Quantifying variability in neural responses and its application for the validation of model predictions

Anne Hsu[1]**, Alexander Borst**[2] **and Frédéric E Theunissen**[3,4]

[1] Department of Physics, University of California, Berkeley, CA, USA
[2] Department of Systems and Computational Neurobiology, Max-Planck-Institute of Neurobiology, Martinsried, Germany
[3] Department of Psychology, University of California, Berkeley, CA, USA
[4] Helen Wills Neuroscience Institute, University of California, Berkeley, CA, USA

E-mail: fet@socrates.berkeley.edu

Abstract

A rate code assumes that a neuron's response is completely characterized by its time-varying mean firing rate. This assumption has successfully described neural responses in many systems. The noise in rate coding neurons can be quantified by the coherence function or the correlation coefficient between the neuron's deterministic time-varying mean rate and noise corrupted single spike trains. Because of the finite data size, the mean rate cannot be known exactly and must be approximated. We introduce novel unbiased estimators for the measures of coherence and correlation which are based on the extrapolation of the signal to noise ratio in the neural response to infinite data size. We then describe the application of these estimates to the validation of the class of stimulus–response models that assume that the mean firing rate captures all the information embedded in the neural response. We explain how these quantifiers can be used to separate response prediction errors that are due to inaccurate model assumptions from errors due to noise inherent in neuronal spike trains.

(Some figures in this article are in colour only in the electronic version)

1. Introduction

Knowledge of spiking variability is essential to the understanding of neuronal function. At the cellular level, the characterization of neuronal spike train variability is important for studying the biophysical mechanisms underlying neural behavior. At the systems level, spiking variability limits the capacity of a neuron to convey information about a stimulus, potentially restricting perceptual accuracy. In addition, the nature of spiking variability is directly related to the neural code.

Many models of spiking neurons assume that the deterministic part of the neural response can be fully characterized by a time-varying rate function (Gabbiani 1996, Johnson 1996, Baddeley *et al* 1997, Svirskis and Rinzel 2000, Barbieri *et al* 2001) and it has been shown that in many sensory systems this time-varying mean rate fully characterizes the information in neural response (Baker *et al* 1991, Oram *et al* 1999). Recent experimental and computational work has also demonstrated that neurons are capable of a high degree of spike timing precision (Bair and Koch 1996, Lestienne 2001) and high rates of information transfer (de Ruyter van Steveninck *et al* 1997, Reinagel and Reid 2000, Borst and Haag 2001, Lewen *et al* 2001). The high spiking precision observed experimentally does not necessarily violate the rate coding assumption. Most cases of precise spiking observed in sensory systems have been in response to a stimulus with relatively fast dynamics. In these cases, the reliable spiking pattern is locked in time relative to stimulus onset and presumably phase locked to the stimulus dynamics (Theunissen and Miller 1995, Oram *et al* 1999, Lestienne 2001). These coding regimes are adequately described by a mean rate function with rapid fluctuations. The rate coding model, however, is an incomplete description of more complex neuronal response properties, where spike patterns which are not phase locked to the stimulus can encode specific stimulus attributes. For example, in the olfactory system of the locust, temporal patterns of neural activity that are not phase locked to stimulus dynamics have been shown to carry information on odor identity (Wehr and Laurent 1996).

Under the rate coding assumption, neuronal variability is considered as noise. This neuronal noise can be characterized by estimating the expected difference between a single example of neuronal response (a spike train) and its time-varying mean rate function. This difference can be quantified by calculating the correlation coefficient (CC) between the time-varying mean response and individual sample spike trains. More generally, the coherence between sample spike trains and the mean rate will quantify the signal to noise ratio in the neuronal response as a function of frequency (Brillinger 1975). There are two issues with the estimation of these quantifiers. First, the presence of noise in high frequencies due to the spike shapes or the representation chosen for spikes can result in an overestimation of the coherence and an underestimation of the correlation coefficient. Second and more importantly, since the actual mean rate is not known, sampling methods must be used to obtain unbiased estimates of both the correlation coefficient and the coherence. We propose a simple methodology for these estimations. We use both model neurons and experimental data to demonstrate that the method converges rapidly and that reliable estimates of correlation and coherence can be obtained with a relatively small number of trials.

We then show how this methodology can be used to judge the accuracy of models of stimulus–response functions for single neurons. We have applied our validation methodology to spectro-temporal receptive fields, a linear stimulus–response model that we have used extensively in our laboratory. However, our procedure is equally suitable for the validation of all (including nonlinear) models of stimulus response functions that assume the neuronal response is entirely described by a time-varying mean firing rate. Spectro-temporal receptive fields in audition and spatial-temporal receptive fields in vision are linear models often used to describe the response properties of high-level sensory neurons (Theunissen *et al* 2000, 2001). Nonlinear models include predictions obtained from neural networks (e.g. Lau *et al* 2002) or from specific analytical models (Victor and Knight 1979). To validate these models of stimulus–response functions, the predicted response must be compared with the actual response, the time-varying mean firing rate being estimated from real spike data. In previous studies, this comparison was often quantified using the correlation coefficient or the coherence between the predicted response and an estimation of the actual response obtained by averaging the spike data (e.g. Eggermont *et al* 1983, Dan *et al* 1996, Theunissen *et al* 2000).

However, since the observed response of the neuron is a combination of a predictable or deterministic response plus noise, an estimation of the actual response by averaging the spike data will be inaccurate. More recently, methods have been proposed to accurately estimate the power in the deterministic response using the assumptions of additive and independent noise (Sahani and Linden 2002). For stimulus–response model validation, this deterministic power can then be compared to the predictive power. Using similar assumptions, we show how to estimate the coherence and correlation between the predicted mean firing rate and a single sample spike train. As long as the neuron's response to a stimulus can be captured by its time-varying mean firing rate, the prediction can be validated by comparing these measures with those obtained between the actual mean firing rate and a single sample spike train. We are then able to separate prediction errors due to the limitations of the stimulus–response model function from those due to the intrinsic noise in the response.

2. Correlation coefficient, coherence and coherency

The correlation coefficient, r, measures the linear relationship between two variables, and for two time-varying signals $s_1(t)$ and $s_2(t)$, it is often calculated with

$$r = \frac{\langle (s_1(t) - \overline{s_1(t)})(s_2(t) - \overline{s_2(t)}) \rangle_t}{\sqrt{\langle (s_1(t) - \overline{s_1(t)})^2 \rangle_t \langle (s_2(t) - \overline{s_2(t)})^2 \rangle_t}}, \tag{1}$$

where the angular brackets and overbar designate expected values taken by averaging over time.

The coherence, $|\gamma|^2$, between the two signals $s_1(t)$ and $s_2(t)$, is given by (Brillinger 1975):

$$|\gamma_{S_1 S_2}(\omega)|^2 = \frac{\langle S_1(\omega) S_2^*(\omega) \rangle \langle S_1^*(\omega) S_2(\omega) \rangle}{\langle S_1(\omega) S_1^*(\omega) \rangle \langle S_2(\omega) S_2^*(\omega) \rangle}, \tag{2}$$

where $S_1(\omega)$ is the Fourier transform for $s_1(t)$ and the angular brackets designate the expected value obtained by averaging over sections of data: the Fourier components of time signals s_1 and s_2 are obtained by windowing the data into multiple sections. Given enough data, the time window should be chosen to be on the order of the length of correlations between the two signals. In multi-tapered coherence estimation methods, multiple orthogonal windows are used for each section resulting in increased robustness of the estimation (Thomson and Chave 1991).

The coherence measures the degree of linear relationship between S_1 and S_2 at each frequency ω. It is the generalization of the correlation coefficient for time series and is effectively the correlation coefficient squared for each frequency component of the two signals. It is a superior measure to the correlation coefficient because it addresses many possible scenarios in time series where the correlation coefficient would give an erroneous description of the linear relationship of the two signals (i.e. when the two signals are related by a linear filter). For example, the two signals could be linearly related but with a particular delay resulting in a very low (or zero) correlation coefficient. The correlation coefficient can also give erroneously high measures of correlations when the two signals have significant autocorrelation. In that situation, an actual correlation at one given point in time between the two signals could lead to additional non-independent correlations at later times. Nonetheless, the correlation coefficient is a simpler measure than coherence, and has been used extensively in neuroscience. For cases in which there is no delay and in which the autocorrelations of the signals are small, the two measures will yield very similar results.

The measures of correlation and coherence are affected both by the degree of linearity between the two signals and by the noise. This relationship is more rigorously described by

examining the coherence between the two time-varying signals under different conditions. If the two signals s_1 and s_2 are linearly related and noise free, their coherence is equal to 1, as demonstrated in the following short proof: if $s_2(t) = \int_{-\infty}^{\infty} h(t - \tau)s_1(\tau) \, d\tau^*(\omega)$, then $S_1(\omega) = H(\omega)S_2(\omega)$ (convolution theorem). Thus,

$$\left|\gamma_{S_1 S_2}(\omega)\right|^2 = \frac{\langle S_1(\omega)H^*(\omega)S_1^*(\omega)\rangle \langle S_1(\omega)^* H(\omega)S_1(\omega)\rangle}{\langle S_1(\omega)S_1^*(\omega)\rangle \langle H(\omega)S_1(\omega)H^*(\omega)S_1^*(\omega)\rangle} = 1,$$

since $H(\omega)$ is a constant and can be pulled out of the expected values.

The coherence is less than 1 if the two signals are either nonlinearly related or if the two signals are linearly related but corrupted by noise. In these cases, the coherence is a function of the signal to noise ratio of the signals (Marmeralis and Marmeralis 1978). In the general case, for two noisy signals $s_1' = s_1 + n_1$ and $s_2' = s_2 + n_2$, the coherence between s_1' and s_2' is

$$\left|\gamma_{S_1' S_2'}(\omega)\right|^2 = \frac{\langle S_1(\omega)S_2^*(\omega)\rangle \langle S_1^*(\omega)S_2(\omega)\rangle}{\langle S_1(\omega)S_1^*(\omega) + N_1(\omega)N_1^*(\omega)\rangle \langle S_2(\omega)S_2^*(\omega) + N_2(\omega)N_2^*(\omega)\rangle} < 1.$$

In a stimulus–response paradigm, we can hypothesize an encoding scenario in which the neural response, $s_2'(t)$, is noisy but with a mean, $s_2(t)$ that is linearly related to the noiseless stimulus, $s_1'(t)$ (i.e. $n_1(t) = 0$), then

$$\left|\gamma_{S_1' S_2'}(\omega)\right|^2 = \frac{\langle|S_1(\omega)|^2\rangle}{\langle|S_1(\omega)|^2\rangle + \frac{\langle|N_2(\omega)|^2\rangle}{|H(\omega)|^2}} = \frac{\langle|S_2(\omega)|^2\rangle}{\langle|S_2(\omega)|^2\rangle + \langle|N_2(\omega)|^2\rangle} < 1.$$

If the neural response is not noisy, and not linearly related to the stimulus then

$$\left|\gamma_{S_1' S_2'}(\omega)\right|^2 = \frac{\langle S_1(\omega)S_2^*(\omega)\rangle \langle S_1^*(\omega)S_2(\omega)\rangle}{\langle S_1(\omega)S_1^*(\omega)\rangle \langle S_2(\omega)S_2^*(\omega)\rangle} < 1,$$

and if the neural response is both noisy and not linearly related to the stimulus then:

$$\left|\gamma_{S_1' S_2'}(\omega)\right|^2 = \frac{\langle S_1(\omega)S_2^*(\omega)\rangle \langle S_1^*(\omega)S_2(\omega)\rangle}{\langle S_1(\omega)S_1^*(\omega)\rangle \langle S_2(\omega)S_2^*(\omega) + N_2(\omega)N_2^*(\omega)\rangle} < 1.$$

In all three cases the coherence between the stimulus and the neural response will be less than 1. To separate the nonlinear encoding situation from the noisy neural response situation, estimations of the neural noise, n_2, need to be obtained. We will show how to estimate the neural noise and validate both linear and nonlinear stimulus–response models of neurons under the time-varying mean rate coding assumption. To validate the stimulus–response models we will calculate the coherence, not between the stimulus and the response, but between the predicted response obtained from the model and a single spike train. This coherence will then be compared to the coherence between a single spike train and the actual time-varying mean firing rate. This latter coherence effectively quantifies the neural noise. If the two coherences are equal, one can conclude that the proposed stimulus–response model describes the response property of the neuron. If a linear stimulus–response model is used, then the degree of discrepancy between the two coherences quantifies the degree of nonlinearity in the actual stimulus–response function.

The coherence can also be obtained from the correlation function calculated in the time domain. The correlation function, $R(\tau)$, measures the variance of the product between two signals with a time lag of τ:

$$R_{S_1 S_2}(\tau) = \langle s_1(t)s_2(t + \tau)\rangle_t, \tag{3}$$

where the angular brackets mean the average across time t. The Fourier transform (FT) of $R(\tau)$ is given by

$$\mathrm{FT}\left\{R_{S_1 S_2}(\tau)\right\} = \langle S_1^*(\omega)S_2(\omega)\rangle \qquad \text{(correlation theorem)},$$

where we now average across the windowed segments of data. Therefore, the coherence between S_1 and S_2 can be written as

$$\left|\gamma_{S_1 S_2}(\omega)\right|^2 = \frac{\text{FT}\{R_{S_1 S_2}(\tau)\}\text{FT}\{R_{S_2 S_1}(\tau)\}}{\text{FT}\{R_{S_1 S_1}(\tau)\}\text{FT}\{R_{S_2 S_2}(\tau)\}}.$$

A time series measure of coherence can also be obtained in the time domain by taking the inverse Fourier transform (IFT) of the correlation function between the two signals normalized by the square root of the product of the autocorrelation functions. This measure is called the coherency and is given by

$$\gamma_{S_1 S_2}(\tau) = \text{IFT}\left\{\frac{\text{FT}\{R_{S_1 S_2}(\tau)\}}{\sqrt{\text{FT}\{R_{S_1 S_1}(\tau)\}\text{FT}\{R_{S_2 S_2}(\tau)\}}}\right\},$$

where τ is a time delay. In addition to quantifying the degree of linear relationship between s_1 and s_2, the coherency also shows the phase relationship or, equivalently, the overall deterministic delay between the two signals. Note that the coherency at delay time zero is generally not equal to the correlation coefficient. The two measures will be equal only if the signals have zero mean and if the autocorrelation functions are delta functions.

Correlation and coherency are defined between two signals that are functions of time. Spike data can be represented in the time domain as a sum of discrete time delta functions. (Each time bin will have either 1 or 0 for spike or no spike.) The resolution of the discrete time should be chosen to preserve the order of spike arrival times, i.e. no more than one spike in each time bin. For our calculations we chose a discrete time resolution of 1 ms. Under our assumptions that neural information is conveyed in a time-varying mean firing rate, it is useful to estimate spike trains to be smoother functions of time by convolving the spike train with a smoothing window. Subsequent data processing results will depend on the size and shape of the smoothing window. While smoothing spike trains is a useful pre-processing step for the calculation of CC, the calculation of coherence does not require this smoothing, as explained with further details in the next section.

3. Quantifying neural noise with correlation and coherence

The correlation coefficient and coherence can both be used to quantify the neural noise in spike trains when the signal is assumed to be the time-varying mean firing rate. Under that assumption, the noise in rate-coding neurons is the difference between a single spike train and the time-varying mean rate. Since the raw spike train is a sum of delta functions, the neural noise calculated in this fashion often has broadband energy with significant amounts of energy at high frequencies where signal power is relatively small. This characteristic neural noise spectrum is due to the discrete nature of spikes and the small variations in spike timing. In these situations, calculating the CC without smoothing the neural responses will result in very low values since all frequencies will be weighted equally in the calculation. In order to reduce the effect of high-frequency noise in the estimate of the CC, the data can be low-pass filtered with a smoothing window. The choice of window can have strong effects on the CC measure. Windows that are too large will yield values of CC arbitrarily close to 1 because in the limit of an infinite sized window, the smoothed signals approach constant values in time. Windows that are too small will yield arbitrarily small CCs due to the confounding presence of the high-frequency noise mentioned above. As we will describe in more detail below, an optimal window can be found when the CC is used to quantify the goodness of fit of the proposed

stimulus–response functions for the neurons. In such situations, the appropriate time scale of the smoothing window can be obtained by only smoothing the neural response. The predicted response from the stimulus–response function is *not* smoothed and therefore determines the appropriate time scale of the mean response signal: the resulting smoothing time window will take into account the amount of prediction power relative to response noise power at different frequencies.

In contrast to the CC, the coherence is a function of frequency which naturally separates signal from noise. Therefore, the spike trains do not need to be smoothed and are left in their original time representation, a sum of discrete delta functions. The resolution of discrete time chosen to represent the spike train does determine the coherence function's Nyquist frequency, the highest frequency at which signal power can be detected. However, by choosing a small enough time resolution (e.g. 1 ms), coherence power of neuronal data will be almost zero or largely noise well before the Nyquist frequency limit. A more critical choice, with a similar effect to that of a smoothing window, is the choice of a cutoff frequency above which the estimated coherence can be disregarded and small positive values attributed entirely to sampling error. Analogous to the choice of smoothing window for the CC calculation, one wants to choose a cutoff which keeps all frequencies at which there are significant signal dynamics and which neglects all frequencies that are mostly corrupted by noise. Again, there is the risk of cutting out too much signal power, or including too much noise. We suggest finding the lowest frequency at which lower bound estimates of coherence equal zero, and setting coherence values equal to zero from this point on. Lower and upper bounds on coherence can be obtained by using a jack-knife method (Thomson and Chave 1991).

The coherence is a function of frequency, and it can be useful to collapse this function into a number that summarizes the overall strength of correlation and effectively obtain a single measure akin to the CC. One possibility is to calculate the average coherence by integrating over all relevant frequencies. In its discrete formulation, the overall linear correlation could then be obtained by averaging all values of the coherence over the desired frequency range. Equivalently, invoking the discrete form of Parseval's theorem, the sum of the coherency squared would yield the same result:

$$\sum_{i=1}^{M} (\gamma(\tau_i))^2 = \frac{1}{M} \sum_{i=1}^{M} \gamma^2(f_i),$$

where M is the number of frequency bins or time bins used in the fast Fourier transform (FFT). If the upper frequency is not bounded and if the signal is corrupted with high-frequency noise, then this overall measure would approach zero as the upper frequency (or Nyquist limit) increases. This has the same effect as making the smoothing window arbitrarily small in the calculation of the correlation coefficient.

A second measure of overall coherence is to calculate a quantity, which we call the *normal mutual information* defined as follows:

$$I = -\int_0^\infty \log_2(1 - \gamma^2(\omega)) \, d\omega, \tag{4}$$

where $\gamma^2(\omega)$ is the coherence as a function of frequency. If the coherence in this formula is between a signal and another function that is the signal plus additive noise, the integral in equation (4) is related to Shannon's mutual information, given certain conditions. First, both signal and noise are assumed to be stationary. Second, both the signal and the noise are Gaussian. Under these conditions, the *normal mutual information* equals the mutual information, and thus the name. If the mean rate is Gaussian but the noise is not, equation (4) gives a lower bound estimate on the mutual information. Conversely, if the mean rate is not

Gaussian and the noise is, the same equation can be used to estimate the channel capacity, an upper bound on mutual information (Shannon and Weaver 1963, Borst and Theunissen 1999). Because of the discrete nature of spike trains, neural noise usually has a probability distribution that is far from Gaussian. However, if the mean firing rate of the neuron is Gaussian, a lower bound on the mutual information can be established. This can be useful when data limitations prevent reasonable estimations of mutual information in spike trains using a direct method for estimating the mutual information (Strong *et al* 1998). In the case that neither is normal, nothing can be said in terms of upper or lower bounds on mutual information. However, in all cases, the *normal mutual information* quantifies the noise power level and can be used as an alternate measure to the correlation coefficient or the average coherence.

4. Estimating the coherence and CC for neural data

In this section, for conciseness, we provide equations using the coherence and not the CC. However, the following examples of estimating coherence can be applied analogously to estimations of the CC squared. This is because the coherence is precisely the squared CC for each individual frequency component of the signals and the properties of uncorrelated noise are identical in time and frequency for stationary signals: by definition $n(t)$ is uncorrelated with $a(t)$ if and only if $\langle a(t)n(t+\tau)\rangle_t = 0$ for all τ or equivalently $\langle A(\omega)N(\omega)\rangle = 0$ for all ω.

The calculation of the coherence is done in the frequency domain. Since we assume that the noise is uncorrelated with the mean rate, we can divide the response power, R^2, into a deterministic part, the power of the actual mean rate, A^2, and the power of the neural noise, N^2. If the second-order statistics are stationary, we then have for each frequency, ω: $R^2(\omega) = A^2(\omega) + N^2(\omega)$, where for notational simplicity, we use $R^2(\omega) = \langle R(\omega)R^*(\omega)\rangle$ and similarly for A^2 and N^2.

In experiments, the actual mean rate is not known but, by averaging responses over multiple trials, one can obtain a post-stimulus time histogram (PSTH) that will converge to the real rate in the limit of infinite trials. If we call the PSTH obtained with M trials, $\overline{R}_M$, we then have

$$\overline{R}_M^2(\omega) = A^2(\omega) + N^2(\omega)/M.$$

The noise power in $\overline{R}_M$ is reduced by a factor, M, because, as long as A captures the deterministic part of the response, noise, by definition, will be independent from trial to trial. The coherence between the actual mean rate and the noise corrupted response, γ^2, can also be thought of as a normalized measure of the signal to noise $\frac{A^2(\omega)}{N^2(\omega)}$ (Marmeralis and Marmeralis 1978)

$$\gamma_{AR}^2(\omega) = \frac{\langle R(\omega)A^*(\omega)\rangle\langle R^*(\omega)A(\omega)\rangle}{\langle A(\omega)A^*(\omega)\rangle\langle R(\omega)R^*(\omega)\rangle} = \frac{A^2(\omega)}{A^2(\omega) + N^2(\omega)}. \tag{5}$$

Following the nomenclature of Borst (Haag and Borst 1997, Borst and Theunissen 1999), we call the above quantity the expected coherence when R is a single spike train.

We will show that a robust measure for the expected coherence shown in equation (5) can be obtained by calculating the coherence directly from the PSTH. Since the coherence between the actual rate and the PSTH is given as

$$\gamma_{A\overline{R}_M}^2(\omega) = \frac{\langle \overline{R}_M(\omega)A^*(\omega)\rangle\langle \overline{R}_M^*(\omega)A(\omega)\rangle}{\langle A(\omega)A^*(\omega)\rangle\langle \overline{R}_M(\omega)\overline{R}_M^*(\omega)\rangle} = \frac{A^2(\omega)}{A^2(\omega) + N^2(\omega)/M},$$

we then have

$$\gamma_{AR}^2 = \frac{\gamma_{A\overline{R}_M}^2}{\gamma_{A\overline{R}_M}^2 + M\left(1 - \gamma_{A\overline{R}_M}^2\right)}.$$ (6)

With equation (6), we can calculate the expected coherence (between one spike train and the actual mean rate) directly from the PSTH. Equation (6) yields the same results as averaging the coherence with mean rate obtained from individual spike trains. It also has comparable convergence properties with a number of spike train trials. Thus, equation (6) does not offer us anything new except saving computational time as the number of trials becomes large. However, equation (6) serves as a stepping stone for the case described next, where the mean rate is unknown.

To verify the validity of our approach and to estimate how well it estimated the expected coherence as a function of data size, we implemented a spike train model where we could specify the actual mean firing rate to be an arbitrary function of time. We modeled the stochastic aspect of the neuronal response by generating spike trains with point process statistics with gamma orders 1 (Poisson) and 3. Gamma processes are a specific class of point processes, defined by a varying mean rate and a gamma order constant. The higher the gamma order, the less variable the point process. The results of the coherence calculation using equation (6) for such a model neuron are shown in figure 1. We found that our estimate of the expected coherence converged rapidly with the number of trials. As expected, the more reliable gamma order 3 neuron shows greater coherence than the Poisson neuron.

The small values of expected coherence, shown in the middle panels of figure 1, result from taking the coherence between the underlying mean rate, a smooth function of time, and a single realization of the stochastic point process with that rate, a particular spike train. In other words, a single spike train is a poor representation of the actual time-varying mean firing rate. In order to obtain a good representation of this mean firing rate, neural systems must average across redundant neurons with similar time-varying mean firing rates. As indicated by the PSTH coherence, by averaging over as few as ten trials of spike trains (or ten redundant neurons), the coherence with the mean rate goes up significantly (figure 1, upper panels). Therefore, the spike train responses converge relatively quickly to the time-varying mean rate.

Cortical neurons show wide range variance versus mean firing rate noise behavior but are for the most part well modeled with Poisson statistics (Shadlen and Newsome 1998, Kara *et al* 2000). Mammalian peripheral visual neurons (Kara *et al* 2000) and visual neurons in the fly (de Ruyter van Steveninck *et al* 1997, Borst 2003) are remarkably more reliable than Poisson. Avian auditory forebrain neurons are, for the majority, well described by a Poisson process. However, a small fraction of neurons (8/185) are better described by a gamma process of order 2 or greater (Hsu and Theunissen, unpublished observations).

Equation (6) can be used for spike train modeling purposes when the actual mean rate function is known. For data obtained from an experiment, the actual rate function is unknown. Estimating the actual rate function to be the PSTH leads to positive biases in the coherence because the individual spike trains will be more correlated with the PSTH than with the actual rate function. Alternatively, one can estimate the actual rate from a PSTH from which the test spike train has been deleted. Other possibilities include implementing a bootstrap or any similar re-sampling technique. These methods, however, are computationally intensive for long data sets. Instead, we propose the following solution: estimate the expected coherence directly from the data composed of M trials by calculating the coherence between two PSTHs, each obtained from $M/2$ different trials.

$\overline{R}_{1,M/2}^2(\omega) = A^2(\omega) + N^2(\omega)2/M$ and $\overline{R}_{2,M/2}^2(\omega) = A^2(\omega) + N^2(\omega)2/M$. The coherence between these two PSTHs is
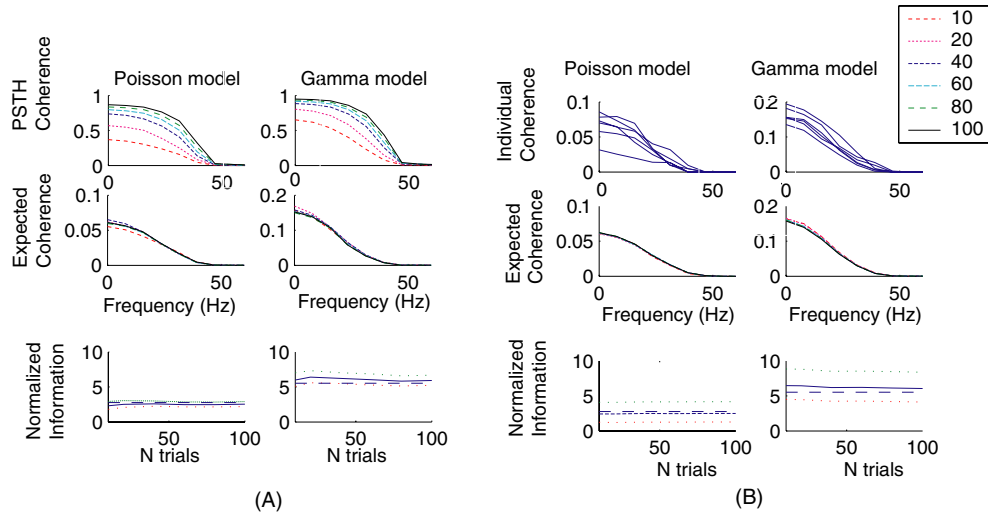
Figure 1. (A) Expected coherences estimated from the PSTH as a function of the number of trials with known mean rate. The expected coherence between a single spike train and its time-varying mean firing rate, γ^2_{AR}, can be estimated from the coherence between a PSTH of M trials and the mean rate, $\gamma^2_{A\overline{R}_M}$. The PSTHs were obtained from spike trains modeled with a Poisson process (left panels) and with a less noisy third-order gamma process (right panels). The results for $\gamma^2_{A\overline{R}_M}$ estimated with different numbers of trials (10, 20, 40, 60, 80, 100) are shown on the top panels. The estimated coherences, calculated with equation (6), γ^2_{AR}, are shown in the middle panels. The lines obtained for different numbers of trials overlap completely demonstrating the robustness of the estimation even for a small number of trials. The bottom panels show the normal mutual information as a function of the number of trials. The dashed line indicates the value estimated with 1000 trials. The dotted lines are two standard deviation error bars formed from jack-knifing the coherence. In these simulations, the spike trains were generated from a random 0–25 Hz band limited Gaussian mean rate function. The coherence cutoff occurred at 65 Hz. (B) Reliable estimates of the expected coherence are obtained with a small number of trials, converging as quickly and reliably as taking the average coherence of individual spike trains with the mean rate function. The calculations are from the same two model spike train sets as in (A). The top panel shows coherences of ten individual spike trains with the known mean rate. The middle panel shows coherence averages with different numbers of trials. The normal mutual information is shown in the bottom panel.

$$\gamma^2_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}(\omega) = \frac{\langle \overline{R}_{1,M/2}(\omega)\overline{R}^*_{2,M/2}(\omega)\rangle\langle \overline{R}^*_{1,M/2}(\omega)\overline{R}_{2,M/2}(\omega)\rangle}{\langle \overline{R}_{1,M/2}(\omega)\overline{R}^*_{1,M/2}(\omega)\rangle\langle \overline{R}_{2,M/2}(\omega)\overline{R}^*_{2,M/2}(\omega)\rangle} = \frac{A^2}{A^2 + \frac{4N^2}{M} + \frac{4N^4}{M^2A^2}}. \qquad (7)$$

We then get an estimate of the expected coherence, given by (appendix A):

$$\frac{1}{\gamma^2_{AR}} - 1 = \frac{1}{2}\left(-M + M\sqrt{\frac{1}{\gamma^2_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}}}\right). \qquad (8)$$

With this method, the expected coherence can be estimated with one simple calculation. We validated this approach both for a real neuron's responses and for modeled spike trains that had mean rates equal to those estimated from the real neuron's PSTH. The real neuron was the H1 visual interneuron of the fly. The real neural spike trains were obtained in response to a randomly drifting grating (Borst 2003). Figure 2 shows the convergence of different types of coherence estimates for data from the two model neurons as well as the H1 neuron.
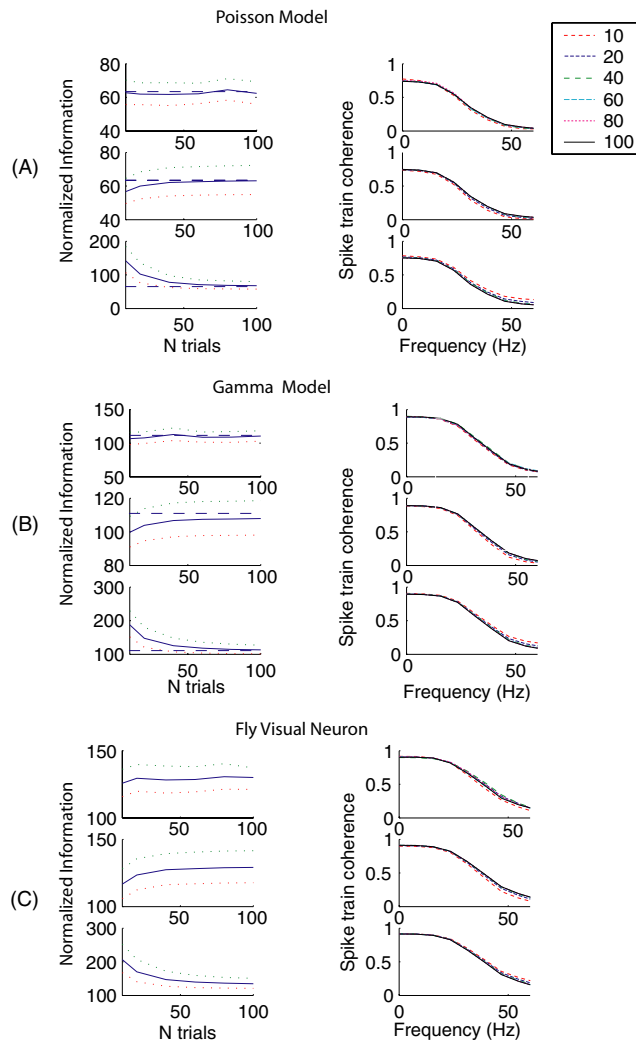
RIGHTSLINK()

Figure 2. Estimates of expected coherence when the actual mean rate is unknown, as a function of the number of trials. The calculation was performed for three cases: (A) model spike trains using a Poisson process; (B) model spike trains using a gamma order 3 process; and (C) for data obtained experimentally from the H1 neuron of the fly visual system (Borst 2003). The model spike trains were generated to have the same time-varying firing rate as was estimated from the PSTH of the fly data. The top panels in each case show the normal mutual information (left-top) obtained by estimating the expected coherence (right-top) using two PSTHs, each with half the total number of trials, calculated with equation (8). Reliable estimates can be obtained with a very small number of trials. The middle panels show expected coherence and normal mutual information estimates when the mean rate is replaced by the PSTH that excludes the spike train used to estimate the noise. In this case, the calculation results in a negative bias. The bottom panels show the estimates when the mean rate is replaced by the PSTH without deleting the trial spike train. This calculation method results in severe overestimates. Note the different scale of normalized information in the bottom panel for each case. The dashed line for the model data shows the normal mutual information estimated with 1000 trials of model data. The gamma order 3 model is a better fit to the real data, and hence its normal information values are more similar to the real data than that of the Poisson model.

Plotted as a function of the number of trials used are coherence estimates for each data set using equation (8) (top panels), using the PSTH with a deleted trial as the mean rate function (middle panels) and using the entire PSTH as the mean rate function (bottom panels). The left panels show the estimated normalized information while the right panels show the estimated coherence function. It can be seen that equation (8) provides reasonable unbiased estimates of the expected coherence with as few as twenty to forty trials. Using the PSTH as the mean rate function results in large positive biases. Using the PSTH with the trial deleted results in negative bias.

Sometimes it is useful to obtain the coherence between spike trains and some function other than the actual mean, as described in the following section on neural stimulus–response model validation. The coherence between a spike train R and any arbitrary rate function B is given by

$$\gamma_{BR}^2(\omega) = \frac{\langle R(\omega)B^*(\omega)\rangle\langle R^*(\omega)B(\omega)\rangle}{\langle B(\omega)B^*(\omega)\rangle\langle R(\omega)R^*(\omega)\rangle} = \frac{|AB|^2}{B^2(A^2 + N^2)}. \tag{9}$$

Again, the coherence between trial R and B, γ_{BR}^2, can be obtained from the coherence between a PSTH of M trials, $\overline{R}_M$ and B, $\gamma_{B\overline{R}_M}^2$. As previously, the power in the PSTH with M trials is equal to: $\overline{R}_M^2(\omega) = A^2(\omega) + N^2(\omega)/M$. Thus, we have for the coherence between $\overline{R}_M$ and B:

$$\gamma_{B\overline{R}_M}^2(\omega) = \frac{\langle \overline{R}_M(\omega)A^*(\omega)\rangle\langle \overline{R}_M^*(\omega)B(\omega)\rangle}{\langle B(\omega)B^*(\omega)\rangle\langle \overline{R}_M(\omega)\overline{R}_M^*(\omega)\rangle} = \frac{|AB|^2}{B^2\left(A^2 + \frac{N^2}{M}\right)}. \tag{10}$$

It can be seen that the ratio $\gamma_{BR}^2 / \gamma_{B\overline{R}_M}^2 = \gamma_{AR}^2 / \gamma_{A\overline{R}_M}^2$. That is the ratio between the coherence of M trials with B and the coherence of one trial with B is the same as the ratio between the coherences of M trials and one trial with the actual mean rate. In the real experimental situation where the actual mean rate is unknown, the ratio $\gamma_{AR}^2 / \gamma_{A\overline{R}_M}^2$ is shown in appendix B to be

$$\frac{\gamma_{AR}^2}{\gamma_{A\overline{R}_M}^2} = \frac{1 + \sqrt{1/\gamma_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}^2}}{-M + M\sqrt{1/\gamma_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}^2} + 2}, \tag{11}$$

where $\gamma_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}^2$ is the coherence of the two halves of the PSTH as mentioned before. As in the case when the actual mean rate is known, this calculation can also be done by averaging the coherence of each spike train with the function B. However, again, the use of equation (11) reduces computation time and reduces numerical artifacts due to the smoother nature of a PSTH versus a single spike train.

5. Validating stimulus–response models

A useful application of our coherence and CC estimation is the validation of stimulus–response models for single neurons under the time-varying rate coding assumption.

A stimulus–response model attempts to map neural stimuli to a neuronal response. These models can include both static nonlinearities, such as saturation, or dynamic non-linearities, such as adaptation. The stimulus–response model is validated by the quality of its predictions of neural responses to novel stimuli (i.e. using stimulus–response data that were not used to fit the model parameters). The goodness of fit of this predicted response is quantified by comparing it with the actual response of the neuron. In the case where the neural response can be described by a time-varying mean firing rate, the rate predicted by the model would be compared with the actual rate of the neuron. Note that this time-varying mean rate assumption assumes additive noise in a neuron's response, but does not place any restrictions on the nature of the stimulus–response model. It only requires that the model (linear or nonlinear) predicts a time-varying rate as the neuronal response to stimuli.

In practice, actual mean firing rate responses can only be estimated from PSTHs constructed from limited size noisy spike data. Until recently, researchers have mostly validated the stimulus–response models using the CC and coherence between predicted rates and actual mean estimated from the PSTH (e.g. Eggermont *et al* 1983, Dan *et al* 1996, Theunissen *et al* 2000, 2001). It is not clear whether deviations from unity in these CC and coherence values arise from inaccuracies of the stimulus–response model's predicted rates or from the inaccurate estimate of the actual mean rate from the noisy PSTH. A solution is to not compare the predicted rate with the actual rate directly. Instead we use our method (equation (8)) to estimate the expected CC between a spike train and its *actual* mean rate without explicitly estimating the mean rate. Because of the noise in the neuronal spike train this expected CC will be less than 1. We then also estimate the expected CC between a spike train and the predicted rate from the model. If the predicted rate is exactly equal to the mean rate, the two expected CCs (spike train with actual rate and with predicted rate) should be identical. In other words, the model fully describes the stimulus–response function of the neuron, if the ratio between these two values equals 1.

We can use both coherence and the CC to quantify the goodness of fit of model predictions. First, we calculate the expected coherence or CC between each spike train and the inferred mean rate in the validation set, using equation (8). These can be thought of as the ceiling values for the coherence and CC that we then obtain from the predictions of the stimulus–response model. Because neural data are inherently noisy, these values are often far from unity. We then compare these ceiling values to the coherence or CC values that are calculated between individual spike trains and the time-varying mean rate predicted by the model. For that second calculation, we use equation (11) with B equal to the predicted rate function. The closer the predicted rate to the actual mean rate, the closer the ratio of these two sets of CCs (CC ratio) and coherence values (coherence ratio) will be to unity.

As mentioned above, the calculation of the CC requires a choice of a smoothing window. We use the following procedure to choose a smoothing window in a principled fashion. First, we calculate the CC for a wide range of windows. For the calculation of the CC inherent in the spike data, both halves of the PSTH are smoothed with the same window. Thus, as the two halves are continually smoothed, they will approach a constant mean value in time, and their CC will approach 1. For the second calculation, involving equation (11), only the spike trains are smoothed and the predicted mean rate is not. Thus, the CC between the spike trains and predicted mean will reach a peak value and then decrease as the smoothing window increases. We take note of two CC values, one at the maximum CC ratio, and one at a constant smoothing window. The constant smoothing window is imperative for the comparison of CCs across different cells since the absolute quality of fit as described by CCs and CC ratio depends on the smoothing window. A good choice for the fixed window for an ensemble of cells would be given by the average window of the maximum CC ratios obtained in the data set.

Alternatively, one might want to report CC ratios for a reasonable range of time windows in order to compare data sets obtained in different laboratories and different systems.

6. An example linear model: the STRF

Spatio-temporal receptive fields (STRFs) are one type of stimulus–response model. STRFs are linear filters used to estimate the stimulus–response characteristics of neurons. Our method of estimating STRFs uses a generalized linear regression formula that is normalized for autocorrelations in the stimulus (Theunissen *et al* 2001). In this section, we will apply the model validation techniques described above to quantify the goodness of fit of the predictions obtained from STRFs in our laboratory. Note, however, that the same techniques can be applied to validate nonlinear stimulus–response models.

We first applied our methods of STRF validation to the hypothetical idealized situation of a perfectly linear neuron. We chose an arbitrary auditory STRF and a set of natural sound stimuli, which consisted of twenty different zebra finch songs. We used this STRF to define the stimulus–response transfer function of our idealized linear model neuron. We then obtained the time-varying mean firing rate of our model neuron by convolving the spectrograms of our zebra finch songs with this STRF. From this mean rate, we created model spike train data by generating multiple trials of spike trains using a Poisson process. Then, we proceeded as with real neural data, and estimated the model neuron's STRF from our stimuli and model data using the methodology described in Theunissen *et al* (2001). Figure 3 shows the original STRF that we used for generating model data, the STRFs estimated from these data and estimates of the validation measures. These estimates are shown using 10, 50 and 100 trials of model data. The non-normalized coherence and CC values are much less than 1. As mentioned above, these small CC and coherence values result from comparing the mean firing rate with a single spike train. On the other hand, the coherence and CC ratios for this perfectly linear model neuron are very close to 1 and indicate the accuracy of the model. The CC and coherence ratios are not exactly 1 because the predicted rate is obtained from an estimated STRF, which deviates slightly from the original STRF used to generate the data. The noise in the predicted STRF arises from the limited data size. The validation measures improve with data size as expected. For comparison, we also show in figure 3 the CC obtained between the prediction and the PSTH, which is the common validation technique (dot-dashed line). The estimate of the goodness of fit is significantly lower because of the noise in the PSTH. As can be seen, the noise in the PSTH significantly lowers the CC and requires longer smoothing time windows. Since the actual information transmitted is proportional to both absolute value of the CC and the inverse of the smoothing time window, the common validation technique underestimates the predictive power of the model when used with limited data size.

Finally, we applied the same procedures to the validation of an STRF calculation performed on actual auditory neurons of zebra finches. In these experiments ten spike trials were obtained in response to 20 different zebra finch songs, each approximately 2 s long. For this neuron, the maximum CC ratio is approximately 0.8 and the coherence ratio is above 0.5 for a large frequency range showing that the linear STRF is capable of capturing a significant portion of the neural response. For this neuron, the original validation measure based on the correlation between predicted rate and PSH is very close to the new measure of CC ratio (dot-dashed line of figure 4). The CC original validation measure yields smaller values than the CC ratio at the shorter smoothing windows, where more noise remains in the PSTH. At longer smoothing windows, less PSTH noise remains, but more signal is also removed by smoothing. Most of the neurons in our data have much smaller CC and coherence ratios than the neuron
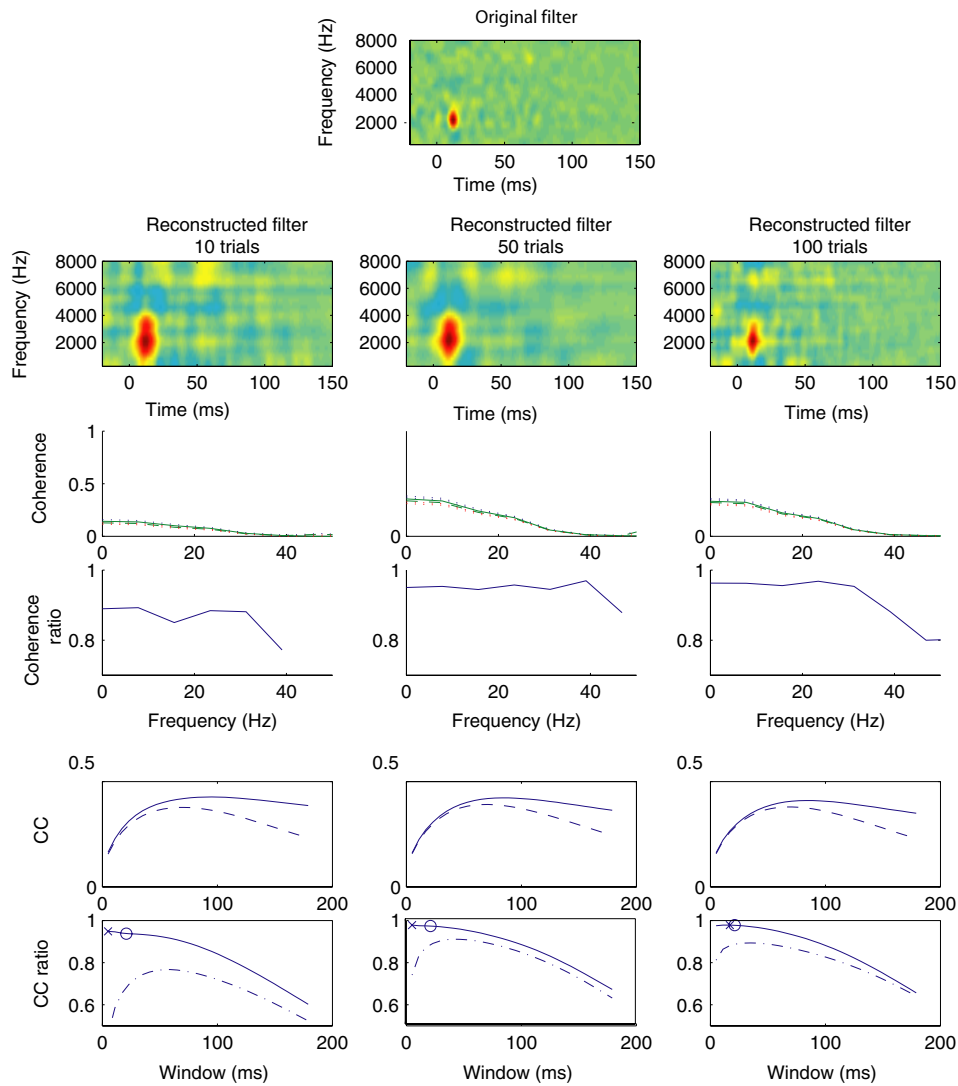
Figure 3. STRFs estimated from model data from a linear model neuron, with 10, 50 and 100 trials. The top image shows the STRF used to generate the mean rate function for the linear model neuron. The second row of images shows the STRFs estimated from the model data. The plots in the next two rows show the different coherence measures used in the validation of the STRF and the plots in the bottom two rows show the equivalent CC measures as a function of the width of the smoothing window. In the third row for the coherence plots and in the fifth row for the CC plots, the solid lines show the expected coherence and CC, respectively, between spike trains and their mean rate, γ_{AR}^2, calculated with equation (8). The dashed lines show the coherence and CC, respectively, between the spike trains and the predicted mean rate, $\gamma_{B\bar{R}}^2$, calculated by estimating $\gamma_{B\bar{R}_M}$ from the data and then using equation (11). The ratio of these two coherences or CC values is shown, respectively, in the fourth and sixth rows. In the CC plots, the '×' marks the smoothing window at which the ratio between the two CCs is maximum. This is the optimal smoothing window for this particular data set. The 'O' marks the CC ratio at the constant smoothing window of 21 ms, which our lab uses to compare across data sets. As expected, the coherence ratio or CC ratio is close to 1 for this linear but noisy model neuron. The CC between the smoothed PSTH and the predicted rate is shown as a dot-dashed line on the same plots as the CC ratio. This CC is significantly less than 1, and reflects error in estimating the actual mean rate from the PSTH as well as inaccuracies of the estimated STRF.
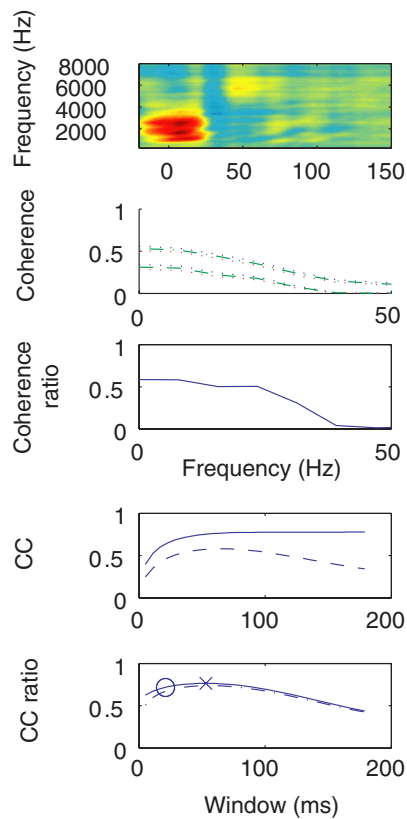
Figure 4. The STRF and the corresponding validation plots as shown in figure 3, calculated for a real neuron from the auditory midbrain of a zebra finch. The coherence ratio shows that, for this particular neuron, the linear STRF model captures slightly more than half of the response variance at low frequencies. The square of the maximum CC ratio ($0.75^2 \simeq 0.56$) obtained with a smoothing window of 50 ms gives the same result. The squared CC for a given smoothing window is related to the coherence averaged for the low frequencies up to a frequency cutoff approximately given by $1/(2 \times$ width of the smoothing window), or 10 Hz for a 50 ms window. As in figure 3, the CC between the smoothed PSTH and the predicted rate is shown with a dot-dashed line on the same plot as the CC ratio. For this neuron, the noise in the PSTH estimation is relatively small compared to the noise in the STRF as can be seen by the only small difference between the CC ratio and the CC between predicted rate and PSTH. The noise in PSTH estimation is greater at higher frequencies, as seen by the greater difference at smaller smoothing windows, where high frequencies are preserved.

example in figure 4. In figures 5 and 6, we compare the CC ratio measure to the original validation measure for a representative subset of neurons in the avian forebrain (analyzed in Theunissen *et al* 2001). In general, the CC ratios obtained are higher than the CC between the predicted and PSTH used in the original validation method but these values are still much smaller than 1. As shown by the modeling example above, even with only ten trials, the CC and coherence ratios for linear neurons should be near unity. We can therefore conclude that our linear STRF model is unable to completely characterize the neuronal response for any of the high-level auditory neurons in our data set. Our validation techniques enabled us to separate the influence of noise from the limitations in our stimulus–response model.
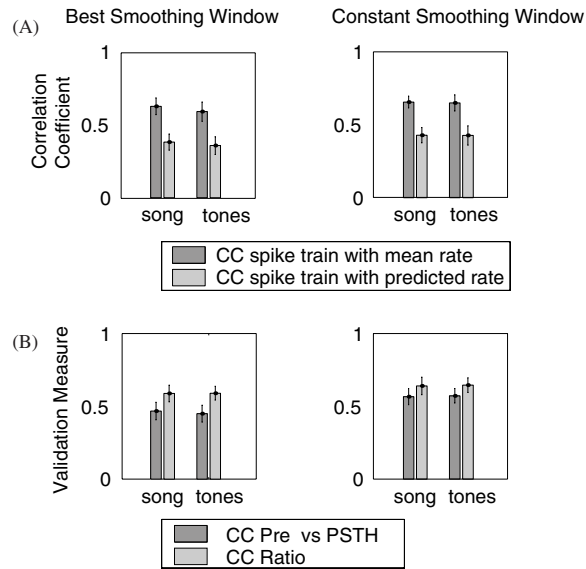
(A)

Best Smoothing Window   Constant Smoothing Window

(B)

Figure 5. (A) Expected CC between a spike train and the mean firing rate and expected CC between a spike train and the predicted rate obtained from the STRF of a representative set of avian auditory forebrain neurons. For each neuron STRFs were estimated using bird-song and pure tone pips stimuli, as explained in more detail in Theunissen *et al* (2000) and Theunissen *et al* (2001). The left panel shows these values for the best smoothing window (resulting in the highest CC) obtained independently for each neuron. The right panel is for a constant smoothing window of 21 ms for all neurons. (B) Validation of the linear STRF model using CC ratio (the ratio of the CCs shown in panel A). This proposed validation method is compared to the CC between PSTH and the predicted rate. The same population of neurons as in (A) is used. The new validation numbers are slightly higher and reflect the portion of the neurons' responses that can be described by the linear STRF model. In general, the STRF is able to predict about half of the neuron's response.
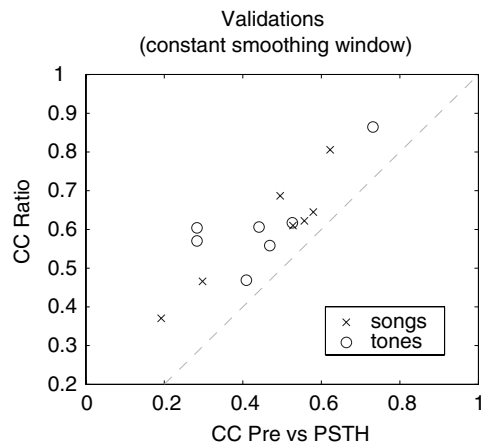

Figure 6. Scatter plot of CC ratio (proposed validation quantifier) versus CC between PSTH and predicted rate (original validation method) for the predictions obtained from the STRFs for a representative population of avian auditory forebrain neurons. For all neurons, the proposed validation measures are higher than the previously used measures and are more accurately quantifying the fraction of a neuron's response that is captured by the stimulus–response model.

7. Conclusion

The presence of neural noise combined with limited data size makes the separation of the deterministic from the stochastic part of a neural response a difficult task. In particular, under the assumptions of rate coding, the deterministic part of the response cannot be obtained reliably by simple averaging of experimental data with fewer than 100 trials. To address this issue, we have proposed a re-sampling method for the estimation of the coherence and correlation coefficient that uses the mean rate implicitly by taking advantage of the properties of independent and additive noise. With this method we are able to quantify the signal to noise ratio in a neural response by the coherence or CC between a single spike trial and the actual mean rate. When the neuronal response is not described by a rate code, these measures would not apply. However, under the rate coding assumptions, these quantifiers provide a measure of the reliability in neural responses and a measure of the information capacity of the neural channel when there is not enough data to calculate the mutual information (Strong *et al* 1998). When comparing the overall coherence amongst different neurons and stimuli, it is convenient to summarize the coherence with one number, such as the normal mutual information. If enough data are available, this normal mutual information can be compared to the direct mutual information and assessments can be made about the normality of the response distribution and the descriptiveness of the rate coding assumption. The same methodology can also be used to validate neural response models. The ratio of coherence and CCs provide bounded measures that separate the effects of limited noisy data from the effects of limitations in the model itself.

Acknowledgments

Appendix A. Estimation of coherence from PSTH

The coherence of a spike train and signal is given by

$$\gamma_{AR}^2(\omega) = \frac{\langle R(\omega)A^*(\omega)\rangle\langle R^*(\omega)A(\omega)\rangle}{\langle A(\omega)A^*(\omega)\rangle\langle R(\omega)R^*(\omega)\rangle} = \frac{A^2(\omega)}{A^2(\omega) + N^2(\omega)}.$$

For two halves of a PSTH $\overline{R}_{1,M/2}^2(\omega)$ and $\overline{R}_{2,M/2}^2(\omega)$ where $\overline{R}_{1,M/2}^2(\omega) = A^2(\omega) + N^2(\omega)2/M$ and $\overline{R}_{2,M/2}^2(\omega) = A^2(\omega) + N^2(\omega)2/M$, the coherence between them is

$$\gamma_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}^2(\omega) = \frac{\langle \overline{R}_{1,M/2}(\omega)\overline{R}_{2,M/2}^*(\omega)\rangle\langle \overline{R}_{1,M/2}^*(\omega)\overline{R}_{2,M/2}(\omega)\rangle}{\langle \overline{R}_{1,M/2}(\omega)\overline{R}_{1,M/2}^*(\omega)\rangle\langle \overline{R}_{2,M/2}(\omega)\overline{R}_{2,M/2}^*(\omega)\rangle} = \frac{A^2}{A^2 + \frac{4N^2}{M} + \frac{4N^4}{M^2A^2}}$$

(ω from now on will be assumed).

Define

$$\beta = \frac{N^2}{A^2}.$$

So now we have

$$\beta = \frac{1}{\gamma_{AR}^2} - 1 \qquad \text{and} \qquad \gamma_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}^2 = \frac{1}{1 + \frac{4\beta}{M} + \frac{4\beta^2}{M^2}}.$$

Solving for β:

$$\left(\gamma^2_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}} - 1\right) + \left(\frac{4}{M}\gamma^2_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}\right)\beta + \left(\frac{4\gamma^2_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}}{M^2}\right)\beta^2 = 0$$

$$\beta = \frac{1}{2}\left(-M + M\sqrt{\frac{1}{\gamma^2_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}}}\right) = \frac{1}{\gamma^2_{AR}} - 1.$$

Appendix B. Ratio of coherence with *M* trials to one trial with arbitrary function

Again, define

$$\beta = \frac{N^2}{A^2}.$$

So now we have

$$\beta = \frac{1}{\gamma^2_{AR}} - 1 \qquad \text{and} \qquad \frac{\beta}{M} = \frac{1}{\gamma^2_{A\overline{R}_M}} - 1.$$

From appendix A we have

$$\beta = \frac{1}{2}\left(-M + M\sqrt{\frac{1}{\gamma^2_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}}}\right) = \frac{1}{\gamma^2_{AR}} - 1,$$

which means that

$$\gamma^2_{AR} = \frac{2}{-M + M\sqrt{1/\gamma^2_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}} + 2}.$$

Using the relationship between β and β/M, we also have

$$\gamma^2_{A\overline{R}_M} = \frac{2}{1 + \sqrt{1/\gamma^2_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}}}.$$

Thus the ratio:

$$\frac{\gamma^2_{AR}}{\gamma^2_{A\overline{R}_M}} = \frac{1 + \sqrt{1/\gamma^2_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}}}{-M + M\sqrt{1/\gamma^2_{\overline{R}_{1,M/2}\overline{R}_{2,M/2}}} + 2}.$$

References

Baddeley R *et al* 1997 Responses of neurons in primary and inferior temporal visual cortices to natural scenes *Proc. R. Soc. Lond.* B **264** 1775–83

Bair W and Koch C 1996 Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey *Neural Comput.* **8** 1185–202

Baker C L Jr *et al* 1991 Optimal spatial displacement for direction selectivity in cat visual cortex neurons *Vis. Res.* **31** 1659–68

Barbieri R *et al* 2001 Construction and analysis of non-Poisson stimulus-response models of neural spiking activity *J. Neurosci. Methods* **105** 25–37

Borst A 2003 Noise, not stimulus entropy, determines neural information rate *J. Comput. Neurosci.* **14** 23–31

Borst A and Haag J 2001 Effects of mean firing on neural information rate *J. Comput. Neurosci.* **10** 213–21

Borst A and Theunissen F E 1999 Information theory and neural coding *Nat. Neurosci.* **2** 947–57

Brillinger D R 1975 *Time Series: Data Analysis and Theory* (New York: Holt, Rinehart and Winston)

Dan Y *et al* 1996 Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory *J. Neurosci.* **16** 3351–62

de Ruyter van Steveninck R R *et al* 1997 Reproducibility and variability in neural spike trains *Science* **275** 1805–8

Eggermont J J *et al* 1983 Quantitative characterisation procedure for auditory neurons based on the spectro-temporal receptive field *Heart Res.* **10** 167–90

Gabbiani F 1996 Coding of time-varying signals in spike trains of linear and half-wave rectifying neurons *Network* **7** 61–85

Haag J and Borst A 1997 Encoding of visual motion information and reliability in spiking and graded potential neurons *J. Neurosci.* **17** 4809–19

Johnson D H 1996 Point process models of single-neuron discharges *J. Comput. Neurosci.* **3** 275–99

Kara P, Reinagel P and Reid R C 2000 Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons *Neuron* **27** 635–46

Lau B, Stanley G B and Dan Y 2002 Computational subunits of visual cortical neurons revealed by artificial neural networks *Proc. Natl Acad. Sci. USA* **99** 8974–9

Lestienne R 2001 Spike timing, synchronization and information processing on the sensory side of the central nervous system *Prog. Neurobiol.* **65** 545–91

Lewen G D *et al* 2001 Neural coding of naturalistic motion stimuli *Network* **12** 317–29

Marmeralis P and Marmeralis V 1978 *Analysis of Physiological Systems. The White Noise Approach* (New York: Plenum)

Oram M W *et al* 1999 Stochastic nature of precisely timed spike patterns in visual system neuronal responses *J. Neurophysiol.* **81** 3021–33

Reinagel P and Reid R C 2000 Temporal coding of visual information in the thalamus *J. Neurosci.* **20** 5392–400

Sahani M and Linden J F 2002 How linear are auditory cortical responses *NIPS 2002 Proc.* at press

Shadlen M N and Newsome W T 1998 The variable discharge of cortical neurons: implications for connectivity, computation, and information coding *J. Neurosci.* **18** 3870–96

Shannon C E and Weaver W 1963 *The Mathematical Theory of Communication* (Chicago, IL: University of Illinois Press)

Strong S P *et al* 1998 Entropy and information in neural spike trains *Phys. Rev. Lett.* **80** 197–200

Svirskis G and Rinzel J 2000 Influence of temporal correlation of synaptic input on the rate and variability of firing in neurons *Biophys. J.* **79** 629–37

Theunissen F and Miller J P 1995 Temporal encoding in nervous systems: a rigorous definition *J. Comput. Neurosci.* **2** 149–62

Theunissen F E *et al* 2000 Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds *J. Neurosci.* **20** 2315–31

Theunissen F E *et al* 2001 Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli *Network* **12** 1–28

Thomson D and Chave A 1991 Jackknifed error estimates for spectra, coherences, and transfer functions *Advances in Spectrum Analysis and Array Processing* ed S Haykin (Englewood Clifs, NJ: Prentice-Hall)

Victor J D and Knight B W 1979 Nonlinear analysis with an arbitrary stimulus ensemble *Q. Appl. Math.* **37** 112–36

Wehr M and Laurent G 1996 Odour encoding by temporal sequences of firing in oscillating neural assemblies *Nature* **384** 162–6