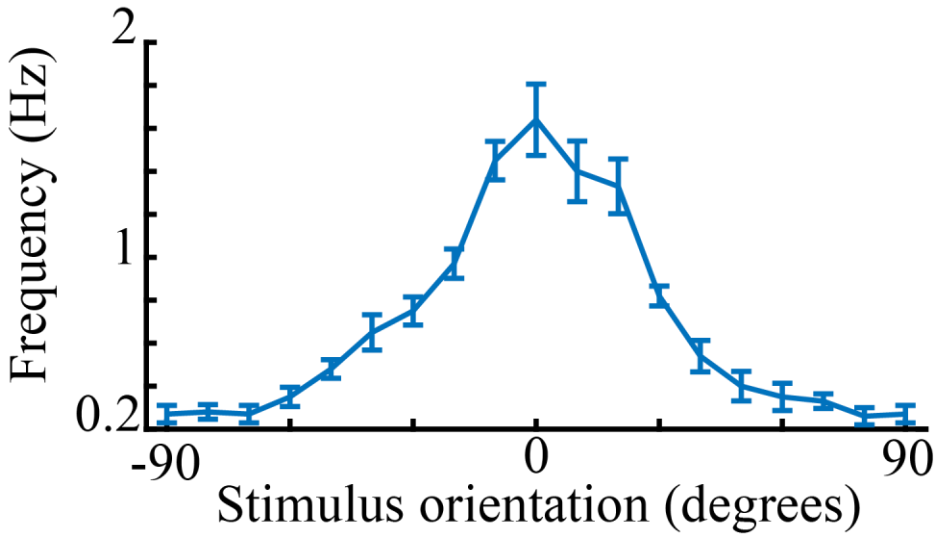

Figure 9. *Orientation tuning curve of the model neuron. Frequency is averaged across 10 repetitions for each orientation. Error bars: Standard error of the mean.*

3.2. Evaluation of Dendritic Non-Linearities

To ensure the highest possible degree of biophysical accuracy, we need to ensure that our dendrites exhibit electrogenic activity – dendritic spikes. Hence, we need to verify that both sodium and NMDA spikes are present in our model, thus rendering our dendrites capable of nonlinear integration of synaptic input. To that end, we used a paired-pulse stimulation protocol (see "*Methods*"). We find that both sodium and NMDA spikes are present in our model (*Figure 10*), but the synapse count required to elicit such events in each dendrite varies.

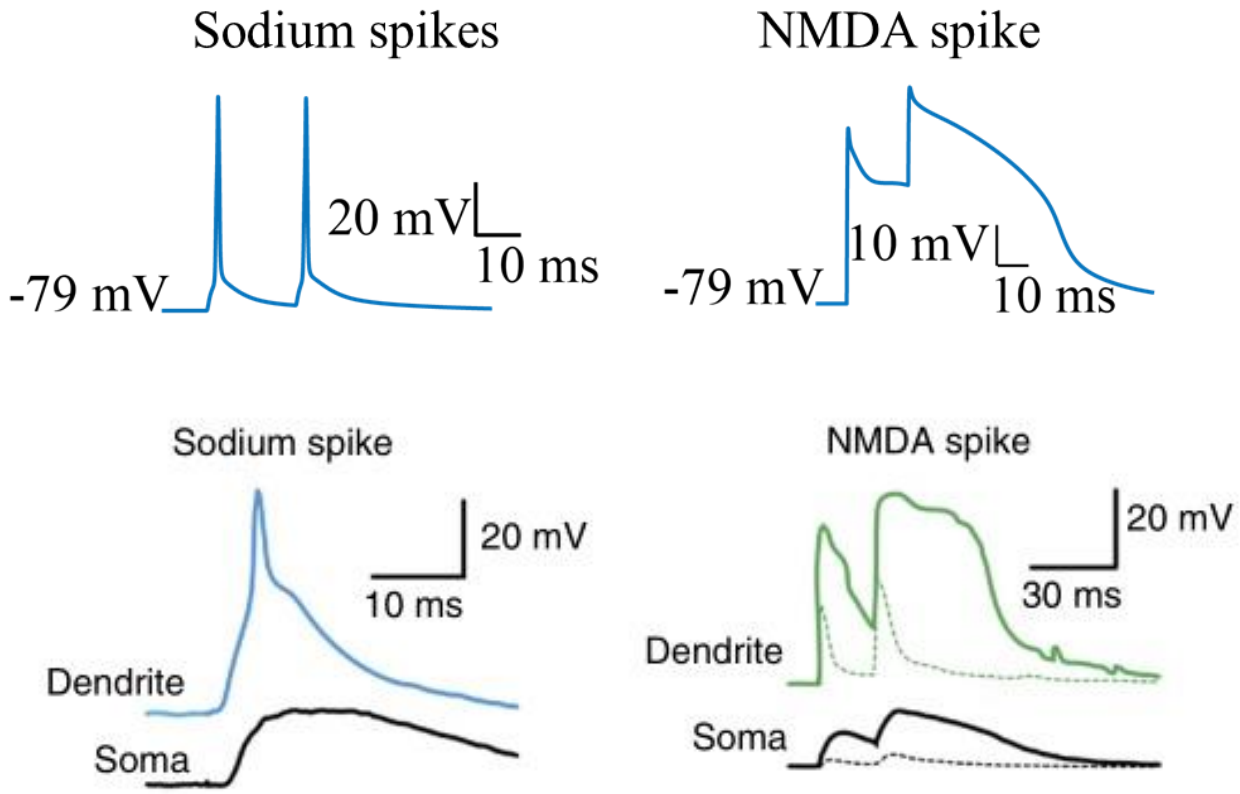

***Figure 10**. Sodium and NMDA spike samples as exhibited by the model neuron, recorded from apical dendrite 35 (A35) (top), compared to experimental recordings (bottom) reproduced from Stuart & Spruston, 2015.*

Next, we systematically examined all voltage traces from each dendrite, attempting to find the threshold for each type of dendritic nonlinearity as the minimum number of synapses

required to elicit a spiking event of the corresponding type – either a sodium spike or an NMDA spike (*Figure 11, Supplementary Figure 3*). Following that, we separated the dendrites into two broad categories: low-threshold (i.e. apical 35, basal 5) and high-threshold (i.e. apical 1, basal 0), with the cutoff arbitrarily set at 50 synapses, half of the examined maximum number.

In order to ascertain whether electrogenic activity allows the dendrites to perform nonlinear integration of synaptic input, we use an iterative paired pulse protocol (*see "Methods"*). To compare the output of each dendrite to the synaptic input received, we used the "Expected vs. Actual" plot (*see "Methods"*), with the maximum postsynaptic depolarization amplitude acting as the output signal (*Figure 12, Supplementary Figures 2-A and 2-B*). However, while the expected vs. actual plot clearly highlights dendritic non-linearities in the control case, it fails to give useful information when sodium channels are blocked. This can be amended by choosing a different type of output signal - in our case, we selected the duration of the EPSP following the second pulse, measured as the width at half maximum EPSP amplitude. By plotting this metric against the number of synapses, the NMDA-derived dendritic nonlinearities can be qualitatively observed (*Figure 13*).

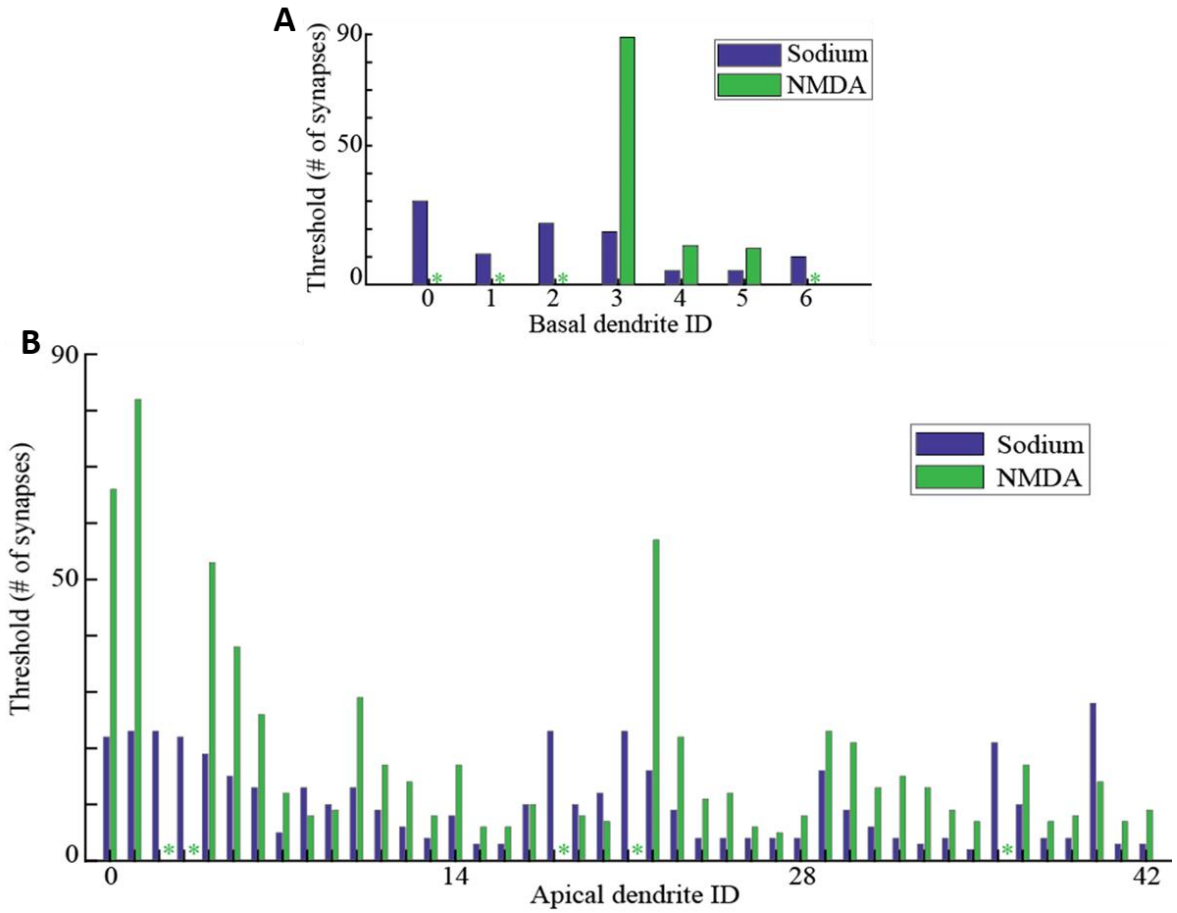

Figure 11. Number of synapses required to elicit sodium (blue) and NMDA (green) spikes for basal (A) and apical (B) dendrites. Asterisks denote a dendrite that did not exhibit that type of non-linearity for up to 100 simultaneously activated afferent synapses.

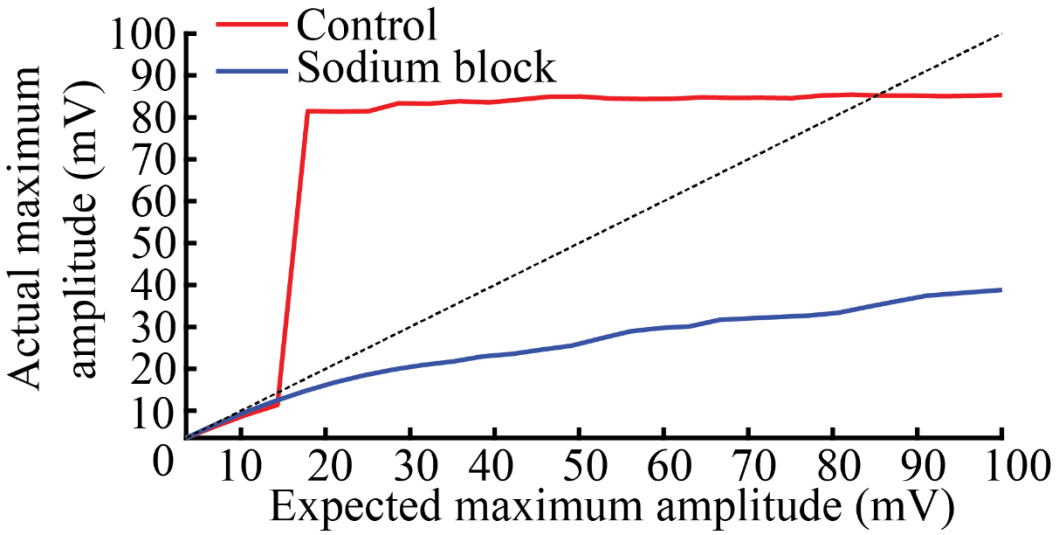

Figure 12. Expected vs Actual plot for basal dendrite 5 (B5). Supralinear behavior is lost if sodium channel conductance is nullified (sodium block).

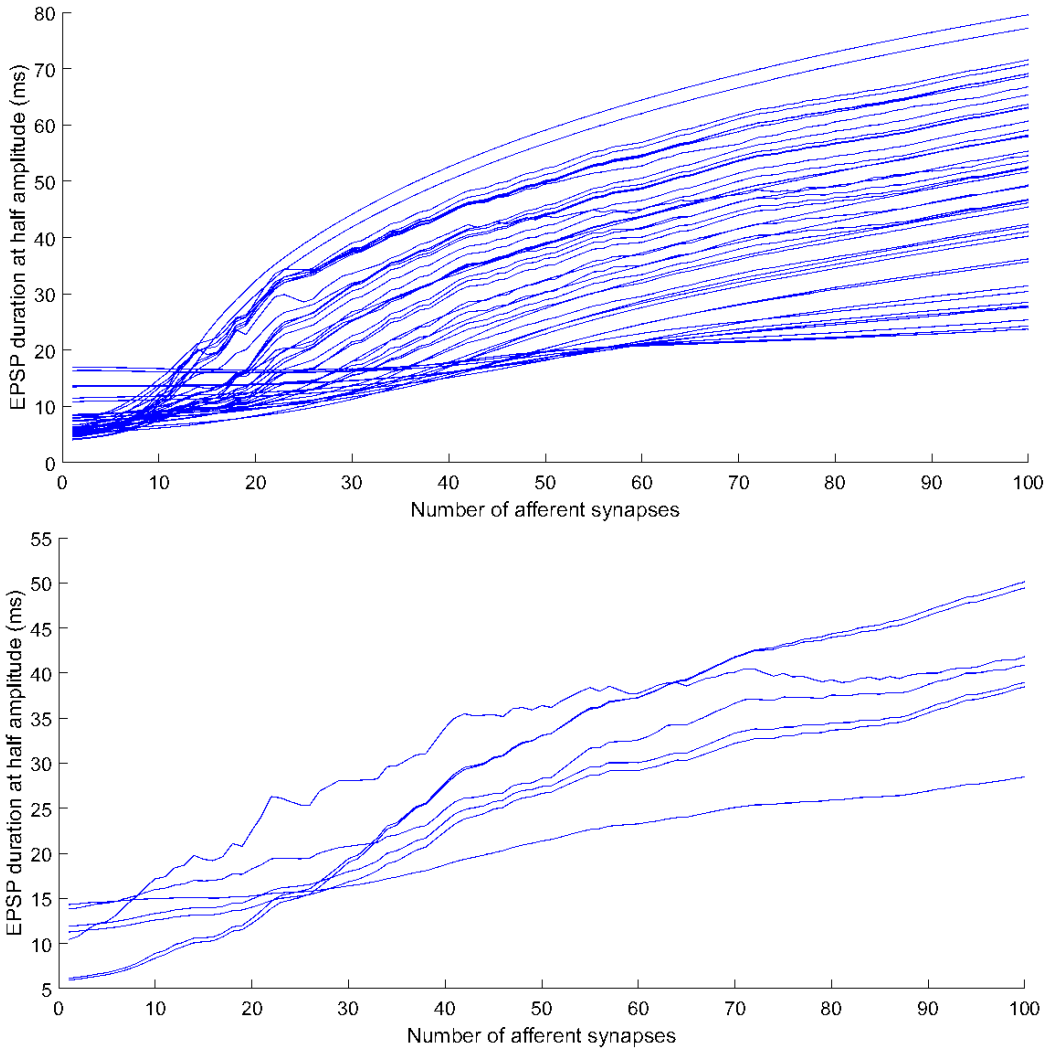

Figure 13*. NMDA-derived EPSP duration at half max amplitude allows visualization of NMDA non-linear behavior. Each trace represents the behavior of an individual apical (top) or basal (bottom) dendritic branch.*

Finally, we wanted to quantitatively characterize the nonlinear behavior of each dendrite. To do this, we used the Nonlinearity Relative to Linear Extrapolation (NRLE) metric (*see "Methods"*). We used the total area under the voltage trace as the output signal, measured both from the dendrite as well as from the soma. Nevertheless, many types of output signals can be used to derive an NRLE value for each dendrite. Results indicate that the dendrites of the model neuron all exhibit supra-linear input-output relations, as indicated by their NRLE values exceeding 1 (*Figure 14*).

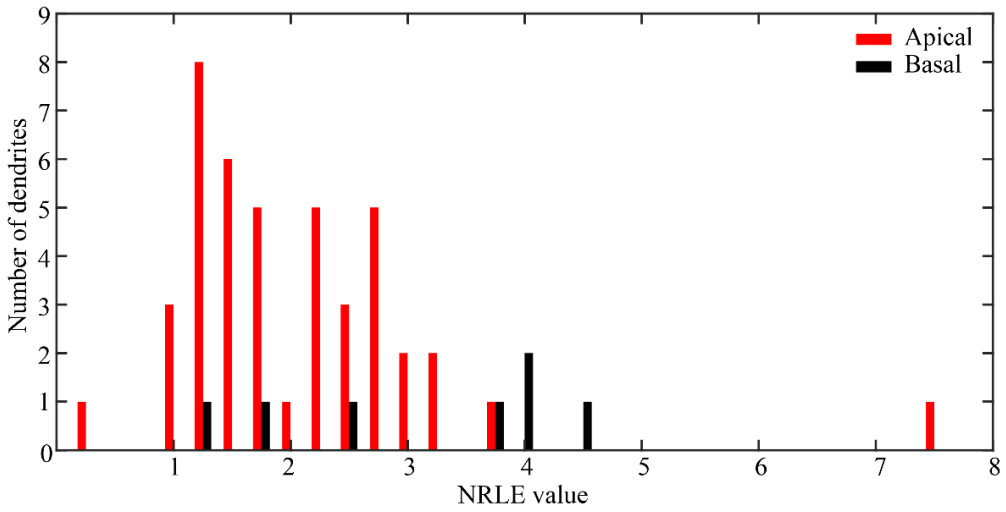

Figure 14. *Frequency histogram of NRLE values for apical (red) and basal (black) dendrites. Values of NRLE are generally lower for apical dendrites.*

3.3. Dendritic Contribution to Orientation Selectivity

Having thoroughly characterized dendritic nonlinearities in our model, we want to determine the roles of the apical and basal trees in shaping the orientation preference of the neuron. To achieve this, we perform a regular stimulation protocol simulation (see "*Methods*"), with sodium channels at the soma being blocked, preventing backpropagating action potentials from interfering with dendritic recordings. To investigate the role of the dendritic trees, we perform four such simulations, in which either the apical tree, basal tree, both trees *(negative control)* or none of the trees *(positive control)* have their sodium channels blocked.

We observe that somatic spiking activity seems to occur when both the apical and basal tree exhibit dendritic sodium spikes or depolarization within a brief time window prior to the somatic spike - a *bilateral input coincidence* (*Figure 15-A*). Interestingly, it appears that the apical dendrites drive the somatic spiking activity, as their activity temporally precedes both the somatic as well as the basal activity (*Figure 15-A, detail*).

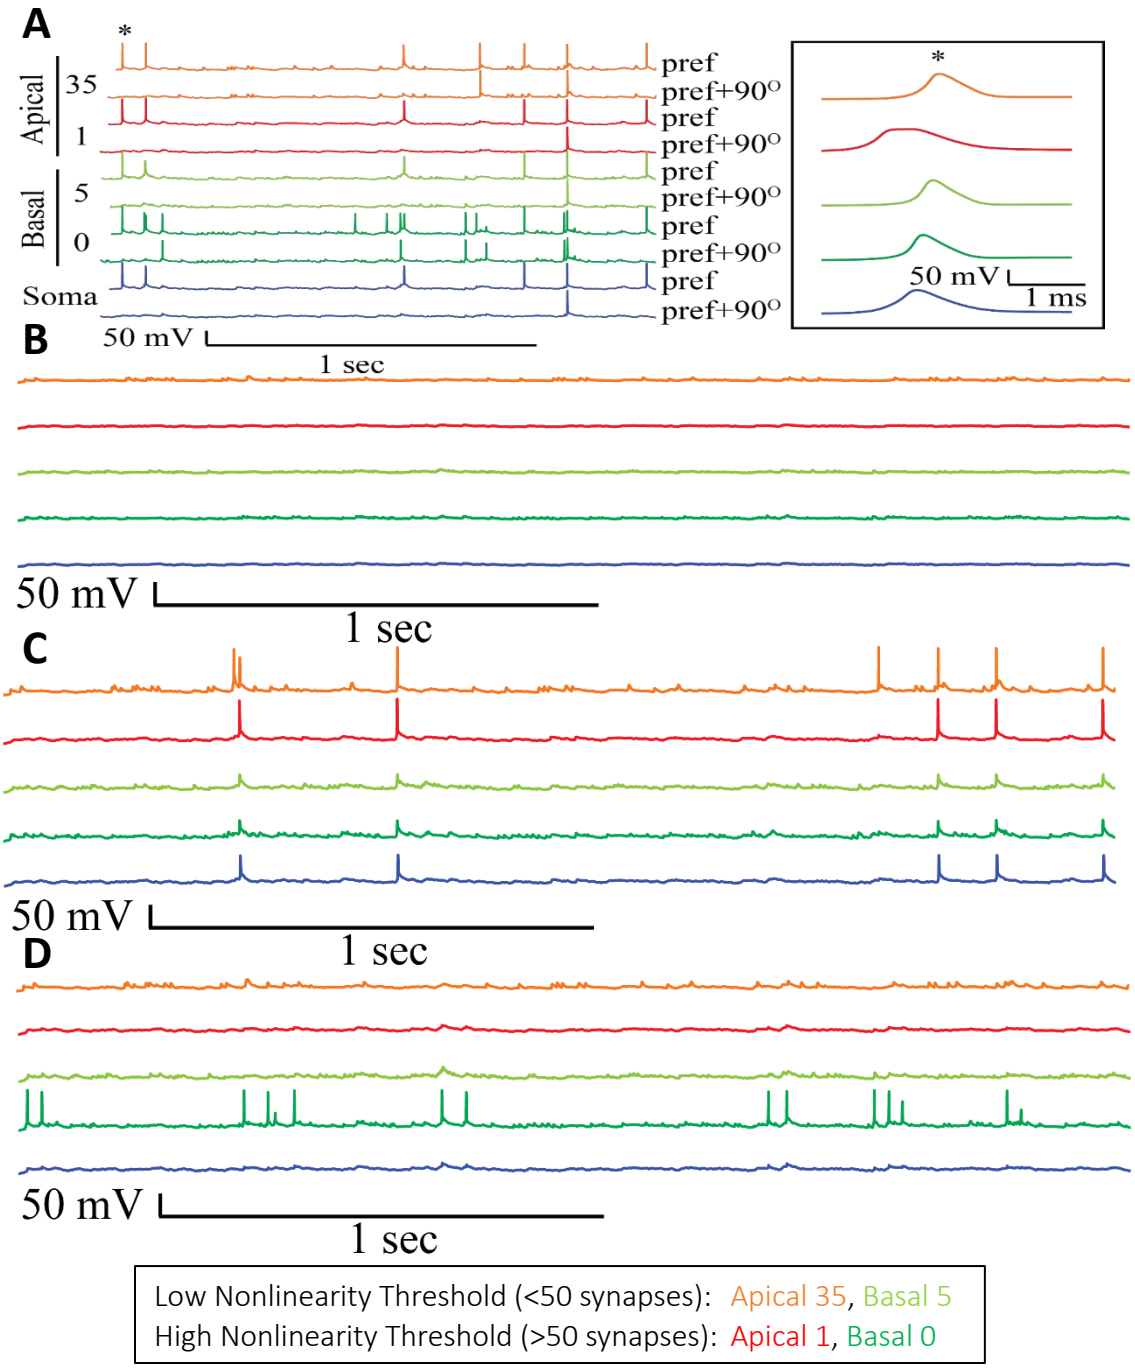

Figure 15. *Dendritic contribution to neuronal output. A: Somatic sodium channels are blocked. Coincident spiking (bilateral input coincidence) of the apical and basal dendrites appears to produce somatic spiking. Detail (asterisk): apical sodium spikes temporally precede basal and somatic spikes. B: All sodium channels blocked. No activity. C: Somatic and basal sodium channels blocked. Apical dendrites still exhibit spiking activity, and somatic spiking is reduced, but not extinct. D: Somatic and apical sodium channels blocked. Basal dendrites still exhibit spiking activity, but both apical and somatic spiking is completely extinct.*

584

585 Furthermore, if the apical tree sodium channels are blocked, the basal tree is completely unable
586 to elicit spiking activity in the soma, and basal spiking is reduced (*Figure 15-D*). Inversely, if the
587 basal tree is likewise treated, the apical tree is still capable of producing some spiking activity in
588 the soma, albeit greatly reduced. It also suffers a reduction of its own spiking activity (*Figure 15-*
589 *C)*. These results, however, are not sufficient evidence from which to draw a conclusion.
590
591 **3.4. Orientation Tuning Dominance**

592 To identify whether the effect of dendritic tree tuning on somatic orientation selectivity is biased
593 towards a specific dendritic tree, we used a disparity protocol and analyzed the resulting
594 orientation preference data (see "*Methods*"). When using the even model, neuronal tuning
595 favors the apical tree orientation preference. However, this trend is reversed in the biased model,
596 with the neuron exhibiting basal dominance (*Figure 14*).

597

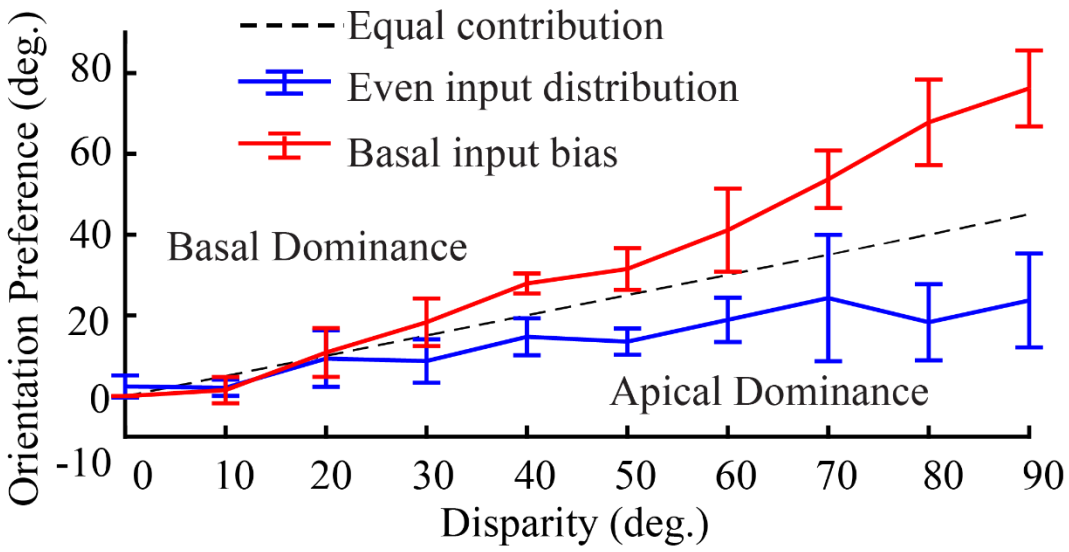

Figure 16*. Shift in neuronal orientation preference with increasing apical-basal tuning disparity. In the even input distribution model, orientation preference is dictated by the apical tree. Trend reverses if using the biased model, with the basal tree now dictating neuronal orientation preference. Error bars: Standard error of the mean (SEM).*

To differentiate between synaptic count and synaptic potency as determinants of dendritic dominance, we changed sodium conductance on the basal tree dendrites. When increasing sodium conductance in the basal tree by 20% using the even model, the curve shifts towards basal dominance. Inversely, a 20% decrease in basal sodium conductance in the biased model shifts the curve towards apical dominance (*Supplementary Figure 2*). This indicates that overall synaptic effectiveness, rather than synaptic count, is the deciding factor in shaping orientation tuning on a neuronal level.

3.5. Causal Interventions

In order to elucidate the exact contribution of the two dendritic arbors to somatic output, we needed to be able to clearly label each occurrence of an action potential (or suprathreshold somatic depolarization, in the case of somatic sodium blockage) in terms of causal instigation. As such, we used a causal classification protocol for simulations (see "*Methods*"), obtaining a large amount of spiking data from both the even and biased model configurations (n_{total} = 2388, n_{even} = 872, n_{bias} = 1516). This data was used in conjunction with a causal intervention protocol (see "*Methods*") in order to ascertain the most probable causal trigger of recorded somatic spikes. We find that the vast majority of somatic spikes are causally instigated by apical tree dendritic spikes (even: 79.12%; biased: 81.47%), which is surprising, considering that feedforward visual

input reaches mostly the basal tree, with few afferents reaching distal apical dendrites. However, non-zero percentages of basally-driven (even: 18.24%; biased: 7.52%), unstable (even: 0.92%; biased: 4.02%) and bistable (even: 1.72%; biased: 6.99%) spikes exist as well (*Figure 17*).

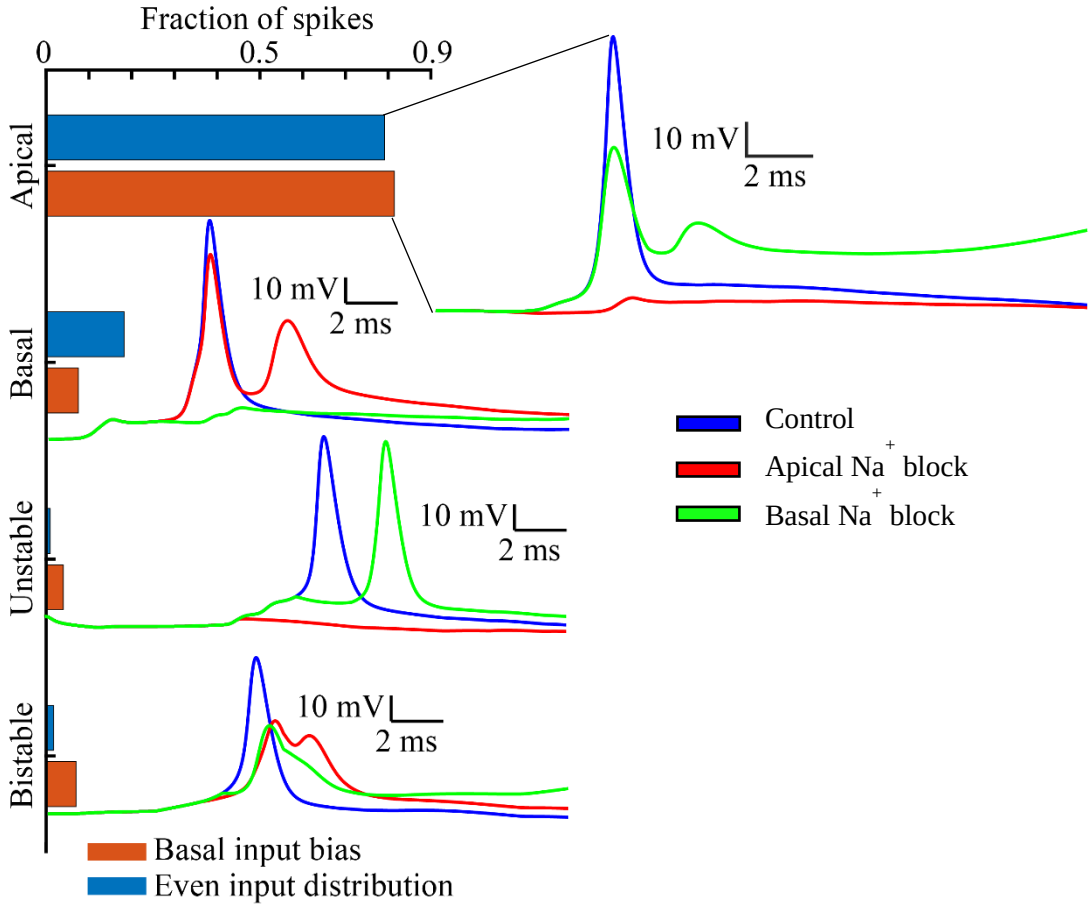

Figure 17. *Classification of all somatic spikes, and sample recordings demonstrating causal intervention results, for both the even and biased models. Most spikes are apically-driven in both models. Rightmost set of recordings represent an instance of an apically-driven spike: First, the blue trace is recorded. Then, when apical sodium conductance is nullified (red trace), there is no somatic activity. Last, if basal sodium conductance is nullified, the amplitude of the spike changes, but not its timing.*

We next used a simple classification algorithm (see "*Methods*") in an effort to broadly separate spiking data into high-confidence and low-confidence sets (*Figure 18*), characterized by single-tree spiking and coincident spiking (on both trees), respectively. These datasets characterized somatic activity solely in terms of being apically-driven or basally-driven, ignoring the remaining two categories. Then, we used the high-confidence data as a training set to train a Machine Learning model to classify causal triggers of neuronal activity in the low-confidence data, via an RS-NCV protocol (see "*Methods*"). Results indicated supra-trivial but still marginally improved performance (Accuracy: 0.8577 trivial, 0.7967 trained; auROC: 0.5 trivial, 0.6183 trained) (*Figure 19*). Given that the true labels were known through the causal intervention experiments, we tried using the low-confidence data as the training set instead, testing on the remaining data. This resulted in classification performance that was notably supra-trivial in terms of auROC (Accuracy: 0.8816 trivial, 0.7698 trained; auROC: 0.5 trivial, 0.9052 trained), despite inferior accuracy (*Figure 17*). As auROC is a better metric of true performance, this indicates a

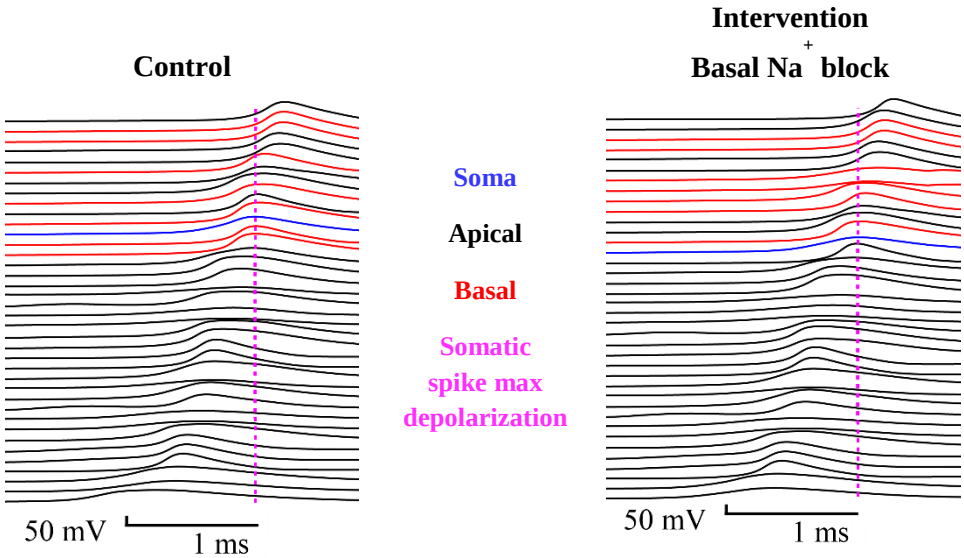

Figure 18. *Example sorted stack of traces from a low-confidence spike (left), and the sorted stack of traces after causal intervention that allows us to assign a label with high confidence regardless.*

good level of performance in this classification task. However, our trivial classifier that always labels each sample as apically-driven has superior accuracy. The reason the trivial classification has such a high accuracy is merely because apically-driven spikes are overrepresented, leading to an accuracy that is equal to the fraction of apically-driven spikes in the dataset. Thus, accuracy is a misleading metric in this case. In cases where all outcomes have identical prior probabilities (equal amounts of each in the samples), accuracy is more reliable.

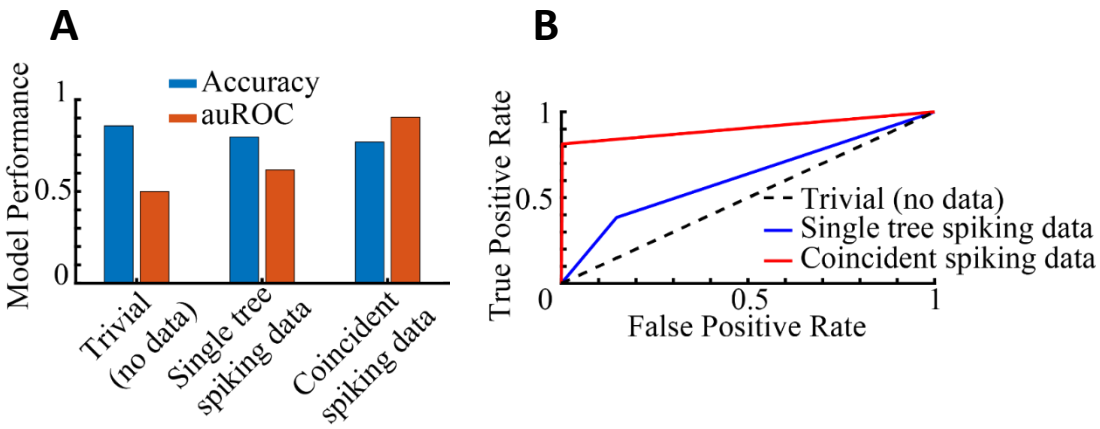

Figure 19*. Causal classification performance when training on single tree spiking data or coincident spiking data. A: accuracy and auROC for all cases. Trivial accuracy shown corresponds to testing on low-confidence data. B: ROC curves for all classifiers. Coincident spiking data offers the best performance as a training set.*

4. Discussion

Neuronal computation involves the spatiotemporal integration of disparate signals. It is well known that visual (feedforward) input reaches the basal tree of L2/3 V1 pyramidal neurons via afferent connections from L4 of V1 and the LGN. At the same time, attention- and prediction-related signals are received by the apical tree of these neurons, propagated from higher-order cortical areas such as V2, V3, LM, PFC and others (Coogan & Burkhalter, 1990; M. Larkum, 2013). The combination of these different types of input with the non-linear integration characteristics

shown to exist on a majority of dendrites (Hausser et al., 2000; Spruston, 2008) gives rise to a large spectrum of possible computations. In addition, it has been demonstrated that despite being fully capable of generating dendritic spikes, single dendrites are generally incapable of producing a significant somatic response (Smith et al., 2013). This would by necessity entail that multiple dendrites need to be activated in order to generate a spike at the soma, which in turn further expands the space of possible neuronal computations.

In this work, we have attempted to delineate the dendritic constituents of neuronal computation in L2/3 of V1 pyramidal cells using a detailed computational model. Our primary goal was to identify the relative contribution of the two dendritic arbors to somatic output. Results indicated that apical tree dendritic spikes instigate somatic spiking in the vast majority of cases (*Figure 15*). Thus, we can surmise that neuronal output is primarily determined by predictive and attentional inputs. Meanwhile, however, we also find that the orientation preference of a model cell is a function of the synaptic efficacy of its dendritic arbors – by increasing either the number of stimulus-driven synapses on the basal tree (basal bias model) or their weight, the neuron follows the orientation preference of the basal tree. In the even distribution case, or when apical synaptic weights are increased, the neuron follows the apical tree orientation preference instead. These results indicate that even though most somatic spikes are instigated through apical tree dendritic spikes, they are also heavily influenced by basal tree activity, to the point that the overall orientation preference of the cell can follow what the basal tree dictates.

Close examination of the performance results returned by our Machine Learning models reveals yet another interesting peculiarity. Comparing the model trained on high-confidence data

to the model trained on low-confidence data, there is a significant difference, with the latter being clearly superior in terms of auROC (see "*Results*"). This would mean that there is some sort of information content (Shannon information) in the low-confidence data that is missing from the high-confidence data. However, both datasets contain examples of apically- and basally-driven somatic spikes, caused by apical and basal dendritic spikes, respectively. The only difference in these two datasets is that the latter consists of examples in which dendrites from both trees fire in close temporal proximity – coincident spiking. As such, the missing information content must lie within this dendritic spiking coincidence. Given that the dendrites in question belong to different trees, this points towards the existence of *intra-tree dendritic cooperativity* – a synergistic effect between apical and basal dendrites that exhibit dendritic spikes in relative synchrony. This inference is further supported by the fact that there exist *unstable* somatic spikes that become extinct when either of their dendritic components is lost.

Finally, causal intervention results indicate that it is indeed possible to classify causal triggers of neuronal output in terms of dendritic origin. However, the task is complicated enough that our simple algorithm could not confidently classify any case exhibiting coincident dendritic spiking (see "*Results*"). Interestingly, given an adequate training set, we found that a Machine Learning model can be used to accurately discern the causal origins of neuronal activity. This opens up the possibility of using calcium imaging data from dendrites alongside such Machine Learning models to classify causal triggers of neuronal activity *in vivo*. Before that can happen, however, testing of this approach in simulated models using calcium signals rather than membrane voltage is needed.

Driven by our observations, we hypothesize that somatic output is determined through a

form of coincidence detection we call *bilateral input coincidence*. The basal tree receives visual feedforward input, representing the information therein as a series of hyperpolarizations, depolarizations, and occasional dendritic spikes. Thus, visual input defines a basal "backdrop" of depolarizations that represents visual information. At the same time, predictive and attentional signals from higher-order cortices reach the apical tree of the neuron, causing the generation of dendritic spikes that propagate to the soma and are temporally summed with any concurrent basal depolarization. In the event that visual input is non-existent, or of a non-preferred orientation, the depolarization is minimal or zero. Thus, the apical dendritic spike will not be significantly augmented through summation and will fail to produce a somatic response in the vast majority of cases. If there is visual input, however, especially of an orientation matching the preferred orientation of the basal tree, the "backdrop" will include multiple sub-threshold depolarizations, perhaps even dendritic spikes. The apical dendritic spikes will thus be temporally summed with these depolarizations and will be more likely to generate a somatic action potential. As such, even though most somatic spikes will be generated through an apical tree dendritic spike, the cases in which this is possible in the first place will be dictated by the backdrop of depolarizations provided by the basal tree.

Our hypothesis can also explain how L2/3 V1 pyramidal cells can assist in performing basic feature extraction from the visual input. It is obvious that the near-infinite amount of information contained in even the most rudimentary visual scene could never be fully represented using the finite space of the brain. As such, it is necessary for the visual system to extract salient features from the input and recombine them in such a way as to create an adequate representation of the true visual scene in the brain. This "simplification" of visual perception can be explained

718 through predictive coding (Rao & Ballard, 1999). Our hypothesis predicts that salient stimuli are
719 either attended to or predicted in advance, so that the corresponding signals that reach the apical
720 tree will "highlight" the appropriate parts of the "backdrop" generated by the visual input
721 reaching the basal tree. This would result in a neuron that is activated only when specific features
722 are present in its receptive field, rendering non-salient stimuli invisible. Such phenomena have
723 been observed experimentally, and it has been indeed hypothesized that they are caused by the
724 effects of attention, or lack thereof (Simons & Chabris, 1999).

725 It has long been proposed (de-Wit, Machilsen, & Putzeys, 2010; Petro & Muckli, 2016; Rao
726 & Ballard, 1999) that the visual cortex operates by relying heavily on predictive signals. In these
727 predictive models, a linear stimulus that is perceived at the level of V2 would generate feedback
728 signals from V2 to the corresponding pyramidal neurons in V1, causing them to generate action
729 potentials in response to that stimulus, even if they had not originally perceived it. This can
730 potentially explain phenomena such as the perception of a triangle in the negative space of the
731 Kanisza illusion (*Figure 19*), which would therefore operate by the perception of an existing line
732 segment propagating from V1 to higher-order areas, which would in turn activate V1 neurons
733 along the extrapolated direction of the line, thus creating the perception of a linear stimulus
734 where there is none. Simply put, the visual system "expects" a line to be present where none
735 exists, and thus one is perceived.

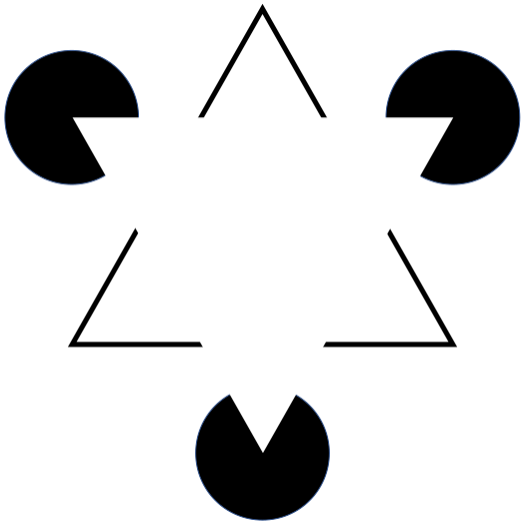

Figure 20. The Kanisza Triangle optical illusion. An inverted triangle can be perceived in the negative space in between the partial circles and chevron-shaped lines.

However, multiple questions still remain unanswered. First of all, the exact nature of the attentional and predictive signals received by the apical tree remains to be demystified. Separating one from the other, and clearly defining the origin and function of each, will greatly improve our understanding of the network-level computations of the visual system. Secondly, the formation of the visual "backdrop" as a result of basal tree depolarization merits further study. Encoding of visual information using mostly subthreshold depolarizations despite noisy inputs is an interesting problem, most likely resolved through the effects of intra-tree dendritic cooperativity, where apical spikes sharpen responses to true input rather than noise. The role of basal AMPA and NMDA receptors is also a field of possible study, as our manipulations when silencing a dendritic tree were limited to nullification of sodium channel conductances. As such, AMPA and NMDA receptors could still be activated. In fact, the large percentage of apically-driven spikes we find are most likely the result of intra-tree cooperativity between apical

dendritic spikes and basal AMPA-derived depolarizations. Finally, perhaps the most interesting unanswered question is whether these computations take place elsewhere in the cortex as well, be it in the visual system or not. The large amounts of information the brain must process necessitate the existence of a simplifying mechanism to render this intractable task possible. Predictive coding as a means of stimulus compression has already been proposed as a way to simplify visual perception (Rao & Ballard, 1999), and this might also be the case for other sensory or cognitive tasks. Regardless, further investigation is required in order to unravel the Gordian knot that is visual perception.

5. Acknowledgments

This work was supported by the ERC Starting Grant "dEMORY: Dissecting the role of dendrites in memory", ERC-2012-StG-311435.

Special thanks to Dr. Panayiota Poirazi, Dr. Sidiropoulou Kyriaki and Dr. Tsamardinos Ioannis for their participation in the three-member committee and invaluable advice.

6. References

1. Adesnik, H., Bruns, W., Taniguchi, H., Huang, Z. J., & Scanziani, M. (2012). A neural circuit for spatial summation in visual cortex. *Nature*, *490*(7419), 226–231. https://doi.org/10.1038/nature11526
2. Behabadi, B. F., Polsky, A., Jadi, M., Schiller, J., & Mel, B. W. (2012). Location-dependent excitatory synaptic interactions in pyramidal neuron dendrites. *PLoS Computational Biology*, *8*(7), 40–42. https://doi.org/10.1371/journal.pcbi.1002599
3. Chen, T.-W., Wardill, T. J., Sun, Y., Pulver, S. R., Renninger, S. L., Baohan, A., … Kim, D. S. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature*, *499*(7458), 295–300. https://doi.org/10.1038/nature12354
4. Cho, K.-H., Jang, H.-J., Lee, E.-H., Yoon, S. H., Hahn, S. J., Jo, Y.-H., … Rhie, D.-J. (2008). Differential Cholinergic Modulation of Ca2+ Transients Evoked by Backpropagating Action Potentials in Apical and Basal Dendrites of Cortical Pyramidal Neurons. *Journal of*

Neurophysiology, *99*(6), 2833–2843. https://doi.org/10.1152/jn.00063.2008
5. Coogan, T. A., & Burkhalter, A. (1990). Conserved patterns of cortico-cortical connections define areal hierarchy in rat visual cortex. *Experimental Brain Research*, *80*(1), 49–53. https://doi.org/10.1007/BF00228846
6. Cruz-Martín, A., El-Danaf, R. N., Osakada, F., Sriram, B., Dhande, O. S., Nguyen, P. L., … Huberman, A. D. (2014). A dedicated circuit links direction-selective retinal ganglion cells to the primary visual cortex. *Nature*, *507*(7492), 358–361. https://doi.org/10.1038/nature12989
7. de-Wit, L., Machilsen, B., & Putzeys, T. (2010). Predictive Coding and the Neural Response to Predictable Stimuli. *Journal of Neuroscience*, *30*(26), 8702–8703. https://doi.org/10.1523/JNEUROSCI.2248-10.2010
8. DeFelipe, J., & Fariñas, I. (1992). The pyramidal neuron of the cerebral cortex: morphological and chemical characteristics of the synaptic inputs. *Prog Neurobiol*, *39*(6), 563–607. https://doi.org/10.1016/0301-0082(92)90015-7
9. Edwards, G., Vetter, P., McGruer, F., Petro, L. S., & Muckli, L. (2017). Predictive feedback to V1 dynamically updates with sensory input. *Scientific Reports*, *7*(1), 1–12. https://doi.org/10.1038/s41598-017-16093-y
10. Gidon, A., & Segev, I. (2012). Principles Governing the Operation of Synaptic Inhibition in Dendrites. *Neuron*, *75*(2), 330–341. https://doi.org/10.1016/j.neuron.2012.05.015
11. Gilbert, C. D. (1977). Laminar differences in receptive field properties of cells in cat primary visual cortex. *The Journal of Physiology*, *268*(2), 391–421. Retrieved from http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1283670&tool=pmcentrez&rendertype=abstract
12. Haider, B., Häusser, M., & Carandini, M. (2013). Inhibition dominates sensory responses in the awake cortex. *Nature*, *493*(7430), 97–100. https://doi.org/10.1038/nature11665
13. Hausser, M., Spruston, N., & Stuart, G. J. (2000). Diversity and dynamics of dendritic signaling. *Science*, *290*(5492), 739–744. https://doi.org/10.1126/science.290.5492.739
14. Hines, M. L., & Carnevale, N. T. (2001). Neuron: A Tool for Neuroscientists. *The Neuroscientist*, *7*(2), 123–135. https://doi.org/10.1177/107385840100700207
15. Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, *160*(1), 106–154. https://doi.org/10.1113/jphysiol.1962.sp006837
16. Iacaruso, M. F., Gasler, I. T., & Hofer, S. B. (2017). Synaptic organization of visual space in primary visual cortex. *Nature*. https://doi.org/10.1038/nature23019
17. Jia, H., Rochefort, N. L., Chen, X., & Konnerth, A. (2010). Dendritic organization of sensory input to cortical neurons in vivo. *Nature*, *16*(3), 236–237. https://doi.org/10.1038/nature08947
18. Kuffler, S. W. (1953). Discharge Patterns and Functional of Mammalian Retina. *Journal of Neurophysiology*, (16), 37–68.
19. Larkum, M. (2013). A cellular mechanism for cortical associations: An organizing principle for the cerebral cortex. *Trends in Neurosciences*, *36*(3), 141–151. https://doi.org/10.1016/j.tins.2012.11.006
20. Larkum, M. E., Zhu, J. J., & Sakmann, B. (1999). A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature*, *398*(6725), 338–341. https://doi.org/10.1038/18686
21. Ohki, K., & Reid, R. C. (2007). Specificity and randomness in the visual cortex. *Current Opinion in Neurobiology*, *17*(4), 401–407. https://doi.org/10.1016/j.conb.2007.07.007
22. Palmer, L., Murayama, M., & Larkum, M. (2012). Inhibitory Regulation of Dendritic Activity in vivo. *Frontiers in Neural Circuits*, *6*(May), 1–10. https://doi.org/10.3389/fncir.2012.00026
23. Papoutsi, A., Park, J., Ash, R.T., Smirnakis, S.M., Poirazi, P., Modeling orientation preference in the apical and basal trees of L2/3 V1 neurons, *CNS* (2017), Antwerp, Belgium.

24. Petro, L. S., & Muckli, L. (2016). The brain's predictive prowess revealed in primary visual cortex. *Proceedings of the National Academy of Sciences*, *113*(5), 1124–1125. https://doi.org/10.1073/pnas.1523834113
25. Poirazi, P., Brannon, T., & Mel, B. W. (2003). Pyramidal neuron as two-layer neural network. *Neuron*, *37*(6), 989–999. https://doi.org/10.1016/S0896-6273(03)00149-1
26. Priebe, N. J. (2016). Mechanisms of Orientation Selectivity in the Primary Visual Cortex. *Annual Review of Vision Science*, *2*(1), 85–107. https://doi.org/10.1146/annurev-vision-111815-114456
27. Priebe, N. J., & Ferster, D. (2012). Mechanisms of Neuronal Computation in Mammalian Visual Cortex. *Neuron*, *75*(2), 194–208. https://doi.org/10.1016/j.neuron.2012.06.011
28. Raman, R., & Sarkar, S. (2016). Predictive coding: A possible explanation of filling-in at the blind spot. *PLoS ONE*, *11*(3), 1–17. https://doi.org/10.1371/journal.pone.0151194
29. Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87. https://doi.org/10.1038/4580
30. Roth, M. M., Dahmen, J. C., Muir, D. R., Imhof, F., Martini, F. J., & Hofer, S. B. (2015). Thalamic nuclei convey diverse contextual information to layer 1 of visual cortex. *Nature Neuroscience*, *19*(2), 299–307. https://doi.org/10.1038/nn.4197
31. Roth, M. M., Dahmen, J. C., Muir, D. R., Imhof, F., Martini, F. J., & Hofer, S. B. (2016). Thalamic nuclei convey diverse contextual information to layer 1 of visual cortex. *Nature Neuroscience*, *19*(2), 299–307. https://doi.org/10.1038/nn.4197
32. Salin, P.-A., & Bullier, J. (1995). Corticocortical connections in the visual system: structure and function. *Physiological Reviews*, *75*(1), 107–154. https://doi.org/10.1152/physrev.1995.75.1.107
33. Scholl, B., Tan, A. Y. Y., Corey, J., & Priebe, N. J. (2013). Emergence of Orientation Selectivity in the Mammalian Visual Pathway. *Journal of Neuroscience*, *33*(26), 10616–10624. https://doi.org/10.1523/JNEUROSCI.0404-13.2013
34. Schuz, A., & Palm, G. (1989). Density of Neurons and Synapses. *Journal of Comparative Neurology*, *455*, 442–455.
35. Shai, A. S., Anastassiou, C. A., Larkum, M. E., & Koch, C. (2015). Physiology of Layer 5 Pyramidal Neurons in Mouse Primary Visual Cortex: Coincidence Detection through Bursting. *PLoS Computational Biology*, *11*(3), 1–19. https://doi.org/10.1371/journal.pcbi.1004090
36. Silver, R. A. (2010). Neuronal arithmetic. *Nature Reviews. Neuroscience*, *11*(7), 474–489. https://doi.org/10.1038/nrn2864
37. Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception*, *28*(28), 1059–1074.
38. Smith, S. L., Smith, I. T., Branco, T., & Häusser, M. (2013). Dendritic spikes enhance stimulus selectivity in cortical neurons in vivo. *Nature*, *503*(7474), 115–120. https://doi.org/10.1038/nature12600
39. Spratling, M. W. (2012). Predictive coding as a model of the V1 saliency map hypothesis. *Neural Networks*, *26*, 7–28. https://doi.org/10.1016/j.neunet.2011.10.002
40. Spruston, N. (2008). Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews Neuroscience*, *9*(3), 206–221. https://doi.org/10.1038/nrn2286
41. Stuart, G. J., & Sakmann, B. (1994). Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature*, *367*(6458), 69–72. https://doi.org/10.1038/367069a0
42. Stuart, G. J., & Spruston, N. (2015). Dendritic integration: 60 years of progress. *Nature Neuroscience*, *18*(12), 1713–1721. https://doi.org/10.1038/nn.4157
43. Stuart, G., Spruston, N., Sakmann, B., & Hausser, M. (1997). Action potential initiation and back propagation in neurons of the mammalian central nervous system. *Trends in Neurosciences*,

20(3), 125–131. https://doi.org/10.1016/S0166-2236(96)10075-8
44. Sun, W., Tan, Z., Mensh, B. D., & Ji, N. (2016). Thalamus provides layer 4 of primary visual cortex with orientation- and direction-tuned inputs. *Nature Neuroscience*, *19*(2), 308–315. https://doi.org/10.1038/nn.4196
45. Thomas Jessell, Siegelbaum, S., & Hudspeth, A. J. (2000). *Principles of neural science* (Vol. 4, pp. 1227-1246). E. R. Kandel, J. H. Schwartz, & T. M. Jessell (Eds.). New York: McGraw-hill.
46. Tibshirani, R. J., & Tibshirani, R. (2009). A bias correction for the minimum error rate in cross-validation. *Annals of Applied Statistics*, *3*(2), 822–829. https://doi.org/10.1214/08-AOAS224
47. Tsamardinos, I., Rakhshani, A., & Lagani, V. (2015). Performance-Estimation Properties of Cross-Validation-Based Protocols with Simultaneous Hyper-Parameter Optimization. *International Journal on Artificial Intelligence Tools*, *24*(05), 1540023. https://doi.org/10.1142/S0218213015400230

7. Supplementary Information

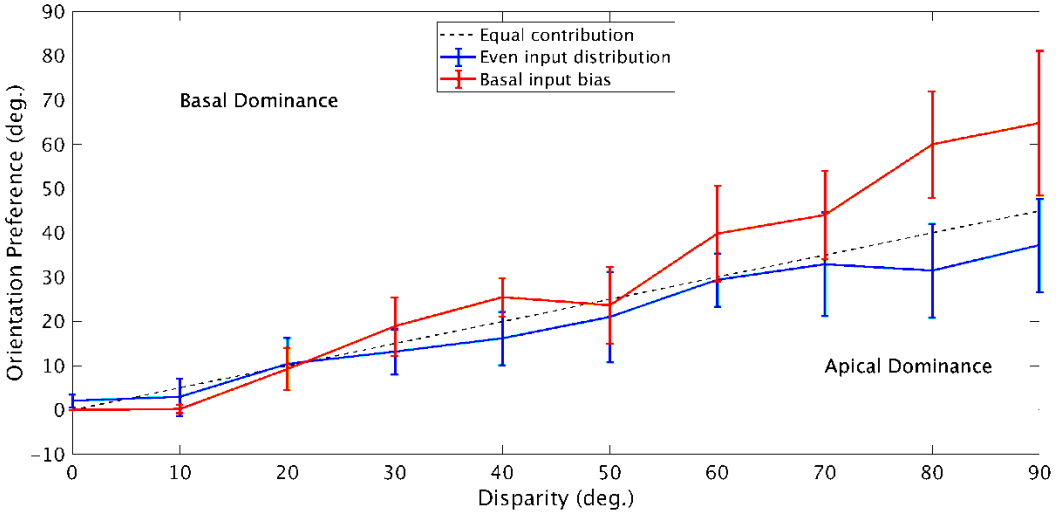

Figure S1. *Orientation tuning dominance curves shift towards the opposing side when increasing sodium conductance of the least favored dendritic tree by 20%. Error bars: Standard error of the mean (SEM).*

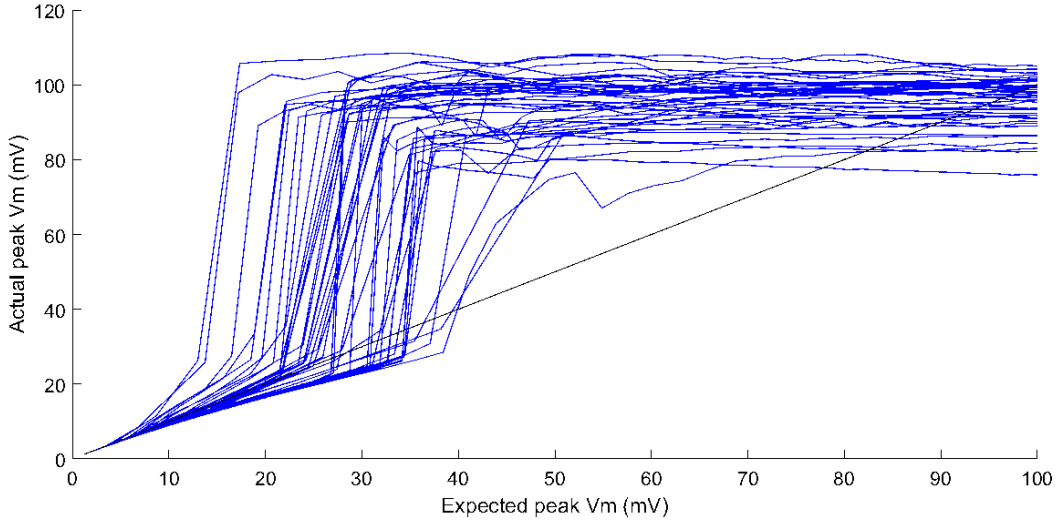

Figure S2-A. *Expected vs. Actual plot for all apical dendrites. All dendrites exhibit supralinear behaviour due to the generation of dendritic sodium spikes.*

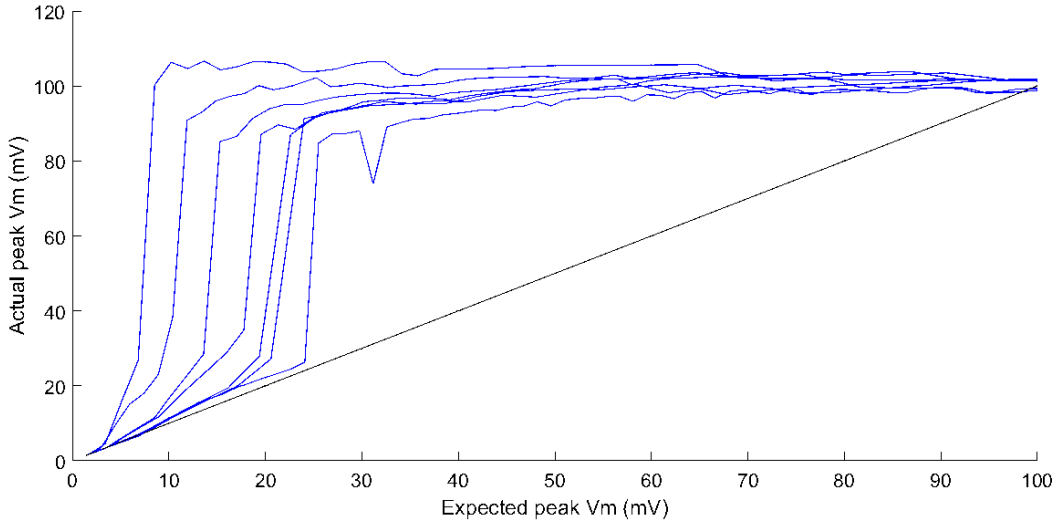

Figure S2-B. *Expected vs. Actual plot for all basal dendrites. All dendrites exhibit supralinear behaviour due to the generation of dendritic sodium spikes.*

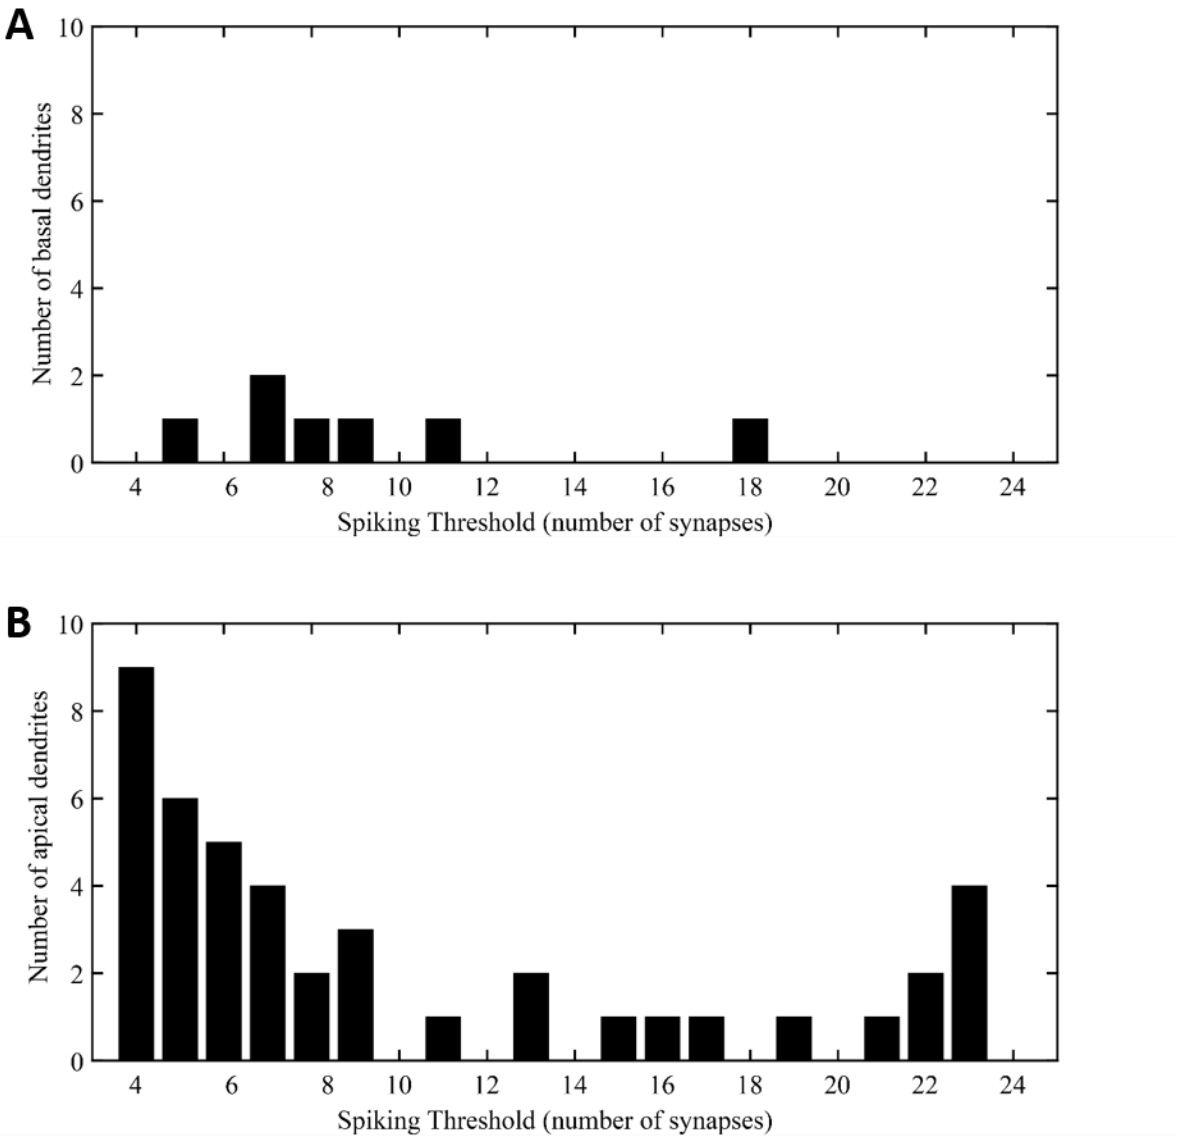

908

909 ***Figure S3****. Sodium spike threshold frequency histograms for the basal (A) and apical (B) trees. Sodium spike thresholds*
910 *are generally lower for basal tree dendrites.*

911