

From Voltage to Wiring:
Synaptic connectivity inference from neural voltage recordings

Tomas Fiers

Doctoral thesis

Supervisors: Mark Humphries & Matias Ison
Assessors: Rüdiger Thul & Dan Goodman

University of Nottingham
School of Psychology

January 2024

Corrected version as of June 5th, 2024

Acknowledgements	2
Code availability	3
1 Introduction	4
1.1 Overview	4
1.2 Recording neural activity	4
1.3 Voltage imaging	5
1.4 Inferring wiring from activity	6
1.5 Voltage imaging in more detail	8
1.6 Conclusion	11
2 Simulation details	12
2.1 The AdEx neuron	12
2.2 Alternative neuron models	15
2.3 The Izhikevich neuron	16
2.4 Synapse model	20
2.5 Model summary	22
2.6 Input in the N-to-1 setup	23
2.7 Spike ceiling	26
2.8 Voltage imaging	27
2.9 Software implementation	28
3 Spike-triggered averaging	29
3.1 Ceiling and clipping	33
3.2 Window length	34
3.3 Performance quantification	35
3.4 Recording duration & noise	39
3.5 EI balance	41
3.6 Computational cost of STA test	48
3.7 Conclusion	50
4 Network model	51
4.1 Introduction	51
4.2 Connectivity structure	51
4.3 External input	51
4.4 EI balance	52
4.5 Subsampling	54
4.6 Connection testing	54
4.7 False positive detections	55
4.8 Conclusion	57
5 New connection inference methods	58
5.1 Introduction	58
5.2 STA Template correlation	59
5.3 Fitting a full STA model	61
5.4 Linear regression of the upstroke	63
5.5 Conclusion	68
6 Discussion	69
6.1 Future work	69
6.2 Conclusion	72
References	73

ACKNOWLEDGEMENTS

I want to thank the following people for their help in making this thesis a reality.

For their expert guidance, and their patience and kindness, my supervisors Mark and Matias.

And for their moral support in many different ways: Lucia, Steven, Sara, Evelien, my siblings Pieter and Silke, my mom and dad, and my friends in Holsbeek (Lies, Emily, Koen, Tim, Arnoud, Katia) and in Nottingham (Frederica, Helen, Emily). Without them, this thesis would have never been.

CODE AVAILABILITY

Figures and tables in this thesis are often accompanied by a link to the Jupyter Notebook that generated them. These notebooks contain more information than can be conveyed in just the caption. These links look like, e.g: [2021-12-06__local_HH_dV_shape](#). To visit these links, in the digital version of this document, just click them. In the paper version, convert them to an URL, following this example: https://tfiers.github.io/phd/nb/2021-12-06__local_HH_dV_shape.html.

The Jupyter notebooks almost always rely on functions developed in previous notebooks, and factored out to separate files. These files can be found at <https://github.com/tfiers/phd> (under the pkg/ subdirectory). This GitHub repository is currently private, but the reader will be granted access after sending their GitHub email or username to tomas.fiers@gmail.com.

Introduction

The first section here gives an overview of the idea of this thesis, and its background. The second section delves a bit deeper into voltage imaging technology.

1.1 Overview

Systems neuroscience studies the links between (1) an animal's behaviour, (2) the activity of its neurons, and (3) how these neurons are connected. Currently, only the first two can be observed simultaneously, using *in vivo* recordings of neural activity. Observing the connections between neurons, on the other hand, requires imaging brain slices, and thus killing the animal. In addition, such a wire-tracing process is costly and time-consuming.

In this thesis, we aim to develop algorithms that infer the connections between neurons based on recordings of their voltages, instead of post-mortem imaging. We believe this is possible because the activity of neurons is mainly determined by the connections between them, and because recent advances in recording technology are yielding, for the first time, the necessary quality of data to solve this problem.

Such an algorithm then allows for fast and cheap estimation of the neural wiring in behaving animals, throughout their lifetime and across experiments. This could allow systems neuroscientists – whether they study e.g. memory, addiction, or movement disorders – to find answers to their questions linking brain wiring and brain (dys)function in a manner more straightforward than before.

1.2 Recording neural activity

There are currently two methods in popular use to record the activity of multiple individual neurons, *in vivo*: calcium imaging and extracellular electrode recordings [CS15; SH09]. The strength of each method is the weakness of the other. Both are performed through a small, surgically created hole in the skull.

In calcium imaging, neurons are filled with a calcium indicator or “dye” – a molecule that becomes fluorescent when calcium binds to it. When a neuron sends an action potential (a “spike”), its cell body is briefly flooded with calcium. A dyed neuron that has just sent a spike thus becomes fluorescent for a short while. To record activity, laser light is focused in a point and scanned across a volume of brain tissue. Neurons that recently spiked will send light back, which is captured to yield a three-dimensional movie of neural activity. This allows scientists to observe large numbers of neurons – namely all active neurons in the volume. An additional advantage is that calcium indicators can be genetically targeted so that they only occur in specific neuron types of interest, providing a focused view. A major disadvantage however is that multiple spikes fired from a neuron in rapid succession cannot be easily distinguished, as the calcium effects of each spike are slow and combine non-linearly.

Extracellular electrode recordings on the other hand have a much finer time resolution and every spike is individually distinguishable. They work by inserting long, thin implants in the brain, that have many electrodes exposed on their surface. Each electrode measures the local electric field potential, and thereby picks up the spikes from nearby neurons. The increased time resolution comes at the cost of only sampling a small subset of the neurons in the areas of interest, not knowing exactly where those neurons are, and not being able to target neural subtypes specifically.

Calcium imaging thus provides good spatial information but has a low temporal resolution, whereas extracellular electrode recordings present the opposite trade-off: precise spike counts and timings, but limited spatial information and sampling of neurons. Recently, a recording technique is emerging that combines the advantages of both.

1.3 Voltage imaging

Voltage imaging is very similar to calcium imaging: all or a genetically selected subset of neurons are made fluorescent, and these are scanned with a focused laser, to yield three-dimensional movies of neural activity [KS19]. The difference is that the indicator molecules used in voltage imaging fluoresce in direct proportion to the membrane potential of the cell, instead of its calcium concentration. This then allows to directly observe the membrane potential of all neurons of interest in the field of view (figure 1).

Although voltage imaging has existed for a long time, the recorded signal has long been too weak to distinguish it from background noise (unless animals with very large neurons are used, or the activity of many co-firing neurons is pooled together).¹ In recent years however, multiple labs have been iteratively refining the voltage indicator molecules. Together with the improvements in fluorescence imaging technology, driven by calcium imaging, this has made voltage imaging now powerful enough to image multiple individual

¹One reason for the relative noisiness of voltage imaging is that the indicator molecules are embedded in the cell membrane, whereas in calcium imaging, the indicator molecules float around the entire cell volume. The dendrites and axons of a neuron make up most of the cell’s surface area, whereas the soma makes up most of its volume. Calcium indicators are thus much more localized than voltage indicators.




Figure 1: **Voltage imaging.**

Left: schematic of an example voltage indicator molecule embedded in the cell membrane. Middle: the imaging setup and an example neuron image. Right: an example voltage imaging trace and a simultaneous intracellular electrode recording.

Adapted from [Vil+19; Abd+19].

neurons in vivo in common model animals. The signal-to-noise ratio has improved to the point that not only individual spikes, but also subthreshold voltage fluctuations can be observed.² As explained next, it is precisely this level of detail that we believe enables in vivo connection mapping.

² Section 1.5 further down expands further on the current capabilities of voltage imaging technology.

1.4 Inferring wiring from activity

The potential to infer the wiring from neural activity rests on the basic link between the two (figure 2): an excitatory neuron that sends a spike will slightly increase the voltage of all its downstream neurons (this small increase is called the excitatory postsynaptic potential, or EPSP). When a neuron has received enough spikes, its voltage crosses a threshold, and it will send a spike itself. To estimate neural wiring, the idea is then to invert this reasoning: if neuron B often shows activity right after neuron A has fired, then neuron A is likely to be connected to neuron B.

As both calcium imaging and extracellular electrode recordings yield (at best) spike timing data only, existing activity-to-wiring approaches have been based only on spike timing [MYD18; CMT18; DF20], and not on more detailed measurements of neural activity. The problem with this is that the correlation between two neurons being connected and them spiking together close in time is quite tenuous. For one, most neurons need to receive many spikes – each of which can come from any of its hundreds to tens of thousands of input neurons – before it fires a spike itself. Second, many neurons have long time constants, meaning that a spike can influence spiking in its receiving neurons up to hundreds of milliseconds later.

As a result, spike-based wiring inference methods require long

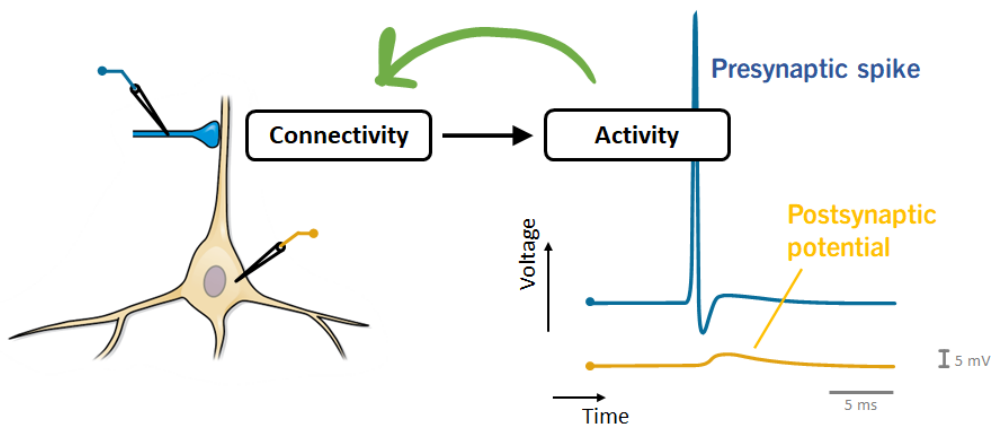


Figure 2: **The causal link between neural connectivity and activity.** On the left, a cartoon of a synaptic connection. The axon of presynaptic neuron M (in blue) impinges on postsynaptic neuron N (in brown). The electrode icons indicate that their membrane voltages are recorded (shown on the right). A successful spike in neuron M will elicit a small but precisely-timed voltage bump in neuron N (the postsynaptic potential, PSP). There is thus a causal relationship between 1) the existence of a connection $M \rightarrow N$ and 2) both neurons' membrane voltages. This causal relationship (black arrow) is exploited to perform network inference from voltage recordings (green arrow).

Drawings adapted from Purves et al.'s "Neuroscience" textbook, 6th edition, 2018.

recording durations to obtain some confidence on the wiring between even small numbers of neurons.³ During these long recordings, the connectivity may have already changed. And long recordings are not possible for fluorescence imaging, as dyes require recovery after each recording session.

When we can observe the subthreshold increases in voltage occurring directly after each spike however, we might be able to accurately reconstruct connectivity from recordings on the timescale of individual *in vivo* experiments. The recent advances in voltage imaging provide exactly this kind of data.

³ E.g. in [Orl+17], one-hour long recordings were used. This was long enough for a lot of the spike-based methods to recover edges, but *not* the directionality of connections.

1.5 Voltage imaging in more detail

What follows is a short literature review of the current capabilities of voltage imaging technology. The main goal is to be able to build simulations in this thesis relevant to reality.

Working principle

Voltage imaging is a type of fluorescence microscopy, itself a type of light microscopy. A light source such as a laser or an LED shines on the imaged object – such as a slice of brain tissue, a young transparent zebrafish, or part of a mouse brain made accessible through a hole in the skull. Most light passes through the sample of interest (brain tissue is mostly transparent for the used wavelengths). However, some parts of the sample ‘reflect’ light back.

These reflecting parts are so called voltage indicator molecules. Such molecules are introduced by the experimenter and gather spontaneously in cell membranes. The amount by which they reflect light back depends on the voltage placed over them. By capturing the reflected light from a sample, repeatedly over time, we thus get a movie showing 1) where cell membranes are, and 2) how the voltages over them change over time.

Physical detail

The above is a simplified description of the mechanics. This section gives a slightly more detailed description of what happens.

The actual mechanism of light ‘reflection’ is fluorescence, in which




Figure 3: **Camera view in voltage imaging.**

Fluorescent neurons visible through a cranial window. This is just one imaging plane (i.e. there are more neurons visible above and below this plane). Voltage imaging recordings are thus 3D videos of neural membrane voltages. Scalebars 1 mm and 0.1 mm, respectively. Adapted from [Abd+19; KS19].

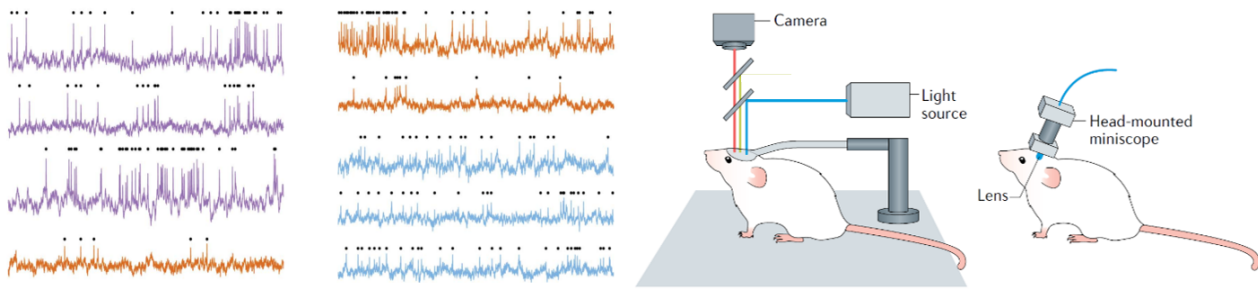


Figure 4: Left: More example VI traces, with different spike-signal-to-noise ratios. Colours correspond to neurons with coloured numbers in figure 3. Right: in-vivo imaging setups. Adapted from [Abd+19; KS19].

molecules absorb a photon, hold on to it for about one nanosecond, and then re-emit it again, at a longer wavelength. This process is about a million times slower than 'true' reflection (that is, photon scattering), but still a million times faster than the timescales we are interested in (milliseconds, corresponding to the duration of an action potential) and at which we image (with movie frame rates of 500 - 1500 Hz) [VB12; Cox19].

Furthermore, a change in voltage does not necessarily lessen the total fluorescence. Rather, the emission spectrum may shift. If you however only look at a fixed narrow band of the emitted light (as is done in fluorescence microscopy), the measured light will indeed seem to decrease or increase on voltage changes. 'Instantaneous reflection modulated by voltage' is thus a sufficient mental model for our purpose.

Genetic targeting

Since the late 90's, cells have been coerced into creating (parts of) indicator molecules themselves, by delivering transgenes into the cell [SI97]. This has multiple advantages. Most importantly, it allows the indicator molecules to be constrained to only certain cell types, by placing the transgene under the control of promoters that are only active in the cell type of interest. By doing so, scientists can avoid labelling glial cells. This decreases the background signal. Going further with the same principle, they can selectively label one type of neuron (for example, only interneurons in one hippocampal area) [Hoc+14].

Another advantage, also decreasing the background signal, is that indicator molecules can be constrained to the membrane of only the cell body, and not the membranes of dendrites and axon branches. This is done by adding small cell-body-targeting signalling sequences to the transgene.

System performance

The history of voltage imaging has mostly been the history of finding better indicator molecules. Earlier versions (starting from their

invention in the late 60's) changed their reflectance only weakly and slowly in response to a change in voltage. In addition, they were toxic, disrupting the normal functioning of cells, or altogether destroying them [SDC73; Zoc+00]. Modern indicators are no longer toxic, and their voltage sensitivity and speed have improved substantially (by more than ten- and hundredfold, respectively) [Mil16; TW18].

Fidelity

Current voltage indicators can have sub-millisecond time constants, meaning they can replicate the shape of an action potential as well as an electrode recording. During an action potential in a brain slice, the measured fluorescence of current voltage indicators increases by 30% (\pm about 20%)⁴ [Pia+19; Ada+19] relative to baseline. In live, head-fixed mice, this fluorescence increase is about 8% (\pm 2%) [Abd+19; Vil+19].

More important is the change in brightness with respect to noise. For voltage imaging, a signal-to-noise ratio (SNR) is often defined as the height of the fluorescence signal during an action potential, divided by the standard deviation of the baseline fluorescence signal.⁵ The action-potential-SNR is around 30 (\pm 7) in brain slices [Ada+19; Pia+19], and around 10 (\pm 3) in live, head-fixed mice [Abd+19; Vil+19].

There has been no quantification yet of how well voltage indicators track subthreshold voltages. Estimating visually from published figures however, the correspondence between simultaneous electrode and optical recordings is substantial, and seems good enough to calculate with, even in recordings from live animals.

Yield

The number of simultaneously voltage-imaged neurons in live mice varies between 4 and 46 in the latest studies (with frame rates of mostly 500 Hz) [Ada+19; Vil+19; Abd+19; Pia+19].

For comparison, calcium imaging yields between 200 and 1000 neurons at a frame rate of 30 Hz, and up to 10,000 at 2 Hz [Pac+17]. These higher yields come of course at the cost of less information per neuron: voltage imaging tracks subthreshold voltages and detects nearly all spikes, while calcium imaging, especially at such frame rates, yields only spike detections, imprecise both in time and in number.

The number of simultaneously imaged neurons is in any case bound to increase for both modalities, as microscopic scanning systems get faster.

Voltage imaging can only be performed in relatively short continuous bouts: the same mechanism that makes the molecules fluorescent – namely, excitation on impact of a photon – also makes the molecules more chemically reactive, making them spontaneously break down ('photobleach') in reaction with their environment. The total fraction

⁴ Round brackets in this section indicate a rough indication of variation between experimental setups. See [this spreadsheet](#) for concrete data from some representative VI studies.

⁵ Note that such a definition is not compatible with the standard definition of SNR, where the denominator is most often the squared standard deviation of noise, and the numerator is the averaged squared value of *all* samples, of some denoised 'signal' time series.

of not-yet photobleached indicator molecules decays exponentially – and so do the fluorescence and the spike SNR. The time constant of this decay (i.e. duration after which SNR has decreased by 63%) is around 10 ± 5 minutes [Pia+19; Abd+19]. Most voltage imaging sessions are therefore not longer than this.

Cells replace broken indicator molecules relatively fast however, and experiments have shown that imaged photobleached cells show complete recovery after two days. Even shorter intervals between imaging sessions might be possible.

Because modern indicator molecules are not toxic, cells can be intermittently imaged over long periods – one study followed the same neurons in vivo over more than a month.

1.6 Conclusion

In this first chapter, we described voltage imaging technology, and introduced the idea to use the recordings it generates to reconstruct neural connectivity.

Our aim in this thesis is to test whether this idea is feasible in principle: can we perform network inference from voltage signals? To test this, we will work with simulated data: We'll generate neural activity from a known connectivity, and develop and test different connection detection methods on this simulated data.

[Chapter 2](#) describes the simulation: the choices we make for the used neuron model, synapse model, inputs, ..; and their parameters. Having set up our test environment, [chapter 3](#) introduces the first and simplest connection detection method, and different ways to evaluate its performance. Up to that point, only a simplified network is used (where N Poisson inputs impinge upon a single simulated neuron). [Chapter 4](#) expands the test to a full recurrent network of simulated neurons. Finally, in [chapter 5](#), we introduce and evaluate three new methods for network inference based on voltage signals. [Chapter 6](#) concludes the thesis with an overview of possible future research directions and a summary of our findings.

Simulation details

In this chapter, we describe our experimental setup: the neuron model we simulate, its inputs, and how we simulate voltage imaging.

2.1 The AdEx neuron

We choose to simulate the ‘AdEx’ point neuron model, or the ‘adaptive exponential integrate-and-fire’ neuron [BG05], with conductance-based synaptic currents. The AdEx neuron is a leaky-integrate-and-fire (LIF) neuron model, with two additions. First, the full upstroke of each spike is simulated, as an exponential runoff. Second, an extra dynamic variable is added: the adaptation current. This current allows the simulation of many non-linear effects of real neurons, like spike-rate adaptation, intermittent bursting, after-hyperpolarization, postinhibitory rebound, resonance, bistability, class 2 excitability¹, and more.

The AdEx model consists of two differential equations (1 and 2), and a discontinuous update after a spike is generated (3). One equation simulates the membrane voltage V , and one the adaptation current w :

$$C \frac{dV}{dt} = -g_L(V - E_L) + g_L \Delta_T \exp\left(\frac{V - V_T}{\Delta_T}\right) - I_{\text{syn}} - w \quad (1)$$

$$\tau_w \frac{dw}{dt} = a(V - E_L) - w \quad (2)$$

Parameters are described in [table 1](#).

The first term of equation (1) is the restorative force pulling the voltage back to the resting (or leak) potential E_L . The second, exponential term is what generates the spike upstrokes. It is practically zero over most of the sub-threshold regime, and only becomes large (and then very large) near the firing threshold. (This firing threshold is characterised further on).

I_{syn} is the synaptic current, explained in [section 2.4](#). We use the sign convention of inter alia Dayan & Abbott² where membrane currents

¹ Class 1 excitable neurons can fire repetitively at arbitrarily slow rates. Class 2 neurons cannot; they have a minimum firing rate (r). In other words, their graph of firing rate versus input current (" $f - I$ " curve) discontinuously jumps up from 0 Hz to r at some threshold input.

As we will see later (in [figure 15](#)), the parameters we use for our AdEx model make it class 1 excitable.

² [DA01], ch. 5.3, p. 162

are defined as positive when positive charges flow *out* of the cell. I.e. a positive I_{syn} decreases the membrane voltage (itself defined as the electric potential inside minus outside the cell).

The adaptation current w decays exponentially to zero on its own ($-w$ in equation (2)), and is influenced by voltage deviations from equilibrium: for $a > 0$, w acts on V in the opposite direction of the deviation, and for $a < 0$, w acts in the same direction. Izhikevich calls the former a resonant current, and the latter an amplifying one.³

³ [Izh07], section 5.2.4. Note the different notation used; see the translation in table 2.

Name	Description	Value
V	Membrane voltage	(in mV)
w	Adaptation current	(in pA)
C	Membrane capacitance	104 pF
g_L	Input / leak conductance	4.3 nS
E_L	Resting / leak potential	-65 mV
Δ_T	Threshold slope factor	0.8 mV
V_T	Location of minimum of $\frac{dV}{dt}$	-52 mV
τ_w	Time constant of adaptation current	88 ms
a	Sensitivity of adaptation current to V	-0.8 nS
θ	Spike definition threshold	40 mV
V_r	Reset voltage after spike	-53 mV
b	Adaptation current bump after spike	65 pA

Table 1: **Quantities and parameters of the AdEx neuron, equations (1) to (3).** Values are from a model fit to a cortical regular spiking (RS) neuron, from [Nau+08]. By defining the location of $\frac{dV}{dt}$'s minimum, V_T also co-determines the location of the firing threshold.

In this chapter, we will analyse $\frac{dV}{dt}$ as a function of V , i.e. analyse it as a dynamical system: will the voltage increase or decrease at the current voltage? For conciseness in later analysis, we call this function $F(V)$. I.e. $F(V) = \frac{dV}{dt}$ = the right-hand-side of equation (1) here, scaled by $1/C$. We'll mostly analyse F in the absence of synaptic and adaptation currents, i.e. for I_{syn} and w both zero. Figure 5 shows the $F(V)$ curve for an AdEx neuron fit to a real neuron. Figure 7 compares the $F(V)$ curve of an AdEx neuron with that of another two-dimensional neuron model, the Izhikevich neuron.

In addition to the two differential equations, the AdEx model also consists of an instantaneous reset condition. When the membrane voltage V reaches a certain threshold θ , a spike is recorded, V is reset, and w is increased:

$$\begin{aligned} \text{if } V > \theta \text{ then: } & V \leftarrow V_r \\ & w \leftarrow w + b \end{aligned} \quad (3)$$

This bump of the adaptation current is what provides the spike rate adaptation: the more spikes the neuron has recently fired, the higher the adaptation current w and the more it drives down the voltage (equation (1)), away from the firing threshold.




Figure 5: A linear-plus-exponential model (red) fit to data from a cortical pyramidal neuron (black), from [Bad+08]

For a full description of the different firing patterns that the AdEx neuron can exhibit, for different parameter values (i.e, a characterisation of its bifurcations), see [Nau+08], “Firing Patterns in the Adaptive Exponential Integrate-and-Fire Model”. For an even fuller description, including a derivation of two-dimensional neuron models from higher-dimensional, Hodgkin-Huxley-like models, see Eugene Izhikevich’s book, *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting* [Izh07].

In terms of parameter choice, we choose to simulate a cortical regular spiking (RS) neuron, i.e. a ‘standard’ excitatory neuron that does not display e.g. bursting or fast-spiking behaviour. We take our parameter values from [Nau+08]⁴, where they fitted an AdEx model to recordings from real cortical RS neurons injected with different step currents. Parameter values are listed in [table 1](#).

Examples of the signals V and w that this model generates are given later, in [figure 13](#) and [figure 16](#).

Analysis

Further on, we will compare the AdEx neuron with another neuron model, the Izhikevich neuron. We will do this by comparing their fixed points, taking inspiration from the dynamical systems approach as used in e.g. [Str94] and [Izh07].

Where are the fixed points of the dynamical system $\frac{dV}{dt} = F(V)$? I.e, where is $F = 0$? Like other neuron models, there are two fixed points: a stable one at the leak potential E_L , and an unstable one at the instantaneous firing threshold (which we’ll call E_T).

We see in equation (1) (for $I_{\text{syn}} = 0$ and $w = 0$) that the leak potential E_L is, numerically speaking (though not strictly mathematically), indeed a fixed point: at realistic parameter values, the exponential term is negligibly small at $V = E_L$.⁵ The second fixed point has no straightforward expression. The exact solutions for F ’s roots need the so called Lambert W or ‘product logarithm’ functions W_0 and W_{-1} (see [figure 6](#)). The roots are found at:

$$V = E_L - \Delta_T W_k \left(-\exp \left(\frac{E_L - V_T}{\Delta_T} \right) \right), \quad (4)$$

where $k = 0$ gives the resting potential, and $k = -1$ the instantaneous firing threshold.⁶ For our cortical RS neuron, this gives us an instantaneous threshold of $E_T = -49.6$ mV.

Also of interest – especially when comparing with the Izhikevich neuron later – is the slope of AdEx’s $F(V)$, i.e. its derivative with respect to V :

$$\begin{aligned} \frac{dF}{dV} &= \frac{d}{dV} \left(-g_L(V - E_L) + g_L \Delta_T \exp \left(\frac{V - V_T}{\Delta_T} \right) \right) \\ &= -g_L + g_L \exp \left(\frac{V - V_T}{\Delta_T} \right) \end{aligned} \quad (5)$$

⁴ Section 6, and Table 1, row 3

⁵ For the parameter values of [table 1](#) e.g, the exponential term in $F(V = E_L)$ works out to 3×10^{-19} mV/ms. The machine epsilon for 64-bit floating point numbers in Julia e.g. is 2×10^{-16} .




Figure 6: The Lambert W functions for real numbers.

⁶ Here too we see that the true resting potential is almost identical to E_L : W_0 passes through zero, and its argument is ≈ 0 , because the exponential’s argument is negative. For W_{-1} however, [figure 6](#) shows that the second term is not negligible around 0.

This derivative is zero at $V = V_T$. I.e, unlike what is suggested by the ‘ T ’ subscript, and the name ‘effective threshold potential’ given to it in the AdEx literature [BG05; Nau+08], V_T is not the instantaneous threshold potential (it is not a zero of F), but rather the minimum of F (it is the zero of $\frac{dF}{dV}$).

From equation (5), we also calculate the slope of F at its two roots. At $V = E_L$, the slope (also known as the leak or input conductance here) is $-g_L$, plus a negligibly small exponential term. The negative sign shows that this is a stable fixed point: in a linearization of F around this point, at voltages below the resting potential, F (i.e. $\frac{dV}{dt}$) is positive and thus V will increase. At voltages above E_L , F is negative and V will decrease. Small deviations on either side of the resting potential will thus decay back to this resting potential.

At the firing threshold $V = E_T$ (i.e. the second solution to equation (4)), the slope is:

$$g_L \left(\exp \left(\frac{E_T - V_T}{\Delta_T} \right) - 1 \right), \quad (6)$$

which is positive (as $E_T > V_T$), indicating that this is an unstable fixed point: a small deviation of the voltage above E_T will blow up to infinity (i.e, a spike is generated).

2.2 Alternative neuron models

Why did we choose the AdEx model to simulate neuron voltages? In short, because it strikes a good balance between realism and complexity. We briefly consider here two alternative neuron models: the simple leaky-integrate-and-fire (LIF) neuron, and the more complex Hodgkin-Huxley (HH) neuron. In the next section, we go into more depth on a third alternative, the very similar Izhikevich neuron.

A simpler model than AdEx would be the well-known LIF neuron:

$$C \frac{dV}{dt} = -g_L(V - E_L) - I_{\text{syn}}$$

$$\text{if } V > \theta, \text{ then: } V \leftarrow V_r$$

As is apparent from comparing this with equations (1) and (3), the AdEx model is an extension of the LIF model. The LIF neuron lacks a simulation of the upstroke of spikes (the exponential term in equation (1)), and the slower time-scale adaptation current (equation (2)), which allows the simulation of many qualitatively different real neuron types.

Would this thesis have been very different had we used LIF neurons instead? Probably not, though it might depend on the mean voltage level of the simulated neuron: if it is well below the firing threshold, both LIF and AdEx are linear (the exponential term is negligible),

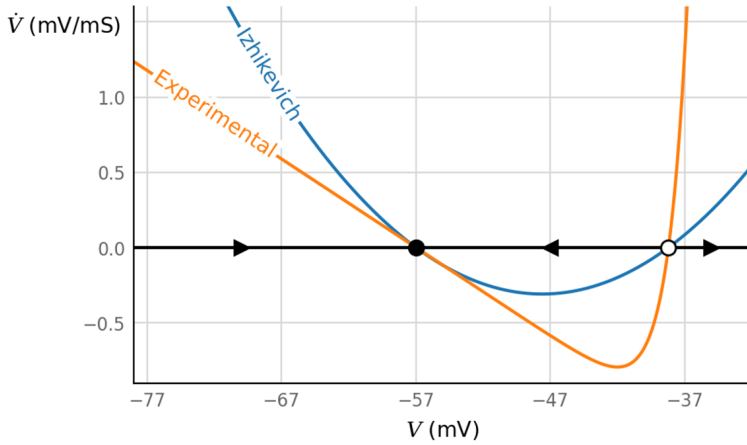


Figure 7: Neuron models as dynamical systems: a comparison of the $F(V)$ curves of the Izhikevich and AdEx neurons. ‘Experimental’ is an AdEx neuron fit to a cortical pyramidal neuron (using data from [Bad+08]). The Izhikevich neuron’s parameters were chosen to match the fixed points and the leak conductance. Arrows indicate whether the voltage will increase or decrease (1D flow field). $\rightarrow \bullet \leftarrow$ is a stable fixed point (resting potential), $\leftarrow \circ \rightarrow$ is an unstable fixed point (spike threshold). $\dot{V} \equiv \frac{dV}{dt}$.

and they behave quasi identically. When a spike is generated in the AdEx model, the exponential feedback makes the upstroke very fast, and thus not many timesteps in the simulation are spent on it, versus the linear regime.

On the other hand, when the neuron would continuously teeter just below its firing threshold, the LIF and AdEx models do not behave similarly. LIF’s $F(V)$ curve is still fully linear, while AdEx’s is not, and AdEx will behave more like a real neuron – see figure 5.

Another well-known alternative neuron model is the class of Hodgkin-Huxley (HH)-like neurons. These models simulate the full trajectory of a spike: both its upstroke and its downstroke. Unfortunately they also have many free parameters. They also take a bit longer to simulate, being higher dimensional (having more differential equations), and containing many more exponential terms, which take the brunt of the time when numerically evaluating a differential expression.

2.3 The Izhikevich neuron

Another alternative neuron model is the Izhikevich neuron, which is exceedingly similar to the AdEx neuron. These are the Izhikevich equations, using the same symbols as used before (in equations (1) to (3)):

$$C \frac{dV}{dt} = k(V - E_L)(V - E_T) - I_{\text{syn}} - w \quad (7)$$

$$\tau_w \frac{dw}{dt} = a(V - E_L) - w \quad (8)$$

$$\text{if } V > \theta, \text{ then:} \quad (9)$$

$$V \leftarrow V_r$$

$$w \leftarrow w + \Delta w$$

We have introduced two new parameters not present in the AdEx equations: the steepness of the parabola, k ; and E_T , the instanta-

neous firing threshold (the firing threshold in the absence of any synaptic or adaptation currents).

The only difference with AdEx is in equation (7), where the $F(V)$ curve is not made by a linear plus exponential term, as in AdEx; but rather by a quadratic (a parabola). Its two zeros (the fixed points) are readily apparent, as E_L and E_T .

Correspondences with AdEx

In Izhikevich's book⁷, different names are used for the same quantities:

⁷ [Izh07], section 5.2.4, equations 5.7 & 5.8

$$C \frac{dv}{dt} = k(v - v_r)(v - v_t) - u + I \quad (10)$$

$$\frac{du}{dt} = a(b(v - v_r) - u) \quad (11)$$

$$\text{if } v > v_{\text{peak}}, \text{ then:} \quad (12)$$

$$v \leftarrow c$$

$$u \leftarrow u + d$$

Table 2 compares both notation conventions.

AdEx	Izh	Description	Units
V	v	Membrane voltage	V
w	u	Adaptation current	A
τ_w	$1/a$	Time constant of adaptation current	s
E_L	v_r	Resting / leak potential	V
V_r	c	Reset voltage after spike	V
a	b	Sensitivity of adapt. current to V	S
b	d	Adaptation current bump after spike	A

Table 2: **Translating between Izhikevich and AdEx.** Different symbols used for the same quantities, in [BG05] and in most of this thesis ('AdEx'), and in [Izh07] ('Izh'). Membrane capacitance C (in farad) is the same in both notations.

Beside these straightforward correspondences, there are some parameters in either model that have no direct equivalent in the other: k and v_t in Izhikevich, and g_L , Δ_T , and V_T in AdEx. For those, we'll look at the shape of Izhikevich's $F(V)$, as we've done for the AdEx neuron before.

First, the AdEx parameter g_L . This is the input conductance, a.k.a. the leak conductance, and the slope of $F(V)$ around the leak potential. We can find this same conductance for the Izhikevich neuron by taking the derivative with respect to v of the right hand side of equation (10), at $w = 0$, $I_{\text{syn}} = 0$, and $v = v_r$. We find:

$$\left. \frac{d}{dv} (k(v - v_r)(v - v_t)) \right|_{v=v_r} = k(v_r - v_t) \quad (13)$$

(this value is negative: the leak potential is a stable fixed point. This corresponds to equation (1), where we find ' $-g_L$ '). Thus, our first

nontrivial correspondence:

$$g_L = k(v_t - v_r) \quad (14)$$

We've seen that V_T is the minimum of AdEx's F . The minimum of Izhikevich's F is easily found as the average of the parabola's two zeros. I.e, V_T corresponds to $(v_r + v_t)/2$.

Finally, Δ_T co-determines the slope of AdEx's F at the firing threshold (equation (6)). Given that Izhikevich's F is a parabola, with slopes equal in magnitude at both roots, we already know the firing threshold slope: it is the same as the leak conductance, equation (14). Here, the AdEx model is more expressive than Izhikevich's: the slope of $\frac{dV}{dt}$ at the firing threshold can be independently tweaked from the leak conductance; in Izhikevich these two are clung together by the form of the quadratic equation.

Comparison with AdEx

The Izhikevich and AdEx models are very similar. Their phase spaces are topologically identical: the adaptive current equation is identical (up to a renaming of the variables); and the $F(V)$ -graph has the same shape, with two fixed points: a stable fixed point at the resting potential, and an unstable one at the firing threshold (figure 7).

They differ in the exact shape: Izhikevich's $F(V)$ is a parabola, while AdEx is the more realistic 'linear subthreshold, and then transitioning to an exponential' (see figures 5 and 7). As a result, Izhikevich neurons have an unrealistically slow spike upstroke, examples of which can be seen in figure 8.

A second issue is Izhikevich's subthreshold nonlinearity. The effects of this can be seen in figure 9. Positive input currents produce stronger responses than equally large negative input currents. This is explained by the quadratic $\frac{dV}{dt}$ shape: positive deviations are attenuated less, and negative deviations more, than a linear neuron would. Real and AdEx neurons do not suffer this asymmetry (figure 5).

This nonlinearity is not visible for small voltage deviations, which is what the postsynaptic potentials we are interested in in this thesis tend to be. There is however an effect of the neuron's average voltage: if this voltage is constantly on the higher side, then inputs – both negative and positive – will cause larger responses than if the median voltage was lower.




Figure 8: **Two neuron models behave differently for (near) identical parameters and input.**

The AdEx neuron's parameters are from [Nau+08], for a cortical regular spiking neuron. The Izhikevich neuron's parameters are copied from the AdEx neuron wherever they correspond directly. The other parameters are chosen so both models have the same resting and threshold potentials, and the same leak conductance, using the correspondences found earlier in this section. Both models receive EI-balanced synaptic input from 6500 Poisson spike trains with lognormal firing rates. The AdEx neuron was given stronger inputs ($\Delta g_{\text{exc}} = 12.2$ pS) than the Izhikevich neuron ($\Delta g_{\text{exc}} = 4$ pS), so as to obtain the same number of output spikes. (In both cases, $\Delta g_{\text{inh}} = 4 \Delta g_{\text{exc}}$). For more details, see [2023-06-23__Vm_traces_AdEx_Izh_Brian](#).

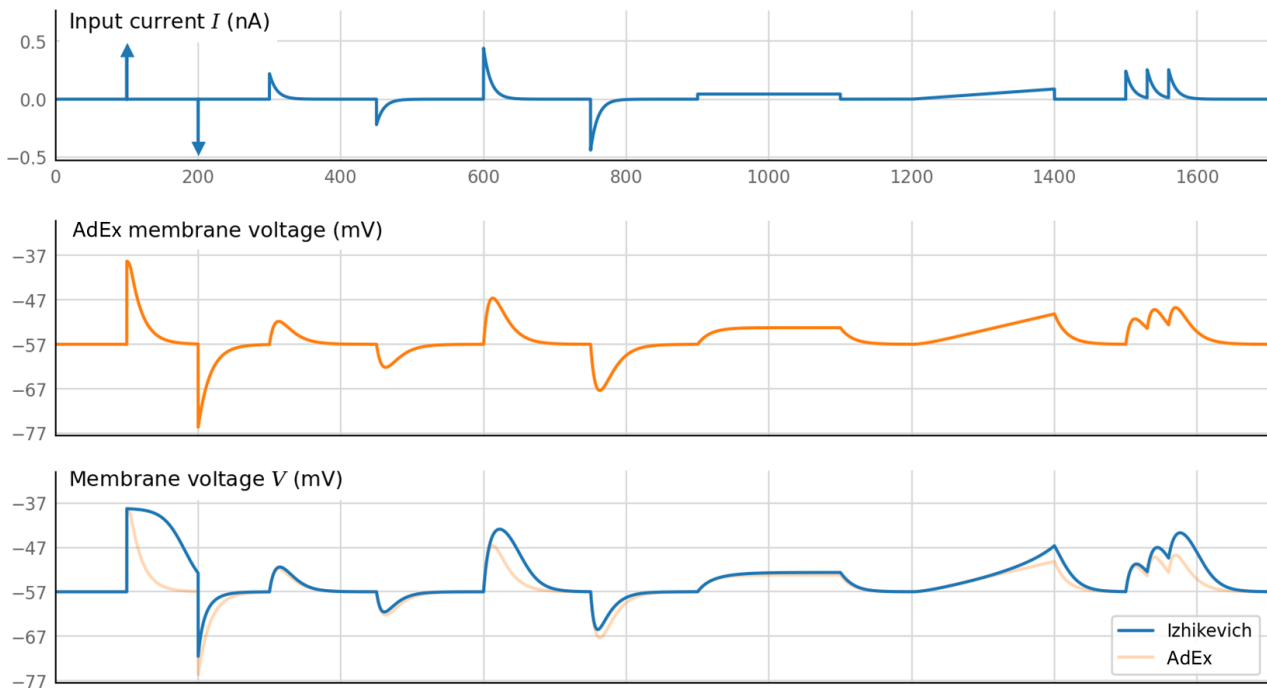


Figure 9: **The nonlinear response of Izhikevich neurons to sub-threshold input currents.**

Adaptation currents are negligibly small for both models in this test scenario. Source: [2021-12-08__biology_vs_Izh_subhtr](#)

2.4 Synapse model

One as of yet unexplained term in our neuron model, equation (1), is the synaptic current I_{syn} . This is the following sum over all input synapses i of the neuron:

$$I_{\text{syn}}(t) = \sum_i g_i(t) (V(t) - E_i), \quad (15)$$

where V is the global membrane voltage of the neuron, E_i is the reversal potential of that synapse, and g_i is the local synaptic conductance, which is modulated by presynaptic spikes.

For an excitatory synapse, $E_i > V(t)$, making $g_i(t) (V(t) - E_i)$ negative, increasing the membrane voltage according to the sign convention for I_{syn} in equation (1).

We simulate the synaptic conductances g_i as exponentially decaying signals (with time constant τ_g), and bump them up instantaneously on arrival of a presynaptic spike:

$$\frac{dg_i}{dt} = -g_i/\tau_g \quad (16)$$

On incoming presynaptic spike:

$$g_i \leftarrow g_i + \Delta g_i \quad (17)$$

Note that these are not the so called alpha-synapses. Those have *two* terms (and corresponding time constants) in their differential equation (instead of one, as in equation (16)): they also have an exponential rise, instead of just an exponential decay. (For an infinitely fast rise though, these models are of course the same). Simulating a full alpha synapse might increase the realism of our voltage traces, for a small simulation cost. We did not try this however. Foremost because alpha synapses fit to real data often have very fast rise times that are almost indistinguishable from instantaneous jumps.

For efficiency, we give all our excitatory synapses the same reversal potential, E_{exc} . Idem for the inhibitory synapses, with E_{inh} . This allows us to factor the synaptic current sum (equation (15)) as follows:

$$I_{\text{syn}}(t) = (V(t) - E_{\text{exc}}) \sum_{\text{exc } i} g_i(t) + (V(t) - E_{\text{inh}}) \sum_{\text{inh } i} g_i(t) \quad (18)$$

The sums of conductance signals $g_i(t)$ can also be simplified. Say that the values of g_i at $t = 0$ are G_i . The solution to equation (16) (at least in the time until a new presynaptic spike arrives) is then

$$g_i(t) = G_i e^{-t/\tau_g} \quad (19)$$




Figure 10: Example synaptic conductance trace $g_1(t)$, with a single incoming spike at $t = 20$ ms.




Figure 11: Another example trace $g_2(t)$, with spikes at $t = 10$ ms and 30 ms, and a smaller Δg .

With this, and when all synapses have the same time constant τ_g , the two sums in equation (18) can be factored as follows:⁸

$$\sum_i g_i(t) = \sum_i \left(G_i e^{-t/\tau_g} \right) = \left(\sum_i G_i \right) e^{-t/\tau_g} \quad (20)$$

This means that we need to only keep track of two conductance signals: g_{exc} and g_{inh} , each the sum of all excitatory or all inhibitory synaptic conductances.

Our synaptic current sum then becomes simply:

$$I_{\text{syn}}(t) = g_{\text{exc}}(t) \cdot (V(t) - E_{\text{exc}}) + g_{\text{inh}}(t) \cdot (V(t) - E_{\text{inh}}), \quad (21)$$

and we only need to simulate two differential equations, instead of one for every synapse:

$$\begin{aligned} \frac{dg_{\text{exc}}}{dt} &= -g_{\text{exc}}/\tau_g \\ \frac{dg_{\text{inh}}}{dt} &= -g_{\text{inh}}/\tau_g, \end{aligned} \quad (22)$$

where on arrival of a spike at synapse i either g_{exc} or g_{inh} is instantaneously increased by a value Δg_i , depending on whether that synapse is excitatory or inhibitory.

We choose an excitatory reversal potential of $E_{\text{exc}} = 0$ mV, an inhibitory one of $E_{\text{inh}} = -80$ mV, and a time constant for the synaptic conductance decay of $\tau_g = 7$ ms. These values are rather arbitrary, but in line with other simulation studies (e.g. [Bre+07]).

⁸ This is only valid in the time before any new spikes arrive. But the ‘summability’ still holds after a new spike. To see this, the given reasoning can be repeated, but simply with different values for the G_i (all decayed by an amount $e^{-t_{\text{spike}}/\tau_g}$, and one increased by a bump Δg_i), and then redefining t_{spike} to be $t = 0$.




Figure 12: A third synaptic conductance trace $g_3(t)$, with three input spikes at the same times and strengths as in figures 10 and 11. This signal is simulated independently, but turns out to be equal to the sum of the two others: $g_3(t) \equiv g_1(t) + g_2(t)$.

2.5 Model summary

Combining equations (1) to (3) with equations (21) and (22), our complete AdEx point neuron with conductance-based synaptic current is:

$$C \frac{dV}{dt} = -g_L(V - E_L) + g_L \Delta_T \exp\left(\frac{V - V_T}{\Delta_T}\right) - \underbrace{g_{\text{exc}}(V - E_{\text{exc}}) - g_{\text{inh}}(V - E_{\text{inh}})}_{-I_{\text{syn}}} - w$$

$$\tau_w \frac{dw}{dt} = a(V - E_L) - w$$

$$\tau_g \frac{dg_{\text{exc}}}{dt} = -g_{\text{exc}}$$

$$\tau_g \frac{dg_{\text{inh}}}{dt} = -g_{\text{inh}}$$

$$\begin{aligned} \text{When } V > \theta: \quad & V \leftarrow V_r \\ & w \leftarrow w + b \end{aligned}$$

$$\begin{aligned} \text{On input spike at..} \\ \text{exc. synapse } i: \quad & g_{\text{exc}} \leftarrow g_{\text{exc}} + \Delta g_i \\ \text{inh. synapse } j: \quad & g_{\text{inh}} \leftarrow g_{\text{inh}} + \Delta g_j \end{aligned}$$

We solve these equations numerically using first-order (Euler) integration, with a timestep Δt of 0.1 ms. This timestep is sufficiently small with respect to the different time constants in the model⁹. See section 2.9 below for more details on the numeric implementation.

Figure 13 shows the impulse response of this model, using a single spike, coming from either an excitatory or an inhibitory input neuron. The PSP bump is visible in the second panel from the bottom ('membrane voltage, V '). Note its tiny size, of about 0.04 mV. Our task will be to detect this tiny signal, in a sea of voltage imaging noise and PSP bumps of other input neurons.

Note also that, even though the inhibitory input is four times as strong as the excitatory one, its synaptic current and PSP bumps are not larger. (In fact, they are a smidge smaller). This is due to the fact that the neuron's resting potential (at -65 mV) lies closer to the inhibitory reversal potential ($E_{\text{inh}} = -80$ mV) than the excitatory one ($E_{\text{exc}} = 0$ mV), making the the inhibitory term of the synaptic current I_{syn} weaker than the excitatory term (equation (21)).

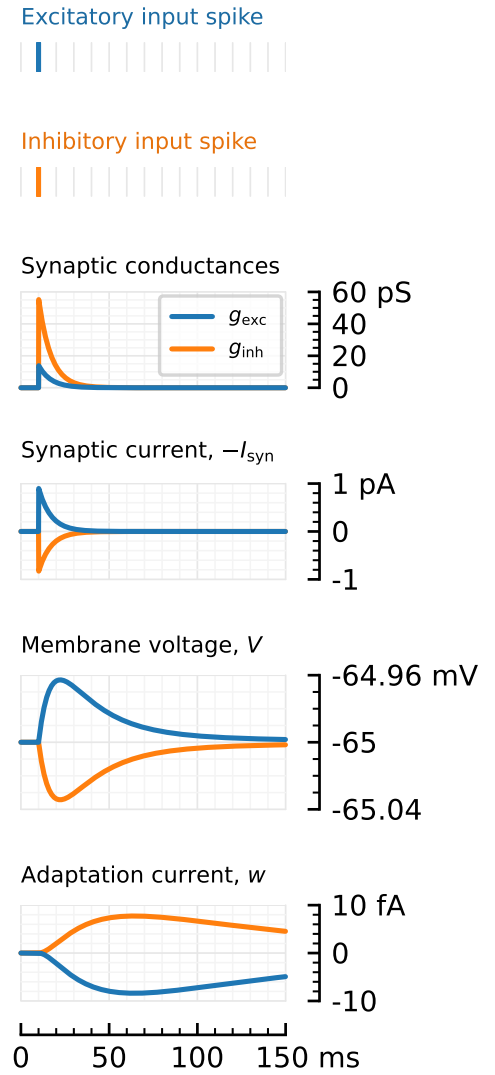


Figure 13: **Impulse response of the conductance-based AdEx neuron.** Overlay of two independent simulations, each using a single spike arriving at $t = 10$ ms, with synaptic weight of either $\Delta g_{\text{exc}} = 14$ pS, or $\Delta g_{\text{inh}} = 56$ pS. Parameters as in table 1, with $E_{\text{exc}} = 0$ mV, $E_{\text{inh}} = -80$ mV and $\tau_g = 7$ ms. For the inhibitory impulse response, $g_{\text{exc}}(t) \equiv 0$, and for the excitatory impulse response, $g_{\text{inh}}(t) \equiv 0$ (neither is shown). More details at [2023-09-05__Inhibitory_impulse_response_PSP](#).

⁹ $\tau_g = 7$ ms,
 $\tau_w = 88$ ms,
 $C/g_L = 24.2$ ms.




Figure 14: **Input spike trains are given log-normally distributed firing rates.**

The distribution used in our simulations is shown in bold. It has mean $\mu_x = 4$ Hz and variance of the underlying Gaussian $\sigma^2 = 0.6$. The two light distributions are from the literature. Note that the Roxin et al. distribution is slightly more heavy tailed than this study's: it has both more low firing and more high firing neurons.

2.6 Input in the N-to-1 setup

In our simplest experimental setup, we simulate just one AdEx neuron. Its input is provided by an array of N Poisson neurons, i.e. they each generate spike trains according to a Poisson process.¹⁰ We call this the 'N-to-1' setup.

Log-normal Poisson spike trains

The inter-event intervals of a Poisson process follow an exponential distribution. We use that fact to generate spike trains: we draw samples from $\text{Exp}(\lambda)$ (with λ the desired firing rate), and cumulatively sum up these intervals to obtain spike times. This is done until we have reached the desired input train duration.

The firing rates of real neurons often follow a long-tailed distribution: most neurons do not fire much at all, while a few fire a lot [MB13; HDZ08; Sha+07]. We recapitulate this in the firing rates λ of our Poisson input neurons, by drawing them from a log-normal distribution. We look to the literature for realistic parameters for this distribution, namely to the modelling paper of Roxin et al. [Rox+11], and the experimental sources it cites [HDZ08; OCo+10].

Log-normal distributions are usually parametrized with μ and σ : the location and scale of the underlying normal distribution, i.e. after log-transforming the input domain. In the above sources however, the mean μ_x of the data distribution itself is given. We can find μ , if σ is known, as $\mu = \ln(\mu_x) - \sigma^2/2$.

Roxin et al. use a mean rate μ_x of 5 Hz and a variance of the logarithm of the rate $\sigma^2 = 1.04$ in their figure 2. Hromádka et al. [HDZ08] recorded neurons in the auditory cortex of awake rats during acoustic stimulation. They find a mean firing rate μ_x of 6.2 Hz and a

¹⁰ This is both a simple and relatively accurate model of real cortical pyramidal firing statistics: the time between two spikes is almost completely independent of preceding inter-spike intervals [SK93].

median of 2.4 Hz; no numeric variances are given. O'Connor et al. [OCO+10] recorded neurons in the barrel (whisker) cortex of behaving mice. Ensembled over all layers of the cortex they recorded from, they report $\mu_x = 7.4$ Hz, $\sigma_x = 12.6$ Hz, a median of 1.5 Hz, and an interquartile range of 9.5 Hz. This corresponds to a Gaussian variance of $\sigma^2 = \ln(1 + \sigma_x^2/\mu_x^2) = 0.30$. The Roxin and O'Connor distributions are also shown in figure 14.

Given these data, we choose our parameters μ and σ so that the distribution lies roughly halfway – on a logarithmic scale – between O'Connor's and Roxin's, while making sure our firing rates are rather low than high: as our idea for connection inference rests on the number of spikes that can be used to calculate spike-triggered averages (see chapter 3), we don't want to overestimate the number of available input spikes and obtain overly optimistic results.

Our log-normal distribution has a median of 2.96 Hz, very close to Roxin's 2.97 Hz. O'Connor's has 6.4 Hz, which is markedly different from the median of 1.5 Hz they reported (hinting that their data distribution might not be log-normal).

Synaptic weights

We choose to simulate EI-balanced input, where we simulate four times as many excitatory as inhibitory input spiketrains; but with the inhibitory inputs four times as strong as the excitatory inputs.

I.e. if our total number of input spiketrains is $N = 6500$, we have $N_{\text{exc}} = 5200$ excitatory inputs, and $N_{\text{inh}} = 1300$ inhibitory inputs. In the N-to-1-setup, we give every input of the same type the same synaptic weight. I.e. in the model summary (section 2.5), $\forall i : \Delta g_i = \Delta g_{\text{exc}}$ and $\forall j : \Delta g_j = \Delta g_{\text{inh}}$. If we would thus choose $\Delta g_{\text{exc}} = 10$ pS, then, for 4:1 EI-balanced input, Δg_{inh} would be 40 pS.

To choose a value Δg_{exc} for the excitatory input strength, we look at the desired output firing rate of our one simulated neuron. We want it to spike at a realistic, average rate. Thus, we try for it to have the same output firing rate as the mean input firing rate: $\mu_x = 4$ Hz. This rate is achieved by testing a range of different input drives Δg_{exc} : see figure 15. For $N = 6500$ inputs, we find an average output firing rate of 4.0 Hz at $\Delta g_{\text{exc}} = 15$ pS (and thus $\Delta g_{\text{inh}} = 60$ pS).

Two asides on figure 15. First, from the bottom panel, we see that the AdEx model, with parameters for a cortical RS (regular spiking) neuron, is a so called 'Type I' neuron:¹¹ the firing rate can be made arbitrarily low, and there is no discontinuous jump from 0 Hz to some minimum firing rate.

Second, why does the output activity level increase for increasing excitatory input, if the inhibitory input becomes stronger by the same amount? The answer lies in the synaptic reversal potentials, which are $E_{\text{exc}} = 0$ mV and $E_{\text{inh}} = -80$ mV. The median membrane voltage (top panel in figure 15), at about -60 mV, lies closer to the

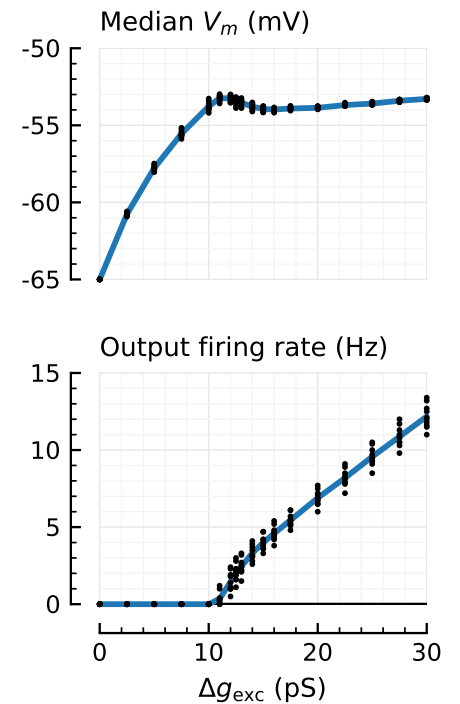


Figure 15: Activity level of the output neuron in the N-to-1 setup, for increasing input strength Δg_{exc} .

With $N = 6500$ EI-balanced inputs, and $\Delta g_{\text{inh}} = 4 \cdot \Delta g_{\text{exc}}$.

The blue line is the average over 10 different simulations (black dots), each with a different random seed for the input spike train generation.

Increasing input first heightens the voltage level without increasing the firing rate, and later heightens the firing rate without increasing the voltage level further.

V_m = membrane voltage V .

Source: 2023-08-05__AdEx_Nto1_wesweep.

¹¹ <https://neurondynamics.epfl.ch/online/Ch4.S4.html>

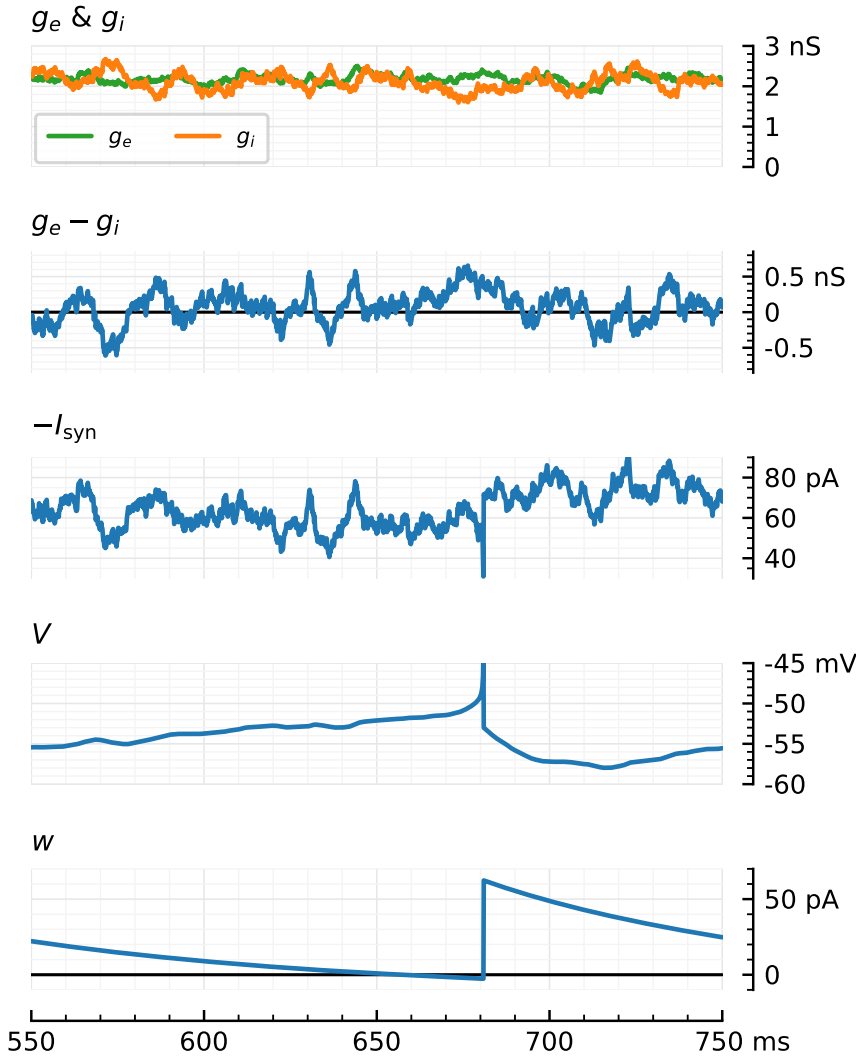


Figure 16: Signals generated by the neuron model, for 6500 EI-balanced input spiketrains.

Example extract from the signal traces of the conductance-based AdEx neuron. N-to-1 setup, with $\Delta g_{\text{exc}} = 15$ pS, and Poisson inputs with lognormally distributed firing rates. Neuron parameters for a cortical regular spiking neuron.

Note that the synaptic current signal $-I_{\text{syn}}$ is very similar to the difference of the excitatory and inhibitory synaptic conductances, as long as the membrane voltage V does not vary too much; on a spike (here at $t \approx 680$ ms), I_{syn} spikes as well.

Also note that the neuron spikes at a time when the inhibitory input randomly falls below the excitatory input, for a certain time. Spikes are 'fluctuation driven' in this input regime.

$g_e \equiv g_{\text{exc}}$ and $g_i \equiv g_{\text{inh}}$.

The V and $-I_{\text{syn}}$ signals are cut-off by the plot boundaries at the spike time.

Source: [2023-07-26_AdEx_Nto1_we_I_syn](#).

inhibitory reversal potential than the excitatory potential, making the inhibitory term of the synaptic current I_{syn} weaker than the excitatory term (equation (21)).

The fact that this EI-balanced input leads to a net-excitatory effect can also be seen in figure 16: the signal $g_{\text{exc}} - g_{\text{inh}}$ hovers around zero, but the signal I_{syn} — even though it has largely the same shape — is nowhere near zero.

Later, we will compare the performance of a connection-detection algorithm across different numbers of inputs N . We choose the approximately evenly log-spaced sequence

$$N \in [10, 20, 45, 100, 200, 400, 800, 1600, 3200, 6500].$$

For each of these N , we want the output firing rate to be the same, namely 4 Hz. To find the synaptic strengths Δg_{exc} that accomplish this, we perform an iterative search.¹² As an initial guess, it would make sense for the required input strength to scale inversely proportional to the number of inputs. We find however that there is a slight deviation from this linear expectation (figure 17): the less inputs, the less strong each input should be to reach the same output

¹² We use SciPy's `root_scalar` function, which uses 'Brent's method', which is like bisection but has faster convergence. We seed the algorithm with the bracket $[w_0/4, w_0 \cdot 4]$, where w_0 is the initial linear guess for Δg_{exc} . In about 8 iterations, we come within 0.01 Hz of the desired 4 Hz.

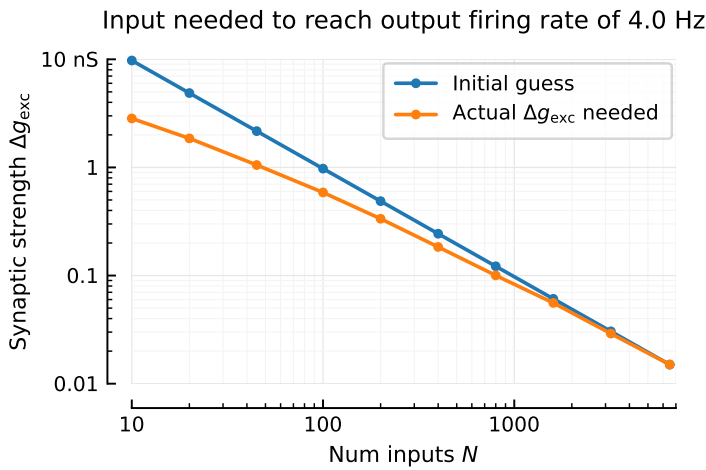


Figure 17: Finding the right input strength in the N-to-1 setup, for different numbers of inputs N .

Δg_{exc} is the synaptic conductance increase per incoming spike from an excitatory input. For each N , the right Δg_{exc} value — namely the one for which the output neuron fires at 4 Hz — was found using an iterative search procedure. For every Δg_{exc} value evaluated, ten different 10-second simulations were run (each with a different random seed for the input spiketrain generation). The average output firing rate of these ten simulations was taken, and compared with the goal firing rate of 4 Hz. Δg_{exc} was then adjusted until the simulated firing rate was sufficiently close. In every case, Δg_{inh} was four times Δg_{exc} . Source: [2023-08-05__AdEx_Nto1_we_sweep](#).

firing rate, compared to a linear extrapolation from our finding of $\Delta g_{exc} = 15$ pS for $N = 6500$. E.g. the linear expectation for $N = 10$ would be $\Delta g_{exc} = 9.75$ nS ($= 15$ pS $\cdot 6500/10$), but we find we need only $\Delta g_{exc} = 2.83$ nS to reach an output firing rate of 4 Hz.

2.7 Spike ceiling

Our neuron model (section 2.5) consists of 1) a set of continuous differential equations, and 2) a discontinuous, instantaneous jump (namely whenever the membrane voltage V crosses the spike-definition threshold θ). Because we simulate our model using a discrete and finite timestep, this discontinuity introduces variability in the height of our spikes: the simulated voltage will never exactly equal the threshold θ in a given timestep, but will be either somewhere below it (where it will become the simulated spike height), or above it (in which case a spike is defined and the voltage is reset). This variability, or jitter, in the height of spikes can be seen in figure 18 (orange trace).

Real neurons generally do not show such variability: their spike heights are quite consistent (remarkably so). To make our simulated voltage traces look more realistic, we set the voltage to a fixed height at spike events. We call this modification "spike ceiling". It is illustrated in figure 18 (blue trace).

This spike ceiling is a common technique. In Izhikevich's book for




Figure 18: Example voltage trace, before (orange) and after (blue) spikes are ceiled.

Source: [2023-09-05__ceil_spikes](#).

example, the same technique is applied (there it is called "spike padding"): see [listing 1](#). There is one difference however. Izhikevich's code pads the spikes during the simulation, and overwrites the spike height of the previous timestep, before the threshold crossing. We ceil our spikes after the simulation has fully run. We ceil the voltage one timestep later than Izhikevich: after the threshold crossing; i.e. the timestep where the voltage was previously reset to V_r . This difference in approach is merely a personal aesthetic preference, and was not found to yield any difference in network inference performance in an informal test.

The act of spike ceiling in itself does have an effect on network inference performance. This is shown later, in [section 3.1](#).

```

1 for i = 1:n-1           % forward Euler method
2   v(i+1) = v(i)+tau*(k*(v(i)-vr)*(v(i)-vt)-u(i)+I(i))/C;
3   u(i+1) = u(i)+tau*a*(b*(v(i)-vr)-u(i));
4   if v(i+1) >= vpeak   % a spike is fired!
5     v(i) = vpeak;      % padding the spike amplitude
6     v(i+1) = c;        % membrane voltage reset
7     u(i+1) = u(i+1)+d; % recovery variable update
8   end;
9 end;

```

Listing 1: Matlab code simulating the Izhikevich neuron. Lines pertinent to spike ceiling/padding are highlighted. Source: Eugene Izhikevich's 2007 book, *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting*, section 8.4.1 ("Simple Model of Choice"), p. 274 ([[Izh07](#)]).

2.8 Voltage imaging

The signals detected by a light microscope in a voltage imaging setup are not the same as the real membrane voltage traces of which they are a reflection.

We model this lossy transformation by simply adding Gaussian noise to our simulated membrane voltage.¹³ As in the voltage imaging literature, we quantify the amount of this noise by a 'spike-SNR' measure (spike signal-to-noise ratio). This is defined as the height of an average spike relative to the standard deviation of the noise:

$$\text{spike-SNR} = \frac{\text{spike height}}{\sigma_{\text{noise}}} \quad (23)$$

In other words, if we call our output signal y , and with V_i the voltage samples resulting from the numeric integration of equation (1), we have:

$$y_i = V_i + \varepsilon_i \quad (24)$$

$$\varepsilon_i \sim \mathcal{N}(0, \sigma_{\text{noise}})$$

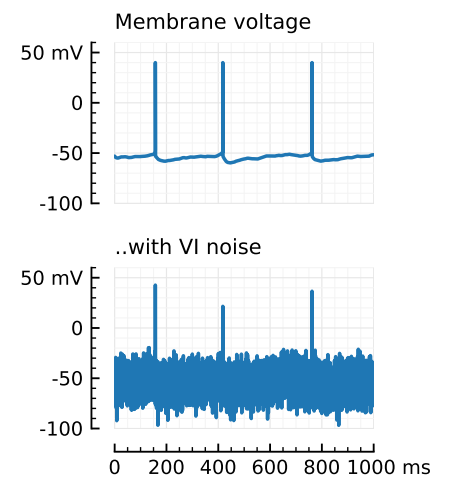


Figure 19: Simulated voltage trace (same as in [figure 8](#), right), without and with 'voltage imaging' (VI) noise added ($\sigma_{\text{noise}} = 10.5$ mV).

¹³ Adding noise with a fixed variance each timestep makes the power of the added noise dependent on the chosen timestep. In this study we do not vary the timestep, but if we did, we would have to account for this when choosing the noise variance.

A typical but conservative level of noise in voltage imaging recordings has a spike-SNR of 10. If we take as spike height the spike detection threshold minus the leak potential ($\theta - E_L$, see [table 1](#)), we obtain a σ_{noise} of 10.5 mV.

A more realistic model of the voltage imaging-transformation would also incorporate the exponential decay over time of the SNR (with a time constant of about 10 minutes), and the short-term ‘smearing in time’ of voltage indicators.¹⁴ The latter could be done by passing the voltage signal through a linear filter with some non-instantaneous impulse response.

2.9 Software implementation

To numerically simulate the model described in this chapter, different software implementations have been used throughout the PhD. Initially, a custom simulator was written in Python, with numerically expensive parts of the code accelerated using the Numba just-in-time compiler [[LPS15](#)]. Later on, we switched to a custom simulator written in Julia [[Bez+17](#)].¹⁵ Finally, we also implemented our model in Brian [[SBG19](#)]. After a comparison of compilation and runtimes of the Brian and Julia implementations ([table 3](#)), we decided to use the software written in Julia.¹⁶

Software	Language	Model	Time (seconds)	
			Compilation	Simulation
Brian2	Cython + Python	full	61	11
		merged	86	11
	C++	full	15	6.5
		merged	18	0.35
own	Julia	full	0.68	0.036

¹⁴ See e.g. the example voltage imaging trace in [figure 1](#): the sharp and high peaks of the electric signal are diluted in time through the low-pass filtering effect of the voltage indicator. See also the ‘kinetics’ column in the [spreadsheet](#) with genetically-encoded voltage indicator (GEVI) characteristics: the listed durations are the time constants of bi-exponential functions fit to the responses of the voltage indicator molecules to a rectangular voltage pulse.

¹⁵ A slideshow introducing Julia and how it compares to Python and Matplotlib can be found [here](#).

Table 3: Time taken to simulate 10 seconds of the N-to-1 experiment with $N = 6500$ Poisson inputs, and one conductance-based AdEx neuron. In the ‘merged’ models, only 200 Poisson spiketrains are independently simulated (instead of all 6500, as in the ‘full’ model). All other inputs are replaced by two single Poisson inputs (one excitatory and one inhibitory), with an equivalent effect. (That is, we use Brian’s ‘PoissonInput’). For details, see [2023-07-10__AdEx_Nto1_Brian_speedtest](#) (Cython), [2023-08-02__speedtest_brian_standalone_AdEx_Nto1](#) (C++), and [2023-08-09__Poisson_cquantile_upperbound](#) (Julia).

¹⁶ The Julia implementation is considerably faster than the implementation in ‘standalone’ Brian, which uses C++ and should thus be comparable in speed. This might be due to me not using Brian optimally, or to the fact that Brian is a general purpose spiking neural network simulator with much more capabilities; and our software was written specifically for the N-to-1 AdEx problem.

Spike-triggered averaging

As shown in [figure 2](#), our idea for connection inference rests on the causal link ‘presynaptic spike’ → ‘postsynaptic voltage bump’. I.e. we want to know for which neuron pairs a spike in one is reliably followed by such a bump in the other. The problem is that these bumps (the postsynaptic potentials or PSPs) are minute, and are easily drowned out by (1) other PSPs, (2) postsynaptic spikes, and (3) voltage imaging noise.

So, as is often done in neuroscience, we take the *average* over many instantiations, so as to hopefully find a signal in the noise. Specifically, we take spike-triggered averages, or STAs, of neurons’ voltage traces. If there is a connection from a neuron ‘M’ to a neuron ‘N’, then an STA of neuron N’s voltage imaging signal, based on neuron M’s spikes, would hopefully show the PSP.

And indeed, when we construct a few such STAs, we do see something resembling a PSP bump: [figure 20](#). We also find that the higher the firing rate of the presynaptic neuron, the cleaner the PSP-like shape is. This is of course because there are more presynaptic spikes and thus more windows to average over, which decreases the noise on the result. Finally, we see that inhibitory inputs result in downwards bumps in their STA, and excitatory inputs in upwards bumps.

Note that in this chapter – and the next one – we only look at the so called N-to-1 case ([figure 21](#)), where we simulate the voltage of one neuron, impinged on by N independent Poisson spiketrains. This is done for simplicity; it is only in the "Networks" chapter later on that we look at full networks, where inputs might be correlated with one another.

To use spike-triggered-averages as an actual connection test, we look specifically at the height of an STA, and compare it to a distribution of STA heights that we’d expect were the two neurons not connected. This is illustrated and explained in more detail in [figure 22](#). This so called ‘shuffle’ test yields the proportion p of how many shuffled (random) spiketrains yield an STA with a larger height than the real STA. In a following section ([Performance quantification](#)), we’ll use




Figure 20: **Example STAs in the 10' simulation with 6500 inputs.**

Note that every panel has a different y-axis (voltage) scale. The STA of the most active input is repeated in every panel (in faded blue), to allow a visual scale comparison nonetheless.

The inset legends indicate with how many presynaptic spikes the STA was calculated, and whether the input was an excitatory or inhibitory one.

The top right panel shows STAs of the 1st and 100th fastest spiking inputs, both within the excitatory inputs (blue shades), and within the inhibitory inputs (orange shades).

Source: [2023-09-13__Clippin_and_Ceilin](#).

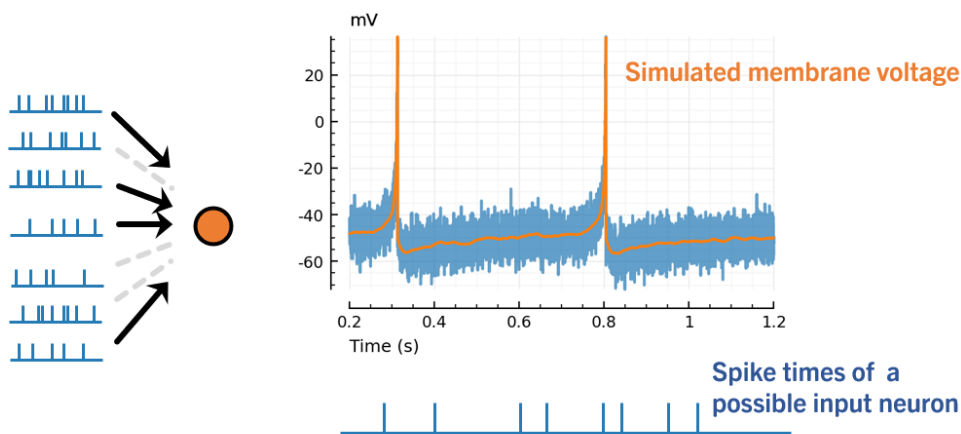


Figure 21: **The ‘N-to-1’ problem.**

Left: A neuron N (orange circle), and the spike trains of other neurons in the network (blue). Some of these other neurons impinge directly on N (black arrows), while others are not (directly) connected (dashed gray lines). Given only neuron N 's voltage signal and the other neurons' spiketrains, we want to detect the direct inputs, while rejecting the not-directly-connected spiketrains.

Right: The simulated membrane voltage of the impinged-upon neuron (orange), and the same signal with Gaussian noise added, to simulate a voltage imaging signal (blue). Underneath the plot, one of the possible input spiketrains, time-aligned to the voltage signal. This alignment is used later to extract spike-triggered windows from the voltage signal.

this number (as $t = 1 - p$) to make predictions and compare them to the ground truth.

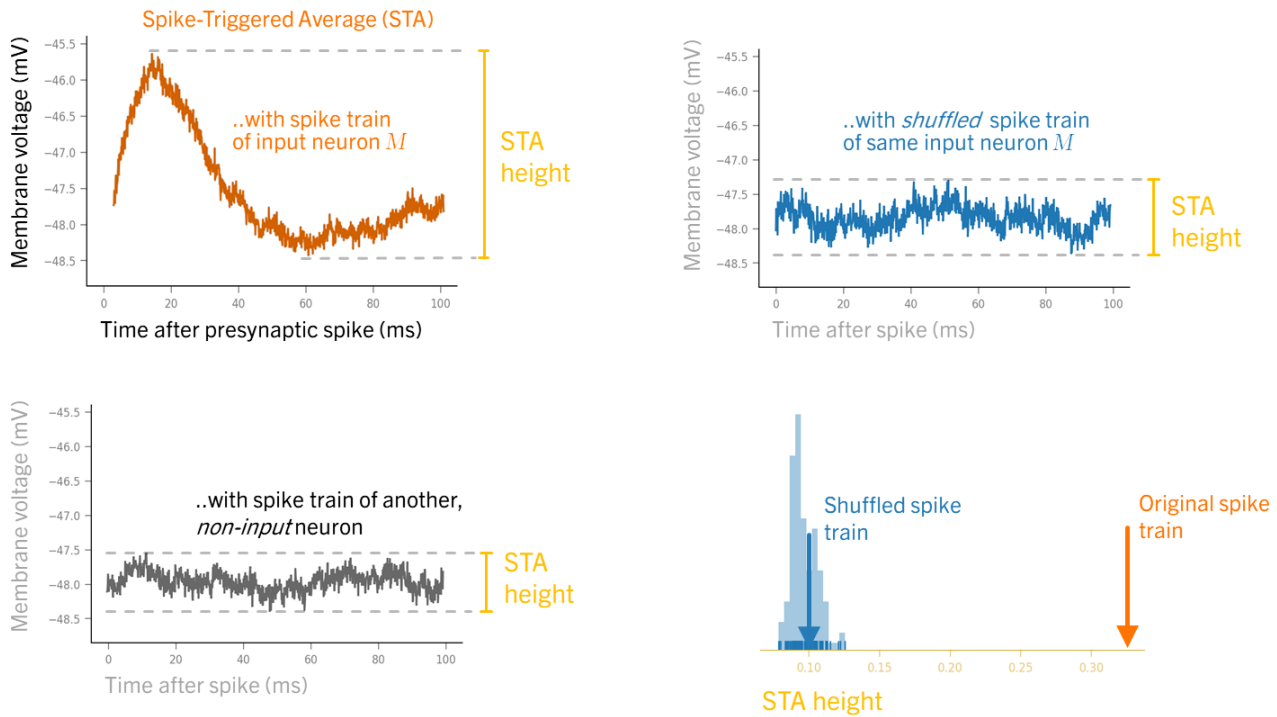


Figure 22: **A simple connection test: STA height with shuffle control.**

The spikes of a possible input neuron are aligned to the voltage trace of the neuron of interest N , as in figure 21. For every such spike, a 100-ms long window is cut out of the voltage of N . The average of all these windows is called the spike-triggered average (STA).

Left: Two example STAs of neuron N 's membrane voltage: one for an actually connected input neuron, M (top, orange); and one for a non-input neuron (below, gray). Given an STA signal x , we will use its height $h = \max(x) - \min(x)$ (also known as 'peak-to-peak' or 'ptp') to test whether two neurons are connected.

Right: An STA of N 's membrane voltage using a shuffled version of M 's spike times (which is made by randomly permuting the inter-spike-intervals of M). This 'shuffled STA height' provides a control for the STA height connection test statistic: "what do we expect the STA height to be if there is *no* connection $M \rightarrow N$ ". By calculating different such shuffles, we obtain a null-distribution for the STA height test statistic. And by comparing the real STA height to this distribution, we can calculate a p -value. Here, the real STA is larger than all shuffle controls, of which there are 100. So $p < 0.01$, and at $\alpha = 0.05$, we conclude there is indeed a connection $M \rightarrow N$.

3.1 Ceiling and clipping

As explained in [section 2.7](#), we modify our simulated voltage trace so that spikes have a consistent height. This modification has an effect on STAs, as is illustrated in [figure 23](#): the blue trace is the signal without spike ceiling, the orange one with spike ceiling. Their corresponding STAs are shown on the right. Note that the orange STA (made with ceiled spikes) is much noisier than the blue STA (from the unmodified voltage trace).

This suggests a relatively easy intervention to drastically de-noise STAs, and presumably increase their effectiveness for network inference: namely to remove the spikes from the signal.

We tried this ‘spike clipping’ and it indeed drastically denoised the STA; see the green signal and STA in [figure 23](#). We show that this decreased noise in the STA does indeed lead to an increase in network inference performance, by running a connection detection test without and with this spike clipping. The results are shown in [figure 24](#): detection performance increases from an AUC of 0.56 for the non-clipped voltage trace, to an AUC of 0.79 for the voltage trace with clipped spikes.

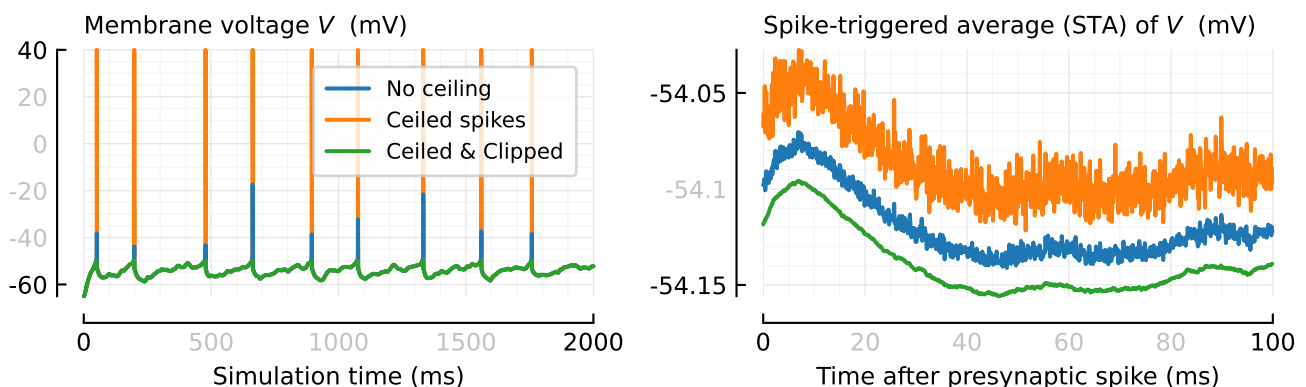


Figure 23: Example voltage traces and corresponding STAs, where the only difference is the height of spikes. In blue, the unmodified simulated voltage trace. In orange, the same, but with ceiled spikes (as in [section 2.7](#)). In green, the same as orange, but with the spikes clipped again after the ceiling (as explained in this section).

Source: [2023-09-13__Clippin_and_Ceilin](#).

STA connection test performance, for different voltage signal types

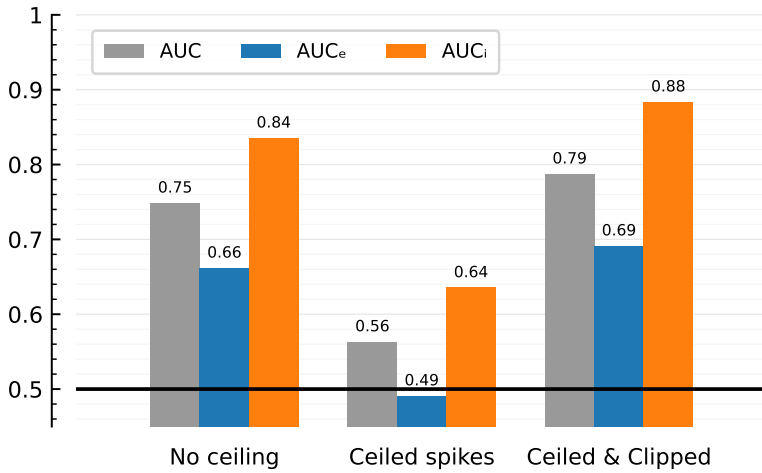


Figure 24: Connection detection performance for the three ways of handling spikes shown in figure 23: not modifying them; ceiling them; and clipping them. The area-under-the-curve or AUC measure is explained in the next section.

Source: [2023-09-13__Clippin_and_Ceilin](#).

3.2 Window length

We can freely choose the length of the windows that are cut from the postsynaptic voltage signal when computing an STA. In previous STA figures e.g, we chose 100 ms. This parameter is an important choice: a longer window linearly increases the computation time of the connection inference algorithm; and we find that window length choice strongly influences detection performance, as shown in figure 25.

Why does the shorter window length (20 ms versus 100 ms) lead to such a strong increase in detection performance? One reason is the heuristic we use to determine if a detected connection is excitatory

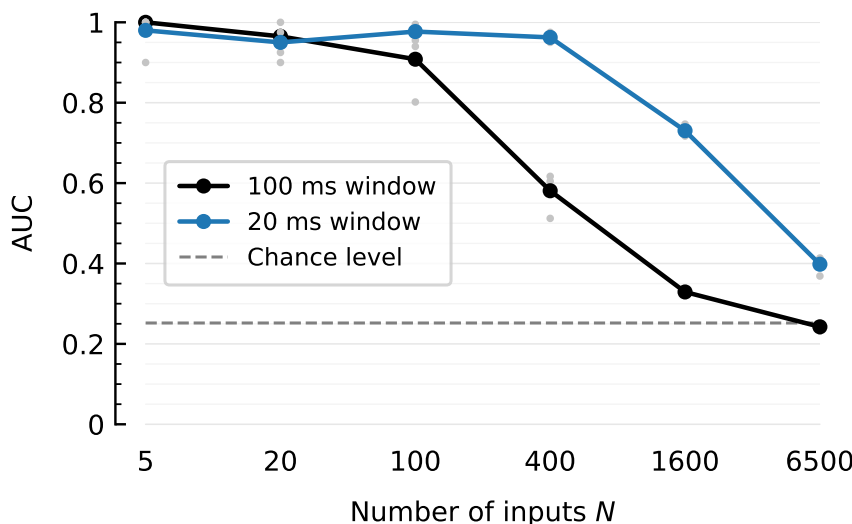


Figure 25: Connection detection performance of the STA method, for two different STA window lengths. Area under the ROC-curve, AdEx neuron, 10 minute recording, no voltage imaging noise, 5 different random seeds for input spiketrain generation (gray dots; larger dots are their means). All N inputs are tested (instead of just the highest firing ones), in addition to N random unconnected spikettrains.

Source: [2024-05-29__STA-winlength](#).

or inhibitory: the STA is referenced to its starting value and is then integrated. I.e. the area below the starting value is subtracted from the area above. If the result is positive, we say the connection is excitatory, and inhibitory otherwise. This heuristic works well in the low N (number of inputs) case. But for higher N , an interesting phenomenon happens: the broad downwards bump happening after the EPSP¹ becomes larger than the positive bump of the EPSP itself; see the top-left panel of [figure 20](#) for an example. This defeats the heuristic, and makes us classify excitatory connections as inhibitory and vice-versa.²

Changing to a shorter window length restores the heuristic's usefulness (in addition to decreasing computation time by a factor of four). A question that we did not yet explore is whether a longer window, but with a smarter polarity-heuristic, could perform better still.

3.3 Performance quantification

It is not easy to express in a single number how good a network inference algorithm is. Depending on what you find important as a user, different measures make more sense than others. This section looks at some measures to quantify the performance of our algorithms, and discusses the merits and disadvantages of each.

All the algorithms that we look at in this and the following chapter eventually output a single number per tested neuron pair (A, B): "How strongly do I believe there is a connection $A \rightarrow B$?" (And: "Is that connection excitatory or inhibitory?": the sign of the number). We will call this connected-ness number " t ".

To get actual predictions out of the algorithm, we must apply a threshold θ to these measures. If $|t| > \theta$, we classify the pair as connected, and as unconnected otherwise. For the detected pairs, we classify them as excitatory if t is positive, and inhibitory if it is negative.

Note that we use the same threshold for both excitatory and inhibitory connections. We could in fact use a different threshold – and we briefly look at this in [figure 30](#) – but for simplicity, we apply the same threshold for both types of connection.

Each threshold chosen yields a different tradeoff between recall and precision ([figure 26](#)). At low thresholds, we can detect more connections ("true positives"); but we will also detect more non-inputs as being connected (false positives). This also lowers our precision.

Some definitions:

- True positive rate (TPR), aka recall, sensitivity, and power: out of all true connections tested, how many did we correctly classify?

¹ This is what we model with the Gaussian 'dip' in [section 5.3](#).

² This is the reason that in an earlier version of this thesis, the template-based detection method of [section 5.2](#) had a performance *lower* than chance level for high N : inputs were detected correctly as connected, but with the wrong polarity. Changing to a lower window length has also drastically increased the performance of the simple STA-based method, to a point where it is almost as good as our newly developed methods ([figures 53](#) and [56](#)); in the previous version of this thesis, the difference in performance was much larger.




Figure 26: Lower detection thresholds increase both true and false positives. On the right, true positive rates are plotted against the false positive rate, to obtain the so called receiver operating characteristic or ROC curve. As the FPR increases more or less linearly with the decreasing detection threshold, both graphs look very similar. Source: [2023-09-13__Clippin_and_Ceilin](#).

- False positive rate (FPR): out of all distractors we added to our test (randomly generated spiketrains), how many did we wrongly classify as an actual input?
- Precision, aka positive predictive value (PPV): out of all the neuron pairs that we classified as connected, how many are actually connected?

There are more measures that quantify the performance of a binary classifier at a given threshold than those three (such as negative predictive value, false discovery rate, false omission rate, ...).³ But recall, FPR, and precision are commonly used ones.

Note that true positive rate (TPR) and precision are similar, in that they both count correct classifications. (They both have the number of true positives in their numerator). But recall (TPR) looks at the number of true positives from the point of view of the ground truth (how many did we find), and precision looks at it from the perspective of the experimenter (out of what this algorithm gives us, how much is correct?).

False positive rate and precision are also similar, in that they both measure the number of false positives. One advantage of using FPR over precision though, is that FPR does not depend on the number of distractors (unconnected spiketrains) that we add to our tests.⁴ Whereas we can arbitrarily increase precision by including less unconnected trains in our test – up to the limit of 100% precision, when we do not add any distractors and all tested trains are actually connected.

³ We do not actually perform binary classification: there are three classes (excitatory, inhibitory, unconnected). But it is not pure ternary (multiclass) classification either: we first classify as connected or not, and then (for the connected ones only), as excitatory or inhibitory. We could thus call it some kind of nested binary classification.

⁴ Besides that, the more distractors we test, the more accurate our estimate of the FPR will be.

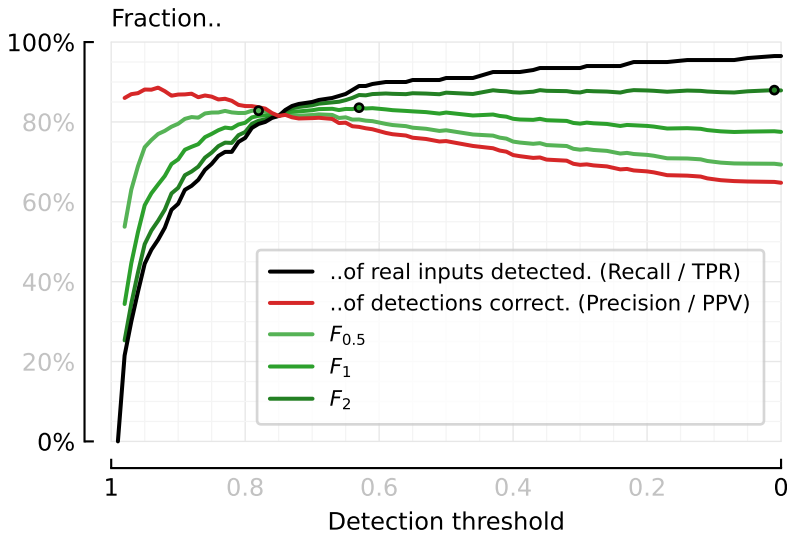


Figure 27: **Lower detection thresholds trade-off higher recall for lower precision.**

The F_β scores interpolate between the two measures. $F_{\rightarrow\infty}$ is recall, $F_{\rightarrow 0}$ is precision. The black circles on the F -curves indicate their maxima. Different trade-offs between precision and recall (different β -values) thus dictate different optimal detection thresholds.

In our tests, we choose the number of unconnected trains rather arbitrarily. (For example, when we test 100 excitatory and 100 inhibitory inputs, we also generate and test 100 unconnected spiketrains). A better way to choose this number of distractors might be to estimate what a realistic fraction of unconnected neurons would be in a typical voltage imaging experiment. Given some patch of brain tissue and one of the neurons in it, how many of the other recorded neurons in that patch will be connected to it? This is an interesting research question – and it is likely that answers can be found in the literature – but we do not explore it here.

In [figure 26](#), we have looked at TPR and FPR, both as a function of the detection threshold and as a function of each other. In [figures 27 to 29](#), we look at TPR (recall) and precision, again as a function of the detection threshold and as a function of each other.

Because recall and precision both increase for ‘better’ detectors, we might combine them into one measure. This is what the F -scores do: they are the harmonic mean of recall and precision, with recall and precision weighted differently depending on a parameter β . The F_β score attaches β times as much weight to precision P as to recall R :

$$F_\beta = \frac{(1 + \beta^2) \cdot P \cdot R}{\beta^2 \cdot P + R} \quad (25)$$

For $\beta = 1$, precision and recall are weighted equally. The F_1 -score is also the most widely used of the F -scores.

When we plot recall against precision, we get the so called PR-curves, shown in [figure 28](#) for both excitatory and inhibitory inputs together,

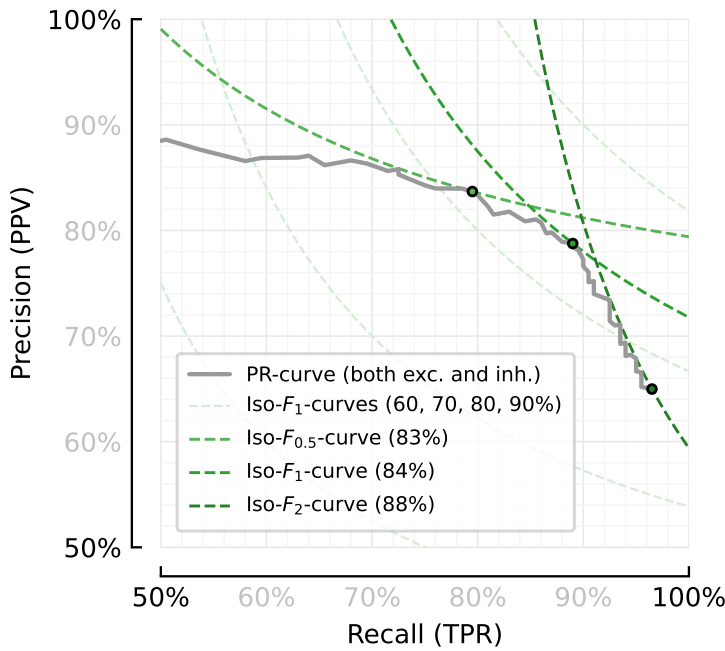


Figure 28: Precision plotted against recall for the STA-test in the $N=6500$ inputs, 10-minute-recording experiment. Black dots indicate where three different F_β -scores reach their respective maximum values.

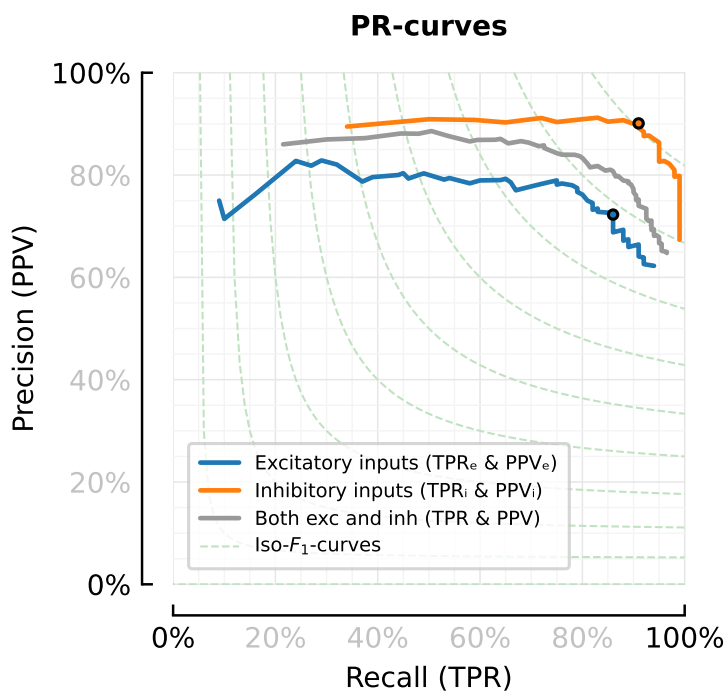


Figure 29: Same as in [figure 28](#), but with excitatory and inhibitory inputs analysed separately. Note that we can detect more inhibitory inputs than excitatory inputs for the same precision value (or for the same false positive rate, as shown in [figure 26](#)).

and in [figure 29](#) for both types separately.

Different thresholds yield different F_β -scores; but there is one threshold where your chosen F -score is maximal. This max F_1 -score is a good candidate for the single "how good is this detector" measure we were looking for. We specifically choose F_1 as it weighs precision and recall equally (and we have no a-priori reason to prefer any one over the other), and because it is the most frequently used.

Another common single measure to quantify a classifier's perfor-

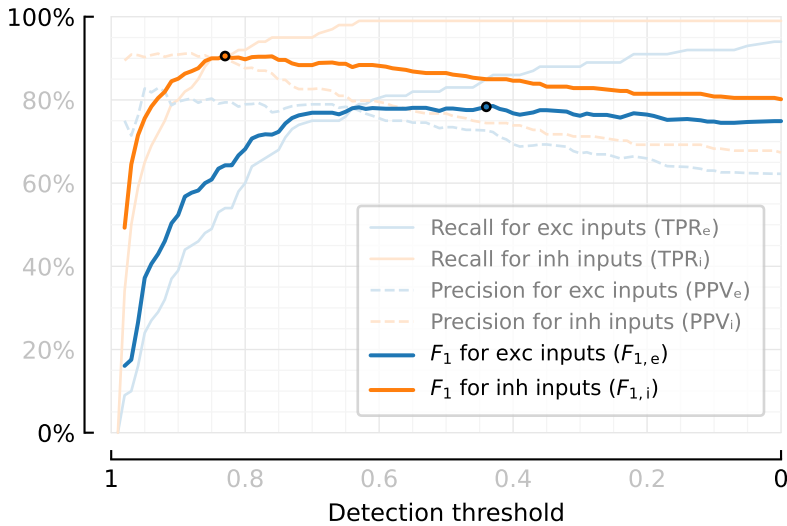


Figure 30: **Excitatory and inhibitory inputs reach $\max F_1$ at different thresholds.**

But for simplicity, whenever we use $\max F_1$ to evaluate a classifier, we will use only one threshold for both types of inputs, which will be a compromise between these two thresholds.

Performance of random connection classifier

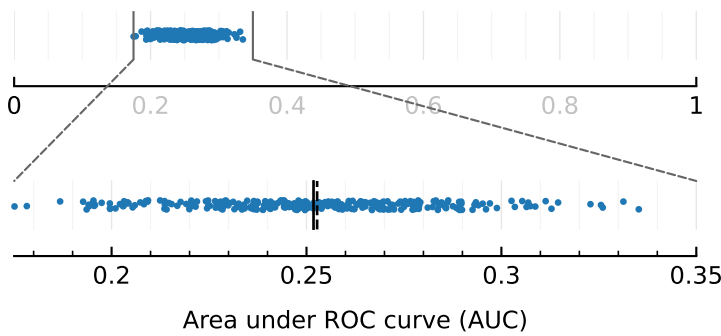


Figure 31: **Randomly classifying connections yields a chance level $AUC < 0.5$.**

300 random test results, and their performance as connection detector, quantified as area under their ROC curves. Solid black line is the mean, dashed line is the median.

In every of the 300 simulations, every connection (100 exc, 100 inh, and 100 unconnecteds) was assigned a random ‘t-value’ uniformly between -1 and 1 ; and then the classification threshold was swept over these t-values.

Source: [2023-09-20__STA_comntest_for_diff_recording_quality_n_durations.](https://doi.org/10.21203/rs.3.rs-2881111/v1)

mance is the area under its ROC-curve, or AUC, already shown in figure 26. A disadvantage of the AUC is that it is less interpretable as a number than the $\max F_1$ score (which immediately gives a rough idea of how many true connections you’ll detect, and how many of your detections are correct). A disadvantage of the $\max F_1$ score is that it uses the precision, which, as discussed above, is rather arbitrary in our setup. AUC does not suffer this problem: the FPR is independent of how many distractors are added to the test.

To make the AUC scores somewhat more interpretable, we compare them to the AUC of a random classifier: a detector that classifies possible connections randomly. For simple binary classification, this results in an AUC of 0.5. But because we have three classes (unconnected, excitatory, and inhibitory), the random AUC will be lower. We simulated such random classifiers to find the chance-level AUC (figure 31), which turns out to be about 0.252.

3.4 Recording duration & noise

In this section we look at how longer or less noisy voltage imaging recordings improve connection inference. In figure 32, the signal-to-




Figure 32: Performance drops to chance level for noisier signals.

All simulations were 10 minutes long. Signal-to-noise (SNR) values on the x-axis are approximately (but not exactly) log-spaced. An SNR of ' ∞ ' corresponds to no noise (i.e. the voltage signal straight out of the simulation, without any noise added). For every SNR value, five different simulations were run (gray dots), each with a different RNG seed for input firing rate and spiketrain generation. The mean performances of these five simulations are plotted with larger dots and are line-connected. Only the 100 highest-firing excitatory and inhibitory inputs were tested. An additional 100 unconnected spiketrains were generated and tested, with similar firing rates as those 200 high-firing real inputs. AUC chance level determined as in figure 31.

Source: [2023-09-20__STA_comntest_for_diff_recording_quality_n_durations.](#)

noise ratio (SNR) is varied, and in figure 33, we vary the recording duration.

As might be expected, noisier and/or shorter recordings decrease detection performance, down to chance level in the limit (namely: for noise levels almost as high as the spikes; and for recordings shorter than a minute). As to recording duration, interestingly, we do not yet see any flattening off of the detection performance curve for longer recording durations, up to the durations that we simulated (up to 1 hour).

For a concrete example of what a neuroscientist might expect from the STA-based connection test, we find that for a 10-minute recording with an SNR of 40 (which are more or less realistic for voltage imaging), the maximum F_1 -score – for the 200-highest firing inputs, and an additionally tested 100 unconnected spiketrains – is about 65%.

I.e. in the N-to-1 setup with 6500 inputs, for the 200 highest-firing of those inputs, we detect approximately 130 of them as being connected ($\pm 65\%$ recall). And of the spiketrains that we detect as inputs, about 65% are correctly classified ($\pm 65\%$ precision). I.e. 35% of them are either excitatory connections classified as inhibitory and vice-versa; or they are random unconnected spiketrains classified as real inputs.

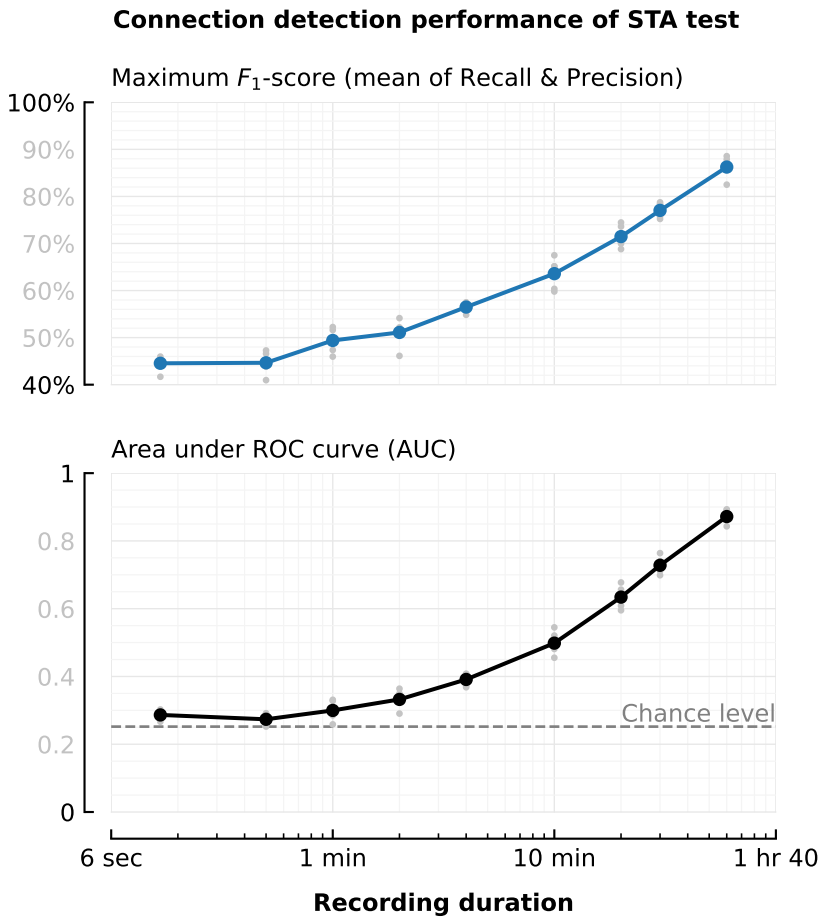


Figure 33: **Longer recordings allow more accurate connection inference.**

All simulations had a voltage imaging noise level (spike-SNR) of 40. The simulation ('recording') durations are on a logarithmic axis. (The first two data points are at 10 and 30 seconds; the last one is at 1 hour). For more, see figure 32's caption.

The area under the TPR/FPR-curve (AUC) for this recording duration and quality is about 0.50 (compared to the chance level of 0.25 – see figure 31).

3.5 EI balance

In the rest of this thesis, we choose to have a fixed EI-balance (4:1 excitatory to inhibitory neurons, with inhibitory neurons 4× stronger) and fixed excitatory and inhibitory reversal potentials. In this section however, we investigate these parameters, and wonder whether changing them has any effect on network inference performance.

Note: The results in this section were obtained before we switched to using an AdEx neuron model. They still use the Izhikevich neuron model (but with comparable parameters – which were calculated using the correspondences found in section 2.3). Second, in the rest of this chapter, we use a high number N of input spiketrains (6500 in total). In this section however, we use just 27 input spiketrains.⁵ As might be expected, using a lower number of inputs makes connection inference much easier, as can be seen in figure 34 (note that here, we test *all* input spiketrains, and not just the highest firing ones as before. This gives a lower overall performance).

In the first subsection below, inhibitory and excitatory connections




Figure 34: **STA performance for different N .**

10-minute 4:1 EI-balanced simulations using the Izhikevich neuron. 16 simulations per condition, for different random number generator seeds. Source: 2022-05-02__STA_mean_vs_peak-to-peak.

⁵ Part of the reason for this initially lower number of spiketrains is that we still used our simulator written in Python and Numba here, which was slower than our later Julia simulator. We also still connection-tested *all* inputs, and not just the highest firing (or a random subset). This made iterating with high N prohibitively slow.

are equally strong ($\Delta g_{\text{exc}} = \Delta g_{\text{inh}}$). In the second subsection, the strength of inhibitory connections is varied.

Influence of inhibitory reversal potential

Connections are still detected by spike-triggered averaging (STA) and testing the STA height of a spike train against the STA height of shuffled versions of the train – both for excitatory and inhibitory inputs. Performance of this algorithm is shown in [figure 35](#), for varying proportions of inhibitory versus excitatory inputs, and for two different inhibitory synaptic reversal potentials.

We find that we can reliably detect the inhibitory as well as the excitatory connections – but only when the inhibitory reversal potential lies below the neuron’s resting potential, or when there are no or very few excitatory inputs.

[Figure 36](#) shows why this is: when the neuron is spiking, the reversal potential $E_{\text{inh}} = -50$ mV lies around the neuron’s median voltage (which is -50.5 mV for $p_{\text{inh}} = 0.1$, and which we’ll call v_{rest} here).⁶ Any inhibitory spike will thus not have any influence on the total synaptic current and hence the neuron’s voltage. This is reflected in a flat STA, and hence, undetectable inhibitory connections.

⁶ Note that this is not the "resting potential" parameter v_r of the Izhikevich simulation, which is -60 mV.

When the inhibitory reversal potential lies below this median voltage however, an inhibitory STA is visible even when the neuron spikes. Similarly, when the neuron does not spike, its median voltage is lower than E_{inh} , and an *upwards* inhibitory STA is visible.

We can thus conclude that, for connections to be detectable via the STA test, their synaptic reversal potential may not lie too close to the neuron’s median voltage.

In real neurons, the reversal potential of inhibitory inputs lies below the neuron’s median membrane voltage. We thus set E_{inh} to -65 mV in the rest of this section. This is the approximate reversal potential of Cl^- , the ion for which GABA_A channels – the main inhibitory receptor – are permeable [[DA01](#); [Kan+13](#)].

Different synaptic conductances

The synaptic conductance bump Δg_{syn} is now split into two different values: Δg_{exc} and Δg_{inh} , for excitatory and inhibitory connections respectively. Δg_{exc} is fixed at 0.4 nS. This value is chosen so that the output neuron has a realistic firing rate (less than one spike per second, but more than zero). Δg_{inh} is varied to be different multiples of this.

The ratio of inputs of either type is fixed at the physiological value of 4:1 excitatory:inhibitory, i.e. $p_{\text{inh}} = 0.2$. The fraction of connected inputs $p_{\text{connected}}$ is set to 0.7 . All other simulation parameters stay the same.

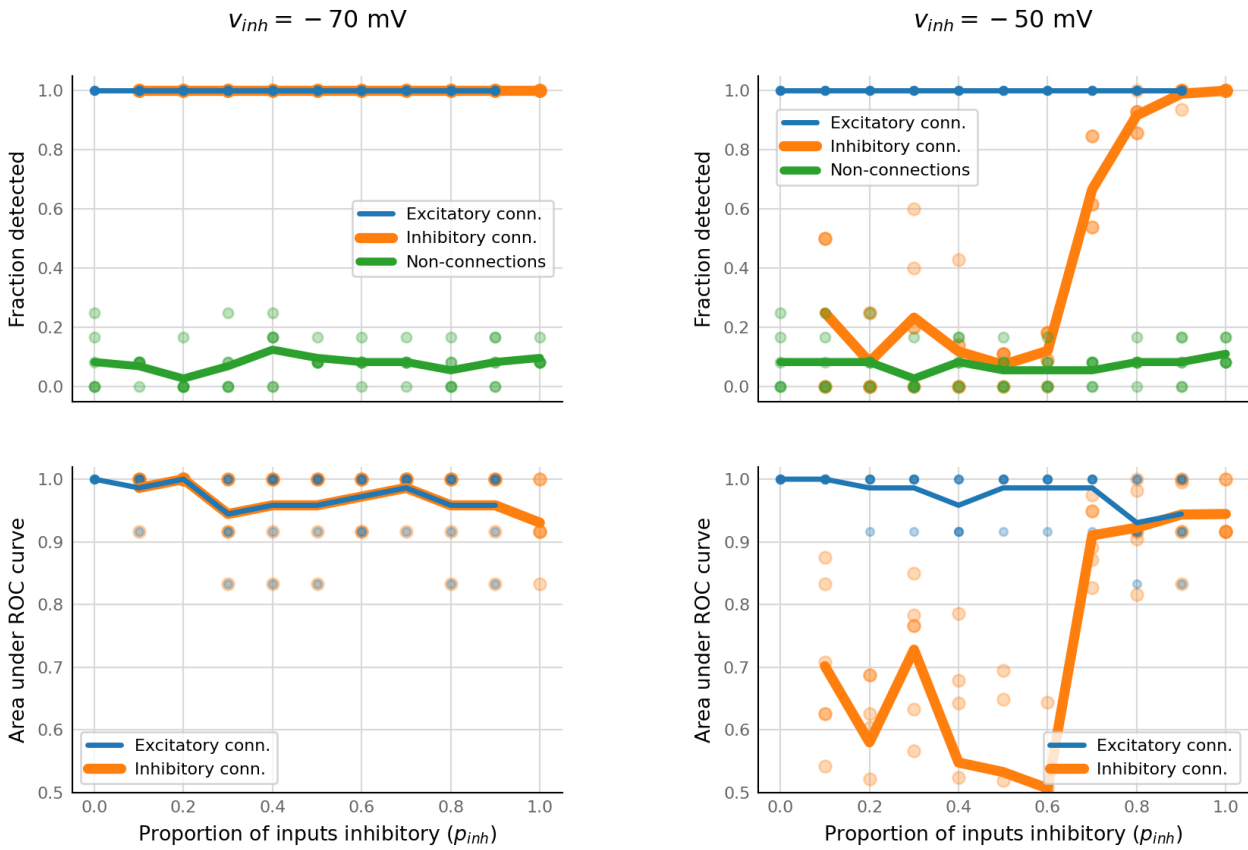


Figure 35: **STA detects inhibitory connections, but only when $E_{inh} < v_{rest}$ or when there are no excitatory inputs.**

Connection detection performance of spike-triggered averaging at varying proportions of inhibitory inputs, and for two different inhibitory synaptic reversal potentials E_{inh} . *Top*: Performance at a fixed detection threshold of $p < 0.05$ (where p is the proportion of shuffled spike trains with an STA larger than the real spike train). Plotted are the true positive rates (separately for excitatory and inhibitory connections) and the false positive rate. *Bottom*: Performance over all detection thresholds, as area under the ROC curve. As TPR is calculated separately for excitatory and inhibitory connections, there are also two ROC curves. For each condition, the simulation is run six times, each time with a different random number generator seed. The semi-transparent circles are individual simulations. The thick lines average over the seeds.

Source: [2021-09-16__vary_E_vs_I.](#)

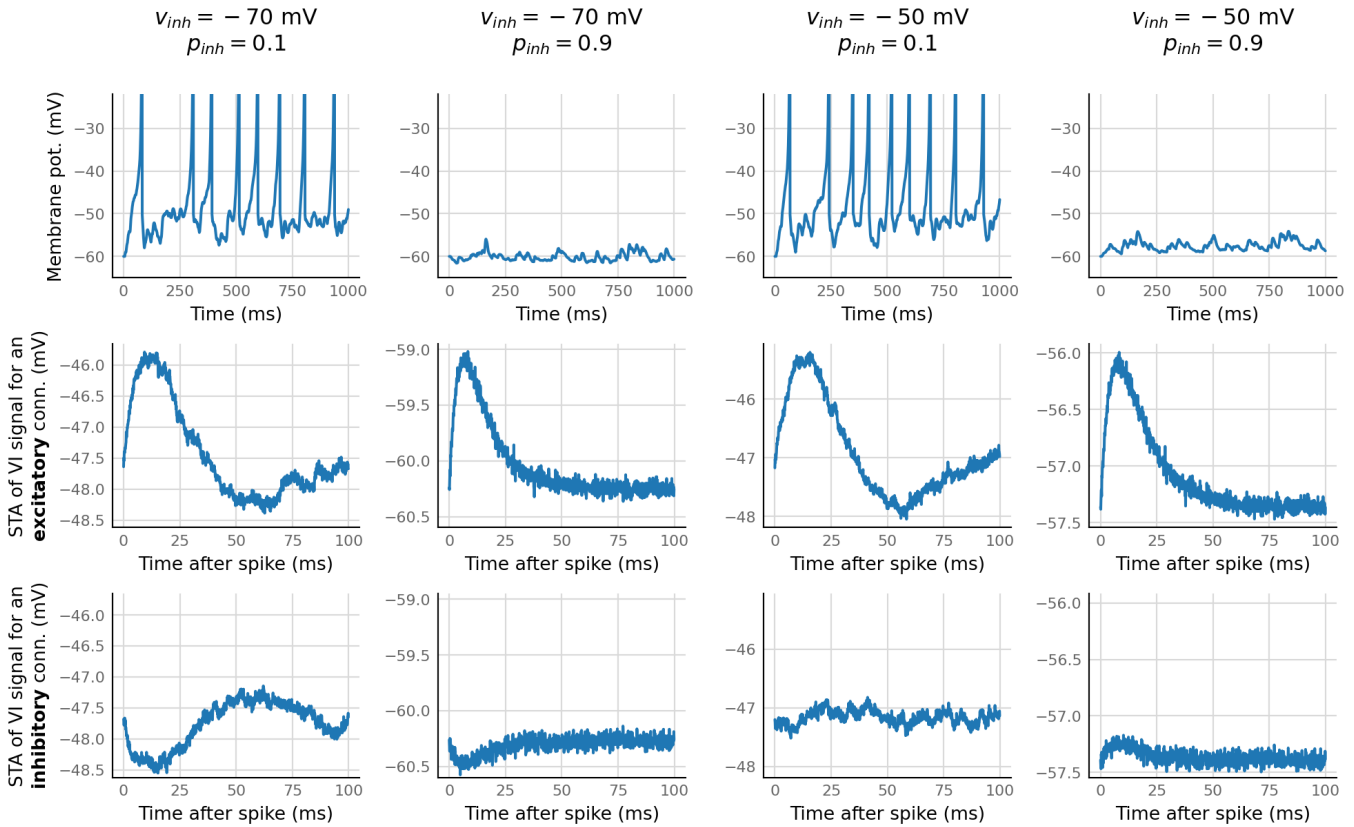


Figure 36: **The inhibitory STA is flat when the reversal potential is close to the neuron’s median voltage.**

Excerpts from the Izhikevich simulation output (top row), and spike triggered averages of the voltage imaging signal (bottom two rows). The top voltage traces are shown without any VI noise for clarity. The actual STA’s are in fact calculated on the noisy VI signal. Note that the y-axis extents of the STA’s differ between columns (but are consistent within one).

Source: [2021-09-16__vary_E_vs_I.](#)

We find that the STA method can detect connections when Δg_{inh} is at least as strong as Δg_{exc} (figure 37). This is straightforwardly explained by larger synaptic conductances causing larger PSP’s, and thus larger STA’s (figure 38).

More connections versus stronger connections

In the previous two subsections, we varied either the proportion of inhibitory inputs p_{inh} , or the inhibitory synaptic conductance Δg_{inh} . Here, we vary both simultaneously.

We fix the number of excitatory inputs N_{exc} to be 24, and set the number of inhibitory inputs N_{inh} to be a (sub)multiple of that. (Note that this changes the total number of inputs, which was fixed at 30 up until now). The synaptic conductance bump per excitatory spike Δg_{exc} is fixed at 0.4 nS as in the previous subsection, and Δg_{inh} is set to a multiple of that.

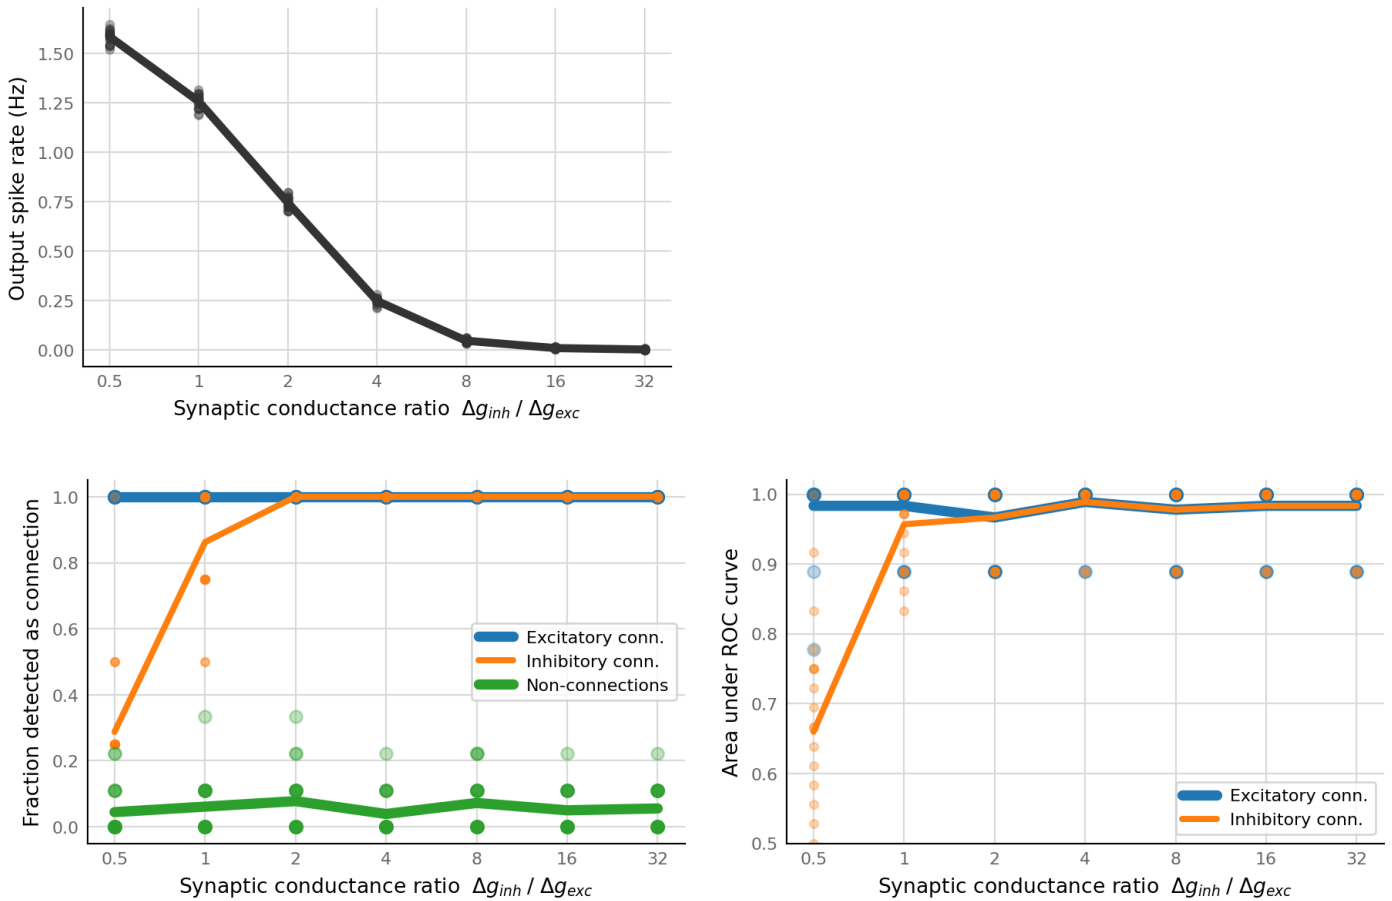


Figure 37: **STA detects inhibitory connections for most conductance ratios, and with or without output spikes.**

Legend as in figure 35, but now with twenty simulations per condition, instead of six. Note that at a p -value threshold of 0.05 (bottom left panel), we expect to see a false positive ratio of 0.05, as can indeed be seen for most conductance ratios. With lower thresholds, we can decrease this FPR without decreasing recall, as evidenced by the AUC scores ≈ 1 . Source: [2021-11-05__vary_syncond_ratio](#).

The multiples are chosen so that on the diagonal of figure 39, the ratio of inputs is the inverse of the ratio of connection strengths: $\Delta g_{inh} / \Delta g_{exc} = 4$ is matched with $N_{inh} / N_{exc} = 1/4$, etc.

We find that the spike triggered averaging method can detect both excitatory and inhibitory connections, over the entire parameter grid – except, as found in the previous subsection, for weak inhibitory connections. In cortex, inhibitory connections are generally stronger than excitatory connections, so we should not find ourselves in this regime. Note also that we can detect incoming connections even when the neuron does not spike at all. This is a major advantage of voltage based connection detection versus spike-to-spike connection detection.

Focusing on the diagonal of figure 39 (highlighted in figure 40), we notice that, notwithstanding the balancing, the firing rate increases for stronger but fewer connections. The reason is shown in figure 41:

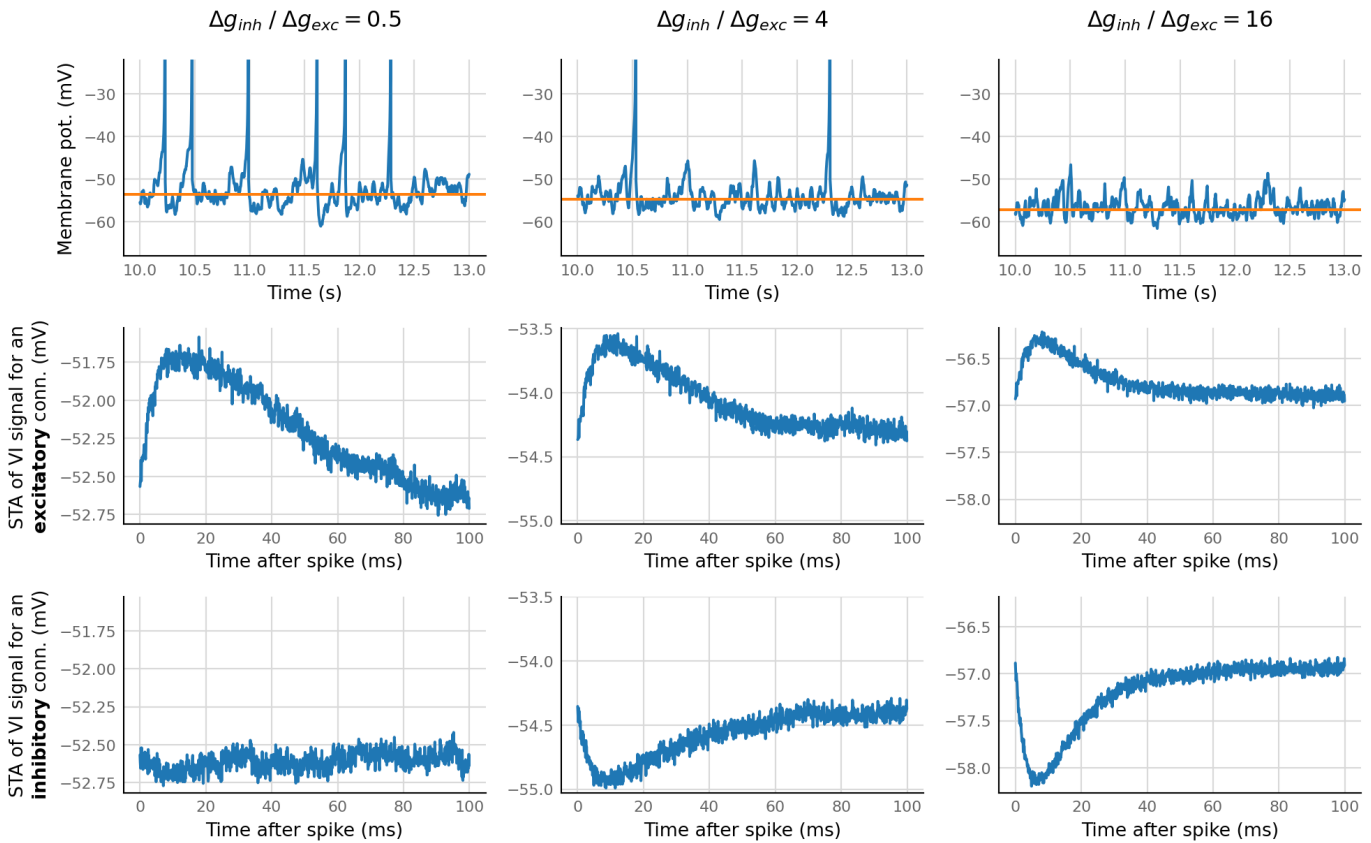


Figure 38: STA height increases with synaptic conductance. Orange lines indicate the neuron’s median voltage over the entire recording duration.

Source: 2021-11-05__vary_syncond_ratio.

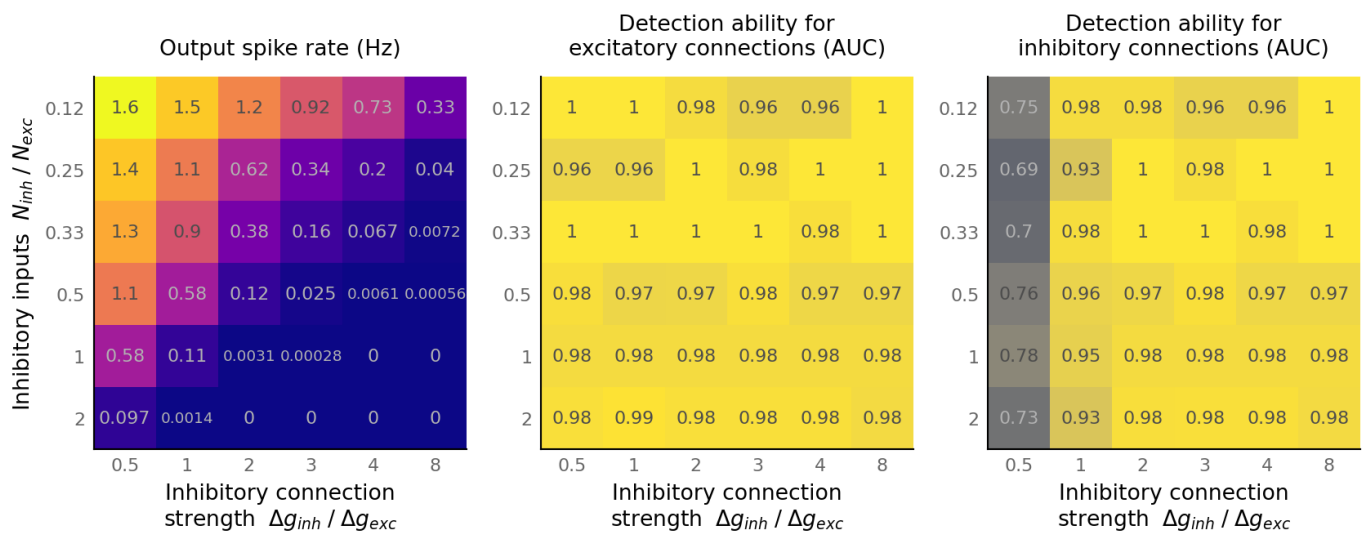


Figure 39: STA detects connections for weak or strong, sparse or dense inhibition.

Output firing rate (left) and connection detection performance (center and right) for different amounts and strengths of inhibitory inputs. Every cell is the mean of 6 simulations.

Source: 2021-11-11__vary_both_inh_strength_and_proportion.

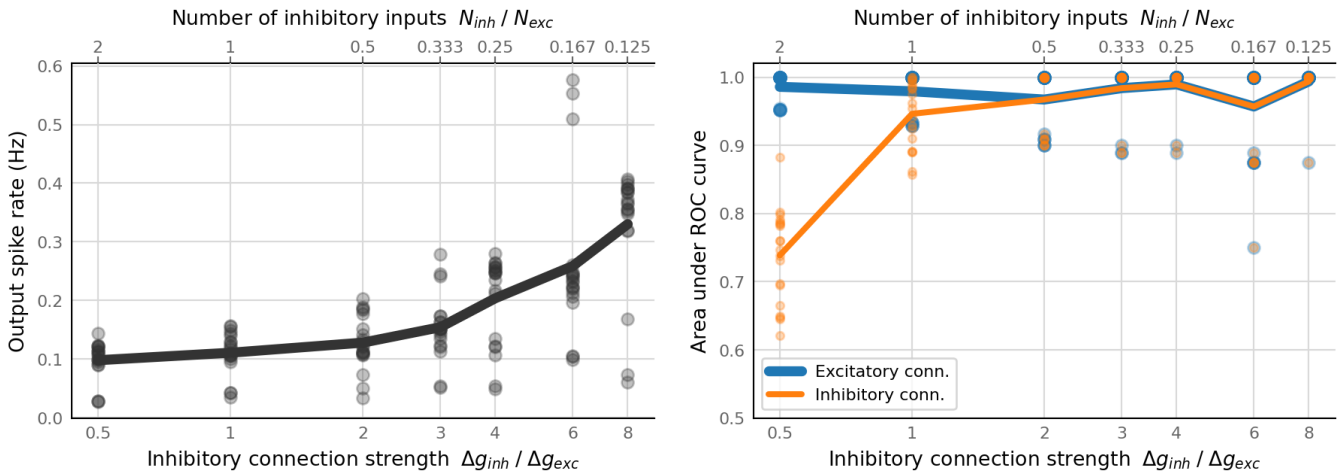


Figure 40: **Firing rate increases with stronger (but fewer) inhibitory inputs.**

Excerpt of figure 39 along the diagonal (but with 20 simulations per condition rather than 6). Dots are individual simulations, lines are their means.

Source: [2021-11-11__vary_both_inh_strength_and_proportion.](#)

in these balanced regimes, the neuron can only fire when the inhibition is temporarily low enough for the excitation to overcome it, and make the neuron spike. When there are fewer inhibitory inputs, there will be more and longer such inhibition gaps (even though when there is inhibition, it is stronger, making the mean inhibition equal, as shown figure 41 right). This nonlinear phenomenon is typical in spiking simulations. In a mean-field model, where only firing rates are simulated, this would not occur.

Looking at individual columns or rows of figure 39, we find that the firing rate gently decreases above the diagonal, but below the diagonal it falls sharply to zero. This is again a consequence of simulating spikes: increasing the net excitation changes nothing about a zero firing rate, until the firing threshold is reached, above which net excitation and firing rate do change in lockstep.

Finally, we explain the high spread / multimodality in the firing rates of figure 40. At the beginning of each simulation, when determining the number of input spike trains of each type (connected or not connected; excitatory or inhibitory), stochastic rounding is used. Say that there are 3 inhibitory inputs (i.e. 1/8th of the 24 excitatory inputs, as on the far right of figure 40). With our current connection probability $p_{connected}$ of 0.7, we should have $0.7 \times 3 = 2.1$ connected inhibitory inputs. With stochastic rounding, in 90% of simulations there will be 2 connected inhibitory inputs, and in 10% of simulations there will be 3. Those 10% simulations with 3 inhibitory connections are responsible for the cluster at the bottom right of the firing rate graph of figure 40.




Figure 41: **The neuron spikes during gaps in inhibitory input.**

Excerpts from the membrane potential (left) and the total synaptic conductance of all inhibitory inputs (right) of two simulations. In the blue simulation, inhibitory inputs are 8 times as strong as excitatory inputs, but there are 1/8th times as many. In the orange simulation, the inhibitory inputs are as strong and as numerous as the excitatory inputs. Horizontal lines are means of the conductances over the entire simulation duration. Note the blue conductance gap around $t = 10.4$ seconds, allowing a spike.

Source: [2021-11-11__vary_both_inh_strength_and_proportion.](#)

3.6 Computational cost of STA test

The brunt of the time in using the STA as a connection test is in actually calculating these STAs – and the STAs of the shuffled versions of the tested spiketrains.

[Table 4](#) lists different parameters that determine how long a connection test takes to run.

The chosen inputs to test (300 high firing trains) have a median firing rate of 16 Hz. I.e. at a simulation duration of 10 seconds, there are about 160 presynaptic spikes per tested connection. There are 101 times that many STAs to calculate per connection: once for the real spiketrain, and a 100 times for shuffles of it. For a 10 second simulation, there are thus about 16k STAs to calculate per connection. For 10 minutes: 967k STAs. For 1 hour: 5.8M STAs.

The computation time of the STA-based test thus scales linearly with the voltage imaging recording duration. We find that testing 300 possible spiketrain inputs to one neuron takes about one-fifth of the time of the recording ([figure 42](#)): about 2 minutes of computation time for a 10-minute recording.

Factor	Description	Value	Unit
N_{post}	Number of analyzed voltage signals (number of postsynaptic neurons)	1	voltage signals
N_{pre}	Number of tested spiketrains (possible presynaptic neurons) per 'post' neuron	300	spiketrains
$N_{\text{conn}} = N_{\text{post}} \cdot N_{\text{pre}}$	Number of connections tested	300	connections
T	Duration of simulation or recording	10	minutes
λ	Firing rate of presynaptic neuron	16	spikes / second
$N_w = T \cdot \lambda$	Number of windows per tested conn.	9600	windows
T_w	Window length	20	ms
Δt	Timestep of simulation or recording	0.1	ms
$f_s = 1/\Delta t$	Sample rate	10	samples / ms
$M = T_w \cdot f_s$	Number of samples per window	200	samples

Table 4: Factors in how long a connection test takes to run. The listed values are ones typically used in this thesis.




Figure 42: Test time scales linearly with voltage signal duration. Simulation timestep ('sample time') of 0.1 ms. STA length of 20 ms; i.e. 200 samples. Black dots are the means over five simulations per duration. Compute times for individual simulations are plotted with gray dots; but the variation is so small that these gray dots are hidden behind the black means. Gray dashed line is the $y = x$ identity. Source: [2023-09-20_STA_comntest_for_diff_recording_quality_n_durations](#).

3.7 Conclusion

Using spike-triggered-averages of the postsynaptic voltage is an obvious way to look for post-synaptic potential bumps, and to thus detect neuron connections.

In this chapter we have found that this STA method indeed does significantly better than chance at detecting connections, under realistic conditions (6500 EI-balanced inputs to one neuron, 10-minute recording, voltage imaging SNR of 40). Of the highest firing inputs (100 excitatory, 100 inhibitory), we detect about 65%, at a precision level where 65% of our detections are correct. The AUC is about 0.50.

This is significantly better than chance (AUC at chance level is 0.252), but there is a lot of room for improvement. In a later chapter ([chapter 5](#)), we will look for improved connection detection methods. But first, we take another look at our test setup.

In this chapter, only one neuron's voltage was simulated, while being stimulated by N independent Poisson spiketrains. Real neural networks are often more recurrent and less feedforward than this, and different inputs to the same neuron can be correlated. This is why, in the next chapter, we simulate the voltages of multiple neurons, connected to each other in a fully recurrent neural network. In particular, we want to test the problem of indirect connections that are detected as direct connections – a well-known problem in the spike-to-spike connection inference literature [[Orl+17](#); [DF20](#)].

Network model

4.1 Introduction

In the previous experiments, only one neuron's voltage was simulated. The inputs were Poisson spike trains. In the next experiments, we simulate the voltages of a full network of neurons, which are recurrently connected to each other. The goal is to investigate the effect on network inference of potentially correlated inputs and indirect connections.

4.2 Connectivity structure

We choose the simple and common 'fully random' connectivity rule,¹ where any neuron has a connection to another with a uniform random probability (we choose $p_{\text{conn}} = 0.04$). After generating an adjacency matrix this way ($A = \text{rand}(N, N) \cdot p_{\text{conn}}$, where rand draws from $\sim U[0, 1]$), we remove autapses. We choose the number of neurons $N = 1000$.

A property of fully random networks is that they are strongly interconnected. In our network, any neuron is reachable from any other in at most three steps (three synapses); most are reachable in just two. This is exemplified in [figure 43](#) and [figure 44](#): one selected neuron (neuron '1' here) is reachable in two hops from more than 700 of the 1000 total neurons in the network. And when we compute the shortest path between every possible neuron pair, we find a very similar distribution ([figure 45](#)).

4.3 External input

As we no longer have Poisson spike trains providing input to our neurons, we need another way of bootstrapping activity in the network.

Instead of external spikes, each neuron is provided with external input by adding Gaussian noise to its membrane voltage. Every time step ($\Delta t = 0.1$ ms), a sample drawn from a normal distribution with mean -0.5 pA and $\sigma = 5$ pA is added to the membrane current. (As




Figure 44: **A selected neuron is reachable in at most three hops.** Shortest path lengths from every other neuron in the network.




Figure 45: **Every neuron is reachable in at most three hops.** Shortest path lengths calculated using the Floyd-Warshall algorithm. Source: [2022-07-14_Unconnected-but-detected.html](#).

¹Other common choices for connectivity structure are scale-free networks, and 'local' networks.

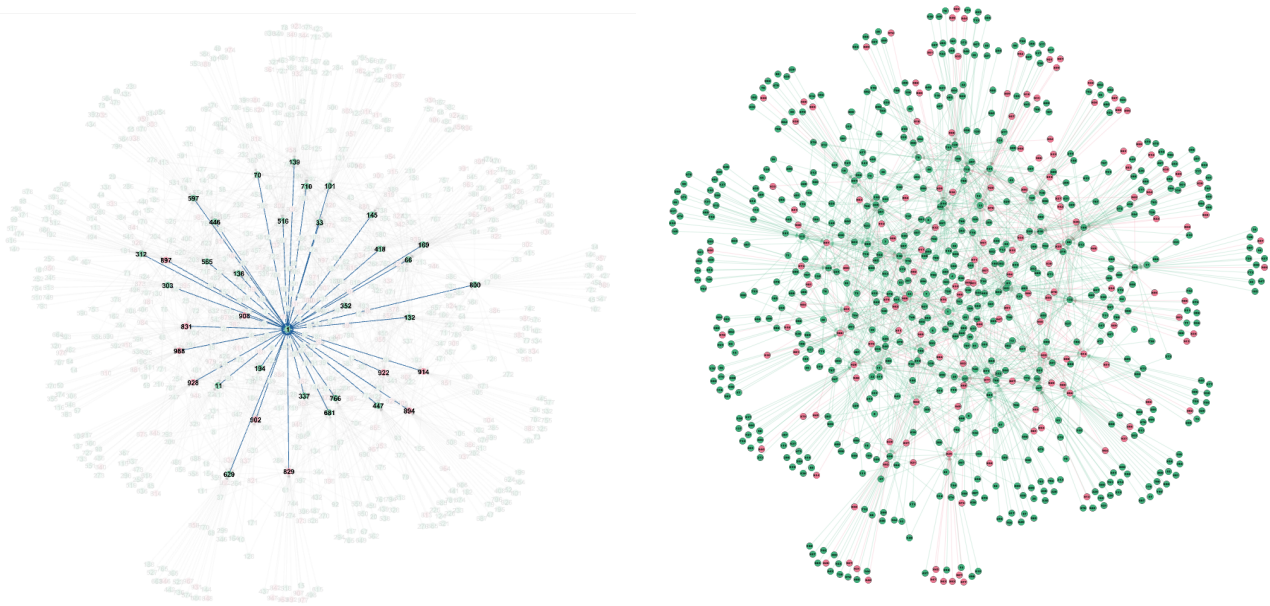


Figure 43: **A neuron is reachable in two hops from most other neurons.** Selected neurons and connections in our random network. Left shows the direct inputs to one of the neurons. Right additionally shows the direct inputs to these inputs. This subnetwork already contains more than 700 of the 1000 total neurons in the network.

Excitatory neurons in green, inhibitory in red. The tangle in the middle consists of neurons that synapse onto multiple of the direct inputs of our selected neuron.

Visualization using the ‘Gephi’ software, with the ‘Yifan-Hu’ layout algorithm, and default parameters otherwise. Source: [2022-08-29__Visualizing_subnets](#).

membrane current is by convention negative, this corresponds to an on-average positive influence on membrane voltage).

Like in the first footnote in [section 2.8](#), this way of adding noise makes the noise power dependent on the timestep. A more principled approach to injecting noise would be to replace our ODE for the neuron voltage (equation (1)) by a stochastic differential equation (SDE); though we did not explore this here.

4.4 EI balance

Similar to the N-to-1 experiments, we make 1 out of 5 neurons inhibitory. As before, this is done by setting the synaptic reversal potential at the outputs of inhibitory neurons to -80 mV (instead of the 0 mV for excitatory neurons). To make sure that each neuron receives a balanced mix of excitation and inhibition, and given that there are 4:1 excitatory to inhibitory neurons, we make excitatory neurons 4x weaker: their synaptic strength (Δg , the instantaneous increase in postsynaptic conductance g on spike arrival) is 4x as small as that of inhibitory neurons.²

² Note that this is only an “EI-balance” at the population level: the input to individual neurons will not be exactly balanced and randomly have more excitatory or more inhibitory input.




Figure 46: **Firing rates in the network with '1144' weights.**
 Left: rasterplot showing all spikes in the first 10 seconds of the simulation.
 Right: histogram of firing rates.
 Source: [2022-07-01__g_EI](#).




Figure 47: **Firing rates in the network with 'Roxin2011' weights.**
 See [figure 46](#).

[Figure 46](#) shows results of the network simulation as spiketrains and firing rate distributions, using the described synaptic weights rule (inhibitory synapses $4\times$ stronger than excitatory ones). We find that the resulting firing rate distribution is fairly symmetrical, and that excitatory and inhibitory neurons have very similar firing rates.

In real neural networks however, inhibitory neurons often have higher firing rates than their excitatory neighbours. In addition, real firing rate distributions are often heavy-tailed and not symmetrical, as we saw back in [section 2.6](#). So, to coax more realistic firing rate distributions out of our network, we looked at the previously mentioned modelling paper of Roxin *et al*, "On the Distribution of Firing Rates in Networks of Cortical Neurons", [[Rox+11](#)].³ In this paper different synaptic weights are used than the ones described above. Normalizing excitatory-to-excitatory (E→E) connections to 1, Roxin *et al*'s synaptic weights are:

³This paper seeks to explain how heavy-tailed firing rate distributions emerge in randomly connected spiking neural networks. We wanted to emulate the simulations in this work, to obtain a heavy-tailed firing rate distribution.

E→E: 1
 E→I: 18 (instead of 1)
 I→E: 36 (instead of 4)
 I→I: 31 (instead of 4)

We will call these weights ‘Roxin2011’. The naive, ‘balancing’ weights used in [figure 46](#), we will call ‘1144’.

Simulating a network with the ‘Roxin2011’ weights indeed increases the firing rates of inhibitory neurons with respect to those of excitatory neurons ([figure 47](#)). The two firing rate distributions are however both as symmetrical as before.⁴

This imbalance between E→E and E→I synapses is not that rare: another modelling paper that also investigated EI-balance, [SC21], has the following synaptic weights. If our original weights are [1, 1, 4, 4], then Sadeh and Clopath’s weights are [1, k , k , k], with $k = 4$ (in their “strong E-I coupling regime”).

In the results that follow, we use both ‘1144’ and the ‘Roxin2011’ synaptic weights.

Finally, we also looked at drawing synaptic weights from a distribution (specifically, a lognormal one) instead of setting all weights of one type to be equal. This did not change the symmetry of the resulting firing rate distributions, so we did not pursue this further.

4.5 Subsampling

We simulate all 1000 neurons’ voltages, but, to save memory and disk space, do not record all these traces.⁵ In most experiments here, we recorded the voltage traces of 40 excitatory and 10 inhibitory neurons (5% of all neurons).

Additionally, when performing connection tests on the inputs of a recorded neuron, we do not test the spiketrains of all 999 other neurons. Instead, we test only a (biased) sample of the possible inputs, to save processing time. This sample is constructed as follows. We test all the a-priori known true direct inputs – both excitatory and inhibitory – and add a random sample of 40 not-directly-connected neurons.

On average, each neuron has ± 32 excitatory inputs and ± 8 inhibitory inputs. This means that, from the 1000×1000 possible connections, we only test about 4000, or 0.4%.⁶

4.6 Connection testing

We tested connections using the STA height (or ‘peak-to-peak’) test. In the chronology of the PhD, the more advanced methods presented in the previous chapter were developed only after the network tests

⁴ When collating the two firing rate distributions from both neuron types, the resulting single distribution is heavy-tailed, and even vaguely looks log-normal (but it is not). In [Rox+11], similar-looking, collated firing rate distributions are shown, and are called lognormal (but this is never quantified in the paper).

⁵ For a 10-minute simulation with a timestep of 0.1 ms, one voltage trace takes 48 MB (at 64 bit per sample). Our 1000 neurons thus take 48 GB – and that is just for one simulation (one set of parameters).

We *do* record the spike trains of all neurons. Saving just spike times takes considerably less space: a neuron spiking at 10 Hz for 10 minutes emits 6000 spikes, which, at 64 bit per timestamp, takes just 48 kB. In other words, at a 0.1 ms sample rate, a spike train occupies about $1000 \times$ less memory than the corresponding voltage trace.

⁶ Calculation behind these numbers:
 - 1000 neurons \times 80% excitatory \times 4% probability of an input connection = 32 excitatory inputs on average.
 - 50 ‘post’ neurons (40 excitatory and 10 inhibitory voltage-recorded neurons) \times 80 ‘pre’ neurons (± 40 connected + 40 unconnected) = 4000 tested connections.
 - To be precise, instead of 1000×1000 , there are rather $1000^2 - 1000$ possible connections, as we would not test for autapses.




Figure 48: Performance of the STA ‘peak-to-peak’ connection test in the random network. 10 minute recording, with ‘1144’ weights. Each dot is the true positive rate for one ‘post’ neuron. The gray dotted line indicates the p-value threshold α of 0.05.

Source: [2022-09-01__1144_weights](#).

in this chapter were done. Applying the newer methods to the full network (and not just the N-to-1 setup) is a topic for further work.

Figure 48 shows the performance of using the STA height as connection test in the network. The detection rates (one for each tested ‘post’ neuron and ‘pre’ neuron type) are shown at a fixed detection threshold α of 0.05. That is, a connection was classified as real if its STA was larger than 95 of its 100 ‘control’ STAs, which were made by randomly shuffling the presynaptic spiketrains.⁷

Detection rates are broken down per type of both the presynaptic and postsynaptic neuron (excitatory or inhibitory). We find that inhibitory inputs are significantly easier to detect (independent of the type of the postsynaptic neuron). This is due to the synaptic weight of inhibitory inputs – and thus their PSPs – being $4\times$ stronger than those of excitatory inputs.

The most interesting category in figure 48 are the ‘Unconnected’ detection rates (i.e. the false positive rates, FPRs), which are almost all higher than the detection threshold α . In theory, the FPR should equal the detection threshold α ; and indeed this is what we roughly found in the N-to-1 experiments with STA-height shuffle test.⁸ In the network, we thus detect unconnected inputs – or to be more precise: not-directly-connected inputs – at a rate higher than chance. This deserves more scrutiny, which we do in the next section.

4.7 False positive detections

As a summary of the above: when we test unconnected spiketrains as inputs to some neuron in the network, we falsely classify them as connected at a higher rate than would be expected if these spiketrains

⁷ A spiketrain is shuffled by taking its inter-spike-intervals (ISIs), randomly shuffling those, and reconstructing a new spiketrain from the resulting shuffled ISIs.

⁸ This is true almost by definition of the shuffle test: an unconnected spiketrain in the N-to-1 experiment has randomly generated spiketimes. Shuffling this random spiketrain creates more random spiketrains. Looking at the distribution of STA heights of these random spiketrains, the chance that the ‘real’ STA (of the unconnected spiketrain) is in the top- α fraction of STA heights – and that thus, it would be wrongly classified as a true input, i.e. a false positive – is exactly α .

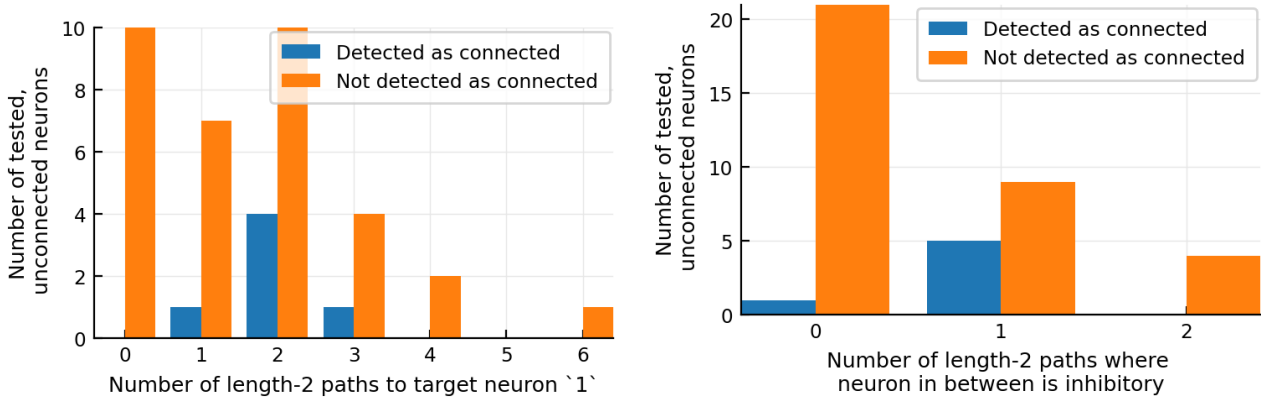


Figure 49: **Characterizing what is different about false positive inputs.**

All tested unconnected inputs to a selected target neuron in the network, with inputs correctly classified as not-connected in orange, and the false positives in blue. Source: [2022-07-14__Unconnected-but-detected](#).

were fully random. So, they have a stronger-than-random STA. Or in other words, their spikes seem to have some influence on the target neuron’s voltage, even though they are not direct inputs to it. In this section, we examine what is special about these false positive input neurons.


First, we found that all false positives (i.e. all not-directly-connected inputs that were nonetheless classified as connected) had a shortest path to the target neuron that consisted of only one in-between neuron. However, as we saw in [figure 45](#), this is not that special in this highly-interconnected random network.

But maybe there are *more* such length-2 paths between the input neuron and the target neuron? We indeed found that this was the case when examining one target neuron and all its tested unconnected inputs. In [figure 49](#) (left), the blue distribution (number of length-2 paths to the target neuron, for false positive inputs) is higher than the orange distribution (same, but for true negatives).

Next, we looked at the type of in-between neurons on the shortest path. We found that for the false positive inputs, the in-between neuron was more likely to be inhibitory ([figure 49](#), right). This makes sense, as inhibitory inputs are stronger (they are also detected at a higher rate, as seen in [figure 48](#)).

Finally, we looked at the firing rate of the in-between neurons ([figure 50](#)). For false positives inputs, the in-between neurons were more active than for true negatives.

In summary, we have some evidence (though it should be tested more thoroughly) that our higher-than-expected false positive rate is due to indirect inputs: neurons that synapse onto direct inputs. False positives synapse onto more such direct inputs, and those direct inputs have a stronger influence on the target neuron: a higher firing



rate, and more likely to be inhibitory.

Further work could adapt our network topology to be more sparse, so that indirect connections are much rarer, and their effect easier to study.

4.8 Conclusion

In this chapter, we extended the N-to-1 testing setup to a full network of simulated neurons. We found that the STA height method could still detect connections. However, unlike in the N-to-1 setup – with its fully independent inputs – in the network, the false positive rate exceeded the detection threshold α . We hypothesized that this is due to indirect connections being detected as direct inputs, and found some preliminary evidence that that is indeed the case.

One of the avenues pursued during the investigation of the indirect connections, was to plot their STA and compare it to the ‘ideal’ STA⁹. This ideal STA (figure 51) was made by using the ground truth connectivity and averaging all the STAs of the true direct connections. This clean STA gave us the idea to (1) model this shape, and use this model as some sort of ‘prior’ for a stronger connection test; and (2) use it as a template to correlate STAs with. This led us to the new connection tests discussed in the next chapter.

Figure 50: False positives synapse onto high-firing direct inputs.

See figure 49.




Figure 51: Average STA of all true excitatory connections to recorded neurons in the network. Source: [2022-09-09__Comntest_with_template_matching](#).

⁹The hypothesis was that we would see a voltage bump twice as late as for direct connections, i.e. indicating a disynaptic connection. This investigation was inconclusive, but it did give us ideas for new connection detection methods.

New connection inference methods

5.1 Introduction

This chapter introduces three new methods to detect synaptic connections from voltage signals.

The first two methods are still based on the STA (spike-triggered average), as in [chapter 3](#). And, also as before, they still use shuffled presynaptic spiketrains¹ to provide a null-distribution for the test statistic

The difference is in how the STA is used. In [chapter 3](#) we used the height of the STA ('peak-to-peak', ptp) as a test statistic.

The first new method here instead *correlates* the STA with some template of what the STA of a true connection would look like, to calculate a test statistic.

The second method fits an idealized function to the STA, and uses the goodness-of-fit as a test statistic.

The third method steps away entirely from the STA, and instead uses the data of the different spike-triggered windows directly, without averaging them together in an STA.

As the focus was on comparing these new methods with each other and with the first STA test, we tested them in the N-to-1 setup. A further investigating could also test them on the fully recurrent network, to compare their effectiveness against indirect connections.

¹ A note on terminology: when we write 'presynaptic' or 'postsynaptic' in this thesis, we often mean the neuron that we are testing as a *possible* presynaptic / postsynaptic connection.




Figure 52: **Correlating the STA with a template.** *Left:* An example STA, and the template it will be correlated with to calculate the connection test statistic.

Right: Two possible STA templates: in blue, the average of the STAs of many true excitatory connections (namely of 50 neurons – part of a 1000-neuron recurrent network – that all had their voltage recorded; and all their excitatory inputs). In orange, the average STA of the excitatory connections detected with a strict ‘peak-to-peak’ test.

5.2 STA Template correlation

The idea behind this connection test is as follows: the more our actual STA looks like some idealized, ‘clean’ STA, the more likely the connection exists. We quantify this ‘looking like’ here by a simple Pearson correlation.

The question then remains what to correlate our real STA with. An ideal template to correlate with would be the average STA of all true connections in the network.² The knowledge of these connections is of course not available a-priori (it is what we are trying to infer). But it turns out we can obtain a signal that is very close to this ideal template with a two-step approach.

In the first step, we use the previous test statistic (peak-to-peak height of the STA); but with a stricter p -value threshold (a higher α). This gives us a sample of true connections with few false positives (but with many missed connections). The STAs of these found connections are then averaged, to obtain an estimate for the STA template. We find that this template matches the ideal template closely (see [figure 52](#), right).

In the second step, we correlate this found template with the STAs of all possible connections (and with their shuffle-control counterparts). The correlation values are then used as test statistic, with the original $\alpha = 0.05$ threshold.

We find that this correlation-based test statistic outperforms the simple peak-to-peak measure ([figure 53](#)).

² Specifically, either all excitatory connections, or all inhibitory connections. When correlating some connection’s STA with an ‘excitatory’ (upwards) template, the sign of the correlation tells us whether the connection is excitatory (+) or inhibitory (–).




Figure 53: **Correlation outperforms STA height as connection test.** Compare also with [figure 56](#).

AdEx neuron, 10-minute recording, no voltage imaging noise, 5 different seeds (blue dots). All N inputs are tested (instead of just the highest firing ones), in addition to N random unconnected spiketrains. AUC chance level is at 0.252 (as in [figure 31](#)).

Source: [2024-05-26__Fix_template-based_method](#).

5.3 Fitting a full STA model

Instead of a data-driven approach for the ‘ideal’ STA shape, for this method we manually design a function (a 7-parameter function that we’ll call f), which we will fit to the actual STAs.

One part of this STA model function is shaped like a postsynaptic potential bump (PSP, figure 54). It is the impulse response of two linear integrators placed in series. Or, in other words, the convolution of two ‘step-and-decay’ functions.^{3, 4} It is the postsynaptic potential in the simplest linear neuron model ($\frac{dv}{dt} \cdot \tau_m = -v + I$) where the synaptic current I is also linear ($\frac{dI}{dt} \cdot \tau_s = -I + s$).⁵

Solving for v and with a single input spike at $t = 0$, we have, for $t \geq 0$:

$$\text{PSP}(t) = \begin{cases} t e^{-t/\tau_m} & \tau_m = \tau_s \\ \frac{\tau_m \tau_s}{\tau_m - \tau_s} (e^{-t/\tau_m} - e^{-t/\tau_s}) & \tau_m \neq \tau_s \end{cases} \quad (26)$$

We find that this function approximates the shape of the *simulated* postsynaptic potential in our actual neuron model well – even though our neuron model is not so linear.⁶

But even though equation (26) models our simulated PSPs well, it does not resemble our STAs⁷; see for example the average STAs in figure 52. First, the bump in the STA does not occur immediately after the presynaptic neuron’s spike. This is due to the simulated axonal transmission delay, and it is easily replicated in the model by shifting the PSP function in time. But more interestingly, the STA shows a sort of ‘dip’, where it dives below baseline just after the PSP bump, and flares up again at the ends (a downwards slant in the initial delay period, and an upwards slant at the end of the STA).⁸

The phenomena of the delay and the ‘dip’ thus differentiate an STA from a PSP. We find empirically that the dip can be parsimoniously modelled by subtracting a broad Gaussian curve from the PSP: see figure 55, right.⁹

Our model f is thus the difference of a delayed version of equation (26), and a Gaussian curve. The result is scaled by some factor α , to match the size of the specific STA that is being fit. α can be negative, to model inhibitory connections; f is then flipped upside-down, as in figure 55.

The model function thus has seven parameters: the PSP’s delay, τ_m , and τ_s ; the Gaussian’s location, width, and relative height wrt. the PSP; and the final scaling α . All parameters, except for the Gaussian’s, have a ready biophysical interpretation.¹⁰




Figure 54: A simple PSP model.

³ ‘Step-and-exponential-decay’: $u(t) \cdot e^{-t/\tau}$, with $u(t)$ the Heaviside step function $\mathbf{1}_{t \geq 0}$.

⁵ τ_s and τ_m are the synaptic and membrane time constants, and $s(t) = \sum_i \delta(t - t_i)$ is the train of input spikes i .

⁶ Our model for the membrane potential has a component (the adaptation current u) whose differential equation recurrently depends on v . Additionally, the synaptic currents I_j are *conductance-based*, meaning that they also recurrently depend on v : $I_j = g_j (v - E_j)$ (with E_j the reversal potential at synapse j , and the synaptic conductance g_j a linear integrator of the input spike train $s_j(t)$):

$$\frac{dg_j}{dt} \cdot \tau_{s,j} = -g_j + s_j.$$

⁷ This falsifies our initial implicit hypothesis that STAs are merely noisy reflections of PSPs.

⁸ Figure 36 from chapter 3 gives a clue to the reason for this dip: when there are no postsynaptic spikes (second and fourth column in figure 36), there is also no dip visible in the STAs (second and third rows). This dip thus seems related to these postsynaptic spikes, maybe enhancing the PSP right before and during the spike, and suppressing it right after (when the voltage is reset).

⁹ Note that in figure 55, we fit an *inhibitory* STA, and the model components are thus inverted: the delayed PSP bump (blue) is downwards, and the Gaussian ‘dip’ (orange) is upwards.

¹⁰ Namely: α is proportional to the connection strength, and indicates excitation (+) or inhibition (-). The PSP’s delay, τ_m , and τ_s are estimates of the axonal transmission delay and the synaptic and membrane time constants.

⁴ In computational neuroscience, this function is sometimes called the ‘alpha function’ or ‘alpha synapse’. *Synapse*, because the double integrator model can be used to model synaptic conductance (g), with then a fast rise τ_1 (modelling neurotransmitter release), and a slower decay τ_2 (modelling its dissipation). The synaptic conductance model in our simulation is simpler: the ‘rise’ is instant, i.e. we only have exponential decay.




Figure 55: **Fitting a parametric model to the STA.** *Left:* the model f (orange) fit to the STA of an example inhibitory connection from a network simulation (blue). Both signals are z-scored. The components of the orange model function are shown on the *Right:* the delayed simple PSP model (blue) and the Gaussian ‘dip’ (orange).

We fit the model function to individual STAs using nonlinear least-squares optimization; more specifically using the Levenberg-Marquardt algorithm.¹¹ To make the fitted functions well-behaved and looking like actual STAs, we had to limit the creative freedom of the optimization algorithm, by enforcing box constraints on the parameters.

The final test statistic is the mean squared error (MSE) between z-scored versions of the fitted function and the STA. The normalization by z-scoring is necessary to be able to compare goodness of fit between STAs with different ranges¹²

In comparison with the previous two methods (STA height and STA template correlation), this model-fitting approach did not perform significantly better than STA height, and it was clearly worse than template correlation.

The model-fitting is also very computationally expensive: iteratively fitting a nonlinear function to an STA takes much longer than simply calculating a Pearson correlation between two STA signals.

For these two reasons, we have not performed an extensive performance evaluation of this method, as we have done for the previous and following methods (figures 53 and 56 respectively).

One advantage of this model-fitting method over the other methods though, is that most of the fitted parameters have a ready biophysical explanation: connection strength, transmission delay, and synaptic and membrane time constants.

¹¹ This is the standard method for nonlinear least-squares optimization. For example, in MATLAB’s `lsqcurvefit` function, it is the default algorithm when the number of functions (in our case, 1) exceeds the number of datapoints (in our case, 1000: 100 ms of signal post-spike, at $\Delta t = 0.1$ ms).

We used the Levenberg-Marquardt implementation of the Julia package `LsqFit.jl`.

¹² Most importantly between a real STA and its shuffled versions. Because those shuffled STAs have a smaller range, they will have a lower MSE, even though the fit is visually clearly worse.

5.4 Linear regression of the upstroke

All previous methods are STA-based, i.e. different windows are cut from the postsynaptic voltage signal – one window for every presynaptic spike – and these are averaged into one signal, which is then used for further analysis. The method described in this section does not use STAs. We still construct the spike-triggered windows, but we do not average them. Instead, we use them to construct a gigantic data matrix for use in a linear regression.

Every single timepoint on the x-axis (in relative time after the presynaptic spike) will correspond to multiple voltage values on the y-axis; namely one for every window. Whereas for an STA, every timepoint on the x-axis corresponds to just a single y-value (the average voltage).

Why a linear regression? As we saw with the STAs in previous sections, they do not look like a line. But their first part, the "upstroke", might be approximated by one.

Let's number our spike-triggered windows 1, 2, ..., N (for a presynaptic spiketrain with N spikes), and let's number the voltage values in window i as $[y_{i,1}, y_{i,2}, \dots, y_{i,M}]$ (for a window length of M samples long).¹³ We will then perform the following linear regression:

$$\begin{aligned}
 \mathbf{y} &= \mathbf{X} \boldsymbol{\beta} + \boldsymbol{\varepsilon} & (27) \\
 \begin{bmatrix} y_{1,1} \\ y_{1,2} \\ \vdots \\ y_{1,M} \\ y_{2,1} \\ y_{2,2} \\ \vdots \\ \vdots \\ y_{N,M} \end{bmatrix} &= \begin{bmatrix} 1 & 1 \\ 1 & 2 \\ \vdots & \vdots \\ 1 & M \\ 1 & 1 \\ 1 & 2 \\ \vdots & \vdots \\ \vdots & \vdots \\ 1 & M \end{bmatrix} \begin{bmatrix} \beta_0 \\ \beta_1 \end{bmatrix} + \begin{bmatrix} \varepsilon_{1,1} \\ \varepsilon_{1,2} \\ \vdots \\ \varepsilon_{1,M} \\ \varepsilon_{2,1} \\ \varepsilon_{2,2} \\ \vdots \\ \vdots \\ \varepsilon_{N,M} \end{bmatrix} & (28)
 \end{aligned}$$

¹³ This window length M is an important parameter. The window must be approximately as long as the upstroke of the STA/PSP: not longer (so that the shape is too complex to be fit by a line), nor shorter (so that there would not be enough data for a proper fit).

I.e. we will regress voltage-after-spike against time-after-spike. Note that the noise term here is not just modelling the Gaussian noise we explicitly add ourselves (i.e. the voltage imaging noise); but rather *everything* that is not the increase in the postsynaptic voltage due to a presynaptic spike. This is mainly the effects of other presynaptic neurons, and postsynaptic spikes.

If we assume additive Gaussian noise, i.e.

$$\varepsilon_{i,j} \sim \mathcal{N}(0, \sigma^2) \tag{29}$$

then the maximum-likelihood estimate $\hat{\boldsymbol{\beta}}$ of the regression parameters $\boldsymbol{\beta}$ is obtained through minimizing the mean squared error (MSE)¹⁴ between our observed voltage values \mathbf{y} and the fitted line

¹⁴ <https://statproofbook.github.io/P/slr-mle>

$$\hat{\mathbf{y}} = X\hat{\boldsymbol{\beta}}: \quad \hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta}} \|\mathbf{y} - X\boldsymbol{\beta}\|_2^2. \quad (30)$$

Our linear regression problem has become standard ordinary least squares (OLS) regression, for which a closed-form solution exists: the so-called normal equations,

$$\hat{\boldsymbol{\beta}} = (X^T X)^{-1} X^T \mathbf{y} \quad (31)$$

As is standard practice and for numerical stability ("don't explicitly invert a matrix if not needed"), we instead find the optimal intercept ($\hat{\beta}_0$) and slope ($\hat{\beta}_1$) of our linear fit using QR-factorization, via Julia's left-division operator: $\hat{\boldsymbol{\beta}} = X \setminus \mathbf{y}$.

From this linear fit we thus obtain a slope $\hat{\beta}_1$. To use our regression as a connection test, we will perform a hypothesis test on this slope. The null hypothesis is that this slope is zero ($H_0 : \hat{\beta} = 0$): the voltage of the tested neuron does not react to spikes of the other tested neuron, and on average it stays flat.

If the slope β_1 really were 0, then we would expect the fitted slope $\hat{\beta}_1$ to be distributed normally around 0, as follows:¹⁵

$$\hat{\beta}_1 \sim \mathcal{N}(0, \sigma^2 [Q]_{2,2}) \quad (32)$$

where σ^2 is the variance of the sampling noise (equation (29)), and $[Q]_{2,2}$ is the second diagonal element of the 'cofactor matrix' Q ,¹⁶ which is the inverse of the Gram matrix $X^T X$:

$$Q = (X^T X)^{-1} \quad (33)$$

We do not know what the sampling noise σ^2 is, but we have a maximum-likelihood estimate $\hat{\sigma}^2$ for it,¹⁷ namely the mean squared error (MSE) of the fit:

$$\hat{\sigma}^2 = \frac{1}{n} \|\mathbf{y} - X\hat{\boldsymbol{\beta}}\|_2^2 \quad (34)$$

(where $n = M \cdot N$ is the number of datapoints).

Our final test statistic t to determine whether there is a connection from neuron A to neuron B is then the slope of the fit (of B's voltage after neuron A's spikes), normalized by how noisy the parameter fit is:

$$t = \hat{\beta}_1 / \hat{\sigma}_{\hat{\beta}_1} \quad (35)$$

where $\hat{\sigma}_{\hat{\beta}_1}$ is the standard error of the slope (equation (32) with equation (34)):

$$\hat{\sigma}_{\hat{\beta}_1} = \sqrt{\hat{\sigma}^2 [Q]_{2,2}} \quad (36)$$

Unlike in equation (32), t is no longer strictly normally distributed.¹⁸ But when the number of datapoints is quite large, as is the case here, the t -distribution becomes practically normal, and our test statistic t

¹⁵ <https://gregorygundersen.com/blog/2021/09/09/ols-hypothesis-testing/>

¹⁶ The indices are off by one ($\hat{\beta}_1$ vs $[Q]_{2,2}$), because the intercept in linear regression is conventionally called β_0 , and so we start numbering the elements of $\boldsymbol{\beta}$ from 0; but we number the elements of other vectors/-matrices from 1.

¹⁷ <https://statproofbook.github.io/P/slr-mle>

¹⁸ Instead, it follows a Student's t -distribution, with $n - p$ degrees of freedom, where $n = M \cdot N$ is the number of datapoints, and $p = 2$ is the number of parameters of the regression (β_0 and β_1).

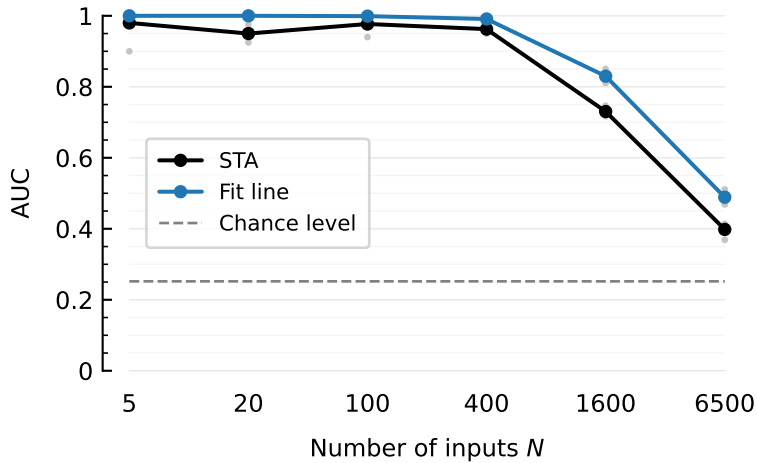


Figure 56: **Line fit outperforms STA height as connection test.**

AdEx neuron, 10-minute recording, no voltage imaging noise, 5 different seeds (blue dots). All N inputs are tested (instead of just the highest firing ones), in addition to N random unconnected spiketrains. AUC chance level is at 0.252 (as in figure 31).

Source: [2024-05-26__Fix_template-based_method.](#)

will very closely follow the standard normal distribution under the null hypothesis ("the slope is zero").

We can (and do) use this t -value as-is in our connection test, i.e. as the test statistic over which the detection threshold is swept.

Additionally, if our assumptions (linear data-generating process, additive Gaussian noise) would be correct,¹⁹ we could also assign an actual probability (a p -value) to observing slopes as extreme as observed. Because t follows the standard normal distribution under the null-hypothesis, and with $\phi(x)$ the cumulative probability function of this distribution, we have: $p = 2 \phi(-|t|)$: the probability that a sample smaller than $-|t|$ or larger than $|t|$ is drawn.

¹⁹ They are not. As one example, consider postsynaptic spikes as one of the sources of noise on the PSP: they are asymmetrical (spikes are always upwards, never downwards). Figure 57 shows that the residuals of a linear regression of an example connection are not symmetric.

As seen in figure 57, our assumption of additive Gaussian noise ε on the regression does not seem valid. But our goal for this regression was not to build a statistical model of the STA, but rather to use the fit as a tool for connection detection, something at which it empirically succeeds, notwithstanding the violated statistical assumptions.

Figure 56 shows the performance of the described linear regression method, for different number of inputs N , and compared to the STA height method. We find a perfect performance (AUC of 1) for up to 400 inputs. After that, the AUC starts to drop, but it stays much higher than the STA height method. At the realistic number of inputs $N = 6500$, the linear fit's AUC is still at 0.5, while the STA height method performs at chance level (≈ 0.25)




Figure 57: **Residuals of an example regression.**

Linear regression of pooled windows of the voltage of an AdEx neuron with 6500 inputs simulated for 10 minutes, for its highest-firing excitatory input; each presynaptic spike corresponds to a 10 ms window. Left: Residuals over time. Right: Residuals marginalized over time; vertical black line indicates both the mean (-5.9×10^{-12} mV) and the median (0.1 mV).

The residuals seem to be homoskedastic, but are not symmetrically distributed.

Source: [2024-05-29_Residuals-linefit](#).

Relationship to fitting the STA

For this last method, we emphasized that it does not operate on the STA, but instead uses the individual windows directly. But an analysis of the normal equations (equation (31)) reveals that fitting a line to all the individual windows pooled together is actually very similar to fitting a line to their average, the STA.

The linear regression on the STA has the following response vector $\bar{\mathbf{y}}$ and design matrix $\bar{\mathbf{X}}$:

$$\bar{\mathbf{y}} = \bar{\mathbf{X}} \bar{\boldsymbol{\beta}} + \bar{\boldsymbol{\varepsilon}} \quad (37)$$

$$\frac{1}{N} \begin{bmatrix} y_{1,1} + y_{2,1} + \dots + y_{N,1} \\ y_{1,2} + y_{2,2} + \dots + y_{N,2} \\ \vdots \\ y_{1,M} + y_{2,M} + \dots + y_{N,M} \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & 2 \\ \vdots & \vdots \\ 1 & M \end{bmatrix} \begin{bmatrix} \beta_0 \\ \beta_1 \end{bmatrix} + \begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \vdots \\ \varepsilon_M \end{bmatrix} \quad (38)$$

We see that the previous design matrix X consists of N stacked repetitions of the design matrix $\bar{\mathbf{X}}$ here.

We can show that both regressions have the same solution for their normal equations, i.e. that

$$\begin{aligned} \hat{\boldsymbol{\beta}} &= (X^T X)^{-1} X^T \mathbf{y} \\ &= \hat{\bar{\boldsymbol{\beta}}} = (\bar{\mathbf{X}}^T \bar{\mathbf{X}})^{-1} \bar{\mathbf{X}}^T \bar{\mathbf{y}} \end{aligned} \quad (39)$$

As an example, working out the second element of $X^T \mathbf{y}$, we find:

$$\begin{aligned}
 [X^T \mathbf{y}]_2 &= [1 \ 2 \ \dots \ M] \cdot \mathbf{y}_1 + [1 \ 2 \ \dots \ M] \cdot \mathbf{y}_2 + \dots + [1 \ 2 \ \dots \ M] \cdot \mathbf{y}_N \\
 &= [1 \ 2 \ \dots \ M] \cdot (\mathbf{y}_1 + \mathbf{y}_2 + \dots + \mathbf{y}_N) \\
 &= [1 \ 2 \ \dots \ M] \cdot N\bar{\mathbf{y}} \\
 &= N \cdot [\bar{X}^T \bar{\mathbf{y}}]_2
 \end{aligned} \tag{40}$$

(where \cdot is the inner or dot product, and $\mathbf{y}_i = [y_{i,1} \ y_{i,2} \ \dots \ y_{i,M}]$).

The same is true for the first element of $X^T \mathbf{y}$ (just with $[1 \ 1 \ \dots \ 1]$ instead of $[1 \ 2 \ \dots \ M]$), so that

$$X^T \mathbf{y} = N \cdot \bar{X}^T \bar{\mathbf{y}} \tag{41}$$

A similar argument can be made for each of the four elements of the Gram matrix, and we find that $X^T X = N \cdot \bar{X}^T \bar{X}$.

That makes the cofactor matrix:

$$(X^T X)^{-1} = \frac{1}{N} (\bar{X}^T \bar{X})^{-1} \tag{42}$$

Taking together equations (41) and (42), we indeed find the exact same solutions $\hat{\beta}$ for the line fit (equation (39)), whether regressing all individual windows, or their average.

Even though the fitted slope $\hat{\beta}_1$ is the same, there is a difference in its standard error (equation (36)), and thus also in the test statistic t we'd use as a connection test (equation (35)).

We can write the standard error on the slope in the first regression (on all the individual windows) as follows:²⁰

²⁰ where c is shorthand for $\sqrt{[Q]_{2,2}/M}$.

$$\hat{\sigma}_{\hat{\beta}_1} = \frac{1}{N} c \|\mathbf{y} - X\hat{\beta}\|_2 \tag{43}$$

$$= \frac{1}{\sqrt{N}} c \sqrt{\sum_{j=1}^M \frac{1}{N} \sum_{i=1}^N (y_{i,j} - (\hat{\beta}_0 + \hat{\beta}_1 \cdot j))^2} \tag{44}$$

For the regression of the STA, we have:

$$\hat{\sigma}_{\hat{\beta}_1} = c \|\bar{\mathbf{y}} - \bar{X}\hat{\beta}\|_2 \tag{45}$$

$$= c \sqrt{\sum_{j=1}^M \left(\left(\frac{1}{N} \sum_{i=1}^N y_{i,j} \right) - (\hat{\beta}_0 + \hat{\beta}_1 \cdot j) \right)^2} \tag{46}$$

In other words, the standard error of the STA regression uses the squared error between the fitted line and the average measured voltage, whereas for the regression on all windows, we use the average of the squared errors between the fitted line and all of the individual measured voltages in one timepoint.

These last equations show the advantage of regression with individual windows, versus regressing the STA: the more windows are used

(higher N) – i.e. the more presynaptic spikes this possible connection has – the lower the standard error on the fitted slope will be, and the higher the test statistic $t = \hat{\beta}_1 / \hat{\sigma}_{\hat{\beta}_1}$. This is thanks to the scaling by $1/\sqrt{N}$ in equation (44), which is absent in equation (46).

This is a quantification of a heuristic that could go something like: “the more ‘presynaptic’ spikes there are, the more certain we can say there is or isn’t a connection here”.²¹ And when using just the STA, the knowledge of how many windows were used to calculate it is lost.

Two notes here: first, as the fitted parameters β are the same when regressing either all windows or the STA, we could just use the STA for fitting, and then manually add a correction factor $1/\sqrt{N}$ to the test statistic t . This is interesting computationally, as you then don’t have to construct the very long design matrix X and response vector y .

Second, one could argue the standard error of the STA regression *does* already account for the number of windows used: the more windows are averaged together, the less the STA will be noisy and the more it will look like a straight line (at its upstroke portion at least); and the lower the standard error will be. Under this view, an additional correction factor in the test statistic to explicitly account for the number of windows / presynaptic spikes, is not necessary.

5.5 Conclusion

In this chapter we wondered whether we could improve upon the most obvious voltage-based connection detection method (namely measuring the height of the STA). We developed and tested three new methods. We found that two of them – template correlation and linear regression of the upstroke – indeed performed significantly better as a connection test. The third method we developed (fitting a parametric model to the STA shape) did not perform better as a connection test, but it does provide a method to estimate biophysical parameters of a connection.

²¹ E.g. if two possible connections would have the same fitted slope for their spike-triggered voltages, but one of them has more presynaptic spikes, then that connection would have a higher test statistic, and would be favoured when classifying the connections.

Discussion

We have developed and tested three new methods for voltage-based connection inference. In the process, we set-up a simulation and performance evaluation pipeline, both in a simplified ‘N-to-1’ setup, and a full network simulation.

We find that two of our new methods (template-based correlation, and linear regression of the upstroke), outperform the simplest voltage-based test, namely measuring the height of the spike-triggered average.

6.1 Future work

Possible extensions to this thesis fall in three categories: how the network inference methods are tested and evaluated; comparing them with existing methods, and improving on our methods.

Improving the test setup

Most importantly, our methods should be tested on real data, and not only simulated data. The problem is knowing the ground-truth connectivity: the connectomes of e.g. mice are unknown and vary from animal-to-animal (and over time). A starting point is to first test on simpler organisms, where the connectivity *is* known and constant. An example is the stomatogastric ganglion in decapods, whose connectivity and activity patterns are well characterized. (Though the neural activity is very oscillatory, which might (1) not be representative of circuits in general, and (2) might confuse network inference algorithms).

But before moving to real data, there is a lot more realism that could be added to our simulation, both on the neural side as the voltage imaging side.

For voltage imaging, as mentioned before we could simulate photobleaching over time by exponentially decreasing the SNR, and we could simulate the varying time-resolutions of different voltage sensors by convolving the voltage signal with e.g. a subtle low pass

filter. Another difficulty that fluorescence imaging poses is light scattering. Photons emitted by one neuron do not all go straight to the microscope sensor, but can bounce off other neurons and intermediate tissue, thus appearing to come from the wrong source. In the Neural Connectomics Challenge (which we introduce in the next section), this was simulated by spatially embedding the simulated neurons, and mixing the signals emitted by each neuron by a fraction of the signals emitted by neighbouring neurons.

On the neural side, many straightforward additions are possible. We have simulated only one neuron type, with fixed parameters (namely the cortical regular spiking neuron). Instead, we could (1) pick neuron parameters so the dynamical system bifurcates to another neuron type (bursting, accelerating, transient spiking, . . .) [Nau+08], and/or (2) draw neuron model parameters from distributions.

An important aspect that we have kept fixed up to now is the input firing rate distribution. It would be informative to quantify the effect of different distribution types and parameter values on connection inference performance, and especially on bias towards higher-firing inputs.

Some neural phenomena that we have not yet modelled but that are likely to also have an effect on connection inference are (1) short-term synaptic plasticity, (2) global bursting (a bit more on that in the next section), and (3) oscillations.

On the network side, we only used the simple uniform random connectivity. But the effect on inference of other (and more realistic) topologies such as scale-free networks should be explored too.

As far as the evaluation of methods goes, a path we did not pursue is to predict and evaluate synaptic *strengths*, and not just the existence (yes/no) and polarity (excitatory/inhibitory) of a synaptic connection. This is not often done in network inference studies, but is not unheard of [Zha+17]. One way could be to quantify how similar the real and the predicted connectivity strength matrices are.

Existing network inference methods

There is a vast literature on spike-to-spike (i.e. purely event-based) network inference methods. A lot of these methods are based on information-theoretic measures, such as mutual information and ‘transfer entropy’ between two spiketrains. In practice, this boils down to counting bit-patterns: spiketrains are binned in time, as e.g. $[0, 0, 0, 1, 0, 1, 0, 0, \dots]$. Two such spiketrains are then overlaid (possibly shifted in time relative to one another), and the patterns that the bits of both form around each spike are gathered and tallied up. (For example: $([0, 1, 0], [0, 0, 1])$, for a window length of one bin before and one bin after the spike). If one of these patterns is more common than random (i.e. the distribution of patterns has low entropy), then one spiketrain might be connected to another.

These methods generally don’t work very well [DF20], and are very

data-intensive (the number of possible bit-patterns increases exponentially the longer your windows are). In 2014, a competition was held on spike-based network inference: the Neural Connectomics Challenge or NCC [Guy+14; Orl+17]¹. The contestants using information-theoretic measures like transfer entropy did not rank high. The winning team used ‘partial correlation’ (PC) between binned spiketrains [Sut+17]. PC quantified the association between two spiketrains after the influence of all other spiketrains was removed. The PCs could be efficiently calculated by inverting the correlation matrix between the spiketrains (to obtain the ‘precision’ matrix). Trying this method on our data is a straightforward next step.

¹ <https://connectomics.chalearn.org>.

What all the winning teams had in common was that they reduced the influence of network-wide activity bursts: in the given simulated data, the network would often explode with activity, where more or less all neurons fire simultaneously. These moments give little information on connectivity, and filtering them out improved their algorithms’ performance. Our simulations did not contain global bursts, but their influence on our voltage-based methods should certainly be tested.

Because our voltage-based methods have more information and operate on the principle of the directly-spike-triggered PSP², we assumed we would be less plagued by the problem of indirect connections, from which spikes-only methods often suffer. But as we saw in chapter 4, we have an excess of false positives, likely due to such indirect connections. A direct comparison with spike-based methods would thus be fruitful to gauge how much it helps to have access to voltage data.

² Post-synaptic potential

There is one published paper that we know off that also explicitly researches voltage-based network inference: Zhang et al. 2017, [Zha+17]. Their method is called “spike-triggered regression”, which at first sight sounds similar to the linear regression of the upstroke in spike-triggered windows that we do. But whereas we regress voltage against time-after-spike, they regress voltage against its own history (i.e. an autoregressive voltage model), and binned spiketrains of other neurons. The regression coefficients of these other spiketrains on the ‘post’ neuron’s voltage are used for the connection test. Implementing Zhang et al.’s method and testing it on our data would be instructive.

New network inference methods

A deep-learning based network inference approach might be feasible, because we have sheer-infinite training data: we can generate and simulate as many different neuronal networks as needed. This would fall under the ‘simulation-based inference’ (SBI) nomenclature [CBL20]. While SBI techniques have been applied in neuroscience [Gon+20], they have to our knowledge not yet been used to reconstruct connectivity.

6.2 Conclusion

In this thesis, we set out to answer the question: "Is it feasible, in principle, to perform network inference from neural voltage signals?". We have tested this idea with simulated data, and developed three new voltage-based connection detection methods. We also extensively tested a simple existing algorithm – namely measuring the height of the spike-triggered average (STA), which is a reflection of the postsynaptic potential – under varying excitatory-inhibitory (EI) conditions, and both in a simplified 'N-to-1' setup, and in a fully recurrent network.

We find that voltage-based network inference seems feasible to a limited extent, under realistic voltage imaging conditions (6500 EI-balanced inputs to one neuron, 10-minute recording, voltage imaging SNR of 40): the simple algorithm performs considerably better than chance.

Two of our newly developed inference methods perform better than the simple algorithm. Our best-performing algorithm correlates STAs with a 'template', which is obtained through a first pass of the simple algorithm. This detection method reaches an AUC value of 0.53 for 6500 inputs (chance level AUC: 0.25). At a false positive rate of 5%, we detect 33% of the neuron's inhibitory inputs, and 13% of its excitatory inputs.

In absolute terms, this performance is not stellar. But we must note the high number of inputs that we used here (up to 6500). Studies of spikes-only connection-detection methods typically only test networks with about 100 inputs per neuron [Ito+11; Guy+14]. Our methods have near perfect detection performance for up to 400 inputs per neuron.

We conclude that voltage-based network inference seems useful for (1) inferring simple networks (low number of inputs per neuron), and (2) finding the high-firing connections in more complex networks.

References

- [Abd+19] Ahmed S. Abdelfattah et al. “Bright and Photostable Chemigenetic Indicators for Extended in Vivo Voltage Imaging”. In: *Science* (Aug. 16, 2019). DOI: [10.1126/science.aav6416](https://doi.org/10.1126/science.aav6416). pmid: [31371562](https://pubmed.ncbi.nlm.nih.gov/31371562/).
- [Ada+19] Yoav Adam et al. “Voltage Imaging and Optogenetics Reveal Behaviour-Dependent Changes in Hippocampal Dynamics”. In: *Nature* (May 2019). DOI: [10.1038/s41586-019-1166-7](https://doi.org/10.1038/s41586-019-1166-7).
- [Bad+08] Laurent Badel, Sandrine Lefort, Thomas K. Berger, Carl C. H. Petersen, Wulfram Gerstner, and Magnus J. E. Richardson. “Extracting Non-Linear Integrate-and-Fire Models from Experimental Data Using Dynamic I–V Curves”. In: *Biological Cybernetics* (Nov. 15, 2008). DOI: [10.1007/s00422-008-0259-4](https://doi.org/10.1007/s00422-008-0259-4).
- [Bez+17] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah. “Julia: A Fresh Approach to Numerical Computing”. In: *SIAM Review* (Jan. 2017). DOI: [10.1137/141000671](https://doi.org/10.1137/141000671).
- [BG05] Romain Brette and Wulfram Gerstner. “Adaptive Exponential Integrate-and-Fire Model as an Effective Description of Neuronal Activity”. In: *Journal of Neurophysiology* (Nov. 2005). DOI: [10.1152/jn.00686.2005](https://doi.org/10.1152/jn.00686.2005).
- [Bre+07] Romain Brette et al. “Simulation of Networks of Spiking Neurons: A Review of Tools and Strategies”. In: *Journal of Computational Neuroscience* (Dec. 1, 2007). DOI: [10.1007/s10827-007-0038-6](https://doi.org/10.1007/s10827-007-0038-6).
- [CBL20] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. “The Frontier of Simulation-Based Inference”. In: *Proceedings of the National Academy of Sciences* (Dec. 2020). DOI: [10.1073/pnas.1912789117](https://doi.org/10.1073/pnas.1912789117).
- [CMT18] Jose Casadiego, Dimitra Maoutsa, and Marc Timme. “Inferring Network Connectivity from Event Timing Patterns”. In: *Physical Review Letters* (2018). DOI: [10.1103/PhysRevLett.121.054101](https://doi.org/10.1103/PhysRevLett.121.054101).

- [Cox19] Guy Cox. *Fundamentals of Fluorescence Imaging*. CRC Press, Apr. 23, 2019. 476 pp. ISBN: 978-1-351-12939-8. Google Books: [Y10WDwAAQBAJ](#).
- [CS15] Matt Carter and Jennifer C. Shieh. *Guide to Research Techniques in Neuroscience*. 2nd edition. Academic Press, 2015.
- [DA01] Peter Dayan and L. F. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. Computational Neuroscience. The MIT Press, 2001. 460 pp. ISBN: 978-0-262-04199-7.
- [DF20] Abhranil Das and Ila R. Fiete. "Systematic Errors in Connectivity Inferred from Activity in Strongly Recurrent Networks". In: *Nature Neuroscience* (Oct. 2020). DOI: [10.1038/s41593-020-0699-2](#).
- [Gon+20] Pedro J Gonçalves, Jan-Matthis Lueckmann, Michael Deistler, Marcel Nonnenmacher, Kaan Öcal, Giacomo Bassetto, Chaitanya Chintaluri, William F Podlaski, Sara A Haddad, Tim P Vogels, David S Greenberg, and Jakob H Macke. "Training Deep Neural Density Estimators to Identify Mechanistic Models of Neural Dynamics". In: *eLife* (Sept. 17, 2020). Ed. by John R Huguenard, Timothy O'Leary, and Mark S Goldman. DOI: [10.7554/eLife.56261](#).
- [Guy+14] I. Guyon, D. Battaglia, A. Guyon, V. Lemaire, J. G. Orlandi, B. Ray, M. Saeed, J. Soriano, A. Statnikov, and O. Stetter. "Design of the First Neuronal Connectomics Challenge: From Imaging to Connectivity". In: *2014 International Joint Conference on Neural Networks (IJCNN)*. 2014 International Joint Conference on Neural Networks (IJCNN). July 2014. DOI: [10.1109/IJCNN.2014.6889913](#).
- [HDZ08] Tomáš Hromádka, Michael R. DeWeese, and Anthony M. Zador. "Sparse Representation of Sounds in the Unanesthetized Auditory Cortex". In: *PLOS Biology* (Jan. 29, 2008). DOI: [10.1371/journal.pbio.0060016](#).
- [Hoc+14] Daniel R. Hochbaum et al. "All-Optical Electrophysiology in Mammalian Neurons Using Engineered Microbial Rhodopsins". In: *Nature Methods* (Aug. 2014). DOI: [10.1038/nmeth.3000](#).
- [Ito+11] Shinya Ito, Michael E. Hansen, Randy Heiland, Andrew Lumsdaine, Alan M. Litke, and John M. Beggs. "Extending Transfer Entropy Improves Identification of Effective Connectivity in a Spiking Cortical Network Model". In: *PLOS ONE* (Nov. 15, 2011). DOI: [10.1371/journal.pone.0027431](#).
- [Izh07] Eugene M. Izhikevich. *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting*. Cambridge, Mass.: The MIT Press, 2007. 464 pp. ISBN: 978-0-262-51420-0.

- [Kan+13] E.R. Kandel, T.M. Jessell, J.H. Schwartz, S.A. Siegelbaum, and A.J. Hudspeth. *Principles of Neural Science*. Fifth edition. McGraw-Hill Education, 2013. ISBN: 978-0-07-139011-8.
- [KS19] Thomas Knöpfel and Chenchen Song. “Optical Voltage Imaging in Neurons: Moving from Technology Development to Practical Tool”. In: *Nature Reviews Neuroscience* (Dec. 2019). DOI: [10.1038/s41583-019-0231-4](https://doi.org/10.1038/s41583-019-0231-4).
- [LPS15] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. “Numba: A LLVM-based Python JIT Compiler”. In: *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC. LLVM ‘15*. New York, NY, USA: Association for Computing Machinery, Nov. 15, 2015. ISBN: 978-1-4503-4005-2. DOI: [10.1145/2833157.2833162](https://doi.org/10.1145/2833157.2833162).
- [MB13] Kenji Mizuseki and György Buzsáki. “Preconfigured, Skewed Distribution of Firing Rates in the Hippocampus and Entorhinal Cortex”. In: *Cell Reports* (Sept. 12, 2013). DOI: [10.1016/j.celrep.2013.07.039](https://doi.org/10.1016/j.celrep.2013.07.039). pmid: [23994479](https://pubmed.ncbi.nlm.nih.gov/23994479/).
- [Mil16] Evan W Miller. “Small Molecule Fluorescent Voltage Indicators for Studying Membrane Potential”. In: *Current Opinion in Chemical Biology. Chemical Genetics and Epigenetics * Molecular Imaging* (Aug. 1, 2016). DOI: [10.1016/j.cbpa.2016.06.003](https://doi.org/10.1016/j.cbpa.2016.06.003).
- [MYD18] Ildefons Magrans de Abril, Junichiro Yoshimoto, and Kenji Doya. “Connectivity Inference from Neural Recording Data: Challenges, Mathematical Bases and Research Directions”. In: *Neural Networks* (June 1, 2018). DOI: [10.1016/j.neunet.2018.02.016](https://doi.org/10.1016/j.neunet.2018.02.016).
- [Nau+08] Richard Naud, Nicolas Marcille, Claudia Clopath, and Wulfram Gerstner. “Firing Patterns in the Adaptive Exponential Integrate-and-Fire Model”. In: *Biological Cybernetics* (Nov. 15, 2008). DOI: [10.1007/s00422-008-0264-7](https://doi.org/10.1007/s00422-008-0264-7).
- [OCo+10] Daniel H. O’Connor, Simon P. Peron, Daniel Huber, and Karel Svoboda. “Neural Activity in Barrel Cortex Underlying Vibrissa-Based Object Localization in Mice”. In: *Neuron* (Sept. 23, 2010). DOI: [10.1016/j.neuron.2010.08.026](https://doi.org/10.1016/j.neuron.2010.08.026). pmid: [20869600](https://pubmed.ncbi.nlm.nih.gov/20869600/).
- [Orl+17] Javier Orlandi, Bisakha Ray, Demian Battaglia, Isabelle Guyon, Vincent Lemaire, Mehreen Saeed, Alexander Statnikov, Olav Stetter, and Jordi Soriano. “First Connectomics Challenge: From Imaging to Connectivity”. In: *Neural Connectomics Challenge*. Ed. by Demian Battaglia, Isabelle Guyon, Vincent Lemaire, Javier Orlandi, Bisakha Ray, and Jordi Soriano. Cham: Springer International Publishing, 2017. ISBN: 978-3-319-53069-7 978-3-319-53070-3. DOI: [10.1007/978-3-319-53070-3_1](https://doi.org/10.1007/978-3-319-53070-3_1).

- [Pac+17] Marius Pachitariu, Carsen Stringer, Mario Dipoppa, Sylvia Schröder, L. Federico Rossi, Henry Dalgleish, Matteo Carandini, and Kenneth D. Harris. “Suite2p: Beyond 10,000 Neurons with Standard Two-Photon Microscopy”. In: *bioRxiv* (July 20, 2017). DOI: [10.1101/061507](https://doi.org/10.1101/061507).
- [Pia+19] Kiryl D. Piatkevich, Seth Bensussen, Hua-an Tseng, Sanaya N. Shroff, Violeta Gisselle Lopez-Huerta, Demian Park, Erica E. Jung, Or A. Shemesh, Christoph Straub, Howard J. Gritton, Michael F. Romano, Emma Costa, Bernardo L. Sabatini, Zhanyan Fu, Edward S. Boyden, and Xue Han. “Population Imaging of Neural Activity in Awake Behaving Mice”. In: *Nature* (Oct. 2019). DOI: [10.1038/s41586-019-1641-1](https://doi.org/10.1038/s41586-019-1641-1).
- [Rox+11] Alex Roxin, Nicolas Brunel, David Hansel, Gianluigi Mongillo, and Carl van Vreeswijk. “On the Distribution of Firing Rates in Networks of Cortical Neurons”. In: *Journal of Neuroscience* (Nov. 9, 2011). DOI: [10.1523/JNEUROSCI.1677-11.2011](https://doi.org/10.1523/JNEUROSCI.1677-11.2011). pmid: 22072673.
- [SBG19] Marcel Stimberg, Romain Brette, and Dan FM Goodman. “Brian 2, an Intuitive and Efficient Neural Simulator”. In: *eLife* (Aug. 20, 2019). Ed. by Frances K Skinner, Ronald L Calabrese, Frances K Skinner, Fleur Zeldenrust, and Richard C Gerkin. DOI: [10.7554/eLife.47314](https://doi.org/10.7554/eLife.47314).
- [SC21] Sadra Sadeh and Claudia Clopath. “Excitatory-Inhibitory Balance Modulates the Formation and Dynamics of Neuronal Assemblies in Cortical Networks”. In: *Science Advances* (Nov. 3, 2021). DOI: [10.1126/sciadv.abg8411](https://doi.org/10.1126/sciadv.abg8411).
- [SDC73] B. M. Salzberg, H. V. Davila, and L. B. Cohen. “Optical Recording of Impulses in Individual Neurones of an Invertebrate Central Nervous System”. In: *Nature* (Dec. 1973). DOI: [10.1038/246508a0](https://doi.org/10.1038/246508a0).
- [SH09] Massimo Scanziani and Michael Häusser. “Electrophysiology in the Age of Light”. In: *Nature* (Oct. 2009). DOI: [10.1038/nature08540](https://doi.org/10.1038/nature08540).
- [Sha+07] M. Shafi, Y. Zhou, J. Quintana, C. Chow, J. Fuster, and M. Bodner. “Variability in Neuronal Activity in Primate Cortex during Working Memory Tasks”. In: *Neuroscience* (May 25, 2007). DOI: [10.1016/j.neuroscience.2006.12.072](https://doi.org/10.1016/j.neuroscience.2006.12.072). pmid: 17418956.
- [SI97] Micah S Siegel and Ehud Y Isacoff. “A Genetically Encoded Optical Probe of Membrane Voltage”. In: *Neuron* (Oct. 1, 1997). DOI: [10.1016/S0896-6273\(00\)80955-1](https://doi.org/10.1016/S0896-6273(00)80955-1).
- [SK93] W. R. Softky and C. Koch. “The Highly Irregular Firing of Cortical Cells Is Inconsistent with Temporal Integration of Random EPSPs”. In: *Journal of Neuroscience* (Jan. 1, 1993). DOI: [10.1523/JNEUROSCI.13-01-00334.1993](https://doi.org/10.1523/JNEUROSCI.13-01-00334.1993). pmid: 8423479.
- [Str94] Steven H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Perseus Books, 1994.

- [Sut+17] Antonio Sutera, Arnaud Joly, Vincent Francois-Lavet, Zixiao Aaron Qiu, Gilles Louppe, Damien Ernst, and Pierre Geurts. “Simple Connectome Inference from Partial Correlation Statistics in Calcium Imaging”. In: *Neural Connectomics Challenge*. Ed. by Demian Battaglia, Isabelle Guyon, Vincent Lemaire, Javier Orlandi, Bisakha Ray, and Jordi Soriano. Cham: Springer International Publishing, 2017. ISBN: 978-3-319-53069-7 978-3-319-53070-3. DOI: [10.1007/978-3-319-53070-3_2](https://doi.org/10.1007/978-3-319-53070-3_2).
- [TW18] Yusuke Tomina and Daniel Wagenaar. “Dual-Sided Voltage-sensitive Dye Imaging of Leech Ganglia”. In: *BIO-PROTOCOL* (2018). DOI: [10.21769/BioProtoc.2751](https://doi.org/10.21769/BioProtoc.2751).
- [VB12] Bernard Valeur and Mário Nuno Berberan-Santos. *Molecular Fluorescence: Principles and Applications*. John Wiley & Sons, May 29, 2012. 593 pp. ISBN: 978-3-527-32837-6. Google Books: [NdySsFt8TN8C](https://books.google.com/books?id=NdySsFt8TN8C).
- [Vil+19] Vincent Villette, Mariya Chavarha, Ivan K. Dimov, Jonathan Bradley, Lagnajeet Pradhan, Benjamin Mathieu, Stephen W. Evans, Simon Chamberland, Dongqing Shi, Renzhi Yang, Benjamin B. Kim, Annick Ayon, Abdelali Jalil, François St-Pierre, Mark J. Schnitzer, Guoqiang Bi, Katalin Toth, Jun Ding, Stéphane Dieudonné, and Michael Z. Lin. “Ultrafast Two-Photon Imaging of a High-Gain Voltage Indicator in Awake Behaving Mice”. In: *Cell* (Dec. 12, 2019). DOI: [10.1016/j.cell.2019.11.004](https://doi.org/10.1016/j.cell.2019.11.004).
- [Zha+17] Yaoyu Zhang, Yanyang Xiao, Douglas Zhou, and David Cai. “Spike-Triggered Regression for Synaptic Connectivity Reconstruction in Neuronal Networks”. In: *Frontiers in Computational Neuroscience* (2017). DOI: [10.3389/fncom.2017.00101](https://doi.org/10.3389/fncom.2017.00101).
- [Zoc+00] M Zochowski, M Wachowiak, Cx Falk, Lb Cohen, Yw Lam, S Antic, and D Zecevic. “Imaging Membrane Potential with Voltage-Sensitive Dyes”. In: *The Biological Bulletin* (Feb. 2000). DOI: [10.2307/1542798](https://doi.org/10.2307/1542798).