

On the role of noise in signal detection

Dennis Huben

August 2, 2024

Contents

1	Suprathreshold stochastic resonance	2
1.1	Introduction	2
1.2	Methods	4
1.3	Simulations with more neurons	6
1.4	Different frequency ranges	26
1.5	Narrow-/wideband	26
1.6	Discussion	27
2	Theory	31
2.1	Firing rates	31
2.2	Refractory period	36
3	Electric fish	38
3.1	Introduction	38
3.2	Methods	38
3.3	How to determine noisiness	39
3.4	Results	48
3.5	Discussion	48
4	Discussion: Combining experiment and simulation	48
5	Literature	48

1 Suprathreshold stochastic resonance

1.1 Introduction

In any biological system, there is a limit to the precision of the components making up that system. This means that even without external input the spike times of each individual neurons will have some variation and will not be perfectly regular. Increasing the precision has a cost in energy requirement (Schreiber et al., 2002) but may not even be desirable.

In populations of neurons, representation of a common stimulus can be improved by population heterogeneity (Ahn et al., 2014). The source of this heterogeneity could for example be a different firing threshold for each neuron. Alternatively, the improvement can be achieved by adding noise to the input of neurons (Shapira et al., 2016). The effect of adding noise to a sub-threshold signal, a phenomenon known as "stochastic resonance" (SR), has been very well investigated during the last decades (Benzi et al., 1981; Gammaitoni et al., 1998; Shimokawa et al., 1999). The noise added to a signal makes it more likely that the signal reaches the detection threshold so that it triggers a spike in a neuron. But often in nature the goal is not to simply detect a signal but to discriminate between two different signals as well as possible. For example in auditory communication it is not sufficient to detect the presence of sound but instead the goal is to encode an auditory stimulus so that an optimal amount of information is gained from the stimulus. Another example is the electrosensory communication between conspecifics in weakly electric fishes. Those fish need to for example differentiate aggressive and courtship behaviors.

More recently it has been shown that for populations of neurons the beneficial role of noise can also be true for signals which already are above the threshold (Stocks, 2000a,b, 2001; Stocks and Mannella, 2001; Beiran et al., 2018), a phenomenon termed "Suprathreshold Stochastic Resonance" (SSR). Despite the similarity in name, SR and SSR work in very different ways. The idea behind SSR is that in case of no or very weak individual noise the different neurons in the population react to the same features of a common input. Additional noise that affects each cell differently desynchronizes the response of the neurons. The spiking behavior of the neurons becomes more probabilistic than deterministic in nature. However, if the noise is too strong, the noise masks the signal and less information can be coded than would be ideally possible. In the case of infinite noise strength, no information about the signal can be reconstructed from the responses. Because some noise is beneficial and too much noise isn't, there is a noise strength where performance is best. This thesis investigates populations of neurons reacting to input signals with cutoff frequencies over a large range. Population sizes range from a single neuron to many thousands of neurons.

Here we use the Integrate-and-Fire model to simulate neuronal populations receiving a common dynamic input. We look at linear coding of signals by different sized populations of neurons of a single type, similar to the situation in weakly electric fish. We show that the optimal noise grows with population size and depends on properties of the input. We use input signals of varying frequencies widths and cutoffs,

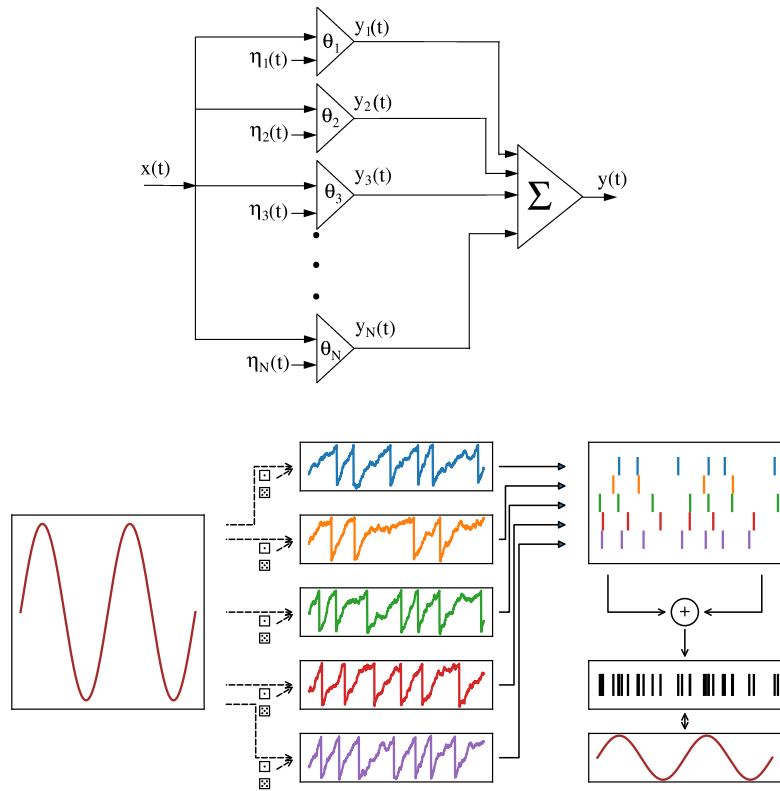


Figure 1: Array of threshold systems as described by Stocks.

along with changing the strength of the signal.

We also present the results of electrophysiological results in the weakly electric fish *Apteronotus leptorhynchus*. Because it is not obvious how to quantify noisiness in the receptor cells of these fish, we compare different methods and find that using the activation curve of the individual neurons allows for the best estimate of the strength of noise in these cells. Then we show that we can see the effects of SSR in the real world example of *A. leptorhynchus*.

1.2 Methods

We use a population neuron model using the Leaky-Integrate-And-Fire (LIF) neuron, described by the equation

$$V_t^j = V_{t-1}^j + \frac{\Delta t}{\tau_v}((\mu - V_{t-1}^j) + \sigma I_t + \sqrt{2D/\Delta t}\xi_t^j), \quad j \in [1, N] \quad (1)$$

with $\tau_v = 10ms$ the membrane time constant, $\mu = 15.0mV$ or $\mu = 10.5mV$ as offset. σ is a factor which scales the standard deviation of the input, ranging from 0.1 to 1 and I the previously generated stimulus. ξ_t are independent Gaussian distributed random variables with mean 0 and variance 1. The Noise D was varied between $1 * 10^{-12}mV^2/Hz$ and $3mV^2/Hz$. Whenever V_t was greater than the voltage threshold (10mV) a "spike" was recorded and the voltage has been reset to 0mV. V_0 was initialized to a random value uniformly distributed between 0mV and 10mV. For the first sets of simulations there was no absolute refractory period¹. In a later chapter I show that qualitatively results don't change with an added refractory period. Simulations of up to 8192 neurons were done using an Euler method with a step size of $\Delta t = 0.01ms$. Typical firing rates were around 90Hz for an offset of 15.0mV and 35Hz for an offset of 10.5mV. Firing rates were larger for high noise levels than for low noise levels.

As stimulus we used Gaussian white noise signal with different frequency cutoff on both ends of the spectrum. By construction, the input power spectrum is flat between 0 and $\pm f_c$:

$$S_{ss}(f) = \frac{\sigma^2}{2|f_c|} \Theta(f_c - |f|). \quad (2)$$

A Fast Fourier Transform (FFT) was applied to the signal so it can serve as input stimulus to the simulated cells. The signal was normalized so that the variance of the signal was 1mV and the length of the signal was 500s with a resolution of 0.01ms.

Analysis

For each combination of parameters, a histogram of the output spikes from all neurons or a subset of the neurons was created. The coherence $C(f)$ was calculated (Lindner, 2016) in frequency space as the fraction between the squared cross-spectral

¹ Absolute refractory period means a time in which the cell ignores any input and can't spike.




Figure 2: Snapshots of 200ms length from three example simulations with different noise, but all other parameters held constant. Black: Spikes of 32 simulated neurons. The green curve beneath the spikes is the signal that was fed into the network. The blue curve is the best linear reconstruction possible from the spikes. The input signal has a cutoff frequency of 50Hz. If noise is weak, the neurons behave regularly and similar to each other (A). For optimal noise strength, the neuronal population follows the signal best (B). If the noise is too strong, the information about the signal gets drowned out (C). D: Example coding fraction curve over the strength of the noise. Marked in red are the noise strengths from which the examples were taken.

density $|S_{sx}^2|$ of input signal $s(t) = \sigma I_t$ and output spikes $x(t)$, $S_{sx}(f) = \mathcal{F}\{s(t) * x(t)\}(f)$, divided by the product of the power spectral densities of input ($S_{ss}(f) = |\mathcal{F}\{s(t)\}(f)|^2$) and output ($S_{xx}(f) = |\mathcal{F}\{x(t)\}(f)|^2$), where $\mathcal{F}\{g(t)\}(f)$ is the Fourier transform of $g(t)$.

$$C(f) = \frac{|S_{sx}(f)|^2}{S_{ss}(f)S_{xx}(f)} \quad (3)$$

The coding fraction γ (Gabbiani and Koch, 1996; Krahe et al., 2002) quantifies how much of the input signal can be reconstructed by an optimal linear decoder. It is 0 in case the input can't be reconstructed at all and 1 if the signal can be perfectly reconstructed (Gabbiani et al., 1996). It is defined by the reconstruction error ϵ^2 and the variance of the input σ^2 :

$$\gamma = 1 - \sqrt{\frac{\epsilon^2}{\sigma^2}}. \quad (4)$$

The variance is

$$\sigma^2 = \langle (s(t) - \langle s(t) \rangle)^2 \rangle = \int_{f_{low}}^{f_{high}} S_{ss}(f) df. \quad (5)$$

The reconstruction error is defined as

$$\epsilon^2 = \langle (s(t) - s_{est}(t))^2 \rangle = \int_{f_{low}}^{f_{high}} S_{ss} - \frac{|S_{sx}|^2}{S_{xx}} = \int_{f_{low}}^{f_{high}} S_{ss}(f)(1 - C(f)) df \quad (6)$$

with the estimate $s_{est}(t) = h * x(t)$. h is the optimal linear filter which has Fourier Transform $H = \frac{S_{sx}}{S_{xx}}$ (Gabbiani, 1996).

We then analyzed coding fraction as a function of these cutoff frequencies for different parameters (noise strength, signal amplitude, signal mean/firing rate) in the limit of large populations. The limit was considered reached if the increase in coding fraction gained by doubling the population size is small (4%)(??). For the weak signals ($\sigma = 0.1mV$) combined with the strongest noise ($D = 10^{-3} \frac{mV^2}{Hz}$), convergence was not reached for a population size of 2048 neurons for both threshold values. The same is true for the combination of the weak signal, close to the threshold ($\mu = 10.5mV$) and high frequencies (200Hz).

1.3 Simulations with more neurons

Noise makes neurons' responses different from each other

If noise levels are low (fig. 2 a)), neurons within a population with behave very similarly to each other. There is little variation in the spike responses of the neurons to a signal, and recreating the signal is difficult. If the strength of the noise is increasing, at some point the coding fraction will also begin increasing. The signal recreation will become better as the responses of the different neurons begin to deviate from each other. When noise strength is increased even further at some point a peak coding




Figure 3: Coherence for a signal with $f_{cutoff} = 200 \text{ Hz}$. Coherence for a small and a large population, each at weak and strong noise values. For weak noise, the curves are indistinguishable from one another. For strong noise an increase in population size allows much better reconstruction of the input. For the small population size weak noise in the simulated neurons allows for better signal reconstruction. The line marks the average firing rate (about 91 Hz) of the neurons in the population.

fraction is reached. This point is the optimal noise strength for the given parameters (fig. 2 b)). If the strength of the noise is increased beyond this point, the responses of the neurons will be determined more by random fluctuations and less by the actual signal, making reconstruction more difficult (fig. 2 c)). At some point, signal encoding breaks down completely and coding fraction goes to 0.

Large population size is only useful if noise is strong

We see that an increase in population size leads to a larger coding fraction until it hits a limit which depends on noise. For weak noise the increase in coding fraction with an increase in population size is low or non-existent. This can be seen in figure 1.3 c) where the red ($10^{-5} \frac{mV^2}{Hz}$) and orange ($10^{-4} \frac{mV^2}{Hz}$) curves (relatively weak noise) saturate for relatively small population size (about 8 neurons and 32 neurons respectively). An increase in population size also leads to the optimum noise level moving towards stronger noise (green dots in figure 1.3 a)). A larger population can exploit the higher noise levels better. Within the larger population the precision of the individual neurons becomes less important. After the optimum noise where peak coding fraction is reached, an increase in noise strength leads to a reduction in coding fraction. If the noise is very strong, coding fraction can reach approximately 0. This happens earlier (for weaker noise) in smaller populations than in larger populations. Together those facts mean that for a given noise level and population size, coding fraction might already be declining; whereas for larger populations, coding fraction




Figure 4: Rasterplots and reconstructed signals for different population sizes; insets show signal spectrum. Rasterplots show the responses of neurons in the different populations. Blue lines show the reconstruction of the original signal by different sets of neurons of that population size. A: Each blue line is the reconstructed signal from the responses of a population of 4 neurons. B: Each blue line is the reconstructed signal from the responses of a population of 16 neurons. C: The same for 64 neurons. D: The same for 256 neurons. Larger population sizes lead to observations which are not as dependent on random fluctuations and are therefore closer to each other. [DH: langsames signal hier nehmen(!?)]

can still be increasing. A given amount of noise can lead to a very low coding fraction in a small population, but to a greater coding fraction in a larger population. (figure 1.3 c), blue and purple curves). The noise levels that work best for large populations are in general performing very bad in small populations. If coding fraction is supposed to reach its highest values and needs large populations to do so, the necessary noise strength will be at a level, where basically no encoding will happen in a single neurons or small populations.

Influence of the input is complex

Two very important variables are the mean strength of the signal, equivalent to the baseline firing rate of the neurons and the strength of the signal. A higher baseline firing rate leads to a larger coding fraction. In our terms that means that a mean signal strength μ that is much above the signal will lead to higher coding fractions than if the signal strength is close to the threshold (see figure 1.3 b), orange curves are above the green curves). The influence of the signal amplitude σ is more complex. In general, at small population sizes, larger amplitudes appear to work better, but with large populations they might perform as well or even better than stronger signals

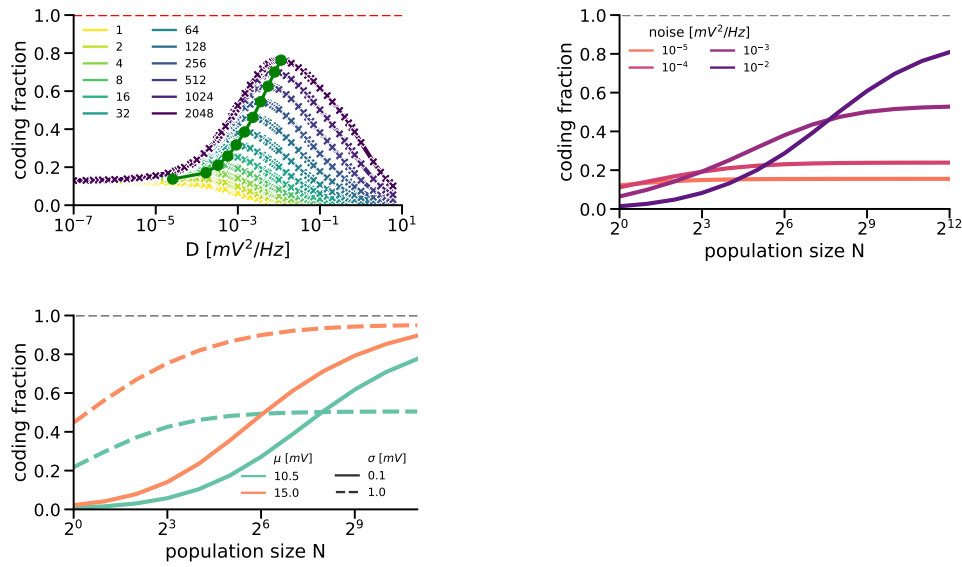


Figure 5: A: Coding fraction as a function of noise for different population sizes. Green dots mark the peak of the coding fraction curve. Increasing population size leads to a higher peak and moves the peak to stronger noise. B: Coding fraction as a function of population size. Each curve shows coding fraction for a different noise strength. C: Peak coding fraction as a function of population size for different input parameters. [DH: needs information about noise]

(figure 1.3 c), dashed curves vs solid curves.)

Slow signals are more easily encoded

To encode a signal well, neurons in a population need to keep up with the rising and falling of the signal. Signals that change fast are harder to encode than signals which change more slowly. When a signal changes more gradually, the neurons can slowly adapt their firing rate. A visual example can be seen in figure 13. When all other parameters are equal, a signal with a lower frequency is easier to recreate from the firing of the neurons. In the rasterplots one can see especially for the 50Hz signal (bottom left) that the firing probability of each neuron follows the input signal. When the input is low, almost none of the neurons fire. The result are the “stripes” we can see in the rasterplot. The stripes have a certain width which is determined by the signal frequency and the noise level. When the signal frequency is low, the width of the stripes can’t be seen in a short snapshot. For the 50Hz signal in this example we can clearly see a break in the firing activity of the neurons at around 25ms. The slower




Figure 6: Rasterplots, input signal and reconstructed signals for different cutoff frequencies; insets show each signal spectrum. Shown here are examples taken from 500s long simulations. Rasterplots show the firing of 64 LIF-neurons. Each row corresponds to one neuron. Blue lines below the rasters are the input signal, the orange line the reconstruction, calculated by convolving the spikes with the optimal linear filter. Reconstruction is closer to the original signal for slower signals than for higher frequency signals. The different time scales lead to spike patterns which appear very different from each other.

changes in the signal allow for the reconstruction to follow the original signal more closely. For the 200Hz signal there is little structure to be seen in the firing behaviour of the population and instead that behaviour looks chaotic. Something similar can be said for the 1Hz signal. Because the peaks are about 1s apart from each other, a snapshot of 400ms cannot capture the structure of the neuronal response. Instead what we see is a very gradual change of the firing rate following the signal. Because the change is so gradual, the reconstructed signal follows the input signal very closely.

Fast signals are harder to encode - noise can help with that

For low frequency signals, the coding fraction is almost always at least as large as the coding is for signals with higher frequency. For the parameters we have used there is very little difference in coding fraction for a random noise signal with frequencies of 1Hz and 10Hz respectively (figure 1.3, bottom row). For all signal frequencies and amplitudes a signal mean much larger than the threshold ($\mu = 15.0mV$, with the threshold at $10.0mV$) results in a higher coding fraction than the signal mean closer to the threshold ($\mu = 10.5mV$). Firing rates of the neurons is much higher at the large input: about 90 Hz vs. 30 Hz for the lower signal mean. We also find that for the signal mean which is further away from the threshold for the loss of coding fraction

from the 10Hz signal to the 50Hz signal is smaller than for the lower signal mean. This is partially explained by the firing rate of the neurons: Around the firing rate the signal encoding is weaker (see figure 3). In general, an increase in signal frequency and bandwidth leads to a decrease in the maximum achievable coding fraction. This decrease is smaller if the noise is stronger. In some conditions, a 50 Hz signal can be encoded as well as a 10 Hz signal (fig. 1.3 d)).

[DH: TODO: frequency vs optimum noise; For slower signals, coding fraction converges faster in terms of population size (figure 1.3). This (convergence speed) is also true for stronger signals as opposed to weaker signals. For slower signals the maximum value is reached for weaker noise.]

A tuning curve allows calculation of coding fraction for arbitrarily large populations

To understand information encoding by populations of neurons it is common practice to use simulations. However, the size of the simulated population is limited by computational power. We demonstrate a way to circumvent these limitations, allowing to make predictions in the limit case of large population size. We use the interpretation of the tuning curve as a kind of averaged population response. To calculate this average, we need relatively few neurons to reproduce the response of an arbitrarily large population of neurons. This allows the necessary computational power to be greatly reduced.

At least for slow signals, the spiking probability at a given point in time is determined by the signal power in this moment. The population response should simply be proportional to the response of a single neuron. This average firing rate is reflected in the tuning curve. We can look at the average firing rate for the input to find the spiking probability and how this probability changes with noise.

For faster signals, the past of the signal plays a role: after a spike there is a short period where the simulated neuron is unlikely to fire again, even if there is no explicit refractory period. If the next spike falls into that period, fewer neurons will spike than they would have without the first spike. We have also seen before that faster signals aren't encoded as well as slower signals; but the results we receive from using the tuning curve this way is frequency-independent.

The noise influences the shape of the tuning curve, with stronger noise linearizing the curve. The linearity of the curve is important, because coding fraction is a linear measure. For strong input signals (around 13mV) the curve is almost linear, resulting in coding fractions close to 1. For signal amplitudes in this range firing rate is almost independent of noise strength. This tells us that the increase in coding fraction that follows a change in noise strength we saw in previous chapters is not simply due to the neurons spiking more frequently. For slow signals (1Hz cutoff frequency, up to 10Hz) the results from the tuning curve and the simulation for large populations of neurons match very well (figure 1.3) over a range of signal strengths, base inputs to the neurons and noise strength. This means that the LIF-neuron tuning curve gives us a very good approximation for the limit of encoded information that can be achieved

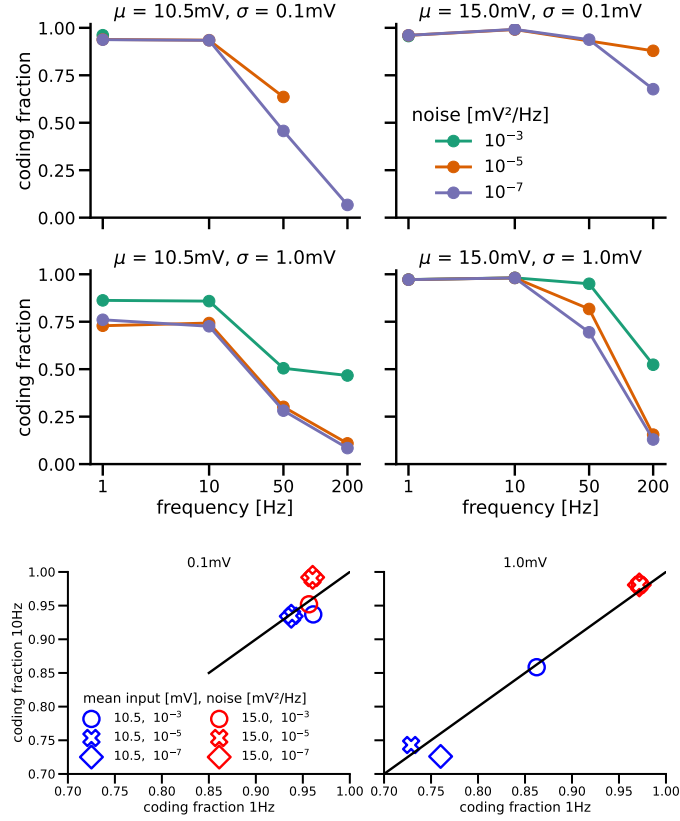


Figure 7: **A-D:** Coding fraction in the large population limit as a function of input signal frequency for different parameters. Each curve represents a different noise strength. Points are only shown when the coding fraction increased by less than 2% when population size was increased from 1024 to 2048 neurons. For small amplitudes ($\sigma = 0.1\text{mV}$, A & B) there was no convergence for a noise of $10^{-3}\text{mV}^2/\text{Hz}$. Coding fraction decreases for faster signals (50Hz and 200Hz). In the large population limit, stronger noise results in coding fraction at least as large as for weaker noise. **E, F:** Comparison of the coding fraction in the large population limit for a 1Hz signal and a 10Hz signal. Shapes indicate noise strength, color indicates mean signal input (i.e. distance from threshold). Left plot shows an amplitude of $\sigma = 0.1\text{mV}$, the right plot shows $\sigma = 1.0\text{mV}$. The diagonal black line indicates where coding fractions are equal.

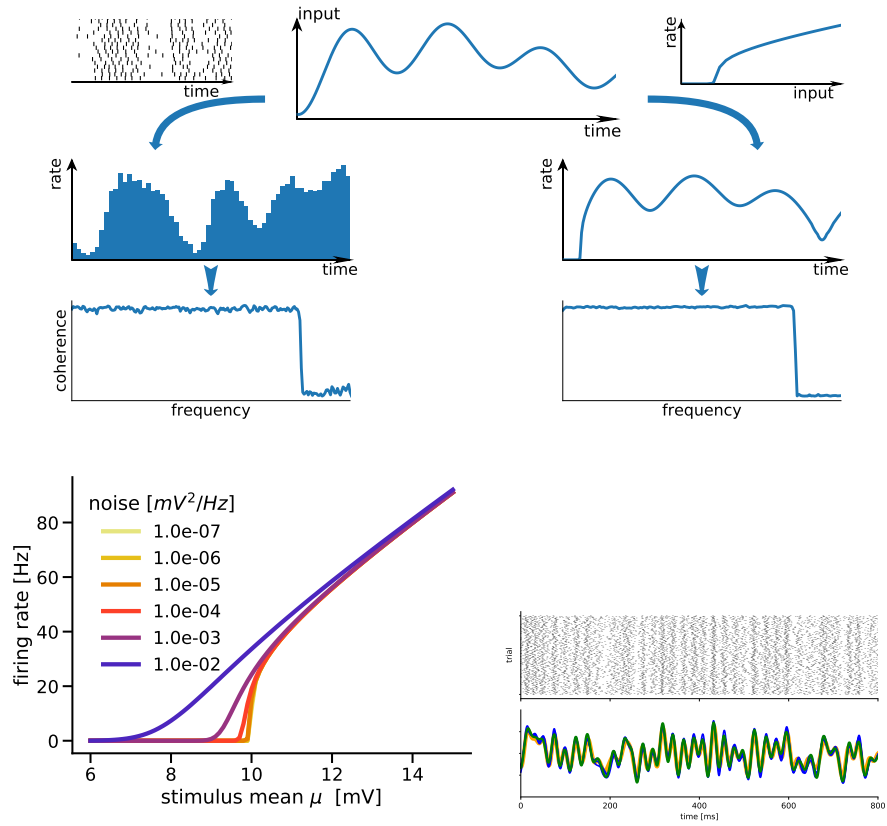


Figure 8: Two ways to arrive at coherence and coding fraction. Left: The input signal (top, center) is received by LIF-neurons. The spiking of the neurons is then binned and coherence and coding fraction are calculated between the result and the input signal. Right: Input signal (top, center) is transformed by the tuning curve (top right). The tuning curve corresponds to a function $g(V)$, which takes a voltage as input and yields a firing rate. Output is a modulated signal. We calculate coherence and coding fraction between input voltage and output firing rate. If the mean of the input is close to the threshold, as is the case here, inputs below the threshold all get projected to 0. This can be seen here at the beginning of the transformed curve. Bottom left: Tuning curves for different noise levels. x-Axis shows the stimulus strength in mV, the y-axis shows the corresponding firing rate. For low noise levels there is a strong non-linearity at the threshold. For increasing noise, firing rate becomes larger than 0 for progressively weaker signals. For strong stimuli (roughly 13mV and more) there is little different in the firing rate depending on the noise.




Figure 9: Tuningcurve works for 10Hz but not for 200Hz.

by summing over independent, identical LIF-neurons with intrinsic noise. For faster signals, the coding fraction calculated through the tuning curve stays constant, as the tuning curve only deforms the signal. As shown in figure 1.3 e) and f), the coding fraction of the LIF-neuron ensemble drops with increasing frequency. Hence for high frequency signals the tuning curve ceases to be a good predictor of the encoding quality of the ensemble.

For high-frequency signals, the method does not work. The effective refractory period prohibits the instantaneous firing rate from being useful, because the neurons spike only in very short intervals around a signal peak. They are very unlikely to immediately spike again, so that the psth is focused around the input peaks, but there is little nuance.

We use the tuning curve to analyse how the signal mean and the signal amplitude change the coding fraction we would get from an infinitely large population of neurons (fig. 15, bottom two rows). We can see that the stronger noise always yields a larger coding fraction. This is expected because the tuning curve is more linear for stronger noise. It matches that we are observing the limit of an infinitely large population, which would be able to “average out” any noise. For coding fraction as a function of the mean we see zero or near zero coding fraction if we are far below the threshold. If we increase the mean at one point we can see that coding fraction starts to jump up. This happens earlier for stronger noise (= more linear tuning curve). The increase in coding fraction is much smoother if we use a larger amplitude (right figure). We also notice some sort of plateau, where increasing the mean does not lead to a larger coding fraction, before it begins rising close to 1. The plateau begins earlier for straighter tuning curves. For coding fraction as a function of signal amplitude we see very different results depending on the parameters. Again, we see that stronger noise leads to higher coding fraction. If we are just above or at the threshold (cen-

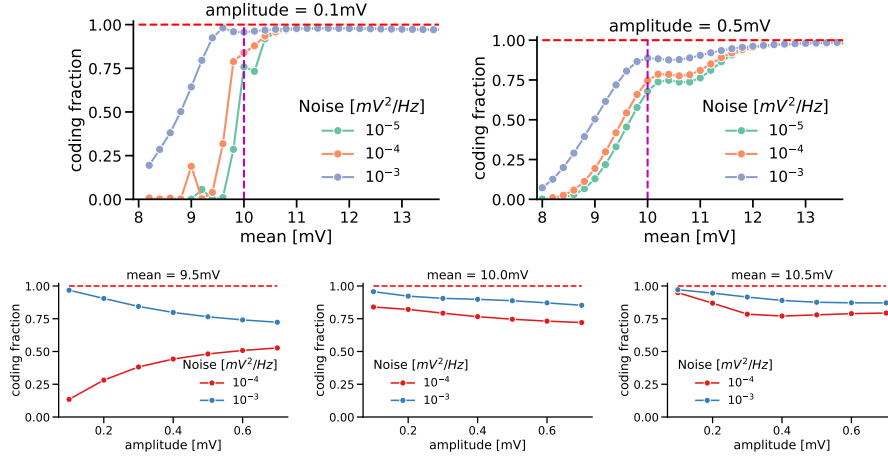


Figure 10: **A,B**: Coding signal as a function of signal mean for two different frequencies. There is little to no difference in the coding fraction. A: $\sigma = 0.1\text{mV}$. Each curve shows coding fraction as a function of the signal mean for a different noise level. The vertical line indicates the threshold. **C-E**: Coding fraction as a function of signal amplitude for different tuning curves (noise levels). Three different means, one below the threshold (9.5mV), one at the threshold (10.0mV), and one above the threshold (10.5mV).

ter and right), an increase in signal amplitude leads to a lower coding fraction. This makes sense, as more of the signal moves into the very non-linear area around the threshold. This means that for increasing amplitude an increasing fraction of the signal gets into the range of the tuning curve with 0Hz firing range, i.e. where there is no signal encoding. A very interesting effect happens if we have a mean slightly below the threshold (left): while for a strong noise we see the same effect as at or above the threshold, for weaker noise we see the opposite. This can be explained as the reverse of the effect that leads to decreasing coding fraction. Here, a larger amplitude means that the signal moves to the more linear part of the tuning curve more often. On the other hand, an increase in amplitude does not lead to worse encoding because of movement of the signal into the 0Hz part of the tuning curve – because the signal is already there, so it can't get worse. This can help explain why the coding fraction seems to saturate near 0.5: In an extreme case, the negative parts of a signal would not get encoded at all, while the positive parts would be encoded linearly.

Results

Noise makes neurons' responses different from each other

If noise levels are low (fig. 2 a)), neurons within a population with behave very similarly to each other. There is little variation in the spike responses of the neurons to a signal, and recreating the signal is difficult. If the strength of the noise is increasing, at some point the coding fraction will also begin increasing. The signal recreation will become better as the responses of the different neurons begin to deviate from each other. When noise strength is increased even further at some point a peak coding fraction is reached. This point is the optimal noise strength for the given parameters (fig. 2 b)). If the strength of the noise is increased beyond this point, the responses of the neurons will be determined more by random fluctuations and less by the actual signal, making reconstruction more difficult (fig. 2 c)). At some point, signal encoding breaks down completely and coding fraction goes to 0.

Large population size is only useful if noise is strong

We see that an increase in population size leads to a larger coding fraction until it hits a limit which depends on noise. For weak noise the increase in coding fraction with an increase in population size is low or non-existent. This can be seen in figure 1.3 c) where the red ($10^{-5} \frac{mV^2}{Hz}$) and orange ($10^{-4} \frac{mV^2}{Hz}$) curves (relatively weak noise) saturate for relatively small population size (about 8 neurons and 32 neurons respectively). An increase in population size also leads to the optimum noise level moving towards stronger noise (green dots in figure 1.3 a)). **A larger population can exploit the higher noise levels better. Within the larger population the precision of the individual neurons becomes less important.** [DH: Is this discussion?] After the optimum noise where peak coding fraction is reached, an increase in noise strength leads to a reduction in coding fraction. If the noise is very strong, coding fraction can reach approximately 0. This happens earlier (for weaker noise) in smaller populations than in larger populations. Together those facts mean that for a given noise level and population size, coding fraction might already be declining; whereas for larger populations, coding fraction can still be increasing. A given amount of noise can lead to a very low coding fraction in a small population, but to a greater coding fraction in a larger population. (figure 1.3 c), blue and purple curves). The noise levels that work best for large populations are in general performing very bad in small populations. **If coding fraction is supposed to reach its highest values and needs large populations to do so, the necessary noise strength will be at a level, where basically no encoding will happen in a single neurons or small populations.** [DH: Discussion?]

Influence of the input is complex

Two very important variables are the mean strength of the signal, equivalent to the baseline firing rate of the neurons and the strength of the signal. A higher baseline firing rate leads to a larger coding fraction. In our terms that means that a mean

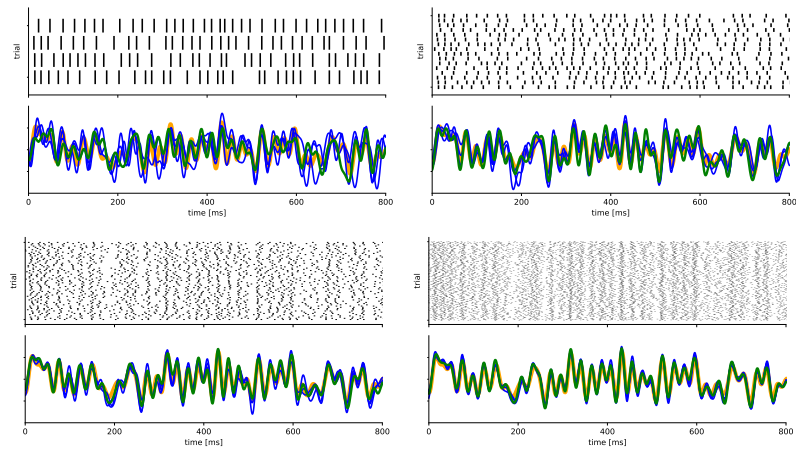


Figure 11: Rasterplots and reconstructed signals for different population sizes; insets show signal spectrum. Rasterplots show the responses of neurons in the different populations. Blue lines show the reconstruction of the original signal by different sets of neurons of that population size. A: Each blue line is the reconstructed signal from the responses of a population of 4 neurons. B: Each blue line is the reconstructed signal from the responses of a population of 16 neurons. C: The same for 64 neurons. D: The same for 256 neurons. Larger population sizes lead to observations which are not as dependent on random fluctuations and are therefore closer to each other. [DH: langsames signal hier nehmen(!?)]




Figure 12: A: Coding fraction as a function of noise for different population sizes. Green dots mark the peak of the coding fraction curve. Increasing population size leads to a higher peak and moves the peak to stronger noise. B: Coding fraction as a function of population size. Each curve shows coding fraction for a different noise strength. C: Peak coding fraction as a function of population size for different input parameters. [DH: needs information about noise]

signal strength μ that is much above the signal will lead to higher coding fractions than if the signal strength is close to the threshold (see figure 1.3 b), orange curves are above the green curves). The influence of the signal amplitude σ is more complex. In general, at small population sizes, larger amplitudes appear to work better, but with large populations they might perform as well or even better than stronger signals (figure 1.3 c), dashed curves vs solid curves.)

Slow signals are more easily encoded

To encode a signal well, neurons in a population need to keep up with the rising and falling of the signal. Signals that change fast are harder to encode than signals which change more slowly. When a signal changes more gradually, the neurons can slowly adapt their firing rate. A visual example can be seen in figure 13. When all other parameters are equal, a signal with a lower frequency is easier to recreate from the firing of the neurons. In the rasterplots one can see especially for the 50Hz signal (bottom left) that the firing probability of each neuron follows the input signal. When




Figure 13: Rasterplots, input signal and reconstructed signals for different cutoff frequencies; insets show each signal spectrum. Shown here are examples taken from 500s long simulations. Rasterplots show the firing of 64 LIF-neurons. Each row corresponds to one neuron. Blue lines below the rasters are the input signal, the orange line the reconstruction, calculated by convolving the spikes with the optimal linear filter. Reconstruction is closer to the original signal for slower signals than for higher frequency signals. The different time scales lead to spike patterns which appear very different from each other.

the input is low, almost none of the neurons fire. The result are the “stripes” we can see in the rasterplot. The stripes have a certain width which is determined by the signal frequency and the noise level. When the signal frequency is low, the width of the stripes can’t be seen in a short snapshot. For the 10Hz signal in this example we can clearly see a break in the firing activity of the neurons at around 50ms. **[DH: There is another break at about 350ms, but the inset overlays...]**. The slower changes in the signal allow for the reconstruction to follow the original signal more closely. For the 200Hz signal there is little structure to be seen in the firing behaviour of the population and instead that behaviour looks chaotic. Something similar can be said for the 1Hz signal. Because the peaks are about 1s apart from each other, a snapshot of 400ms cannot capture the structure of the neuronal response. Instead what we see is a very gradual change of the firing rate following the signal. Because the change is so gradual, the reconstructed signal follows the input signal very closely. **it is possible for neurons to encode signals which have a higher frequency than their own firing rate (Knight 73)**

Fast signals are harder to encode - noise can help with that

For low frequency signals, the coding fraction is always at least as large as the coding is for signals with higher frequency. For the parameters we have used there is very little difference for a random noise signal with frequencies of 1Hz and 10Hz respectively (figure 1.3, bottom row). For all signal frequencies and amplitudes a signal mean much larger than the threshold ($\mu = 15.0mV$, with the threshold at $10.0mV$) results in a higher coding fraction than the signal mean closer to the threshold ($\mu = 10.5mV$). We also find that for the signal mean which is further away from the threshold for the loss of coding fraction from the 10Hz signal to the 50Hz signal is smaller than for the lower signal mean. For the fast signal (200Hz) we always find a large drop in coding fraction. The drop is less pronounced for stronger noise. Coding fractions for the 1Hz and the 10Hz signal are almost identical (fig. 1.3, bottom row). It seems a reasonable assumption that for the parameters in our simulations coding fraction can be considered converged at a frequency of 10Hz. Analysis?

[DH: Is there frequency vs. optimum noise missing? For slower signals, coding fraction converges faster in terms of population size (figure 1.3). This (convergence speed) is also true for stronger signals as opposed to weaker signals. For slower signals the maximum value is reached for weaker noise.]

A tuning curve allows calculation of coding fraction for arbitrarily large populations

To understand information encoding by populations of neurons it is common practice to use simulations. However, the size of the simulated population is limited by computational power. We demonstrate a way to circumvent these limitations, allowing to make predictions in the limit case of large population size. We use the interpretation of the tuning curve as a kind of averaged population response. To calculate this average, we need relatively few neurons to reproduce the response of an arbitrarily large population of neurons. This allows the necessary computational power to be greatly reduced. Is it possible to approximate an arbitrarily large population size in some way? Most importantly, we would like to know towards which value coding fraction converges for given parameters in the limit $N \rightarrow \infty$.

At least for slow signals, the frequency response at a given point in time is determined by the signal power in this moment (for faster signals, past plays a role; we have also seen before that faster signals aren't encoded as well as slower signals; but tuning curve is frequency-independent). Then, population response should simply be proportional to the response of a single neuron. We can look at the average firing rate for the input and how it changes with noise. This average firing rate is reflected in the tuning curve.

The noise influences the shape of the tuning curve, with stronger noise linearizing the curve. The linearity of the curve is important, because coding fraction is a linear measure. For strong input signals (around 15mV) the curve is almost linear, resulting in coding fractions close to 1. For slow signals (1Hz cutoff frequency, up to 10Hz) the results from the tuning curve and the simulation for large populations of neu-

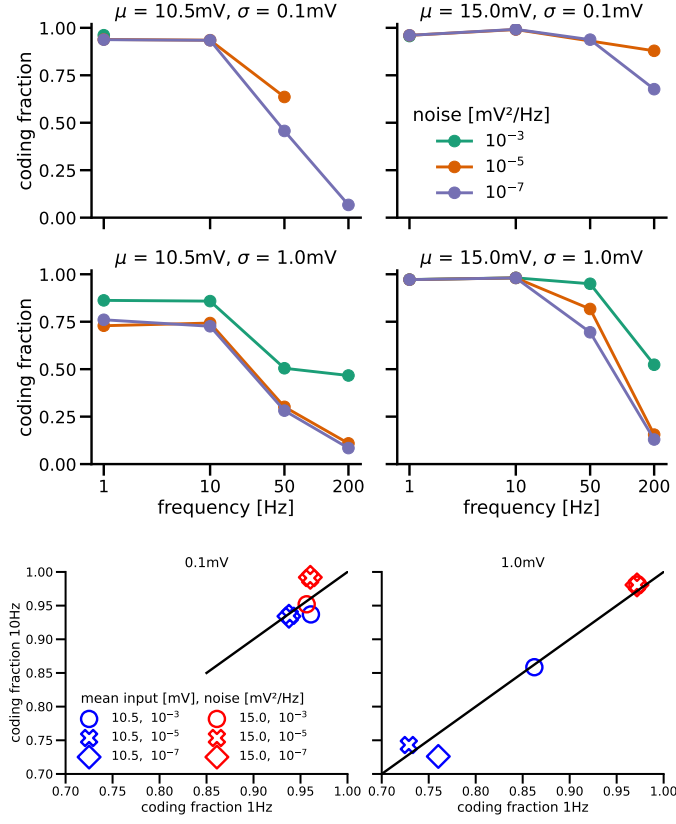


Figure 14: **A-D**: Coding fraction in the large population limit as a function of input signal frequency for different parameters. Each curve represents a different noise strength. Points are only shown when the coding fraction increased by less than 2% when population size was increased from 1024 to 2048 neurons. For small amplitudes ($\sigma = 0.1\text{mV}$, A & B) there was no convergence for a noise of $10^{-3}\text{mV}^2/\text{Hz}$. Coding fraction decreases for faster signals (50Hz and 200Hz). In the large population limit, stronger noise results in coding fraction at least as large as for weaker noise. **E, F**: Comparison of the coding fraction in the large population limit for a 1Hz signal and a 10Hz signal. Shapes indicate noise strength, color indicates mean signal input (i.e. distance from threshold). Left plot shows an amplitude of $\sigma = 0.1\text{mV}$, the right plot shows $\sigma = 1.0\text{mV}$. The diagonal black line indicates where coding fractions are equal.

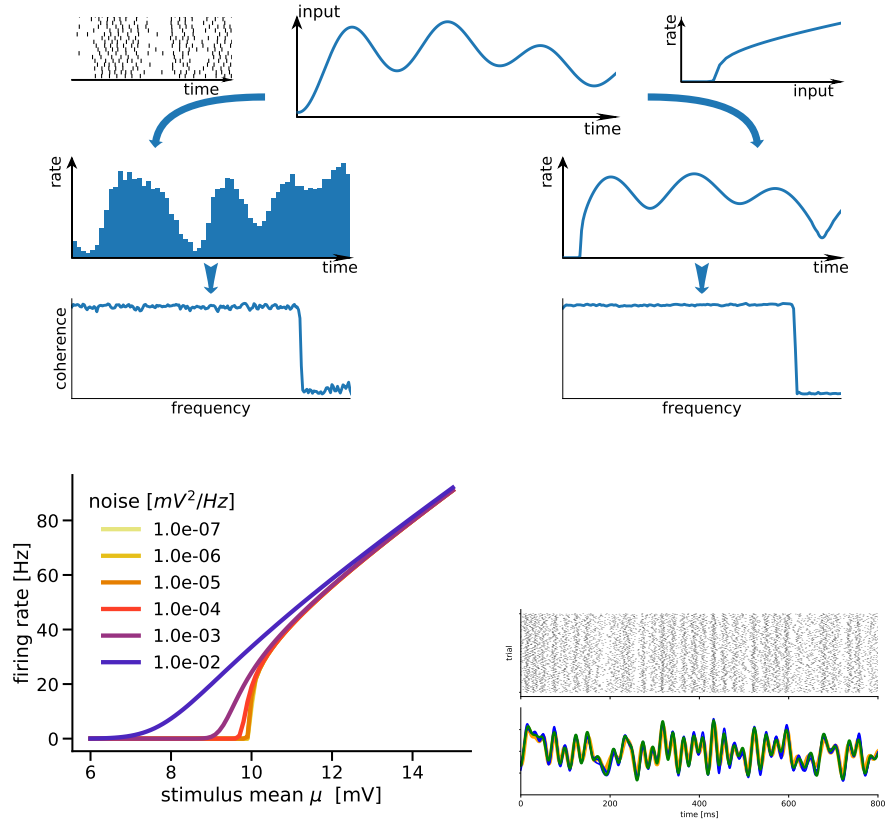


Figure 15: Two ways to arrive at coherence and coding fraction. Left: The input signal (top, center) is received by LIF-neurons. The spiking of the neurons is then binned and coherence and coding fraction are calculated between the result and the input signal. Right: Input signal (top, center) is transformed by the tuning curve (top right). The tuning curve corresponds to a function $g(V)$, which takes a voltage as input and yields a firing rate. Output is a modulated signal. We calculate coherence and coding fraction between input voltage and output firing rate. If the mean of the input is close to the threshold, as is the case here, inputs below the threshold all get projected to 0. This can be seen here at the beginning of the transformed curve. Bottom left: Tuning curves for different noise levels. x-Axis shows the stimulus in mV, the y-axis shows the corresponding firing rate. For low noise levels there is a strong non-linearity at the threshold. For increasing noise, firing rate increases particularly




Figure 16: Tuningcurve works for 10Hz but not for 200Hz.

rons match very well (figure 1.3) over a range of signal strengths, base inputs to the neurons and noise strength. This means that the LIF-neuron tuning curve gives us a very good approximation for the limit of encoded information that can be achieved by summing over independent, identical LIF-neurons with intrinsic noise. For faster signals, the coding fraction calculated through the tuning curve stays constant, as the tuning curve only deforms the signal. As shown in figure 1.3 e) and f), the coding fraction of the LIF-neuron ensemble drops with increasing frequency. Hence for high frequency signals the tuning curve ceases to be a good predictor of the encoding quality of the ensemble.

For high-frequency signals, the method does not work. The effective refractory period prohibits the instantaneous firing rate from being useful, because the neurons spike only in very short intervals around a signal peak. They are very unlikely to immediately spike again, so that the psth is focused around the input peaks, but there is little nuance.

We use the tuning curve to analyse how the signal mean and the signal amplitude change the coding fraction we would get from an infinitely large population of neurons (fig. 15, bottom two rows). We can see that the stronger noise always yields a larger coding fraction. This is expected because the tuning curve is more linear for stronger noise. It matches that we are observing the limit of an infinitely large population, which would be able to “average out” any noise. For coding fraction as a function of the mean we see zero or near zero coding fraction if we are far below the threshold. If we increase the mean at one point we can see that coding fraction starts to jump up. This happens earlier for stronger noise (= more linear tuning curve). The increase in coding fraction is much smoother if we use a larger amplitude (right figure). We also notice some sort of plateau, where increasing the mean does not lead to a larger coding fraction, before it begins rising close to 1. The plateau begins earlier




Figure 17: **A,B**: Coding signal as a function of signal mean for two different frequencies. There is little to no difference in the coding fraction. A: $\sigma = 0.1\text{mV}$. Each curve shows coding fraction as a function of the signal mean for a different noise level. The vertical line indicates the threshold. **C-E**: Coding fraction as a function of signal amplitude for different tuning curves (noise levels). Three different means, one below the threshold (9.5mV), one at the threshold (10.0mV), and one above the threshold (10.5mV).

for straighter tuning curves. For coding fraction as a function of signal amplitude we see very different results depending on the parameters. Again, we see that stronger noise leads to higher coding fraction. If we are just above or at the threshold (center and right), an increase in signal amplitude leads to a lower coding fraction. This makes sense, as more of the signal moves into the very non-linear area around the threshold. This means that for increasing amplitude an increasing fraction of the signal gets into the range of the tuning curve with 0Hz firing range, i.e. where there is no signal encoding. A very interesting effect happens if we have a mean slightly below the threshold (left): while for a strong noise we see the same effect as at or above the threshold, for weaker noise we see the opposite. This can be explained as the reverse of the effect that leads to decreasing coding fraction. Here, a larger amplitude means that the signal moves to the more linear part of the tuning curve more often. On the other hand, an increase in amplitude does not lead to worse encoding because of movement of the signal into the 0Hz part of the tuning curve – because the signal is already there, so it can't get worse. This can help explain why the coding fraction seems to saturate near 0.5: In an extreme case, the negative parts of a signal would not get encoded at all, while the positive parts would be encoded linearly.

Discussion

In this paper we have shown the effect of Suprathreshold Stochastic Resonance (SSR) in ensembles of neurons. We detailed how noise levels affect the impact of population size on the coding fraction. We looked at different frequency ranges and could show that the encoding of high-frequency signals profits particularly well from SSR. Using the tuningcurve we were able to provide a way to extrapolate the effects of SSR for very large populations. Because in general analysis of the impact of changing parameters is complex, we investigated limit cases, in particular the slow stimulus limit and the weak stimulus limit. For low-frequency signals, i.e. the slow stimulus limit, the tuningcurve also allows analyzing the impact of changing signal strength; in addition we were able to show the difference in sub-threshold SR and SSR for different noise levels. For the weak stimulus limit, where noise is relatively strong compared to the signal, we were able to provide an analytical solution for our observations.

(Hoch et al., 2003) also shows that SSR effects hold for both LIF- and HH- Neurons. However, Hoch et al. have found that optimal noise level depends "close to logarithmically" on the number of neurons in the population. They used a cutoff frequency of only 20Hz for their simulations. **[DH: Hier fehlt ein plot, der Population size und optimum noise in Verbindung setzt]**

We investigated the impact of noise on homogeneous populations of neurons. Neurons being intrinsically noisy is a phenomenon that is well investigated (Grewe et al 2017, Padmanabhan and Urban 2010). In natural systems however, neuronal populations are rarely homogeneous. Padmanabhan and Urban (2010) showed that heterogeneous populations of neurons carry more information than heterogeneous populations. Beiran et al. (2017) investigated SSR in heterogeneous populations of neurons. They made a point that heterogeneous populations are comparable to homogeneous populations where the neurons receive independent noise in addition to a deterministic signal. They make the point that in the case of weak signals, heterogeneous population can encode information better, as strong noise would overwhelm the signal. **[DH: Unterschiede herausstellen!]** Similarly, Hunsberger et al. (2014) showed that both noise and heterogeneity linearize the tuning curve of LIF neurons. In summary, while noise and heterogeneity are not completely interchangeable. In the limit cases we see similar behaviour.

(?) Sharafi et al. (2013) had already investigated SSR in a similar way. However, they only observed populations of up to three neurons and were focused on the synchronous output of cells. They took spike trains, convolved those with a gaussian and then multiplied the response of the different neurons. In our simulations we instead used the addition of spike trains to calculate the coherence between input and output. Instead of changing the noise parameter to find the optimum noise level, they changed the input signal frequency to find a resonating frequency, which was possible for suprathreshold stochastic resonance, but not for subthreshold stochastic resonance. For some combinations of parameters we also found that coding fraction does not decrease monotonically with increasing signal frequency (fig. 1.3). It is especially notable for signals that are far from the threshold (fig 1.3 E,F (red markers)). That we don't see the effect that clearly matches Sharafi et al.'s observation that in the

case of subthreshold stochastic resonance, coherence monotonically decreased with increasing frequency. Pakdaman et al. (2001) [DH: Besser verknüpfen als das Folgende (vergleichen über Größenordnungen; vergleichen mit Abbildung 5??; mehr als Sharafi zitieren Stichwort "Coherence Resonance")]

Similar research to Sharafi et al. was done by (de la Rocha et al. 2007). They investigated the output correlation of populations of two neurons and found it increases with firing rate. We found something similar in this paper, where an increase in μ increases both the firing rate of the neurons and generally also the coding fraction [DH: Verknüpfen mit output correlation](fig. 1.3). Our explanation is that coding fraction and firing rate are linked via the tuningcurve. In addition to simulations of LIF neurons de la Rocha et al. also carried out *in vitro* experiments where they confirmed their simulations.

[DH: Konkreter machen: was machen die Anderen, das mit uns zu tun hat und was genau hat das mit uns zu tun?] [DH: Vielleicht nochmal Stocks, obwohl er schon in der Einleitung vorkommt? Heterogen/homogen] [DH: Dynamische stimuli! Bei Stocks z.B. nicht, nur z.B. bei Beiran. Wir haben den Übergang.]

Examples for neuronal systems that feature noise are P-unit receptor cells of weakly electric fish (which paper?) and ...

In the case of low cutoff frequency and strong noise we were able to derive a formula that explains why in those cases coding fraction simply depends on the ratio between noise and population size, whereas generally the two variables have very different effects on the coding fraction.

1.4 Different frequency ranges

1.5 Narrow-/wideband

Narrowband stimuli

Using the $f_{cutoff} = 200\text{Hz}$ signal, we repeated the analysis for only a part of the spectrum. We did so for two "low frequency" (0–8Hz, 0–50Hz) and two "high frequency" (192–200Hz, 150–200Hz) intervals. We then compared the results to the results we get from narrowband stimuli, with power only in those frequency bands. To keep the power of the signal inside the two intervals the same as in the broadband stimulus, amplitude of the narrowband signals was less than that of the broadband signal. For the 8Hz intervals, amplitude (i.e. standard deviation) of the signal was 0.2mV, or a fifth of the amplitude of the broadband signal. Because signal power is proportional to the square of the amplitude, this was appropriate for a stimulus with a spectrum 25 times smaller. Similarly, for the 50Hz intervals we used a 0.5mV amplitude, or half of that of the broadband stimulus. As the square of the amplitude is equal to the integral over the frequency spectrum, for a signal with a quarter of the width we need to half the amplitude to have the same power in the interval defined by the narrowband signals.

Smaller frequency intervals in broadband signals

We want to know how good encoding works for different frequency intervals in the signal. When we take out a narrower frequency interval from a broadband signal, the other frequencies in the signal serve as common noise to the neurons encoding the signal. In many cases we only care about a certain frequency band in a signal of much wider bandwidth. In figure 18 A we can see that SSR has very different effects on certain frequencies inside the signal than on others. In blue we see the case of very weak noise ($10^{-6}\text{mV}^2/\text{Hz}$). Increasing the population size from 4 neurons to 2048 neurons has practically no effect. Around the average firing rate of the neurons, coherence becomes almost zero. When we keep population size at 4 neurons, but add more noise to the neurons (green, $2 \cdot 10^{-3}\text{mV}^2/\text{Hz}$), encoding of the low frequencies (up to about 50Hz) becomes worse, while encoding of the higher frequencies stays unchanged. When we increase the population size to 2048 neurons we have almost perfect encoding for frequencies up to 50Hz. Coherence is still reduced around the average firing rate of the neurons, but at a much higher level than before. For higher frequencies coherence becomes higher again. In summary, the high frequency bands inside the broadband stimulus experience a much greater increase in encoding quality than the low frequency bands, which were already encoded quite well.

Narrowband Signals vs Broadband Signals

In nature, often an external stimulus covers a frequency range that starts at high frequencies, so that only using broadband white noise signals as input appears to be insufficient to describe realistic scenarios. Therefore, we investigate the coding of narrowband signals in the ranges described earlier (0–50Hz, 150–200Hz). Comparing the results from coding of broadband and coding of narrowband signals, we see several differences.

For both low and high frequency signals, the narrowband signal can be resolved better than the broadband signal for any amount of noise and (figure 19, bottom left). That coding fractions are higher when we use narrowband signals can be explained by the fact that the additional frequencies in the broadband signal are now absent. In the broadband signal they are a form of "noise" that is common to all the input neurons. Similar to what we saw for the broadband signal, the peak of the low frequency input is still much more broad than the peak of the high frequency input. To encode low frequency signals the exact strength of the noise is not as important as it is for the high frequency signals which can be seen from the wider peaks.

1.6 Discussion

The usefulness of noise on information encoding of subthreshold signals by single neurons has been well investigated. However, the encoding of supra-threshold signals by populations of neurons has received comparatively little attention and different effects play a role for suprathreshold signals than for subthreshold signals (?).




Figure 18: Coherence for broad and narrow frequency range inputs. a) Broad spectrum. At the frequency of the firing rate (91Hz, marked by the black bar) and its first harmonic (182Hz) the coding fraction breaks down. For the weak noise level (blue), population sizes $n=4$ and $n=4096$ show indistinguishable coding fraction. In case of a small population size, coherence is higher for weak noise (blue) than for strong noise (green) in the frequency range up to about 50Hz. For higher frequencies coherence is unchanged. For the case of the larger population size and the greater noise strength there is a huge increase in the coherence for all frequencies. b) Coherence for two narrowband inputs with different frequency ranges. Low frequency range: coherence for slow parts of the signal is close to 1 for weak noise. SSR works mostly on the higher frequencies (here >40 Hz). High frequency range: At 182Hz (first harmonic of the firing frequency) there is a very sharp decrease in coding fraction, especially for the weak noise condition (blue). Increasing the noise makes the drop less clear. For weak noise (blue) there is another break down at 182-(200-182)Hz. Stronger noise seems to make this sharp drop disappear. Again, the effect of SSR is most noticeable for the higher frequencies in the interval.

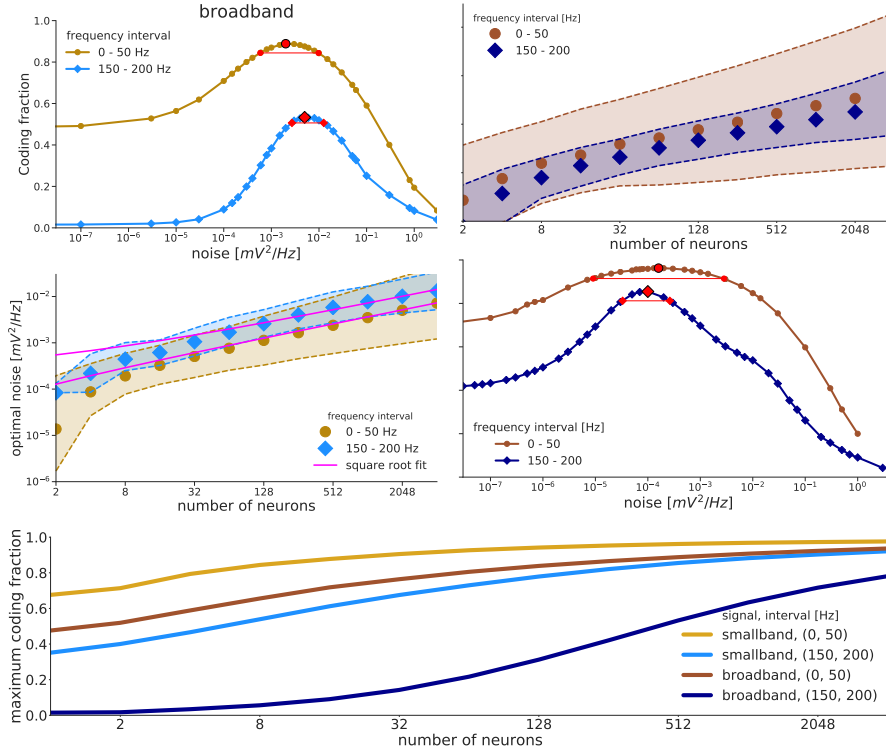


Figure 19: A: Input signal spectrum of a broadband signal. The colored area marks the frequency ranges considered here. B: Two narrowband signals (red and blue). The broadband signal from A (grey) is shown again for comparison. C and D: Best amount of noise for different number of neurons. The dashed lines show where coding fraction still is at least 95% from the maximum. The width of the peaks is much larger for the narrowband signals which encompasses the entire width of the high-frequency interval peak. Optimum noise values for a fixed number of neurons are always higher for the broadband signal than for narrowband signals. In the broadband case, the optimum amount of noise is larger for the high-frequency interval than for the low-frequency interval and vice-versa for the narrowband case. E and F: Coding fraction as a function of noise for a fixed population size ($N=512$). Red dots show the maximum, the red line where coding fraction is at least 95% of the maximum value. G: An increase in population size leads to a higher coding fraction especially for broader bands and higher frequency intervals. Coding fraction is larger for the narrowband signal than in the equivalent broadband interval for all neural population sizes considered here. The coding fraction for the low frequency intervals is always larger than for the high frequency interval. Signal mean $\mu = 15.0\text{mV}$, signal amplitude $\sigma = 1.0\text{mV}$ and $\sigma = 0.5\text{mV}$ respectively.

This paper delivers an important contribution for the understanding of suprathreshold stochastic resonance (SSR). We simulate populations of leaky integrate-and-fire neurons to answer the question how population size influences the optimal noise strength for linear encoding of suprathreshold signals. We are able to show that this optimal noise is well described as a function of the square root of population size. This relationship is independent of frequency properties of the input signal and holds true for narrowband and broadband signals.

In this paper, we show that SSR works in LIF-neurons for a variety of signals of different bandwidth and frequency intervals. We show that signal-to-noise ratio is for signals above a certain strength sufficient to describe the optimal noise strength in the population, but that the actual coding fraction depends on the absolute value of signal strength.

We contrast how well the low and high frequency parts of a broadband signal can be encoded. We take an input signal with $f_{cutoff} = 200$ Hz and analyse the coding fraction for the frequency ranges 0 to 50 Hz and 150 to 200 Hz separately. The maximum value of the coding fraction is lower for the high frequency interval compared to the low frequency interval. This means that inside broadband signals higher frequencies intervals appear more difficult to encode for each level of noise and population size. The low frequency interval has a wider peak (defined as 95% coding fraction of its coding fraction maximum value), which means around the optimal amount of noise there is a large area where coding fraction is still good. The noise optimum for the low frequency parts of the input is lower than the optimum for the high frequency interval (Fig. ??). In both cases, the optimal noise value appears to grow with the square root of population size.

In general, narrowband signals can be encoded better than broadband signals. narrowband vs broadband

Another main finding of this paper is the discovery of frequency dependence of SSR. We can see from the shape of the coherence between the signal and the output of the simulated neurons, SSR works mostly for the higher frequencies in the signal. As the lower frequency components are in many cases already encoded really well, the addition of noise helps to flatten the shape of the coherence curve. In the case of weak noise, often there are border effects which disappear with increasing strength of the noise. In addition, for weak noise there are often visible effects from the firing rate of the neurons, in so far that the encoding around those frequencies is worse than for the surrounding frequencies. Generally this effect becomes less pronounced when we add more noise to the simulation, but we found a very striking exception in the case of narrowband signals. Whereas for a firing rate of about 91Hz the coding fraction of the encoding of a signal in the 0-50Hz band is better than for the encoding of a signal in the 150-200Hz band. However, this is not the case if the neurons have a firing rate about 34Hz. We were thus able to show that the firing rate on the neurons in the simulation is of critical importance to the encoding of the signal.

2 Theory

2.1 Firing rates

Limit case of large populations

For large population sizes and strong noise, coding fraction becomes a function of their quotient

For the linear response regime of large noise, we can estimate the coding fraction. From Beiran et al. 2018 we know the coherence in linear response is given as

$$C_N(\omega) = \frac{N|\chi(\omega)|^2 S_{ss}}{S_{x_i x_i}(\omega) + (N1)|\chi(\omega)|^2 S_{ss}} \quad (7)$$

where $C_1(\omega)$ is the coherence function for a single LIF neuron. Generally, the single-neuron coherence is given by (?)

$$C_1(\omega) = \frac{r_0 \omega^2 S_{ss}(\omega)}{D} \frac{\left| \mathcal{D}_{i\omega-1}\left(\frac{\mu-v_T}{\sqrt{D}}\right) - e^{\Delta} \mathcal{D}_{i\omega-1}\left(\frac{\mu-v_R}{\sqrt{D}}\right) \right|^2}{\left| \mathcal{D}_{i\omega}\left(\frac{\mu-v_T}{\sqrt{D}}\right) \right|^2 - e^{2\Delta} \left| \mathcal{D}_{i\omega}\left(\frac{\mu-v_R}{\sqrt{D}}\right) \right|^2} \quad (8)$$

where r_0 is the firing rate of the neuron,

$$r_0 = \left(\tau_{ref} + \sqrt{\pi} \int_{\frac{\mu-v_r}{\sqrt{2D}}}^{\frac{\mu-v_t}{\sqrt{2D}}} dz e^{z^2} \operatorname{erfc}(z) \right)^{-1}$$

. In the limit of large noise (calculation in the appendix) this equation evaluates to:

$$C_1(\omega) = \sqrt{\pi} D^{-1} \frac{S_{ss}(\omega) \omega^2 / (1 + \omega^2)}{2 \sinh\left(\frac{\omega\pi}{2}\right) \Im\left(\Gamma\left(1 + \frac{i\omega}{2}\right) \Gamma\left(\frac{1}{2} - \frac{i\omega}{2}\right)\right)} \quad (9)$$

From eqs.7 and 9 it follows that in the case $D \rightarrow \infty$ the coherence, and therefore the coding fraction, of the population of LIF neurons is a function of $D^{-1}N$. We plot the approximation as a function of ω (fig. 2.1). In the limit of small frequencies the approximation matches the exact equation very well, though not for higher frequencies. We can verify this in our simulations by plotting coding fraction as a function of $\frac{D}{N}$. We see (fig. 2.1) that in the limit of large D, the curves actually lie on top of each other. This is however not the case (fig. 2.1) for stimuli with a large cutoff frequency f_c , as expected by our evaluation of the approximation as a function of the frequency.




Figure 20: Top row: Coding fraction as a function of noise. Bottom row: Coding fraction as a function of the ratio between noise strength and population size. For strong noise, coding fraction is a function of this ratio. Left: signal mean 10.5mV, signal amplitude 0.5mV, f_c 10Hz. Middle: signal mean 15.0mV, signal amplitude 0.5mV, f_c 50Hz. Right: signal mean 15.0mV, signal amplitude 1.0mV, f_c 200Hz.

High and low firing rates

One key factor that determines how well a neuron can encode a given signal is the firing rate of the neuron. Though it has been shown to be possible for neurons to encode signals with frequencies above their firing rate (?), in general higher firing rates lead to a better encoding of the signal. In our simulations the firing rate of the neurons depends on the noise added to the neurons. The effect can be seen in figure ??, which shows average firing rate as a function of the mean input for different noise strengths. The largest differences can be seen for inputs with a mean around the firing threshold (10mV) and below. While for weak noise (gray) there is an obvious non-linearity at the firing threshold, increasing noise strength linearizes the average firing rate. This illustrates subthreshold SR quite well, as the noise induces firing of the neurons for signals which would otherwise be too weak to elicit a response. For stronger mean inputs like the 15mV (second vertical black line) we use in our simulations firing rate is roughly linear with input strength and is not sensitive to changes in the noise strength. Therefore, the effect of noise on the average firing rate of the neurons plays at most a weak role in the explanation of SSR. However, the mean input strength is very important because of its effects on the average firing rate, as we will show below.² We found that in our simulations, the amplitude of the signal has a negligible influence on the firing rate.

Previously we looked at the optimum noise value for a given a population size.

²could use a plot comparing high/low directly

Now we look at optimal population sizes for a given noise value: We define a "maximum" coding fraction as the coding fraction at a population size of 4096 neurons, if the coding fraction at this population size is no more than 2% greater than the coding fraction for a population of 2048 neurons. This ensures that coding fraction has reasonably converged at this point. To see why this is important, compare the yellow line (noise strength $10^{-3} \text{ mV}^2/\text{Hz}$) in figure 3 B to the other lines in that figure. The coding fraction is still rising with increasing population size and there is no reliable way to estimate for which population size it will ultimately converge and what the maximum coding fraction will be. We are only using the noise strengths for which coding fraction has converged in the following analysis. Then, we define the "optimum" population size as the size where coding fraction is 95% of the maximum coding fraction, using linear interpolation between the different population sizes. Exponential interpolation yielded essentially the same results. We call this population size "optimal" because we assume that, for efficiency reasons, population size should be as small as possible while having very good encoding capabilities.

Strong average input (high firing rates)

First, we consider the case of strong input (15 mV) which leads to a an average firing rate of about 91Hz. As before, we see great differences for the different frequency intervals. In figure 21 A we see population sizes necessary to reach an encoding quality close to the maximum. To encode the high frequency parts of the spectrum much larger populations are required than for the low frequency parts. As expected, the size of the optimal population increases with increasing noise strength. To encode the broadband signal over its entire spectrum, a population size between the population sizes for the narrowband intervals is optimal. This can be understood because the broadband signal contains both the "easy" and the "difficult" intervals and should therefore fall in between. The same is true for the maximum coding fraction (figure 21 C): Again, the broadband signal falls in between the intervals. As we could see before (figure 19), the lower frequency interval is easier to encode than the higher frequency interval and the whole broadband signal. We also considered the relative effect of increasing population size. To quantify this, we divided the maximum coding fraction by the coding fraction of a single neuron. Figure 21 E shows that the increase in coding fraction gained by increasing population size is larger for the higher frequency interval. In contrast, a single neuron can encode the broadband signal or the low frequency interval about as well as a larger population can. That relative improvement increases for stronger noise is a consequence of the reduction in the encoding capability of a single neuron (compare figure 3 C).

For the narrowband signals we see a similar picture: Except for very weak noise, optimal population size to encode the high frequency signal is larger than that for the low frequency signal (figure 21 B). Optimal population size for encoding of the low frequency signal now starts at a much higher level than before. Maximum coding fraction again is larger for the low frequency signal (figure 21 D). For the high frequency signal, coding fraction is at a much higher level than it is for the high frequency interval in the broadband signal for all noise strengths. With increasing noise




Figure 21: An overview of the effect of different noise strengths on maximum coding fraction and population size. We only considered noise values where the difference in coding fraction between population sizes $n=4096$ and $n=2048$ is less than 2%. Average firing rate of the neurons was about 91Hz. Input strengths were chosen so that the power of the signal in the corresponding bands is the same (1.0mV for the broadband and 0.5mV for the narrowband signals). Top to bottom: a) Minimum population size needed to have coding fraction be at least 95% of the maximum as a function of noise. Optimal population size grows with increasing noise. Optimal population size is larger for the narrowband signals (dots) than for the broadband signal (crosses). For weak noise and narrowband signals, there is little difference in the optimal population size for the high frequency interval (brown) and the low frequency interval (blue). As noise becomes stronger, optimal population size is larger for the higher interval. For the broadband signal, optimum population size is always largest for the higher frequency interval, then for the broadband signal and finally the low frequency narrowband signal. That minimum population size is larger for the narrowband signals than for the broadband signal can be explained by the fact that the maximum coding fraction (b) is higher for the narrowband signals. b) Maximum coding fraction is higher for lower intervals and narrower bands. In the case of the narrowband signal and the lower interval, maximum coding fraction is close to 1 for all noise strengths. For weak noise, the low interval in the broadband signal and the higher narrowband signal have very similar maximum coding fractions. With increasing noise, the maximum coding fraction rises much faster for the high frequency narrowband signal. The broadband signal and the high frequency interval inside that signal have very low maximum coding fraction for weak noise. Increasing the strength of the noise at some point coding fraction begins to increase rapidly. c) Frequency band, not signal bandwidth appears to be the main factor in the relative improvement of coding fraction with increasing population size. For the slow narrowband signal the improvement is mostly not because the maximum becomes higher, as it changes very little. Instead the reason is that a single neuron has a diminished ability to encode the signal with increasing noise levels (see fig. 3 B and C). For the whole broadband signal and the low frequency interval within, for weak noise there is almost no improvement in coding fraction through increasing population size. Interestingly, with stronger noise the relative improvement for the low frequency interval is the same for both the narrowband and the broadband signal, even though maximum coding fraction is very different, at least for intermediate noise strength (10^{-4} to 10^{-3} mV²/Hz).




Figure 22: Coherence curves for broad and narrow frequency range inputs. The average firing rate of the cells is marked with a black vertical line. a) Broad spectrum. Coherence for low frequencies is much lower than for the case of higher mean input (fig. 18 A). For the weak noise level (blue), population sizes $n=4$ and $n=4096$ show indistinguishable coding fraction. In case of a small population size, coherence is higher for stronger noise (green), contrary to what we have seen before. This can not be explained by an increase in firing rate, as the difference in firing rate is only about 1%. b) Narrow band inputs for two frequency ranges. Low frequency range: coherence for slow parts of the signal is similar to those in the broadband signal. High frequency range: In contrast to the case of higher mean input (fig. 18) the high-frequency signal is encoded better than the low-frequency signal. Even for comparatively weak noise (blue), an increase in population size offers better encoding.

strength the coding fraction increases rapidly and almost reaches the level achieved for the low frequency signal for the population sizes considered here. The relative increase in coding fraction (figure 21 F) is similar to what we saw for the broadband signal. For high frequencies the increase in coding fraction from a single neuron to a population of neurons becomes apparent even at a relatively low noise strength. Whereas for the low frequency signal the relative improvement starts at much higher noise levels and is mostly explained by the decrease in encoding capabilities of a single neuron.

For both narrowband and broadband signals, the results here qualitatively do not depend on the input amplitude. See *appendix* for more information.

Weak average input (low firing rates)

We also consider the case of a weaker mean input (10.5mV) which leads to lower average firing rates (about 34Hz). Results can be seen in figure ?? . For the broadband signal, in general results are similar to the results in the strong average input case. We again see (fig. ?? A) that optimum population size is larger for the high frequency interval than for the low frequency interval, with the entire broadband signal somewhere in between. Now the optimal population sizes are much closer than in the

case of high average firing rates. The values being closer together is also true for the maximum coding fraction (fig. ?? C). The value for the low frequency interval is now much lower for very weak noise than for the high mean input. The curves for both intervals and the broadband signal now look very similar to each other. Maximum coding fraction for the broadband signal and the high frequency interval are almost equal for noise strengths greater than $10^{-3} \text{mV}^2/\text{Hz}$. Relative improvement (fig. ?? E) is very similar for all intervals. For increasing noise strength relative improvement starts to show the pattern we have seen before, with improvement being greatest for the high frequency interval, followed by the broadband signal and then the low frequency interval.

For the narrowband signals we see some striking differences to what we saw for the high mean input. Optimal population sizes are now very different for the low frequency signal and the high frequency signal for all noise strengths (fig. ?? B). The high frequency interval signal again needs a larger population for encoding being close to optimal. The optimal population size for the low frequency signal is now very similar to the optimal population size for the low frequency interval in the broadband signal. This contrasts with the high input case (fig. 21 A & B). For weak noise the optimal population size was much larger for the narrowband signal than for the interval in the broadband signal.

A striking change happens for the maximum coding fraction (fig. ?? D). As opposed to every case we have looked at before, now the higher frequency signal appears to be more easily encoded than the low frequency signal. This can be explained by looking at the coherence curves of the two signals (fig. 22 B). As the firing rate of the neurons is inside the frequency range of the low frequency signal, encoding of those frequencies is suppressed. This is not the case for the high frequency signal, where the coherence behaves similarly to what we have seen for the strong input case (fig. 18 B). Maximum coding fraction for the low frequency signal is very similar to the maximum coding fraction for the low frequency interval of the broadband signal. This is also different to what we have seen before, as for every other case coding fraction can get much higher for the narrowband signals.

Relative improvement is higher for the high frequency signal than for the low frequency signal (fig. ?? F). For the high frequency signal, the improvement is greater than for the equivalent interval in the broadband signal for weak noise. Relative improvement is also greater than it was for high average input (fig. 21 E & F). Here, even for weak noise increasing the population size has a large effect. In other words, to encode the high frequency signal in the case of a low average firing rate, population size is very critical to the quality of the encoding. For the low frequency signal there is no large difference in relative improvement to the other cases.

2.2 Refractory period

Refractory Periods

We analyzed the effect of non-zero refractory periods on the previous results. We added a 1ms or a 5ms refractory period to each of the LIF-neurons. Then, we re-

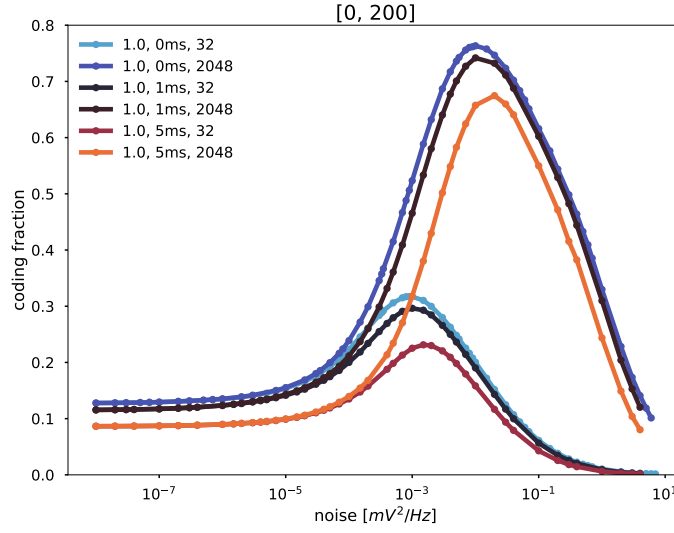


Figure 23: Repeating the simulations adding a refractory period to the LIF-neurons shows no qualitative changes in the SSR behaviour of the neurons. Coding fraction is lower the longer the refractory period. The SSR peak moves to stronger noise; cells with larger refractory periods need stronger noise to work optimally.

peated the same simulations as before. Results are summarized in figure 23. Results change very little for a refractory period of 1ms, especially for large noise values. For a refractory period of 5ms resulting coding fraction is lower for almost all noise values. Paradoxically, for high frequencies in smallband signals and very small noise, coding fraction actually is larger for 5ms refractory period than for 1ms. In spite of this, coding fraction is still largest for the LIF-ensembles without refractory period.

We also find all other results replicated even with refractory periods of 1ms or 5ms: Figure (??) shows that the optimal noise stills grows with \sqrt{N} for both the 1ms and the 5ms refractory period. We see an increase in the value of the optimum noise with an increase of the refractory period. The achievable coding fraction is lower for the neurons with refractory periods, especially at the maximum. In the limit of large noise, the neurons with 1ms refractory period and the ones with no refractory period also result in similar coding fractions, over a wide range of population sizes. However, this is not true for the neurons with 5ms refractory period.

3 Electric fish

3.1 Introduction

3.2 Methods

Electrophysiology

We recorded electrophysiological data from X cells from Y different fish.

Surgery. Twenty-two *E. virescens* (10 to 21 cm) were used for single-unit recordings. Recordings of electroreceptors were made from the anterior part of the lateral line nerve. Fish were initially anesthetized with 150 mg/l MS-222 (PharmaQ, Fordingbridge, UK) until gill movements ceased and were then respiration with a constant flow of water through a mouth tube, containing 120 mg/l MS-222 during the surgery to sustain anesthesia. The lateral line nerve was exposed dorsal to the operculum. Fish were fixed in the setup with a plastic rod glued to the exposed skull bone. The wounds were locally anesthetized with Lidocaine hydrochloride 2% (bela-pharm, Vechta, Germany) before the nerve was exposed. Local anesthesia was renewed every 2 h by careful application of Lidocaine to the skin surrounding the wound. After surgery, fish were immobilized with 0.05 ml 5 mg/ml tubocurarine (Sigma-Aldrich, Steinheim, Germany) injected into the trunk muscles. Since tubocurarine suppresses all muscular activity, it also suppresses the activity of the electrocytes of the electric organ and thus strongly reduces the EOD of the fish. We therefore mimicked the EOD by a sinusoidal signal provided by a sine-wave generator (Hameg HMF 2525; Hameg Instruments, Mainhausen, Germany) via silver electrodes in the mouth tube and at the tail. The amplitude and frequency of the artificial field were adjusted to the fish's own field as measured before surgery. After surgery, fish were transferred into the recording tank of the setup filled with water from the fish's housing tank not containing MS-222. Respiration was continued without anesthesia. The animals were submerged into the water so that the exposed nerve was just above the water surface. Electroreceptors located on the parts above water surface did not respond to the stimulus and were excluded from analysis. Water temperature was kept at 26°C.³

Recording. Action potentials from electroreceptor afferents were recorded intracellularly with sharp borosilicate microelectrodes (GB150F-8P; Science Products, Hofheim, Germany), pulled to a resistance between 20 and 100 M and filled with a 1 M KCl solution. Electrodes were positioned by microdrives (Luigs-Neumann, Ratingen, Germany). As a reference, glass microelectrodes were used. They were placed in the tissue surrounding the nerve, adjusted to the isopotential line of the recording electrode. The potential between the micropipette and the reference electrode was amplified (SEC-05X; npi electronic) and lowpass filtered at 10 kHz. Signals were digitized by a data acquisition board (PCI-6229; National Instruments) at a sampling rate of 20 kHz. Spikes were detected and identified online based on the peak-detection algorithm proposed by Todd and Andrews (1999). The EOD of the fish was measured between the head and tail via two carbon rod electrodes (11 cm long, 8-mm diam-

³From Stöckl et al. 2014

eter). The potential at the skin of the fish was recorded by a pair of silver wires, spaced 1 cm apart, which were placed orthogonal to the side of the fish at two-thirds body length. The residual EOD potentials were recorded and monitored with a pair of silver wire electrodes placed in a piece of tube that was put over the tip of the tail. These EOD voltages were amplified by a factor of 1,000 and band-pass filtered between 3 Hz and 1.5 kHz (DPA-2FXM; npi electronics). Stimuli were attenuated (ATN-01M; npi electronics), isolated from ground (ISO-02V; npi electronics), and delivered by two carbon rod electrodes (30-cm length, 8-mm diameter) placed on either side of the fish parallel to its longitudinal axis. Stimuli were calibrated to evoke defined AM measured close to the fish. Spike and EOD detection, stimulus generation and attenuation, as well as preanalysis of the data were performed online during the experiment within the RELACS software version 0.9.7 using the efish plugin-set (by J. Benda: <http://www.relacs.net>).⁴

Stimulation. White noise stimuli with a cutoff frequency of 300Hz defined an AM of the fish's signal. The stimulus was combined with the fish's own EOD in a way that the desired AM could be measured near the fish. Amplitude of the AM was 10% (?) of the amplitude of the EOD. Stimulus duration was between 2s and 10s, with a time resolution of X.

3.3 How to determine noisiness

Determining noise in real world

While in simulations we can control the noise parameter directly, we cannot do so in electrophysiological experiments. Therefore, we need a way to quantify "noisiness". One such way is by using the activation curve of the neuron, fitting a function and extracting the parameters from this function. Stocks (2000) uses one such function to simulate groups of noisy spiking neurons:

$$\frac{1}{2} \operatorname{erfc} \left(\frac{\theta - x}{\sqrt{2}\sigma} \right) \quad (10)$$

where σ is the parameter quantifying the noise (figure ??). A neuron with a σ of 0 would be a perfect thresholding mechanism. Firing probability for all inputs below the threshold is 0, and firing probability for all inputs above is 1. Large values mean a flat activation curve. Neurons with such an activation curve will likely fire even for some signals below the firing threshold, while it will sometimes not fire for inputs above the firing threshold. Its firing behaviour is influenced less by the signal, which indicates noisiness. We also tried different other methods of quantifying noise commonly used (citations), but none of them worked as well as the errorfunction fit (fig. ?? b-d)).

⁴From Stöckl et al. 2014

Methodology

The signal values were binned in 50 bins. The result is a discrete Gaussian distribution around 0mV, the mean of the signal, as is expected from the way the signal was created. We have to account for the delay between the moment we play the signal and when it gets processed in the cell, which can for example depend on the position of the cell on the skin. We calculate the cross correlation between the signal and the discrete output spikes. The position of the peak of the correlation is the time shift for which the signal influences the result of the output the most. Then, for every spike, we note the value of the signal at the time of the spike minus the time shift. The result is a histogram, where each signal value bin has an associated number of spikes. This histogram is then normalized by the distribution of the signal. The result is another histogram, whose values are firing frequencies for each signal value. Because those frequencies are just firing probabilities multiplied with equal time steps, we can fit a Gaussian error function to those probabilities.

Simulation

To confirm that the σ parameter estimated from the fit is indeed a good measure for the noisiness, we validated it against D , the noise parameter from the simulations. We find that there is a strictly monotonous relationship between the two for different sets of simulation parameters. Other parameters often used to determine noisiness (citations) such as the variance of the spike PSTH, the coefficient of variation (CV) of the interspike interval are not as useful. In figure 27 we see why. The variance of the psth is not always monotonous in D and is very flat for low values of D .

We tried several different bin sizes (30 to 300 bins) and spike widths. There was little difference between the different parameters (see appendix).

Electrophysiology

We can see from figure 28 that the fits look very close to the data. For very weak and very strong inputs, the firing rates themselves become noisy, because the signal only assumes those values rarely. This is especially noticeable for strong inputs, as there are more spikes there, and therefore large fluctuations, while there is very little spiking anyway for weak inputs.

When we group neurons by their noise and plot coding fraction as a function of population size for averages of the groups, we see results similar to what we see for simulations. Noisier cells have a lower coding fraction for small populations. For increasing population size, coding fraction increases for all groups, but the increase is much larger for noisy cells. For large population sizes the noisy cells show a better linear encoding of the signal than the more regular cells.




Figure 24: a) The parameter σ as a function of the noise parameter D in LIF-simulations. There is a strictly monotonous relationship between the two, which allows us to use σ as a substitute for D in the analysis of electrophysiological experiments. b-d) different other parameters commonly used to quantify noise. None of these functions is strictly monotonous and therefore none is useful as a substitute for D . b) Peri-stimulus time histogram (PSTH) of the spikes with a bin width of 1ms, normalized by c) PSTH of the spikes with a bin width of 5ms. d) coefficient of variation (cv) of the interspike-intervals.




Figure 25: Histogram of spike count distribution (firing rate) and errorfunction fits. 50 bins represent different values of the Gaussian distributed input signal [maybe histogram in background again]. The value of each of those bins is the number of spikes during the times the signal was in that bin. Each of the values was normalized by the signal distribution. To account for delay, we first calculated the cross-correlation of signal and spike train and took its peak as the delay. The lines show fits according to equation (11). Left and right plots show two different cells, one with a relatively narrow distribution and one with a distribution that is more broad, as indicated by the parameter σ . Different amounts of bins (30 and 100) showed no difference in resulting parameters.




Figure 26: Left: Coding fraction as a function of population size for all recorded neurons. Color are by σ from the fit of the function in equation 11, so that there are roughly an equal number of neurons in each category. Red: $\sigma = 0$ to 0.5, pink: 0.5 to 1.0, purple: 1.0 to 1.5, blue: 1.5 and above. Thick colored lines are average of the neurons in each group. For a population size of 1, coding fraction decreases on average with increasing σ . As population sizes increase, coding fraction for weak noise neurons quickly stops increasing. Strong noise neurons show better coding performance for larger population sizes (about 8 to 32 neurons). Right [missing]: Increase in coding as a function of sigma. y-axis shows the difference in coding fraction between $N=1$ and $N=32$,

Determining the strength of noise in a real world example

While in simulations we can control the noise parameter directly, we cannot do so in electrophysiological experiments. Therefore, we need a way to quantify the intrinsic noise of the cell from the output of the measured cells. Common measures to quantify noisiness of neuronal spike trains are not directly correlated with intrinsic noise strength (figure 27). An example for such a measure is the coefficient of variation (cv) of the interspike interval (ISI) (White et al., 2000; Goldberg et al., 1984; Nowak et al., 1997). The ISI is the time between each consecutive pair of spikes. The coefficient of variation is then defined as the standard deviation of the ISI divided by the mean ISI. Even though it is frequently used, we find that the cv as a function of the intrinsic noise in our LIF-simulations is not necessarily monotonously related. In addition, for different membrane constants, which determine how quickly a neuron reacts to inputs, the same intrinsic noise results in widely different cv-values. Refractory periods also have an influence on the cv. [DH: their/holt1996 cv2 looks interesting.] Another measure which has been used before is the standard variation of the peri-stimulus spike histogram (Mainen and Sejnowski, 1995) [DH: can't find any paper which did something like we did here, even though Schreiber et al. 2003 A new correlation-based measure of spike timing reliability - claim it's frequently done with psth]. This approach also does not work well, as it also depends on the membrane constant and to a lesser extend the refractory period.

The approach used here uses the activation curve of the neuron, fitting a function to it and extracting the parameters from the fitted function. It is assumed that the neurons show Gaussian noise. The mean of the distribution is the activation threshold and the width of the Gaussian is a measure for noise. The probability of spiking as a function of the input is then the integral over the Gaussian, i.e. an error function. Stocks (2000) uses one such function to simulate groups of noisy spiking neurons:

$$\frac{1}{2} \operatorname{erfc} \left(\frac{\theta - x}{\sqrt{2}\sigma} \right) \quad (11)$$

where σ is the parameter quantifying the noise (figure ?) A neuron with a σ of 0 would be a perfect thresholding mechanism. Firing probability for all inputs below the threshold is 0, and firing probability for all inputs above is 1. If σ is greater than 0, a neuron with such an activation curve will fire even for some signals below the firing threshold, while it will sometimes not fire for inputs above the firing threshold. For large values of σ the activation curve becomes flatter, meaning the probability for inputs below the threshold eliciting a spike is large and the probability that an input above the threshold does not lead to firing is also large. The firing behaviour of such a cell is influenced less by the signal, which indicates noisiness. However, for strong noise ($> 10^{-2} \frac{mV^2}{Hz}$), results are not monotonous anymore. This happens at a point where σ becomes large. Therefore, we excluded all values of the unit-less σ larger than two from the following analyses.

Methodology

The signal was binned according to its amplitude. The result is a discrete Gaussian distribution around 0mV, the mean of the signal, as is expected from the way the signal was created. After accounting for time delays in signal processing, we make a histogram which contains the distribution of spikes according to signal amplitude. This histogram is then normalized by the distribution of the signal. The result is another histogram, where values are firing frequencies for each signal value. Because those frequencies are just firing probabilities multiplied with equal time steps, we can fit a Gaussian error function to those probabilities.

Simulation

To confirm that the σ parameter estimated from the fit is indeed a good measure for the noisiness, we validated it against D, the noise parameter from the simulations. We find that there is a strictly monotonous relationship between the two for different sets of simulation parameters.

We tried several different bin sizes (30 to 300 bins) and spike widths. There was little difference between the different parameters (see appendix).

Electrophysiology

We find that the fits match the experimental data very well (figure 28). For very weak and very strong inputs, the firing rates themselves become noisy, because the signal only assumes those values rarely. This is especially noticeable for strong inputs, as there are more spikes there, and therefore large fluctuations, while there is very little spiking anyway for weak inputs.

When we group neurons by their noise and plot coding fraction as a function of population size for averages of the groups, we see results similar to what we see for simulations (figure 29 a)): Noisier cells (larger σ , purple) have a lower coding fraction for small populations. However, coding fraction mostly stops increasing with population sizes once a population size of about 16 is reached. The increase is much larger for noisy cells (orange). The averages of the coding fraction for the noisy cells does not increase above the coding fraction of the less noisy cells for the population sizes investigated here (N=128). In contrast to the more regular cells, coding fraction is still improving for the noisy cells, so it is plausible that at a certain population size the noisy cells can outperform the less noisy cells. Indeed, if results are not averaged and single cells are considered, we find that for large population sizes the noisy cells show a better linear encoding of the signal than the more regular cells (figure 29 b), red).




Figure 27: a) The parameter σ as a function of the noise parameter D in LIF-simulations. There is a strictly monotonous relationship between the two, which allows us to use σ as a substitute for D in the analysis of electrophysiological experiments. b-e) Left to right: σ , CV and standard deviation of the psth with two different kernel widths as a function of D for different membrane constants (4ms, 10ms and 16ms). The membrane constant τ determines how quickly the voltage of a LIF-neuron changes, with lower constants meaning faster changes. Only σ does not change its values with different τ . The CV (c) is not even monotonous in the case of a timeconstant of 4ms, ruling out any potential usefulness. f-i) Left to right: σ , CV and standard deviation of the psth with two different kernel widths as a function of D for different refractory periods (0ms, 1ms and 5ms). Only σ does not change with different refractory periods.

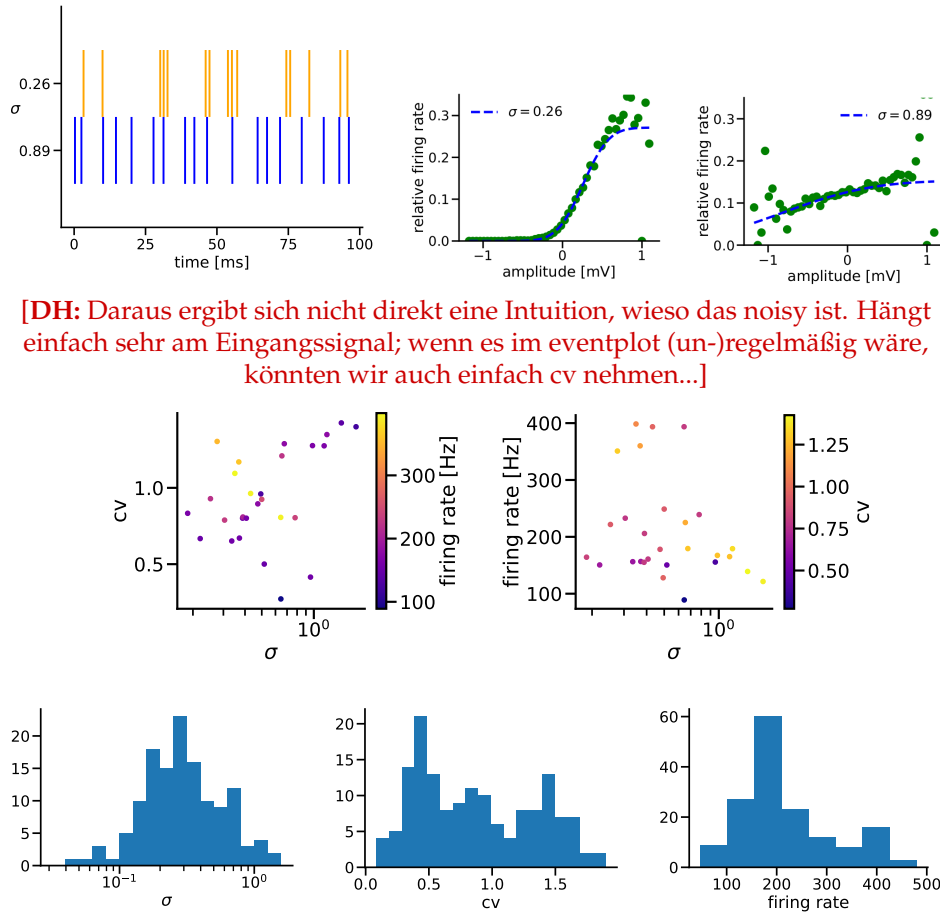


Figure 28: Histogram of spike count distribution (firing rate) and errorfunction fits. 50 bins represent different values of the Gaussian distributed input signal [maybe histogram in background again]. The value of each of those bins is the number of spikes during the times the signal was in that bin. Each of the values was normalized by the signal distribution. For very weak and very strong inputs, the firing rates themselves become noisy, because the signal only assumes those values rarely. To account for delay, we first calculated the cross-correlation of signal and spike train and took its peak as the delay. The lines show fits according to equation (11). Left and right plots show two different cells, one with a relatively narrow distribution (left) and one with a distribution that is more broad (right), as indicated by the parameter σ . An increase of σ is equivalent to a broader distribution. Cells with broader distributions are assumed to be noisier, as their thresholding is less sharp than those with narrow distributions. Different amounts of bins (30 and 100) showed no difference in resulting parameters.

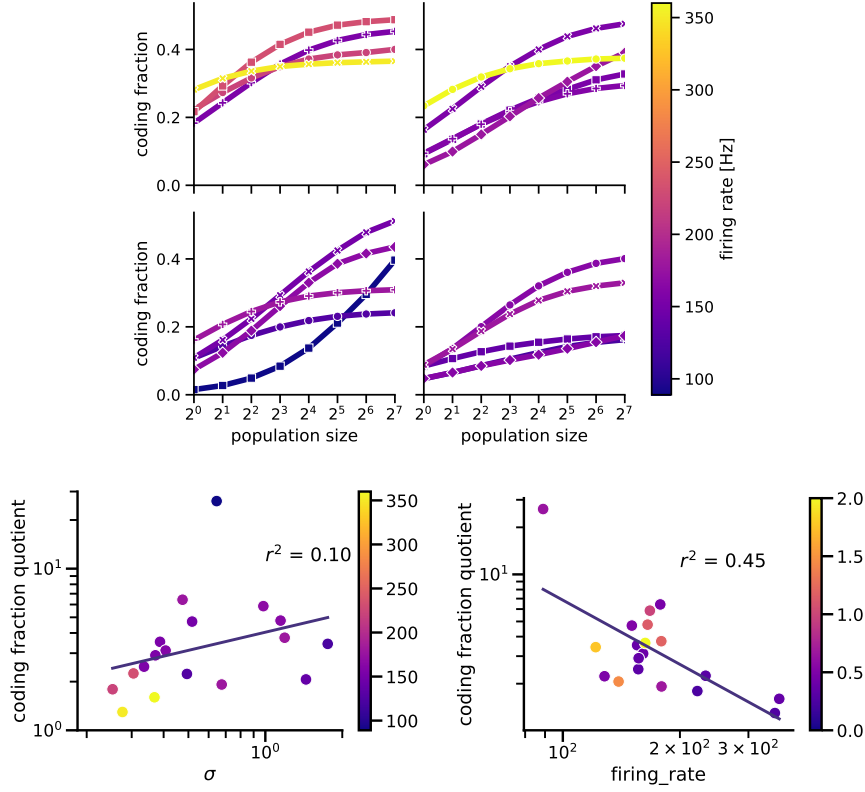


Figure 29: Left: Coding fraction as a function of population size for all recorded neurons. Cells are grouped by σ from the fit of the function in equation 11. Lines are averages over three cells each, with the shading showing the standard deviation. For stronger noise, coding fraction is far smaller for a single neuron. With increasing population size, coding fraction increases much faster for the noisy cells than for the less noisy cells. Right: Examples for the two cells with lowest, intermediate and highest σ . For a population size of $N=1$, the cell with the largest σ (brown) has the lowest coding fraction out of all the cells here. The coding fraction of that cell increases hugely with population size. At a population of $N=128$, coding fraction is second highest among the pictured cells.

3.4 Results

Electric fish as a real world model system

To put the results from our simulations into a real world context, we chose the weakly electric fish *Apteronotus leptorhynchus* as a model system. *A. leptorhynchus* uses an electric organ to produce electric fields which it uses for orientation, prey detection and communication. Distributed over the skin of *A. leptorhynchus* are electroreceptors which produce action potentials in response to electric signals.

These receptor cells ("p-units") are analogous to the simulated neurons we used in our simulations because they do not receive any input other than the signal they are encoding. Individual cells fire independently of each other and there is no feedback.

Results

Figure 3.4 A,B and C show three examples for coherence from intracellular measurements in *A. leptorhynchus*. Each cell was exposed to up to 128 repetitions of the same signal. The response was then averaged over different numbers of trials to simulate different population sizes of homogeneous cells. We can see that an increase in population size leads to higher coherence. Similar to what we saw in the simulations, around the average firing rate of the cell (marked by the red vertical lines), coherence decreases sharply. We then aggregated the results for 31 different cells (50 experiments total, as some cells were presented with the stimulus more than once). Figure ?? D shows that the increase is largest inside the high frequency intervals. As we could see in our simulations (figures 21 C and ?? C), the ratio of coding fraction in a large population to the coding fraction in a single cell is larger for higher frequencies.

3.5 Discussion

We also confirmed that the results from the theory part of the paper play a role in a real world example. Inside the brain of the weakly electric fish *Apteronotus leptorhynchus* pyramidal cells in different areas are responsible for encoding different frequencies. In each of those areas, cells integrate over different numbers of the same receptor cells. Artificial populations consisting of different trials of the same receptor cell show what we have seen in our simulations: Larger populations help especially with the encoding of high frequency signals. These results are in line with what is known about the pyramidal cells of *A. leptorhynchus*: The cells which encode high frequency signals best are the cells which integrate over the largest number of neurons.

4 Discussion: Combining experiment and simulation

5 Literature

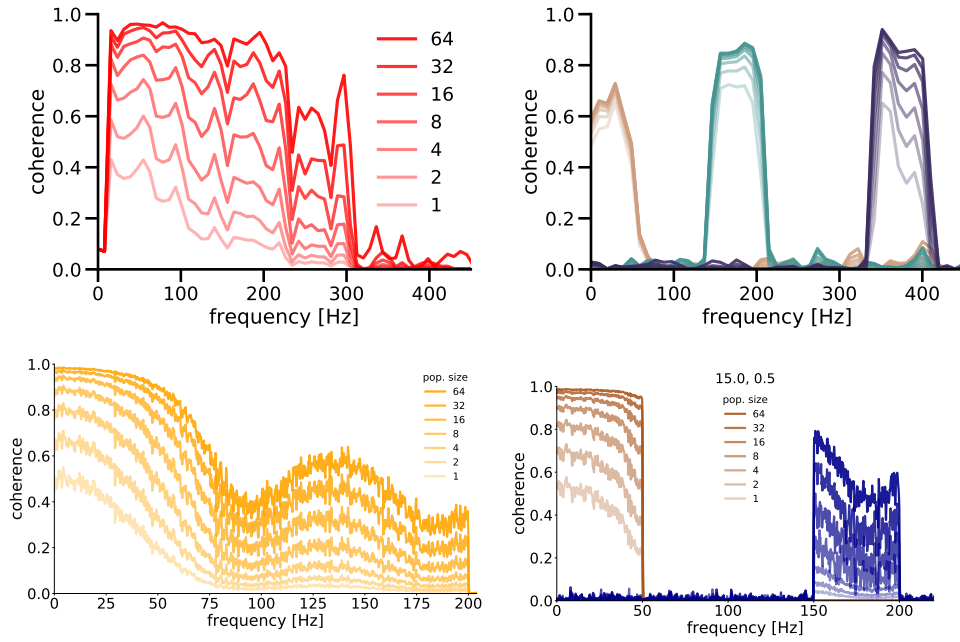
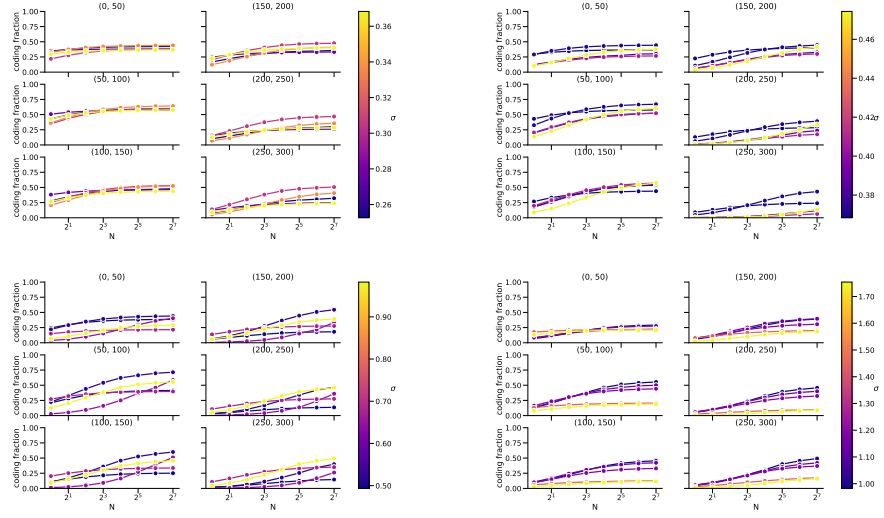



Figure 30: A,B,C: examples of coherence in the p-Units of *A. leptorhynchus*. Each plot shows the coherence of the response of a single cell to a stimulus for different numbers of trials. Like in the simulations, increased population sizes lead to a higher coherence. D: Encoding of higher frequency intervals profits more from an increase in population size than encoding of lower frequency intervals. The ratio of the coding fraction for the largest number of trials divided by the coding fraction for a single trial for each of six different frequency intervals. Shown here are the data for all 50 experiments (31 different cells). The orange line signifies the median value for all cells. The box extends over the 2nd and 3rd quartile.


broad




broad



narrow





[DH: needs figure 3.6 from yue and equivalent]

Figure 31: This is about frequency and how it determines Δf_c . In other paper I have used $quot_c f$.

References

- Jheeyae Ahn, Lauren J Kreeger, Susan T Lubejko, Daniel A Butts, and Katrina M MacLeod. Heterogeneity of intrinsic biophysical properties among cochlear nucleus neurons improves the population coding of temporal information. *Journal of neurophysiology*, 111(11):2320–2331, 2014.
- Manuel Beiran, Alexandra Kruscha, Jan Benda, and Benjamin Lindner. Coding of time-dependent stimuli in homogeneous and heterogeneous neural populations. *Journal of computational neuroscience*, 44(2):189–202, 2018.
- Roberto Benzi, Alfonso Sutera, and Angelo Vulpiani. The mechanism of stochastic resonance. *Journal of Physics A: mathematical and general*, 14(11):L453, 1981.
- Fabrizio Gabbiani. Coding of time-varying signals in spike trains of linear and half-wave rectifying neurons. *Network: Computation in Neural Systems*, 7(1):61–85, 1996.
- Fabrizio Gabbiani and Christof Koch. Coding of time-varying signals in spike trains of integrate-and-fire neurons with random threshold. *Neural Computation*, 8(1):44–66, 1996.
- Fabrizio Gabbiani, Walter Metzner, Ralf Wessel, Christof Koch, and Others. From stimulus encoding to feature extraction in weakly electric fish. *Nature*, 384(6609):564–567, 1996.
- Luca Gammitoni, Peter Hänggi, Peter Jung, and Fabio Marchesoni. Stochastic resonance. *Rev. Mod. Phys.*, 70(1):223–287, 1998. doi: 10.1103/RevModPhys.70.223. URL <http://link.aps.org/doi/10.1103/RevModPhys.70.223>.
- JM Goldberg, CHARLES E Smith, and C Fernandez. Relation between discharge regularity and responses to externally applied galvanic currents in vestibular nerve afferents of the squirrel monkey. *Journal of neurophysiology*, 51(6):1236–1256, 1984.
- Thomas Hoch, Gregor Wenning, and Klaus Obermayer. Optimal noise-aided signal transmission through populations of neurons. *Physical Review E*, 68(1):011911, 2003.
- Rüdiger Krahe, Gabriel Kreiman, Fabrizio Gabbiani, Christof Koch, and Walter Metzner. Stimulus encoding and feature extraction by multiple sensory neurons. *Journal of Neuroscience*, 22(6):2374–2382, 2002.
- Benjamin Lindner. Mechanisms of information filtering in neural systems. *IEEE Transactions on Molecular, Biological and Multi-Scale Communications*, 2016.
- Zachary F Mainen and Terrence J Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506, 1995.
- Lionel G Nowak, Maria V Sanchez-Vives, and David A McCormick. Influence of low and high frequency inputs on spike timing in visual cortical neurons. *Cerebral cortex* (New York, NY: 1991), 7(6):487–501, 1997.

- Susanne Schreiber, Christian K Machens, Andreas V M Herz, and Simon B Laughlin. Energy-efficient coding with discrete stochastic events. *Neural Computation*, 14(6): 1323–1346, 2002.
- Einat Shapira, Rémy Pujol, Michael Plaksin, and Eitan Kimmel. Sound-induced motility of outer hair cells explained by stochastic resonance in nanometric sensors in the lateral wall. *Physics in Medicine*, 2016.
- T Shimokawa, A Rogel, K Pakdaman, and S Sato. Stochastic resonance and spike-timing precision in an ensemble of leaky integrate and fire neuron models. *Physical Review E*, 59(3):3461, 1999.
- N G Stocks. Suprathreshold stochastic resonance in multilevel threshold systems. *Physical Review Letters*, 84(11):2310, 2000a.
- NG Stocks. Suprathreshold stochastic resonance in multilevel threshold systems. *Physical Review Letters*, 84(11):2310, 2000b.
- NG Stocks. Information transmission in parallel threshold arrays: Suprathreshold stochastic resonance. *Physical Review E*, 63(4):041114, 2001.
- NG Stocks and Riccardo Mannella. Generic noise-enhanced coding in neuronal arrays. *Physical Review E*, 64(3):030902, 2001.
- John A White, Jay T Rubinstein, and Alan R Kay. Channel noise in neurons. *Trends in Neurosciences*, 23(3):131–137, 2000.