



Published in final edited form as:

Nature. 2010 July 1; 466(7302): 123–127. doi:10.1038/nature09086.

Sensitivity to perturbations *in vivo* implies high noise and suggests rate coding in cortex

Michael London¹, Arnd Roth¹, Lisa Beeren¹, Michael Häusser¹, and P.E. Latham²

¹Wolfson Institute for Biomedical Research and Department of Neuroscience, Physiology and Pharmacology, University College London, Gower Street, London WC1E 6BT, UK

²Gatsby Computational Neuroscience Unit University College London, Gower Street, London WC1E 6BT, UK

Abstract

It is well known that neural activity exhibits variability, in the sense that identical sensory stimuli produce different responses, but it has been difficult to determine what this variability means. Is it noise, or does it carry important information – about, for example, the internal state of the organism? We address this issue from the bottom up, by asking whether small perturbations to activity in cortical networks are amplified. Based on *in vivo* whole-cell recordings in rat barrel cortex, we find that a perturbation consisting of a single extra spike in one neuron produces ~28 additional spikes in its postsynaptic targets, and we show, using simultaneous intra- and extra-cellular recordings, that a single spike produces a detectable increase in firing rate in the local network. Theoretical analysis indicates that this amplification leads to intrinsic, stimulus-independent variations in membrane potential on the order of $\pm 2.2 - 4.5$ mV – variations that are pure noise, and so carry no information at all. Therefore, for the brain to perform reliable computations, it must either use a rate code, or generate very large, fast depolarizing events, such as those proposed by the theory of synfire chains – yet in our *in vivo* recordings, we found that such events were very rare. Our findings are consistent with the idea that cortex is likely to use primarily a rate code.

The brain, like all physical devices, operates in the presence of noise. Nevertheless, it performs complex computations with amazing speed and accuracy, in some cases reaching fundamental physical limits set by its sensors¹. Clearly, the brain has devised computational strategies, and a neural code, that are robust to noise, and understanding the structure of that noise should shed light on both.

The traditional experimental approach to studying noise in cortical sensory areas is to repeatedly present the same stimulus to an organism while recording neuronal responses. Such recordings always show trial-to-trial variability²⁻⁴. However, interpreting that variability has been difficult, as there are two possible sources for it. One is the variability

Users may view, print, copy, download and text and data- mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use: http://www.nature.com/authors/editorial_policies/license.html#terms

Address for editorial correspondence: Peter E. Latham Gatsby Computational Neuroscience Unit University College London 17 Queen Square London WC1N 3AR UK tel. +44-(0)20-7679-1178 fax +44-(0)20-7679-1173 pel@gatsby.ucl.ac.uk.

associated with truly random events, such as ion channel noise and stochastic synaptic release. This is intrinsic noise: intrinsic because it cannot be eliminated, and noise because it contributes to the neuronal variability but carries no information whatsoever. The other source of variability is activity from other brain areas. That activity might provide information about, say, the degree of arousal or some other internal state, but it would not be related to the stimulus. This variability is signal, even though it would look like noise to an observer trying to relate the neural activity to the stimulus.

Here we determine a lower bound on the level of intrinsic noise in cortical networks. The lower bound we consider is the trial-to-trial variability that would be observed in a deterministic network that received identical input, down to the last spike, on multiple trials, except for one very brief random event. If the dynamics of the network is such that small differences in activity associated with the single random event lead to very large differences in patterns of neuronal activity, then trial-to-trial variability would, necessarily, be high. If, on the other hand, small differences in activity lead to even smaller differences in patterns of neuronal activity, then trial-to-trial variability can be low.

In our analysis and experiments, the random event is a single extra spike added to a randomly chosen excitatory neuron, as in Fig. 1a. This one extra spike (magenta arrow at time $t=0$ in Fig. 1a) can produce other extra spikes in its postsynaptic targets (magenta arrows). If it produces more than one, on average, then perturbations would be amplified, and noise would be high (see the steady state regime in Figs. 1a and b). If, on the other hand, one extra spike produces less than one extra spike, on average, then perturbations would decay, and noise could be small.

To determine the average number of extra postsynaptic spikes produced by a single extra presynaptic spike, we note that it is the product of two numbers: the average number of connections made by each neuron, and the average probability that a unitary synaptic event will cause an extra spike. The first number is known from anatomical studies^{5,6} to be between 1,000 and 2,000 (a synaptic connection which gives rise to a unitary EPSP can consist of multiple synaptic contacts; here we assume that a neuron makes 5 synaptic contacts per connection⁷). Thus, one extra spike produces, on average, 1500 ± 500 extra EPSPs in the network.

The second number, the probability that a unitary synaptic input produces an extra spike, was determined experimentally. We made whole-cell recordings from layer 5 pyramidal neurons in the barrel cortex of anesthetized rats while injecting current pulses to generate postsynaptic currents (injected PSCs) of various amplitudes. For each current amplitude we constructed a PSTH (post-stimulus time histogram) triggered on the time of the injected current pulses, and used it to deduce the probability of an extra spike. A typical experiment is illustrated in Fig. 2a. As the expanded trace at the bottom of Fig. 2a shows, on a trial-by-trial basis it is very difficult to tell whether an individual injected PSC has an effect on the probability of an extra spike. The PSTHs, however, reveal a clear signal (Fig. 2b). Integrating the PSTH over a time window of 5 ms, we find that a single input with an amplitude of +25 pA causes the probability of observing a spike to increase by 0.004 (top

panel of Fig. 2b), and an input with an amplitude of -25 pA causes the probability to decrease by 0.001 (bottom panel).

Figure 3a shows, for a range of positive and negative injected PSCs, the pooled data from the 40 cells in our dataset. There are clear peaks in the PSTHs for positive currents and clear valleys for negative currents. Moreover, the cumulative probability of an extra spike (the integral of the PSTH relative to the mean firing rate) shown below the PSTHs does not return to baseline, indicating that the peaks are due to changes in net spike output, rather than shifted spikes (all cumulative probabilities are significantly different from zero at the $p=0.05$ level, except for the last two points of the 25 pA histogram). In Fig. 3b we combine the data from the different current amplitudes and plot probability versus the total charge in the injected current pulses (we use charge rather than amplitude for reasons we discuss shortly).

The results of these experiments provide us with a relationship between the size of injected PSCs and the probability of extra spikes *in vivo*. However, what we need is the relationship between physiological synaptic inputs (which are generated by conductance changes, typically on dendrites) and the probability of an extra spike. To determine this, we constructed a detailed compartmental model of a pyramidal neuron (see Supplementary Information, Sec. 4), and used it to simulate the effect of conductance changes on the probability of an extra spike. These simulations show that the probability of an extra spike depends on the total charge arriving at the soma, regardless of whether it is caused by a current injection or a conductance change, and is relatively insensitive to the location of the input (Fig. 3c; see also refs. 8,9).

Combining this result with the fact that there is an approximately linear relationship between charge and probability (Fig. 3b), we see that the average probability that a unitary synaptic input causes an extra spike is the product of the slope of the positive part of the regression line in Fig. 3b and the average charge associated with a unitary synaptic input *in vivo*. The former we have measured; it is 0.061 ± 0.010 , in units of probability/pC (Fig. 3b). The latter we estimated from published data on paired recordings^{7,10-12} (data that includes synaptic failures). Taking an average across these reports, we find that the average charge associated with a unitary synaptic input *in vivo* is 0.31 ± 0.07 pC, corresponding to an EPSP of approximately 1 mV (see Supplementary Information, Sec. 4). The product of the two numbers, denoted \bar{p}_e , is given by

$$\begin{aligned} \bar{p}_e &= (0.061 \pm 0.010 \text{ probability/pC}) \times (0.31 \pm 0.07 \text{ pC/unitary synaptic input}) \\ &= 0.019 \pm 0.0053 \text{ probability/unitary synaptic input} \end{aligned}$$

(see Supplementary Information, Sec. 8, for a derivation of confidence limits).

Multiplying \bar{p}_e (0.019 ± 0.0053) by the average number of connections made by each neuron (1500 ± 500) yields 28 ± 13 extra spikes/spike. Thus one extra spike in an excitatory neuron causes, on average, 28 of its postsynaptic neurons to emit an extra spike. This implies very rapid amplification of perturbations, and should quickly disrupt spike patterns across the

network. In fact, if the perturbations were to grow unchecked, in just 5 integration time steps, there would be about 17 million extra spikes in the network.

Perturbations do not, of course, grow unchecked. In fact, after the initial extra spike, the growth rate decreases. This is because extra and missed excitatory and inhibitory spikes interact: if an extra excitatory and extra inhibitory spike have the same postsynaptic target, they will at least partially cancel each other, thus reducing the probability of either a missed or extra postsynaptic spike (points 9 and 10 in Fig. 1a). As the number of missed and extra spikes grows, cancellation becomes more likely, and eventually the network reaches a steady state in which missed spikes are produced at the same rate, on average, as extra ones (see Fig. 1b and Supplementary Figure S9).

These results suggest that a single extra spike should have a measurable effect on network firing rate. That effect, though, should be small: the approximately 2% increase in firing rate that we saw for connected pairs of neurons (summarized by p_e^- above) is reduced by the low connectivity in somatosensory cortex (approximately 4%¹³), and spread out by axonal delays, dendritic filtering and latency to spike. Taking these into account quantitatively (see Supporting Information, Sec. 2), we find that a single spike should lead to a local increase in firing rate of 0.04-0.08 Hz.

To test this prediction, we conducted a second series of *in vivo* experiments in which we triggered single spikes in a presynaptic neuron via a whole-cell patch-clamp recording, and simultaneously recorded spikes from a population of neurons in the local network using a 16 channel extracellular electrode array placed in the somatosensory cortex (Figs. 4a and b). We then constructed a PSTH of extracellular spikes triggered on the stimulus (Fig. 4c). As predicted, we observed an increase in firing rate on the extracellular electrodes following the single spike in the stimulated neuron. The increase, as assessed by the cumulative increase in the probability of an extra spike, is statistically significant at the $p=0.01$ level for greater than 100 ms (Fig. 4d; $n=10$ experiments). Moreover, in the first 10-20 ms, the increase was 0.03-0.065 Hz (Fig. 4c, inset); very close to the increase of 0.04-0.08 Hz predicted above. Thus, not only does a single extra spike introduced into somatosensory cortex produce a measurable effect on the network – one that lasts for more than 50 ms – it produces an effect whose size is predicted by our single-neuron current injection experiments.

So far we have focused on the effect of a single spike; now we turn to the steady state, where the perturbations have stopped growing (so that missed and extra spikes occur at the same rate), and ask: how big are the trial-to-trial voltage fluctuations associated with the ongoing missed and extra spikes? This is a critical question, since it is, ultimately, these voltage fluctuations that limit spike timing precision. To answer it, we assume that in steady state there are m missed and extra spikes in a time window that corresponds to a typical neuron's mean integration time, compute m self-consistently (by demanding that the presynaptic and postsynaptic probabilities of a spike are the same), and then relate it to the size of the voltage fluctuations. The presynaptic probability, denoted, p_{pre} , is equal to m / K where K is the average number of presynaptic connections. The postsynaptic probability, denoted p_{post} , is approximately \sqrt{m} times larger than p_e^- , the probability associated with one presynaptic spike (assuming that missed and extra spikes are reasonably uncorrelated,

invoking central limit type arguments, and using the linearity of probability versus charge in Fig. 3b). Consequently, $p_{post} \propto \sqrt{m} \bar{p}_e$. Self-consistency tells us that $p_{pre} = p_{post}$, and so we arrive at

$$\sqrt{m} \propto K \bar{p}_e. \quad (1)$$

Again assuming that the m missed and extra spikes received by a neuron are reasonably uncorrelated, the amplitude of intrinsic voltage fluctuations associated with missed and extra spikes, denoted σ_V , is proportional to $\sqrt{m} \bar{V}_{PSP}$ where \bar{V}_{PSP} is the average PSP size. Combining this observation with Eq. (1), we have

$$\sigma_V \propto K \bar{p}_e \bar{V}_{PSP}. \quad (2)$$

The key point to extract from Eq. (2) is that the voltage fluctuations are proportional to the growth rate of the perturbations, $K \bar{p}_e$. Since $K \bar{p}_e \approx 28$, we expect the intrinsic voltage fluctuations associated with the growth of perturbations to be large. And, indeed, they are: in Supplementary Information, Sec. 6, we carry out a much more extensive analysis which takes into account correlations among excitatory and inhibitory neurons as well as saturation of the number of missed and extra spikes, and we verify the analysis with large-scale network simulations. What we find is that the constant of proportionality in Eq. (2) ranges in our data from 0.08 to 0.16, and so σ_V ranges from about 2.2 to 4.5 mV (Supplementary Figure S7).

These large trial-to-trial voltage fluctuations indicate that there is a high level of intrinsic noise in the cortex. Therefore, rapid depolarizing events, like the ones postulated for synfire chains¹⁴ or polychronization¹⁵, would be required to produce very precisely timed spikes. More quantitatively, for a given value of the voltage fluctuations, σ_V , a spike will occur with a precision of $\delta\tau$ ms if the membrane potential changes by $2\sigma_V$ mV in $\delta\tau$ ms (the factor of 2 is necessary to ensure reliability: since σ_V is the standard deviation, the voltage can, with reasonable probability, range from σ_V below the mean to σ_V above it). To determine how often there are voltage excursions that would allow spike timing of precision $\delta\tau$, we examined experimentally-recorded voltage traces (Fig. 5a), counted how often the membrane potential changed by $2\sigma_V$ mV in $\delta\tau$ ms, and averaged over our uncertainty in σ_V . The results are shown in Fig. 5b. This figure indicates that precisely timed events are very rare: events with a precision of $\delta\tau=1$ ms occur, on average, once every 10,000 seconds per neuron (about once every three hours); events with a precision of 5 ms occur less than once every 100 seconds per neuron, and events with a precision of 10-20 ms occur less than once every 40 seconds per neuron. In addition, not only are precisely timed events rare, the rates shown in Fig. 5b constitute an upper bound, since precisely timed events occur by chance at nonzero rates.

Taken together, our results indicate that there is a large amount of intrinsic noise in cortex, and that this noise puts severe constraints on spike timing codes. One potential caveat is that our results were obtained under anaesthesia. The effect of anaesthesia on our estimate of the number of extra spikes/spike is likely to be small, since responsiveness of barrel cortex

neurons under the awake and anesthetized states is similar^{16,17}. The frequency of large voltage excursions, on the other hand, may be more sensitive to anaesthesia; in particular, it is possible that the rate of large, fast depolarizing events (Fig. 5b) is simply higher in awake than anesthetized animals (although they would have to be much higher to affect our conclusions). A second caveat is that our analysis assumed linear synaptic integration. Recent studies show that precisely timed input to dendritic branches can yield precisely timed output spikes in the axon without large somatic sub-threshold voltage excursions¹⁸⁻²¹. These mechanisms, though, have been demonstrated only *in vitro*, and only when input to the dendrites was carefully regulated in both space and time; it remains unclear to what extent these conditions are satisfied *in vivo* (but see ref. 22).

Our study is in line with previous theoretical work suggesting that neuronal networks are chaotic²³⁻²⁶. It is, however, the first experimental demonstration of the sensitivity to perturbations *in vivo*. We are also the first to explore the consequences of these results for the level of noise in the cortex and its likely effect on the precision of spike timing.

What do our results imply for neural coding? Superficially, it seems natural to conclude that if every spike has a large effect on network activity, then every spike should count, and the brain must be using a very sophisticated neural code in which the time and identity of every spike carries meaningful information. In fact, our results imply just the opposite. This is because network activity is bounded, so growth of perturbations in some dimensions (e.g., as measured by trial-to-trial difference in membrane potential) necessarily implies contraction in others²⁷. It is this contraction that causes networks to rapidly forget their past. Thus, while an extra spike can radically modify patterns of activity, patterns of activity cannot encode which extra spike caused the modification. The implication, then, is not that rat barrel cortex (and, we suspect, other areas of cortex and other species) must be using a very sophisticated spike timing code, but that it is likely to be using a code that is robust to perturbations, such as a rate code in which it is the average firing rate over large populations of neurons that carries information.

Finally, the fact that studies have found millisecond timing in both anaesthetized and awake animals in the rat barrel cortex²⁸ as well as in other cortical regions^{3,29} is not inconsistent with our results. The precise timing in those studies is associated with a feedforward sweep of activity caused by a rapidly time-varying stimulus. Our results, on the other hand, apply to slowly varying stimuli and higher order computations, and suggest that in those cases the cortex does not rely on spike timing.

Methods Summary

Sprague Dawley rats (postnatal day 18 to 25, average 45.6 g) were anesthetized with urethane (Sigma; 1.5 g/kg, i.p.), and *in vivo* whole-cell recordings were made using blind patch techniques. Recordings were made at a depth of $1180 \pm 165 \mu\text{m}$, and neurons were identified as layer 5 pyramidal cells by input resistance, firing properties, and in 5 cases by cell morphology ($n=5/5$). For the single-cell experiments, EPSC waveforms separated by either 100 or 200 ms were injected into the recorded neuron. The rise and decay times of the EPSC waveforms were 0.3 and 1.7 ms, respectively, and the amplitudes alternated between

(+50, -25, +25, -50 pA) or (+100, -50, +50, -100 pA). In some cases, we performed continuous whisker stimulation; the stimulus had no measurable effect on the probability of an extra spike (paired *t*-test, $p > 0.9$), so we pooled all data. For the combined intracellular-extracellular recording experiments, we used in addition a silicon extracellular multi-site probe (probe type A-1-16-3mm-50-177, Neuronexus Technologies, MI USA) which was lowered to 1200 μm from the brain surface at an angle of 60 degrees. The patch pipette was inserted and lowered to 100-300 μm from the probe. Brief current pulses (2.5-10 ms, 1-3 nA) were injected via the patch pipette at 200-400 ms intervals to trigger spikes, and both the intracellular and extracellular signals were recorded. A total of 13,000 stimuli were injected into 10 neurons. When the same spike appeared on more than one extracellular recording site, only one spike from one site was selected for analysis.

Methods

Animals and surgery

The care and experimental manipulation of the animals was carried out in accordance with institutional and national guidelines. For all experiments, Sprague Dawley rats (postnatal day 18 to 25, average 45.6 g) were anesthetized with urethane (Sigma; 1.5 g/kg, i.p.). A small craniotomy (1 mm²) in a region overlying S1 (centered 5.5 mm lateral, 2.5 mm caudal of Bregma) was made and a small opening (~0.1 mm²) was made in the dura. Supplemental urethane (10% of original dose, i.p.) was given whenever limb withdrawal responses were present or whisker movements were observed. Body temperature was maintained at 37° C with a feedback-controlled heating blanket.

In vivo recording (patch-clamp)

For the current injection experiments (summarized in Fig. 3, main text), whole-cell recordings were made using blind patch techniques³¹. Standard borosilicate glass patch pipettes (5.5 M Ω) were filled with internal solution containing the following (in mM): K-methanesulfonate 110, KCl 15, HEPES 10, Mg-ATP 4, Na₂GTP 0.3, Na-phosphocreatine 10, and 0.3% biocytin; pH 7.2, 285 mOsm. Access resistance was typically 20-40 M Ω at the start of the recording, and degraded with time. Traces with access resistance above 100 M Ω were not included in the analysis. Recordings were made at a depth of 1180 \pm 165 μm , and neurons were identified as layer 5 pyramidal cells by input resistance, firing properties, and in 5 cases by cell morphology ($n=5/5$). Data were filtered at 3-10 kHz and acquired at 50 kHz using Axograph software (Axon Instruments) and an ITC-18 interface (Instrutech). Input resistance was calculated by fitting a linear function to the steady-state I/V curve (obtained from voltage deflections during 400 ms current steps from -300 pA to +500 pA in 100 pA steps). EPSC waveforms were generated by a double exponential current injection with $\tau_{\text{rise}}=0.3$ ms and $\tau_{\text{decay}}=1.7$ ms (ref. 32). The experimental protocol consisted of a train of 50 EPSC waveforms separated by 200 ms with alternating amplitudes of either (+50, -25, +25, -50 pA) or (+100, -50, +50, -100 pA). In six of the experiments, diffuse whisker stimulation was delivered by a custom-made motor rotating at 1 Hz driving a 1 cm diameter disk covered with #0 sandpaper. The sandpaper touched the primary whisker and the two adjacent whiskers on the same row (all other whiskers were trimmed). The primary whisker was identified prior to recording by monitoring the local field potential with a low resistance

(1 M Ω) electrode filled with ACSF. Stimuli with and without whisker stimulation were interleaved. The stimulus had no measurable effect on the probability of an extra spike (paired *t*-test, $p>0.9$), so we pooled all data.

***In vivo* recording (combined patch-clamp and extracellular)**

For the combined patch-clamp/extracellular recordings, a silicon extracellular 16 site linear probe (type A-1-16-3mm-50-177, Neuronexus Technologies, MI USA) was lowered 1200 μm from the brain surface at an angle of 60 degrees. Following identification of clear spiking units, a patch pipette was inserted and lowered to 100-300 μm from the probe (distance calculated based on the distance and the angle between the two electrodes using a custom made system by Luigs & Neumann GmbH). We then searched for cells in 2 μm steps. After establishing a whole-cell recording, brief current pulses (2.5-10 ms, 1-3 nA) were injected via the patch pipette into the recorded neuron at 200-400 ms intervals, and both the intracellular and extracellular signals were recorded using an Rx5 Pentusa system (Tucker Davis, FL USA). Ideally, we would have liked to choose the amplitude and duration of the current pulses so that each stimulus produced exactly one spike. However, because of up and down states, this was not possible: a stimulus strong enough to always trigger a spike in the down state would be strong enough to regularly produce more than one spike in the up state. Thus, we adjusted the strength of the stimulus so that it rarely produced more than one spike. We were largely successful: out of the 13,000 stimuli we delivered (in 10 experiments), only 498 (3.8%) produced 2 spikes. Given the linearity between presynaptic and postsynaptic spikes (Fig. 3b, main text), the small fraction of stimuli that elicited 2 spikes should have virtually no effect on our results.

Because spikes were far less likely to occur in the down state than the up state, the number of spikes in the intracellularly recorded neuron was correlated with ongoing activity, and thus correlated with spikes on the extracellular electrodes. Thus, to ensure a flat baseline, we constructed PSTHs triggered on stimulus onset, not on intracellular spike times. As can be seen in Fig. 4c of the main text, this strategy was successful, as the PSTH was, indeed, flat before stimulus onset.

The peak in the PSTH after stimulus onset (Fig. 4c) was small, but the fact that 16 consecutive 5 ms bins were more than one standard deviation above the mean suggests that it is significant. To test this, we constructed a cumulative PSTH by subtracting the baseline (the mean firing rate between -100 and 0 ms) from the PSTH and integrating the difference (starting from $t=0$ and integrating in both directions). To compute error bars, we used a bootstrap method in which we constructed surrogate PSTHs by randomly sampling 200 ms epochs from the extracellular spike trains. We constructed 10,000 surrogate PSTHs, which we turned into cumulative PSTHs as described above. This gave us a null-distribution of cumulative PSTHs at each time point; we used those to construct the error bars shown in Fig. 4d of the main text.

Spike detection

The extracellular signal was separated online into an LFP component (0-0.3 kHz) sampled at 3 kHz, and a high frequency component (0.3-5 kHz) sampled at 25 kHz. Spike detection was

done offline using the high frequency component of the extracellular signal, as described in Supplementary Information, Sec. 1.

Biophysical models and network simulations

Simulations using a detailed model of a Layer 5 pyramidal neuron were based on a published model³³ with slight modification, and were performed in the NEURON simulation environment³⁴; see Supplementary Information, Sec. 4. The network simulations were also based on a previously published model³⁵, but with modifications to allow the addition and detection of extra spikes; see Supplementary Information, Sec. 7.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

We would like to thank Peter Dayan for many valuable discussions, Ingrid van Welle and Peter Dayan for comments on the manuscript, and Hermann Cuntz for helpful comments on the spike detection algorithm. PEL was supported by the Gatsby Charitable Foundation and by the US National Institute of Mental Health grant R01 MH62447. ML, AR, LB and MH were supported by the Wellcome Trust, Gatsby Charitable Foundation, Engineering and Physical Sciences Research Council and Medical Research Council.

References

1. Hecht S, Shlaer S, Pirenne M. Energy, Quanta, and Vision. *J. Gen. Physiol.* 1942; 25(6):819–840. [PubMed: 19873316]
2. Richmond BJ, Optican LM, Podell M, Spitzer H. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. I. Response characteristics. *J Neurophysiol.* 1987; 57(1):132–146. [PubMed: 3559668]
3. Victor JD, Purpura KP. Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J Neurophysiol.* 1996; 76(2):1310–1326. [PubMed: 8871238]
4. Tolhurst DJ, Movshon JA, Dean AF. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* 1983; 23(8):775–785. [PubMed: 6623937]
5. Binzegger T, Douglas RJ, Martin KA. A quantitative map of the circuit of cat primary visual cortex. *J Neurosci.* 2004; 24(39):8441–8453. [PubMed: 15456817]
6. Braitenberg, V.; Schüz, A. *Anatomy of the Cortex.* 1991.
7. Markram H, Lübke J, Frotscher M, Roth A, Sakmann B. Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex. *J Physiol (Lond).* 1997; 500(Pt 2):409–440. [PubMed: 9147328]
8. Herrmann A, Gerstner W. Noise and the PSTH response to current transients: I. General theory and application to the integrate-and-fire neuron. *J. Comput. Neurosci.* 2001; 11:135–151. [PubMed: 11717530]
9. Richardson M. Firing-rate response of linear and nonlinear integrate-and-fire neurons to modulated current-based and conductance-based synaptic drive. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 2007; 76:021919–021919. [PubMed: 17930077]
10. Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB. Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol.* 2005; 3(3):e68. [PubMed: 15737062]
11. Barbour B, Brunel N, Hakim V, Nadal JP. What can we learn from synaptic weight distributions? *Trends Neurosci.* 2007; 30(12):622–629. [PubMed: 17983670]
12. Thomson AM, Lamy C. Functional maps of neocortical local circuitry. *Front. Neurosci.* 2007; 1(1):19–42. [PubMed: 18982117]

13. Holmgren C, Harkany T, Svennenfors B, Zilberter Y. Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. *J Physiol.* 2003; 551(Pt 1):139–153. [PubMed: 12813147]
14. Abeles, M. *Corticonics: neural circuits of the cerebral cortex.* 1991.
15. Izhikevich E. Polychronization: computation with spikes. *Neural Comput.* 2006; 18:245–282. [PubMed: 16378515]
16. Crochet S, Petersen CC. Correlating whisker behavior with membrane potential in barrel cortex of awake mice. *Nat Neurosci.* 2006; 9(5):608–610. [PubMed: 16617340]
17. Ferezou I, Bolea S, Petersen CC. Visualizing the cortical representation of whisker touch: voltage-sensitive dye imaging in freely moving mice. *Neuron.* 2006; 50(4):617–629. [PubMed: 16701211]
18. Larkum ME, Zhu JJ, Sakmann B. A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature.* 1999; 398(6725):338–341. [PubMed: 10192334]
19. Ariav G, Polsky A, Schiller J. Submillisecond precision of the input-output transformation function mediated by fast sodium dendritic spikes in basal dendrites of CA1 pyramidal neurons. *J Neurosci.* 2003; 23(21):7750–7758. [PubMed: 12944503]
20. Häusser M, Mel B. Dendrites: bug or feature? *Curr Opin Neurobiol.* 2003; 13(3):372–383. [PubMed: 12850223]
21. London M, Häusser M. Dendritic Computation. *Annu. Rev. Neurosci.* 2005; 28(1):503–532. [PubMed: 16033324]
22. Murayama M, et al. Dendritic encoding of sensory stimuli controlled by deep cortical interneurons. *Nature.* 2009; 457(7233):1137–1141. [PubMed: 19151696]
23. van Vreeswijk C, Sompolinsky H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science.* 1996; 274:1724–1726. [PubMed: 8939866]
24. van Vreeswijk C, Sompolinsky H. Chaotic balanced state in a model of cortical circuits. *Neural Comput.* 1998; 10:1321–1371. [PubMed: 9698348]
25. Banerjee A, Seriès P, Pouget A. Dynamical constraints on using precise spike timing to compute in recurrent cortical networks. *Neural Comput.* 2008; 20:974–993. [PubMed: 18085984]
26. Izhikevich EM, Edelman GM. Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences.* 2008; 105(9):3593.
27. Smale S. Differentiable dynamical systems. *Bulletin of the American Mathematical Society.* 1967; 73:747–817.
28. Arabzadeh E, Panzeri S, Diamond M. Deciphering the spike train of a sensory neuron: counts and temporal patterns in the rat whisker pathway. *J Neurosci.* 2006; 26(36):9216–9226. [PubMed: 16957078]
29. Bair W, Koch C. Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput.* 1996; 8(6):1185–1202. [PubMed: 8768391]
30. de Kock CP, Sakmann B. Spiking in primary somatosensory cortex during natural whisking in awake head-restrained rats is cell-type specific. *Proc Natl Acad Sci U S A.* 2009; 106(38):16446–16450. [PubMed: 19805318]
31. Margrie TW, Brecht M, Sakmann B. In vivo, low-resistance, whole-cell recordings from neurons in the anaesthetized and awake mammalian brain. *Pflügers Arch.* 2002; 444(4):491–498. [PubMed: 12136268]
32. Häusser M, Roth A. Estimating the time course of the excitatory synaptic conductance in neocortical pyramidal cells using a novel voltage jump method. *J Neurosci.* 1997; 17(20):7606–7625. [PubMed: 9315883]
33. Mainen ZF, Joerges J, Huguenard JR, Sejnowski TJ. A model of spike initiation in neocortical pyramidal neurons. *Neuron.* 1995; 15(6):1427–1439. [PubMed: 8845165]
34. Hines ML, Carnevale NT. The NEURON simulation environment. *Neural Comput.* 1997; 9(6):1179–1209. [PubMed: 9248061]
35. Latham PE, Richmond BJ, Nelson PG, Nirenberg SN. Intrinsic dynamics in neuronal networks: I. Theory. *J. Neurophysiol.* 2000; 83:808–827. [PubMed: 10669496]

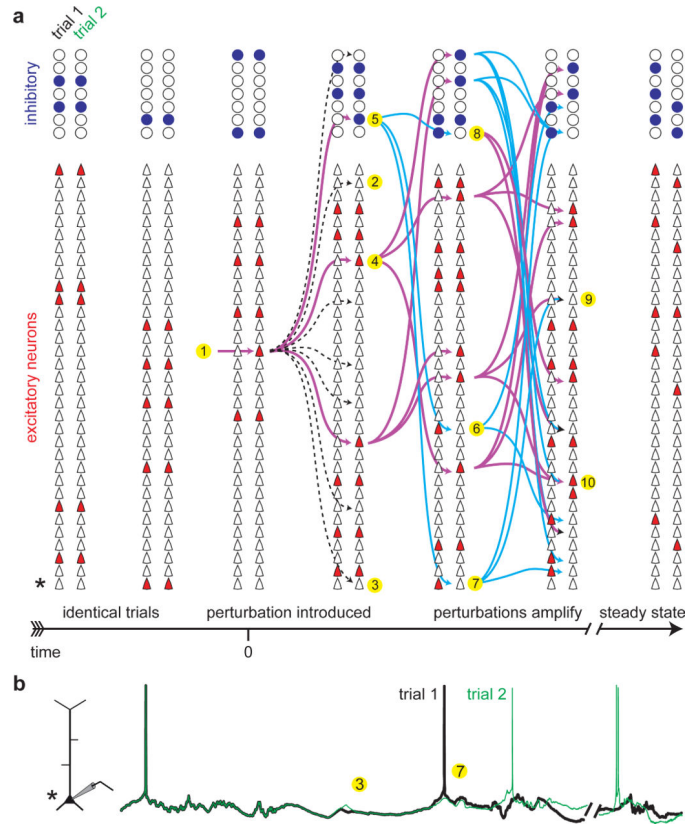


Figure 1. The effect of an extra spike

(a) Schematic showing the propagation of missed and extra spikes in a recurrent network. Each two-column block represents a snapshot of the activity of a population of excitatory and inhibitory neurons on two different trials. Filled neurons are spiking. The trials are identical until time $t=0$, at which point an extra spike is added to a neuron in trial 2 (point 1). The extra spike has no effect on the majority of its postsynaptic targets (dashed arrows and 2,3; omitted subsequently for clarity), but it triggers an extra spike in a fraction of them (magenta arrows and 4,5). These extra spikes cause a cascade of extra and, as soon as inhibitory neurons are recruited, missed spikes (6-8). Extra inhibitory spikes (5) and missed excitatory spikes (6,7) are indicated with cyan arrows. The perturbation amplification rate decreases when collisions occur (9,10), and eventually missed and extra spikes occur at the same rate, resulting in a steady state (rightmost column). (b) Membrane potential of the bottom neuron (*) on the two trials. The membrane potential is identical until an extra presynaptic spike causes a slight divergence (3). As missed and extra spikes accumulate the difference grows (7), until it eventually reaches steady state.

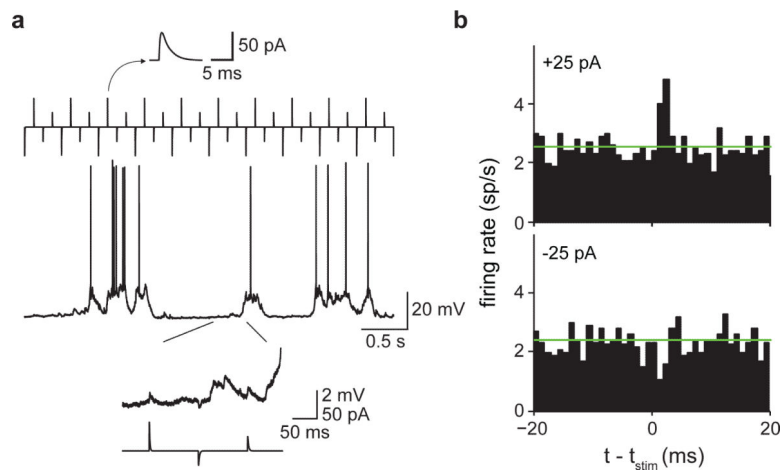


Figure 2. Small perturbations affect spiking probability

(a) Positive and negative current pulses (upper trace) were injected via a whole-cell patch-clamp electrode into a neuron, and the accompanying membrane potential was recorded (lower trace). (b) PSTHs triggered on the +25 pA (top) and -25 pA (bottom) current pulses, binned at 1 ms; the green line shows the average firing rate. For this neuron the probability of an extra spike is 0.004 for +25 pA and the probability of a missed spike is 0.001 for -25 pA.

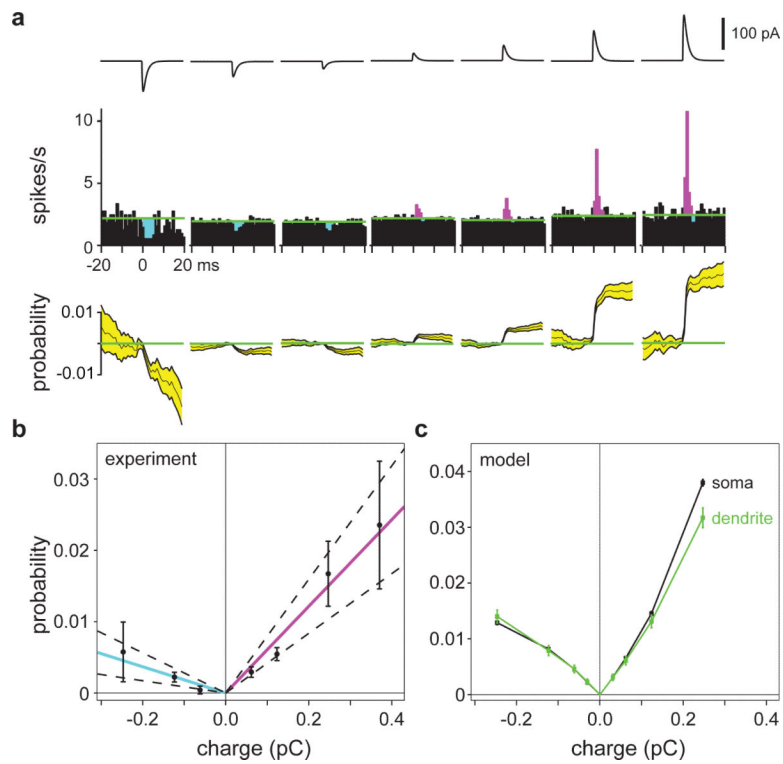


Figure 3. Determining the sensitivity of neurons to small perturbations

(a) Top: injected PSCs. Center: combined PSTHs from 40 experiments, triggered on the current pulses ($t=0$) and binned at 1 ms. Cyan: missed spikes; magenta: extra spikes; green lines: mean firing rate. Bottom: cumulative probability of an extra spike; yellow indicates one standard deviation. (b) Probability of an extra spike within 5 ms of the current pulse versus total injected charge. Magenta and cyan lines: least squares fit to the data; dashed lines: 95% confidence intervals; error bars: standard error of the mean. Positive charge: slope= 0.061 ± 0.010 probability/pC ($p=5 \times 10^{-9}$); negative charge: slope= -0.018 ± 0.0049 probability/pC ($p=4 \times 10^{-4}$). (c) Same as B but from simulations of a layer 5 pyramidal neuron (Supplementary Information, Sec. 4 and Fig. S3). Black: current injected at the soma; green: current injected at the distal dendrites, 403 μm from the soma.

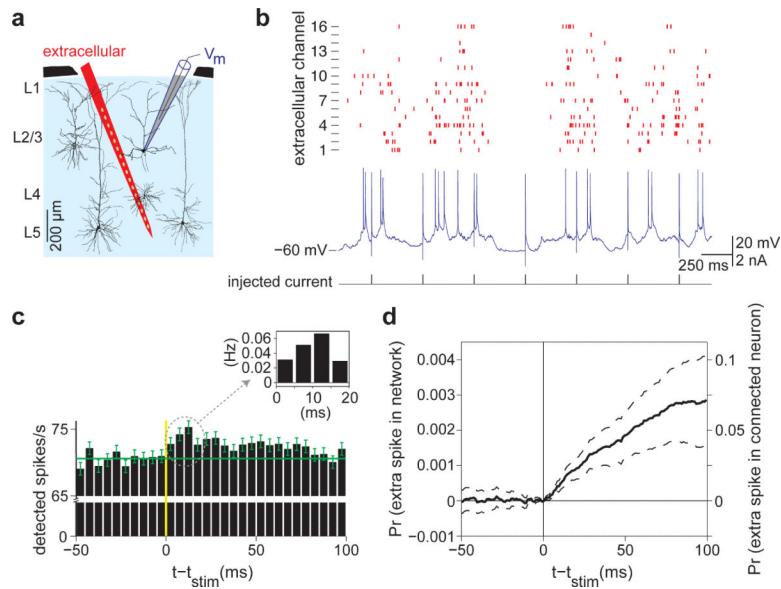


Figure 4. The effect of one extra spike on network activity

(a) The recording configuration. The extracellular silicon probe (red) contained 16 recording sites spaced 50 μm apart. The patch electrode (blue) was used to trigger spikes via brief depolarizing current pulses. (b) Extracellular spikes (top) and intracellular membrane potential (bottom). (c) PSTH triggered on the stimulus and binned at 5 ms; includes all extracellular spikes on all electrodes from 10 experiments. Error bars are one standard deviation. Inset: change in firing rate per neuron, assuming an average firing rate of 1 Hz (ref. 30). (d) Cumulative probability of an extra spike, averaged over all recorded neurons, again assuming an average firing rate of 1 Hz. Dash lines indicate one standard deviation, obtained using bootstrap sampling. Left scale: probability of an extra spike in a randomly chosen neuron. Right scale: probability of an extra spike between connected pairs, found by dividing the left side by 0.04, corresponding to the 4% connectivity observed in somatosensory cortex¹³.




Figure 5. Precisely timed events are rare

(a) Top: example membrane potential recording. A steady current between -250 and -450 pA was applied to reduce the frequency of action potentials. Bottom: expanded view of an up state. The coloured line segments indicate events of various amplitudes and rise times. Green: amplitude 3 mV, rise time 12 ms; cyan: 6 mV, 20 ms; magenta: 6 mV, 5 ms. (b) The rate of events with precision $\delta\tau$ (events for which the voltage change by at least $2\sigma_V$ mV in time $\delta\tau$, averaged over our uncertainty in σ_V , the latter given in Supplementary Figure S7). Data from nine cells recorded *in vivo*, with events counted only during up states (boxed regions in panel a). Three of the cells received, on alternate trials, stimulation via a slowly rotating (1 Hz) drum of sandpaper. There was no statistically significant difference between event rates with and without whisker stimulation.