Chapter 7. Appendix A: Methods for signal-processing, analysis of spike trains, and local field potentials (LFPs)

7.1 Signal acquisition

All extracellular recordings were acquired continuously using Neuralynx Hardware (Neuralynx Inc, Tucson, AZ). We used two generations of systems: An Analog Cheetah system (25 kHz sample rate) with 32 channels, and a Digital Cheetah system (32 kHz sample rate) with 64 channels. In both systems, signals were first pre-amplified as close as possible to the source (pre-amplifiers were placed on the head of the patient). After pre-amplification, signals were fed into the acquisition system, which was located in the room of the patient (but several meters away). The acquisition system amplified the analog signals (with gain in the range of 2000–50000) and fed them into an A/D converter (analog system) or directly fed them to the A/D converter (digital system, no analog amplification). Spike times were determined offline after the recording (see spike detection and sorting chapter for details). All parts of the system that were in contact with the patient were powered by DC batteries to avoid safety problems as well as to reduce line-noise interference. The interface between the acquisition system and the recording PC (acquisition card) was optical. We used the Cheetah software to acquire all data (Neuralynx Inc, Tucson, AZ).

Each macroelectrode contained 8 microwires (see Introduction for details). One of these wires was used as ground. The choice of ground wire (based on background noise levels and impedance) occurred on the first day of recording for every patient. Special care was taken to identify a ground wire that had very low levels of electrical activity, as otherwise the activity on the groundwire would be recorded on all other wires as a signal. All our recordings were locally grounded. Thus, the measured voltage (the output of the amplifier) is the difference between the two inputs to the amplifier (differential amplification, relative): the measuring wire and the ground wire. Thus, the signal S_i (output of the amplifier) represents $S_i = V_i - V_G$, where V_i is the voltage on each microwire measured relative to a distant ground (i.e., the skull). All microwires of the same macroelectrode are located very closely together (spatially, typically < 1 mm). This kind of differential recording thus allows the measurement of very local electrical activity. All activity that is common to both wires (such as global line noise, long-range oscillations) is cancelled from the signal due to the subtraction. This has implications for the LFP signal recorded from these electrodes: It is very different from a traditional iEEG signal (see the LFP chapter for details).

Signals were acquired with the widest bandpass filter settings possible (given the level of background noise). However, emphasis was placed on recording spikes rather then LFP. Thus, if the dynamic range of the low-frequency components was too large to have appropriate amplification to see clear spikes (given the limited dynamic range), a bandpass filter was used to allow appropriate increases in gain. All gain and filter settings were determined before recording started. This limitation only applies to the first-generation system (analog) that we used. The second-generation system did not have this constraint due to the increased dynamic range of the A/D converter, which has 18 effective bits. With this system, we could always record the entire frequency band (1 Hz–9000 Hz bandpass filter).

7.2 The origin and structure of the extracellular signal

The wideband extracellular signal recorded from a microwire electrode with relatively high impedance (200 kOhm–1 MohM) and small surface area contains a mixture of electrical signals from many different sources. Electrical events in neurons occur on two fundamental timescales: i) spikes are fast events that last 0.4 – 1 ms and ii) excitatory and inhibitory post synaptic potentials (EPSPs and IPSPs), on the other hand, are slow events that last from 10–100 ms. These two timescales are reflected in the structure of the wideband extracellular signal. The high-frequency components (> 300 Hz) are dominated by spikes, whereas the low-frequency components (< 300 Hz) are dominated by synaptic events. Simulations show that spikes contribute dominantly to the 300–3000 Hz frequency band and have negligible power at lower frequencies (See Figure 15 in (Logothetis, 2002) for an insightful illustration of this fact). This is the justification for using the 300-3000Hz frequency band for extracting spikes from the extracellular signal. Simulated synaptic potentials, on the other hand, have their dominant power at frequencies lower than 150 Hz (Logothetis, 2002).

The shape of the waveform of the spikes of a particular neuron (see spike sorting chapter for examples) depends on many factors such as the location, surface shape, and impedance of the electrode, as well as neuronal morphology and type, and the expression of different ion channels (Gold, 2007; Gold et al., 2006). For this reason waveforms from different neurons recorded on the same electrode are different. We exploit this fact to attribute each waveform to a particular neuron (see spike sorting chapter). These differences can also be exploited to infer properties of the recorded neuron from the shape of the action potential. For example, inhibitory neurons have sharper waveforms than excitatory neurons (McCormick et al., 1985). This fact can be used to infer the identity of the neuron recorded from (Buzsaki and Eidelberg, 1982; Csicsvari et al., 1999; Fox and Ranck, 1981; McCormick et al., 1985; Mitchell et al., 2007; Viskontas et al., 2007). Extracellularly recorded waveforms have amplitudes on the order of 50–200 μV (peak-to-peak). Background noise levels are in the range of 5–20 μV (RMS). Given these noise levels and the fact that the amplitude decays linearly with the distance from the source, it is estimated that an extracellular electrode can record spikes from neurons within a radius of perhaps up to 140 uM (Buzsaki, 2004; Gold, 2007; Gold et al., 2006; Henze et al., 2000; Holmgren et al., 2003).

The origin of the low-frequency components of the extracellular field, the local field potential (LFP), are much less clear (Bullock, 1997). It is thought that the LFP is mostly composed of the sum of large numbers of postsynaptic discharges. It is estimated that the LFP from a single extracellular electrode is influenced by potentials within a radius of 0.5–3 mm (Juergens et al., 1999; Logothetis, 2002; Mitzdorf, 1985). Due to its (predominantly) synaptic origin, the LFP can be independent of the spiking output measured at a particular location (Logothetis, 2002). The LFP is dominated by synchronized synaptic/dendritic components of neurons that are oriented in space such that their potentials add rather then cancel. The organization of cortical pyramidal neurons yields a particularly large LFP because neurons are parallel, with dendrites in one direction and axons in the other direction. This yields an open field geometrical arrangement (Mitzdorf, 1985). Due to its dominantly synaptic origin, the LFP is thought to represent the synaptic input as well as local processing. Some have used this to argue that spikes measure the output and LFP the input to a particular area. There are cases, however, where this strict distinction does not hold. Also, the exact origin of the LFP (in general) remains unknown, and making this argument thus requires detailed knowledge about the neuronal architecture of the area under investigation.

7.3 Signal processing

7.3.1 Filtering

All filters were 4th order zero-phase-lag Butterworth filters unless otherwise noted. For spike extraction, signals were bandpass filtered between 300–3000 Hz. For LFP, signals were down-sampled to 1000 Hz sampling rate and lowpass filtered < 300 Hz. To extract specific LFP frequencies (for example 4–8 Hz), a narrow bandpass filter was applied.

7.3.2 Local field potential (LFP)

The LFP is the sum of all oscillations that influence the extracellular electrical field at the point of space where the electrode is placed. There are many different forms of oscillations of widely varying frequencies. Some of these oscillations are known to have distinct physiological mechanisms. For example, oscillations of some frequencies are only present during sleep or during motor movement. The LFP bands are traditionally (and arbitrarily) decomposed into the following frequency bands (Buzsáki, 2006; Penttonen and Buzsaki, 2003): Delta (δ, 0–4 Hz), theta (θ, 4–8 Hz) , Alpha (α, 8–12 Hz), Beta (β, 12–24 Hz), Gamma (γ, 24–100 Hz or higher for high gamma). The frequency of a particular oscillation, however, can vary substantially depending on brain state (wake, sleep, drowsy) as well as between species (Steriade et al., 1993). For example, the frequency of theta is slower in larger mammals (such as primates or cats; 3–5

Hz) compared to rodents (6–9 Hz) (Robinson, 1980). These terms should thus only be used as guidelines but not as fixed entities.

Since the recorded LFP is a mixture of many frequencies (a voltage as a function of time), it is necessary to decompose the signal into a different representation which is a function of both frequency and time W(t,f). The fundamental technique to achieve this is the Fourier transform (FT), which transforms a function of time x(t) into a function of frequency x(f) (and vice versa). While this is useful to calculate a power spectrum, all time resolution is lost. One technique to circumvent this is to split the data into small time bins and calculate the FT for each (windowed fourier transform, WFT) (Teolis, 1998). Due to the small window in time, this technique will prevent estimation of frequencies whose wavelength is less than the window size. A more sophisticated version of WFT is wavelet analysis. Wavelets (see below) are functions which are well localized in both time and frequency. Their effective window size is adapted based on the frequency and is thus always optimal. Here, wavelets or the Hilbert transform were used to compute a continuous estimate of power and phase as a function of time.

Time-frequency decomposition using wavelets: The raw signal $S(t)$ was decomposed into a function of frequency and time using the continuous wavelet transform (cwt). In the following I am using the notation developed in (Torrence and Compo, 1998). The mother wavelet used was always a complex Morlet wavelet:

$$\psi_0(\eta) = \frac{1}{\sqrt{\pi\omega}} \exp(i2\pi\eta f_0)\exp(-\frac{\eta^2}{2})$$

The two parameters are the center frequency f_0 and the number of cycles. We used $f_0 = 1$ and $\omega = 4$ cycles, unless mentioned otherwise (see below).

The cwt of the raw signal $S(t)$ is a function of both scale (frequency) and time: $W(t,s)$. It is computed by convolving the raw signal (of length N) with the wavelet function $\psi_0(\eta)$ for a number of different frequencies (scales) s.

$$W(t,s) = \sum_{t'=0}^{N-1} S(t')\psi^*\left[\frac{(t'-t)\Delta t}{s}\right]$$

$\psi^*(\eta)$ is the complex conjugate of the wavelet function $\psi(\eta)$. $\psi(\eta)$ is a normalized version of the wavelet $\psi_0(\eta)$. See (Torrence and Compo, 1998) for details.

The effective resolution of the Morlet wavelet depends on the center frequency f_0 and the scale s. If δT is the spacing between two sampled points (due to the sampling rate), the effective frequency of a Morlet wavelet at scale s is $f = \frac{f_0}{s\delta T}$. Thus, the higher the scale, the lower the frequency. The resolution is measured separately in terms of the standard deviation in time σ_t and frequency σ_f. Time resolution at scale s is $a\delta T$ and frequency resolution is $\frac{\sigma_f}{a}$. Thus, the better the resolution in time the worse it is in frequency and vice versa (uncertainty principle, a fundamental limit, dictates $\sigma_t \sigma_f \leq \frac{1}{2\pi}$). The time width of a wavelet is defined as (Najmi and Sadowsky, 1997):

$$\sigma_t^2 = \frac{\int_{-\infty}^{\infty} t^2 \psi^2(t)dt}{\int_{-\infty}^{\infty} \psi^2(t)dt}$$

Thus, $\sigma_f = \frac{1}{2\pi s \sigma_t}$. To illustrate this trade-off, Figure 7-1 shows Morlet wavelets in both time and frequency space for 3 different parameter combinations. The time and frequency resolution for the same 3 wavelets are shown in Figure 7-2. Notice the trade-off between accuracy in time and frequency clearly visible from the size of the error bars in Figure 7-2 (bottom row). Since the width in frequency space increases as a function of frequency, the frequencies at which the wavelets are calculated are logarithmically scaled. This leads to an even sampling in frequency space (Figure 7-2). Here, we sampled at frequencies of $f = 2^x$ with $x \in [2:2:52]/8$ (not all are shown in Figure 7-2).

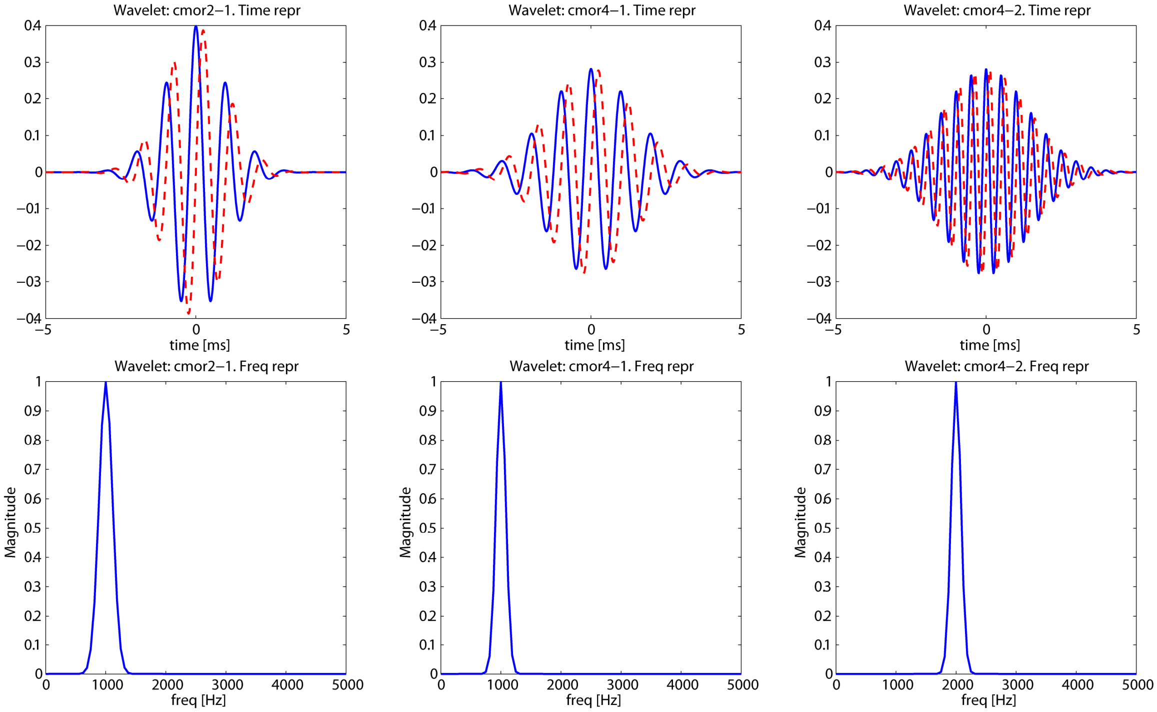

Figure 7-1. Illustration of the complex morlet wavelet.
The wavelet is illustrated for 3 different combinations of parameters of cycle number and center frequency: (4,1), (6,1), and (6,2) (from left to right). The top row shows the wavelet in time (blue: real part; red: complex part) and the bottom row in frequency (Fourier transform of the above). Notice the tight tuning in both time and frequency.

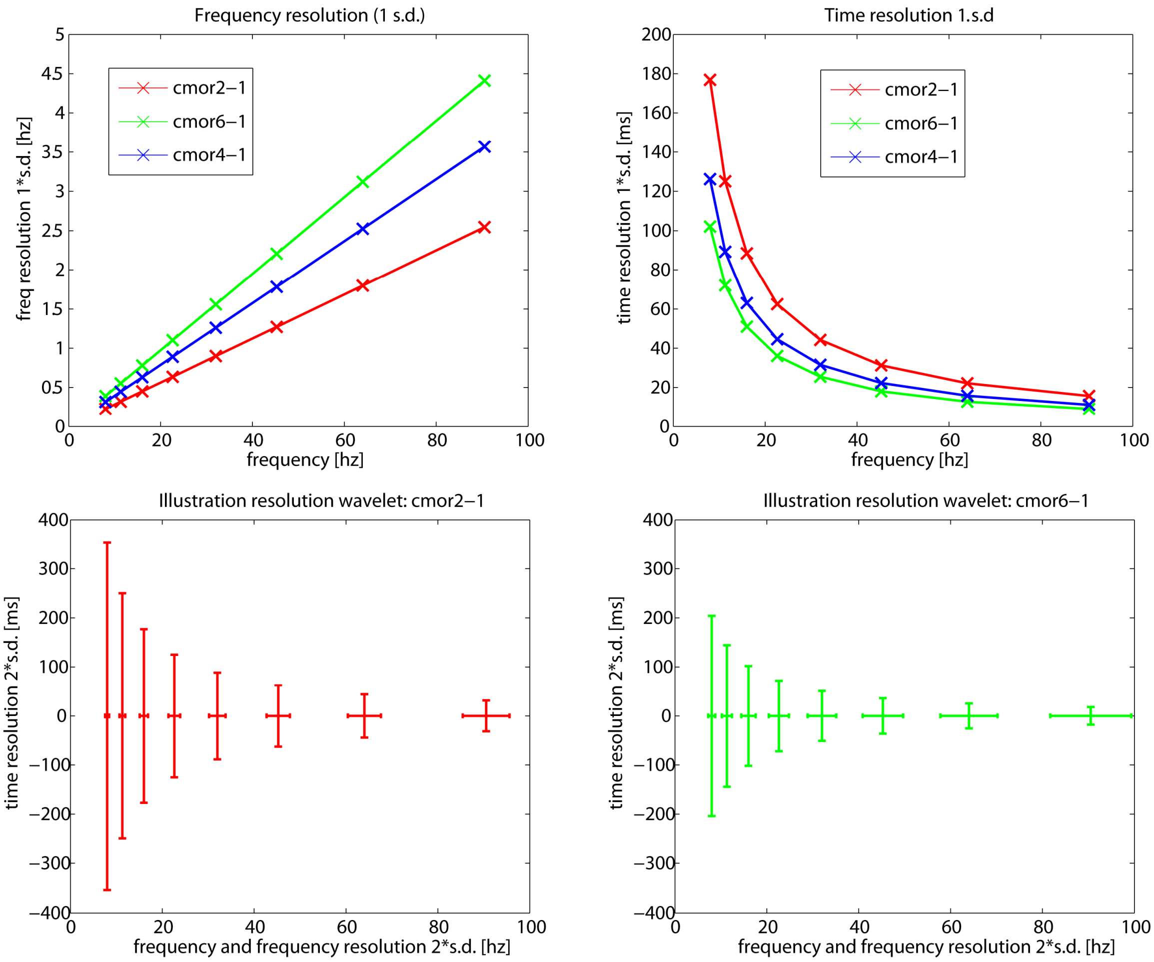

Figure 7-2. Illustration of the trade-off between specificity in time and frequency space.

Illustrated is the complex morlet wavelet for three different parameters (cycle number and center frequency): (2,1), (4,1), and (6,1). The top row shows the frequency resolution (left) and the time resolution (right). The y axis shows one standard deviation as a function of frequency. Note that whenever one wavelet has better time resolution (left, red) it has worse frequency resolution (right, red) and vica-versa. The bottom row illustrates this property by showing the 95% confidence interval (±2*s.d.) for both time (y axis) and frequency (y axis). Note how the left wavelet (2,1) has better resolution in frequency compared to the wavelet on the right (6,1). However, the left wavelet has better time resolution. Only a subset of the frequencies used for the analyis are shown (every second is shown). Note that in the bottom row, the y axis is in log units, and thus the error bars appear to be of equal length.

Computation of the analytic signal with the Hilbert transform: To estimate the phase and power of a narrowly bandpass-filtered signal without using the wavelet transform (such as in the theta band), the Hilbert transform was used to calculate an analytical signal. The analytical signal $X(t)$ is complex and can be used to calculate the phase/power with the same methods as for wavelet coefficients (see below). The Hilbert transform $S_H(t)$ is equal to the signal phase shifted by 90°. The real part of the analytic signal equals the raw signal and the complex part is the Hilbert transformed signal.

$$X(t) = S(t) + iS_H(t)$$

Estimation of instantaneous phase and power (energy): Given a complex signal as a function of time $X(t)$, the following methods were used to estimate the instantaneous phase $\phi(t)$ and power $R(t)$. $X(t)$ is either the result of a Hilbert transform or a continuous wavelet transform (see above). In the following, $\Re\{X\}$ and $\Im\{x\}$ refer to the real and imaginary part of X, respectively.

$$R(t) = \Re\{X(t)\}^2 + \Im\{X(t)\}^2$$

$$\phi(t) = \arg(X(t)) = \operatorname{atan}2(\Im\{X(t)\}, \Re\{X(t)\})$$

Wavelet power spectrum and distribution of wavelet power: The wavelet power spectrum is equivalent to $R(t)$ as defined above, as a function of frequency. The real and the imaginary parts of the wavelet coefficients $W(t,s)$ for any particular scale are normally distributed random variables with mean 0 and unknown variance. Since the wavelet power $R(t)$ is the squared sum of the real and imaginary part of $W(t,s)$, $R(t)$ is χ^2 distributed with 2 degrees of freedom. Since the variance is unknown (but not 1), however, the mean of this variable is unknown. Since the LFP is 1/f distributed, the mean of this distribution is a function of the frequency. The χ^2 distribution needs to be scaled appropriately to allow statistical tests (Caplan et al., 2003).

An alternative approach is to normalize the real and complex part of $W(t,s)$ to a variance of 1 independently before calculating the power. This removes the 1/f frequency dependency and allows easy statistical comparisons with a unscaled χ^2 distribution. I estimated the variance of $W(t,s)$ (separately for the real and imaginary part) for each scale for the entire experiment and normalized (divided) all samples by this value to assure that each are normally distributed with mean zero and variance 1. This allows the construction of a flat power spectrum (instead of 1/f) where peaks correspond to a deviation from the null hypothesis of no signal.

Statistics of phase locking: All statistics related to phases were performed using circular statistics (Batschelet, 1981; Fisher, 1993). The phase was measured (in radians) in the range $[-\pi ... \pi]$ (-180°–+180°) with 0 equal to the peak and $-\pi / \pi$ equal to the through of the oscillation. Statistics of a sample of n angles θ_i (phases) were calculated based on the mean resultant vector:

$$C = \sum_{i=1}^{n} \cos(\theta_i)\,,\ S = \sum_{i=1}^{n} \sin(\theta_i)\,,\ R^2 = C^2 + S^2\,,\ \overline{R} = \frac{R}{n}$$

The mean angle $\overline{\theta}$ is also calculated from above measures:

$$\cos(\overline{\theta}) = \frac{C}{R},\ \sin(\overline{\theta}) = \frac{S}{R}$$

The larger the length of the mean resultant $\overline{R}$ (range 0–1), the stronger the phase locking of the sample. The sample circular variance is $V = 1 - \overline{R}$. To test whether a neuron is significantly phase locked, the sample of all phase angles was compared against uniformity using a Rayleigh test. The Rayleigh test is based on the length of $\overline{R}$:

$$Z = n\overline{R}^2\,,\ P = \exp(-Z)\left[1 + \frac{2Z - Z^2}{4n} - \frac{24Z - 132Z^2 + 76Z^3 - 9Z^4}{288n^2}\right]$$

If P is sufficiently small, the null hypothesis of uniformity can be rejected. The alternative hypothesis is that the data is unimodal (one mean direction). To quantify the distribution of a sample of phase values that was significantly non-uniform, we fit a Von Mises distribution to the data using maximum likelihood. The Von Mises distribution is the normal distribution adapted for circular data. The following is its density function:

$$f(\theta) = \frac{1}{2\pi I_0(\kappa)} \exp(\kappa \cos(\theta - \mu)$$

It is fully specified by a mean direction μ and a concentration parameter κ. The concentration parameter is analogous to the standard deviation of a normal distribution, although of opposite direction: The larger κ, the more concentrated the distribution (the smaller its variance). For $\kappa = 0$, the Von Mises distribution is equivalent to the uniform distribution on the circle. $I_0()$ is the modified Bessel function of order zero. A definition of it can be found in (Fisher, 1993).

Simulated LFP: For systematic evaluation of our methods we used artificially simulated LFP which has a phase spectrum similar to real data ("red noise", i.e., "1/f"). Such LFP was simulated using sinusoidal pink noise (Cohen, 1995; Rohani et al., 2004):

$$X(t) = \sum_{i=1}^{N/2} \sin(\frac{2\pi t}{i}) + \phi_i)\sqrt{\frac{i}{N}}$$

This generates a time series of length N. The phase is sampled randomly from a uniform distribution, i.e., $\phi_i \in U[0,2\pi]$.

7.4 Spike train analysis

7.4.1 Single neurons: Spike times and the distribution of interspike intervals

In the following I describe the statistical properties of a series of spikes (a “spike train”). In probability theory, this is commonly referred to as a point process. For details and proofs refer to (Dayan and Abbott, 2001; Gabbiani and Koch, 1999; Kass et al., 2005; Koch, 1999).

For the purposes of analysis, spikes are treated as unitary events that occur at a particular point of time t_i. Spikes emitted by real neurons last 0.5–1.5 ms and are thus not restricted to a single point of time. Here, the peak (maximal deviation from baseline) of the waveform is used as the point of time the neuron spikes. The measurement accuracy of t_i is restricted by the sampling rate and uncertainty in determining the peak. Here, the accuracy is estimated to be on the order of 0.1 ms. Time is measured relative to a fixed reference point, such as the start of the experiment or trial. The unit of time is usually assumed to be milliseconds (ms), but any units can be used. Observing the N spikes emitted by a single identified neuron leads to a set of spike times $T = \{t_1, t_2, ..., t_N\}$. T is thus a list of events emitted by a point process. Observing the properties of T allows us to make inferences about the properties of this point process (which here is equal to a single neuron). The most important measure to quantify the behavior of a point process is the interspike interval (ISI). The interspike intervals are defined as the times between two neighboring spikes, i.e., $I_1 = t_2 - t_1, I_2 = t_3 - t_{2,...}$. The set of ISIs $I = \{I_1, I_2, ..., I_{N-1}\}$ is the set of all differences between neighboring spikes. The shape of the distribution of the ISIs can be used to infer a great number of properties about the neuron that emitted the spikes. Examples are:

inferences about the firing rate (the mean), the variability of the firing rate, bursting behavior, or whether the neuron fires periodically. Also, the shape of the ISI can be used to judge whether the set of spikes used for calculating it could have been emitted by a single neuron or not. This can be used to judge spike sorting quality (see the spike sorting chapter for details).

The spikes fired by a neuron are, in the great majority of cases, Poisson distributed (Dayan and Abbott, 2001; Holt et al., 1996; Softky and Koch, 1993). Due to biophysical constraints (such as the refractory period), neurons cannot fire at extremely high firing rates. Thus, it is unlikely that a neuron will fire more than one spike within approximately 3 ms (although there are cell types which have a shorter refractory period). For this reason, the firing probability at any particular point of time is a function of the time since the last spike. Such a process is modeled as a renewal process. For this reason the intervals deviate systematically from a pure Poisson distribution.

Given a homogenous Poisson process with rate r, the probability of observing n spikes within a time period T is:

$$P_T[n] = \frac{(rT)^n}{n!} e^{-rT}$$

The Poisson process is entirely defined by the rate r. Given a Poisson process, the waiting times between two spikes are exponentially distributed:

$$P_{ISI}(\tau) = re^{-r\tau}$$

The above function specifies the probability that, given a spike at $t = 0$, no spike will have occurred in the interval $t + \tau$. For a homogenous Poisson process this is the expected shape of the ISI distribution. Given a sample of ISIs, the mean ISI is $\langle ISI \rangle = \frac{1}{r}$ and the variance is $\sigma^2 = \left\langle \left[ISI - \langle ISI \rangle \right]^2 \right\rangle = \frac{1}{r^2}$. The mean and variance of the underlying Poisson process can thus be calculated from the ISI distribution.

If an absolute refractory period t_{ref} is introduced, this function is shifted (to the right on the time axis) by t_{ref}: $P_{ref}(\tau) = P_{ISI}(\tau - t_{ref})$. Another possibility for expression of the ISI distribution of a neuron with a refractory period is to use a gamma distribution:

$$P_{ISI}(\tau) = \frac{r(r\tau)^k e^{-r\tau}}{k!}$$

This representation has two parameters: the rate r and the parameter k, which is the order of the gamma distribution. When $k = 0$ an exponential distribution results. With $k > 0$ (estimated from the data using maximum likelihood), this typically provides a very good fit for ISI distributions. Also note that the mean of a gamma process with $k > 0$ is the same as for $k = 0$. The estimate of the mean rate as $\langle ISI \rangle = \frac{1}{r}$ remains thus valid, regardless of the order of the gamma process.

Computationally, a random series of spike times that are Poisson distributed can be generated by sampling randomly from an exponential distribution (and discarding the ones which are less than the refractory period). The returned numbers are wait times (interspike intervals). This is the strategy that we used whenever random spike trains were generated.

Also note that (under the Poisson assumption) the standard deviation of the ISI distribution is equal to its mean (the rate). Thus, the variance is not independent of the mean for neurons. The ratio of the standard deviation to the mean of the ISI of a perfectly homogenous Poisson process is thus equal to 1. This ratio is the coefficient of variation (CV):

$$CV_{ISI} = \frac{\sigma}{r}$$

The CV is an important measure of the regularity of firing of a neuron. A neuron that fires perfectly at a single rate has a CV of 0. A neuron that fires perfectly according to a Poisson distribution has a CV of 1.0. A neuron with highly irregular firing (for example, complex spikes, bursts) will have a CV > 1. The CV is routinely calculated for recorded neurons. There is a wide range of observed CV values. For neocortical neurons, it is often close to 1, as expected (Britten et al., 1993; Shadlen and Newsome, 1998; Softky and Koch, 1993; Tomko and Crapper, 1974). The measured relationship between mean rate and variance is approximately 1.5 (whereas the theoretical prediction is 1.0) (Shadlen and Newsome, 1998). A refractory period will impose some form of regularity and thus lowers the CV. A perfect Poisson neuron with a refractory period will thus have a CV < 1.