

Limits to the temporal fidelity of cortical spike rate signals

Mark E. Mazurek and Michael N. Shadlen

Howard Hughes Medical Institute, Department of Physiology and Biophysics, Regional Primate Research Center, University of Washington, Box 357290, Seattle, Washington 98195-7290, USA

Correspondence should be addressed to M.N.S. (shadlen@u.washington.edu)

Published online: 1 April 2002, DOI: 10.1038/nn836

The cerebral cortex processes information primarily through changes in the spike rates of neurons within local ensembles. To evaluate how reliably the average spike rate of a group of cortical neurons can represent a time-varying signal, we simulated an ensemble with realistic spike discharge behavior. We found that weak interneuronal correlation, or synchrony, allows the variability in spike rates of individual neurons to compromise the ensemble representation of time-varying signals. Brief cycles of sinusoidal modulation at frequencies above 115 Hz could not be represented by an ensemble of hundreds of neurons whose interneuronal correlation mimics that of the visual cortex. The spike variability and correlation assumed in our simulations are likely to apply to many areas of cortex and therefore may constrain the fidelity of neural computations underlying higher brain function.

Neurons in many regions of the cerebral cortex alter their rate of spike discharge in response to particular sensory stimuli, body movements or cognitive states. These event-dependent changes in spike rate are typically reproducible, especially when averaged across experimental repetitions. However, the precise timing and number of spikes produced by a single cortical neuron show substantial variability, even under highly controlled sensory or behavioral conditions^{1–3}. This variability in single neurons must be attenuated for the cortex to represent, compute with and transmit information with reasonable speed and accuracy.

One way to distill an accurate rate signal from noisy neurons is to average the responses from many such neurons, thereby obtaining a high-fidelity estimate of the mean spike rate at any moment in time. This seems a reasonable solution, because the cortex is organized into columns comprised of neurons that exhibit similar rate modulations during sensory and behavioral events. This columnar architecture is recognized as a ubiquitous feature of the cortex^{4,5}. It implies that first, any given rate signal is represented by many neurons, and second, each neuron receives input from many of its neighbors, thus obtaining many samples of the signals it uses for computation. In principle, such redundancy could allow the cortex to average away much of the variability seen in individual neurons.

In reality, however, the gain in fidelity achieved by averaging activity from many neurons is likely to be limited by the fact that variability is shared among different neurons. In simultaneous recordings from two nearby neurons, it is often observed that a spike in one neuron predicts a small increase in the probability of a spike occurring in the other (within a few milliseconds). Such spike time correlation, or synchrony, is an active area of research for many laboratories interested in neural coding and computation^{6,7}. One of the many possible functions of correlated spike timing is to curtail the benefits of averaging that would normally accrue for independent signals^{8–10}. For example, measurements

of the number of spikes produced by neighboring pairs of neurons in the visual cortex (area MT) during repeated presentation of visual stimuli yields correlation coefficients of $r \sim 0.15–0.20$ (refs. 9,10). These studies conclude that the accuracy of the average spike count calculated from such weakly correlated neurons would be only 2–3 times better than the accuracy obtained from a single neuron, even if arbitrarily large numbers of neurons were used to compute the average. Similar correlation levels have been reported for nearby neurons in other regions of the cortex^{11–15}, suggesting that weak correlation—and its limiting effect on the fidelity of average spike rate signals—is a common property of cortical ensembles.

Here, we extend this principle to time-varying signals. The spike rate is not a stationary quantity but varies as a function of time to represent changes in the environment, the dynamics of motor function and the evolution of neural computations. We investigated the consequences of variability and correlation on the ability of cortical ensembles to represent time-varying signals through their average spike rate. We developed modeling techniques for simulating the spike discharge from ensembles of neurons whose single-neuron variability and pair-wise correlation in spike timing approximate the known statistics of cortical neurons. We show that correlation within an ensemble causes the rate signal obtained by averaging across the ensemble to retain much of the temporal variability of single neurons. This variability limits the fidelity with which time-varying rate signals can be represented and transmitted in the cortex.

RESULTS

Neural variability and correlation

We devised an efficient procedure to simulate the spike discharge from an ensemble of cortical neurons with realistic single-neuron and pair-wise statistical behavior. Our goal was not to develop a realistic model of synaptic integration, but to mimic the

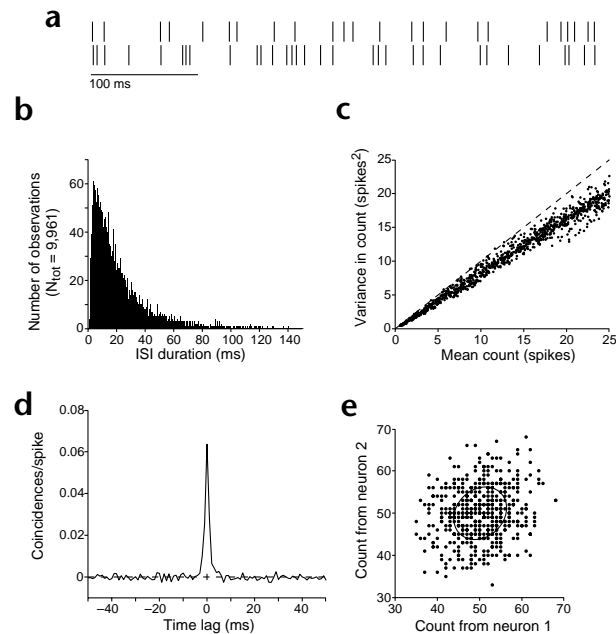


Fig. 1. Model neurons mimic the variability and correlation of neurons in cortex. The examples in this figure were generated using model neurons that discharged at an average rate of 50 spikes/s. **(a)** A pair of spike trains with realistic variability and correlation. Although the simulated spike rate is constant on average, there is considerable variability in spike timing. **(b)** Inter-spike interval (ISI) distribution. The average interval was 20 ms (s.d. = 17.6 ms). **(c)** Relationship between mean and variance in the number of spikes produced in an epoch. Each point shows the variance and mean calculated for 100 repetitions of a fixed simulation epoch; the epoch duration is chosen randomly for each point (range, 10–500 ms). The dashed line of equality is the expected relationship for a Poisson process. **(d)** Correlation in spike timing. The CCG plots the probability of coincidence between spikes from a pair of neurons as a function of the time lag between spikes. The probabilities are normalized so that chance coincidence is assigned a value of 0. The area under the central peak (r_{CCG}) yields a total correlation of 0.2 (see ref. 10). **(e)** Correlation in spike count. Each point represents the number of spikes produced by each of two neurons during a single 1-s epoch. The correlation ellipse is drawn at one standard deviation ($r = 0.2$).


statistics of cortical spike discharge. The method is an extension of so-called diffusion (or random walk) models developed by several investigators to explain the variability of the inter-spike interval^{2,3,16–20}. It is based on the idea that cortical spike times can be likened to the arrival times, or first passage times, of a particle diffusing toward a barrier. Our principal innovation here was to construct n ($\gg 1$) correlated neural responses from the same number of weakly coupled diffusion processes.

For each simulated neuron, a variable broadly analogous to membrane voltage undergoes random steps toward and away from spike threshold in the manner of a one-dimensional random walk. The voltage steps simulate perturbations resulting from synaptic currents arriving at the soma. Each step for the n neurons is drawn from an n -dimensional Gaussian distribution with a mean of zero, thereby mimicking a condition in which excitation and inhibition are in approximate balance³. We manipulated the rate of threshold crossing (spike rate) of a model neuron by controlling the n standard deviation terms, thus altering the average step size of the random walks. Modulation of the standard deviation of voltage step size is analogous to modulation of the input spike rate, again assuming approximate balance of excitation and inhibition. By adjusting the covariance of the multivariate Gaussian distribution, we controlled the tendency for any pair of processes to cross threshold in approximate synchrony. Additional details of the model are described in Methods.

The model neurons showed variability in spike timing similar to that of real cortical neurons. The spike times and intervals produced by model neurons were quite variable, even when the expected rate was held constant as a function of time (Fig. 1a). The inter-spike interval distribution (Fig. 1b) was approximated by an exponential function except at the shortest intervals, which are rare because of refractoriness. Like real cortical neurons, the standard deviation of this distribution was slightly less than the mean interval (c.v. (coefficient of variation) = 0.88; see ref. 1). The number of spikes counted in an epoch showed a constant ratio between the variance and mean (geometric mean ratio = 0.81; Fig. 1c), also consistent with cortical spike statistics. The magnitude of variability was, however, lower by approximately a factor of two than that seen in real cortical neurons^{3,12,21,22} (see Discussion concerning this discrepancy; we note, however, that our underestimation of variability in populations of neurons renders our conclusions more conservative).

The model neurons also mimicked the temporal correlation in spike timing seen in real cortical neurons. This is shown by the cross-correlogram (CCG), which plots the relative probability that two neurons will emit spikes together as a function of the time separating the spikes (Fig. 1d). The narrow central peak cen-

Fig. 2. Correlation causes fluctuations in the ensemble representation of a signal with a constant spike rate. In these examples, each neuron discharges at an average rate of 50 spikes/s. The spike rate is intended to remain constant as a function of time. **(a)** Raster plot (above) and ensemble rate histogram (below) from 100 uncorrelated neurons. The dots in each row of the raster represents the spike times from a single neuron in the ensemble. The rate histogram is computed in 4-ms bins. **(b)** Same for 500 uncorrelated neurons. **(c)** Same for 100 weakly correlated neurons (mean $r = 0.2$ for this simulation). Ensemble rate values above 150 spikes/s are truncated. **(d)** Same for 500 correlated neurons (mean $r = 0.2$).



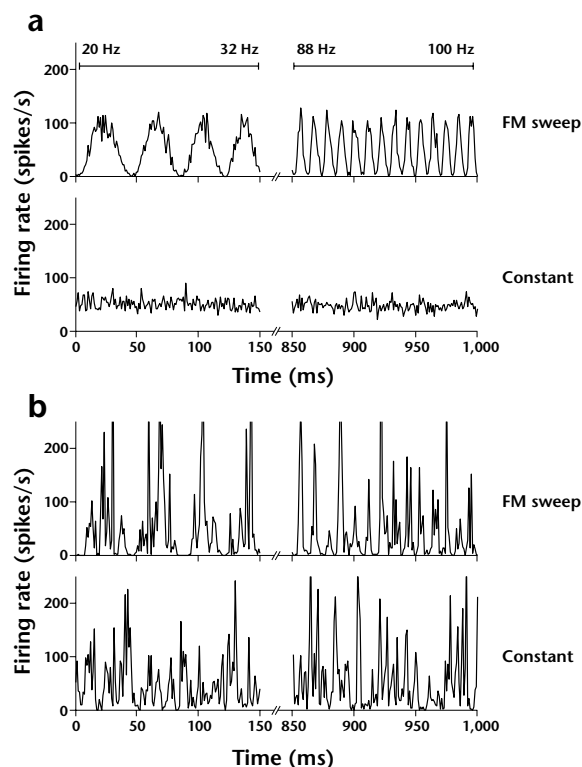


Fig. 3. Random fluctuations in the ensemble spike rate can mimic a time-varying signal. **(a)** Average rate from an ensemble of 500 uncorrelated neurons. Top, ensemble spike rate when the expected rate modulates sinusoidally between 0 and 100 spikes/s. The frequency of the sinusoid sweeps from 20 to 100 Hz in 1 s; only the first and last 150 ms of the sweep are shown. Bottom, same ensemble firing rate when the expected rate is constant at 50 spikes/s. **(b)** Average rate from 500 weakly correlated neurons ($r = 0.2$). Top, ensemble spike rate when the expected rate modulates according to the FM sweep described in **(a)**; bottom, ensemble response when the expected rate is constant at 50 spikes/s.

assumed to reflect the population response of a cortical ensemble on a single experimental trial. We will show, however, that this equivalency rests on the unrealistic assumption that variability between neurons is uncorrelated.

When a realistic degree of correlation is incorporated into the ensemble (Fig. 2b, top), shared fluctuations in spike timing caused the ensemble's average spike rate to undergo large temporal excursions that dominate the 50 spikes/s signal. Moreover, because spike-time variability is shared among the neurons, the variability of this ensemble's output was not reduced substantially by adding more neurons (Fig. 2b, bottom). The magnitude of ensemble rate fluctuations depends on the average pair-wise correlation among neurons in the ensemble. Although each pair of neurons is equally correlated in this example, the fluctuations would be as large if some neurons were more strongly coupled than others, provided the total covariance across the ensemble remained the same. Because the neurons within the ensembles illustrated here show fairly realistic spiking statistics (variability and correlation), the noisy spike rate functions depicted in Fig. 2b represent the kind of signals that might actually arise in cortical ensembles.

Temporal fluctuations degrade time-varying rate signals

We have shown that an ensemble spike rate value that should remain steady at 50 spikes/s instead shows marked temporal variation. Real neural computation involves quantities that change over time, so how do the observed spike rate fluctuations affect the ensemble's representation of time-varying signals? This problem may be illustrated by a neural ensemble representing a sinusoidally modulating signal whose frequency gradually increases (an FM sweep; Fig. 3). To successfully represent the signal, modulations in firing rate must be distinguishable from stochastic fluctuations in the absence of a signal. When the neurons in the ensemble are uncorrelated (Fig. 3a), the modulations in ensemble firing rate are readily apparent, even at the highest modulation rates (100 Hz in this example), and this representation could become arbitrarily more precise with larger ensembles.

However, in the presence of correlation (Fig. 3b), random fluctuations in the average rate obscure the representation of the sinusoidal signal. Slow modulations can still be distinguished from a constant rate, but, at high frequencies, the representation of the sinusoid is difficult to distinguish from the representation of a constant spike rate signal (compare 'FM sweep' to 'constant' in Fig. 3b). In principle, given precise knowledge about the frequency and phase of the sinusoid and the freedom to look over multiple cycles, the FM signal could be extracted from the noisy representation in Fig. 3b. When each cycle must be transmitted without prior knowledge of the underlying signal, however, the random fluctuations in the ensemble firing rate limit the brain's ability to distinguish a rapidly changing signal from one that is constant.

To address this issue more quantitatively, we studied the ensemble representation of brief modulations in rate (Fig. 4). The problem can be formulated as an exercise in signal detec-

tered at zero lag indicates that when one neuron emitted a spike, there was an increased tendency for the other neuron to do so at roughly the same time. The width of the correlogram measured at half height is 2 ms, and the tendency for correlated spike production fell to chance at time lags of ± 5 ms. Paired recordings from nearby cortical neurons often have broader peaks^{10,11,13} (see final section of Results for discussion of the impact of widening the timescale of correlation). One consequence of correlation was that the spike counts from a pair of neurons tended to co-vary in any epoch. The scatter plot (Fig. 1e) shows spike counts from a pair of simulated neurons measured in successive 1-second epochs. The correlation coefficient ($r = 0.2$) is consistent with multiunit recordings from nearby neurons^{9,10,12,14,15,23,24}.

Ensemble rate signals show temporal fluctuations

Our primary aim was to investigate the impact of variability and correlation in spike timing on the representation of time-varying spike rate signals. To illustrate the problem, we first consider the representation of a stationary signal (Fig. 2). Each neuron in the ensemble discharges at a rate of 50 spikes/s on average, but each shows sizeable variability in spike timing as is typical of cortical neurons. If the ensemble is comprised of 100 independent neurons (correlated only by chance), then the instantaneous average spike rate obtained from the ensemble reveals little of the single-neuron variability (Fig. 2a, top). There are small fluctuations in the ensemble rate shown here, but increasing the number of neurons ameliorates this (Fig. 2a, bottom), and the fidelity will continue to improve by \sqrt{n} as more neurons are added. Note that the signal obtained by averaging the output of uncorrelated neurons is equivalent to the average response that is obtained experimentally by recording a single neuron over multiple trials, assuming the same rate function underlies each trial. Thus, the peristimulus time histogram (PSTH) derived from independent repetitions in an experiment is commonly

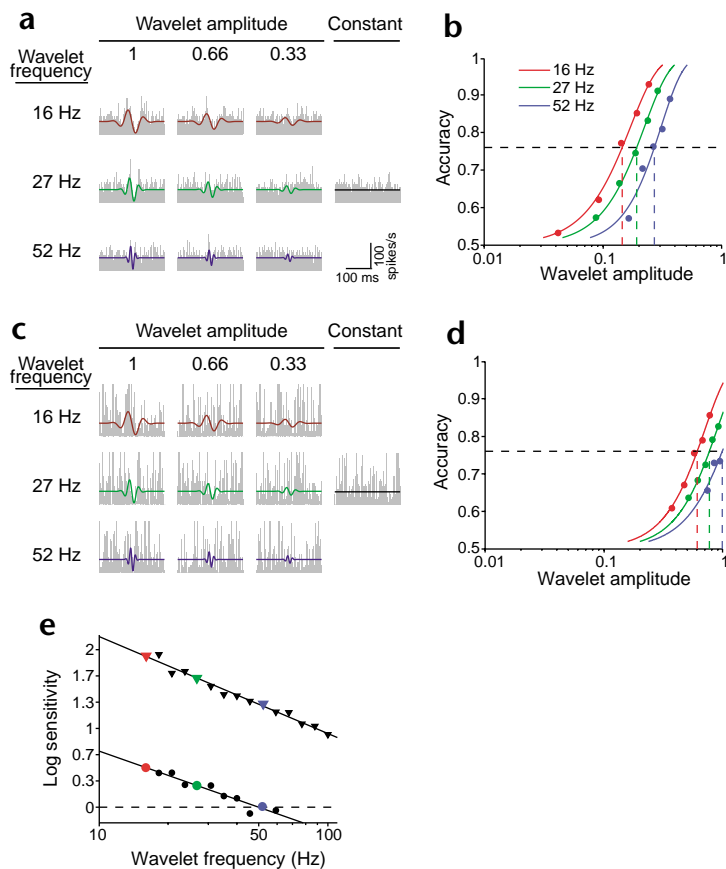


Fig. 4. Correlation-induced fluctuations limit the temporal fidelity of the ensemble spike rate. **(a)** Representation of brief sinusoidal signals by an ensemble of 100 uncorrelated neurons. The solid curves show the expected firing rate at three frequencies and amplitudes of the Gabor Wavelet modulation. Unmodulated rate from a stationary signal is also shown (Constant). Gray histograms show ensemble spike rate. **(b)** Discrimination accuracy using an ensemble of uncorrelated neurons. The frequency with which the Gabor wavelet would be accurately discriminated from a constant spike rate signal was calculated using signal detection theory for each of the wavelets in **(a)**. Threshold amplitude (A_{th}) values were derived from these functions by finding the wavelet amplitude that supports 76% accuracy (colored dashed lines). **(c)** Representation of brief sinusoidal signals by an ensemble of weakly correlated neurons (mean $r = 0.2$). Same conventions as in **(a)**. **(d)** Discrimination accuracy using an ensemble of weakly correlated neurons. Same conventions as in **(b)**. **(e)** Fidelity of the spike rate signal diminishes with briefer, higher-frequency signals. The sensitivity of the brain to ensemble rate fluctuations is estimated from the accuracy curves **(b and d)**. The log of sensitivity ($\log [(A_{th})^{-1}]$) is plotted as a function of wavelet frequency, f . The colored points correspond to the frequencies illustrated in **(a–d)** (uncorrelated ensemble, ∇ ; weakly correlated ensemble, \bullet). Lines are least-square fits to the data; x -intercepts estimate the highest-frequency wavelet that the ensemble can represent with 76% accuracy when the spike rate modulates between 0 and 100 spikes/s. For 100 uncorrelated neurons, the maximum detectable frequency is 576 Hz; this value would increase if the ensemble contained more neurons. The maximum detectable frequency for 100 correlated neurons is 49.8 Hz; this value would not increase with larger ensembles.

tion. If a neural ensemble discharges steadily at a mean firing rate of 50 spikes/s, and an event—a stimulus or the result of some computation—modulates the spike rate up and down with respect to the background rate, how reliably can such modulation be distinguished from the ensemble spike rate fluctuations that arise stochastically?

The intended spike rate modulation in this exercise (the ‘signal’) consists of a brief sinusoid multiplied by a Gaussian envelope (a Gabor wavelet):

$$s(t) = \mu \left[1 + Ae^{-\frac{(t-t_0)^2}{2\sigma^2}} \sin(2\pi ft - \phi) \right] \quad (1)$$

The spike rate, $s(t)$, is μ spikes/s except for an epoch centered at t_0 when it modulates sinusoidally about μ . Therefore, μ determines the dynamic range of the signal, which is at most 2μ spikes/s. A controls the depth of modulation: when $A = 0$, there is no modulation and when $A = 1$, $s(t)$ varies sinusoidally between -0 and $\sim 2\mu$ spikes/s. The frequency and phase of the modulation are controlled by f and ϕ , respectively, and the duration of the modulation is controlled by σ . We used signal durations equal to about 1 period of sinusoidal modulation (see Methods). The Gabor wavelet has several properties that make it suited to studying the representation of time-varying spike rate signals. First, it imposes a balanced increase and decrease in spike rate that would average to zero change if sufficiently blurred. Second, it can be brief enough to introduce a single period of modulation, which is essential for investigating the direct representation of a modulating signal. Third, it is conveniently characterized by the frequency of the underlying sinusoidal

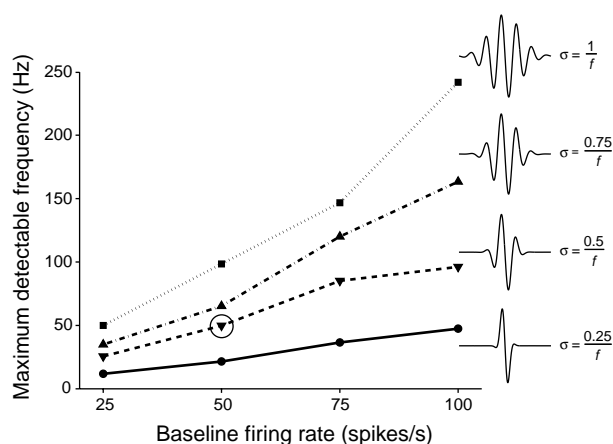
modulation, allowing us to describe the fidelity of the ensemble spike rate as a function of temporal frequency.

Modulation of the ensemble spike rate can be represented faithfully across a broad range of frequencies if neurons emit spikes independently, but not if neurons are weakly correlated. For example, Gabor wavelets at 16, 27 and 52 Hz are all apparent in the discharge of an ensemble of uncorrelated neurons and are readily distinguishable from an unmodulated discharge, even when the modulation amplitude is small (Fig. 4a). In contrast, the wavelets represented by the correlated ensemble tend to be obscured by prominent rate fluctuations (Fig. 4c). These fluctuations make it difficult to distinguish the modulated ensemble rates from the ensemble rate associated with a stationary signal, especially for high-frequency modulations (Fig. 4c). With this ambiguity in the representation of a rapid modulation, the underlying signal is unavailable to neurons reading out the ensemble rate.

To quantify the discriminability of brief sinusoidal modulations from the random variations associated with a stationary rate, we created a detector that takes the ensemble rate as input and produces an output value that reflects the strength of sinusoidal modulation. The detector is optimal for sensing the Gabor wavelet in the spike rate: it has precise knowledge of the time of the wavelet (t_0) and the shape of the waveform (f and σ). Only the phase (ϕ) of the sinusoidal modulation is unknown. The detector produces a scalar output value that measures the amplitude of rate modulation at frequency f in the time window defined by t_0 and σ (namely, an envelope detector; see Methods for additional details). Whether a detector of this sort could be implemented in the brain remains an open question; our aim here was to create a detector at least as sensitive as any neural mechanism. ‘Accurate detection’ was met when our detector



Fig. 5. Ensemble fidelity is affected by firing rate dynamic range and signal width. The graph shows the highest frequency modulation that can be reliably represented by an ensemble of weakly correlated neurons plotted as a function of the baseline spike rate (μ). Each neuron in the ensemble modulates its firing rate about the mean firing rate to represent the Gaussian damped sinusoidal signal (Gabor wavelet; right). The maximum detectable sinusoidal frequency occurs when the full dynamic range of spike rate—from 0 to twice the mean rate—is required to support discrimination of the signal from the mean spike rate with 76% accuracy in our forced-choice exercise. Each curve uses a different ratio of signal width (σ) to sinusoidal period (f^{-1}). For larger ratios (top), the sinusoidal signal contains 2 or more repetitions. The circled point is the frequency limit derived using the parameters shown in Fig. 4.



reported a higher value for an ensemble rate signal containing the Gabor wavelet than for a rate signal of comparable magnitude and duration but containing no modulation (Fig. 4a and c, 'Constant'). The exercise was thus designed to study the spike rate representation of a single cycle of modulation that would be lost if it were smeared over the epoch in which it occurs.

We measured the accuracy of detection across a range of frequencies and amplitudes to find the threshold amplitude (A_{th}) for each frequency, defined as the wavelet amplitude needed to produce detector responses that exceed the noise level (response to no wavelet) by 1 standard deviation, resulting in a discrimination rate of 76% (ref. 25; Fig. 4b and d). A threshold amplitude of 1 implies that the spike rate must vary over the entire available range from 0 to twice the mean rate in order to transmit the wavelet with 76% accuracy; the level of discharge in this exercise (a dynamic range of 50 ± 50 spikes/s) was intended to mimic the high firing rates seen in optimally activated cortical neurons. Hence $A_{th} = 1$ can be interpreted as a rough upper limit for the encoding of a wavelet within a specified dynamic range. The A_{th} values estimated for the uncorrelated ensemble (Fig. 4b) are substantially less than 1 for the frequencies used here, indicating that these wavelets can be represented faithfully. In contrast, the A_{th} values measured

for the correlated ensemble are much higher, approaching 1 at 52 Hz (Fig. 4d, blue curve), indicating that the full dynamic range of the neurons would be required to represent this wavelet.

We used this detection exercise to estimate the temporal fidelity of an ensemble rate code, given our assumptions. To do this, we calculated A_{th} at a variety of modulation frequencies and plotted the reciprocals (sensitivity) against f on a log scale (Fig. 4e). The graph shows that sensitivity falls by the square root of frequency (slope $\approx -1/2$); this is a simple consequence of there being fewer spikes in the briefer signals. Nevertheless, the uncorrelated ensemble in Fig. 4e reliably transmits the rate modulation up to the highest frequencies used here (100 Hz), and the sensitivity would continually improve as more neurons are added to the ensemble (see Fig. 2). In contrast, the correlated ensemble in Fig. 4e exhibits lower sensitivity at all frequencies, failing to reliably transmit any signals faster than about 52 Hz. Because adding more correlated neurons does not improve the fidelity, the observed limitation can be interpreted as an upper bound on the ability of a cortical ensemble like those we have simulated to represent fast changes in spike rate within a range of 0–100 spikes/s. We next consider the dependency of this limitation on the assumptions in our model.

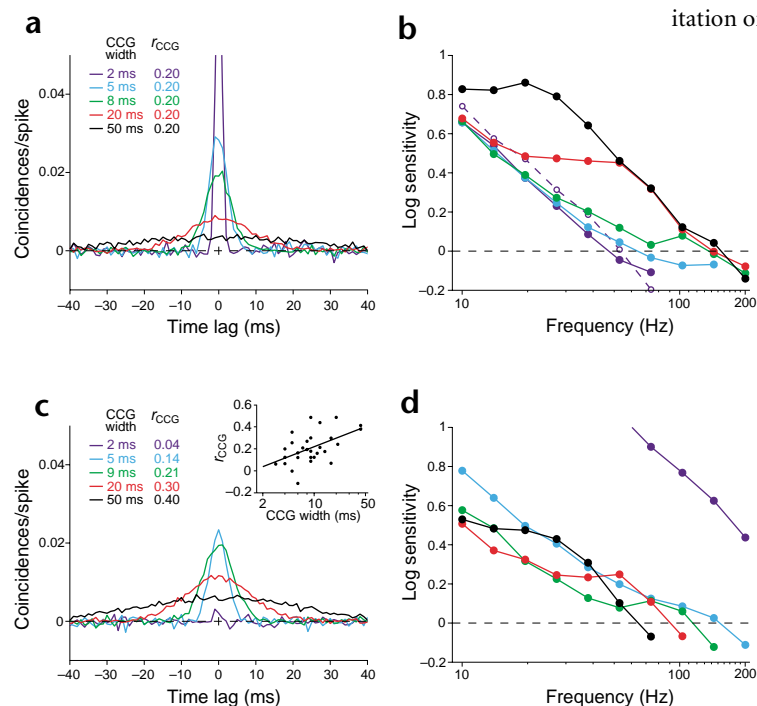


Fig. 6. Ensemble fidelity depends on the time scale and magnitude of correlation. (a) Average CCG for pairs of neurons in each of the five ensembles used to generate the sensitivity plots in (b). The CCG widths at half-height are given in the inset. Each CCG has the same area under the central peak, yielding the same total correlation ($r_{CCG} = 0.2$). Simulations use the method of random spike sequences (see Methods). (b) Fidelity of the ensemble representations of brief sinusoidal rate modulations using neurons whose CCGs are represented in (a). Sensitivity is derived using the same signal detection exercise as in Fig. 4 and plotted as a function of wavelet frequency (same conventions as in Fig. 4e). Solid curves were obtained using CCGs represented in (a) (corresponding colors). The dashed curve shows the result using the linked random walk method (width = 2 ms, $r_{CCG} = 0.2$). (c) Average CCG for pairs of neurons in five ensembles. Wider CCGs have more area such that the total correlation between neuron pairs increases with CCG width. Inset, data from pairs of neurons recorded in area MT; regression line shows the relationship between r_{CCG} and width used to generate the ensembles shown in this panel. (d) Fidelity of the ensemble representations of brief sinusoidal rate modulations using same neurons as in (c).



Temporal sensitivity with more spikes

A decline in sensitivity is inevitable for signals that make use of fewer spikes across the ensemble. We therefore investigated the sensitivity of the correlated ensemble over a range of mean firing rates and wavelet durations. As shown in Fig. 5, sensitivity improved with increased firing rates and with broader wavelets. The improvement along each of these dimensions was a consequence of having more spikes available for performing the discrimination.

The improvement in fidelity obtained with the broadest wavelets (Fig. 5, $\sigma = 0.75/f$ and $\sigma = 1/f$) provides some insight into how neural ensembles could escape the fidelity limits we have derived. For these wavelets, our detector effectively averaged the spike rate across several cycles of sinusoidal modulation and is thus a specialized envelope detector capable of detecting amplitude changes that last for several periods of modulation. For signals of only one cycle, however, an ensemble of model neurons showing variable spike timing and a weak tendency to produce synchronous spikes is unable to reliably represent modulations occurring faster than ~ 100 Hz, even when allowed to modulate between 0 and 200 spikes/s. In the next section, we further refine this limitation by exploring its dependency on the time course of interneuronal correlation.

Temporal sensitivity with longer correlation time

The fidelity limitations exposed in the preceding analyses result from the inability to average out noise that is shared by neurons in the ensemble. This shared noise arises from the tendency of neurons to emit a fraction of their spikes at approximately the same time. Up to now, we have assumed that this tendency occurs over an epoch of a few milliseconds. Although such precisely-timed correlation is observed in cortex^{7,13}, correlation is also commonly found to occur over broader time scales^{10,11,26–28}. Spread out in time, correlated spikes might be less of an impediment for signal averaging within short epochs. To address this issue, we devised a second method for generating ensemble spike discharge. We extended an existing algorithm²⁹ that generates random spike sequences with time-varying rate, to create an ensemble of correlated spike sequences (see Methods).

We used this method, which allows any cross-correlation function, to explore the effect of a broader time-course of correlation on temporal sensitivity (Fig. 6). As a variety of correlation time scales have been reported^{10,11,13,26–28}, we varied this parameter between 2 and 50 ms, measured as the full width at half-height of the CCG (Fig. 6a and c). The literature so far is also unclear on the overall shape of the correlogram and hence the total amount of shared variability (r). We therefore performed our analysis under two sets of assumptions: (i) constant correlogram area in which broader correlation functions are obtained at the expense of height (Fig. 6a–b), and (ii) increasing correlogram area with increasing width (Fig. 6c–d).

The cross correlation functions shown in Fig. 6a encompass the range of correlation widths reported for pairs of nearby neurons^{10,13,14,27}. These functions are scaled along the vertical axis so that the total pair-wise correlation (r_{CCG}), measured as the area under the central peak of the CCG¹⁰, is 0.2. When correlation is nearly synchronous (Fig. 6a and b, blue curves), the sensitivity falls off rapidly as a function of modulation frequency and drops to 0 at ~ 50 Hz (Fig. 6b, dashed curve), similar to the random walk model (Fig. 1e). When correlation is spread out in time, sensitivity declines more gradually, allowing the ensemble to represent higher frequency modulations. For the ensemble with the broadest correlogram (Fig. 6a and b, black curves), rate modulations can be accurately represented up to ~ 150 Hz.

In theory, the cortex could spread out correlated spikes so that their impact on temporal fidelity would be attenuated even further (as long as the total shared variability remained constant). When spikes are correlated over a broader time scale, however, there is likely to be a stronger correlation between neurons. Although this relationship has not been studied systematically in different regions of cortex, measurements from neighboring neurons in area MT show that as the time scale of correlation becomes broader, the magnitude of correlation tends to increase (Fig. 6c, inset; unpublished analysis of data in ref. 10; W. Bair, personal communication). When correlation time and magnitude are coupled in this way (Fig. 6c), the effect of broader correlation is reversed: ensembles with the broadest correlograms exhibit the worst sensitivity because they have the largest magnitude of correlation (Fig. 6d), and the narrowest correlation produces little impairment in fidelity because here r is nearly 0. The correlograms shown here encompass the range observed in area MT, although no correlograms were seen to be as narrow as our narrowest CCG (2 ms). The geometric mean of the data gives a CCG width of 9 ms with $r = 0.21$, yielding a frequency limit of 115 Hz (Fig. 6c and d, green curves). If these data prove to be representative of pairs of neurons in other brain areas, then the cortex would be unable to represent signals that change substantially faster than this value.

DISCUSSION

The weak tendency of neighboring cortical neurons to discharge spikes in a correlated manner has been observed in several brain areas^{9,14,22,30}. Correlation is probably a consequence of shared connectivity^{3,11,12}, but may be important for synaptic integration^{1,31–33} and perceptual and motor function^{34–36}. Whatever its ontogeny and function, correlation represents a departure from independence and is therefore a barrier to signal averaging from multiple neurons, limiting the accuracy of a neural ensemble to represent stimulus features^{3,8–10}.

Our findings extend this principle to time-varying signals. We have shown that the average rate signal derived from weakly correlated neurons undergoes large fluctuations as a function of time, even when the intended signal is changing slowly (Figs. 2 and 3). These fluctuations resemble the brief modulations in spike rate that are generally believed to represent changes in sensory, motor or cognitive signals. If such a signal were to change faster than about one cycle in 10 ms (~ 100 Hz), it would be indistinguishable from noise. Given the rather generous assumptions built into our detection scheme and the conservative values used to derive this limit, our estimate of fidelity may be somewhat exaggerated.

The result can be interpreted as a temporal acuity limit for spike rate signals in cortex. Analogous to acuity measures in other domains, it implies that two events that are represented by changes in the rate of discharge in an ensemble of neurons will be blurred into one if they occur too close together in time (see Supplementary Methods online for further elaboration of this point). This could explain the inability to perceive high rates of light flicker³⁷ or flutter vibration³⁸, and it might pose a limitation to motor dexterity³⁹. Similarly, it could explain why motion-sensitive neurons in area MT fail to represent temporal frequencies greater than about 60 Hz (ref. 40).

Our findings expose a commonly overlooked inconsistency between signals carried in the cortex and data derived from single-neuron recording experiments. Most cortical physiology experiments involve recording from a neuron during repetitions of sensory stimuli or behavioral responses. The spike discharge from the neuron is typically averaged across many trials to generate a

peristimulus time histogram (PSTH) that furnishes an estimate for the instantaneous firing rate as a function of time. Because of the redundancy of synaptic organization in the cortex, the spike rate function from one neuron probably approximates the expectation of the rate-versus-time function for each of the neurons in a local ensemble. However, the PSTH would depict an actual signal present in the brain if and only if the repeated measurements from one neuron are equivalent to the average response from many neurons on a single trial. Weak correlation between neurons implies that this assumption (ergodicity) is incorrect; thus the PSTH substantially underestimates the variability of the ensemble.

Our conclusions rest on several approximations that warrant examination. First, the variability in spike discharge from single neurons was approximated by a renewal process (linked random walks) or by random distribution of events (random spike sequences). These two simulation methods differ in their mechanics, and each has its advantages: the random walk method simulates a process analogous to membrane voltage and assigns no distinct status to correlated versus uncorrelated spikes, whereas the random spike sequence method allows correlated spikes to be distributed in time. The similarity in the performance of these models (Fig. 6b) suggests that our findings do not depend on the details of the simulation. Neither approximation is correct⁴¹, however, and the random walk method, in particular, underestimates the variability measured in cortex. Real neurons are likely to exhibit greater variability, in part because spike rate is determined from inputs that are themselves variable³. In principle, this discrepancy is conservative with respect to our conclusions because any additional variability should further limit the ability of the cortical ensemble to represent time varying changes in rate.

Second, for most calculations, we allowed firing rates to vary between 0 and 100 spikes/s. It is possible that neurons could achieve higher firing rates and thereby escape the limitations derived here (Fig. 5). As most cortical neurons do not achieve this range except under ideal circumstances, our choice of dynamic range was also conservative. Real ensembles of weakly coupled neurons are likely to respond over a narrower dynamic range, on average, and therefore achieve lower temporal fidelity than we estimated.

Third, we mainly considered a narrow time scale for correlated spike discharge, corresponding loosely to what is termed synchrony. As shown in Fig. 6, if pairs of neurons share variability over a more prolonged time interval, then the effective correlation is lessened in any brief epoch, enabling the ensemble to represent rapid changes in rate. Data from area MT, however, shows that neurons with broader correlation times tend to also have higher total correlation, leading to lower overall fidelity (Fig. 6c and d). The correlations we used in our simulations encompass the range observed in area MT, as the correlations between neurons in the same column are relevant to this study.

Finally, both simulation methods made the conservative approximation that correlation is achieved without temporal smearing. An alternative method for broadening the correlation would be to introduce temporal correlation in the individual spike trains (for example, by smoothing the steps in the random walk¹⁹). In separate simulations (data not shown), we have seen that broadening the correlation time in this way produces lower temporal sensitivity, as smoothing directly attenuates the representation of high-frequency signals.

The limitations in signal fidelity that we describe here apply to any computation in cortex that uses the average spike rate from a network of weakly correlated neurons. Our findings do not imply that neurons are unable to represent signals at higher fre-

quencies, only that they must do so using means other than the ensemble spike rate (such as auditory processing in brainstem nuclei⁴²). Indeed, it has been suggested that computations more sophisticated than averaging could actually benefit from the presence of correlated variability⁴³. So why worry about a fidelity limit for averaged spike rates? There are two reasons. First, evidence from recording and stimulation experiments indicates that sensory signals encoded in the spike rate of nearby neurons directly affect perception^{44–47}, and it is among such ensembles of nearby neurons that correlated discharge has been demonstrated^{9,13,14,23}. Second, temporal modulations in spike rate cannot be detected during the inter-spike interval from a single neuron. Thus, whenever spike rate enters a computation, however sophisticated, it is likely to do so as an ensemble average.

The fidelity limits exposed here contrast with the speeds of processing typically seen in computing devices. These limits arise from the variability and correlation in spike timing that result from the highly interconnected architecture of cortical columns^{2,3}. Why would the brain employ a computational architecture possessing such a sluggish CPU? The answer has been alluded to already: the abundance of connections within a column causes each cortical neuron to receive a barrage of synaptic input for each spike it produces⁴⁸, which furnishes it with a nearly continuous estimate of incoming rate signals, even when individual inputs are firing at low rates. This, in turn, allows cortical neurons to perform powerful and flexible computations in real time using signals produced by other cortical neurons. Limited temporal fidelity thus seems to be a necessary consequence of the remarkable computational complexity in the cerebral cortex.

METHODS

Random walk model. The spike train from each simulated neuron is modeled by the first passage times of a diffusion process defined as follows: a state variable, $V_k(t)$, analogous to the membrane voltage for neuron k , wanders randomly between a lower reflecting boundary and an upper absorbing barrier representing spike threshold. Every 0.1 ms, the value of V_k changes by a random value $z_{k,t}$ chosen from a normal distribution with mean 0 and standard deviation ζ_k , abbreviated $N(0, \zeta_k)$. In addition, the state variable decays toward the lower boundary with time constant of 10 ms. Whenever $V_k(t)$ crosses the upper threshold, a spike occurs and the state variable is reset to the lower reflecting boundary. A brief refractory period is imposed by raising the threshold barrier and triggering a first-order decay (time constant 0.5 ms). The rate of threshold crossings is controlled by changing the magnitude of ζ_k ; the exact relationship between spike rate and ζ_k was derived empirically through simulation, and controlled as a function of time to produce a desired time-varying spike rate signal. The spike times from this model approximate first passage times of a previously defined stochastic process^{17,19}.

We produced an ensemble of neural spike trains by simulating n diffusion processes at once ($k = 1 \dots n$). Using bold to represent n -dimensional vectors, we simply update all the $v(t)$ at each time step by picking \mathbf{z}_t from an n -dimensional multivariate normal distribution, $N(0, \Sigma)$, defined by the density

$$f(\mathbf{z}) = (2\pi)^{-\frac{n}{2}} |\Sigma|^{-\frac{1}{2}} e^{-\frac{1}{2}\mathbf{z}\Sigma^{-1}\mathbf{z}} \quad (2)$$

where Σ is an n by n covariance matrix,

$$\Sigma = \begin{pmatrix} \zeta_1^2 & \rho_{1,2}\zeta_1\zeta_2 & \cdots & \rho_{1,n}\zeta_1\zeta_n \\ \rho_{1,2}\zeta_1\zeta_2 & \zeta_2^2 & & \rho_{2,n}\zeta_2\zeta_n \\ \vdots & & \ddots & \\ \rho_{1,n}\zeta_1\zeta_n & \rho_{2,n}\zeta_2\zeta_n & & \zeta_n^2 \end{pmatrix} \quad (3)$$

Each of the $z_{k,t}$ that provides the random steps for any one diffusion process

is the k^{th} marginal distribution of the multivariate normal. By choosing suitable values for ρ_{ij} , we can impose a desired degree of correlation among the diffusion processes. Intuitively, ρ_{ij} is related to the fraction of inputs shared by neurons, but it should not be confused with the correlation coefficient, r , that describes the relationship between spike counts from pairs of neurons. An efficient method for choosing \mathbf{z} in each time step is to choose a vector \mathbf{x} of n independent values from the normal distribution, $N(0,1)$, and multiplying \mathbf{x} by the matrix square root of Σ :

$$\mathbf{z} = \Sigma^{\frac{1}{2}} \mathbf{x} \quad (4)$$

The elements of \mathbf{z} (distributed as $N(0,1)$ but with covariance Σ) are then used to perturb each of the k diffusion processes. An advantage of the random walk method is that spiking correlation arises as a result of covariance in the underlying Gaussian perturbations; hence, correlated and uncorrelated spikes are produced by the same process. A disadvantage is that it gives rise to synchronous spikes over a short time scale (~ 2 ms).

Random spike sequences. We used a statistical procedure for generating random sequences of spike times that conform to a desired function of spike rate versus time²⁹. The procedure distributes spike times according to a probability density defined by the desired spike rate function. This is achieved by using inverse probability⁴⁹ (generating uniform deviates and solving numerically for the inverse of the cumulative spike rate distribution) and repeating the procedure to produce the desired number of spikes. To impose a refractory period, a spike that falls within t_{ref} of another spike is redrawn with probability p_{ref} that is greatest for small values of t_{ref} and falls to zero after 2 ms. We have extended this procedure to simulate an ensemble of neurons and modified it to achieve correlation between the neurons. This modified procedure allows us to distribute correlated spikes over any desired time scale.

First, the total number of spikes is selected for each neuron. The spike count for the k^{th} neuron, denoted c_k , is normally distributed with expectation and variance equal to the mean spike rate (m) multiplied by the simulated duration, T . Correlation in spike count across the ensemble is achieved by drawing \mathbf{c} (the n -dimensional vector of counts, c_k) from a multivariate normal distribution, $N(mT, \Sigma)$, where mT is an n -dimensional vector of mT values and Σ is an n by n covariance matrix with terms mT on the diagonal and $0.2T\sqrt{(m_i m_j)}$ for the off-diagonal (covariance) terms. This ensures that on average the counts have variance equal to the mean and correlation coefficient equal to 0.2.

Each neuron's spike sequence is then constructed from two sources. A proportion of the neuron's spikes (u_k) are assigned random times using the method just described; hence, the times of these spikes are independent from one neuron to the next. To generate correlated spikes, the remaining $v_k = c_k - u_k$ spikes are assigned times by drawing them from a common source. To obtain this common source of spikes, a dummy spike sequence with mT spikes is generated in the manner just described; this dummy sequence is not included in the ensemble rate. For each neuron, we choose v_k random spikes from the dummy sequence. The fraction of spikes chosen from the dummy train (v_k/c_k) is determined empirically so that the area under the central peak of the correlogram accounts for the magnitude of the correlation between spike counts (r)¹⁰. For $r = \{0.04, 0.14, 0.2, 0.21, 0.3, 0.4\}$ as shown in Fig. 6, we use $v_k/c_k = \{0.08, 0.42, 0.48, 0.51, 0.565, 0.62\}$. To broaden the timescale of correlation, we perturb the times of the spikes chosen from the dummy sequence by adding a random offset drawn from $N(0, \sigma_c)$. The standard deviation of this normal distribution, σ_c , is determined empirically to achieve desired CCG widths. For the CCGs shown in Fig. 6, width = $\{2, 5, 8, 9, 20, 50\}$ ms requires $\sigma_c = \{0.5, 1.7, 2.5, 2.9, 6.2, 14.8\}$ ms. Thus the spike train from each neuron consists of some spikes that are independent (except by virtue of their shared spike rate function) of other neurons in the ensemble, and some spikes that are related to a common source.

Signal detection in ensemble firing rates. To measure how accurately a simulated neural ensemble represents a Gabor wavelet in its average rate, we generated wavelets of width ~ 1 cycle of modulation. To detect the presence of such a wavelet in the ensemble firing rate, we constructed an 'envelope detector' that assumes prior knowledge of the frequency and time

course of the wavelet; only the phase is assumed to be unknown. Application of the detector yields a scalar value, d , that can be used to gauge the strength of the signal in ensemble rate when it is modulated by the Gabor wavelet $s(t)$ with parameters f, A, μ , and σ (Eq. 1):

$$d = \int_{t_0 - 3\sigma}^{t_0 + 3\sigma} \sqrt{[r(t)s(t)]^2 + [r(t)s'(t)]^2} dt \quad (5)$$

where $r(t)$ is the ensemble-average firing rate, $s(t)$ is the Gabor wavelet and $s'(t)$ is the Gabor wavelet with its phase shifted by 90° . Note that $s(t)$ serves a dual role: it specifies the shape of the Gabor wavelet and serves as a template for our detector. We performed ~ 500 simulations to estimate the distribution of d for a set of parameters (f, A, μ , and σ , see Eq. 1) and compared the distribution to the d values obtained using the same template but using a constant spike rate μ with no signal (for example, $A = 0$). The probability of a correct detection of the wavelet is given by the area under a receiver operator characteristic (ROC) curve constructed from the distributions of d under signal and no signal²⁵. This is simply the probability that the detector gives a larger value when the wavelet is present than when it is absent.

Note: Supplementary information is available on the Nature Neuroscience website.

Acknowledgments

Supported by HHMI, NIH grants RR00166 and EY11378 and the McKnight Foundation. M.E.M. was supported by an NIH training grant (GM07108), a Poncin grant and an ARCS fellowship. We thank C. Brody, J. Ditterich, J. Gold, M. Leon and M. McKinley for comments and discussion.

Competing interests statement

The authors declare that they have no competing financial interests.

RECEIVED 22 JANUARY; ACCEPTED 1 MARCH 2002

- Softky, W. R. & Koch, C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.* 13, 334–350 (1993).
- Shadlen, M. N. & Newsome, W. T. Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* 4, 569–579 (1994).
- Shadlen, M. N. & Newsome, W. T. The variable discharge of cortical neurons: implications for connectivity, computation and information coding. *J. Neurosci.* 18, 3870–3896 (1998).
- Braitenberg, V. & Schuz, A. *Anatomy of the Cortex: Statistics and Geometry* (Springer, Berlin, 1991).
- Albright, T. D., Jessell, T. M., Kandel, E. R. & Posner, M. I. Neural science: a century of progress and the mysteries that remain. *Neuron* 25 Suppl, S1–S55 (2000).
- Shadlen, M. N. & Movshon, J. A. Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* 24, 67–77 (1999).
- Singer, W. Neuronal synchrony: a versatile code for the definition of relations? *Neuron* 24, 49–65 (1999).
- Johnson, K. O. Sensory discrimination: neural processes preceding discrimination decision. *J. Neurophysiol.* 43, 1793–1815 (1980).
- Zohary, E., Shadlen, M. N. & Newsome, W. T. Correlated neuronal discharge rate and its implications for psychophysical performance [published erratum appears in *Nature* 1994 Sep 22;371(6495):358]. *Nature* 370, 140–143 (1994).
- Bair, W., Zohary, E. & Newsome, W. T. Correlated firing in Macaque visual area MT: time scales and relationship to behavior. *J. Neurosci.* 21, 1676–1697 (2001).
- Fetz, E., Toyama, K. & Smith, W. in *Cerebral Cortex* (eds. Peters, A. & Jones, E. G.) 1–47 (Plenum, New York, 1991).
- Gawne, T. J. & Richmond, B. J. How independent are the messages carried by adjacent inferior temporal cortical neurons? *J. Neurosci.* 13, 2758–2771 (1993).
- Das, A. & Gilbert, C. D. Receptive field expansion in adult visual cortex is linked to dynamic changes in strength of cortical connections. *J. Neurophysiol.* 74, 779–792 (1995).
- Lee, D., Port, N. L., Kruse, W. & Georgopoulos, A. P. Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J. Neurosci.* 18, 1161–1170 (1998).
- Salinas, E., Hernandez, A., Zainos, A. & Romo, R. Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. *J. Neurosci.* 20, 5503–5515 (2000).



16. Gerstein, G. & Mandelbrot, B. Random walk models for the spike activity of a single neuron. *Biophys. J.* 4, 41–68 (1964).
17. Ricciardi, L. & Sacerdote, L. The Ornstein-Uhlenbeck process as a model for neuronal activity. *Biol. Cybern.* 35, 1–9 (1979).
18. van Vreeswijk, C. & Sompolinsky, H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274, 1724–1726 (1996).
19. Shinomoto, S., Sakai, Y. & Funahashi, S. The Ornstein-Uhlenbeck process does not reproduce spiking statistics of neurons in prefrontal cortex. *Neural Comput.* 11, 935–951 (1999).
20. Salinas, E. & Sejnowski, T. J. Impact of correlated synaptic input on output firing rate and variability in simple neuronal models. *J. Neurosci.* 20, 6193–6209 (2000).
21. Werner, G. & Mountcastle, V. B. The variability of central neural activity in a sensory system, and its implications for the central reflection of sensory events. *J. Neurophysiol.* 26, 958–977 (1963).
22. Gawne, T. J., Kjaer, T. W., Hertz, J. A. & Richmond, B. J. Adjacent visual cortical complex cells share about 20% of their stimulus-related information. *Cerebral Cortex* 6, 482–489 (1996).
23. van Kan, P. L. E., Scobey, R. P. & Gabor, A. J. Response covariance in cat visual cortex. *Exp. Brain Res.* 60, 559–563 (1985).
24. Lampl, I., Reichova, I. & Ferster, D. Synchronous membrane potential fluctuations in neurons of the cat visual cortex. *Neuron* 22, 361–374 (1999).
25. Green, D. M. & Swets, J. A. *Signal Detection Theory and Psychophysics* (Wiley, New York, 1966).
26. Vaadia, E. *et al.* Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature* 373, 515–518 (1995).
27. Brody, C. D. Disambiguating different covariation types. *Neural Comput.* 11, 1527–1535 (1999).
28. Brody, C. D. Correlations without synchrony. *Neural Comput.* 11, 1537–1551 (1999).
29. Oram, M. W., Wiener, M. C., Lestienne, R. & Richmond, B. J. Stochastic nature of precisely timed spike patterns in visual system neuronal responses. *J. Neurophysiol.* 81, 3021–3033 (1999).
30. Ts'o, D. Y., Gilbert, C. D. & Wiesel, T. N. Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.* 6, 1160–1170 (1986).
31. Murthy, V. N. & Fetz, E. E. Effects of input synchrony on the firing rate of a three-conductance cortical neuron model. *Neural Comput.* 6, 1111–1126 (1994).
32. Stevens, C. F. & Zador, A. M. Input synchrony and the irregular firing of cortical neurons. *Nat. Neurosci.* 1, 210–218 (1998).
33. Usrey, W. M., Alonso, J. M. & Reid, R. C. Synaptic interactions between thalamic inputs to simple cells in cat visual cortex. *J. Neurosci.* 20, 5461–5467 (2000).
34. Engel, A. K., Roelfsema, P. R., Fries, P., Brecht, M. & Singer, W. Role of the temporal domain for response selection and perceptual binding. *Cerebral Cortex* 7, 571–582 (1997).
35. Riehle, A., Grün, S., Diesmann, M. & Aertsen, A. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278, 1950–1954 (1997).
36. Di Lorenzo, P. M. & Lemon, C. H. The neural code for taste in the nucleus of the solitary tract of the rat: effects of adaptation. *Brain Res.* 852, 383–397 (2000).
37. Lumer, E. D., Friston, K. J. & Rees, G. Neural correlates of perceptual rivalry in the human brain. *Science* 280, 1930–1934 (1998).
38. Mountcastle, V. B., Talbot, W. H., Sakata, H. & Hyvärinen, J. Cortical neuronal mechanisms in flutter-vibration studied in unanesthetized monkeys. Neuronal periodicity and frequency discrimination. *J. Neurophysiol.* 32, 452–484 (1969).
39. Harris, C. M. & Wolpert, D. M. Signal-dependent noise determines motor planning. *Nature* 394, 780–784 (1998).
40. Bair, W. & Koch, C. Temporal precision of spike trains in extrastriate cortex. *Neural Comput.* 8, 1185–1202 (1996).
41. Teich, M. C., Heneghan, C., Lowen, S. B., Ozaki, T. & Kaplan, E. Fractal character of the neural spike train in the visual system of the cat. *J. Opt. Soc. Am. A* 14, 529–546 (1997).
42. Oertel, D. The role of timing in the brain stem auditory nuclei of vertebrates. *Annu. Rev. Physiol.* 61, 497–519 (1999).
43. Abbott, L. F. & Dayan, P. The effect of correlated variability on the accuracy of a population code. *Neural Comput.* 11, 91–101 (1998).
44. Britten, K. H., Shadlen, M. N., Newsome, W. T. & Movshon, J. A. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* 12, 4745–4765 (1992).
45. Salzman, C. D., Murasugi, C. M., Britten, K. H. & Newsome, W. T. Microstimulation in visual area MT: effects on direction discrimination performance. *J. Neurosci.* 12, 2331–2355 (1992).
46. Hernandez, A., Zainos, A. & Romo, R. Neuronal correlates of sensory discrimination in the somatosensory cortex. *Proc. Natl. Acad. Sci. USA* 97, 6191–6196 (2000).
47. Romo, R., Hernandez, A., Zainos, A., Brody, C. D. & Lemus, L. Sensing without touching: psychophysical performance based on cortical microstimulation. *Neuron* 26, 273–278 (2000).
48. Destexhe, A. & Pare, D. Impact of network activity on the integrative properties of neocortical pyramidal neurons *in vivo*. *J. Neurophysiol.* 81, 1531–1547 (1999).
49. Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. *Numerical Recipes in C: The Art of Scientific Computing* (Cambridge Univ. Press, 1992).

Appendix

Here, we examine the representation of transient elevations in spike rate (pulses) by ensembles of spiking neurons. We develop an exercise that complements the representation of brief sinusoidal modulations in spike rate considered in the main text. Like the exercise in the main text, we compare the ensemble representation of modulating and stationary signals over a time scale that the brain might use for computations.

The signals that are compared in this exercise consist of either a single pulse of variable duration, or a pair of pulses, each 2 ms wide, separated by a gap of variable duration (Δt). The average spike rate is 100 spikes/s during the pulse(s) and 0 spikes/s otherwise. These two signal types are illustrated at the top of **Fig. S1**.

The problem we need to consider is not whether the pulse (P) and pulse-gap-pulse (PgP) are discriminable—they are discriminable on the basis of total spikes—but whether the ensemble rate contains the right number of bumps. Specifically, for the ensemble representation to be useful it should be possible to detect two modes in the ensemble spike histogram representation of PgP and just one in the representation of P. Of course it is hard to detect 2 modes when Δt is small and it is hard to detect only one mode when Δt is large. In the present example, we consider a narrow range of Δt from 3 to 50 ms. We make the assumption that neural circuits must be capable of using both P and PgP signals across this time span. In other words, we consider neural computations that would utilize signals that change faster than ~ 20 Hz.

For each Δt , we generated ensemble representations of P and PgP using the method of random spike sequences (see Methods). The ensembles contained 500 neurons with a CCG shape represented by the green trace in **Fig. 6c** ($r = 0.2$, width = 9 ms). **Fig.**

S1 shows the ensemble representation of P and PgP for gaps ranging from 10 to 50 ms. Four samples of the ensemble discharge are shown for each signal. We examined the ensemble rate functions for evidence of bimodality by applying Harnad's dip test¹. To apply the test, the ensemble rate function is interpreted as a frequency distribution of spike times. If the test is positive ($p < 0.05$) the representation is classified as multimodal (red in **Fig. S1**). We repeated the test 100 times under each condition to estimate the probability that the ensemble representation would contain more than one mode (values in parentheses in **Fig. S1**). The dip test is ideally suited for this problem because it is based on the null hypothesis that the unimodal distribution is uniform, like P. It does not make any assumptions about the bin size in the frequency histograms because it works on the cumulative distribution of spike times, each of which is represented with a precision of 0.1 ms. Also, for present purposes it helps that the test is fairly conservative (e.g., it fails to detect bimodality in a samples of 1000 random numbers described by an equal mixture of two Gaussian distributions unless they are separated by at least 3σ).

The rate histograms in **Fig. S1** are colored red if they contain more than one mode ($p < 0.05$). Ideally, all signals on the left side should be red whereas those on the right should be black. When no smoothing of the rate function was performed (**Fig. S1a**), almost all representations of PgP with $\Delta t > 3$ ms gave ensemble responses which were appropriately classified as multimodal. However, the representation of the longer P also contains several erroneous peaks which arise because of the weak correlation between neurons in the ensemble. The P lasting 54 ms (i.e., $\Delta t = 50$) was misclassified in 78% of the samples, and $>90\%$ of pulses lasting 75-100 ms contained at least two distinct modes. This is not acceptable because any brain mechanism that would care about detecting 2

events over short epochs could not afford to be fooled by the existence of similar events in a 50 ms time span. The problem is easily remedied.

A simple means for attenuating the erroneous peaks in the ensemble response is to blur the rate signal. To achieve this in our simulations, we replace each spike that occurs at t_i with one that occurs at $t_i + \xi$, where ξ is a random number drawn from an exponential distribution with mean τ . The process is equivalent to smoothing the rate function by a filter with time constant τ . After smoothing with $\tau = 10$ ms, the longer pulse is almost always classified as unimodal (**Fig. S1b**). However, this smoothing leads to a higher probability that PgP will be misclassified as unimodal when Δt is short. Accordingly, the neural ensemble can now represent signals lasting at least 50 ms but it can no longer represent changes occurring in less than 10 ms. Thus we have achieved accurate representation at long time scales at the cost of accurately representing brief changes. We obtained similar results using ensemble sizes of 100, 200 and 500 neurons and using a narrower CCG width. For the narrowest case (2 ms half width, identical to that used in **Figs. 1-5** of the main text) the erroneous modes are narrower and more abundant for $\Delta t = 50$ ms. Nevertheless, smoothing with a time constant of 10 ms permits accurate representation of the longer pulses.

This illustration complements the exercise with brief sinusoidal modulations pursued in the main text. In both cases, fluctuations in ensemble spike rate resulting from the weak coupling between neurons masquerade as fast changes in spike rate. In the exercise using sinusoidal wavelets, we found that modulations occurring faster than ~ 100 Hz would not be distinguishable from the fluctuations accompanying a constant spike rate. In the present example, we smoothed away these noise fluctuations at the cost of

representing brief signals separated by 10 ms or less. We may conclude from this exercise that the brain would do best to ignore rate fluctuations on this time scale. Importantly, neither exercise implies any limitation on temporal comparisons or on the precision in estimating the time of a signal. For example, the center of a blurred pulse can be estimated with arbitrary precision. Both exercises expose a limit on the ability to encode fluctuations in variables that are represented by ensembles of neurons: any signal that would be lost upon averaging over an epoch of ~10 ms will fail to be represented or transmitted by the cortex.

1. Hartigan, J. A. & Hartigan, P. M. The dip test of unimodality. *The Annals of Statistics* **13**, 70-84 (1985).

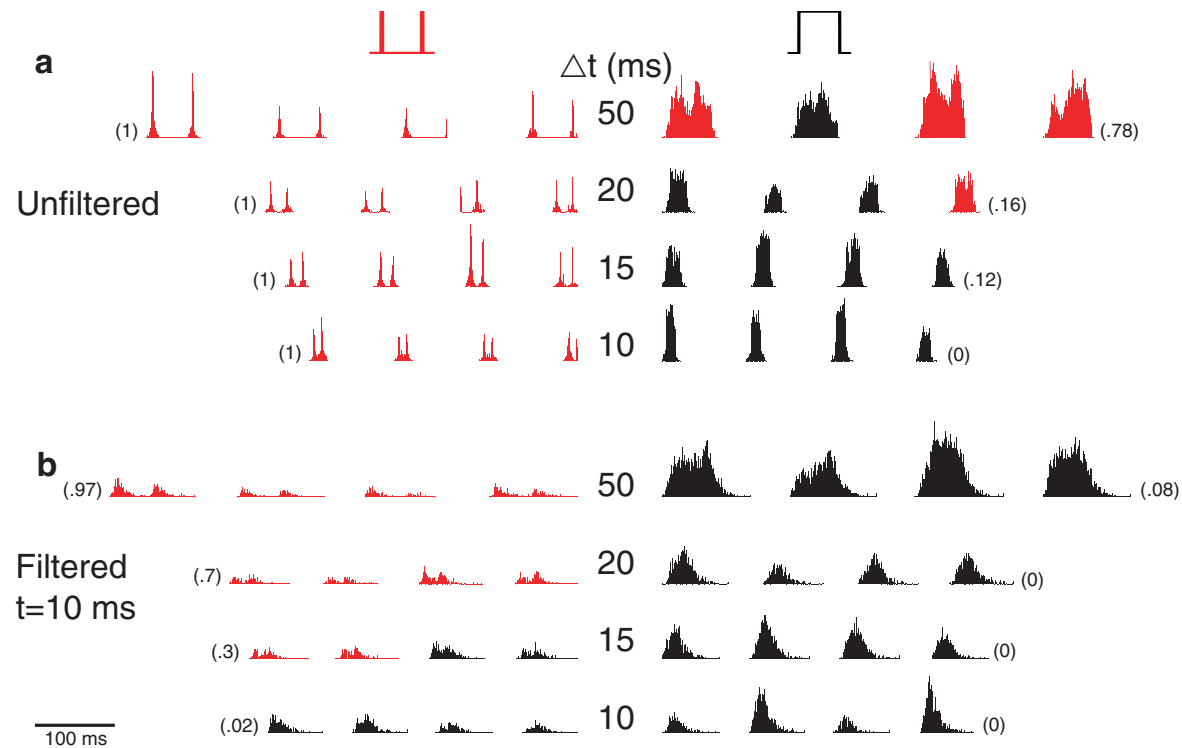


Fig. S1. Simulated ensemble representation of pulse-gap-pulse and single-pulse signals. Numbers in the center column indicate the gap duration. The single pulse signals have matching duration to gap plus surrounding 2 ms pulses. Ensemble spike histograms were generated using 500 weakly correlated neurons using the method of random spike sequences ($r = 0.2$; CCG width = 9 ms). Numbers in parentheses indicate the probability that the distribution of ensemble spike times contains two or more modes (Harnad's dip test, $p < 0.05$). Histogram binwidth is 1 ms. a, Raw responses. Four examples are shown for the two signal types at each value of Δt . Those with two distinct modes are colored red. Notice that many responses on the right side exhibit more than one mode. b, Smoothed responses. Smoothing is achieved by delaying each spike by a random time drawn from an exponential distribution ($t = 10$ ms). This degree of smoothing attenuates the fluctuations in the ensemble representation of longer pulses at the expense of the shortest pulse-gap-pulse.