CellPress
REVIEWS

Opinion

Noisy Synaptic Conductance: Bug or a Feature?

Dmitri A. Rusakov [iD],[1],* Leonid P. Savtchenko,[1],* and Peter E. Latham[2],*

More often than not, action potentials fail to trigger neurotransmitter release. And even when neurotransmitter is released, the resulting change in synaptic conductance is highly variable. Given the energetic cost of generating and propagating action potentials, and the importance of information transmission across synapses, this seems both wasteful and inefficient. However, synaptic noise arising from variable transmission can improve, in certain restricted conditions, information transmission. Under broader conditions, it can improve information transmission per release, a quantity that is relevant given the energetic constraints on computing in the brain. Here we discuss the role, both positive and negative, synaptic noise plays in information transmission and computation in the brain.

Neuronal Communication is Noisy

Neurons communicate primarily through chemical synapses, and that communication is critical for proper brain function. However, chemical synaptic transmission appears unreliable: for most synapses, when an action potential arrives at an axon terminal, about half the time, no neurotransmitter is released and so no communication happens. Even strong and reliable synapses, such as the neuromuscular junction or the calyx of Held, operate by using multiple unreliable release sites [1]. Furthermore, when neurotransmitter is released at an individual synaptic release site, the size of the local postsynaptic membrane conductance change is also variable. Given the importance of synapses, the energetic cost of generating action potentials, and the evolutionary timescales over which the brain has been optimized, the high level of synaptic noise seems surprising. Here we ask: how surprised should we be? More to the point: is synaptic noise a bug, or is it a feature? Whilst the role of unreliable release on neural network performance has been explored in multiple studies (see [2–4] for informative reviews), here we expand on these views, in particular giving consideration to noisy **synaptic conductance** (see Glossary).

Noise in Nonlinear Communication Channels

When communicating information, noise would seem to be universally bad. Static on the radio, background noise during a conversation, grainy videos – they all make it harder to extract interesting information. However, for **nonlinear systems** – especially when a threshold is involved, as in spike generation – noise can be helpful [5–7]. Most notably, if a noise-free presynaptic signal is too small, it will never generate a spike and thus will not transmit any information. Add a small amount of noise, and, at least sometimes, a spike will occur. Thus, a subthreshold sensory signal has a better chance of being detected when noise is added [8,9]. This phenomenon is, very broadly, known as **stochastic resonance** [3,10,11], and it is observable in electrophysiological experiments: add a small amount of noise to **subthreshold synaptic currents**, and the probability of a spike can go way up [12]. A simple multisynaptic neuron model (Figure 1A) illustrates this point. With no noise in the synaptic conductance, there are very few spikes (Figure 1B), but adding a small amount of variability in the postsynaptic conductance can produce fairly reliable

Highlights

Synaptic transmission is noisy, partially because neurotransmitter release is unreliable but also because synaptic receptor current varies stochastically.

When synaptic input is below the cell spike threshold, adding noise may push it above the threshold, thus improving information transmission.

Synaptic noise can widen the dynamic range of transmission.

Synaptic noise allows synapses to communicate their degree of uncertainty.

Noise is a ubiquitous feature of brain activity and is likely to be a critical element in any theory of how the brain works.

[1]Queen Square UCL Institute of Neurology, University College London, Queen Square, London, WC1N 3BG, UK
[2]Gatsby Computational Neuroscience Unit, University College London, 25 Howland Street, London, W1T 4JG, UK

*Correspondence:
d.rusakov@ucl.ac.uk (D.A. Rusakov),
leonid.savtchenko@ucl.ac.uk
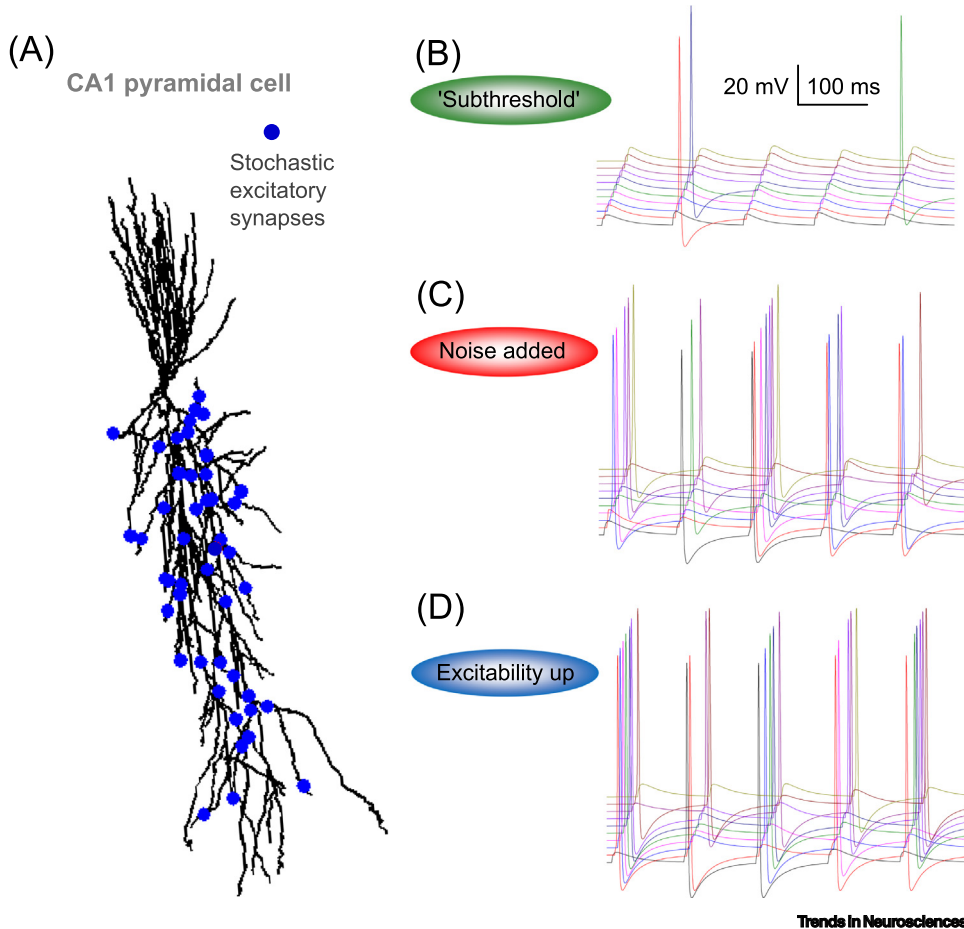(L.P. Savtchenko), and
pel@gatsby.ucl.ac.uk (P.E. Latham).

CellPress REVIEWS

Glossary

Bimodal: having two peaks in the distribution, so that typical values are either, say, high or low.

Nonlinear system: a system in which the relationship between input and output is not linear.

Release probability: average probability with which a synaptic vesicle is released upon the arrival of an action potential.

Stochastic: probabilistic; containing an element of uncertainty.

Stochastic resonance: a phenomenon in which adding noise increases the probability that a signal will reach threshold for detection.

Subthreshold synaptic currents: synaptic receptor-generated membrane current that is too weak to a neuron to fire.

Synaptic conductance: peak cell membrane conductance arising through the opening of postsynaptic receptor channels.

Synaptic drive: integrated effect of all active synaptic input to a neuron.

Figure 1. Increasing Synaptic Input Noise or Postsynaptic Cell Excitability Can Improve Signal Transfer: An Illustrative Example. (A) A basic multisynaptic neuron model: reconstructed CA1 pyramidal cell model equipped with known membrane mechanisms [57] (ModelDB accession number 2796, NEURON database). Fifty excitatory inputs (random scatter, blue dots) generate biexponential conductance change (rise and decay time, 0.1 ms and 1.2 ms, respectively), with the same onset but stochastically, in accord with their release probability, P_r; simulations with NEURON 7.2 [58], variable time step dt, $T = 34°C$. (B) Somatic voltage traces in response to synchronous activation of 50 unreliable synapses ($P_r = 0.5$), five stimuli at 10 Hz, no noise (conductance coefficient of variation, CV = 0), 'near-subthreshold' activation (peak synaptic conductance $G_s = 7.6$ pS; somatic sodium channel conductance $G_{Na} = 0.032$ mS/cm^2). (C) As in (B) but with synaptic conductance noise added (CV = 0.05) leading to increased spiking. (D) As in (B) but with cell excitability increased (somatic sodium channel conductance doubled to $G_{Na} = 0.064$ mS/cm^2). A fully functional model example and illustrative movie files can be downloaded from http://www.sciencebox.org/Neuroalgebra/TINSWeb/

spikes (Figure 1C). Note, however, that decreasing threshold for the generation of an action potential can produce even more reliable spikes (Figure 1D example has a higher somatic sodium conductance, making the cell more excitable compared to Figure 1B example). Although in a wider context we cannot equate more spiking with more information, we will, in the interest of simplicity, assume that is the case here.

Abstractly, the ability to communicate information depends on three things: the range of possible input signals, the noise around those signals, and the transformation from input to output (Box 1 provides a commonly accepted definition of information). For neurons, the transformation from input to output yields either a spike or no spike. In that case, communication is optimal when the integrated input to a neuron is either well above or well below threshold relative to the

Box 1. Defining (Mutual) Information

Spike trains convey information about presynaptic input. Intuitively, this means that if you were to observe spikes from a neuron, you would know more about its synaptic drive than if you did not observe spikes. To quantify this intuition, we need a way to measure knowledge – or, as is much more common, uncertainty. There are many ways to do this, but an especially convenient one is entropy, which is a relatively direct measure of uncertainty; the higher the entropy of a distribution over some variable (i.e., the broader the distribution), the more the uncertainty about that variable (according to Equations I and II as follows). With this choice, the difference in entropy is the Shannon mutual information [59], which indicates how much information about the input signal we can glean from the output signal (Figure I).

The entropy of a continuous distribution $P(x)$, often denoted $H(x)$, is

$$H(x) = -\int dx P(x) \, \log_2 P(x). \tag{I}$$

If $P(x)$ is discrete there is an analogous definition; the only difference is that the integral becomes a sum,

$$H(x) = -\sum_i P(x_i) \, \log_2 P(x_i). \tag{II}$$

To compute mutual information, we need to compare the entropy before and after some observation (which for us is a spike train). Let's use s for the observation, and $P(x \mid s)$ for the distribution of x (which for us is the synaptic drive) after observing a spike train. In that case, the conditional entropy, denoted $H(x \mid s)$ – the average entropy after observing a spike train – is

$$H(x|s) = -\int dx ds P(s) P(x|s) \, \log_2 P(x|s), \tag{III}$$

with integrals replaced by sums for discrete distributions. The mutual information, $I(x,s)$, between x and s is the difference between the two entropies,

$$I(x,s) = H(x) - H(x|s). \tag{IV}$$

This definition is sufficiently flexible that x could be continuous and s discrete, or vice versa. As is not hard to show, information is symmetric; it can also be written

$$I(x,s) = H(s) - H(s|x). \tag{V}$$

Symmetry is important because it tells us that the higher the entropy of a spike train (which typically translates to higher firing rate), the larger the potential for high information. Importantly, though, increasing the firing rate is not guaranteed to increase information, as that may also increase the second term.
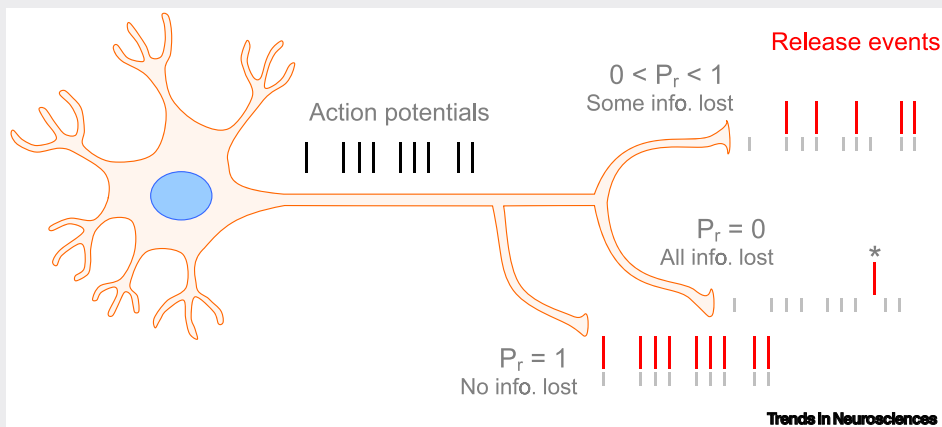

Figure I. Schematic Illustrating the Mutual Information between a Series of Axonal Spikes and the Corresponding Series of Neurotransmitter Release. This schematic illustrates the mutual information, I, between a series of axonal spikes (black bars; indicated with grey near the axonal terminals) and the corresponding series of neurotransmitter release (red bars). Note: this is different from the mutual information between the input and output because release does not guarantee an output spike. Information depends on the release probability, P_r, at each presynaptic site. The asterisk indicates spontaneous (action potential-independent) release.

noise. In such a regime, a spike provides unambiguous information about the input: it tells downstream neurons whether the input was high or low. But typically, the signal and noise cannot be adjusted independently; in particular, the **synaptic drive** to a neuron is not **bimodal** – it is graded. So, the simple, optimal, regime is not realised. Consequently, the question of whether synaptic noise is a bug or a feature depends on how the signal and noise interact. For instance, if increasing noise made synaptic drive more bimodal, it could, at least in principle, increase information transmission. And we have already seen that adding noise can increase the reliability of output spikes (Figure 1C). But we have also seen that increasing cell excitability, and thus lowering the threshold for a spike, can do the same thing (Figure 1D). Which one is better, and when?

Sources of Synaptic Noise

To investigate how synaptic noise affects information transmission in realistic situations, we need to understand the sources of noise. There are two main one. The first, and best studied, are release failures: when an action potential invades the presynaptic terminal, neurotransmitter may or may not be released [13,14]. The average **release probability**, P_r – calculated as 1 minus the probability of failure – ranges from 0.2 to 0.8 at central synapses [15–18], with experiments *in vivo* putting this towards the low end [19]. This generates considerable noise: if ten action potentials arrive at a cell approximately simultaneously, typically somewhere between two and eight of them will cause a release of neurotransmitter, and on any one trial, anywhere from zero to ten will cause neurotransmitter release.

The second source of noise is variability in the synaptic conductance (producing variability in the peak synaptic receptor current) in response to release of one neurotransmitter-filled synaptic vesicle. Fluctuations in the synaptic conductance depend on two main factors. First, synaptic vesicle content can vary depending on prior activity [20,21]. Second, synaptic receptor binding and activation is a **stochastic** process controlled by conformational changes of receptor proteins [22,23]. The lateral mobility of synaptic receptors [24,25] and, possibly, neurotransmitter release sites [26,27], have also been considered a factor. However, recent super-resolution studies have revealed a structural alignment of presynaptic release sites and postsynaptic receptor clusters [28], which should minimise this source of noise, at least on time scales of minutes [29]. Nonetheless, excitatory postsynaptic currents recorded at individual synapses in response to synaptic vesicle release, in basal conditions, vary significantly. A recent study suggested that the architecture of common excitatory synapses is such that it maximises synaptic conductance variability [30]. If so, it would appear that individual excitatory synaptic connections 'prefer' to be noisy.

An additional source of noise affecting synaptic circuit signalling is stochastic fluctuations of the cell membrane potential. Its multiple origins and its effect on performance have recently been reviewed in detail [31] and are outside the scope of this present opinion article.

How Does Synaptic Noise Affect Information Transmission?

Figure 1 shows that if the spike threshold is set too high, noise from variability in the conductance can improve information transmission. This is not true of release failures: lowering the probability of release decreases the drive to a postsynaptic cell. However, if we conserve transmitter expenditure, by increasing the conductance as we lower the release probability so that the product of conductance and P_r stays constant, the story changes. In this regime, the mean is independent of the release probability and the variance scales as $(1 - P_r)/P_r$. The variance thus increases as P_r decreases – exactly what we need to drive a cell when the threshold is set too high. Consequently, both failures and conductance variability could be beneficial for information transfer.

This is illustrated in Figure 2. Figure 2A shows that reliable synapses (P_r = 1) with constant conductance are perfectly efficient, but only when the overall input is above the threshold (compare left and right). In comparison, unreliable synapses (P_r < 1) with conductance scaled to conserve transmitter expenditure (as described earlier) can transfer more information when below threshold, but their performance is worse above threshold (Figure 2B). Not surprisingly, reliable synapses with variability in conductance will behave in a similar way to unreliable synapses with constant conductance (Figure 2C).

The picture that emerges from this analysis is that when the threshold is set too high, noise increases information transmission; otherwise, noise decreases information transmission (see Figure I in Box 2). This suggests that the optimal strategy is to reduce noise as much as possible,
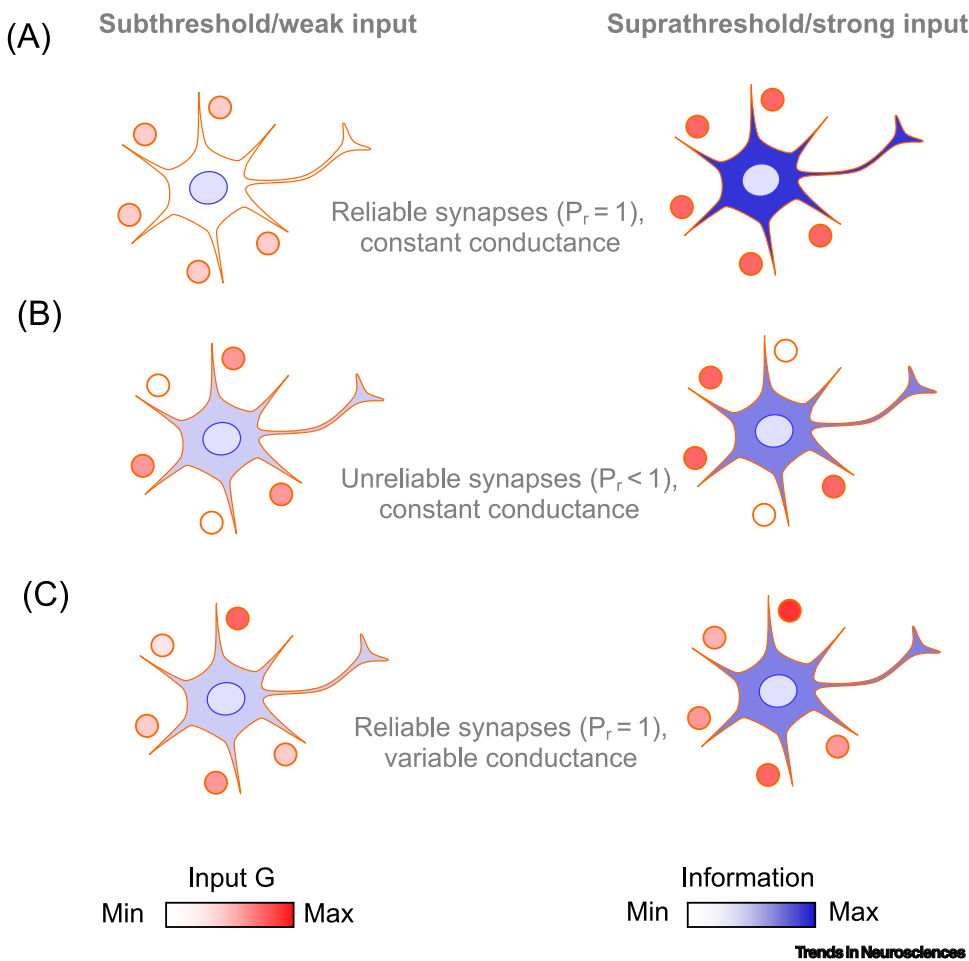

Figure 2. Noisy Synapses Can Be Informationally Advantageous towards the Lower End of the Transmission Dynamic Range. Synaptic inputs are shown by circles, with the intensity of red reflecting their activity. The amount of postsynaptic spiking activity is indicated by the shade of blue, with higher probability of spiking (darker blue) corresponding to higher information rates. (A) Reliable synapses (P_r = 1) and no variability in postsynaptic conductance. If the threshold is too high there are no spikes (left), but if the threshold is lowered slightly, spike transmission is 100% reliable (right). (B) Unreliable synapses (P_r < 1) and no variability in postsynaptic conductance, but conductance is proportional to $1/P_r$ (to conserve the transmitter expenditure). In this case, postsynaptic spiking is possible under weak input, thus transmitting some information (left). However, under strong input, unreliable synapses lose information [right; compare with (A)]. (C) Reliable synapses (P_r = 1), but with variability in postsynaptic conductance. Information transmission is about the same as with unreliable synapses with variability in postsynaptic conductance [compare with (B)].

Box 2. Computing (Mutual) Information

To gain an understanding of how mutual information depends on parameters, consider a simplified scenario in which there is a single presynaptic neuron that either does or does not fire, the output spikes tell us whether or not it fired, time is divided into small bins, say 10 ms, and there are either zero spikes or one spike in each bin.

In this scenario, it is not hard to show that the mutual information (Box 1) is given by

$$I(x,s) = h(qp) - ph(q). \qquad \text{[I]}$$

where p is the probability that the presynaptic neuron fires, q denotes the probability that the postsynaptic neuron emits a spike given that the presynaptic neuron fired, and $h(z)$ is the entropy of a Bernoulli random variable with probability z,

$$h(z) = -z \log_2 z - (1-z) \log_2(1-z). \qquad \text{[II]}$$

Assume the postsynaptic neuron fires when its input exceeds a threshold, θ, and that the input drive is Gaussian with mean μ and standard deviation σ. In that case, q is a cumulative normal function of $(\mu - \theta)/\sigma$. A plot of mutual information versus μ is shown in Figure I for various values of the noise (the thicker the line, the larger the noise). If the mean synaptic drive is above threshold, noise hurts, because it decreases the probability that a presynaptic spike will cause a postsynaptic one; if it is below threshold, noise helps, because it increases the probability that a presynaptic spike will cause a postsynaptic one. In addition, the higher the firing rate of the presynaptic neuron, the higher the mutual information (compare red and blue curves).

These observations suggest that mutual information is maximized when the input is bimodal [either well above or well below threshold, so that $q = 1$ in Equation II, which means $h(q)$ is zero], the noise is small, and the output firing rate is large. However, it is unlikely that these are achievable. For instance, presynaptic spikes are spread out over time. Consequently, the synaptic drive is a continuous function of time, and neurons spike as soon as the drive crosses threshold. Thus, synaptic drive is rarely bimodal. In addition, there are energetic constraints on firing rate, so it can't be arbitrarily large. Note, though, that in the brain, the input information is much larger than the output information – about 1000 times larger, since each neuron receives about 1000 inputs. Consequently, the information lost due to noise added by failures and variability in the size of postsynaptic currents is likely to be small compared with the information that is thrown away by the finite capacity of neurons, an argument made by Levy and Baxter [37].
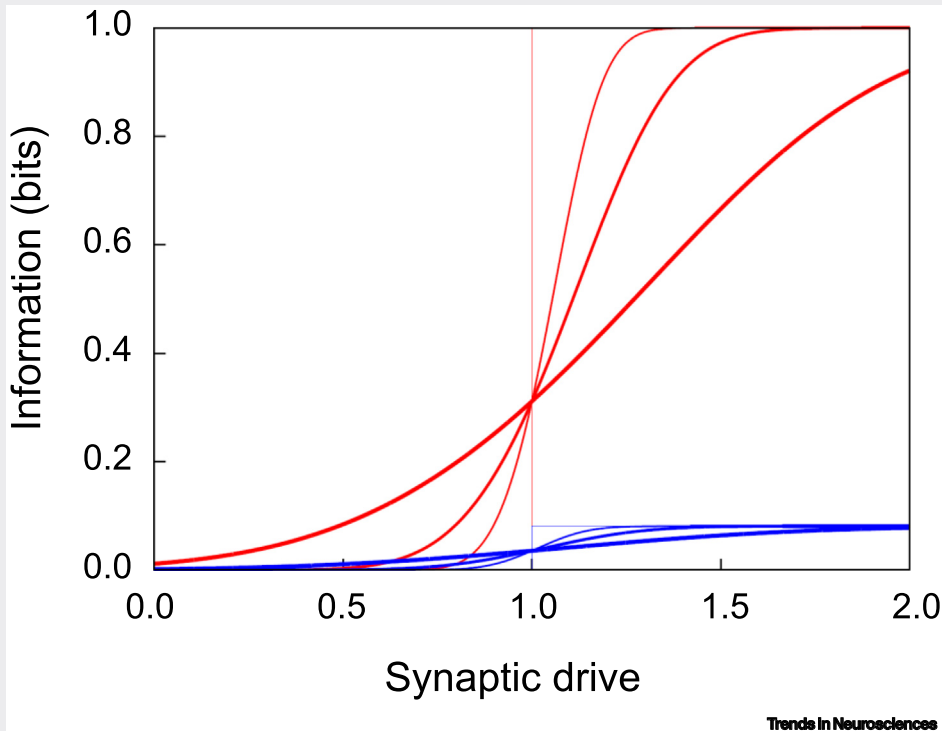

Trends in Neurosciences

Figure I. Information (in Bits) Versus Synaptic Drive, μ. Red: the probability of a presynaptic spike, p, is 0.5, corresponding to an effective presynaptic firing rate of 50 Hz, assuming a time bin of 10 ms. Blue: the probability of a presynaptic spike is 0.01, corresponding to an effective presynaptic firing rate of 1 Hz. Noise levels, σ, are 0.0, 0.1, 0.2, and 0.5 (thin to thick lines). The threshold, θ, was set to 1.0.

then set the threshold for spike generation to maximize information transmission. As we have seen, increasing release probability and reducing conductance variability both reduce noise. In addition, the nervous system has another option: increase the redundancy of synaptic connections – by allowing individual axons to make multiple synaptic connections on a postsynaptic cell [32]. This averages the noise and so partially compensates for synaptic unreliability [33,34]. This is, in fact, the strategy taken by reliable synapses, such as the neuromuscular junction or the calyx of Held, which are equipped with multiple unreliable release sites [1].

Taken to its extreme, if neurons can optimise their threshold, zero noise should always be optimal. However, the continual learning that goes on in the brain involves constant adjustment of synaptic strengths and cell excitability. This should make threshold optimisation a continuously evolving process – one that can never be perfectly realised. Indeed, the large range of timescales over which changes in synaptic strength and in cell excitability occur [35,36] suggest that brain circuits cannot rely on permanently optimised thresholds. Thus, it may be advantageous to use synaptic noise to ensure that some spikes are transmitted rather than relying on an optimal threshold. In this context, release probability is not strictly noise: its use-dependent changes may represent a certain 'informative' operation that contributes to the transmitted message.

So far, we have defined optimality purely with respect to information transmission. However, even if a circuit is optimal from an information transmission standpoint, it is unlikely to be optimal once the cost of neurotransmitter release is considered. There is a relatively simple reason for that: spikes have limited bandwidth (sustained firing rates are almost always well below 100 Hz, with only occasional bursts of a few spikes at 1 kHz), so substantial input information, which occurs when release probability is high and/or input activity is intense, is largely wasted. As shown previously [37], for reasonable output firing rates, release probabilities that maximize information per release are in the range 0.2–0.8; exactly what is seen *in vivo*. Subsequent studies taking into account short-term plasticity [38–40] found the same thing: information per release peaked at release probabilities below one. It is easy to see why, at least in the extreme case: if two consecutive presynaptic spikes occur with a small interspike interval, it is more energy efficient to release neurotransmitter on only one of the spikes but double the conductance change than to release neurotransmitter on both spikes without doubling the conductance change. This has the added benefit of increasing slightly the peak conductance, and thus increasing the probability of a postsynaptic spike.

The studies mentioned earlier, on optimal information per release, considered noise only from synaptic failures. However, the same lessons should apply to variability in synaptic conductance but with an important difference: because conductance variability does not affect release probability, if the threshold is set perfectly, adding conductance variability would not increase information per release, and so would seem to be uniformly bad for information transmission per release. However, the fact that excitatory synapses appear to maximise synaptic conductance variability [30] points to the possibility that theory is missing an important role for this kind of noise.

It is also worth mentioning that noise can affect the input–output function of neurons. Modern theories of recurrent spiking networks suggest that most of the time the mean drive to a neuron is subthreshold, and it is temporal fluctuations around the mean that drive activity [31,41–44]. Whenever this is the case, increasing the (zero mean) noise – no matter how it is done – increases the probability that a postsynaptic cell will fire. Thus, noise, due either to synaptic failures or variability in peak synaptic conductance, is an energetically cheap way to increase the gain of a neuron (albeit at the cost of larger fluctuations). Noise can also decrease the probability of firing;

that happens when the input to a neuron is superthreshold. Thus, the overall effect of noise is to flatten the gain curve (input current to firing rate), and so widen a neuron's dynamic range. This is something that threshold adjustments cannot do.

A Computational Role for Noise

It is reasonable to assume that information transmission is the main function of synapses. However, there is another potential role, one rooted in the Bayesian approach to neuroscience [45–48]. It has become increasingly clear over the last several decades that organisms – from fly larvae to humans – compute and manipulate probabilities [48–51]. This is not surprising; it is impossible to make a good decision of any kind without some measure of uncertainty. For instance, every time we cross the street we have to estimate whether the oncoming car will hit us; because it is impossible to get an exact estimate, we must evaluate how uncertain we are about a collision. Without some estimate of uncertainty, we would not know whether we can cross the street reasonably safely. Much of the work in this area has focused on uncertainty in our perception of the outside world. However, there is another source of uncertainty: the computations the brain performs depend on synaptic strengths, but the correct strengths are not known perfectly. In principle, that degree of uncertainty should be communicated to the rest of the brain. Recently it was proposed that synapses compute their uncertainty and then communicate that uncertainty by adding noise to synaptic weights [52]. With this theory, variability in synaptic strength – caused by failures, variability in conductance size, or both – is a feature, not a bug; it is used by the rest of the brain to estimate uncertainty, and, ultimately, to determine confidence. Whether or not synapses do this is currently not known, but it is an avenue worth exploring, we would argue, and one that raises concrete questions for future research.

In our discussion of information transmission, we have operated under the assumption that noise is useful only if the spike threshold is too low. However, noise can be useful in and of itself. For instance, it can boost behavioural variability – an essential feature in most, if not all, adversarial games. It is also useful for exploratory behaviour, one prominent example being identifiable random number generators in birdsongs [53]. In some instances, however, synaptic noise could be unimportant, simply forced on the brain by biological reality. And the cost of noisy synapses may not be very high: chaotic dynamics [54] already introduces a large amount of noise in neural circuits, so adding a little more probably does not cost much in terms of computational capabilities. For instance, the noise associated with failures should have very little effect on the accuracy of computations [55], and the approximations that the brain makes have a much larger effect on performance than the noise found in either synapses or chaotic dynamics [56].

Concluding Remarks

Two key physiological phenomena contribute to what we consider synaptic noise: stochastic neurotransmitter release and variability in postsynaptic conductance given release. Multiple lines of experimental and theoretical evidence suggest that, at least in certain conditions, synaptic noise improves information handling by brain circuits. Here we attempted to discuss, and illustrate, how and when this is the case. Although much progress has been made, we still lack a deep understanding of the role of noise in the nervous system. In particular, it is currently unknown how noise interacts with plasticity at the synaptic level, learning at the behavioural level, or computations at the circuit level. These are all fundamental issues and represent fertile ground for experimental and theoretical exploration. No matter what its role, noise is a ubiquitous feature of brain activity and is likely to be a critical element in any theory of how the brain works (see Outstanding Questions).

Outstanding Questions

What is the experimental relationship between synaptic noise and synaptic morphology?

How stable are conductance fluctuations at individual synapses over short periods of time?

Is there a relationship between synaptic current variability and a synapse's position on the dendritic tree?

How does synaptic noise influence dendritic signal integration?

Do central synapses change their 'noisiness' with development and aging?

Are there subgroups of synapses or cells with distinct magnitudes of conductance variability?

Are there synaptic circuits or brain regions with distinct magnitudes of noise?

Does the magnitude of synaptic noise correlate with a neural circuit's ability to undergo plastic changes?

Does the magnitude of neural network noise correlate with the animal's ability to learn?

References

1. Tarr, T.B. *et al.* (2013) Are unreliable release mechanisms conserved from NMJ to CNS? *Trends Neurosci.* 36, 14–22
2. Stein, R.B. *et al.* (2005) Neuronal variability: noise or part of the signal? *Nat. Rev. Neurosci.* 6, 389–397
3. McDonnell, M.D. and Ward, L.M. (2011) The benefits of noise in neural systems: bridging theory and experiment. *Nat. Rev. Neurosci.* 12, 415–426
4. Faisal, A.A. *et al.* (2008) Noise in the nervous system. *Nat. Rev. Neurosci.* 9, 292–303
5. Edwards, P.J. and Murray, A.F. (1993) Analog synaptic noise – implications and learning improvements. *Int. J. Neural Syst.* 4, 427–433
6. Murray, A.F. and Edwards, P.J. (1993) Synaptic weight noise during multilayer perceptron training: fault tolerance and training improvements. *IEEE Trans. Neural Netw.* 4, 722–725
7. Varshney, L.R. *et al.* (2006) Optimal information storage in noisy synapses under resource constraints. *Neuron* 52, 409–423
8. Kitajo, K. *et al.* (2003) Behavioral stochastic resonance within the human brain. *Phys. Rev. Lett.* 90, 218103
9. Lu, L.L. *et al.* (2019) Effects of noise and synaptic weight on propagation of subthreshold excitatory postsynaptic current signal in a feed-forward neural network. *Nonlinear Dynam.* 95, 1673–1686
10. Stocks, N.G. (2000) Suprathreshold stochastic resonance in multilevel threshold systems. *Phys. Rev. Lett.* 84, 2310–2313
11. Mino, H. and Durand, D.M. (2010) Enhancement of information transmission of sub-threshold signals applied to distal positions of dendritic trees in hippocampal CA1 neuron models with stochastic resonance. *Biol. Cybern.* 103, 227–236
12. Stacey, W.C. and Durand, D.M. (2001) Synaptic noise improves detection of subthreshold signals in hippocampal CA1 neurons. *J. Neurophysiol.* 86, 1104–1112
13. Redman, S. (1990) Quantal analysis of synaptic potentials in neurons of the central nervous system. *Physiol. Rev.* 70, 165–198
14. Jack, J.J.B. *et al.* (1990) Quantal analysis of excitatory synaptic mechanisms in the mammalian central nervous system. *Cold Spring Harb. Symp. Quant. Biol.* 55, 57–67
15. Oertner, T.G. *et al.* (2002) Facilitation at single synapses probed with optical quantal analysis. *Nat. Neurosci.* 5, 657–664
16. Emptage, N.J. *et al.* (2003) Optical quantal analysis reveals a presynaptic component of LTP at hippocampal Schaffer-associational synapses. *Neuron* 38, 797–804
17. Sylantyev, S. *et al.* (2013) Cannabinoid- and lysophosphatidylinositol-sensitive receptor GPR55 boosts neurotransmitter release at central synapses. *Proc. Natl. Acad. Sci. U. S. A.* 110, 5193–5198
18. Jensen, T.P. *et al.* (2019) Multiplex imaging relates quantal glutamate release to presynaptic Ca^{2+} homeostasis at multiple synapses *in situ*. *Nat. Commun.* 10, 1414
19. Borst, J.G.G. (2010) The low synaptic release probability *in vivo*. *Trends Neurosci.* 33, 259–266
20. Hanse, E. and Gustafsson, B. (2001) Paired-pulse plasticity at the single release site level: an experimental and computational study. *J. Neurosci.* 21, 8362–8369
21. Liu, G. *et al.* (1999) Variability of neurotransmitter concentration and nonsaturation of postsynaptic AMPA receptors at synapses in hippocampal cultures and slices. *Neuron* 22, 395–409
22. Franks, K.M. *et al.* (2003) Independent sources of quantal variability at single glutamatergic synapses. *J. Neurosci.* 23, 3186–3195
23. Zheng, K. *et al.* (2008) Receptor actions of synaptically released glutamate: the role of transporters on the scale from nanometers to microns. *Biophys. J.* 95, 4584–4596
24. Groc, L. *et al.* (2004) Differential activity-dependent regulation of the lateral mobilities of AMPA and NMDA receptors. *Nat. Neurosci.* 7, 695–696
25. Triller, A. and Choquet, D. (2008) New concepts in synaptic biology derived from single-molecule imaging. *Neuron* 59, 359–374
26. Xie, X.P. *et al.* (1997) Novel expression mechanism for synaptic potentiation: alignment of presynaptic release site and postsynaptic receptor. *Proc. Natl. Acad. Sci. U. S. A.* 94, 6983–6988
27. Savtchenko, L.P. and Rusakov, D.A. (2014) Moderate AMPA receptor clustering on the nanoscale can efficiently potentiate synaptic current. *Philos. Trans. R. Soc. Lond. Ser. B Biol. Sci.* 369, 20130167
28. Tang, A.H. *et al.* (2016) A trans-synaptic nanocolumn aligns neurotransmitter release to receptors. *Nature* 536, 210–214
29. Nair, D. *et al.* (2013) Super-resolution imaging reveals that AMPA receptors inside synapses are dynamically organized in nanodomains regulated by PSD95. *J. Neurosci.* 33, 13204–13224
30. Savtchenko, L.P. *et al.* (2013) Central synapses release a resource-efficient amount of glutamate. *Nat. Neurosci.* 16, 10–12
31. Yarom, Y. and Hounsgaard, J. (2011) Voltage fluctions in neurons: signal or noise? *Physiol. Rev.* 91, 917–929
32. Bloss, E.B. *et al.* (2018) Single excitatory axons form clustered synapses onto CA1 pyramidal cell dendrites. *Nat. Neurosci.* 21, 353–363
33. Zador, A. (1998) Impact of synaptic unreliability on the information transmitted by spiking neurons. *J. Neurophysiol.* 79, 1219–1229
34. Manwani, A. and Koch, C. (2001) Detecting and estimating signals over noisy and unreliable synapses: information-theoretic analysis. *Neural Comput.* 13, 1–33
35. Ma, Z. *et al.* (2019) Cortical circuit dynamics are homeostatically tuned to criticality *in vivo*. *Neuron* 104, 655–664
36. Turrigiano, G.G. *et al.* (1998) Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* 391, 892–896
37. Levy, W.B. and Baxter, R.A. (2002) Energy-efficient neuronal computation via quantal synaptic failures. *J. Neurosci.* 22, 4746–4755
38. Salmasi, M. *et al.* (2019) Short-term synaptic depression can increase the rate of information transfer at a release site. *PLoS Comput. Biol.* 15, e1006666
39. Goldman, M.S. (2002) Redundancy reduction and sustained firing with stochastic depressing synapses. *J. Neurosci.* 22, 584–591
40. Goldman, M.S. (2004) Enhancement of information transmission efficiency by synaptic failures. *Neural Comput.* 16, 1137–1162
41. Renart, A. *et al.* (2007) Mean-driven and fluctuation-driven persistent activity in recurrent networks. *Neural Comput.* 19, 1–46
42. Roach, J.P. *et al.* (2018) Resonance with subthreshold oscillatory drive organizes activity and optimizes learning in neural networks. *Proc. Natl. Acad. Sci. U. S. A.* 115, E3017–E3025
43. van Vreeswijk, C. and Sompolinsky, H. (1998) Chaotic balanced state in a model of cortical circuits. *Neural Comput.* 10, 1321–1371
44. van Vreeswijk, C. and Sompolinsky, H. (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274, 1724–1726
45. Haefner, R.M. *et al.* (2016) Perceptual decision-making as probabilistic inference by neural sampling. *Neuron* 90, 649–660
46. Aitchison, L. and Lengyel, M. (2017) With or without you: predictive coding and Bayesian inference in the brain. *Curr. Opin. Neurobiol.* 46, 219–227
47. Knill, D.C. and Pouget, A. (2004) The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* 27, 712–719
48. Pouget, A. *et al.* (2013) Probabilistic brains: knowns and unknowns. *Nat. Neurosci.* 16, 1170–1178

49. Strange, B.A. *et al.* (2005) Information theory, novelty and hippo-campal responses: unpredicted or unpredictable? *Neural Netw.* 18, 225–230

50. Grabska-Barwinska, A. *et al.* (2017) A probabilistic approach to demixing odors. *Nat. Neurosci.* 20, 98–106

51. Navajas, J. *et al.* (2017) The idiosyncratic nature of confidence. *Nat. Hum. Behav.* 1, 810–818

52. Aitchison, L. *et al.* (2017) Probabilistic synapses. *ArXiv* 1410, 1029

53. Masco, C. *et al.* (2016) The Song Overlap Null model Generator (SONG): a new tool for distinguishing between random and non-random song overlap. *Bioacoustics* 25, 29–40

54. London, M. *et al.* (2010) Sensitivity to perturbations *in vivo* implies high noise and suggests rate coding in cortex. *Nature* 466, 123–127

55. Zylberberg, J. and Strowbridge, B. (2017) Mechanisms of persistent activity in cortical circuits: possible neural substrates for working memory. *Annu. Rev. Neurosci.* 40, 603–627

56. Beck, J.M. *et al.* (2012) Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron* 74, 30–39

57. Migliore, M. *et al.* (1999) Role of an A-type K+ conductance in the back-propagation of action potentials in the dendrites of hippocampal pyramidal neurons. *J. Comput. Neurosci.* 7, 5–15

58. Hines, M.L. and Carnevale, N.T. (2001) NEURON: a tool for neuroscientists. *Neuroscientist* 7, 123–135

59. Shannon, C.E. and Weaver, W. (1949) *The Mathematical Theory of Communication*, The University of Illinois Press