LETTERS

Decision-related activity in sensory neurons reflects more than a neuron's causal effect

Hendrikje Nienborg[1] & Bruce G. Cumming[1]

During perceptual decisions, the activity of sensory neurons correlates with a subject's percept, even when the physical stimulus is identical[1–9]. The origin of this correlation is unknown. Current theory proposes a causal effect of noise in sensory neurons on perceptual decisions[10–12], but the correlation could result from different brain states associated with the perceptual choice[13] (a top-down explanation). These two schemes have very different implications for the role of sensory neurons in forming decisions[14]. Here we use white-noise analysis[15] to measure tuning functions of V2 neurons associated with choice and simultaneously measure how the variation in the stimulus affects the subjects' (two macaques) perceptual decisions[16–18]. In causal models, stronger effects of the stimulus upon decisions, mediated by sensory neurons, are associated with stronger choice-related activity. However, we find that over the time course of the trial these measures change in different directions—at odds with causal models. An analysis of the effect of reward size also supports this conclusion. Finally, we find that choice is associated with changes in neuronal gain that are incompatible with causal models. All three results are readily explained if choice is associated with changes in neuronal gain caused by top-down phenomena that closely resemble attention[19]. We conclude that top-down processes contribute to choice-related activity. Thus, even forming simple sensory decisions involves complex interactions between cognitive processes and sensory neurons.

Considerable progress has been made in explaining the neuronal mechanisms underlying decision making[12], which is a major goal in systems neuroscience. For simple perceptual decisions, recent theory proposes that sensorimotor areas accumulate sensory evidence about the physical world, delivered by sensory neurons[10,11,20–22]. Noise in the sensory neurons causes variability in the behavioural response[10–12], resulting in a covariation between the neuronal activity and the behaviour[1–9]. (We note that this causal effect of noise in the sensory representation has only been invoked for sensory areas, not for sensorimotor areas.) However, this covariation could also result from top-down effects[13] in which brain states[23] that are associated with one behavioural response also alter the response of the sensory neurons. A third (bottom-up) possibility is that sensory neurons that themselves have no causal effect on the decision are correlated with sensory neurons that do have a causal effect. These schemes have markedly different implications for the role of sensory neurons in forming decisions. Sensory neurons either only encode the physical stimulus or simultaneously form an integral part of the mechanism used by the brain to decode the sensory information. To distinguish between these views, we combined the measurement of choice-related activity in disparity selective V2 neurons in a disparity-discrimination task, with a stimulus that permitted the use of white-noise analysis. This allowed the simultaneous application of 'subspace mapping'[15], to describe how disparity affects the neuronal response ('disparity subspace map'), and 'psychophysical reverse correlation'[16–18], to extract a kernel describing how disparity affects the subjects' perceptual choices. This comprehensive data set enables us to differentiate among these schemes.

Two macaque monkeys performed a coarse disparity-discrimination task (Fig. 1) while we recorded extracellularly from disparity-selective neurons in their visual areas V2. The stimulus, a circular random dot stereogram, contained a spatially uniform binocular disparity that varied randomly on each video frame. At the end of each trial the monkeys reported in a forced choice task whether the stimulus appeared near ('near' choice) or far ('far' choice). We exploited this random variation to perform psychophysical reverse correlation[16–18], and simultaneously to measure neuronal subspace maps[15] for disparity.

First we examined how the monkeys weight the disparity signal in the stimulus to form their decision[16]. We calculated the difference between the average stimulus preceding the monkeys' 'near' choices

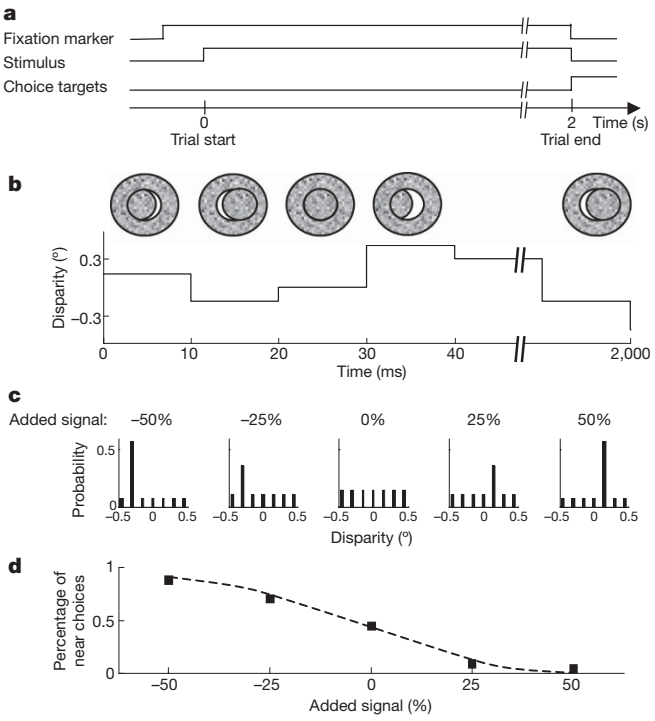

Figure 1 | Methods. a, Sketch of the sequence of events during one trial. **b**, Example time series of the stimulus. **c**, Probability mass distributions of the stimuli for one experiment (probability as a function of disparity), with signal disparities of -0.3° and 0.15°. Each plot depicts one signal condition (negative percentages indicate near signal disparities). **d**, The monkey's performance for this experiment (in percentage 'near' choices as a function of percentage added signal).

[1]Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health, 49 Convent Drive, Bethesda, Maryland 20892, USA.

and the average stimulus preceding the monkeys' 'far' choices. This 'psychophysical kernel' measures the relative probability with which the disparity on any given frame occurred before the monkeys' 'near' choice. The amplitude of the kernel decreased substantially over the course of the trial (Fig. 2a, b). (The Supplementary Information discusses the shape of the psychophysical kernel and shows that this linear analysis adequately captures the monkeys' behaviour.) This means that the monkeys relied predominantly on the stimulus disparities at the start of the trial and progressively less so towards its end. If one considers neurons representing this sensory evidence, their activity early in the trial should have a stronger effect on the decision than activity late in the trial. Thus, if the choice-related activity reflected only the causal effect of the neuronal firing on the choice, the size of the choice-related activity should also decrease over time. This prediction follows directly from the fact that, in the causal explanation, choice-related activity is the effect of noise in the sensory evidence that is used to make a decision.

To test this prediction, we quantified the choice-related signal as 'choice probability'[3]. (Choice probabilities were corrected for the stimulus-induced component; see Supplementary Information.) The time course of the choice-related signal in our data (Fig. 2c) is different from that predicted from the time course of the psychophysical data in the causal-only scheme. Choice probabilities were measured for 76 neurons that had been recorded while the data for the psychophysical kernel were gathered. For 57 of the 76 neurons, for which the choice probability was >0.5, we examined the mean choice probability as a function of time (Fig. 2c). Consistent with previous findings[4], the choice probability plateaus after about 500 ms, which is quite unlike the statistically significant decrease in amplitude of the psychophysical kernel over time (correlation coefficient, $r = -0.81$; $P < 10^{-23}$ between amplitude and time, over the second half of the trials). Although choice-probability time courses for individual neurons are noisy, we addressed the possibility that some neurons behave as if they play a causal role by computing the correlation coefficient, R, between the time course of the choice probability for each individual neuron and the time course of the average psychophysical kernel amplitude (Fig. 2d). We found a significant negative correlation between these coefficients and a neuron's choice probability ($r = -0.28$, $P < 0.05$), indicating that neurons with high choice probabilities tended to show a negative correlation with the time course of the psychophysical kernel amplitude, as expected from the average data (Fig. 2b). This and other features of individual time courses (Supplementary Discussion) are at odds with the causal model.

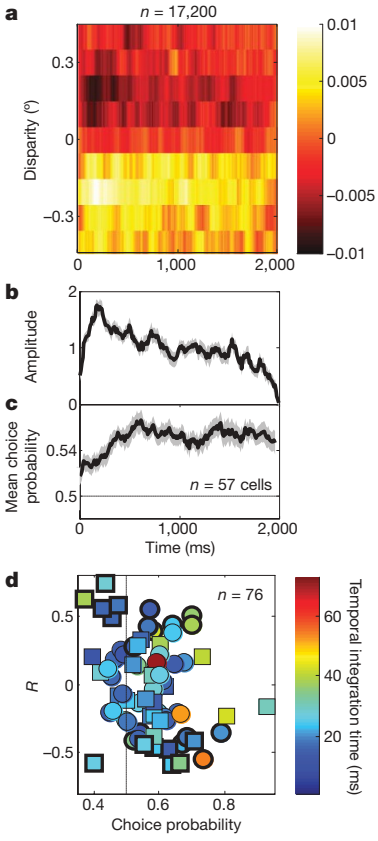

Figure 2 | Psychophysical kernel and choice-related signal have different time courses. **a**, Psychophysical kernel (averaged over 76 experiments; $n = 17{,}200$ trials; two monkeys) as a function of disparity and time. Colour represents amplitude (in occurrences per frame). **b**, Normalized amplitude of the psychophysical kernels decreases over time. **c**, Averaged choice-related signal over time. Shaded grey areas in **b** and **c**, ±1 standard error. **d**, The correlation coefficient, R, over time between choice probability (for individual neurons) and the amplitude of the mean psychophysical kernel, plotted against a neuron's mean choice probability. Colour represents temporal integration time (Supplementary Methods); bold symbols, significant R ($P < 0.05$, by resampling); circles, data from monkey 1; squares, data from monkey 2.

The fact that the results in Fig. 2 are incompatible with the causal-only account, suggests that choice probabilities are at least partly of non-causal, possibly top-down, origin. We therefore sought a signature of possible top-down mechanisms at the level of individual neurons. This could employ a mechanism, similar to attention, that characteristically alters the gain of sensory neurons[19]. To test this possibility, we designed our disparity-varying stimulus such that it permitted the measurement of subspace maps for disparity (Methods).

These subspace maps quantify the effect of each disparity (in the stimulus with no added signal) on the neuron. Calculating subspace maps separately for stimuli associated with 'near' and 'far' choices quantifies any effects of choice on the neuronal response. Choice-related activity itself implies some difference between these subspace maps. If the difference is caused by a change in neuronal gain, the two subspace maps should be scaled versions of each other. Example subspace maps for one neuron (Fig. 3a, b) show that the gain of this neuron increased by 84%, whereas the additive change was close to 0 (−0.032 spikes per frame). A second example shows a more typical gain increase (18%; y offset, 0.005 spikes per frame; Fig. 3c, d).

The distribution of relative gain change as a function of choice probability demonstrates that choice probabilities are associated with choice-related changes in neuronal gain ($n = 76$, $r = 0.44$, $P < 10^{-4}$ (monkey 1: $n = 42$, $r = 0.54$, $P < 0.001$; monkey 2: $n = 34$, $r = 0.32$, $P < 0.07$); Fig. 3e). The geometric mean of the relative gains was 1.16 (1.17 and 1.15 for monkeys 1 and 2, respectively), which is significantly different from 1 ($P < 0.001$, by resampling). Conversely, there was no systematic relationship between the y offset and the choice probability ($r = 0.03$, $P = 0.77$ (monkey 1: $r = -0.18$, $P = 0.25$; monkey 2: $r = 0.18$, $P = 0.31$); mean offset, −0.03 spikes per frame; Fig. 3f). Thus, it is the choice-dependent change in response gain that explains the difference in mean response rate between preferred and null choices. A modest gain change could arise, even in the causal account of choice probability, from the firing properties of cortical neurons (for example Poisson spiking[24]). A shuffling technique showed that this effect was too small to account for the observed gain changes (Supplementary Information).

The gain change suggests the operation of a mechanism similar to feature-selective attention[19], but which varies from trial to trial. This could arise in several ways. First, as the decision is formed, a signal altering the neuronal gain may be sent back to those neurons supporting this decision. Alternatively, this gain change may implement a perceptual working memory[25], or a perceptual bias/expectation: attending to near features increases the response gain of near-preferring neurons and thus makes a 'near' response more likely.

An additional feature of our data provides evidence that at least the latter mechanism operates. The reward size depended systematically on the animals' performance (Methods). This performance was better in trials for which a large reward was available (Fig. 4a and

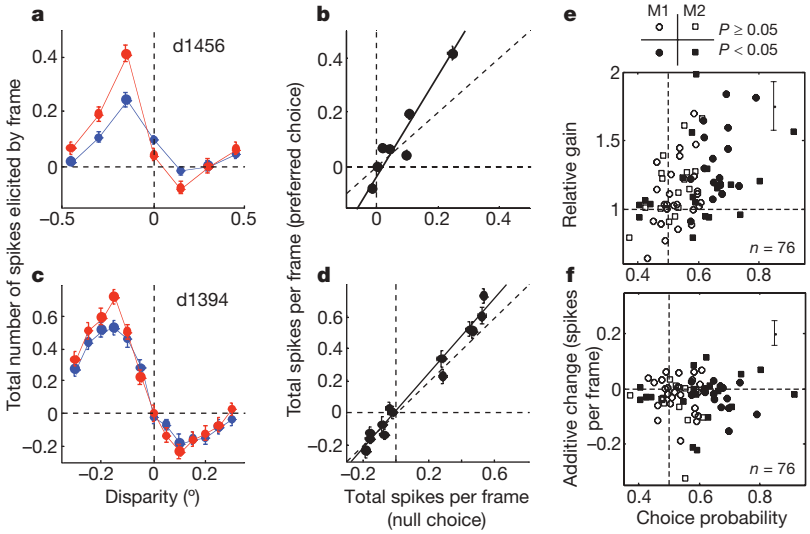

Figure 3 | Choice-dependent gain change. **a**, Subspace maps for preferred (red) and null (blue) choices superimposed (neuron d1456). Dashed lines: 0° disparity, 0 spikes per frame. **b**, Null-choice responses plotted against preferred-choice responses, yielding relative gain (slope, 1.84) and additive change (*y* offset, −0.032 spikes per frame). Dashed lines: unit relative gain, 0 spikes per frame. **c**, **d**, Same format as in **a** and **b** for neuron d1394, the slope (1.18) and *y* offset (0.005 spikes per frame) of which resemble the population mean. Error bars in **a**–**d**, standard errors (by resampling). **e**, Slope and choice probability are correlated. Filled and open symbols respectively represent cells with and without significant choice probability. Circles, data for monkey 1 (M1); squares, data for monkey 2 (M2). Dashed lines: 0.5 choice probability, unit relative gain. **f**, There is no correlation between *y* offset and choice probability (symbols as in **e**). Dashed lines: 0.5 choice probability, 0 spikes per frame. Error bars in **e** and **f**, median standard error.

Supplementary Information), indicating that animals used more information about the stimulus when reward size was large. It allows us to further test causal explanations for choice probability: when the animal uses more stimulus-derived information, choice probability should be larger. Contrary to this prediction, we found that choice probabilities were significantly smaller for trials in which a large reward was available ($P < 0.006$, paired *t*-test; Fig. 4b). This result can be explained by assuming that the animal has some bias (or expectation) at the start of each trial (regardless of reward size), and that this bias engages our proposed top-down mechanism. When the available reward is small and the monkeys make less use of the sensory input (as demonstrated by the psychophysical kernel; Fig. 4a), the bias will have a stronger impact on the behavioural response. Conversely, when a large reward is available, the improved performance implies that any initial bias is more likely to be overridden by the evidence provided by the visual stimulus. Hence, any component of choice probability reflecting a top-down effect of bias will be smaller in large-reward trials when the decision is more strongly driven by the actual stimulus and less by the monkey's initial bias.

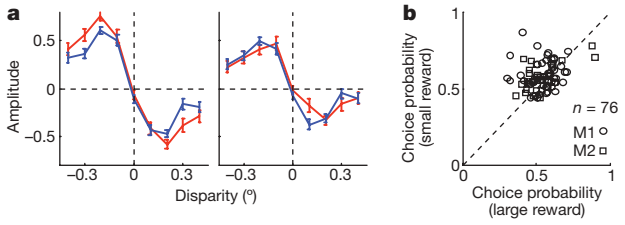

Figure 4 | Reward size affects behaviour and choice-related signal. **a**, Psychophysical kernel as a function of disparity and available reward, in occurrences per 1,000 ms, averaged over the first (left panel) and second (right panel) 1,000 ms of each trial ($n = 6{,}886$ trials for large reward (red); $n = 10{,}314$ trials for small reward (blue)). Improved performance is mainly caused by there being a larger psychophysical kernel in the first 1,000 ms of the trials (kernel difference $P < 0.001$, by resampling), but not the second half (difference not significant). Error bars, standard errors (by resampling). **b**, Choice probability computed for the first 1,000 ms of the trials was larger when a smaller reward was available ($P < 0.006$, $n = 76$). Dashed line: unit slope.

Our results provide three lines of evidence against the view that decision-related activity in sensory neurons reflects only the causal effect of neuronal noise on sensory decisions. First, the time course of the decision-related signal was incompatible with that predicted from the behavioural data in the causal-only scheme. Second, larger rewards systematically improved the animals' behaviour, but reduced choice probability, a result opposite to the expectation from causal explanations. Finally, choice probabilities were associated with choice-dependent gain changes larger than could be explained in the causal scheme. All three phenomena follow naturally from a top-down scheme in which the animals' perceptual state alters the response of sensory neurons. An alternative explanation is that neurons which do not contribute to the decision show choice probability because they are correlated with neurons that do contribute. Such a scheme, if sufficiently rich, may explain the data without invoking a top-down mechanism (Supplementary Discussion), but nonetheless abandons the principle that choice probabilities reflect only the causal effect of sensory noise upon decisions. Given that the choice-dependent gain changes we observe are characteristic of top-down mechanisms such as attention, our top-down scheme is the most parsimonious.

Changes in neuronal gain may facilitate the decoding of neuronal populations by appropriately weighting relevant neurons[26–28]. Implementing such a decoding mechanism at the level of sensory neurons allows the brain extraordinary flexibility to perform sensory decisions in different circumstances. Here we have shown that these gains vary with a subject's choice, within a fixed task. This gain change could implement a perceptual bias or expectation (attending to near or far features), and could also follow the formation of a decision. It may serve to promote perceptual stability in the presence of ambiguous[29] or noisy sensory signals. Whether it reflects an expectation or a post-decision signal, our data suggest that even simple sensory decisions involve top-down mechanisms that combine cognitive processes and the sensory representation in previously unreported ways.

METHODS SUMMARY

All procedures were performed in accord with the US Public Health Service policy on the humane care and use of laboratory animals and all protocols were approved by the National Eye Institute Animal Care and Use Committee. We recorded extracellular activity from disparity-selective single V2 units while two

monkeys (*Macaca mulatta*) performed disparity discrimination. Upon fixation, a stimulus was presented for 2 s and followed by two choice targets. After a saccade to the correct target, the monkeys received a liquid reward. Stimuli were dynamic random dot patterns: a disparity-varying centre (disparity changed randomly on each frame; 96-Hz frame rate) and a surrounding annulus at 0°. The centre disparity was chosen from a set of evenly spaced disparity values centred around 0° (encompassing the preferred and null disparity of each neuron). We introduced a detectable signal by increasing the probability of occurrence of one disparity within some trials. These signal disparities approximately matched each neuron's preferred and null disparity. Signal trials served only to control behaviour: all analyses were restricted to trials with no signal added.

Psychophysical kernels were computed as the difference in the respective mean stimulus matrices before 'near' and 'far' choices. The average kernel was a weighted average of the kernel for each experiment for which neuronal data were included. Choice probabilities were obtained as described previously[3,6], but corrected for the stimulus-induced component (Supplementary Information). For the subspace analysis, the average response of each neuron following one frame of a given disparity, d_i, was computed as a spike density function, $S_i(t)$. We calculated the total number of spikes elicited by one frame of disparity d_i as the sum of the overall mean number of spikes per frame and the integral of the deviation of $S_i(t)$ about this mean. Analysing the trials separately according to the monkeys' choice yield subspace maps for each choice.

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

Received 29 October 2008; accepted 23 January 2009.
Published online 8 March 2009.

1. Logothetis, N. K. & Schall, J. D. Neuronal correlates of subjective visual perception. *Science* **245,** 761–763 (1989).
2. Leopold, D. A. & Logothetis, N. K. Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature* **379,** 549–553 (1996).
3. Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S. & Movshon, J. A. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis. Neurosci.* **13,** 87–100 (1996).
4. Uka, T. & DeAngelis, G. C. Contribution of area MT to stereoscopic depth perception: choice-related response modulations reflect task strategy. *Neuron* **42,** 297–310 (2004).
5. Uka, T., Tanabe, S., Watanabe, M. & Fujita, I. Neural correlates of fine depth discrimination in monkey inferior temporal cortex. *J. Neurosci.* **25,** 10796–10802 (2005).
6. Nienborg, H. & Cumming, B. G. Macaque V2 neurons, but not V1 neurons, show choice-related activity. *J. Neurosci.* **26,** 9567–9578 (2006).
7. Dodd, J. V., Krug, K., Cumming, B. G. & Parker, A. J. Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. *J. Neurosci.* **21,** 4809–4821 (2001).
8. Purushothaman, G. & Bradley, D. C. Neural population code for fine perceptual decisions in area MT. *Nature Neurosci.* **8,** 99–106 (2005).
9. Ress, D. & Heeger, D. J. Neuronal correlates of perception in early visual cortex. *Nature Neurosci.* **6,** 414–420 (2003).
10. Shadlen, M. N., Britten, K. H., Newsome, W. T. & Movshon, J. A. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* **16,** 1486–1510 (1996).
11. Yang, T. & Shadlen, M. N. Probabilistic reasoning by neurons. *Nature* **447,** 1075–1080 (2007).
12. Schall, J. D. Neural correlates of decision processes: neural and mental chronometry. *Curr. Opin. Neurobiol.* **13,** 182–186 (2003).
13. Krug, K. A common neuronal code for perceptual processes in visual cortex? Comparing choice and attentional correlates in V5/MT. *Phil. Trans. R. Soc. Lond. B* **359,** 929–941 (2004).
14. Gilbert, C. D. & Sigman, M. Brain states: top-down influences in sensory processing. *Neuron* **54,** 677–696 (2007).
15. Ringach, D. L., Hawken, M. J. & Shapley, R. Dynamics of orientation tuning in macaque primary visual cortex. *Nature* **387,** 281–284 (1997).
16. Neri, P., Parker, A. J. & Blakemore, C. Probing the human stereoscopic system with reverse correlation. *Nature* **401,** 695–698 (1999).
17. Ahumada, A. J. Perceptual classification images from Vernier acuity masked by noise. *Perception* **26,** 18 (1996).
18. Nienborg, H. & Cumming, B. G. Psychophysically measured task strategy for disparity discrimination is reflected in V2 neurons. *Nature Neurosci.* **10,** 1608–1614 (2007).
19. Treue, S. & Martinez Trujillo, J. C. Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* **399,** 575–579 (1999).
20. Gold, J. I. & Shadlen, M. N. Representation of a perceptual decision in developing oculomotor commands. *Nature* **404,** 390–394 (2000).
21. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30,** 535–574 (2007).
22. Mazurek, M. E., Roitman, J. D., Ditterich, J. & Shadlen, M. N. A role for neural integrators in perceptual decision making. *Cereb. Cortex* **13,** 1257–1269 (2003).
23. Arieli, A., Sterkin, A., Grinvald, A. & Aertsen, A. Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* **273,** 1868–1871 (1996).
24. Dean, A. F. The variability of discharge of simple cells in the cat striate cortex. *Exp. Brain Res.* **44,** 437–440 (1981).
25. Pasternak, T. & Greenlee, M. W. Working memory in primate sensory systems. *Nature Rev. Neurosci.* **6,** 97–107 (2005).
26. Ma, W. J., Beck, J. M., Latham, P. E. & Pouget, A. Bayesian inference with probabilistic population codes. *Nature Neurosci.* **9,** 1432–1438 (2006).
27. Martinez-Trujillo, J. C. & Treue, S. Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr. Biol.* **14,** 744–751 (2004).
28. Jazayeri, M. & Movshon, J. A. Optimal representation of sensory information by neural populations. *Nature Neurosci.* **9,** 690–696 (2006).
29. Leopold, D. A., Wilke, M., Maier, A. & Logothetis, N. K. Stable perception of visually ambiguous patterns. *Nature Neurosci.* **5,** 605–609 (2002).

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

Acknowledgements This research was supported by the Intramural Research Program of the US National Institutes of Health, National Eye Institute. We are grateful to J. A. Movshon and M. Shadlen for discussions and to the members of the Laboratory of Sensorimotor Research for comments on an earlier version of this manuscript. We also thank D. Parker and B. Nagy for excellent animal care.

Author Contributions H.N. designed the project, collected the data, performed the analyses and wrote the paper. B.G.C. supervised the project.

Author Information Reprints and permissions information is available at www.nature.com/reprints. Correspondence and requests for materials should be addressed to H.N. (hnienb@gmail.com).

METHODS

Task-and-reward regimen. Two monkeys were trained in a binary, forced-choice disparity-discrimination task (Fig. 1a). They judged whether the central stimulus region appeared in front or behind the surrounding annulus. Trials started upon fixation (within 0.5° of a fixation marker), initiating a 2-s stimulus presentation followed by the appearance of two choice targets (respectively 3° above and below the fixation marker). If the monkeys indicated their decision by a saccade to the correct choice target within 500 ms of the stimulus disappearance, they received liquid rewards. If the monkeys made correct choices on three consecutive trials, the reward on the fourth and on all subsequent correct trials was approximately three times its normal size, until the monkey made an error. After an error, the reward size was reset to its normal size.

Recordings. We recorded extracellular activity from disparity-selective single units in these monkeys' visual areas V2, as described previously[6,30]. Both animals were implanted with scleral search coils in both eyes[31], head-fixation posts and a recording chamber under general anaesthesia. Positions of both eyes (for 17 of 58 neurons for monkey 2, signals were available only for one eye) were measured with a magnetic scleral search system (CNC Engineering) and digitized at 800 Hz. The monkeys viewed the stimuli on EIZO Flexscan F980 monitors in a Wheatstone stereoscope configuration (89-cm viewing distance). All procedures were performed in accord with the US Public Health Service policy on the humane care and use of laboratory animals and all protocols were approved by the National Eye Institute Animal Care and Use Committee.

Stimulus. All stimuli were circular, dynamic, random dot stereograms (50% black and 50% white dots of 99% contrast; dot density, generally 40%; dot size, 0.09°× 0.09°). Each random dot stereogram had a disparity-varying centre (3–5° in diameter) and a 1–2°-wide surrounding annulus at 0° disparity (Fig. 1b). On each frame, all centre dots had the same disparity, but this disparity value changed randomly from frame to frame (96-Hz frame rate). For the condition with no added signal, the disparity on each frame was drawn at random from a uniform distribution of discrete, equally spaced disparities (symmetrical about 0° disparity (central plot in Fig. 1c), encompassing the preferred and the null disparity of each neuron). Signal disparities (always one near and one far disparity) approximately matched the preferred and null disparities of the neuron. Disparity signal was introduced by increasing the probability of the signal disparity on each frame (Fig. 1c).

Psychophysical reverse correlation. Only trials with no added signal were included in the analysis. Each stimulus trial was summarized in a two-dimensional matrix in which each row corresponds to one disparity and each column to one stimulus frame. For each column in this matrix, there is one entry that is 1, corresponding to the disparity on this frame, and all other entries are 0. We then computed the average matrices preceding the monkey's 'near' and 'far' choices. For each of the 200 stimulus frames, the resulting values correspond to the probabilities with which each disparity preceded a 'near' choice or a 'far' choice, respectively. This yields a two-dimensional (time versus disparity) probability distribution. The difference between the probability distributions preceding 'near' choices and 'far' choices defined the psychophysical kernel for each experiment. (Negative disparities were defined as near.)

The kernel shapes change little between monkeys or as a function of the signal disparities (Supplementary Fig. 2). We therefore collapsed all the data into a single psychophysical kernel to maximize the temporal resolution. The average psychophysical kernel (Fig. 2a) was obtained for all experiments for which the simultaneously recorded neurons passed the inclusion criteria. Because the disparity range was adjusted for each neuron, the psychophysical kernel for each experiment was weighted by the number of disparity values included in this experiment (this ranged between 5 and 13 disparity values) and by the number of trials for this experiment. Only data for disparity values (−0.4°, −0.3°, …, 0.4) were included in this average.

As an estimate of the amplitude of the psychophysical kernel, we computed the inner product of the time-averaged psychophysical kernel and the psychophysical kernel (temporally smoothed, 93-ms boxcar window) for each 10-ms bin, and normalized this inner product by its overall mean. Confidence intervals for all measures were derived by resampling.

All analyses were based on the linear kernel of the psychophysical data. Consistency analyses (Supplementary Information) showed that this linear kernel provides an excellent description of the monkeys' behaviour. Further analyses indicated that second-order interactions were negligible (Supplementary Information).

Subspace analysis. The analysis was based on all trials with no added signal. First, the average response of each neuron following one frame of a given disparity, d_i, was computed as a spike density function, $S_i(t)$, smoothed using a 10-ms boxcar window (coloured solid lines in Supplementary Fig. 3a). As an estimate of the impact of one frame of disparity d_i on the firing rate of the neuron, we calculated the total number of spikes, s_i, elicited by one frame of d_i. This metric corresponds to the sum of the mean number of spikes per frame, μ (black line in Supplementary Fig. 3a) and the integral of the deviation of $S_i(t)$ around this mean:

$$s_i = \mu + \int (S_i(t) - \mu)\mathrm{d}t$$

We plot the disparity subspace map, s_i, as a function of d_i in Supplementary Fig. 3b. Performing this analysis on the trials separated according to the monkeys' choice yields the subspace maps separated by choice (Fig. 3a, c). To quantify the choice-dependent modulation in tuning, we plotted the responses in the null-choice trials against those in the preferred-choice trials (Fig. 3b, d), and estimated (using type II regression) the slope (gain change) and the y offset (additive change). We note that because the spike density function is a mean rate calculated separately for each choice, variations in the disparity content of the stimulus that are associated with choice will not produce differences in the subspace maps.

Analysis of choice probabilities. Choice probabilities were obtained for all trials with no added signal as described previously[3,6]. As the psychophysical kernel demonstrates, there are systematic differences in the stimuli preceding the monkeys' choices. Choice probabilities were corrected for this stimulus-induced component (Supplementary Methods).

30. Cumming, B. G. & Parker, A. J. Binocular neurons in V1 of awake monkeys are selective for absolute, not relative, disparity. *J. Neurosci.* **19,** 5602–5618 (1999).
31. Judge, S. J., Richmond, B. J. & Chu, F. C. Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res.* **20,** 535–538 (1980).