Internship Report, Automated tests for Sea-of-Noise networks on Neuromorphic Hardware, 3.9.-28.10.2018

Frederik Mayer

17.10.2018

Abstract

The neuromorphic hardware BrainScaleS [**1**] at the university of Heidelberg is an attempt to emulate neural processes of the human brain on a largely simplified scale.
Due to the analogue implementation of the neuron circuits, experiments on the BrainScaleS hardware are subject to partly large fluctuations caused by deviations in the hardware components.
Other sources of errors and confusion are the frequent updates of the calibration data and other software components.
In this internship, it was thus tried to implement a series of automated test programs, that assure basic functionalities of the hardware to produce comparable data over time.

1 Spiking Leak over Threshold Neuron

A neuron on the BrainScaleS hardware is specified by some parameters.
In the neuron model used by the emulation, the membrane voltage converges exponentially against a resting potential V_{rest}.
If a certain threshold level V_{th} is reached, the voltage is reset to the reset voltage V_{reset}.
The membrane voltage of a neuron whose resting potential is higher than its threshold value should behave as can be seen in Figure 1.1:

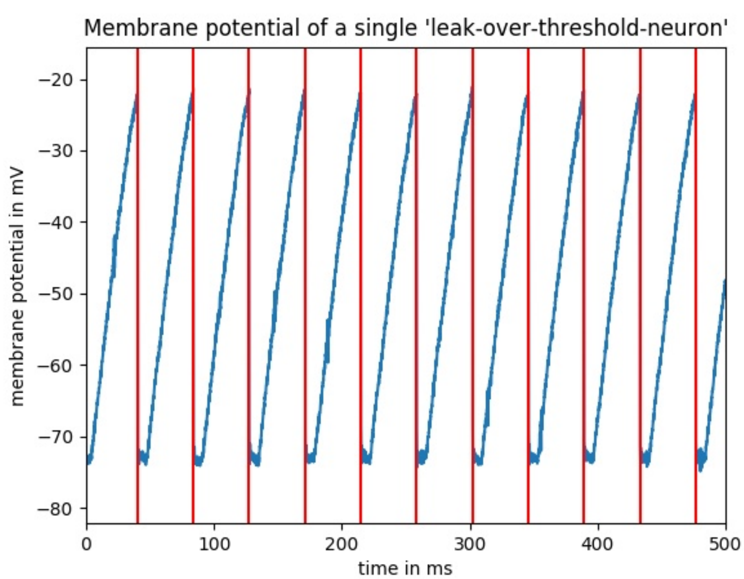

Figure 1.1: Membrane potential of a single firing Leak-over-Threshold(LoT)-neuron

The real value of those voltage levels, as well as the firing rate of the neuron, resulting from those, should not vary too much over time and using different HICANNs.
This is to be tested automatically.
For the calibration, the experimental setup is run in total 400 times, on 2 different wafers and on 10 different HICANNs each.
The mean firing rates are then averaged and saved to a calibration file together with the standard deviation.
As can be seen in Figure 1.2, the biggest deviation can be found due to differences in the HICANN-configurations, whereas there is no significant deviation in firing rates between the two used wafers 21 and 33:

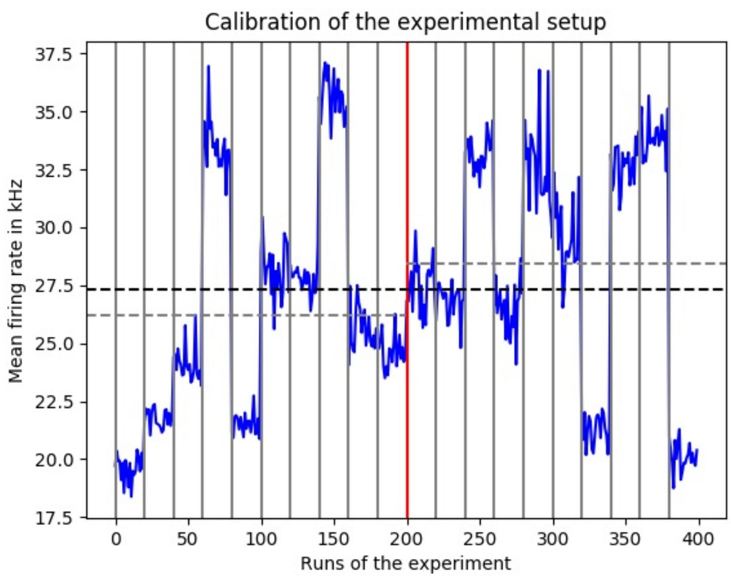

Figure 1.2: Mean firing rate of each experiment run; the grey vertical lines indicate a switch to a different HICANN.
Measurements on the left of the red vertical line are done on wafer 33, to the right on wafer 21.
The grey dashed lines display the mean value of each wafer, the black line the overall average.

The distribution of mean firing rates looks as follows:

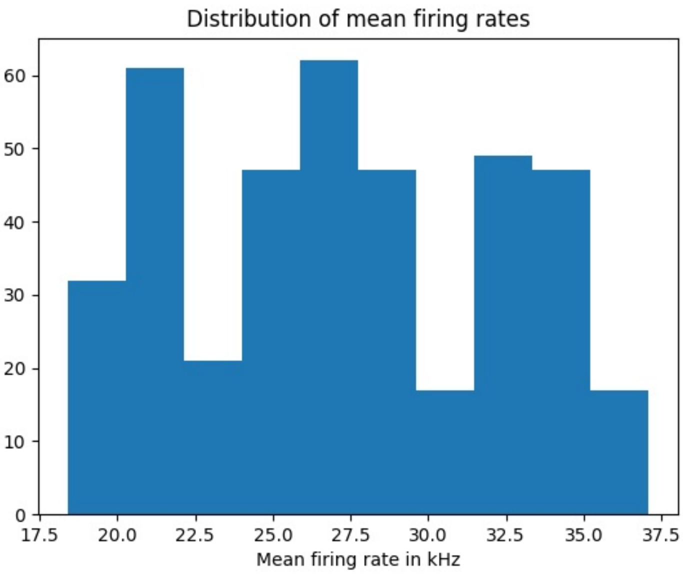

Figure 1.3: Distribution of mean firing rates, measurements on wafers 33 and 21, 10 HICANNs each

The automated test runs the experiment on the given HICANN and wafer, measures the mean firing rate and compares it to the average firing rate saved in the calibration file.
The test is rejected if the deviation is greater than 3σ.

2 LoT neuron connected to receiving neuron

As a next step, we want to implement a setup, where a LoT neuron as in the previous part is connected to a second neuron, whose threshold is higher than its resting potential.
The weight of the connection is varied, in our example it takes the values 5, 10 and 15.
The resulting membrane voltages of both neurons can be seen in Figure 2.1-2.3:

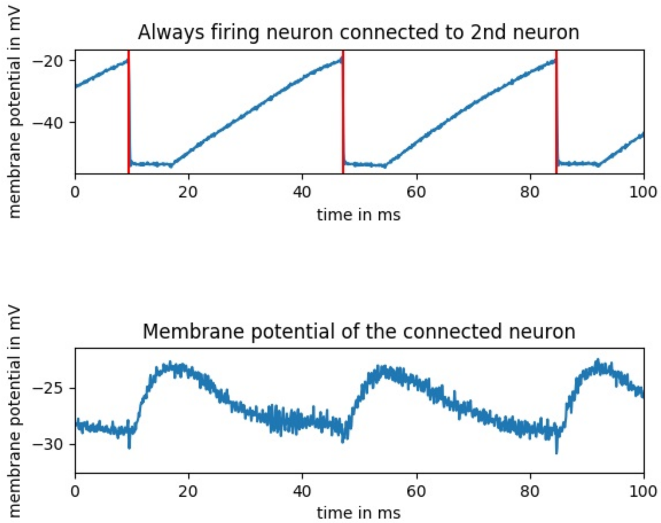

Figure 2.1: Membrane potential for the two neurons, connected with synapse strength 5

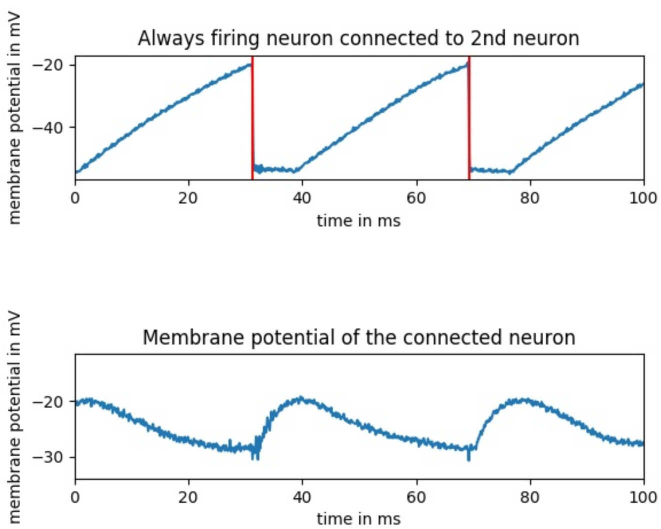

Figure 2.2: Membrane potential for the two neurons, connected with synapse strength 10

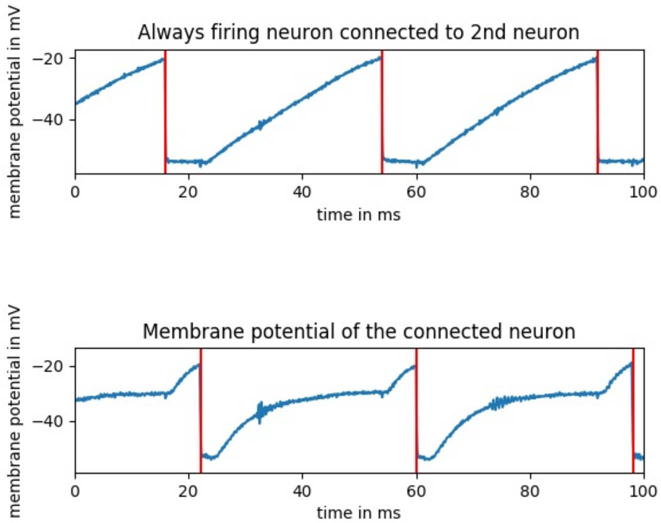

Figure 2.3: Membrane potential for the two neurons, connected with synapse strength 15

As can be seen, the receiving neuron does not spike in the case of the weight set to 10, whereas at weight 15, every spike of the firing neuron causes the receiving neuron to spike exactly once.
This behaviour should be maintained over time and on different HICANNs. The automated test therefore rejects an experiment, if *either* the receiving neuron spikes at least once at synapse weight 10 *or* there is at least one spike of the firing neuron that does not evoke a spike of the receiving neuron at synapse weight 15.
If we examine Figure 2.3, we see that each spike of the receiving neuron is delayed by a delay time $t_{delay} < 10ms \ll t_{isi}$ with t_{isi} being the average inter-spike interval of the firing neuron.
We therefore define a spike of the receiving to be corresponding to a firing neuron spike if it occurs 10 ms or less after this spike.
If a firing neuron spike occurs less than 10 ms before the end of the measurement, there is no corresponding answer spike needed, because it might have occurred after the experiment stop.

3 Sea of Noise network

A sea of noise network is a large population of LoT-neurons which are connected with an amount of randomly picked inhibitory and excitatory synapses. [**2**] In our example, the population size is 100 and a fixed-probability-connector is used.
Each pair of neurons in the population has a probability of 15 % to be connected by an inhibitory synapse and of 1 % to be connected in an

excitatory way.
Without any connections, all the network neurons would spike permanently at a regular frequency that is defined by the neuron parameters.
Some of the spikes are however inhibited by inhibitory incoming spikes from other network neurons.
The large number of connections between random neurons causes the spiking behaviour of a single neuron to be basically unpredictable, thus creating a pseudo-random spiking behaviour of the whole system.
This can be seen in the irregular distances of the spikes in Figure 3.1.

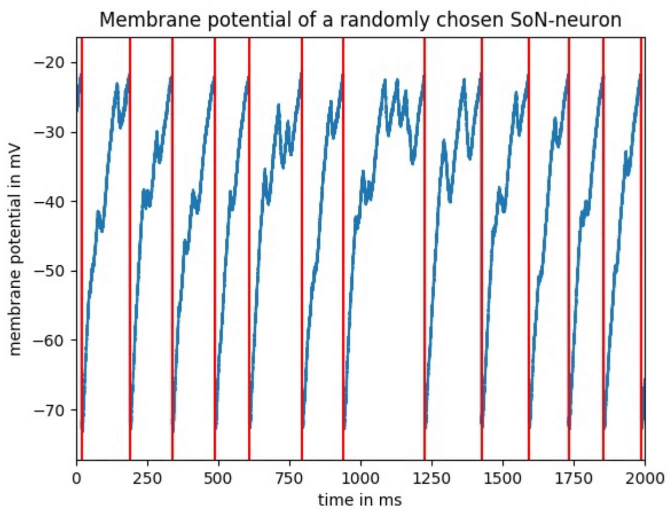

Figure 3.1: Membrane voltage over time of a randomly picked network-neuron

One way to quantify the network-behaviour is to take a look at the coefficient of variation (CV) of the inter-spike-intervals (ISIs).
The CV of a physical quantity is defined through: [**3**]

$$CV = \frac{\sigma}{\mu} \tag{1}$$

If we plot the CV of the ISIs over the mean firing rate for all network neurons, we obtain the following result:

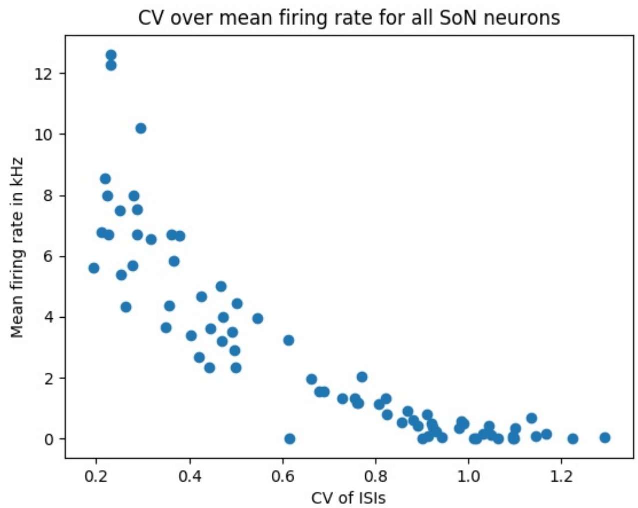

Figure 3.2: CV of ISIs over mean firing rate for all 100 network neurons

In the case of a perfectly random process, a discrete variable would take the form of a Poisson distribution with CV 1.
As can be seen, not all of the neurons' CVs are in this range.
Moreover, the "randomness" of the firing is clearly rate-dependent, a neuron firing less random when its mean firing rate is higher.
This can be explained keeping in mind the way we're producing the "randomness":
Without the inhibitory connections, the neurons would all spike very regularly at a frequency mainly determined by their time constant τ_{mem} and the refractory time τ_{ref}.
These parameters are set to $\tau_{mem} = 20ms$ and $\tau_{ref} = 4ms$ in our set-up.
The time that the neuron takes to reach the threshold level of $V_{thresh} = -20mV$ from the reset level $V_{reset} = -70mV$ with the resting potential being $V_{rest} = +10mV$ is given by solving the solution:

$$V(t) = V_{reset} + (V_{rest} - V_{reset})(1 - \exp(-\frac{t}{\tau_{mem}})) \stackrel{!}{=} V_{thresh} \quad (2)$$

and adding the refraction time, which in total leads to:

$$t = \tau_{ref} - \tau_{mem} * ln(1 - \frac{V_{thresh} - V_{reset}}{V_{rest} - V_{reset}}) \quad (3)$$

For our values, we obtain $\tau_{total} = 23.6ms$, which corresponds to a firing frequency of 42.4 Hz.
Therefore, the neurons that fire with a frequency closer to this value have to receive a very small amount of inhibitory signals in order to achieve such a high firing rate.

This, on the other hand, means that they are much less influenced by the pseudo-random incoming signals and show a firing behaviour closer to the "unconnected case" with small CVs.
In the case of Figure 3.2, the maximum average firing rate ($\approx 12kHz$) lies quite far from the frequency calculated from equation (3).
This is due to the erroneous calibration of the hardware, that causes deviations of all potentials in the order of 10-30mV.
It is likely that the actual maximum firing rate given by (3) and the actual potential values is only slightly higher than the measured 12 kHz and that those neurons are therefore almost uninfluenced by the inhibitory signals.

The automated test for the Sea-of-Noise-Network should test if

a) the average firing rate of the neurons is approximately constant over time and on different HICANNs/wafers and if

b) the characteristic curve in the CV-mean-rate-diagram (Figure 3.2) can be reproduced.

For a), the same test procedure as in the first chapter (Spiking LoT neuron) can be used, with calibration data obtained on different HICANNs and wafers setting a 3σ-range that has to be kept by any future measurement.
For b), we fit a function to the characteristic curve and save the values and covariances of the fit parameters as calibration data. A future measurement will be processed equally, again the test will be rejected if the fit parameters deviate by more than 3σ from the calibration data.

3.1 Problems in implementing the CV-test

When calibrating the target fit parameters for test b), we found ourselves confronted with some problems:
While some of the CV-diagrams could be fitted very accurately (e.g. in Figure 3.3), the fit function failed to display the desired curve shape of a parabola open at the top, even though the CV-distribution apparently showed the desired appearance. (Figure 3.4 & 3.5)

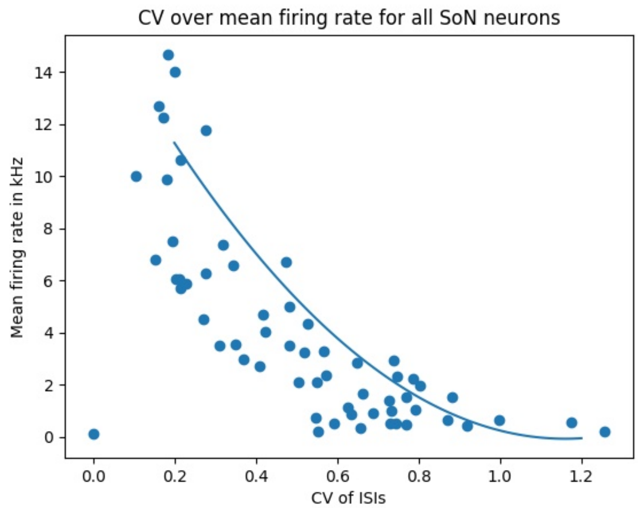

Figure 3.3: CV diagram with correct fit curve

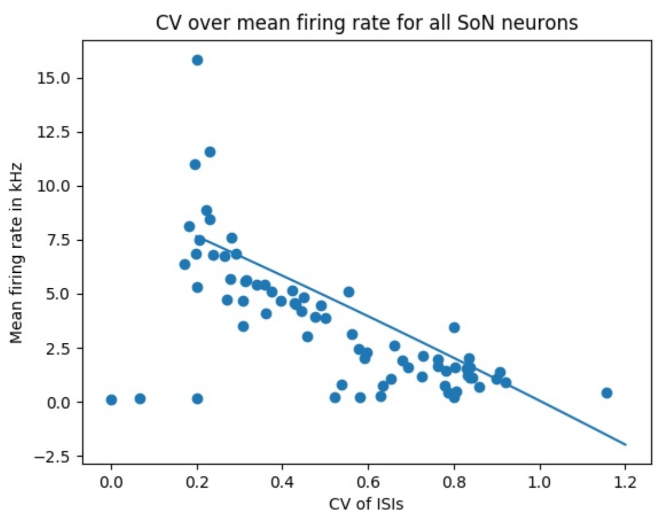

Figure 3.4: CV diagram with failed fit curve

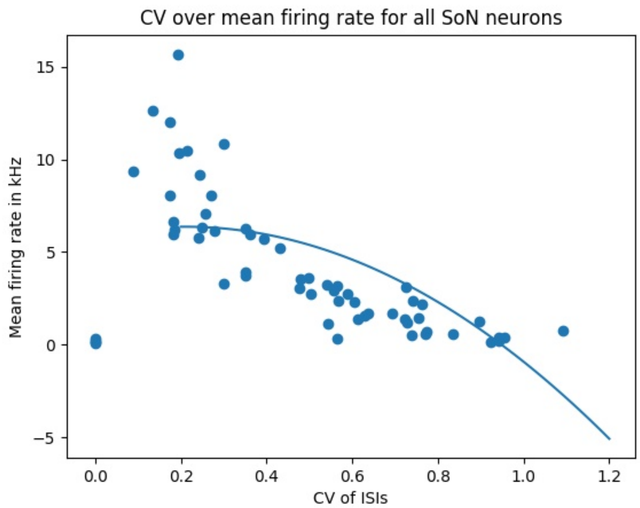

Figure 3.5: CV diagram with failed fit curve

Since this problem would cause the test to fail even though the SoN-network is configured correctly, it was decided to cancel this test and use only the rate-based test a).

Again, we take a look at the calibration data and plot the mean rate over runs on different HICANNs:

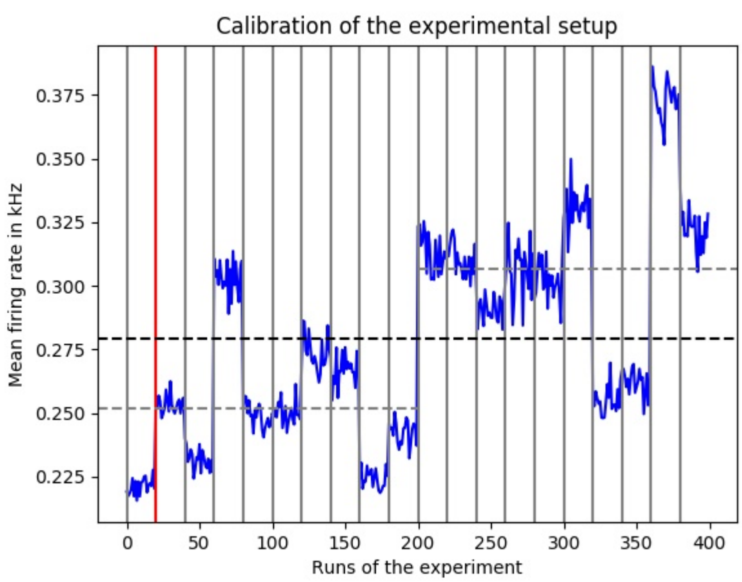

Figure 3.6: Mean firing rate of each experiment run; the grey vertical lines

indicate a switch to a different HICANN.
Measurements on the left of the red vertical line are done on wafer 33, to the right on wafer 21.
The grey dashed lines display the mean value of each wafer, the black line the overall average.

The distribution of mean firing rates looks as follows:

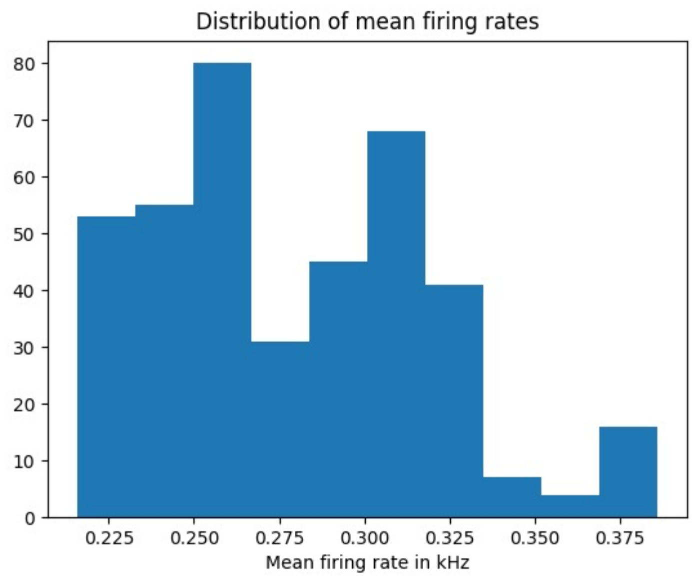

Figure 3.7: Distribution of mean firing rates, measurements on wafers 33 and 21, 10 HICANNs each

In this case, it does look like the 2 wafers would produce different average firing rates.
The deviation is not significant though, also our sample of only 10 HICANNs per wafer is not very big.
That there is a deviation independent of any HICANN change is evident, as the random connections between the neurons are drawn randomly again after every new configuration.
The calculated target rate and deviation are again used to decide if a tested SoN-network is accepted or rejected by the automated test.

4 Conclusion

The analogue implementation of the neuromorphic hardware lead to sometimes unanticipated errors and difficulties.
To be able to distinguish more easily between mistakes in the own

implementation and errors in the configuration of the hardware, automated tests for some standard examples of neural networks were developed.
The question of finding an adequate measure of quality for a given configuration and thus determine when a test should be accepted or rejected had to be answered for each example separately.
Especially in the case of the Sea-of-Noise-Network, it was hard to find a good measure for a good-working-network.
Since the attempt to fit a curve to the CV-rate-diagram failed due to the lack of uniqueness of the fit parameters, it was decided to apply the previously used mean-firing-rate-test as well to this experiment.
The thought behind this decision was, that a correctly configured SoN-network surely has to comply with more conditions than a correct mean firing rate. Nevertheless, the probability of a defective network producing a firing rate in exactly the range given by the calibration data is very low. A constant firing rate therefore is an acceptable measure for problems resulting from changes in the hardware behaviour.

During the time of the test configuration, the experiment results remained constant and the automated tests were passed for all valid HICANN configurations.

References

[1] `https://electronicvisions.github.io/hbp-sp9-guidebook/index.html`

[2] Towards Spikebased Expectation Maximization in a ClosedLoop Setup on an Accelerated Neuromorphic Substrate, Felix Schneider, Bachelorarbeit (2018)

[3] `https://en.wikipedia.org/wiki/Coefficient_of_variation`