

Point neurons with conductance-based synapses in the Neural Engineering Framework

Andreas Stöckel, Aaron R. Voelker, and Chris Eliasmith

Centre for Theoretical Neuroscience, University of Waterloo

September 8, 2017

Abstract

The mathematical model underlying the Neural Engineering Framework (NEF) expresses neuronal input as a linear combination of synaptic currents. However, in biology, synapses are not perfect current sources and are thus nonlinear. Detailed synapse models are based on channel conductances instead of currents, which require independent handling of excitatory and inhibitory synapses. This, in particular, significantly affects the influence of inhibitory signals on the neuronal dynamics. In this technical report we first summarize the relevant portions of the NEF and conductance-based synapse models. We then discuss a naïve translation between populations of LIF neurons with current- and conductance-based synapses based on an estimation of an average membrane potential. Experiments show that this simple approach works relatively well for feed-forward communication channels, yet performance degrades for NEF networks describing more complex dynamics, such as integration.

1 Introduction

As a central assumption, the Neural Engineering Framework (NEF, [1]) relies on linear interaction of pre-synaptic activity in the form of synaptic currents injected into the neuron membrane. This assumption is readily fulfilled by current-based synapse models. Biological synapses however are not perfect current sources. Instead, currents injected into the cell membrane are caused by changes in receptor channel conductance directly or indirectly gated by neurotransmitters released in the pre-synapse. Consequently, current flow into the membrane is modulated both by pre-synaptic spikes, as well as—due to the nature of ion channels—the difference between the neuron membrane potential and the

specific ion species reversal potential. This nonlinear interdependence is modelled by conductance-based synapse models [2, 3].¹

For various reasons, conductance-based synapses are a useful addition to the NEF. Foremost, when modelling functional, biologically plausible neural networks, more realistic synapse models pose additional constraints on the network design and reduce degrees of freedom. For example, conductance-based synapses strictly differentiate between excitatory and inhibitory synapses, allowing the network designer to map their parameters to those derived from neurophysiological evidence. Furthermore, some neuromorphic hardware systems (e.g. Spikey, BrainScaleS; [5, 6]) exclusively implement conductance-based synapses, which must be taken into account when assembling networks for these platforms. Finally, nonlinear synaptic interaction has been argued to be an important element of biological neural networks, and as a site for biological computation [2]. Exploiting these nonlinearities in the NEF increases the computational power per neuron. In the future, conductance-based synapses may thus be utilized to increase the efficiency of certain mathematical operations such as multiplication, requiring fewer neurons at the same precision.

The remainder of this report is structured as follows: section 2 gives a short overview of the relevant equations in the NEF. Section 3 summarizes the biological background of conductance-based synapses and describes the underlying theoretical model. Section 4 elaborates on the actual integration of conductance-based synapses into the NEF and provides an equation highly similar to the original NEF formalism, yet containing additional, voltage-dependent scaling factors. Section 5 provides estimations for these factors based on the expected value for the average membrane potential. Section 7 presents two benchmark experiments, followed by concluding remarks in section 8.

2 Current sources in the Neural Engineering Framework

The NEF is a systematic approach to the representation of low-dimensional vectors in populations of spiking neurons and the approximation of arbitrary functions in the connections between these populations [1]. As an example, consider two populations of n spiking neurons with all-to-all connectivity from the first (pre-synaptic) to the second (post-synaptic) population. Let $\vec{x}(t)$ be the value represented by the pre-synaptic population, and $\vec{y}(t)$ the value represented by the post-synaptic population. Given a function f , the NEF allows us to compute connection weights w_{ij} such that the resulting activity in the post-population represents $\vec{y}(t) \approx f(\vec{x}(t))$.

¹Not to be confused with conductance-based neuron models, such as the Hodgkin-Huxley model of action potential generation [4, 2].

The formalism underlying the NEF assumes that each neuron j receives its input in the form of a current $J(t)$ injected into its (virtual) neuron membrane.² This current is the sum of a constant bias current J^{bias} , and a synaptic current $J^{\text{syn}}(t)$. The latter is calculated from a weighted sum of filtered input spike events: spikes produced by the i th pre-synaptic neuron are modelled as a sum of Dirac pulses $a_i(t)$, and a filter $h(t')$ defines the shape of the post-synaptic current in response to a single input spike:

$$J(t) = J^{\text{bias}} + J^{\text{syn}}(t) = J^{\text{bias}} + \sum_{i=1}^n w_{ij} (a_i * h)(t). \quad (1)$$

Central to the NEF is the idea of factorizing the connection weight matrix W into gain factors α_j , decoding vectors \vec{d}_i , and unit-length encoding vectors \vec{e}_j , where $w_{ij} = \alpha_j \langle \vec{e}_j, \vec{d}_i \rangle$. Correspondingly, eq. (1) can be written as

$$J(t) = J^{\text{bias}} + J^{\text{syn}}(t) = J^{\text{bias}} + \alpha_j \vec{e}_j^T \left(\sum_{i=1}^n \vec{d}_i (a_i * h)(t) \right). \quad (2)$$

The decoders \vec{d}_i linearly project the high-dimensional activity of the pre-synaptic population onto the desired low-dimensional representation $\vec{y}(t) = f(\vec{x}(t))$. Multiplication with unit-length encoders \vec{e}_j corresponds to a projection of this low-dimensional representation onto the preferred direction of each post-synaptic neuron j . Correspondingly, the transformed value is again represented as a high-dimensional population coding, allowing to chain an arbitrary number of neuron populations and transformations.

As a side-effect, factorization of the weight matrix W into decoders and encoders simplifies the calculation of w_{ij} . Instead of optimizing weights w_{ij} individually, we can randomly select bias currents J_j^{bias} , gain factors α_j , and encoders \vec{e}_j , and calculate decoders \vec{d}_j from a least squares linear regression

$$Y = DA \Rightarrow D \approx YA^T (AA^T + I\sigma^2)^{-1}, \quad (3)$$

where the matrix D is composed of the decoding vectors, A is a matrix of neuron responses in the post-synaptic population for a set of sample points, and Y is a matrix of the desired function outputs for these points. The expression $I\sigma^2$ is a regularization term which accounts for uncertainty and noise in the neuron response.

Crucially, the above equations contain the implicit assumption that the input current $J(t)$ —and in turn the effect of each individual input spike—is independent of the neuron state at time t . Furthermore, all input spikes are treated equally (except for the scaling factor w_{ij}), and there is no distinction between excitatory and inhibitory synapses. Having

²Note that $J(t)$ (as well as other currents and conductances) should be read with a subscript j in mind – all equations, unless explicitly stated otherwise, refer to the j th neuron in the post-synaptic population.




Figure 1: Equivalent circuit diagram of a LIF neuron with current-based synapses (subthreshold behaviour; without spike generation/reset mechanism).

synaptic currents which are independent of the neuron state, is commonly referred to as using “current-based synapses” [2, 7].

The subthreshold circuit diagram of a LIF neuron with current-based synapses is depicted in fig. 1. The dynamical system describing this circuit is³

$$C_m \frac{dv(t)}{dt} = g_L(E_L - v(t)) + J^{\text{bias}} + J^{\text{syn}}(t), \quad (4)$$

which has a stable attractor (the equilibrium potential E_{eq}) at

$$E_{\text{eq}}(t) = \frac{J^{\text{bias}} + J^{\text{syn}}(t)}{g_L} + E_L. \quad (5)$$

The synaptic current $J^{\text{syn}}(t)$ drives the membrane potential $v(t)$ towards more positive values for input spikes received at excitatory synapses ($w_{ij} > 0$) and towards more negative values for spikes received at inhibitory synapses ($w_{ij} < 0$). Importantly, the equilibrium potential E_{eq} in eq. (5) can be driven to arbitrary values if the input currents J are large enough. In practice this is not an issue for positive J , as the maximum membrane potential is limited by the neuron threshold potential v_{th} . However, such a natural limitation does not exist for negative currents, and the NEF puts no constraint on the negative current strength, for neither J^{bias} , nor J^{syn} . A solution to this problem is to modify the LIF model, such that the membrane potential is clamped to a minimum value⁴, i.e., the reset potential v_{reset} . Note that all experiments in this report are based on LIF neurons with such a minimum membrane potential.

This problem does not occur for conductance-based synapses. Conductance-based synapses are not only biologically motivated, but eliminate the physically implausible,

³The ratio between the membrane capacitance and the leak channel conductance is the membrane time constant $\tau_{\text{RC}} = C_m/g_L$.

⁴This is for example implemented in the LIF neuron model in the software neuron simulator included in Nengo, the reference software implementation of the NEF.




Figure 2: Illustration of the ion channel reversal potential caused by equalling osmotic and electric forces. Although the intra- and extracellular fluids are electrically neutral (a), an open, selectively permeable ion channel causes a change in the membrane potential (b). Illustration adapted from [8]. See [2, 9] for more details.

perfect current sources altogether. Unfortunately, and what will be the focus of the remainder of this report, conductance-based synapses introduce a dependency between the neuron membrane potential $v(t)$, and the injected current $J^{\text{syn}}(t)$. As a result, the neuron response is no longer linear with respect to the pre-synaptic spikes: consider two excitatory input spikes which do not trigger an action potential in the receiving neuron. Due to the increased membrane potential $v(t)$, the post synaptic current (PSC) induced by the second spike is smaller than the PSC induced by the first spike; the PSCs do not sum linearly. Furthermore, conductance based synapses require excitatory and inhibitory synapses to be handled differently. Both issues are not accounted for in the standard NEF eqs. (1) and (2).

3 Conductance-based synapses

In biology, changes in the neuron membrane potential are primarily caused by variations in the permeability—or conductance—of selective ion channels embedded into the lipid bilayer cell membrane [2, 9]. As depicted in fig. 2a, ions are dissolved in different concentrations in the intra- and extracellular fluid. Ion channels provide a way for a specific ion species to follow the osmotic concentration gradient, which in turn causes a charge imbalance and an electrical force countering the osmotic force (fig. 2b). The charge imbalance itself is measurable as a change in the membrane potential $v(t)$. Once a certain membrane potential is reached, the osmotic and electric forces cancel and the ion




Figure 3: Equivalent circuit diagram of a LIF neuron with conductance-based synapses (without spike generation/reset mechanism), see also [2].

flow stops. Further increase/decrease of $v(t)$ reverses the ion flow through a particular channel, giving rise to the name *reversal potential* for this particular $v(t)$.⁵ The exact value of the reversal potential differs between ion channel types. For example, in mammalian cells, the reversal potential for a sodium channel is $E_{Na^+}^+ \approx 67$ mV, while for potassium the reversal potential is $E_{K^+}^+ \approx -83$ mV. In the resting state, the membrane is slightly permeable for a mixture of ions, resulting in an equilibrium potential⁶ of $E_L \approx -65$ mV [3, 10].

An equivalent circuit diagram of a cell membrane popularized by Hodgkin and Huxley [4] (though in the context of action potential generation, and not synaptic input) with leak, sodium and potassium channel is shown in fig. 3. Post-synaptic currents can be modeled as a change of permeabilities in the Na^+ and K^+ ion channel.⁷ For example, excitatory input spikes could cause the sodium channels to be open for a certain period of time, driving the membrane potential towards E_{Na^+} . Conversely, inhibitory input spikes might cause potassium channels to open, driving the membrane potential towards E_{K^+} . In the following, we refer to E_{Na^+} as E_E (the excitatory reversal potential) and to E_{K^+} as E_I (the inhibitory reversal potential). The dynamics for conductance-based synapses in response to incoming spikes $a(t)$ can be described with a linear filter $h(t')$, where the channel conductance $g(t)$ is $g(t) = (a * h)(t)$. Possible filters include simple exponential filters, the alpha function, or the difference of exponentials. Of course, in biology, channel dynamics involve the release of a neurotransmitter, diffusion of the transmitter through the synaptic cleft, and a chemical cascade causing ion channels to open as soon as the transmitter reaches the post-synaptic neuron. More detailed models of these

⁵Since ion diffusion is a stochastic process, ions will always diffuse in and out of the cell as long as the channel is open. The above claims holds for the average ion flow only.

⁶To clarify the difference between reversal and equilibrium potential: the reversal potential is the potential at which the flow of a *single* ion species reverses direction, the equilibrium potential is the potential at which the sum of *all* currents is zero.

⁷In biology, a large number of channels are embedded into the neuron cell membrane. Individual ion channels are either fully open or fully closed, following a stochastic process. The conductance $g(t)$ can be interpreted as the probability of individual channels to be open.

highly nonlinear processes exist [7]. Nevertheless, in the remainder of this report we assume the synaptic dynamics to be describable as a linear filter, and in particular focus on normalized exponential low-pass filters with time constant τ_{syn} where

$$h(t) = \begin{cases} \frac{1}{\tau_{\text{syn}}} \exp(-t/\tau_{\text{syn}}) & t \geq 0 \\ 0 & t < 0 \end{cases}, \quad \text{and} \quad \frac{dg(t)}{dt} = -\frac{1}{\tau_{\text{syn}}} \left(g(t) + \sum_{i=1}^n w_{ij} a_i(t) \right).$$

Note that while the synaptic dynamics are linear with respect to the input, the post synaptic current is not (see section 2).

The differential equation describing a LIF point neuron with conductance-based channels is

$$C_m \frac{dv(t)}{dt} = g_L \cdot (E_L - v(t)) + g_E(t) \cdot (E_E - v(t)) + g_I(t) \cdot (E_I - v(t)). \quad (6)$$

The equilibrium potential $E_{\text{eq}}(t)$ for this model is a weighted average of the individual reversal potentials

$$E_{\text{eq}}(t) = \frac{g_L \cdot E_L + g_E(t) \cdot E_E + g_I(t) \cdot E_I}{g_L + g_E(t) + g_I(t)}. \quad (7)$$

This implies that—in contrast to current-based models, which, as noted at the end of the previous section, can provide arbitrarily large input—the membrane potential can never be driven to values outside the voltage range spanned by E_L , E_E , and E_I , or, in biological terms, the ion channel reversal potentials.

4 Conductance-based synapses in the NEF

One of the initial challenges when implementing conductance-based synapses in the NEF is the different handling of excitatory and inhibitory synapses. Spikes arriving at excitatory synapses will influence the excitatory ion channel conductance, whereas spikes arriving at inhibitory synapses influence the inhibitory ion channel conductance.

This issue is virtually nonexistent when adhering to Dale's Principle [11]. Each pre-population of neurons will either project with excitatory or inhibitory connections to the post-population, and influence one of the two channel types only.

In the general case, the weight matrix W (section 2) can be split into two non-negative matrices W^+ and W^- . The matrix W^+ contains positive, excitatory weights, whereas W^- contains negative, inhibitory weights with all other weights set to zero:

$$w_{ij}^+ = \max\{0, w_{ij}\}, \quad w_{ij}^- = \max\{0, -w_{ij}\}.$$

This non-negative decomposition of W unfortunately ruins the factorization into encoders and decoders employed by the NEF. For two reasons, this is not a real issue. First, decomposition happens after encoders and decoders have already been computed, maintaining the entire NEF pipeline. Second, to increase the computational efficiency during simulation, artificial decoder and encoder matrices E^\pm , D^\pm can be calculated, e.g., by performing a singular value decomposition of W^+ and W^- :

$$W^\pm = USV, \quad \vec{E}^\pm = U[k], \quad \vec{D}^\pm = V[k] \cdot S[k].$$

Here, k is the number of non-zero eigenvalues or, alternatively, eigenvalues above a certain threshold, and $M[k]$ refers to the first k columns in the matrix M . Under most circumstances, this factorization yields little to no error, since the weight matrix W is an outer product and thus of low rank. Usually, k is equal to the number of dimensions represented by the population plus one (caused by clamping the weights).

As mentioned in the previous section, the channel conductance time-dynamics $g_E(t)$ and $g_I(t)$ can be modeled as a linear filter of the input spike trains. Essentially, we can just reuse the post-synaptic current filter $h(t)$ from eq. (1) and apply it to the individual channel conductances instead. Assuming (as in section 2 and without loss of generality) that all synapses have the same synaptic filter $h(t)$, the conductances $g_E(t)$ and $g_I(t)$ for the j th neuron are

$$g_E(t) = \sum_{i=1}^n w_{ij}^+ \cdot (a_i * h)(t), \quad g_I(t) = \sum_{i=1}^n w_{ij}^- \cdot (a_i * h)(t). \quad (8)$$

As per eq. (6) the synaptic current $J^{\text{syn}}(t)$ injected into a single neuron j with conductance-based synapses is


$$J^{\text{syn}}(t) = g_E(t) \cdot (E_E - v(t)) + g_I(t) \cdot (E_I - v(t)). \quad (9)$$

Combing eqs. (8) and (9) yields


$$J^{\text{syn}}(t) = \sum_{i=1}^n \left(w_{ij}^+ \cdot (E_E - v(t)) + w_{ij}^- \cdot (E_I - v(t)) \right) \cdot (a_i * h)(t). \quad (10)$$

This equation is almost equivalent to eq. (1), except for the dynamic scaling factors $\alpha_E(t) = E_E - v(t)$ and $\alpha_I(t) = E_I - v(t)$. For biologically plausible neuron parameters $v(t) \ll E_E$, so one might argue that $\alpha_E(t)$ is almost constant. However, for E_I , the difference between the reset or resting potential—if any—is just a few millivolts, causing a strong interdependence between the inhibitory current and $v(t)$ [7].

These potential deviations for inhibitory synapses are supported by the experiments




(a) Response curves




(b) Response curve difference

Figure 4: Neuron response curves $a^{\text{cur}}(\lambda_E)$, $a^{\text{cond}}(\lambda_E)$ of a LIF neuron with current-based or conductance-based synapses under varying inhibitory bias currents J^{bias} . The excitatory input spike train is Poisson-distributed with rate λ_E over 10s. Lines represent the mean over a thousand trials, shaded areas the standard deviation.



(a) Response curve



(b) Response curve difference

Figure 5: Neuron response curves $a^{\text{cur}}(\lambda_E)$, $a^{\text{cond}}(\lambda_E)$ of a LIF neuron with current-based or conductance-based synapses for an input spike train arriving at an inhibitory synapse. Both, the excitatory and the inhibitory input spike trains are Poisson-distributed with rates λ_E , λ_I over 10s. Lines represent the mean over a thousand trials, shaded areas the standard deviation.

Table 1: Neuron and synapse parameters used throughout this report unless specified otherwise. “Normalized values” refers to the scaled LIF model commonly used in neuron simulators such as Nengo. Parameters for the conductance-based synapses are scaled according to the linear transition model in section 5, eq. (11). Synaptic filters are assumed to be normalized to unit area – whenever this is not the case divide the synaptic weights by τ_{syn} for the denormalized values.

Neuron and synapse parameters					
<i>Neuron parameters</i>					
<i>Parameter</i>		<i>Value</i>		<i>Normalized value</i>	
Membrane capacitance	C_M	1.00	nF	1.00	
Membrane time constant	τ_{RC}	20.00	ms	20.00	ms
Leak conductivity	g_L	50.00	nS	50.00	s^{-1}
Resting potential	E_L	-65.00	mV	0.00	
Reset potential	v_{reset}	-65.00	mV	0.00	
Threshold potential	v_{th}	50.00	mV	1.00	
Refractory period	τ_{ref}	2.00	ms	2.00	ms
<i>Current-based synapses</i>					
<i>Parameter</i>		<i>Value</i>		<i>Normalized value</i>	
Synaptic time constant	τ_{syn}	5.00	ms	5.00	ms
Excitatory weight	w^+	35.00	pA	2.33	10^{-9}
Inhibitory weight	w^-	-35.00	pA	-2.33	10^{-9}
<i>Conductance-based synapses</i>					
<i>Parameter</i>		<i>Value</i>		<i>Normalized value</i>	
Synaptic time constant	τ_{syn}	5.00	ms	5.00	ms
Excitatory reversal potential	E_E	0.00	mV	4.33	
Inhibitory reversal potential	E_I	-80.00	mV	-1.00	
Excitatory weight	w^+	0.61	nS	40.60	10^{-9}
Inhibitory weight	w^-	1.56	nS	103.70	10^{-9}

presented in figs. 4 and 5. In each individual experiment, a single LIF neuron with either current or conductance-based synapses is simulated (see table 1). Parameters for the conductance-based synapses are matched to the current-based synapses with the linear transition model discussed in section 5, eq. (11), so the neuron should behave similarly in both cases. The neuron receives an excitatory input spike train with Poisson statistics and rate λ_E . In fig. 4 an inhibitory bias current J^{bias} is injected into the neuron, whereas in fig. 5 the neuron receives spikes at an inhibitory synapse with rate λ_I . The inhibitory currents/rates were matched such that the average neuron rate surpasses zero for approximately the same excitatory input (the x -intercepts are in similar locations). Strikingly, for excitatory input only ($J^{\text{bias}} = 0$, $\lambda_I = 0$), there are only minor deviations between the neuron response curves. For an inhibitory J^{bias} , the neuron with conductance-based synapses has a slightly smaller output rate than the neuron with current-based synapses,

though this difference converges to zero for high output rates. However, if the inhibitory bias current is replaced by an inhibitory synaptic input, the deviations from the neuron with current-based synapses are much stronger: the output rate is significantly reduced in the presence of inhibitory synaptic input, especially for large input spike rates. **??** shows the average current flowing into the membrane in the above experiments. For constant inhibitory bias current, the current increases almost linearly with the excitatory spike rate, whereas inhibitory spikes cause a saturation of the input current. This non-linearity is likely caused by inhibitory spikes injecting a larger negative current in the presence of excitatory input due to the increased average membrane potential.

5 Average membrane potential estimation

There are at least two possible approaches regarding the integration of conductance-based synapses into the NEF. A “natural” approach would be to measure the actual neuron response curves A as required for the calculation of decoders in eq. (3). However, the neuron response is no longer defined in terms of a current J but as a two-dimensional function of g_E and g_I , which slightly complicates matters, and is not the focus of this report.

Alternatively, we can try to transform a network of neurons with current-based synapses into networks of neurons with conductance-based synapses. Again, this assumes that the switch to conductance-based synapses does not have a strong effect on the neuronal dynamics. For this approach, we calculate weights W as if the network was still using current-based synapses, but rescale the decomposed W^+ , W^- linearly to account for $\alpha_E(t)$ and $\alpha_I(t)$. Instead of treating these factors as functions time, we estimate expected values $\langle \alpha_E(t) \rangle_t$ and $\langle \alpha_I(t) \rangle_t$. This can be achieved by replacing $v(t)$ with $\langle v(t) \rangle_t$. Clearly, a “one size fits all” average membrane potential does not exist. However, in the NEF neurons are often in either a mode of (semi) regular spiking or a do not spike at all for longer periods of time.

A naïve estimate for $\langle v(t) \rangle_t$ in the spiking mode is to assume a linear transition between v_{reset} and v_{th} . In this case, the average membrane potential is

$$\langle v(t) \rangle_t = \frac{v_{\text{reset}} + v_{\text{th}}}{2}. \quad (\text{Linear transition}) \quad (11)$$

However, the membrane potential transition for a neuron is not linear. The force driving the membrane potential towards E_L gets larger with higher $v(t)$, whereas the force driving towards E_E gets smaller (for conductance-based neurons). Correspondingly, and as illustrated in fig. 6, $v(t)$ progresses more slowly through high voltage regimes, causing the actual $\langle v(t) \rangle_t$ to be larger than the prediction made by the linear model.




Figure 6: Idealized membrane potential traces. Comparison of the membrane potential transition from $v_{\text{reset}} = -65 \text{ mV}$ to $v_{\text{th}} = -50 \text{ mV}$ for a LIF neuron with conductance-based synapses, a LIF neuron with current-based synapses, and a simple linear model. The left and right graphs show different transition times (9 ms and 27 ms). Conductances/currents were assumed to be constant. Dotted lines show the average for each of the three models.

Under the assumption of fast enough regular input spikes we can derive a better approximation for the expected value. For high input rates $g_E(t)$ and $g_I(t)$ are approximately constant. Specifically, for an infinite input spike train $a(t)$ with constant inter-spike interval Δt and a normalized exponential low pass filter $h(t')$ with time constant τ , it holds (see appendix A for the derivation)

$$\langle g(t) \rangle_t = \langle (a * h)(t) \rangle_t = \frac{1}{\Delta t}.$$

For constant conductance, eq. (6) can be written as

$$\begin{aligned} C_m \frac{dv(t)}{dt} &= g_L \cdot (e_L - v(t)) + g_E \cdot (e_E - v(t)) + g_I \cdot (e_I - v(t)) + J^{\text{bias}} \\ &= -v(t) \cdot (g_L + g_E + g_I) + g_L e_L + g_E e_E + g_I e_I + J^{\text{bias}}, \end{aligned}$$

which is brought into a fairly general form by substituting

$$\begin{aligned} \lambda &= \frac{g_L + g_E + g_I}{C_m}, \\ \mu &= \frac{e_L g_L + e_E g_E + e_I g_I + J^{\text{bias}}}{C_m}. \end{aligned}$$

The closed-form solution for an initial value of $v(0) = v_0$ is

$$\frac{dv(t)}{dt} = -\lambda v(t) + \mu, \quad v(t) = \left(v_0 - \frac{\mu}{\lambda}\right) \cdot e^{-\lambda t} + \frac{\mu}{\lambda}.$$

If the synaptic current suffices to generate output spikes, the time-to-spike t_{th} after which the threshold potential v_{th} is reached is given as

$$t_{\text{th}} = -\frac{1}{\lambda} \log\left(\frac{v_{\text{th}} - \frac{\mu}{\lambda}}{v_0 - \frac{\mu}{\lambda}}\right). \quad (12)$$

The average membrane potential in the period between two spikes can now be calculated by setting $v_0 = v_{\text{reset}}$, integrating over the above equation, and dividing by t_{th} :

$$\langle v(t) \rangle_t = \frac{1}{t_{\text{th}}} \int_0^{t_{\text{th}}} v(t) dt = \frac{v_0 - \frac{\mu}{\lambda}}{\lambda t_{\text{th}}} \left(1 - e^{-\lambda t_{\text{th}}}\right) + \frac{\mu}{\lambda}$$


For $g_E \rightarrow \infty$ the expected value $\langle v(t) \rangle_t$ is

$$\langle v(t) \rangle_t = \frac{\log\left(\frac{E_E - v_{\text{th}}}{E_E - E_L}\right) \cdot E_E - E_L + v_{\text{th}}}{\log\left(\frac{E_E - v_{\text{th}}}{E_E - E_L}\right)}. \quad (\text{Conductance-based}) \quad (13)$$


Explicitly, these calculations do not include the refractory period τ_{ref} . The LIF neuron model assumes all currents flowing into the cell membrane to be zero during the refractory state. As a consequence, conductance-based synapses do not change the behavior of the model during the refractory period, so the average membrane potential calculation should be independent of τ_{ref} .

To give a sense for the difference between the two estimates discussed so far: for the neuron parameters listed in table 1 the linear transition model eq. (11) estimates an average value of -57.5 mV, whereas the conductance-based transition model estimates a slightly higher value of -57.17 mV for an infinitely strong excitatory input.

Figure 7 shows numerical measurements of the average membrane potential for the experimental setup presented in the previous section. For a constant bias current J^{bias} and no inhibitory spikes, the membrane potential converges to the estimate from eq. (13) as λ_E increases. In the presence of inhibitory input spikes, $\langle v(t) \rangle_t$ is consistently larger than in the J^{bias} experiment and reaches values clearly above the estimate. Unintuitively, and despite output-rate deviations being higher (figs. 4 and 5), the mean error between the estimated average membrane potential and the numerically measured value is consistently smaller when using inhibitory spikes instead of a bias current J^{bias} . This is especially true for high λ^I in comparison to high J^{bias} (i.e. 5.5 mV RMSE for $\lambda^I = 1000 \text{ s}^{-1}$, 6.5 mV RMSE for $J^{\text{bias}} = -30 \text{ nA}$). Consequently, it might be worthwhile to investigate the effect of not



(a) Inhibitory J^{bias}



(b) Inhibitory input spikes

Figure 7: Estimate and measured average membrane potentials for varying excitatory spike rates λ_E and either a constant bias current J^{bias} or an inhibitory spike train with rate λ_I . Lines represent the mean over a thousand trials, shaded areas the standard deviation.

using an external current source J^{bias} in the NEF altogether without negatively affecting the variety of tuning curves. The next section contains a simple technique achieving exactly this by decoding J^{bias} from the pre-population.

6 Elimination of the bias current

As pointed out in eq. (2), the NEF generally defines the current injected into the membrane $J(t)$ of a neuron j as $J(t) = J^{\text{bias}} + J^{\text{syn}}(t)$. For a population of n neurons, this can be written in matrix notation

$$\vec{J}(t) = \vec{J}^{\text{bias}} + \text{diag}(\vec{\alpha}) \cdot ED \cdot (\vec{a} * h)(t) = \vec{J}^{\text{bias}} + W \cdot (\vec{a} * h)(t),$$

where \vec{J}^{bias} and \vec{a} are n -dimensional vectors, E is the $n \times d$ encoding matrix (d being the dimensionality of the vector represented by the population), D is $d \times m$ decoding matrix (m being the size of the pre-population), W is the $n \times m$ weight matrix, and $(\vec{a} * h)(t)$ is the filtered, m dimensional activity of the pre-population. We can approximate this equation as

$$\vec{J}(t) \approx (D' + \text{diag}(\vec{\alpha}) \cdot ED) \cdot (\vec{a} * h)(t) = W' \cdot (\vec{a} * h)(t),$$

where D' is a $n \times m$ matrix decoding the vector \vec{J}^{bias} from the pre-population. Analogously to the L_2 approximation in eq. (3), we can approximate D' as

$$Y' = \underbrace{(\vec{J}^{\text{bias}}, \dots, \vec{J}^{\text{bias}})}_{m \text{ times}} = D' A \Rightarrow D' = Y' A^T (A A^T + I \sigma^2)^{-1},$$

where Y' is the $n \times N$ dimensional target matrix (with N being the number of evaluation points), and A the $n \times N$ neuron response matrix. Using the adapted weight matrix $W' = D' + W$ instead of J^{bias} should not significantly alter the behaviour of the system. As discussed in section 4, we can split W' into excitatory weights W'^+ and inhibitory weights W'^- for use with conductance-based synapses.

7 Benchmark experiments


In this section we present two benchmark networks: a feed-forward communication channel, and an integrator. The first network tests how well one can decode a function from a conductance-based synapse population with respect to pure feed-forward processing and a constant stream of input. The second network assess how well dynamical systems are realized. We then analyze a potential cause of unexpected behaviour in the recurrent

network by empirically measuring neuron tuning curves.

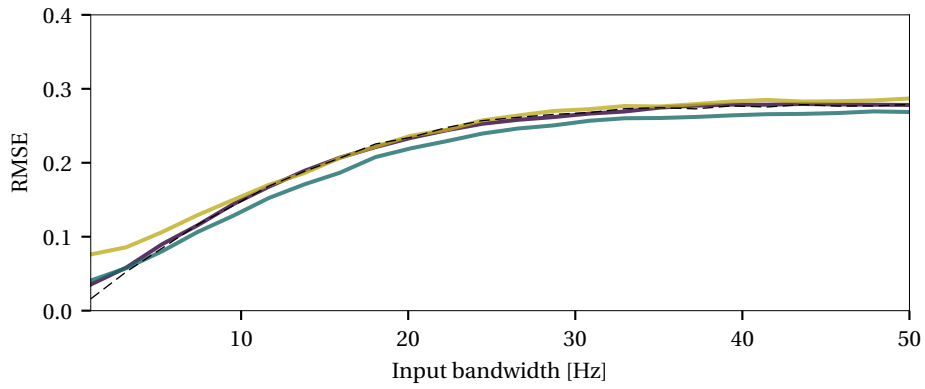
Communication channel experiments The communication channel experiment consists of a chain of ℓ neuron ensembles with one hundred neurons each. All connections compute the identity function. The input to the system is a bandwidth limited white noise signal. As a benchmark measure, we compare the decoded output of the last ensemble in the chain to the input value. The experiment is executed for standard LIF neurons, LIF neurons with conductance-based synapses and eliminated J^{bias} (see section 6), as well as LIF neurons with conductance-based synapses and active J^{bias} . Results are shown in fig. 8.

As is clearly visible, there is no significant difference in error for the three neuron types for short communication channels. Notably however, across all experiments, current-based synapses are always best for inputs with low bandwidth, and using J^{bias} in conjunction with conductance-based synapses increases the error. Apart from this, conductance-based synapses are rather competitive compared to current-based synapses, and as the bandwidth increases, the error is mostly dominated by the baseline error obtained when filtering the input signal with the synaptic filter. Interestingly, current-based synapses are only better in high bandwidth scenarios when increasing the communication channel length to eight. At least the error caused by J^{bias} can be explained with the observations laid out in section 5. Specifically, the average membrane potential deviates from the estimated values for a large input spike rate range.


Integrator experiments In the integrator experiment, a standard NEF integrator with 100 ms synaptic time constant for the recurrent connection is tested. A value $x \in [0, 1]$ is fed into the integrator for one second, theoretically causing the integrator to converge to x . The simulation is continued for nine seconds and the difference between x and the actual value is recorded. Figure 9 shows the results over varying integrator neuron counts. For standard current-based synapses, the error is relatively small and converges to zero with increasing neuron count. The integrator with conductance-based synapses exhibits much larger errors. The error does not significantly decrease with increasing neuron count and is significantly higher when not eliminating J^{bias} . When recording the integrator value trace over time, the integrator often converges to a value of one instead of the desired value, hinting at an additional strong attractor in the underlying dynamical system. Figure 9b depicts the average value the integrator converges to. While the integrator with current-based synapses on average converges to the value 0.5, as expected when sampling uniformly from the interval $[0, 1]$, the integrator with conductance-based synapses converges to significantly larger values.



(a) $n = 2$

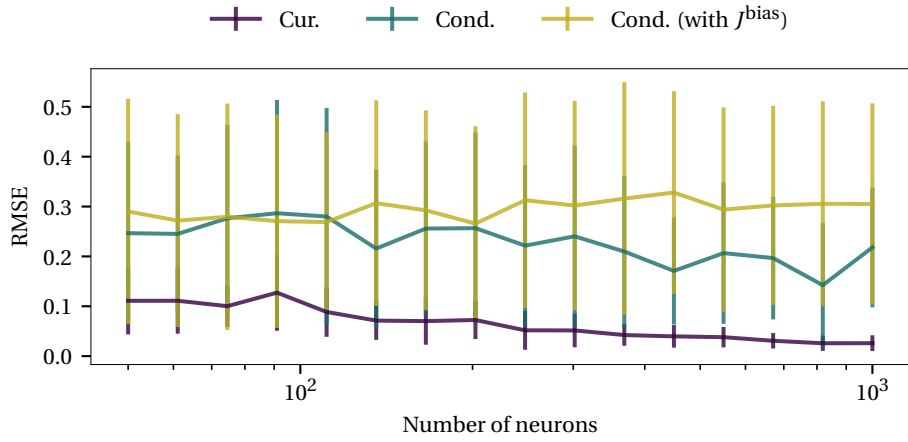


(b) $n = 4$

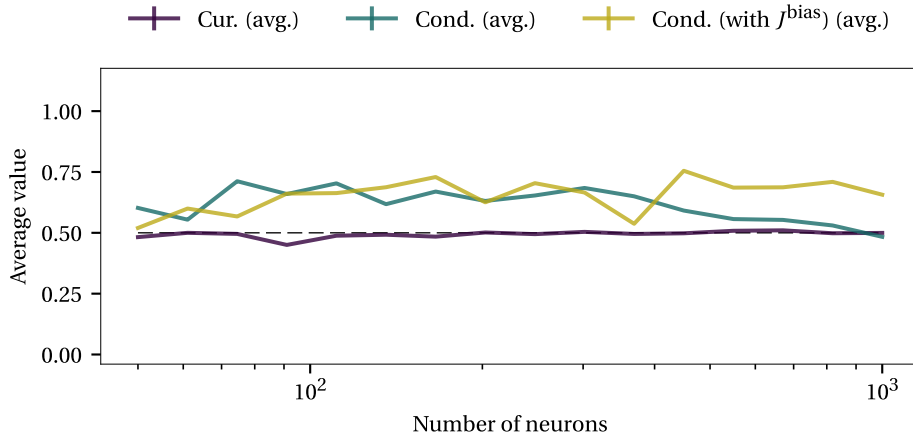


(c) $n = 8$

Figure 8: Communication channel benchmark. A bandwidth limited white noise signal is fed into a series of neuron ensembles (where the number of ensembles n is the “communication channel length”), with a synaptic time constant of 5 ms for all connections. The graphs show the RMSE between input and output signal averaged over 16 runs. The dashed line depicts the baseline error obtained when filtering the input n times with the synaptic filter.




(a) Integration error



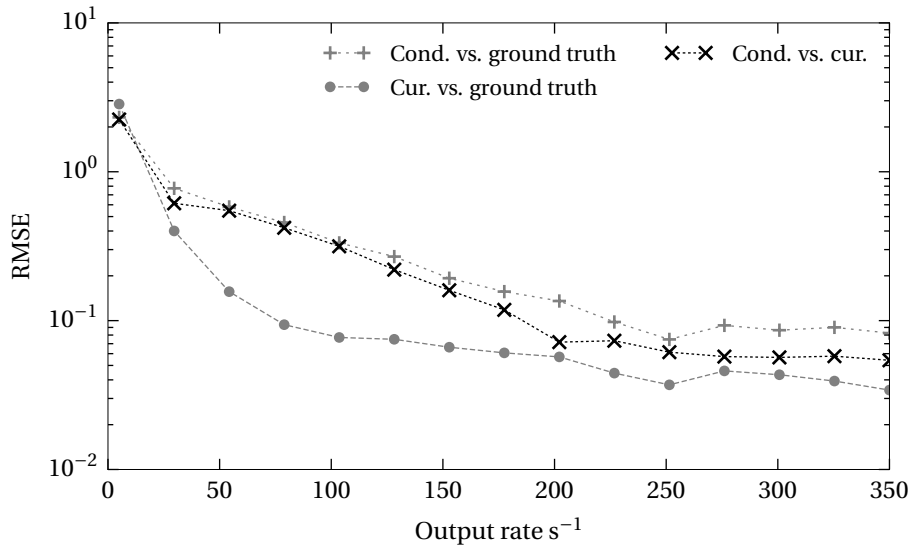
(b) Average integrator value

Figure 9: Results for the integrator benchmark. **(a)** Shows the average integrator representation error over 24 runs in which the integrator is driven to values in the interval $[0, 1]$. Error bars correspond to the standard deviation. **(b)** Shows the average value that the integrator converges to over all 24 runs. The dashed line depicts the expected value of 0.5.

Empirical tuning curve measurement Figure 10a shows empirically measured tuning curves for a LIF population with either current or conductance based synapses. Tuning curves are measured by feeding a bandwidth limited (5 Hz) white noise signal $f(t) = x$ into a standard current-based LIF ensemble with 1000 neurons connected to the target population with 100 neurons. The neural activities $a(x)$ are recorded. Finally, the inverse mapping $x(a)$ is calculated. The measured tuning curves are similar, however, as shown in fig. 10b, they mainly differ for low output rates—the output rates of conductance based neurons systematically are too small. When simulating dynamical systems such



(a) Tuning curves



(b) Error over output rate

Figure 10: Tuning curve experiment results. (a) Empirically measured tuning curves for both current and conductance based synapses. (b) RMSE between the tuning curves for varying output rate regimes. Ground truth refers to the non-spiking LIF firing rate model. In the “cond. vs. cur.” graph, the empirical measurements for the current-based model are used as a reference.

as an integrator, the imbalance caused by the later spike onset (which prevents both excitation and inhibition) may be a reason for the strong attractors imprinted on the system dynamics.

Summary In the feed-forward case, the simplistic transformation presented in section 4 allows neuron ensembles with conductance based synapses to function as well as their current-based counterparts, given that J^{bias} is eliminated from the neuron model (as predicted in section 6). For mostly unknown reasons, the performance of the conductance based synapses breaks down once recurrent networks are implemented. This may be caused by delayed spike onset in the tuning curves not accounted for when computing decoders.

8 Conclusion

In this report, we described both the NEF and LIF neurons with conductance-based synapses. We presented a simple transformation translating existing current-based NEF networks into their conductance-based counterparts. A sequence of benchmark experiments shows that this transformation yields relatively little error for feed-forward networks. However, when implementing dynamics such as integrators, the network implements dynamics with additional strong attractors. Experiments show that the neuron tuning curves differ significantly between neurons with current-based synapses and the transformed conductance-based version. They exhibit a delayed spike onset, which is a potential explanation for the additional attractors in recurrent networks.

Unfortunately, this problem can likely not be solved with a simple transformation of pre-computed weight matrices. Instead, the two-dimensional neuron response curve of a neuron with excitatory and inhibitory conductance based synapses must be taken into account. While the neuron response curve can be described analytically for constant, noiseless input, the non-negativity of weight matrices W^+ and W^- mandates an iterative optimization scheme. Preliminary research indicates that these methods are feasible, yet computationally more expensive than standard NEF solvers, especially because the entire weight matrix W is optimized and not the significantly smaller and reusable decoder D .

References

- [1] Chris Eliasmith and Charles H. Anderson. *Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems*. Cambridge, Massachusetts: MIT press, 2003. 380 pp.
- [2] Christof Koch. *Biophysics of Computation: Information Processing in Single Neurons*. Oxford, United Kingdom: Oxford University Press, 1999. 562 pp.
- [3] Wulfram Gerstner and Werner M. Kistler. *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. LCCB: 2002067657. Cambridge University Press, 2002.

- [4] Alan L. Hodgkin and Andrew F. Huxley. “A Quantitative Description of Membrane Current and Its Application to Conduction and Excitation in Nerve”. In: *The Journal of Physiology* 117.4 (1952), pp. 500–544.
- [5] J. Schemmel et al. “A Wafer-Scale Neuromorphic Hardware System for Large-Scale Neural Modeling”. In: *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*. May 2010, pp. 1947–1950.
- [6] Thomas Pfeil et al. “Six Networks on a Universal Neuromorphic Computing Substrate”. In: *Frontiers in Neuroscience* 7 (2013), p. 11.
- [7] Arnd Roth and Mark C. W. van Rossum. “Modeling Synapses”. In: *Computational Modeling Methods for Neuroscientists*. Ed. by Erik De Schutter. The MIT Press, 2009, pp. 139–159.
- [8] Heinrich Reichert. *Neurobiologie*. 2., neubearb. und erw. Aufl. Stuttgart [u.a.]: Thieme, 2000. VII, 250 S. : Ill., graph. Darst.
- [9] E. Kandel et al. *Principles of Neural Science*. 5th ed. McGraw-Hill Education, 2012.
- [10] Wulfram Gerstner et al. *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge University Press, 2014.
- [11] Christopher Parisien, Charles H. Anderson, and Chris Eliasmith. “Solving the Problem of Negative Synaptic Weights in Cortical Models”. In: *Neural Computation* 20 (2008), pp. 1473–1494.




Figure 11: Example of a low-pass filtered regular spike train. In this example, individual input spikes are $\Delta t = 0.1$ apart. The filter time constant is $\tau = 0.3$. The right graph is a magnification of the left one. Dotted lines correspond to \hat{h} and $\hat{h} - \tau^{-1}$, the dashed line to the final average $\Delta t^{-1} = 10$.

A Expected value of a filtered, regularly spaced spike train

Claim Given an infinite spike train with regular spacing $a(t) = \sum_{i=0}^{\infty} \delta(t - \Delta t \cdot i)$ and a normalized exponential low-pass filter $h(t) = \frac{1}{\tau} \cdot e^{-t/\tau}$ it holds

$$\langle h * a \rangle = \lim_{t \rightarrow \infty} \int_0^t \frac{(h * a)(t')}{t} dt' = \frac{1}{\Delta t}.$$

Proof This directly follows from linear systems theory and the unit-gain of the exponential low-pass. Alternatively, the supremum \hat{h} of $h * a$ can be expressed in terms of a geometric series

$$\hat{h} = \sup(h * a) = \frac{1}{\tau} \cdot \sum_{\ell=0}^{\infty} e^{-\frac{\Delta t \cdot \ell}{\tau}} = \frac{1}{\tau} \cdot \frac{e^{\frac{\Delta t}{\tau}}}{e^{\frac{\Delta t}{\tau}} - 1}.$$

Correspondingly, and as illustrated in fig. 11, after a finite amount of time the filtered spike train approximately oscillates between \hat{h} and $\hat{h} - \tau^{-1}$ every Δt . The function between these two points is an infinite sum of exponentials, which itself is an exponential $g(t) = a \cdot e^{-bt}$, with $g(0) \stackrel{!}{=} \hat{h}$ and $g(\Delta t) \stackrel{!}{=} \hat{h} - \tau^{-1}$. Solving for a, b yields $a = \hat{h}$ and $b = \tau^{-1}$. The average of $g(t)$ over $[0, \Delta t)$ is

$$\begin{aligned} \langle h * a \rangle &= \langle g(t) \rangle = \frac{1}{\Delta t} \cdot \int_0^{\Delta t} g(t) dt = \frac{1}{\Delta t \cdot \tau} \cdot \frac{e^{\frac{\Delta t}{\tau}}}{e^{\frac{\Delta t}{\tau}} - 1} \cdot \int_0^{\Delta t} e^{-\frac{t}{\tau}} dt \\ &= \frac{1}{\Delta t} \cdot \frac{e^{\frac{\Delta t}{\tau}}}{e^{\frac{\Delta t}{\tau}} - 1} \cdot \left(1 - e^{-\frac{\Delta t}{\tau}}\right) = \frac{1}{\Delta t} \cdot \frac{e^{\frac{\Delta t}{\tau}}}{e^{\frac{\Delta t}{\tau}} - 1} \cdot \frac{e^{\frac{\Delta t}{\tau}} - 1}{e^{\frac{\Delta t}{\tau}}} = \frac{1}{\Delta t}. \quad \square \end{aligned}$$