

LETTERS

Temporal precision in the neural code and the timescales of natural vision

Daniel A. Butts^{1,2}, Chong Weng³, Jianzhong Jin³, Chun-I Yeh^{3,4}, Nicholas A. Lesica¹, Jose-Manuel Alonso³ & Garrett B. Stanley¹

The timing of action potentials relative to sensory stimuli can be precise down to milliseconds in the visual system^{1–7}, even though the relevant timescales of natural vision are much slower. The existence of such precision contributes to a fundamental debate over the basis of the neural code and, specifically, what timescales are important for neural computation^{8–10}. Using recordings in the lateral geniculate nucleus, here we demonstrate that the relevant timescale of neuronal spike trains depends on the frequency content of the visual stimulus, and that ‘relative’, not absolute, precision is maintained both during spatially uniform white-noise visual stimuli and naturalistic movies. Using information-theoretic techniques, we demonstrate a clear role of relative precision, and show that the experimentally observed temporal structure in the neuronal response is necessary to represent accurately the more slowly changing visual world. By establishing a functional role of precision, we link visual neuron function on slow timescales to temporal structure in the response at faster timescales, and uncover a straightforward purpose of fine-timescale features of neuronal spike trains.

Figure 1a illustrates one of the many contexts in which millisecond precision has been observed in neuronal responses, showing the response of a geniculate neuron to repeated presentations of a spatially uniform white-noise visual stimulus (SUN). This remarkable precision at millisecond timescales has been observed in the retina^{2,7}, the lateral geniculate nucleus (LGN)^{5,6} and the visual cortex^{1,3,10} as well as in many other sensory systems such as the fly visual system^{4,9}, the electrosensory system of the weakly electric fish¹¹, and the mammalian somatosensory^{12,13} and auditory systems¹⁴. Although the presence of such fine temporal structure in the neuronal response would not be surprising if the sensory stimulus had similar temporal structure, its role is less clear in the mammalian visual system in which relevant visual stimuli are typically on much slower timescales. In particular, visual perception is ultimately limited by the relatively slow integration time of photoreceptors, which, for example, results in the appearance of continuous motion from the flickering images that constitute a movie. As a result, the much finer temporal structure in visual neuron responses has been proposed to be evidence for ‘temporal encoding’, which postulates that particular temporal patterns in the spike train carry additional information about the visual stimulus^{8,15}.

If millisecond temporal structure is important for the representation of visual information, one might expect it to be preserved in more natural stimulus conditions. Figure 1b shows the response of the same neuron in Fig. 1a responding to a ‘natural movie’ stimulus (see Methods)¹⁶. Although the response of the neuron still consists of discrete firing ‘events’, all relevant measures of event timing, ranging from the amount of trial-to-trial variability in timing of the first

spike (‘first-spike jitter’) to the overall event duration, are significantly slower during the natural movie versus during SUN (Supplementary Information 2). However, the increased response timescales are nearly matched by the increased timescales of the stimulus filtered by the receptive field of the neuron, because both the stimulus itself and the temporal filtering of the receptive field become slower in natural movies (Supplementary Information 1). By comparing the timescale of the ‘filtered stimulus’ τ_{FS} with that of the response τ_R during the SUN and the natural movie stimuli (Fig. 1c, d), we observe that ‘relative’ precision (defined as the ratio of the characteristic timescales of the filtered stimulus and neural response) is maintained at a level of 2–4 during both SUN and natural stimuli by all neurons in our study (Fig. 1e, f).

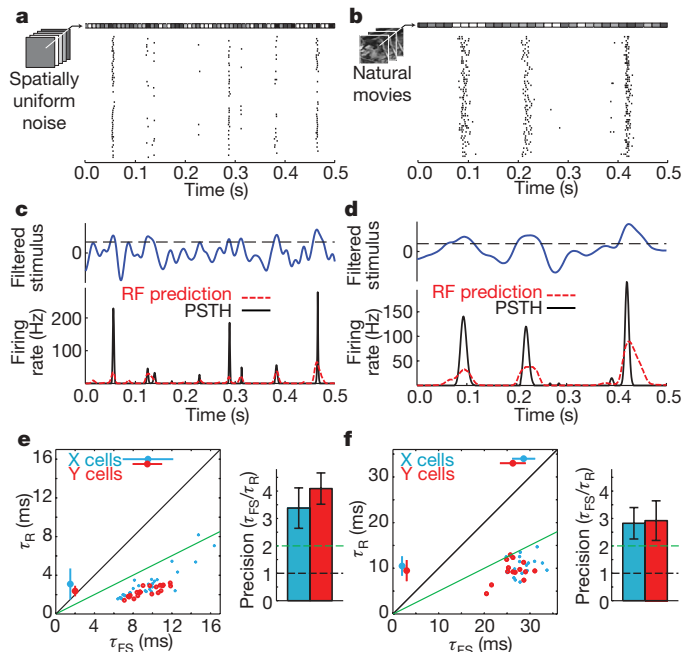


Figure 1 | The timescale of the neuronal response depends on the nature of the visual stimulus, defining relative precision. **a, b**, LGN X-cell responses (spike rasters) over 60 repeated trials in response to SUN (**a**) and natural movie stimuli (**b**). **c, d**, Corresponding time courses of the filtered stimulus (arbitrary units, with threshold shown as dashed line), PSTH (see Methods) and RF (receptive field) prediction. **e, f**, Left: comparison of the filtered stimulus and response timescales (τ_{FS} and τ_R , respectively) across the population of LGN neurons studied, for SUN (**e**, $n = 49$) and natural movies (**f**, $n = 32$). Right: the consistent ratio between these timescales (shown as mean \pm s.d.) defines the level of relative precision.

¹School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA. ²The HRH Prince Alwaleed Bin Talal Bin Abdulaziz Alsaud Institute for Computational Biomedicine, Weill Medical College of Cornell University, New York, New York 10021, USA. ³Department of Biological Sciences, State University of New York College of Optometry, New York, New York 10036, USA. ⁴Department of Psychology, University of Connecticut, Storrs, Connecticut 06269, USA.

Several information-based studies of retinal and geniculate spike trains in response to SUN have demonstrated that millisecond precision conveys additional information about the visual stimulus^{5–7}. However, we observe that—although still ‘precise’ in both stimulus contexts—the response timescale itself nearly quadruples from an average of 2.8 ms in the context of SUN to 10.1 ms in the natural movie. This suggests that any functional role of temporal precision cannot depend on fixed temporal features in the LGN response, and raises the question of whether the extra information that is encoded at millisecond timescales disappears altogether in more natural contexts. To address this question, we must first understand what this millisecond-timescale information represents about the visual stimulus. We investigated this using a ‘realistic model’ that reproduces the observed temporal features and trial-to-trial variability of LGN spike trains in the context of SUN (see Supplementary Information 3)¹⁷. Both the timescale of this model’s filtered stimulus (7.1 ms) and its response (1.6 ms) are in the range of experimental observations (Fig. 1e), giving a relative precision of 4.5; its spike train also carries a typical amount of information⁶.

Notably, the relative precision of the model can be varied systematically (Fig. 2a) from infinite precision (blue) to a relative precision of 4.5 (green; realistic model) and to a relative precision of 1 (red; matching the response timescale to the filtered stimulus timescale), without changing how the stimulus itself is processed by the receptive field. We evaluate the effects of precision on the information content

of the spike train using the single-spike information rate I_{SS} (ref. 18), which is a first-order approximation to the full mutual information^{6,19} and reflects the more general dependence of mutual information measures on response timescales. By systematically varying the precision, this analysis reveals what at first seems to be a paradoxical effect of precision: information about the stimulus grows without bound with increasing precision (Fig. 2b, note the log-scaling of the horizontal axis), even though no information can be created through spike generation that does not already exist at the level of the filtered stimulus²⁰.

To understand the nature of this literally boundless amount of information that becomes available with increasing precision, consider that, for a temporally varying stimulus, there is information to be gained in distinguishing between a stimulus $s(t)$ and its temporally shifted ‘twin’ $s(t + \Delta t)$ for any time shift Δt , regardless of the scale of the stimulus features (Fig. 2c). Despite this conceptual distinction between stimuli, as Δt becomes vanishingly small, $s(t) \approx s(t + \Delta t)$.

To formalize this intuition, it is necessary to apply more traditional measures of information between stimulus and response, based on linear stimulus reconstruction. The quality of the reconstruction is barely affected by the level of variability of the realistic model (Fig. 2d, middle), whereas increasing the temporal jitter such that the relative precision is decreased to 1 significantly degrades the reconstruction (Fig. 2d, bottom). This can be quantified by the reconstruction information I_{REC} (refs 8, 21 and 22), which explicitly measures how much information about the stimulus can be recovered from the observed neuronal response using linear reconstruction. Unlike the single-spike information, the reconstruction information saturates at high precision (Fig. 2e), such that the response precision of the realistic model (green dashed line) is sufficiently high to capture the available I_{REC} .

What is surprising is that, although the stimulus reconstruction is based on the receptive field function at the timescales of the filtered stimulus (Fig. 2e, red dashed line), we see that the neuron’s response must be significantly more precise to recover the full reconstruction information. Conversely, nearly half of the available information is lost if the response timescale is only dictated by the filtered stimulus. Thus, these results recapitulate the dichotomy between the timescales of receptive field function and the precision of the response.

As Fig. 2c implies, the frequency content of the stimulus determines the temporal scale at which the response must be specified to reconstruct the stimulus faithfully. The reconstruction information can be explicitly decomposed as a function of frequency using the coherence $\gamma^2(f)$ between the stimulus and the reconstructed stimulus. The coherence for the precise model neuron is compared to the coherence when the model’s precision is degraded by adding temporal jitter to its spike train (Fig. 3a). Decreased precision has the largest effect on the highest stimulus frequencies represented, and a simple relationship emerges when considering the relationship between the coherence of the precise model spike trains and the coherence of the jittered spike trains (Fig. 3b). As derived in Supplementary Information 4, the frequency scale of the degradation of the reconstruction (denoted here as the ‘attenuation frequency’, f_A) is directly related to the amount of temporal jitter in the neural response. This relationship provides a general guideline for understanding which filtered stimulus frequencies can be reconstructed for a certain amount of response jitter. For example, an LGN neuron recorded in the context of the natural movie has a filtered stimulus with frequency content up to 20 Hz (Fig. 3c and Supplementary Information 4). This frequency maps to a jitter of 6.2 ms (12.4-ms timescale), which nearly matches the observed timescale of the neuron’s response (Fig. 3d). Furthermore, the experimentally measured I_{REC} degrades for larger magnitudes of jitter, suggesting that the precision observed during the natural movie is just enough to avoid disrupting the information available to linear stimulus reconstruction.

Owing to the spatiotemporal nature of LGN receptive fields, the requirement for temporal precision also applies to the ability to

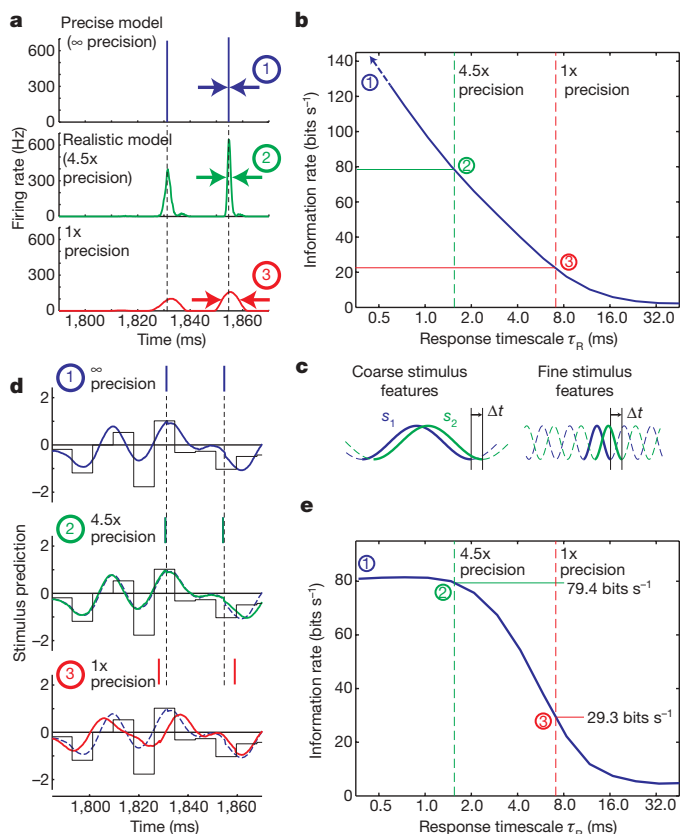


Figure 2 | Precision is necessary to convey information about visual stimuli. **a**, Resulting PSTH as the temporal precision of the LGN spike train is systematically varied. **b**, Single-spike information I_{SS} increases without bound with increasing precision. **c**, Jittering a reconstructed stimulus by Δt has a greater effect for fine (right) than for coarse (left) stimulus features. **d**, Jitter affects the reconstruction quality at a relative precision of 1 (bottom), but has negligible influence at a relative precision of 4.5 (middle). The stimulus is scaled such that its variance over time is unity, establishing the scale of the y axis. **e**, I_{REC} reveals that response timescales much finer than that of the stimulus (red) are necessary to maximize the neuron’s linear representation of the stimulus.

represent spatial aspects of the stimulus. We consider a simple spatiotemporal reconstruction using an array of model neurons, each representing one pixel of the natural movie stimulus considered in this paper¹⁶. Figure 3e shows single frames of the reconstructed movie from this simulation using temporally precise neurons (left), neurons with a relative precision of 2.3 (6.2-ms jitter, middle) and neurons with a relative precision of 1 (14.5-ms jitter, right). Although this stimulus only updates at 30 Hz (corresponding to 33-ms stimulus frames), the degradation of the quality in the image from left to right demonstrates that spatial information about the stimulus is represented at timescales finer than that of the stimulus. Furthermore, at the spatial scale of individual receptive fields (Fig. 3f), a simple ‘edge detector’ demonstrates that boundaries of objects present in the image become disrupted as precision decreases to 1. This has clear implications for the necessary temporal sensitivity of cortical neurons that receive direct input from the LGN, which are thought to encode edges. To have access to the full amount of information from their LGN inputs, they must be sensitive to these inputs at the same level of precision, which in the natural movie is on the order of 10 ms.

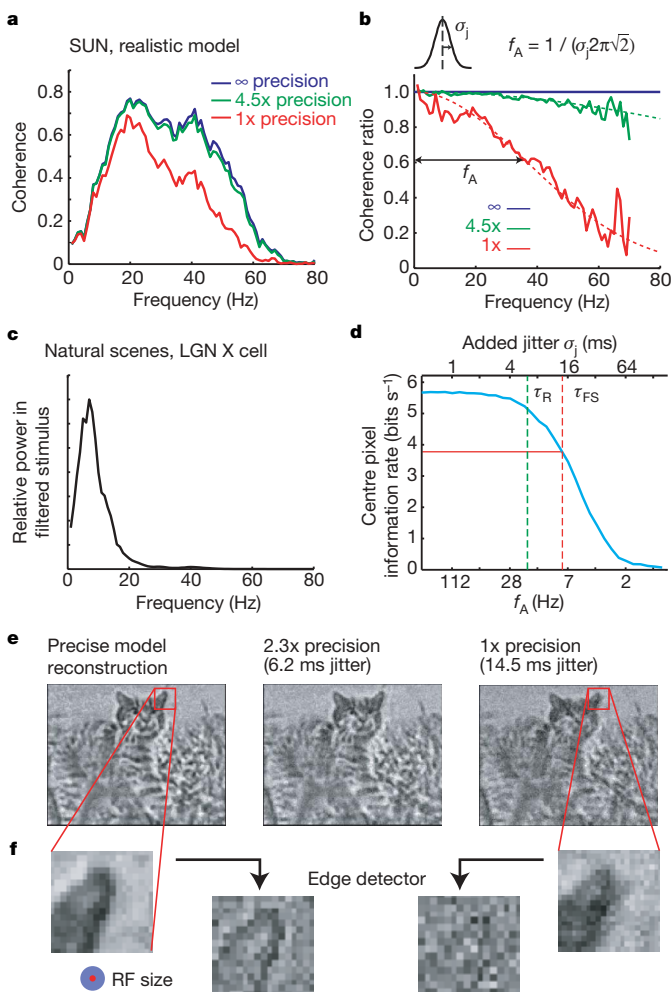


Figure 3 | Precision is necessary to represent relevant stimulus frequencies. **a**, Coherence between SUN and its reconstruction in models with varying amounts of jitter. **b**, Ratio of coherence between jittered and precise models, defining the attenuation frequency f_A . **c**, Power in the natural movie filtered by a measured X-cell receptive field. **d**, Reconstruction information for the X cell at its centre pixel as a function of added jitter (green, neuron’s measured jitter, relative precision of 2.3; red, degraded, relative precision of 1). **e**, Frame of the reconstructed stimulus generated from an array of neurons with different precisions. **f**, Spatial fidelity is compromised by a lack of precision.

Thus, we have demonstrated that the ‘classical’ function of visual neurons—signalling stimuli in their receptive field—actually necessitates the neural precision observed experimentally. Whereas it has long been known that the timescales of visual neuron responses can be modulated by the nature of the visual stimulus, the analysis described here places this finding in the context of their functional role: representation of the natural visual scene. This provides a link between the timescales of natural vision (with frequency content largely <20 Hz owing to the filtering of LGN receptive fields, see Supplementary Information 4) and response timescales on the order of 10 ms, which are increasingly being observed in the cortex (for example, cortical dendritic integration²³, cortical synaptic plasticity²⁴, spike-train-based discrimination^{25,26} and gamma-band modulation of cortical fields²⁷). In this light, LGN inputs to the cortex may be at a suitable timescale for cortical computation, and understanding the temporal structure of these inputs at the appropriate level of precision may lead to a new perspective on cortical processing.

The observation that the timescales of neural responses are not absolute indicates that downstream neurons cannot depend on fixed temporal features of the response to encode the stimulus. Instead, we demonstrate that the information in the spike train (at least at the level of single spikes) is accessible through simple linear decoding, which relies on relative precision rather than fixed temporal patterning. Although this study does not preclude the existence of more complicated explanations for visual neuron precision that extend beyond receptive-field function (that is, ‘temporal encoding’¹⁵) or across populations^{22,28,29}, it demonstrates that precision is not by itself evidence for such ideas. In this sense, our results establish stronger criteria that ‘temporal encoding’ must satisfy to demonstrate a functional role above and beyond straightforward receptive-field-based processing.

Relative precision may be a general feature of sensory neuron communication, in which an analogue input (the sensory stimulus) is encoded by what is essentially a digital signal (the neuron’s spike train). In this context, temporal precision of neuronal responses is conceptually similar to the problem of digital sampling, in which encoding frequencies must be at least double that of the analogue signal information because of the Nyquist limit^{8,30}. From this perspective, the mechanisms that generate neuronal precision (Supplementary Information 3), which seem to make the encoding of visual information more complicated, may actually serve to provide easier means for downstream neurons to decode this information.

METHODS SUMMARY

Data were recorded extracellularly from layer A of the LGN of anaesthetized paralysed cats. Individual units were classified as X or Y according to their responses to counterphase sinusoidal gratings.

The SUN stimuli were 120 Hz gaussian white noise with an r.m.s. (root-mean-squared) contrast of 0.55. The naturalistic movie was generated by a camera mounted to the head of a cat roaming in a forest, provided by the König laboratory (Institute of Neuroinformatics, ETH/UNI Zürich, Switzerland)¹⁶. It was modified so that each frame had the same mean luminance and an r.m.s. contrast of 0.4, and was updated at 60 Hz. All stimuli were presented on a cathode-ray tube display with a 120 Hz monitor refresh rate; we verified that the monitor refresh did not affect our results (Supplementary Information 5).

The receptive field predictions were based on ‘linear–non-linear’ (LN) models, estimated using standard techniques (Supplementary Information 1). The filtered stimulus is the linear convolution of the stimulus with the receptive field. Filtered stimulus and response peri-stimulus time histogram (PSTH) timescales (τ_{FS} and τ_R) were derived from their autocorrelation functions (Supplementary Information 1).

The parameters of the realistic model LGN neuron were provided by M. Meister¹⁷. Precise model responses were generated from the realistic model by identifying individual spikes across trials and shifting each spike to its average time over all trials in which it was present. The amount of each shift was then scaled to change the level of precision (Fig. 2). The spatiotemporal simulation (Fig. 3e, f) consisted of realistic model neurons modified to have measured spatiotemporal receptive fields (Supplementary Information 1) and assembled into two 160×120 overlapping arrays (ON and OFF). Edge detection (Fig. 3f)

was performed by measuring the difference between each pixel and the average of the surrounding pixels.

The single-spike information rate¹⁸ was calculated directly from the PSTH, and the reconstruction information I_{REC} was calculated from the coherence $\gamma^2(f)$ between the stimulus and the reconstructed stimulus: $I_{\text{REC}} = -\int df \log_2[1 - \gamma^2(f)]$.

Full Methods and any associated references are available in the online version of the paper at www.nature.com/nature.

Received 20 February; accepted 16 July 2007.

- Bair, W. & Koch, C. Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput.* **8**, 1185–1202 (1996).
- Berry, M. J. II & Meister, M. Refractoriness and neural precision. *J. Neurosci.* **18**, 2200–2211 (1998).
- Buracas, G. T., Zador, A. M., DeWeese, M. R. & Albright, T. D. Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* **20**, 959–969 (1998).
- Lewen, G. D., Bialek, W. & de Ruyter van Steveninck, R. R. Neural coding of naturalistic motion stimuli. *Network* **12**, 317–329 (2001).
- Liu, R. C., Tzovev, S., Rebrik, S. & Miller, K. D. Variability and information in a neural code of the cat lateral geniculate nucleus. *J. Neurophysiol.* **86**, 2789–2806 (2001).
- Reinagel, P. & Reid, R. C. Temporal coding of visual information in the thalamus. *J. Neurosci.* **20**, 5392–5400 (2000).
- Uzzell, V. J. & Chichilnisky, E. J. Precision of spike trains in primate retinal ganglion cells. *J. Neurophysiol.* **92**, 780–789 (2004).
- Borst, A. & Theunissen, F. E. Information theory and neural coding. *Nature Neurosci.* **2**, 947–957 (1999).
- de Ruyter van Steveninck, R. R., Lewen, G. D., Strong, S. P., Koberle, R. & Bialek, W. Reproducibility and variability in neural spike trains. *Science* **275**, 1805–1808 (1997).
- Mainen, Z. F. & Sejnowski, T. J. Reliability of spike timing in neocortical neurons. *Science* **268**, 1503–1506 (1995).
- Gabbiani, F., Metzner, W., Wessel, R. & Koch, C. From stimulus encoding to feature extraction in weakly electric fish. *Nature* **384**, 564–567 (1996).
- Bolouri, A. R. & Stanley, G. B. The dynamics of spatiotemporal response integration in the somatosensory cortex of the vibrissa system. *J. Neurosci.* **26**, 3767–3782 (2006).
- Phillips, J. R., Johnson, K. O. & Hsiao, S. S. Spatial pattern representation and transformation in monkey somatosensory cortex. *Proc. Natl Acad. Sci. USA* **85**, 1317–1321 (1988).
- Wehr, M. & Zador, A. M. Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature* **426**, 442–446 (2003).
- Theunissen, F. & Miller, J. P. Temporal encoding in nervous systems: a rigorous definition. *J. Comput. Neurosci.* **2**, 149–162 (1995).
- Kayser, C., Salazar, R. F. & Konig, P. Responses to natural scenes in cat V1. *J. Neurophysiol.* **90**, 1910–1920 (2003).
- Keat, J., Reinagel, P., Reid, R. C. & Meister, M. Predicting every spike: a model for the responses of visual neurons. *Neuron* **30**, 803–817 (2001).
- Brenner, N., Strong, S. P., Koberle, R., Bialek, W. & de Ruyter van Steveninck, R. R. Synergy in a neural code. *Neural Comput.* **12**, 1531–1552 (2000).
- Pola, G., Thiele, A., Hoffmann, K. P. & Panzeri, S. An exact method to quantify the information transmitted by different mechanisms of correlational coding. *Network* **14**, 35–60 (2003).
- de Ruyter van Steveninck, R. & Laughlin, S. B. The rate of information transfer at graded-potential synapses. *Nature* **379**, 642–645 (1996).
- Bialek, W., Rieke, F., de Ruyter van Steveninck, R. R. & Warland, D. Reading a neural code. *Science* **252**, 1854–1857 (1991).
- Dan, Y., Alonso, J. M., Usrey, W. M. & Reid, R. C. Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nature Neurosci.* **1**, 501–507 (1998).
- Yuste, R. & Denk, W. Dendritic spines as basic functional units of neuronal integration. *Nature* **375**, 682–684 (1995).
- Fu, Y. X. *et al.* Temporal specificity in the cortical plasticity of visual space representation. *Science* **296**, 1999–2003 (2002).
- Narayan, R., Grana, G. & Sen, K. Distinct time scales in cortical discrimination of natural sounds in songbirds. *J. Neurophysiol.* **96**, 252–258 (2006).
- Chichilnisky, E. J. & Kalmar, R. S. Temporal resolution of ensemble visual motion signals in primate retina. *J. Neurosci.* **23**, 6681–6689 (2003).
- Womelsdorf, T., Fries, P., Mitra, P. P. & Desimone, R. Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* **439**, 733–736 (2006).
- Ahissar, E. & Arieli, A. Figuring space by time. *Neuron* **32**, 185–201 (2001).
- VanRullen, R., Guyonneau, R. & Thorpe, S. J. Spike times make sense. *Trends Neurosci.* **28**, 1–4 (2005).
- Lazar, A. A. Perfect recovery and sensitivity analysis of time encoded bandlimited signals. *IEEE Trans. Circ. Syst.* **51**, 2060–2073 (2004).

Supplementary Information is linked to the online version of the paper at www.nature.com/nature.

Acknowledgements This work was supported by a Charles King Trust Postdoctoral Fellowship (Bank of America, Co-Trustee, Boston; D.A.B.), by the NGIA (D.A.B., N.A.L., G.B.S.), by the NIH and by the SUNY Research Foundation (C.W., J.J., C.-I.Y., J.-M.A.). We thank M. Goldman, M. Meister, G. Desbordes and A. Bolouri for comments on the manuscript, C. Kayser for providing the natural-scene movies, and P. Wolfe for discussions regarding sampling issues.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Correspondence and requests for materials should be addressed to D.A.B. (dab2024@med.cornell.edu).

METHODS

Experimental procedures. Surgical and experimental procedures were performed in accordance with United States Department of Agriculture guidelines and were approved by the Institutional Animal Care and Use Committee at the State University of New York, State College of Optometry. As described in detail in ref. 31, cats were initially anaesthetized with ketamine (10 mg kg⁻¹ intramuscular) followed by thiopental sodium (20 mg kg⁻¹ intravenous during surgery and at a continuous rate of 1–2 mg kg⁻¹ h⁻¹ intravenous during recording; supplemented as needed). A craniotomy and duratomy were performed to introduce recording electrodes into the LGN (anterior, 5.5; lateral, 10.5). Animals were paralysed with atracurium besylate (0.6–1 mg kg⁻¹ h⁻¹ intravenous) to minimize eye movements, and were artificially ventilated. LGN responses were recorded extracellularly within layer A. Recorded voltage signals were conventionally amplified, filtered and passed to a computer running the RASPUTIN software package (Plexon). For each cell, spike waveforms were identified initially during the experiment and were verified carefully off-line by spike-sorting analysis. Cells were classified as X or Y according to their responses to counter-phase sinusoidal gratings. Cells were eliminated from this study if they did not have at least 2 Hz mean firing rates in response to all stimulus conditions, or if the maximum amplitude of their spike-triggered average in spatiotemporal white noise was not at least five times greater than the amplitude outside of the receptive field area.

Visual stimuli. SUN stimuli consisted of spatially uniform luminances randomly selected from a gaussian distribution with zero mean (corresponding to the midway point of the full range of monitor luminance) and a root-mean-squared (r.m.s.) contrast of 0.55, presented at 120 Hz. The naturalistic movie sequence was recorded using a removable lightweight CCD camera mounted to the head of a freely roaming cat in natural environments such as grassland and forest¹⁶. A 48 × 48 windowed area was processed to play at 60 Hz and to have a constant mean and standard deviation of luminance for each frame (contrast held at 0.4 of maximum)³². All stimuli were displayed on a cathode-ray tube display at a resolution of 0.2° per pixel with a monitor refresh rate of 120 Hz. Care was taken to ensure that the potential effects of phase-locking to the monitor refresh did not affect the results (Supplementary Information 5).

Receptive-field-based predictions and the filtered stimulus. The receptive field predictions (Fig. 1c, d) were based on linear–non-linear (LN) models that were estimated using standard techniques (see Supplementary Information 1). Temporal-only receptive fields were estimated directly from responses to SUN, whereas we used 60 Hz spatiotemporal binary noise to map the spatiotemporal receptive field. The filtered stimulus (for example, Fig. 1c, d) is defined as the linear convolution of the stimulus with the receptive field, and represents how similar the stimulus is to the receptive field at a particular time. It is scaled to have a standard deviation of unity for a given stimulus class, and is mapped to a firing

rate through a ‘static non-linearity’ also using standard techniques (Supplementary Information 1).

Timescale of filtered stimulus and PSTH. The PSTH was estimated from 60 or 120 repeated trials, and was subsequently used to estimate response timescales and single-spike information (Supplementary Information 1).

The ‘realistic’ and ‘precise’ models. The parameters of the realistic model used in this paper were provided by M. Meister, originating from a fit to a representative cat LGN cell (neuron number 6 in ref. 17) recorded under similar experimental conditions to those used here. The precise model responses were generated from the realistic model spike times over 1,000 repeated trials by identifying individual spikes across trials and shifting each spike to agree with its average time over all trials in which it appeared. The amount of each shift of each spike from the precise model was then scaled to generate models with any amount of precision (Fig. 2).

Linear reconstruction. The optimal linear reconstruction filter $h(\tau)$ corresponds to the temporal filter that minimizes the mean-squared difference between the stimulus $s(t)$ and the reconstructed stimulus $s^*(t)$ (refs 8 and 21). The reconstruction filters were estimated numerically for a given stimulus and set of experimentally observed spike times.

Information measurements. The single spike information rate¹⁸ is given by:

$$I_{SS}[S, \lambda] = \lambda_0 \int dt \left[\frac{\lambda(t)}{\lambda_0} \right] \log_2 \left[\frac{\lambda(t)}{\lambda_0} \right] \quad (1)$$

In equation (1), $\lambda(t)$ is the instantaneous firing rate characterized by the PSTH, and λ_0 is the mean firing rate. The reconstruction information I_{REC} is calculated from the coherence spectra $\gamma^2(f)$ between stimulus $s(t)$ and reconstructed stimulus $s^*(t)$: $I_{REC} = - \int df \log_2 [1 - \gamma^2(f)]$. In the case of the natural movie reconstruction, only the reconstruction information of the centre pixel of the receptive field is reported.

LGN model neurons and natural movie simulation. We used a 160 × 120 array of precise neurons with the same parameters as the precise model of Fig. 2, except we replaced the purely temporal receptive field with a measured spatiotemporal receptive field. The spatial components were spatially shifted for each neuron to tile a 160 × 120 spatiotemporal stimulus derived from the same natural movie. An ON centre cell and an OFF centre cell were simulated at each position, and each pixel was independently reconstructed into a 160 × 120 image evolving over time using the resulting spike trains.

31. Weng, C., Yeh, C. I., Stoelzel, C. R. & Alonso, J. M. Receptive field size and response latency are correlated within the cat visual thalamus. *J. Neurophysiol.* **93**, 3537–3547 (2005).
32. Lesica, N. A. *et al.* Dynamic encoding of natural luminance sequences by LGN bursts. *PLoS Biol.* **4**, e209 (2006).