

Identifying positive roles for endogenous stochastic noise during computation in neural systems

Mark D. McDonnell¹, Lawrence M. Ward²

Abstract—Information processing in nonlinear systems can sometimes be enhanced by the presence of stochastic fluctuations, or noise. Although the electrical properties of neurons and synapses are known to be influenced by intrinsic stochastic variability, it remains an open question as to whether living systems exploit this noise during neuronal information processing. This is despite various forms of noise-enhanced processing, such as classical stochastic resonance, having been observed in mathematical models of neural systems and in data acquired experimentally. We recently argued that advancing our understanding of the potential roles of random noise in assisting neuronal information processing will require specific focus on a concrete hypothesis about the computational roles of a specific neural system that can then be tested experimentally using signals and metrics relevant to the hypothesis. In this invited symposium paper, we argue why most existing approaches to studying stochastic resonance based on classical definitions and methods are highly limited in their applicability, since they impose an implied computational hypothesis that may have little relevance for real neurobiological systems.

I. INTRODUCTION AND BACKGROUND

Intrinsic stochastic noise is ubiquitous in biological systems, and in particular manifests in numerous ways in brains and nervous systems, at all scales from molecules to neurons to networks to whole brain [1], [2]. Examples include probabilistic release of synaptic neurotransmitter, and fluctuations in neuronal membrane potentials due to stochastic opening and closing of ion channels.

As reviewed in [3], there are many ways in which a specific neural system has been observed to better achieve a putative computational goal in the presence of random fluctuations originating from stochastic biologically relevant noise, than in their absence. Although this defies intuition, it is its interaction with nonlinearities that enables noise to have beneficial effects—sometimes noise can diminish detrimental properties of nonlinearities [4].

We proposed in [3] that whenever noise is observed to have a positive effect on neuronal information processing ‘stochastic facilitation’ can be said to be observed. Although such effects have been observed empirically, there has not yet been any confirmation that intrinsic noise is actually exploited in-vivo, since the noise has usually been artificially

introduced rather than intrinsic [5], [4]. Confirming in-vivo stochastic facilitation may require experimentally changing the properties of biologically relevant noise and measuring the resulting changes in the efficacy of the computation, and achieving this is clearly challenging.

We also proposed in [3] a framework consisting of six sequential steps that future experimental and computational neuroscience approaches might follow in order to aid in identifying new and interesting forms of stochastic facilitation. This framework makes explicit the importance of commencing such studies with a concrete computational hypothesis that is relevant to a specific neural system. It also emphasises the need to make biologically appropriate choices with regard to stimulation of the system, and to use analysis methods relevant to the computational hypothesis.

We proposed this framework because many studies of the potential benefits of stochastic noise in neural systems focus solely on stochastic resonance, and its classical definition [6]. Such an approach, however, predetermines potentially inappropriate choices for stimulation and analysis, because it imposes a restrictive computational hypothesis that is unlikely to be functionally relevant in most neural systems. We argue that to better assess whether stochastic facilitation occurs in a neural system, a biologically appropriate computational role of the system should be identified or proposed as a first step, along with a biologically relevant indicator of performance.

In order to elucidate these arguments, in the following section we discuss classical stochastic resonance, and contrast it with our proposed framework for studying stochastic facilitation generally. Next in Section III we state explicitly how signal-to-noise ratio is defined in classical stochastic resonance, in order that in Section IV we can discuss why this metric imposes an implied computational hypothesis, and argue that its relevance for real neurobiological systems is questionable.

II. CLASSICAL STOCHASTIC RESONANCE VS STOCHASTIC FACILITATION

A. Definition of classical stochastic resonance

In classical stochastic resonance [6], a periodic signal arrives as an input to a non-linear dynamical system. This signal is assumed to be ‘weak’ and/or ‘sub threshold’ such that the system provides an output from which detection of the periodic signal is difficult when noise is absent. Classical stochastic resonance is observed when noise is present and it enables the input signal to be detected statistically, in a better way than when noise is absent. In classical stochastic resonance, the quality of detection is measured by an output

The work of Mark D. McDonnell was supported by an Australian Research Fellowship funded by the Australian Research Council (project number DP1093425), and that of L.M. Ward by NSERC of Canada.

¹M. D. McDonnell is with the Computational and Theoretical Neuroscience Laboratory, Institute for Telecommunications Research, University of South Australia, Mawson Lakes, SA, Australia. mark.mcdonnell@unisa.edu.au

²L. M. Ward is with the Department of Psychology and Brain Research Centre, University of British Columbia, Vancouver, Canada

signal-to-noise ratio (SNR) calculated using the spectral content (power spectral density (PSD)) of the response. Typically, the SNR exhibits a single peak as the power of the noise is varied. Figure 1a illustrates these points.

This classical definition mandates the form of the input signal and the performance metric. Together these two essential features are highly restrictive, both in terms of the relevance of the signal (most biological signals are not strictly periodic, even though many are rhythmic), and the metric, as signal-to-noise ratio has many problematic features, which we discuss in depth below.

B. Six steps for studying stochastic facilitation

Nonclassical variants of stochastic resonance have discarded the requirements of periodic signals and SNR (see [4] for a review), and weak subthreshold signals have been shown to be unnecessary for a simple network of neurons [7], [8]. Also, in these studies metrics other than SNR are used. This makes it clear that interesting facilitative effects of noise can occur that are in many ways dissimilar to classical stochastic resonance. Even so, many papers have appeared in the neuroscience literature recently that take note only of classical stochastic resonance. We aimed to emphasise how restrictive this is, and put forward the following six step framework as a way to better elucidate whether intrinsic stochastic noise in neural systems may be exploited during neuronal information processing (see Figure 1b):

- 1) State a hypothesis re the positive role of stochastic biological noise in facilitating signal processing or a computational task of a specified neural system.
- 2) Specify a neural preparation or mathematical or computational model that can be stimulated by inputs relevant to the hypothesis and produce output responses that can be measured.
- 3) Choose hypothesis-relevant input signals (if necessary for the hypothesis) and noise that can be generated and introduced into, or deleted from, the experimental material or model.
- 4) Acquire relevant output data after introducing the chosen input signals and noise into the experimental rig or simulation of the model,
- 5) Process the output data into a form relevant for assessing the computational hypothesis
- 6) Assess the hypothesis that noise has a positive role, based on the processed data.

Classical stochastic resonance does not follow these steps in this sequence, because the signal in Step 2, and the performance metric in Step 6 are mandated and a computational hypothesis follows only by implication.

In the remainder of this paper we explore the consequences of the computational hypothesis implied in classical stochastic resonance, and argue that unless a model is put forward by which a neurobiological system actually uses the output of the studied system in the manner implied by the processing needed in Step 5, observing classical SR does not mean it is relevant a real neural system.

III. SIGNAL-NOISE-RATIO IN CLASSICAL STOCHASTIC RESONANCE

Any definition of signal-to-noise ratio (SNR) requires definitions of signal power and noise power. Demonstration of classical stochastic resonance requires calculation of the output SNR from the nonlinear system under consideration, i.e. we need to define an output signal, and output noise.

The input signal must be a single frequency sine-wave signal with frequency f_0 Hz. Since all input power is located at a single frequency, the output signal power is defined as the output power spectral density (PSD) at f_0 Hz, and we denote this as $S(f_0)$.

The putative utility of measuring output SNR for a single frequency input signal, is to use it to decide whether that frequency is present in the input. It is the values of the output PSD at frequencies *in the vicinity of the signal frequency*, as measured with the signal *present*, that determines this [6]. The average value of the PSD near the signal frequency, but exclusive of PSD at that frequency, is known as the *noise floor*, and we denote this as $N(f_0)$. Using these definitions, SNR can (after conversion to decibels) be written as

$$\text{SNR} = 10 \log_{10} \left(\frac{S(f_0)}{N(f_0)} \right). \quad (1)$$

In practice, the quantities $S(f_0)$ and $N(f_0)$ are estimated from discrete time sampled data generated by simulations or recorded from experiments. Standard algorithms, such as the periodogram or Welch methods, which make use of the Fast Fourier Transform (FFT) algorithm with N frequency bins, can be used to obtain estimates for the PSD of the output. After obtaining such estimates, the signal power at the output is typically ‘smeared’ across several frequency bins, and a good estimate of $S(f_0)$ involves summing the PSD estimate over several bins centred on f_0 . Similarly, a good estimate of $N(f_0)$ requires averaging the PSD across a range of bins either side of the bins used to estimate $S(f_0)$.

If the system being studied is a neuron model, often it is only the action potential times that are of interest when calculating the SNR in studies of classical stochastic resonance. In this case, a point process may be created from the times at which action potentials occur. The PSD is then estimated after low pass filtering this point process, and ensemble averaging over many trials enables more accurate PSD estimates.

IV. WHY SIGNAL-NOISE-RATIO IS PROBLEMATIC FOR REAL NEURAL SYSTEMS

We argue that there are two main problems in studying classical stochastic resonance in neural systems:

- 1) Rather than first stating a hypothesis regarding a computational role as in the six-step approach we advocate, the choice of SNR as a metric applied to neural systems imposes an implied hypothesis, and this hypothesis lacks plausibility.
- 2) Assessing performance using SNR is only suggestive of *potential* stochastic facilitation, and it could be that the actual system does not produce a response where

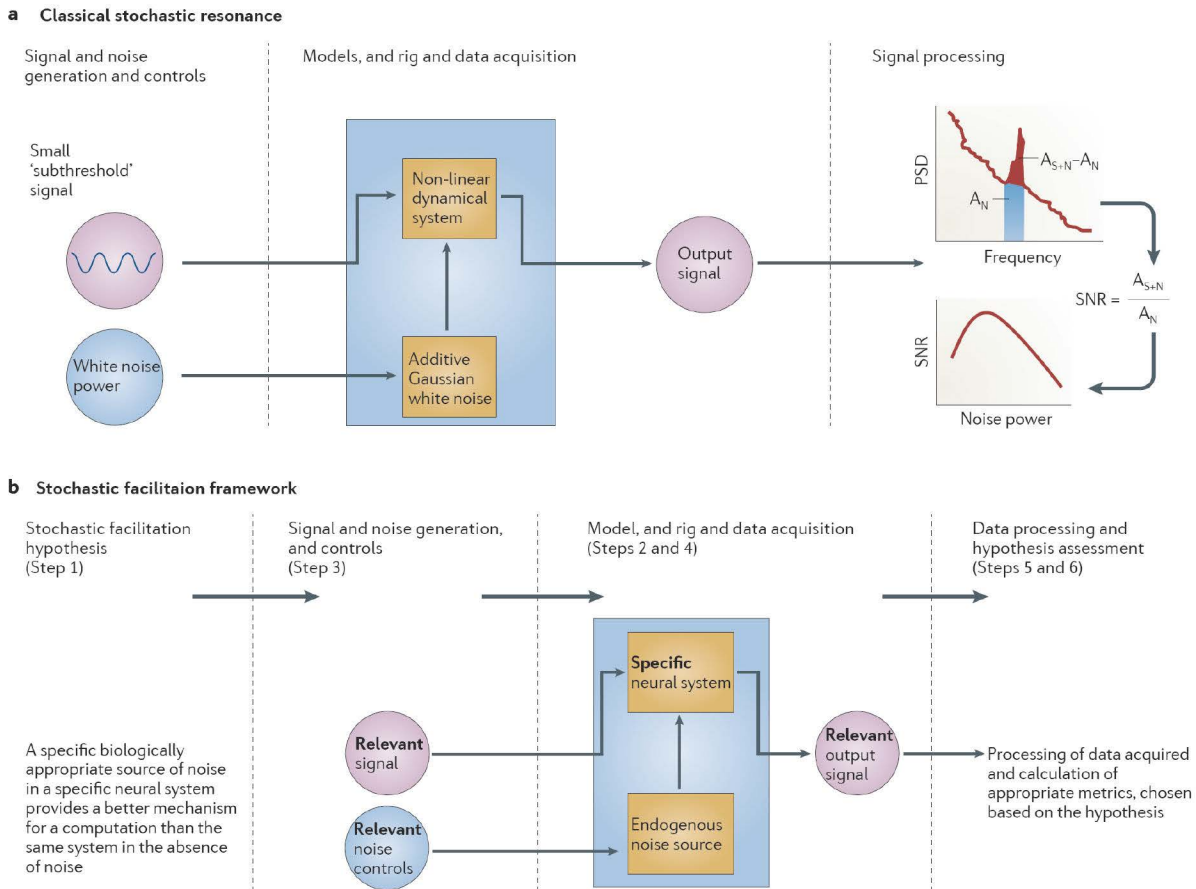


Fig. 1. Comparison of classical stochastic resonance with our six step framework for studying stochastic facilitation. Figure first published in Nature Reviews Neuroscience 12, pp. 415-426 (July 2011), doi:10.1038/nrn3061 [3], Nature Publishing Group, a division of McMillan Publishers Limited

the potential benefits of noise have any consequences relative to the absence of noise, or relative to sub-optimal noise.

We now illustrate these points by describing a typical scenario, which is to begin by choosing a model of a single spiking neuron (step (2) in our framework), and then choosing signals (step (3)) with the following properties:

- a periodic single frequency sine wave as an input signal;
- additive white noise, that randomly perturbs the sine wave;
- the periodic signal has an amplitude such that if the noise is absent, the neuron model does not produce action potentials;
- the range of possible noise levels is such that action potentials are induced to occur that otherwise would not have, given the signal amplitude.

The elements in this example are a perfect choice for an essential requirement in classical stochastic resonance, namely that

- performance as a function of noise intensity is measured by SNR, as in Equation (1), based on the output power spectral density of the stochastic process defined by

action potential timings.

Rather than first stating a hypothesis regarding a computational role, this choice of assumption that SNR should be used to measure performance (step (5)), reverses the conceptual sequence, and imposes an implied hypothesis (step (1)), that may be stated as follows.

- The computational role of the neuron is to produce a sequence of action potentials when a sinusoidal input current at a specific frequency excites the cell, and to produce a statistically distinct pattern of action potentials in the absence of the sinusoidal input current.

If we assume that this computational role is correct, then it suggests that some neurobiological mechanisms must be available for extracting information from the patterns of action potentials that is closely related to that extracted by estimating SNR.

Spectral based SNR can be used to determine the presence of a sinusoid, based on measured data, if there is spectral power at a single frequency that is clearly larger than the noise floor at other frequencies in the vicinity of that single frequency. The SNR is the ratio of the power of that peak, to the power of the noise floor. The larger the SNR, the

more likely it is that the peak is really due to a sinusoid, and is not an artifact, and the more likely it is that the peak will be apparent when the data used to estimate the power spectral density are not ideal. So for our data processing purposes, measuring SNR implies that we should carry out the following:

- Estimate the PSD of the stochastic process consisting of the action potential timings of the neuron. If the estimate of the spectral power at a specified input frequency is larger, by some specified amount, than the noise floor at frequencies in the vicinity, then make the decision that the input signal is present.

If we carry out this algorithm using standard simulation methods, then it is simple to demonstrate that SNR will vary with noise level, and under the stated conditions for the model, the SNR will be largest for some optimal noise level, i.e. stochastic resonance occurs. Whereas this might be a good algorithm for achieving the computational goal, given the mechanisms available when using a digital computer, we propose two reasons why the information obtained may not be accessible in a neurobiological system.

First, the stated algorithm relies on some assumptions. PSD is a mathematical concept that is, by definition, an average quantity whose production requires infinite time. It can be estimated from data of finite duration, but long recordings are required for accurate estimates—the more action potentials, the more accurate the estimate. Furthermore, it is necessary to assume that the process being measured is stationary in time (e.g. the sine wave is either always present or never present, during the whole recording, and the noise statistics remain constant). It is not unusual to require all conditions to remain stationary for periods of the order of 10 seconds. The question of when is this an appropriate neural time scale is not usually addressed, and the time scales of dynamics involved in neuronal computation are generally thought to be much less than 10 seconds.

Second, estimation of PSD usually relies on applying the FFT algorithm to a sequence of stored samples of data, as well as other complicated processing. Whether there are neurobiological processes that can extract exactly the same information that we can extract using the FFT is unlikely, although mechanisms for related processing based on somewhat different SNR metrics may be possible (see below).

V. DISCUSSION

In summary, the choice of algorithm implied by the choice of SNR as a performance measure, as mandated by classical stochastic resonance, is potentially at odds with the mechanisms that are likely available to the neural systems that respond to the results of the hypothetical computation implied by use of SNR. Although our computer based data processing that tells us the SNR, and therefore the potential for stochastic facilitation as the noise varies, it does not necessarily follow that the processing that takes place in the real neural system enables enhanced detection of periodic signals as the noise varies.

Another way of putting it is that even though SNR does in a way measure the encoding capabilities of the neuron that receives the input signal and noise, a more relevant metric for stochastic facilitation would be one that additionally measures how well a second neural system can extract or decode the response of the first system.

On the other hand, demonstration of a large SNR at a range of biologically plausible levels of noise does not falsify the stated hypotheses, and perhaps does lend some support to them. We argue, however, that it is not sufficient to fully assess the hypotheses. The above scenario also ignores several other points: why should the input frequency be known in advance, and why not compare the power spectral density with and without the signal?

So, what is an alternative algorithm? We could hypothesise that the neuron being considered has an axon that allows propagation of action potentials to a synaptic junction with another neuron, and consider the following algorithm.

- Exploit a cell that produces a pattern of action potentials when a sinusoidal signal at a certain frequency is present at its input, and a different pattern otherwise. The first pattern causes the downstream neuron to fire action potentials at a near constant rate. The second pattern causes no firing in the downstream neuron.

Perhaps it is feasible that the synapse and the second neuron provide neural mechanisms that can achieve this algorithm, given that there are various ways in which spectral filtering can occur in the responses of neurons [9], [10], [11]. Whether it is or isn't feasible, however, the point is that consideration of the neuronal mechanisms for carrying out an algorithm should be a necessary component of hypotheses regarding stochastic resonance in neural systems.

REFERENCES

- [1] A. A. Faisal, L. P. J. Selen, and D. M. Wolpert, "Noise in the nervous system," *Nature Reviews Neuroscience*, vol. 9, pp. 292–303, 2008.
- [2] A. Destexhe and M. Rudolph-Lilith, *Neuronal Noise*. Springer, New York, 2012.
- [3] M. D. McDonnell and L. M. Ward, "The benefits of noise in neural systems: bridging theory and experiment," *Nature Reviews Neuroscience*, vol. 12, pp. 415–426, 2011.
- [4] M. D. McDonnell and D. Abbott, "What is stochastic resonance? Definitions, misconceptions, debates, and its relevance to biology," *PLoS Computational Biology*, vol. 5, p. e1000348, 2009.
- [5] L. M. Ward, "Psychophysics of stochastic resonance," *Fluctuation and Noise Letters*, vol. 4, pp. L11–L21, 2004.
- [6] L. Gammaitoni, P. Hänggi, P. Jung, and F. Marchesoni, "Stochastic resonance," *Reviews of Modern Physics*, vol. 70, pp. 223–287, 1998.
- [7] N. G. Stocks, "Suprathreshold stochastic resonance in multilevel threshold systems," *Physical Review Letters*, vol. 84, pp. 2310–2313, 2000.
- [8] M. D. McDonnell, N. G. Stocks, C. E. M. Pearce, and D. Abbott, *Stochastic Resonance: From Suprathreshold Stochastic Resonance to Stochastic Signal Quantisation*. Cambridge University Press, Cambridge, UK, 2008.
- [9] B. Hutcheon and Y. Yarom, "Resonance, oscillation and the intrinsic frequency preferences of neurons," *Trends in Neurosciences*, vol. 23, pp. 216–222, 2000.
- [10] E. M. Izhikevich, N. S. Desai, E. C. Walcott, and F. C. Hoppensteadt, "Bursts as a unit of neural information: selective communication via resonance," *Trends in Neurosciences*, vol. 26, pp. 161–167, 2003.
- [11] L. F. Abbott and W. G. Regehr, "Synaptic computation," *Nature*, vol. 431, pp. 796–803, 2004.