

Passive nonlinear dendritic interactions as a computational resource in spiking neural networks

Andreas Stöckel¹, Chris Eliasmith¹

¹ Centre for Theoretical Neuroscience, University of Waterloo, Waterloo, Ontario, Canada

Keywords: passive dendritic computation, nonlinear synapses, Dale’s principle, Neural Engineering Framework

Abstract

Nonlinear interactions in the dendritic tree play a key role in neural computation. Nevertheless, modeling frameworks aimed at the construction of large-scale, functional spiking neural networks, such as the Neural Engineering Framework, tend to assume a linear superposition of post-synaptic currents. In this paper, we present a series of extensions to the Neural Engineering Framework that facilitate the construction of networks incorporating Dale’s principle and nonlinear conductance-based synapses. We apply these extensions to a two-compartment LIF neuron that can be seen as a simple model of

passive dendritic computation. We show that it is possible to incorporate neuron models with input-dependent nonlinearities into the Neural Engineering Framework without compromising high-level function and that nonlinear post-synaptic currents can be systematically exploited to compute a wide variety of multivariate, bandlimited functions, including the Euclidean norm, controlled shunting, and non-negative multiplication. By avoiding an additional source of spike noise, the function-approximation accuracy of a single layer of two-compartment LIF neurons is on a par with or even surpasses that of two-layer spiking neural networks up to a certain target function bandwidth.

1 Introduction

A central challenge in theoretical neuroscience is to describe how biological mechanisms ultimately give rise to behavior. One way to approach this challenge is to build models of neurobiological systems that generate the behavior of interest to a researcher. Since constructing models that span multiple levels of abstraction is typically difficult, theoretical neuroscientists are working on methods that facilitate mapping high-level behaviour onto neural mechanisms. Such modeling frameworks include the Neural Engineering Framework (NEF) (Eliasmith and Anderson, 2003; Eliasmith, 2013), Efficient, Balanced Spiking Networks (EBN) (Boerlin and Denève, 2011; Boerlin et al., 2013), and FORCE (Sussillo and Abbott, 2009; Nicola and Clopath, 2017). Generally speaking, these approaches describe how to translate dynamical systems—corresponding to some hypothesized behavioral model—into an idealized spiking neural network that adheres to the desired neurophysiological constraints, for example neural tuning, fir-

ing rate distributions, and population-level connectivity (Komer and Eliasmith, 2016; Nicola and Clopath, 2017). This mechanistic grounding facilitates model validation by enabling a direct comparison of simulation results and empirical data (e.g., Stewart et al., 2012; Bekolay et al., 2014a; Duggins et al., 2017; Voelker and Eliasmith, 2018; Gosmann and Eliasmith, 2020).

The frameworks mentioned above primarily rely on two biophysical phenomena as computational primitives: synaptic filtering and the nonlinear relationship between somatic input currents and the neural response. Somatic response models range from leaky integrate-and-fire (LIF) to Hodgkin-Huxley type dynamics (Schwemmer et al., 2015; Eliasmith et al., 2016; Duggins, 2017). Crucially however, these approaches typically assume that post-synaptic currents are a linear superposition of filtered pre-synaptic events. Nonlinear interactions between input channels as they may occur when modeling conductance-based synapses or dendritic structures are typically ignored.

While some research exists that explores the effect of nonlinear post-synaptic currents within these frameworks (Bobier et al., 2014; Thalmeier et al., 2016; Alemi et al., 2018), these nonlinearities are seldom systematically exploited. Yet, empirical and theoretical work suggests that active and passive nonlinear effects within the dendritic tree—and not only the soma—are at least partially responsible for the complex responses observed in some biological neurons, including cortical pyramidal cells (Mel, 1994; Koch, 1999; Polsky et al., 2004). London and Häusser (2005) argue that in addition to voltage-gated ionic currents, fundamental passive effects such as shunting inhibition are worth being investigated as computational resources.

Put differently, current functional modeling frameworks only consider a subset of

the computational resources available in individual neurons and thus underestimate their computational power. Modelers wishing to multiply two signals might for example be forced to introduce an additional layer of neurons, although—in biology—the interplay between excitation and inhibition within the dendrites of a single neuron layer could have the same effect (Koch, 1999). The goal of this paper is to present mathematically tractable methods that allow researchers to take nonlinear post-synaptic currents into account. We demonstrate, as demanded by London and Häusser (2005), that the interactions between passive conductance-based input channels within a single dendritic compartment provide significant computational advantages over standard LIF neurons even within a noisy spiking neural network with low firing rates and small neuron counts.

Specifically, we extend the NEF towards systematically exploiting nonlinear post-synaptic currents. The NEF has been applied to various research areas, including low-level modeling of neurobiological systems (Kuo and Eliasmith, 2005; Tripp, 2009; Bobier et al., 2014), and studying large-scale models of cognitive systems grounded in biology (Eliasmith et al., 2012; Eliasmith, 2013; Eliasmith et al., 2016). A software implementation of the NEF is part of the neural simulation package Nengo (Bekolay et al., 2014b) and has been used as a neural compiler targeting analog and digital neuromorphic hardware (Choudhary et al., 2012; Mundy et al., 2015; Berzish et al., 2016; Blouw et al., 2018; Neckar et al., 2019).

The main contributions of this paper are as follows. First, we present a series of extensions to the NEF that improve its compatibility with more biologically detailed neuron models. We describe how to enforce nonnegative weights in conjunction with Dale’s principle and extend the NEF towards nonlinear post-synaptic currents, thereby

lifting some long-standing limitations of the NEF. Second, we derive a post-synaptic current model for a two-compartment leaky integrate-and-fire (LIF) neuron that can be interpreted as a simple dendritic nonlinearity. Third, we demonstrate that a single layer of two-compartment LIF neurons can compute a wide variety of functions with an error smaller than or on a par with the accuracy achieved by a comparable two-layer spiking neural network, as long as the target function does not surpass a certain bandwidth.

2 Materials and Methods

We begin with a review of relevant portions of the Neural Engineering Framework (NEF) followed by four novel extensions: decoding in current space, nonnegative weights, equality relaxation for subthreshold currents, and the input-dependent nonlinearity model H . These extensions facilitate incorporating biological detail into NEF models, including nonlinear, dendritic post-synaptic currents. We close with a description of the two-compartment LIF neuron, and the derivation of a corresponding nonlinearity H .¹

2.1 The Neural Engineering Framework (NEF)

At its core, the NEF describes three principles that govern the construction of models of neurobiological systems (Eliasmith and Anderson, 2003). These principles apply established concepts from artificial neural networks to time-continuous spiking neural networks in a way that facilitates the integration of neurophysiological constraints.

¹The methods presented in this section have been implemented as a software library extending the Nengo simulation package, see <https://github.com/astoeckel/nengo-bio>. A standalone version used to conduct the experiments in this paper can be found in the supplemental material.

The first principle postulates that populations of neurons represent values in a relatively low-dimensional manifold within their high-dimensional activity space via nonlinear encoding and linear decoding. According to the second principle, network connectivity defines transformations as mathematical functions of the represented values. Finally, the third principle states that recurrent connections approximate arbitrary dynamical systems in which the represented values are state variables. Although we will not discuss recurrent networks in this paper, it should be noted that the methods presented here are fully compatible with the third principle and as such can be used to implement dynamical systems.

Representation

A fundamental assumption of the NEF is that populations of spiking neurons represent d -dimensional vectors $\vec{x} \in \mathbb{R}^d$ by means of nonlinear encoding and linear decoding. In essence, each neuron population is treated as a single hidden layer neural network with linear input and output units. Similar schemes have been proposed since the late 1960s, an early example being Marr’s model of pattern extraction in the cerebellar granule layer (Marr, 1969).

For the encoding process, each neuron i in a population of size n is assigned a tuning-curve $a_i(\vec{x})$ that maps any represented value $\vec{x}$ onto a corresponding activity (Figs 1A and 1B). Mathematically,

$$a_i(\vec{x}) = G[J_i(\langle \vec{x}, \vec{e}_i \rangle)] = G[\alpha_i \langle \vec{x}, \vec{e}_i \rangle + \beta_i], \quad (1)$$

where the encoding vector $\vec{e}_i$ projects the input $\vec{x}$ onto a scalar ξ , which in turn is translated by $J_i(\xi)$ into a somatic current. While the specific current translation function

$J_i(\xi)$ depends on modeling assumptions, it is often defined as a first-order polynomial parametrized by a gain α_i and a bias current β_i . When building models, encoders $\vec{e}_i$, biases β_i , and gains α_i , must be selected in such a way that the population exhibits diverse neural tuning, while at the same time staying within modeling constraints such as maximum firing rates for the range of represented $\vec{x}$. The nonlinear neuron response $G[J]$ defines the spiking neuron model used in the network. For a wide variety of neuron models, this can be characterized by a rate approximation that maps the somatic input current J onto an average firing rate.

Critically, this mapping does not suggest the adoption of a “rate code” by the NEF, but rather is a convenience for solving the synaptic weight optimization problem. That is, eq. (1) is merely normative: it defines the average activity that a modeler *expects* individual neurons to have when the population represents a certain $\vec{x}$. This normative constraint can be dropped entirely, and all optimization done in the spiking domain within the NEF, but the computational costs are significantly higher (MacNeil and Eliasmith, 2011; Eliasmith, 2013). In this paper, we purely focus on offline optimization of the synaptic weights. Hence, the methods we present should in no way be interpreted as a theory of learning within nervous systems, but purely as a method for constructing models of mature, already trained, neural systems.

Notice that for the purpose of modelling neurobiological systems, all functions and parameters listed above can be hand-tuned by the modeler to match neurophysiological data. For example, as explored in Eliasmith and Anderson (2003), two-dimensional encoding vectors $\vec{e}_i$ can be used to reproduce a wide variety of localised tuning curves found in the neuroscience literature. Since NEF models, in contrast to classical multi-




Figure 1: Representation and transformation in the Neural Engineering Framework. Example for a population of $n = 10$ neurons. (A) Randomly selected linear (affine) current translation functions $J_i(\xi) = \alpha_i \xi + \beta_i$ map the encoded value $\xi = \langle \vec{x}, \vec{e}_i \rangle$ onto an input current J_i . The dotted line corresponds to the spike threshold J_{th} . (B) Tuning curves for each neuron in the population after applying the somatic nonlinearity $G[J]$ (eq. 14). (C) Reconstructing the represented value from neuron activity by means of linear decoding. (D) Approximating the function $f(x) = 2x^2 - 1$ with a different linear decoding. Dashed lines correspond to the mathematical function being approximated.

layer neural networks, are typically not globally optimized, tuning curves are maintained throughout the weight optimization process and neural activities can be interpreted at any point in the simulation. Global stochastic gradient descent can be applied to parts of the network (Hunsberger and Eliasmith, 2015) at the cost of negatively impacting the interpretability and biologically informed tuning of individual neuron populations.

Complementary to the encoding operation is decoding, which reconstructs an approximation of the represented value from the momentary population activity at time t by multiplication with a decoding matrix $D \in \mathbb{R}^{d \times n}$, i.e., $\vec{x}(t) \approx D\vec{a}(t)$ (Fig 1C). Finding D can be phrased as a Tikhonov regularized least-squares optimization problem

$$\min_D \sum_{k=1}^N \|D\vec{a}(\vec{x}_k) - \vec{x}_k\|_2^2 + \lambda N \|D\|_F^2 = \min_D \|DA - X\|_F^2 + \lambda N \|D\|_F^2, \quad (2)$$

where $\|\cdot\|_2$ is the L_2 -norm, $\|\cdot\|_F$ is the Frobenius matrix norm, $\vec{x}_k$ is one of N training samples, $A \in \mathbb{R}^{n \times N}$ is a matrix of population responses for each sample, $X \in \mathbb{R}^{d \times N}$ is a matrix of training samples, and λ is a regularization term accounting for spike variability and other sources of error (Eliasmith and Anderson, 2003).

The decoders D are usually optimized offline under the assumption of the rate model $G[J]$. As formulated, the same D can be used to decode represented values through time in spiking networks when defining activity $\vec{a}(t)$ as low-pass filtered population spike trains. Linear low-pass filters are a common model for post-synaptic currents (Roth and van Rossum, 2009) and usually employed in spiking NEF networks (Eliasmith and Anderson, 2003).

Transformation

Nonlinear encoding and linear decoding schemes similar to the one described above have been analyzed in more detail by researchers in the field of machine learning (Broomhead and Lowe, 1988). Neuron population tuning-curves $\vec{a}(\vec{x})$ span a function space with a set of non-orthogonal basis functions. We can approximate any continuous function over the represented values to a desired degree by linearly combining a finite number of these basis functions (Hornik et al., 1989). Specifically, the linear projection D in eq. (2) approximates the identity function. By modifying the loss function and substituting X with a matrix X^f of target vectors $f(\vec{x}_k)$ we can solve for decoders D^f approximating some function f (Fig 1D)

$$\min_{D^f} \sum_{k=1}^N \frac{1}{2} \|D^f \vec{a}(\vec{x}_k) - f(\vec{x}_k)\|_2^2 + \lambda N \|D^f\|_F^2. \quad (3)$$

In order to construct neural *networks*, we need to find synaptic weight matrices $W \in \mathbb{R}^{m \times n}$ that connect from a pre-population of n neurons to a post-population of m neurons. In particular, we would like to find a W that implicitly decodes $f(\vec{x}) = \vec{y}$ from the pre-population and at the same time encodes $\vec{y}$ in the post-population. If we assume that the current translation function $J_i(x)$ is an intrinsic part of the neuron model, or, put differently, each neuron i is assigned its own response-curve $G_i[\langle \vec{e}_i, \vec{x} \rangle] = G[J_i(\langle \vec{e}_i, \vec{x} \rangle)]$, both the encoding and decoding process are linear operators. With this simplification, $W = ED^f$ fulfils the properties listed above, where $E \in \mathbb{R}^{m \times d}$ is a matrix of post-population encoding vectors $\vec{e}_i$, and D^f the desired function decoder for the pre-population (Eliasmith and Anderson, 2003).

Crucially, when combining the input from multiple pre-populations, the post-po-




Figure 2: Multivariate computation in the NEF. (A) In case the current translation function J_i is a part of the individual neuron response curve, NEF networks are additive: summation in activity space corresponds to addition in representational space. (B) Computing nonlinear multivariate functions φ generally requires all variables to be represented in an intermediate population. (C) The dendritic computation scheme discussed in this paper. Two pre-populations project onto a post-population with separate excitatory and inhibitory input channels. The nonlinear interaction between these channels is exploited to compute φ .

pulation will always represent a linear combination of potentially nonlinear functions. To see this, consider two pre-populations projecting onto a common post-population, where the first projection approximates a function f_1 , and the second a function f_2 . Let i be the index of a neuron in the post-population. Then, it holds

$$\begin{aligned} G_i[\langle \vec{w}_i^{f_1}, \vec{a}^{\text{pre}}(\vec{x}_1) \rangle + \langle \vec{w}_i^{f_2}, \vec{a}^{\text{pre}}(\vec{x}_2) \rangle] &= G_i[\langle \vec{e}_i, D^{f_1} \vec{a}^{\text{pre}}(\vec{x}_1) + D^{f_2} \vec{a}^{\text{pre}}(\vec{x}_2) \rangle] \\ &\approx a_i^{\text{post}}(f_1(\vec{x}_1) + f_2(\vec{x}_2)). \end{aligned} \quad (4)$$

As a consequence, if we, for example, try to multiply two scalars x and y , these values must be represented as a two-dimensional vector in a common pre-population. We alleviate this restriction by introducing input-dependent nonlinearities (Figs 2B and 2C).

2.2 Extending the Neural Engineering Framework

The NEF as described above makes several assumptions that are not biophysically plausible. This includes the presence of a bias current β_i in each neuron, and excitatory and inhibitory connections not being treated separately. We describe two extensions that lift the aforementioned assumptions and present an alternative synaptic weight solving procedure that takes subthreshold currents into account, followed by our proposed method for accounting for input-dependent nonlinearities.

Decoding the current translation function

In the previous subsection we assumed that the current translation function $J_i(x)$ is an intrinsic part of the neuron model. Consequently, each neuron is not only assigned a neuron-specific tuning-curve $a_i(\vec{x})$, but also a neuron-specific response-curve $G_i[J]$. This comes at the cost of biological plausibility, since neurons in general do not possess

a strong bias current source.

Tripp and Eliasmith (2007) demonstrate that it is possible to robustly solve for synaptic weight matrices that approximate arbitrary post-synaptic current functions. We can use this insight to directly solve for weights that approximate the current translation function $J_i(\langle \vec{e}_i, \vec{x} \rangle)$, instead of optimizing with respect to represented values $\vec{x}$. Since, for now, we assume that the post-synaptic current is linear in the pre-population activities, we must find a weight vector $\vec{w}_i$ such that the following loss is minimized

$$\min_{\vec{w}_i} \sum_{k=1}^N \left(J_i(\langle \vec{e}_i, f(\vec{x}_k) \rangle) - \langle \vec{w}_i, \vec{a}^{\text{pre}}(\vec{x}_k) \rangle \right)^2 + \lambda N \|\vec{w}_i\|_2^2. \quad (5)$$

This equation can be brought into canonical least squares form and solved as before.

Nonnegative weights and Dale’s principle

So far we have assumed that synaptic weights are real-valued. This is problematic for two reasons. First, the least-squares optimization proposed above arbitrarily assigns algebraic signs to the synaptic weights; we cannot specify which pre-neurons are excitatory, and which inhibitory. Being able to do so is important, since neurons tend to follow Dale’s principle—individual pre-neurons exclusively influence all their post-neurons in either an excitatory or inhibitory manner (Strata and Harvey, 1999). Empirical data suggest that, depending on the modeled brain region, excitatory cells outnumber inhibitory cells by a factor between two and four (Hendry and Jones, 1981; Gabbott and Somogyi, 1986). Second, real-valued weights do not generalize to conductance-based synapses. The concept of negative conductances—in contrast to negative currents—is neither physically nor biologically plausible. Biological correlates of synaptic weights, such as the number of vesicles released from the pre-synapse or the channel density in the

post-synapse, are inherently nonnegative quantities (Roth and van Rossum, 2009).

An extension to the NEF that takes these biological constraints into account has been proposed by Parisien et al. (2008). The *Parisien transform* splits each projection into an excitatory and inhibitory path, where the latter mediates the signal from the excitatory pre-population through a population of inhibitory interneurons. The solution we discuss here does not introduce interneurons, and as such does not affect the structure of the network. Modelers using the following method will have to explicitly specify inhibitory interneuron populations if present in the modeled circuit.²

For the sake of simplicity, we assume in this paper that each population is arbitrarily split into a group of excitatory and inhibitory neurons. We write the somatic input current of post-neuron i in response to pre-synaptic activity as $\langle \vec{w}_i^+, \vec{a}^+ \rangle - \langle \vec{w}_i^-, \vec{a}^- \rangle$, where, $\vec{w}_i^\pm$ are the nonnegative excitatory and inhibitory weight vectors and $\vec{a}^\pm$ the activities of the excitatory and inhibitory neurons in the pre-population. Combining this current term with eq. (5) yields an optimization problem that allows us to solve for weights approximating f for each individual post-neuron i

$$\begin{aligned} \min_{\vec{w}_i^+, \vec{w}_i^-} \quad & \frac{1}{2} \sum_{k=1}^N \left\| \langle \vec{w}_i^+, \vec{a}_k^+ \rangle - \langle \vec{w}_i^-, \vec{a}_k^- \rangle - J_i(\langle \vec{e}_i, f(\vec{x}_k) \rangle) \right\|_2^2 + \lambda N \|\vec{w}_i^+\|_2^2 + \lambda N \|\vec{w}_i^-\|_2^2 \\ & = \frac{1}{2} \|\vec{w}_i' A' - \vec{j}\|_2^2 + \lambda N \|\vec{w}_i'\|_2^2, \text{ where } \vec{w}_i' = (\vec{w}_i^+, \vec{w}_i^-), A' = (A^+, -A^-)^T, \\ & \text{and } (\vec{j})_k = J_i(\langle \vec{e}_i, f(\vec{x}_k) \rangle), \end{aligned} \tag{6}$$

subject to $\vec{w}_i^+ \geq 0, \vec{w}_i^- \geq 0$. This problem is in canonical least-squares form and, as before, can be solved with a standard regularized nonnegative least-squares solver.

²Our software library `nengo-bio` facilitates the process of specifying network topologies with excitatory and inhibitory populations. The library provides special syntactic sugar for networks with inhibitory interneurons built without the aforementioned Parisien transform.

Alternatively, eq. (6) can be phrased as a convex quadratic program (QP), a generalization of least-squares optimization (Wright and Nocedal, 1999)

$$\min_{\vec{w}_i} \frac{1}{2} (\vec{w}'_i)^T ((A')^T A') \vec{w}'_i - \vec{j}^T A' \vec{w}'_i + \lambda N \|\vec{w}'_i\|_2^2, \quad \text{subject to } \vec{w}'_i \geq 0. \quad (7)$$

The global minimum of a convex QP can be computed in polynomial time (Kozlov et al., 1980). We propose a QP similar to eq. (7) to solve for conductance-based synaptic weights in the context of the two-compartment neuron model discussed below.

Equality relaxation for subthreshold currents

Most neurons possess a threshold current J_{th} below which their firing rate is zero. However, we do not take this into account when solving for somatic currents in eqs. (5) to (7)—we optimize for synaptic weights that *precisely* evoke certain post-synaptic currents, despite the magnitude of currents smaller than J_{th} having no effect on the neural activity in a steady-state neural response model. Instead of optimizing for equality, i.e., enforcing that the decoded current J_{dec} must equal the target current J_{tar} , we could relax this condition to an inequality constraint $J_{\text{dec}} \leq J_{\text{tar}}$ whenever $J_{\text{tar}} < J_{\text{th}}$.

We define a new optimization problem based on eq. (5) that treats target currents smaller than J_{th} as an inequality constraint

$$\min_{\vec{w}_i} \sum_{k=1}^N \frac{1}{2} \mathcal{E} \left(J_i(\langle \vec{e}_i, f(\vec{x}_k) \rangle), \langle \vec{w}_i, \vec{a}(\vec{x}_k) \rangle \right)^2 + \lambda N \|\vec{w}_i\|_2^2, \quad (8)$$

$$\text{where } \mathcal{E}(J_{\text{tar}}, J_{\text{dec}}) = \begin{cases} 0 & \text{if } J_{\text{tar}} < J_{\text{th}} \text{ and } J_{\text{dec}} < J_{\text{th}}, \\ J_{\text{th}} - J_{\text{dec}} & \text{if } J_{\text{tar}} < J_{\text{th}} \text{ and } J_{\text{dec}} > J_{\text{th}}, \\ J_{\text{tar}} - J_{\text{dec}} & \text{otherwise,} \end{cases}$$

and N is the number of samples, J_i the current translation function of the i th post-neuron, $\vec{e}_i$ the encoding vector, f the desired target function, $\vec{a}(\vec{x}_k)$ is the pre-population activity for the k th sample point $\vec{x}_k$, λ is the regularization factor, and J_{th} the aforementioned neuronal threshold current.

By splitting the matrix of pre-population activities A , and the vector of target currents $\vec{j}$, we can rewrite eq. (8) as a quadratic program. Let A^+ and $\vec{j}^+$ be the pre-population activities and target currents for superthreshold samples, and A^- be the activities for subthreshold samples. Then, the QP is given as

$$\min_{\vec{w}_i, \vec{s}_i} \frac{1}{2} \vec{w}_i^T ((A^+)^T A^+) \vec{w}_i - ((A^+)^T (\vec{j}^+)) \vec{w}_i + \|\vec{s}_i\|_2^2 + \lambda N \|\vec{w}_i\|_2^2, \quad (9)$$

subject to $A^- \vec{w}_i - \vec{s}_i \leq J_{\text{th}}$, where $\vec{s}_i$ is a vector of slack variables. Of course, the nonnegativity constraint from eq. (7) can be incorporated into this quadratic program.

Extension towards input-dependent nonlinearities

Up to this point we assumed current-based synapses. Their defining property is that the somatic current J is linear in the synaptic weights $\vec{w}$ and the pre-synaptic activities $\vec{a}$. Now, consider a neuron model with ℓ nonlinear input channels. We write the corresponding response curve $\mathcal{G}$ (mapping from neural input onto an average firing rate) for the i th neuron in a population as a multivariate function

$$a_i = \mathcal{G}[g_i^1, \dots, g_i^\ell] = \mathcal{G}[\langle \vec{a}_i^1, \vec{w}_i^1 \rangle, \dots, \langle \vec{a}_i^\ell, \vec{w}_i^\ell \rangle], \quad (10)$$

where $g_i^1, \dots, g_i^\ell$ describe some abstract “input state”, such as the conductance of each synaptic channel in a neuron model with conductance-based synapses. As expressed in the above equation, we assume that on average each g_i^k is linear in the pre-synaptic




Figure 3: Neural response curve decomposition. (A) Illustration of the multivariate neuron response curve $\mathcal{G}(g_E, g_I)$ for a two-compartment LIF neuron with excitatory and inhibitory conductance-based channels. (B,C) The chosen somatic nonlinearity G and its inverse G^{-1} . (D) corresponding input-dependent nonlinearity H . The neuron does not fire in the hatched regions, that is, G^{-1} is ill-defined.

activities $\vec{a}_i^k$. However, we do not make any assumption regarding the effect of g_i^k on the somatic current; more fundamentally, we do not assume that there exists an easily identifiable somatic current in the model at all.

The crucial idea is to mathematically reintroduce a “virtual” somatic current J by decomposing $\mathcal{G}$ into an input-dependent nonlinearity H and a somatic nonlinearity G . We define G and H according to the following equivalence relations

$$\mathcal{G}[g_i^1, \dots, g_i^\ell] = G[H(g_i^1, \dots, g_i^\ell)] \Leftrightarrow H(g_i^1, \dots, g_i^\ell) = J \Leftrightarrow G[J] = \mathcal{G}[g_i^1, \dots, g_i^\ell], \quad (11)$$

where H maps from the input channel state onto an average somatic current and, as before, G maps from a somatic current onto the output activity. In other words, H summarizes nonlinear effects caused by dendrites or conductance-based synapses.

While this formalization does not constrain G and H beyond the above equivalence, a sensible choice for G and H will facilitate solving for synaptic weights. For example, if the neuron model in question is an extension to the current-based LIF neuron model, it makes sense to select G as the standard LIF response curve. Then, H translates the input state into an “LIF-equivalent somatic current”.

As we show in the next section, H , or at least a parametrized surrogate for H , can be derived in closed form in some cases. In case this is not feasible, it is possible to purely rely on empirical data for H by sampling the neuron output rate over varying synaptic states. Assume that we can only observe H indirectly by controlling the input channels of our model neuron and measuring the output activity $\mathcal{G}$. Depending on our choice of $G[J]$, we can apply an inverse mapping G^{-1} to the recorded data to obtain H . See Fig 3 for an illustration.

When solving for weights that approximate a specific function, recall from the above review that the first NEF principle normatively assigns a somatic current $J_i(\vec{x})$ to each post-neuron i and representational state $\vec{x}$. Correspondingly, given H , we can combine eq. (10) with the current-decoding problem discussed above, as well as the nonnegativity constraint and the equality relaxation:

$$\min_{\vec{w}_i^1, \dots, \vec{w}_i^\ell} \sum_{k=1}^N \mathcal{E} \left(J_i(\langle \vec{e}_i, f(\vec{x}_k) \rangle), H(\langle \vec{w}_i^1, \vec{a}_k^1 \rangle, \dots, \langle \vec{w}_i^\ell, \vec{a}_k^\ell \rangle) \right)^2 + \lambda N \sum_{j=1}^\ell \|\vec{w}_i^j\|_2^2, \quad (12)$$

subject to $\vec{w}_i^j \geq 0$, where $\mathcal{E}$ is as defined in eq. (8).

2.3 Two-compartment leaky integrate-and-fire neuron model

In the previous section we established the abstract notion of an input-dependent nonlinearity H . We now derive a concrete H for a biologically plausible extension to the leaky integrate-and-fire (LIF) neuron with nonlinear conductance-based synapses. In particular, we consider a two-compartment version of the LIF neuron originally described in Vu et al. (1993) and subsequently discussed by Koch (1999) and Capaday and Van Vreeswijk (2006). In addition to the active, somatic compartment, the two-compartment LIF model possesses a resistively coupled passive compartment that represents excitatory and inhibitory input into the dendritic tree. Depending on the coupling conductance g_C , the input may either be interpreted as distal or proximal.

We first review the neuron model itself, derive a parametrized surrogate for the nonlinearity H , and finally propose a convex quadratic program for this H that can be used to solve for synaptic weights.

Model description

The sub-threshold dynamics of the conductance-based two-compartment LIF model can be expressed as a two-dimensional system of linear differential equations

$$\begin{aligned}\frac{d}{dt}C_{m,1}v_1 &= g_C(v_2 - v_1) + g_{L,1}(E_L - v_1), \\ \frac{d}{dt}C_{m,2}v_2 &= g_C(v_1 - v_2) + g_{L,2}(E_L - v_2) + g_E(t)(E_E - v_2) + g_I(t)(E_I - v_2),\end{aligned}\tag{13}$$

where state variables v_1, v_2 correspond to the membrane potential of the active somatic compartment and the passive dendritic compartment, respectively. $C_{m,1}, C_{m,2}$ are the compartment capacitances, g_C is the inter-compartment coupling conductance, $g_{L,1}, g_{L,2}$ are the individual compartment leak conductances, $g_E(t), g_I(t)$ are the momentary excitatory and inhibitory conductances of the dendritic compartment as evoked by pre-synaptic spikes, and E_L, E_E, E_I are the conductance-channel reversal potentials. An equivalent circuit diagram of the model is shown in Fig 4A.

In contrast to point-neuron models, and as pointed out by Capaday and Van Vreeswijk (2006), a multi-compartment model mandates an explicit spike model. The strong depolarization in the somatic compartment during spike production propagates backwards into the dendritic compartment and has a significant effect on its membrane potential (Fig 4B). The model accounts for this with a “spike-phase” occurring right before the LIF refractory period. The spike phase is implemented by clamping the somatic compartment to a voltage v_{spike} over a time τ_{spike} whenever the somatic membrane potential crosses the threshold v_{th} . Subsequently, the membrane is clamped to v_{reset} for a period τ_{ref} .




Figure 4: Two-compartment LIF neuron model. (A) Equivalent circuit diagram of the two-compartment LIF neuron model. The “somatic” compartment corresponds to a classical LIF neuron; the “dendritic” compartment is resistively coupled to the somatic compartment. The current flowing between the two compartments corresponds to the somatic current $J = H(g_E, g_I)$. (B) Voltage trace for the two-compartment model state variables v_1 and v_2 during spike production. The artificial depolarization following spike events in the somatic compartment significantly affects v_2 .

Somatic current surrogate model

We assume that $H(g_E, g_I)$ is equivalent to the current flowing from the dendritic into the somatic compartment (Fig 4A), that is $H(g_E, g_I) = g_C(v_2 - v_1)$. Considering the definition in eq. (11), this implies that G is the standard LIF response curve

$$G[J] = \left(\tau_{ref} + \tau_{spike} - \frac{C_m}{g_L} \ln \left(1 - \frac{(v_{th} - E_L)g_L}{J} \right) \right)^{-1}, \quad (14)$$

where τ_{ref} , τ_{spike} are the length of the refractory and spike periods (Fig 4B), C_m is the membrane capacitance, v_{th} is the threshold potential, E_L is the leak reversal or resting potential, and g_L is the leak conductance (Eliasmith and Anderson, 2003). Yet,

in practice, as pointed out by Hunsberger et al. (2014), and as we demonstrate in our experiments below, a rectified linear unit (ReLU) may be a sensible choice for G as well when modeling noisy input.

Unfortunately, when considering both sub- and superthreshold dynamics, there exists no exact closed-form solution for the average somatic current given constant g_E , g_I . Instead, our approach is to select a parametrized surrogate model for H and to fit the model parameters to results obtained from numerical simulations to account for inaccuracies in our derivation.

Assuming the subthreshold dynamics described in eq. (13) are in equilibrium, g_E , g_I are constant, and applying the above definition of $H(g_E, g_I)$, we get

$$H(g_E, g_I) = g_C(v_2 - v_1) = g_C \frac{g_{L,2}(E_L - v_1) + g_E(E_E - v_1) + g_I(E_I - v_1)}{g_C + g_{L,2} + g_E + g_I}. \quad (15)$$

A single-compartment model can be derived by taking the limit of this equation for $g_C \rightarrow \infty$. In this case, and as demonstrated in Stöckel et al. (2017), H is an affine function and less interesting from a computational perspective. In general, we expect H to become more nonlinear as g_C decreases. While a higher degree of nonlinearity may be desirable for computation in a neural network, g_C cannot be made arbitrarily small, as this limits the maximum somatic input current (see below).

Since the somatic membrane potential v_1 is clamped during the spike and refractory phases, the current flowing into the soma during those times does not influence the overall firing rate (ignoring the feedback effect on v_2 discussed earlier). Once the neuron is tonically firing, the somatic membrane potential oscillates between v_{reset} and v_{th} . We can thus substitute v_1 with a constant average membrane potential $\bar{v}_{\text{som}} \approx \frac{1}{2}(v_{\text{reset}} + v_{\text{th}})$ (Fig 5). Under these assumptions, $H(g_E, g_I)$ from eq. (15) can be written as a




Figure 5: Average LIF membrane potential over firing rate. Data corresponds to a standard LIF neuron (i.e., the somatic compartment) over varying output rates. Data measured by capturing 1000 membrane potential traces for a ten second current sweep with superimposed spike noise from a Poisson spike source. Individual lines correspond to the effective rate of the spike source, where smaller input rates are equivalent to a larger amount of noise on the input current. Shaded areas correspond to the 25/75 percentile. (A) Average potential excluding the refractory period and spike production. Except for very low firing rates, the average potential remains relatively constant. Dotted line corresponds to $\bar{v}_{\text{som}} = \frac{1}{2}(v_{\text{reset}} + v_{\text{th}})$. (B) Average potential including the refractory and spike production. Dotted line corresponds to a linear model that takes the relative amount of time spent in the subthreshold, spike, and refractory regime into account.

parametrized rational function

$$H(g_E, g_I) = \frac{b_0 + b_1 g_E + b_2 g_I}{a_0 + a_1 g_E + a_2 g_I}, \quad \text{where } a_0, a_1, a_2, g_E, g_I \geq 0. \quad (16)$$

This equation has one superfluous degree of freedom in the parameter space. Setting $b_1 = 1$ tends to be a numerically stable normalization.

Equation (16) implies an absolute maximum and minimum somatic current; $H(g_E, g_I)$ maps onto an open interval $(J_{\max}, J_{\min})$, where

$$\begin{aligned} J_{\min} &= \lim_{g_I \rightarrow \infty} H(g_E, g_I) = -\frac{b_2}{a_2} = g_C(E_I - \bar{v}_{\text{som}}), \\ J_{\max} &= \lim_{g_E \rightarrow \infty} H(g_E, g_I) = \frac{b_1}{a_1} = g_C(E_E - \bar{v}_{\text{som}}). \end{aligned} \quad (17)$$

In practice, the maximum attainable current for realistic conductance values is significantly smaller than $J_{\max}$, limiting the maximum firing rate. This must be taken into account when selecting the neuron tuning curve.

Model parameters can be estimated by solving the QP

$$\min_{\substack{b_0, b_2, \\ a_0, a_1, a_2}} \sum_{\substack{i, \\ J_i \gg J_{\text{th}}}} (b_0 + b_1 g_E^i + b_2 g_I^i - J^i a_0 - J^i g_E^i a_1 - J^i g_I^i a_2 - g_E^i)^2, \quad (18)$$

subject to the nonnegativity constraints in eq. (16), where $J^i = G^{-1}[\mathcal{G}(g_E^i, g_I^i)]$. The conductances g_E^i, g_I^i should be sampled over the operating range of the neuron, and samples with zero or very small output rates ignored: the inverse of G is not defined for a zero output rate, and H was derived under the assumption of superthreshold dynamics.

Synaptic weights as a quadratic program

Given the nonlinearity model H as defined in eq. (16), our goal is to find weights $\vec{w}_i^E, \vec{w}_i^I$ such that a desired current $(\vec{j})_k$ flows into the soma for every sample k , cf. eq. (6). Due to H being a rational function, we cannot directly minimize the loss function in eq. (12) as

a quadratic program. Instead, we solve a related quadratic program which, in practice, produces reasonable solutions. In particular, we simply equate the desired currents and the input-dependent nonlinearity model H . Let $\div$ and $\circ$ denote elementwise division and multiplication. Then, in vector notation, we get

$$\vec{j}_i = \frac{b_0 + b_1 A \vec{w}_i^E - b_2 A \vec{w}_i^I}{a_0 + a_1 A \vec{w}_i^E + a_2 A \vec{w}_i^I}, \quad \text{where } a_0, a_1, a_2, A, \vec{w}_i^E, \vec{w}_i^I \geq 0,$$

and A is a matrix of pre-population activities. Rearranging into a canonical form yields

$$\begin{aligned} \vec{j}_i \circ (a_0 + a_1 A \vec{w}_i^E + a_2 A \vec{w}_i^I) &= b_0 + b_1 A \vec{w}_i^E - b_2 A \vec{w}_i^I, \\ \Leftrightarrow (a_1 J_i \circ A - b_1 A) \vec{w}_i^E + (a_2 J_i \circ A + b_2 A) \vec{w}_i^I &= b_0 - a_0 \vec{j}_i, \text{ where } (J_i)_{mn} = (\vec{j}_i)_m, \\ \Leftrightarrow \mathfrak{A}_i \vec{w}'_i &= \vec{\mathfrak{b}}_i, \end{aligned} \tag{19}$$

where $\mathfrak{A}_i = (a_1 J_i \circ A - b_1 A, a_2 J_i \circ A + b_2 A)^T$, $\vec{\mathfrak{b}}_i = (b_0 - a_0 \vec{j}_i)$, and $\vec{w}'_i = (\vec{w}_i^E, \vec{w}_i^I)$.

We account for the subthreshold equality relaxation, by splitting $\mathfrak{A}_i$, $\vec{\mathfrak{b}}_i$ according to the samples invoking a zero firing rate (resulting in $\mathfrak{A}_i^-$, $\vec{\mathfrak{b}}_i^-$) and those invoking a positive firing rate (resulting in $\mathfrak{A}_i^+$, $\vec{\mathfrak{b}}_i^+$). Thus, the quadratic program becomes:

$$\begin{aligned} \min_{\vec{w}'_i, \vec{s}_i} \frac{1}{2} (\vec{w}'_i)^T ((\mathfrak{A}^+)^T \mathfrak{A}^+) \vec{w}'_i - (\vec{\mathfrak{b}}_i^+)^T \mathfrak{A}^+ \vec{w}'_i + \lambda N \|\vec{w}'_i\|_2^2 + \|\vec{s}_i\|_2^2, \\ \text{subject to } \mathfrak{A}_i^- \vec{w}'_i - \vec{s}_i \geq \vec{\mathfrak{b}}_i^- \text{ and } \vec{w}'_i \geq 0. \end{aligned} \tag{20}$$

3 Experiments and Results

In this section, we validate the methodology described above. To this end, we perform three experiments. In Experiment 1, we test how well the two-compartment LIF nonlinearity model H predicts simulation data for a single spiking neuron. In Experiment 2, we study the computational properties of H under optimal circumstances when approximating random bandlimited functions. Finally, in Experiment 3, we test whether the

results from Experiment 1 and 2 are still visible in the context of a noisy, feed-forward spiking neural network.

Unless explicitly specified, the neuron model parameters are chosen according to Table A1. We model fast excitatory synapses as an exponential low-pass with a time-constant of 5 ms as found in glutamatergic pyramidal neurons with AMPA receptor (Jonas et al., 1993). The inhibitory pathway is modeled with a 10 ms time-constant as found in inhibitory interneurons with GABA_A receptors (Gupta et al., 2000). We use the `cvxopt` library (Vandenberghe, 2010) as a QP solver for eq. (18), and the `osqp` library (Stellato et al., 2020) to solve eq. (20).³ The source code of the computer programs used to conduct and evaluate the experiments can be found in the supplemental materials.

3.1 Experiment 1: Fitting the surrogate model to simulation data

The surrogate model $H(g_E, g_I)$ in eq. (16) predicts the current flowing into the somatic compartment of a two-compartment LIF neuron for excitatory and inhibitory conductances g_E and g_I . H is characterized by five parameters a_0, a_1, a_2, b_0, b_2 , for which a coarse theoretical estimate can be derived from eq. (15). We compare spike rates measured in numerical simulation to the spike rate prediction obtained before and after fitting the model parameters to empirical data using eq. (18). We do this both for constant conductances g_E, g_I and for artificial temporal spike noise superimposed on the conductance pair.

³`osqp` is used as a part of `libbioneuronqp`, see <https://github.com/astoeckel/libbioneuronqp>

Experiment 1.1: Constant conductances

We consider three two-compartment LIF neurons with different coupling conductances g_C , namely 50 nS, 100 nS, and 200 nS, and measure their output spike rate for constant conductances g_E , g_I on a grid. The conductance range has been selected such that the maximum firing rate is 100 s^{-1} , and the spike onset approximately coincides with the diagonal of the resulting g_E - g_I -rate contour plot. We measure the steady-state firing rate by taking the inverse of the median inter-spike-interval over the simulation period. We compare these data to the current predicted by H according to the theoretical and optimized estimated parameter sets. Parameter optimization is based on a training-set of 200 conductance pairs sampled with uniform probability from the conductance-range. The final prediction error is computed over all 10 000 grid points.

The results for this experiment are depicted in Fig 6A, the fitted parameters can be found in Table A2. When using the theoretical parameter estimate from eq. (15), there is a significant discrepancy between the model prediction and the numerical simulation, especially for large g_C . This discrepancy is greatly reduced after fitting the model parameters. The model prediction almost perfectly fits the empirical data for output spike rates greater than 25 s^{-1} . However, it fails to predict the spike onset correctly, placing it too early with respect to increasing g_E . Furthermore, the predicted slope at the spike onset is less steep than what is actually measured. As discussed in section 2.3, we can see the linearity of H increase as g_C is increased, i.e., the contour lines are more “parallel” for larger g_C . Still, with an overall Root Mean Square Error (RMSE) of about 4 s^{-1} , the model provides a reasonable approximation of the empirical data.




Figure 6: Somatic current model parameter fits. (Continued on next page.)

Figure 6: Somatic current model parameter fits (see previous page). Average spike rates $\mathcal{G}(g_E, g_I)$ measured in simulation are depicted as colored contour plots and solid lines in the cross-section. Dashed lines correspond to the model prediction $G[H(g_E, g_I)]$. Dotted lines indicate the location of the cross-section. E denotes the spike-rate RMSE over the regions where either the measured or predicted spike rate is greater than 12.5 s^{-1} . Columns correspond to different values of g_C . Top half shows the spike measured and modeled spike rates for constant conductance values g_E, g_I . The bottom half corresponds to conductance values with superimposed spike noise. The model fits the simulated response functions well after optimization.

Experiment 1.2: Conductances with artificial temporal spike noise

In a network context, $g_E(t)$ and $g_I(t)$ are usually not constant, but instead modeled as the weighted sum of low-pass filtered spike trains, resulting in a considerable amount of “spike noise” being superimposed onto the signal. In this experiment, we simulate artificial spike noise as two Poisson spike sources (one excitatory, one inhibitory) with rate $1/\lambda = 1800 \text{ s}^{-1}$ for inhibitory synapses and a rate of $1/\lambda = 4500 \text{ s}^{-1}$ for excitatory synapses. These rates were obtained by fitting Poisson distributions to measurements from Experiment 3. Spikes arriving at different synapses are simulated by uniformly sampling a random weight from the unit-interval for each spike event. The time-averaged conductance equals g_E, g_I , respectively. Apart from the simulation period being extended to one hundred seconds, the remaining experimental setup is unchanged from the last experiment.

As can be seen in our results for noisy conductances in Fig 6B, and particularly in the bottom cross-sections of the measured spike-rates in Fig 6B, the steep spike-onsets predicted by the theoretical LIF response curve $G[J]$ (eq. 14) are no longer present. Furthermore, the relationship between g_E and the rate appears to be roughly linear in each cross-section of the g_E - g_I -rate plot, which is not well captured by the standard LIF response curve. Hence, a “soft” version of the LIF response curve that takes noise into account would be a better choice when fitting the parameters (Hunsberger et al., 2014; Capocelli and Ricciardi, 1971; Kreutz-Delgado, 2015). We instead take a pragmatic approach and define G as a rectified linear unit (ReLU), that is $G[J] = \max\{0, \alpha J + \beta\}$. Due to a relatively pronounced noise floor we only consider samples with a spike rate greater than 12.5 s^{-1} for both fitting the model parameters and reporting the RMSE.

Fitting the parameters using the ReLU results in excellent accuracy for spike rates greater than 12.5 s^{-1} (RMSE less than 1 s^{-1}). However, the fitted model does not capture the subtle sigmoid shape of the response curve near the spike onset that is particularly pronounced for larger g_C .

To summarize, the results indicate that, after fitting the model parameters, our surrogate model H can indeed be used to predict the neural response curve with a relatively high accuracy. When taking noise in the input into account, it can be beneficial to choose a different neural response curve $G[J]$.

3.2 Experiment 2: Computational properties of the two-compartment

LIF nonlinearity H

The above experiment is encouraging in that our input-dependent nonlinearity H seems to predict somatic currents well. Hence, in this experiment, we assume that H accurately describes the somatic current and analyze whether there, in theory, could be any advantage to using a neuron model with this kind of input-dependent nonlinearity.

Conceptually, what we would like to do is to measure how well a system can approximate functions of increasing complexity. If systems using input-dependent nonlinearities are able to approximate complex functions with a lower error than linear systems, they are “computationally more powerful”. Since “complexity” is somewhat ill-defined, we choose the spatial frequency content of a function as a proxy.

In our experiment, we randomly generate bandlimited current functions $J_\sigma(x, y)$ over the compact domain $(x, y) \in [-1, 1]^2$ on a 63×63 grid. σ is a spatial low-pass filter coefficient that is inversely proportional to the bandwidth of the function. The functions we generate are normalized such that the mean is zero and the standard deviation equals 1 nA. We measure how well we can approximate this desired post synaptic current for a single post-neuron with a given input-dependent nonlinearity. The pre-neuron populations are as depicted in Fig 2C. Two independent pre-populations with 100 neurons each represent x and y , respectively, and project onto a single post-neuron. The population tuning curves are randomly generated in each trial with a maximum firing rate between 50 and 100 s^{-1} per neuron. All pre-neurons project both inhibitorily and excitatorily onto the post-neuron. We measure the static decoding error, i.e., the difference between the target current function and the output of the surrogate model

assuming ideal pre-population tuning curves without any dynamics. We consider dynamical simulation of a spiking *network* in Experiment 3.

All synaptic weights are computed by solving the QP in eq. (20) for 256 randomly selected training samples. For the conductance-based two-compartment LIF nonlinear input-dependent current function $H_{\text{cond}} = H$ we use the model parameters derived in the last experiment (Table A2). We emulate linear input currents $H_{\text{cur}} = J_E - J_I$ in terms of H by setting the model parameters to $b_0 = a_1 = a_2 = 0$, $a_0 = b_1 = 1$, $b_2 = -1$. The regularisation parameters were selected independently for each setup (Fig A1). The final error is the normalised RMSE (relative to the standard deviation of 1 nA of f_σ) over all 3969 grid points. We add normal distributed noise (zero mean, unit standard deviation) to the pre-activities when computing the error to test how well the computed weights generalize.

As a further point of comparison, we include a “two-layer” neural network setup (Fig 2B). The 200 pre-neurons are tuned to both input dimensions (x, y) , and not x and y independently. This setup should perform the best—but keep in mind that when modeling neurobiological systems, we may be constrained to pre-populations that do not exhibit the kind of multivariate tuning to the input signals we assume here.

Results are depicted in Fig 7. For a current-based neuron, and without subthreshold relaxation (dashed line in the plot), the median error increases linearly on a log-log plot from a 2.5% error for low-frequency—almost linear—functions to an error of about 50% for functions with a spatial cut-off frequency greater than one. Subthreshold relaxation reduces this error by up to 50% for low σ^{-1} .

The error for the conductance-based nonlinearity increases sub-linearly on a log-




Figure 7: Median decoding error for random functions. Continued on next page.

Figure 7: Median decoding error for random multivariate current functions (see previous page). (A) Normalized RMSE between random, two-dimensional current functions and the decoded approximation using different input-dependent nonlinearities. The error measurement does not take subthreshold currents into account, using the function $\mathcal{E}$ defined in eq. (8). The low-pass filter coefficient σ^{-1} is a proxy for the spatial frequency content in the target function. All points correspond to the median over 1000 trials. Dashed lines show results for not taking the subthreshold relaxation into account when solving for weights. The black lines show the results for a linear, current-based H_{cur} ; blue/green lines show the results for the two-compartment conductance-based model H_{cond} with parameters given in Table A2 (without noise). Orange lines correspond a current-based network with two-dimensional pre-neuron tuning (i.e., a two-layer neural network). Shaded areas correspond to the 25/75 percentile for the current-based models and the conductance-based model with $g_C = 50 \text{ nS}$. (B) Exemplary random functions and the corresponding decodings. Deep violet regions (dashed contour lines) correspond to subthreshold currents. Note how the shape of the subthreshold contour lines no longer match the target when subthreshold relaxation is active. As visible, the conductance-based nonlinearity H_{cond} helps to decode some target functions with a drastically smaller error compared to the current-based model H_{cur} , especially when comparing the setups without subthreshold relaxation. H_{cond} does not provide any benefit for target functions with a high bandwidth.

log plot, starting at median errors of about 0.8% for low-frequency functions. It is competitive with the two-layer network (see below) for $\sigma^{-1} < 0.5$. The error function converges to the results for the current-based model for spatial frequencies $\sigma^{-1} > 1$. Overall, the error for the conductance-based model is reduced by up to 65% compared to the current-based model. The benefits of subthreshold relaxation are not as pronounced as for the linear current model.

The large errors for the current-based and the conductance-based models for $\sigma^{-1} > 1$ can be explained by the fact that both functions cannot be used to solve the XOR problem (cf. Appendix C). Functions with $\sigma^{-1} \geq 1$ are likely to possess multiple maxima/minima over $[-1, 1]^2$, akin to XOR, leading to a large error.

The two-layer network setup is able to approximate functions well up to $\sigma^{-1} = 10$, where it reaches the same final error values as the other setups. The complexity of the functions that can be approximated well by this setup is limited by the number of pre-neurons. The two-layer network can be thought of as linearly combining rectified hyperplanes to fit the target function, where each hyperplane is a single pre-neuron response. At a certain σ , the number of hyperplanes is no longer sufficient to reconstruct all the local maxima/minima in the target function.

To summarize, this experiment demonstrates that using a neuron model with the input-dependent nonlinearity H significantly reduces the approximation error of current functions with $\sigma^{-1} < 1$ in a network setup with pre-populations independently representing the input dimensions. It is competitive with a two-layer network for $\sigma^{-1} < 0.5$.

3.3 Experiment 3: Dendritic computation of multivariate functions in spiking network models

Experiment 1 suggests that we can use the non-linear post-synaptic current model H to predict the average current flowing into the somatic compartment. Experiment 2 shows that we can, assuming that H accurately describes the somatic current, approximate a somewhat larger class of random functions well. In our final experiment we study whether we can still observe the reduction in error in a feed-forward spiking neural network when using our model H to solve for weights.

In contrast to previous experiment we do not base our error measurements on the decoded static somatic current for a *single* post-neuron. Instead we decode the represented value from the neural activities of a target *population* over time. Optimally, this decoded value should be $f(x(t), y(t))$, where f is the function that we want to approximate and $x(t), y(t) \in [-1, 1]$ are input signals in turn represented by populations of 100 LIF neurons each. The target population either consists of standard current-based LIF neurons (Fig 2A) or conductance-based two-compartment neurons (Fig 2C).

Just as in the previous experiment, we consider the two-layer topology in Fig 2B as a point of reference. Here, the input is mediated via an additional layer of 200 neurons representing the vectorial quantity $(x(t), y(t))$ over the interval $[-1, 1]^2$.

For all neurons, we generate tuning curves such that the maximum firing rate falls between 50 and 100 s⁻¹ over their represented range. Neurons are randomly marked as either excitatory or inhibitory. The probability of a neuron being inhibitory is 30%. Excitatory and inhibitory synapses are modeled as a first-order exponential low-pass filter with time-constants of $\tau_E = 5$ ms and $\tau_I = 10$ ms, respectively.

The network is simulated over 10 s at a time-resolution of 100 μ s. Inputs $x(t)$ and $y(t)$ are sampled by moving through time along a fourth-order space-filling Hilbert curve over $[-1, 1]^2$. The output of the target population is decoded and filtered with a first-order exponential low-pass at $\tau = 100$ ms. We compute the desired target value $f(x, y)$ from the original input and pass it through the same series of low-pass filters as the spiking signals. We use the average synaptic time-constant of 7.5 ms to emulate the effect of the synaptic low-pass filters. Our final measure E_{net} is the normalized RMSE between the decoded output and the target values over time; the normalization is relative to the standard deviation of the target signal.

All synaptic weights are computed by solving the QP in eq. (20), with the same mapping of the current-based onto the conductance-based model as in Experiment 2. The regularization term λ has been chosen independently for each neuron type and model parameter set such that the network error E_{net} is minimized when computing multiplication (cf. Fig B1).

Experiment 3.1: Random bandlimited functions

We first test our network setup with the random bandlimited functions f_σ we already used in the previous experiment. Results are depicted in Fig 8. Qualitatively, the results are very similar to what we saw before. The reduction in error between the current- and conductance-based models is not quite as large as suggested by the theoretical experiment, with a maximum reduction (in terms of the median) of only 45% (instead of 65% before). While subthreshold relaxation mostly increased the performance of the current-based model in the previous experiment, the improvement in error is now




Figure 8: Median error for computing random bandlimited functions in a feed-forward network over 1000 trials. Measured RMSE E_{net} is the difference between the represented value $D\vec{a}(t)$ and the expected value $f_{\sigma}(x(t), y(t))$ relative to the standard deviation of f_{σ} . See Fig 7A for more detail.

clearly visible for the conductance-based model as well.

Notably, the minimum median approximation error of the two-layer network is about 10%, whereas the single-layer current- and conductance-based models reach minimum errors of about 6.5% and 5%, respectively. The two-layer network clearly surpasses the performance of the two-compartment LIF single-layer network for $\sigma^{-1} > 0.6$. The larger errors are mainly caused by the representation of the two-dimensional quantity $(x(t), y(t))$ begin noisier than the representation of the scalars $x(t), y(t)$ in the two pre-populations. This is because chaining multiple populations of spiking neurons slightly increases the noise floor. Furthermore, to cover the square $[-1, 1]^2$ as densely as two one-dimensional intervals $[-1, 1]$, we would optimally have to square the number of neurons. In our case, we would have to use 10 000 instead of 200 neurons for the intermediate layer, which would not really be comparable to the single-layer setups—

keep in mind that the two-layer network already uses 66% more neurons.

Experiment 3.2: Benchmark functions

While the random functions in the above experiments (see Fig 7B for an example) are useful to systematically characterize the individual setups, it is hard to tell from these data alone what the practical impact of the two-compartment LIF neuron is. To this end, we selected eight mathematical benchmark functions $f(x, y)$ and repeated the experiment. Functions include the maximum $\max(x, y)$, and various forms of multiplication ($\sqrt{x \times y}$, $x \times y$, $(x \times y)^2$; see table A3 for a complete list). Note that we compute all these functions over the interval $[0, 1]^2$ instead of $[-1, 1]^2$ by shifting and scaling the values represented by the neuron populations, i.e., we compute $f((x + 1)/2, (y + 1)/2)$. As mentioned above and proved in Appendix C, we know that we are not be able to solve the XOR problem with the two-conductance LIF neuron, and multiplication over $[-1, 1]^2$ can be thought of as a continuous form of XOR. We should be able to approximate multiplication over $[0, 1]^2$ one quadrant however.⁴

A summary of the results over 256 trials per function and setup is given in Table 1, traces from an example trial are depicted in Fig 9. More detailed results can be found in Table A4. For all but one target function (squared multiplication, which has the highest bandwidth of all tested functions), the conductance-based two-compartment model with a coupling conductance of $g_C = 50$ ns achieves the smallest error E_{net} . Using the surrogate model parameters derived under noise is beneficial when computing multiplicative

⁴The obvious solution to approximating “full” multiplication using two-compartment LIF neurons is to split the target population into four quadrants; however, we wanted to use network setups that are not optimized for a particular problem.




Figure 9: Single spiking neuron experiment showing computation of multiplication using a two-compartment LIF neuron. (A) Top two plots: inputs $x(t)$ and $y(t)$ as represented by the two pre-populations. The input is a fourth order 2D Hilbert curve. Bottom: mathematical target $f(x(t), y(t)) = x(t)y(t)$, filtered target function, as well as the decoded target population output. (B) Spike raster plots corresponding to the spiking activity of each of the populations. Red shaded background corresponds to inhibitory neurons in the pre-populations, all other neurons are excitatory.

Table 1: Spiking neural network approximation errors for function approximations on $[0, 1]^2$. Error values E_{net} correspond to the normalized RMSE (relative to the standard deviation of the target function) and are measured as the difference between the output decoded from the target population and the desired output for a ten second sweep across a 4th order 2D Hilbert curve over the input space. Results are the mean and standard deviation over 256 trials. The best result for a target function is set in bold; darker background colors indicate a worse ranking of the result in the corresponding row. Columns labeled “standard” refer to the default, single layer network setup, “two layers” refers to the two-layer setup, and “noise model” to the single layer network setup with model parameters derived under noise (see Experiment 1.2). Additional tables can be found in Appendix A.

Target	Experiment setup						
	LIF			Two comp. LIF $g_C = 50 \text{ nS}$		Two comp. LIF $g_C = 100 \text{ nS}$	
	standard	standard [†]	two layers [†]	standard [†]	noise model [†]	standard [†]	noise model [†]
$x + y$	$4.2 \pm 0.3\%$	$4.2 \pm 0.3\%$	$8.2 \pm 0.4\%$	$2.3 \pm 0.3\%$	$7.1 \pm 0.8\%$	$3.7 \pm 0.4\%$	$8.3 \pm 0.9\%$
$x \times y$	$26.6 \pm 0.9\%$	$24.6 \pm 0.9\%$	$9.2 \pm 0.5\%$	$7.5 \pm 1.1\%$	$7.4 \pm 1.3\%$	$10.9 \pm 2.0\%$	$9.5 \pm 2.0\%$
$\sqrt{x \times y}$	$13.5 \pm 0.6\%$	$12.5 \pm 0.7\%$	$9.2 \pm 0.4\%$	$5.0 \pm 0.8\%$	$6.2 \pm 0.9\%$	$8.1 \pm 1.7\%$	$7.9 \pm 1.4\%$
$(x \times y)^2$	$45.6 \pm 1.5\%$	$42.6 \pm 1.5\%$	$10.9 \pm 1.1\%$	$19.7 \pm 3.4\%$	$16.0 \pm 3.4\%$	$22.4 \pm 3.9\%$	$18.6 \pm 4.1\%$
$x/(1 + y)$	$5.6 \pm 0.3\%$	$5.4 \pm 0.3\%$	$8.1 \pm 0.5\%$	$2.3 \pm 0.3\%$	$7.9 \pm 1.2\%$	$3.8 \pm 0.5\%$	$9.8 \pm 1.5\%$
$\ (x, y)\ $	$7.6 \pm 0.5\%$	$7.4 \pm 0.5\%$	$8.1 \pm 0.4\%$	$2.2 \pm 0.2\%$	$6.4 \pm 0.7\%$	$2.7 \pm 0.4\%$	$8.9 \pm 0.9\%$
$\text{atan}(x, y)$	$9.4 \pm 0.5\%$	$9.0 \pm 0.5\%$	$9.7 \pm 0.5\%$	$4.0 \pm 0.8\%$	$7.4 \pm 0.8\%$	$6.1 \pm 1.2\%$	$11.6 \pm 1.2\%$
$\max(x, y)$	$14.9 \pm 0.6\%$	$13.8 \pm 0.6\%$	$8.4 \pm 0.3\%$	$6.9 \pm 0.7\%$	$6.4 \pm 0.7\%$	$9.4 \pm 1.1\%$	$9.0 \pm 0.7\%$

[†]With subthreshold relaxation

functions and the maximum. For these target functions, the synaptic connection matrix tends to be sparser, increasing the input noise. Apparently, this increase in noise matches the environment the neuron parameters have been optimized for. Interestingly, a purely current-based, single-layer network is competitive for all functions except for multiplication. The minimum error for the two-layer network is about 8% even for simple functions, matching the observation we made in the random function experiment above.

An effect that could contribute to the superior performance of the two-compartment neuron model in some experiment are the low-pass filter dynamics of the dendritic compartment. These filter the high-frequency spike noise and thus may reduce the target error. We control for this effect in an experiment described in Appendix B, where we add an optimal low-pass filter to each network setup. Results are shown in Table B1. We find that a matched pre-filter consistently reduces the error of all setups by only 1%–2%, which indicates that the low-pass filter dynamics of the dendritic compartment are not the primary source for the reduction in error.

To summarize our experiments, we demonstrate in three stages (validation of the nonlinearity model H_{cond} for a single neuron, purley mathematical properties of H_{cond} , and, finally, performance on a network-level) that we are able to successfully incorporate an—admittedly simple—model of nonlinear passive dendritic interaction into functional modeling frameworks. Instead of reducing the accuracy of our networks, the added detail can be systematically leveraged for computation. Our experiments also suggest that—at least in a biologically informed setting, i.e., using spiking neurons—this type of computation may result in a higher accuracy compared to two-layer archi-

tructures that suffer from an increase in the amount of spike-induced temporal noise due to the additional neuron layer.

4 Discussion

We derived a mathematical model of input-dependent post-synaptic currents in a two-compartment LIF neuron that can be interpreted as a simple form of passive dendritic computation. We experimentally demonstrated that networks with fewer layers but biophysically plausible nonlinearities can compute a broad range of multivariate functions as well as or better than networks typically constructed using functional modeling frameworks. In particular, we proposed a mathematical model H that captures nonlinear interactions between input channels, for example caused by conductance-based synapses or the dendritic tree. By mapping individual channel states onto an average somatic current J , this model can be integrated into mathematical frameworks that classically rely on current-based input channels.

Specifically, we demonstrated how to incorporate the dendritic nonlinearity H into the Neural Engineering Framework (NEF). To this end, we discussed extensions to the NEF that allow us to optimize for nonnegative synaptic weights that invoke a desired somatic current J , and relax the optimization problem by taking subthreshold currents into account. We combined these methods with a specific surrogate model for H in the context of a two-compartment LIF neuron. Finally, we performed a series of spiking neural network simulations that show that our methods allow dendritic nonlinearities to be systematically exploited to efficiently approximate nonlinear multivariate functions

up to a certain spatial bandwidth.

While our approach is a step towards providing a general model of dendritic computation in top-down neurobiological modeling frameworks, it admittedly has several limitations. Most importantly, we treat the dendritic nonlinearity H as time-independent. Correspondingly, we implicitly assume that synaptic time-constants typically dominate the overall neuronal dynamics. However, dendritic trees in biology—especially when considering active channels and dendritic spikes (Koch, 1999)—possess filter properties and adaptation processes that are not accounted for in our model. It would be interesting to incorporate the dynamical properties of dendritic trees into the NEF by employing the recent techniques presented by Voelker and Eliasmith (2018).

A further shortcoming of the derivation of the surrogate model of H for the two-compartment neuron model is the assumption that the average somatic membrane potential is constant. While we are able to alleviate this assumption to some degree by fitting the model parameters to simulation data, the exact model parameters depend on the specific working-regime in which the neuron is used. Deviations from the modeled behavior are particularly apparent in situations with output firing rates smaller than ten spikes per second (cf. Figs 5A and 6). Correspondingly, the dendritic nonlinearity presented in this paper may not be a suitable model for brain areas featuring extremely low maximum firing rates. There are two potential ways to work around this limitation. First, it may be possible to include an input-dependent membrane potential term in the nonlinearity. Or, second, one could directly use a sampled model for H . While these approaches are compatible with the concept of dendritic nonlinearity as introduced above, they both increase the mathematical complexity of the weight optimization problem

to a point where strategies such as stochastic gradient descent are required. These techniques tend to have significantly weaker guarantees regarding finding an optimal solution compared to the convex quadratic programs employed in this paper.

In light of the above limitations, we would like to re-emphasize that, as stated in the introduction, our goal is not to provide a detailed mechanistic model of dendritic computation. Instead, we hope to provide a useful tool that captures essential aspects of dendritic computation—a nonlinear interaction between input channels—while being computationally cheap and mathematically tractable, but still grounded in biophysics. This helps to bridge the gap between purely abstract functional networks and more biophysically grounded mechanisms.

A potential application of our work outside of neurobiological modeling is programming neuromorphic hardware. Neuromorphic computers are inspired by neurobiological principles and promise to reduce the energy consumption of certain computational problems by several orders of magnitude compared to conventional computers (Boahen, 2017). Especially when considering mixed analogue-digital neuromorphic hardware systems, it should be possible to achieve a higher energy efficiency by implementing a more complex model neuron—such as the two-compartment LIF neuron discussed here—and performing local analog computation. Potential future work in this regard would be to validate our methods on a neuromorphic computing platform that implements dendritic trees, such as the *BrainScales 2* system (Schemmel et al., 2017).

Another line of future work is to consider arbitrary configurations of passive dendritic trees beyond the two-compartment LIF model. By applying Kirchhoff’s circuit laws, any passive dendritic tree configuration can be described as a linear dynamical

system. Correspondingly, it is possible to derive the dendritic nonlinearity H . It would be interesting to see whether it is still possible to relatively quickly optimize connection weights and in how far the number of compartments influences the computational power of the dendritic nonlinearity.

In conclusion, we believe that the methods proposed here provide a solid grounding for future work exploring both detailed biophysical mechanisms in the context of functional spiking networks, and improving neuromorphic methods for neural computation. We have shown how to cast the determination of connection weights in a functional network with conductance based synapses as an optimization problem with guaranteed convergence to the minimum. This optimization not only exploits known dendritic nonlinearities, but respects specifiable network topologies that conform to Dale’s Principle. The result are functional spiking networks with improved accuracy and biophysical plausibility using fewer neurons than competing approaches.

Acknowledgments and Funding

The authors would like to thank Aaron R. Voelker for his comments on earlier drafts of this paper, as well as his advice and constructive criticism regarding the conducted experiments. This work was supported by the Canada Research Chairs program (no grant number), NSERC Discovery grant 261453, and AFOSR grant FA9550-17-1-002.

References

- Alemi, A., Machens, C. K., Deneve, S., and Slotine, J.-J. (2018). Learning nonlinear dynamics in efficient, balanced spiking networks using local plasticity rules. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Bekolay, T., Bergstra, J., Hunsberger, E., DeWolf, T., Stewart, T. C., Rasmussen, D., Choo, X., Voelker, A. R., and Eliasmith, C. (2014a). Nengo: A python tool for building large-scale functional brain models. *Frontiers in Neuroinformatics*, 7(48).
- Bekolay, T., Bergstra, J., Hunsberger, E., DeWolf, T., Stewart, T. C., Rasmussen, D., Choo, X., Voelker, A. R., and Eliasmith, C. (2014b). Nengo: A Python tool for building large-scale functional brain models. *Frontiers in Neuroinformatics*, 7(48).
- Berzish, M., Eliasmith, C., and Tripp, B. (2016). Real-Time FPGA Simulation of Surrogate Models of Large Spiking Networks. In Villa, A. E., Masulli, P., and Pons Rivero, A. J., editors, *Artificial Neural Networks and Machine Learning – ICANN 2016*, pages 349–356. Springer International Publishing.
- Blouw, P., Choo, X., Hunsberger, E., and Eliasmith, C. (2018). Benchmarking Keyword Spotting Efficiency on Neuromorphic Hardware. Available from arXiv:1812.01739. Cited 7 February 2019.
- Boahen, K. (2017). A neuromorph’s prospectus. *Computing in Science Engineering*, 19(2):14–28.
- Bobier, B., Stewart, T. C., and Eliasmith, C. (2014). A unifying mechanistic

- model of selective attention in spiking neurons. *PLoS Computational Biology*, 10(6):e1003577.
- Boerlin, M. and Denève, S. (2011). Spike-Based Population Coding and Working Memory. *PLOS Computational Biology*, 7(2):1–18.
- Boerlin, M., Machens, C. K., and Denève, S. (2013). Predictive Coding of Dynamical Variables in Balanced Spiking Networks. *PLOS Computational Biology*, 9(11):1–16.
- Broomhead, D. S. and Lowe, D. (1988). Radial basis functions, multi-variable functional interpolation and adaptive networks.
- Capaday, C. and Van Vreeswijk, C. (2006). Direct control of firing rate gain by dendritic shunting inhibition. *Journal of integrative neuroscience*, 5(02):199–222.
- Capocelli, R. M. and Ricciardi, L. M. (1971). Diffusion approximation and first passage time problem for a model neuron. *Kybernetik*, 8(6):214–223.
- Choudhary, S., Sloan, S., Fok, S., Neckar, A., Trautmann, E., Gao, P., Stewart, T., Eliasmith, C., and Boahen, K. (2012). Silicon Neurons That Compute. In Villa, A. E. P., Duch, W., Érdi, P., Masulli, F., and Palm, G., editors, *Artificial Neural Networks and Machine Learning – ICANN 2012*, pages 121–128. Springer Berlin Heidelberg.
- Duggins, P. (2017). *Incorporating Biologically Realistic Neuron Models into the NEF*. Masters thesis.
- Duggins, P., Stewart, T. C., Choo, X., and Eliasmith, C. (2017). Effects of Guanfacine and Phenylephrine on a Spiking Neuron Model of Working Memory. *Topics in Cognitive Science*.

- Eliasmith, C. (2013). *How to Build a Brain: A Neural Architecture for Biological Cognition*. Oxford Series on Cognitive Models and Architectures. Oxford University Press.
- Eliasmith, C. and Anderson, C. H. (2003). *Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems*. MIT press, Cambridge, Massachusetts.
- Eliasmith, C., Gosmann, J., and Choo, X.-F. (2016). BioSpaun: A large-scale behaving brain model with complex neurons. Available from arXiv:1602.05220. Cited 21 March 2018.
- Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., and Rasmussen, D. (2012). A large-scale model of the functioning brain. *Science*, 338:1202–1205.
- Gabbott, P. L. A. and Somogyi, P. (1986). Quantitative distribution of GABA-immunoreactive neurons in the visual cortex (area 17) of the cat. *Experimental Brain Research*, 61(2):323–331.
- Gosmann, J. and Eliasmith, C. (2020). CUE: A unified spiking neuron model of short-term and long-term memory (accepted). *Psychological Review*.
- Gupta, A., Wang, Y., and Markram, H. (2000). Organizing principles for a diversity of gabaergic interneurons and synapses in the neocortex. *Science*, 287(5451):273–278.
- Hendry, S. and Jones, E. G. (1981). Sizes and distributions of intrinsic neurons incor-

- porating tritiated GABA in monkey sensory-motor cortex. *Journal of Neuroscience*, 1(4):390–408.
- Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366.
- Hunsberger, E. and Eliasmith, C. (2015). Spiking Deep Networks with LIF Neurons. Available from arXiv:1510.08829. Cited 21 March 2018.
- Hunsberger, E., Scott, M., and Eliasmith, C. (2014). The Competing Benefits of Noise and Heterogeneity in Neural Coding. *Neural Computation*, 26(8).
- Jonas, P., Major, G., and Sakmann, B. (1993). Quantal components of unitary EPSCs at the mossy fibre synapse on CA3 pyramidal cells of rat hippocampus. *The Journal of Physiology*, 472(1):615–663.
- Koch, C. (1999). *Biophysics of Computation: Information Processing in Single Neurons*. Oxford University Press.
- Komer, B. and Eliasmith, C. (2016). A unified theoretical approach for biological cognition and learning. *Current Opinion in Behavioral Sciences*, 11:14–20. Computational modeling.
- Kozlov, M., Tarasov, S., and Khachiyan, L. (1980). The polynomial solvability of convex quadratic programming. *USSR Computational Mathematics and Mathematical Physics*, 20(5):223–228.
- Kreutz-Delgado, K. (2015). Mean Time-to-Fire for the Noisy LIF Neuron-A Detailed

- Derivation of the Siegert Formula. Available from arXiv:1501.04032. Cited 16 May 2017.
- Kuo, P. D. and Eliasmith, C. (2005). Integrating behavioral and neural data in a model of zebrafish network interaction. *Biological Cybernetics*, 93(3):178–187.
- London, M. and Häusser, M. (2005). Dendritic Computation. *Annual Review of Neuroscience*, 28(1):503–532.
- MacNeil, D. and Eliasmith, C. (2011). Fine-tuning and the stability of recurrent neural networks. *PLoS ONE*, 6.
- Marr, D. (1969). A theory of cerebellar cortex. *The Journal of physiology*, 202(2):437–470.
- Mel, B. W. (1994). Information processing in dendritic trees. *Neural Computation*, 6(6):1031–1085.
- Mundy, A., Knight, J., Stewart, T. C., and Furber, S. (2015). An efficient SpiNNaker implementation of the Neural Engineering Framework. In *IJCNN*. Institute of Electrical and Electronics Engineers (IEEE).
- Neckar, A., Fok, S., Benjamin, B. V., Stewart, T. C., Oza, N. N., Voelker, A. R., Eliasmith, C., Manohar, R., and Boahen, K. (2019). Braindrop: A Mixed-Signal Neuromorphic Architecture With a Dynamical Systems-Based Programming Model. *Proceedings of the IEEE*, 107(1):144–164.
- Nicola, W. and Clopath, C. (2017). Supervised learning in spiking neural networks with FORCE training. *Nature Communications*, 8(1):2208.

- Parisien, C., Anderson, C. H., and Eliasmith, C. (2008). Solving the problem of negative synaptic weights in cortical models. *Neural Computation*, 20:1473–1494.
- Polsky, A., Mel, B. W., and Schiller, J. (2004). Computational subunits in thin dendrites of pyramidal cells. *Nature Neuroscience*, 7:621.
- Roth, A. and van Rossum, M. C. W. (2009). Modeling Synapses. In Schutter, E. D., editor, *Computational Modeling Methods for Neuroscientists*, pages 139–159. The MIT Press.
- Schemmel, J., Kriener, L., Müller, P., and Meier, K. (2017). An accelerated analog neuromorphic hardware system emulating NMDA- and calcium-based non-linear dendrites. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 2217–2226.
- Schwemmer, M. A., Fairhall, A. L., Denéve, S., and Shea-Brown, E. T. (2015). Constructing Precisely Computing Networks with Biophysical Spiking Neurons. *Journal of Neuroscience*, 35(28):10112–10134.
- Stellato, B., Banjac, G., Goulart, P., Bemporad, A., and Boyd, S. (2020). OSQP: An operator splitting solver for quadratic programs. *Mathematical Programming Computation*.
- Stewart, T. C., Bekolay, T., and Eliasmith, C. (2012). Learning to select actions with spiking neurons in the basal ganglia. *Frontiers in Decision Neuroscience*, 6.
- Strata, P. and Harvey, R. (1999). Dale’s principle. *Brain Research Bulletin*, 50(5):349–350.

- Stöckel, A., Voelker, A. R., and Eliasmith, C. (2017). Point Neurons with Conductance-Based Synapses in the Neural Engineering Framework. Available from arXiv:1710.07659. Cited 20 Oct 2017.
- Sussillo, D. and Abbott, L. (2009). Generating Coherent Patterns of Activity from Chaotic Neural Networks. *Neuron*, 63(4):544–557.
- Thalmeier, D., Uhlmann, M., Kappen, H. J., and Memmesheimer, R.-M. (2016). Learning universal computations with spikes. *PLOS Computational Biology*, 12(6):1–29.
- Tripp, B. (2009). *A Search For Principles of Basal Ganglia Function*. Phd thesis, University of Waterloo.
- Tripp, B. and Eliasmith, C. (2007). Neural populations can induce reliable postsynaptic currents without observable spike rate changes or precise spike timing. *Cerebral Cortex*, 17:1830–1840.
- Vandenberghe, L. (2010). The CVXOPT linear and quadratic cone program solvers.
- Voelker, A. R. and Eliasmith, C. (2018). Improving Spiking Dynamical Networks: Accurate Delays, Higher-Order Synapses, and Time Cells. *Neural Computation*, 30(3):569–609.
- Vu, E. T., Lee, S. C., and Krasne, F. B. (1993). The mechanism of tonic inhibition of crayfish escape behavior: Distal inhibition and its functional significance. *Journal of Neuroscience*, 13(10):4379–4393.
- Wright, S. and Nocedal, J. (1999). Numerical optimization. *Springer Science*, 35(67-68):438–484.