

# Separable Codes for Attention and Luminance Contrast in the Primary Visual Cortex

Arezoo Pooresmaeli,<sup>1</sup> Jasper Poort,<sup>1</sup> Alexander Thiele,<sup>2</sup> and Pieter R. Roelfsema<sup>1,3</sup>

<sup>1</sup>The Netherlands Institute for Neuroscience, an Institute of the Royal Netherlands Academy of Arts and Sciences, 1105 BA Amsterdam, The Netherlands,

<sup>2</sup>Institute of Neuroscience, Newcastle University, Newcastle upon Tyne NE2 4HH, United Kingdom, and <sup>3</sup>Department of Integrative Neurophysiology, Centre for Neurogenomics and Cognitive Research, VU University Amsterdam, 1081 HV Amsterdam, The Netherlands

The visual system encodes the features of visual stimuli as well as their behavioral relevance. Stimuli with a high luminance contrast evoke more activity in the visual cortex than stimuli with a low contrast. At the same time, attended stimuli evoke more activity than nonattended stimuli. There is a debate about how visual features and attention jointly determine neuronal activity in the visual cortex. Some studies suggested that attention increases apparent contrast (Reynolds et al., 2000), others that attention amplifies responses by a constant factor (Williford and Maunsell, 2006), and yet others that attention and contrast have largely additive effects (Buras and Boynton, 2007; Thiele et al., 2009). The influence of attention on contrast sensitivity differs between neurons, raising the possibility that attention and contrast could be coded conjointly in a population of neurons. Here we investigate this possibility by recording neuronal activity at multiple sites in the primary visual cortex of macaque monkeys using multielectrode recording techniques and support vector machines to decode attended stimuli as well as stimulus contrast. We find that many, but not all, V1 neurons are influenced by attention and that the effects of attention and contrast are additive on average. Stimulus contrast can be decoded from neuronal responses not strongly modulated by attention, whereas the attended stimulus can be decoded as the difference in activity of cells that are influenced by attention and cells that are not. The success of the approach suggests that visual attention and stimulus contrast are represented by largely separable codes.

## Introduction

It is not completely understood how visual attention interacts with the contrast of a visual stimulus to determine the strength of neuronal responses in the visual cortex. An increase in the contrast of a visual stimulus enhances neuronal activity, and the same is true for attention shifts because attended stimuli evoke stronger responses than nonattended stimuli. Initial single-cell studies suggested that the effect of attention is best described as an increase in the apparent contrast of the stimulus in the receptive field (RF) of many neurons in area V4 (Reynolds et al., 2000) and area MT (Martínez-Trujillo and Treue, 2002), and similar observations were made for the effect of frontal-eye-field microstimulation on the contrast tuning of functional magnetic resonance imaging (fMRI) signals in many additional visual areas (Ekstrom et al., 2009). However, the effects of attention and luminance contrast are not equivalent for all neurons in the visual cortex. In area V4, for example, neuronal responses usually saturate at higher contrasts so that more contrast hardly increases activity, but Williford and Maunsell (2006) found that attention can in-

crease the activity of some V4 cells at the highest contrasts. Thiele et al. (2009) obtained similar results in area V1 because attention increased neuronal activity relatively independently of contrast, and the average interaction between attention and contrast was additive. These studies may appear incompatible at first sight, but a recent elegant model by Reynolds and Heeger (2009) accounted for some of the discrepancies by considering differences in the size of stimuli and the attention field that was employed (Boynton, 2009; Reynolds and Heeger, 2009).

Here we focus on another finding that unifies previous studies: the effect of attention on contrast tuning varies across neurons. For some cells, the effect of attention is similar to an increase in contrast; however, for other cells the effect of attention is additive or even most prominent at higher contrasts (Williford and Maunsell, 2006; Thiele et al., 2009). There is also variation in the strength of attention effects, with some cells strongly influenced by attention, other cells with weaker effects, and yet other neurons not influenced by attention at all (Vidyasagar, 1998; Roelfsema and Spekreijse, 2001; Roelfsema et al., 2004; Roberts et al., 2007). These heterogeneous effects raise the possibility that attention and contrast are, to some degree, represented in different, albeit overlapping circuits (Fig. 1). Neurons weakly modulated by attention could code stimulus contrast regardless of attention shifts. It might also be possible to isolate a pure attentional signal by subtracting the activity of neurons that are weakly influenced by attention from the activity of neurons with stronger attention effects.

To investigate this possibility, we trained monkeys on a curve-tracing task that is associated with attention shifts

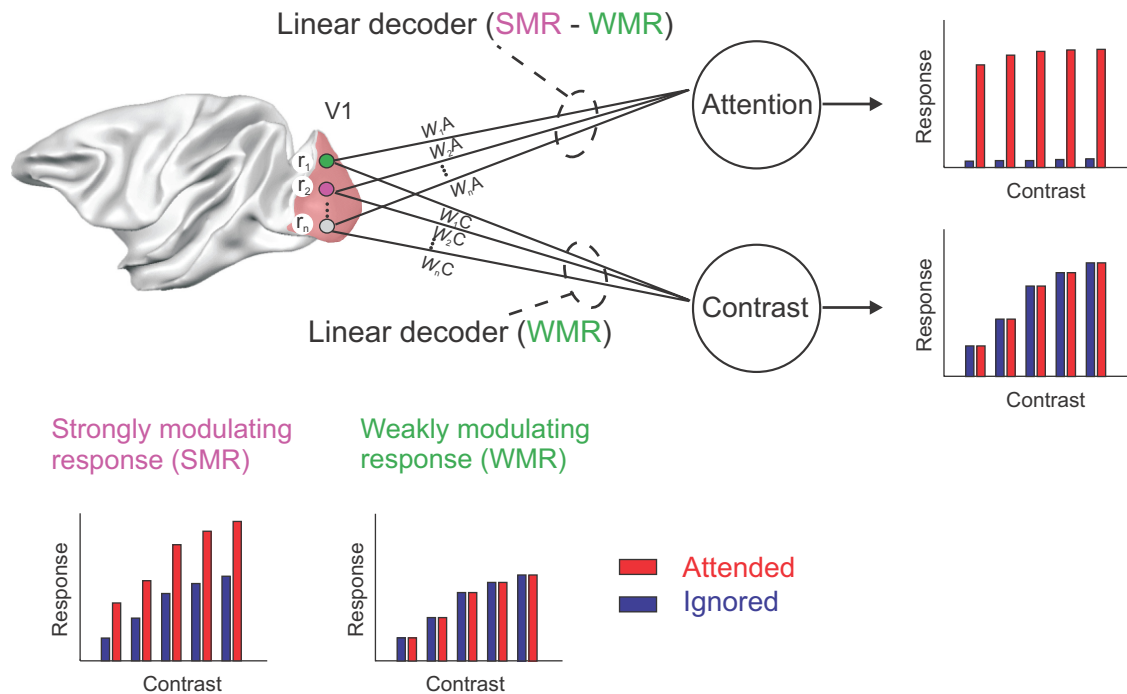
Received March 17, 2010; revised July 7, 2010; accepted July 15, 2010.

This work was supported by a grant from Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO)—Aard en levenswetenschappen, a grant from the Human Frontier Science Program (RGPO07/2007-C), and an NWO-VICI grant awarded to P.R.R. We thank Kor Brandsma, Dave Vleesenbeek, and Anneke Ditevich for biotechnical assistance.

Correspondence should be addressed to Pieter R. Roelfsema, Netherlands Institute for Neuroscience, Meibergdreef 47, 1105 BA Amsterdam, The Netherlands. E-mail: p.roelfsema@nin.knaw.nl.

DOI:10.1523/JNEUROSCI.1388-10.2010

Copyright © 2010 the authors 0270-6474/10/3012701-11\$15.00/0



**Figure 1.** Decoding luminance contrast and the attended stimulus from neuronal responses in area V1. We recorded the responses of neurons at several multiunit recording sites in area V1 ( $r_1, r_2, \dots, r_n$ ) in a task that involves shifts of attention and variations in stimulus contrast. We investigated whether we can decode the attended stimulus and luminance contrast on the basis of the following two types of responses: (1) strongly modulating responses (SMR) that distinguish between attended and nonattended stimuli and are typically also influenced by luminance contrast; and (2) weakly modulating responses (WMR) that are mainly sensitive to stimulus contrast and weakly influenced by attention. We used SVMs to decode the attended stimulus and luminance contrast. We hypothesized that contrast decoding might be possible by focusing on the weakly modulating responses, whereas a relatively pure attention signal might be extracted by evaluating the difference between strongly and weakly modulating responses.

(Roelfsema et al., 1998) and varied stimulus contrast. We recorded the spiking activity of a population of V1 neurons with chronically implanted multielectrode arrays and used support vector machines to decode stimulus contrast as well as the attended stimulus. Our results demonstrate that it is indeed possible to decode both the luminance contrast and the locus of attention from the activity of a population of V1 neurons on a trial-by-trial basis.

## Materials and Methods

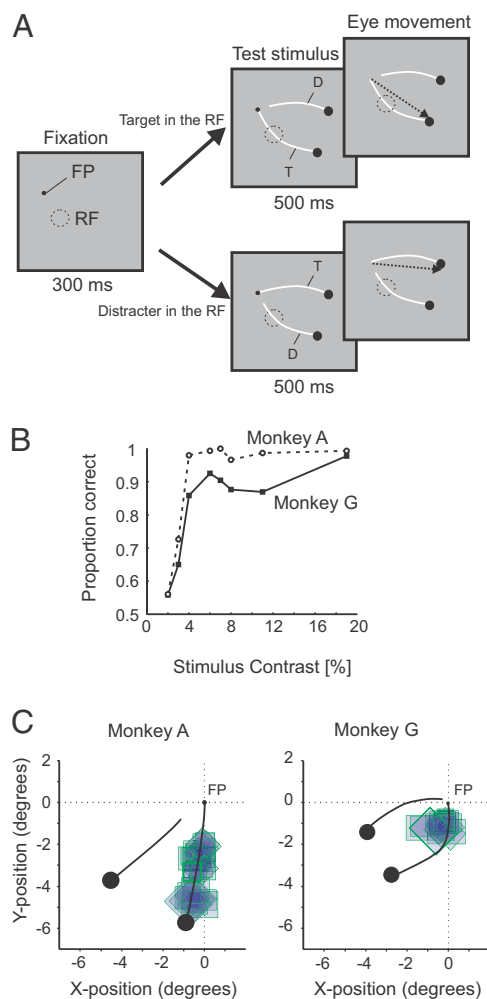
Two male monkeys participated in the present study. In a first operation, a head holder was implanted and a gold ring was inserted under the conjunctiva of one eye for the measurement of eye position. In a separate operation, arrays of  $4 \times 5$  or  $5 \times 5$  electrodes (Cyberkinetics Neurotechnology Systems) were chronically implanted in area V1. The operations were performed under aseptic conditions and general anesthesia. Details of the surgical procedures and the postoperative care have been described elsewhere (Roelfsema et al., 1998, 2007). All procedures complied with the U.S. National Institutes of Health Guidelines for the Care and Use of Laboratory Animals and were approved by the institutional animal care and use committee of the Royal Netherlands Academy of Arts and Sciences.

**Behavioral task.** The animals performed a curve-tracing task where they had to locate a circular target that was connected to the fixation point by a curve (target curve) and to ignore another curve that was a distracter (Fig. 2A). A trial started as soon as the monkey's eye position was within a  $1^\circ \times 1^\circ$  window centered on the fixation point ( $0.2^\circ$  diameter). After an interval of 300 ms, the stimulus, consisting of two curves and two circular targets ( $0.6^\circ$  diameter), appeared on the screen. The background of the display was gray (luminance,  $16.3 \text{ cd/cm}^2$ ), the circular targets and the fixation point were darker than the background (with a luminance of  $9 \text{ cd/cm}^2$ , 28% Michelson contrast), and the curves were brighter than the background. We displayed both curves with the same luminance contrasts: 1.7, 2.6, 4.3, 5.9, 6.8, 8.5, 10.9, or 19.2% (Michelson

contrast). When the stimulus had been in view for 500 ms, the fixation point disappeared and the monkey had to make an eye movement to the circle on the other end of the target curve. All stimulus conditions were randomly interleaved and presented equally often. We recorded at least 40 correct trials for every stimulus in a recording session.

**Recording and data analysis.** Spiking activity was recorded from the chronically implanted multielectrode arrays (Cyberkinetics Neurotechnology Systems). In one animal, we recorded from two arrays at the same time (monkey A), while the activity of one array was recorded in the other animal (monkey G). The electrodes of the arrays had a length of 1 or 1.5 mm. We used TDT (Tucker Davis Technologies) multichannel recording equipment. To record multiunit activity (MUA), we amplified the signal from the electrodes and applied a bandpass filter (300–9000 Hz). This filtered signal (Filt; supplemental Fig. 1, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material) was full-wave rectified, low-pass filtered ( $<200 \text{ Hz}$ ), and sampled at a rate of 760 Hz. The resulting signal is called  $\text{MUA}_E$  (where “E” stands for envelope) because it corresponds to the envelope of the Filt signal, and it provides an instantaneous measure of the number and the size of action potentials of neurons in the vicinity of the electrode tip. The population response obtained with this method is expected to be identical to the population response obtained by pooling across single units.

Another method to record MUA is to set an arbitrary threshold (Schmidt trigger) and to count events crossing this threshold, a signal in spikes per second that has been called  $\text{MUA}_S$ . Supplemental Figure 1 (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material) shows the filtered signal as well as  $\text{MUA}_S$  and  $\text{MUA}_E$  recorded from chronically implanted electrode arrays in a third monkey using a version of the curve-tracing task where the curves had a fixed contrast. Supplemental Figure 1C (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material) illustrates that the shape of the response obtained with the two multiunit signals is similar. Comparable results were obtained for the two monkeys participating in the main experiment (supplemental Fig. 2, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material), although the bandpass-filtered signal was not stored



**Figure 2.** The curve-tracing task and the behavioral performance of the monkeys. **A**, Curve-tracing task. When the monkeys directed their gaze to a fixation point (FP), we presented two curves with circles at their end. One of the curves was connected to the fixation point [target curve (T)], whereas the other one was not [distracter (D)]. The RF of the V1 neurons under study either fell on the target or on the distracter curve. The monkeys had to make an eye movement to the circle at the end of the target curve after a delay of 500 ms. **B**, Behavioral performance of the two monkeys as a function of the luminance contrast of the two curves. **C**, Receptive fields of the recording sites in the two monkeys relative to the two curves.

for this experiment. We conclude that  $MUA_E$  provides a reliable estimate of the average neuronal response, in accordance with previous studies (Cohen and Maunsell, 2009; Supér and Roelfsema, 2005). Recordings with a good signal-to-noise ratio (SNR) were obtained from ~70% of the electrodes (see below for the criterion for inclusion).

RF dimensions of neurons at these electrodes were measured by determining the onset and offset of the visual response to a slowly moving light bar, for each of eight movement directions (Supér and Roelfsema, 2005). The eccentricity of the RFs ranged from  $0.9^\circ$  to  $4.4^\circ$  with an average of  $2.5^\circ$ . The median area of the receptive fields was  $0.8 \text{ deg}^2$  (range,  $0.12$  to  $3.9 \text{ deg}^2$ ), which is slightly larger than the receptive fields of single units at the same eccentricity (Hubel and Wiesel, 1968). This slightly larger RF size is presumably caused by the multiunit recording because the RF of an MUA recording site corresponds to the aggregate of a number of single-unit RFs.

Peristimulus time histograms were calculated in a time window from 300 ms before stimulus onset to 500 ms thereafter, and normalized to the peak response ( $Pe$ ) after subtraction of the spontaneous activity ( $Sp$ , the average activity in the 300 ms fixation interval before stimulus onset). The  $Pe$  was determined as the maximum of the response evoked by the stimuli with the highest contrast, in a time window from 35 to 135 ms

after the stimulus onset. We included a recording site in our analysis if the peak response ( $Pe - Sp$ ) to the highest contrast exceeded the SD of the spontaneous activity,  $\sigma_{Sp}$ , by a factor of three.

We used  $d'$  to quantify how well the neurons at a recording site discriminated between two stimulus conditions.  $d'$  was computed as follows:

$$d' = \frac{(\mu_2 - \mu_1)}{\sigma}, \quad (1)$$

where

$$\sigma = \sqrt{\frac{(n_1 - 1)\sigma_1^2 + (n_2 - 1)\sigma_2^2}{n_1 + n_2 - 2}}, \quad (2)$$

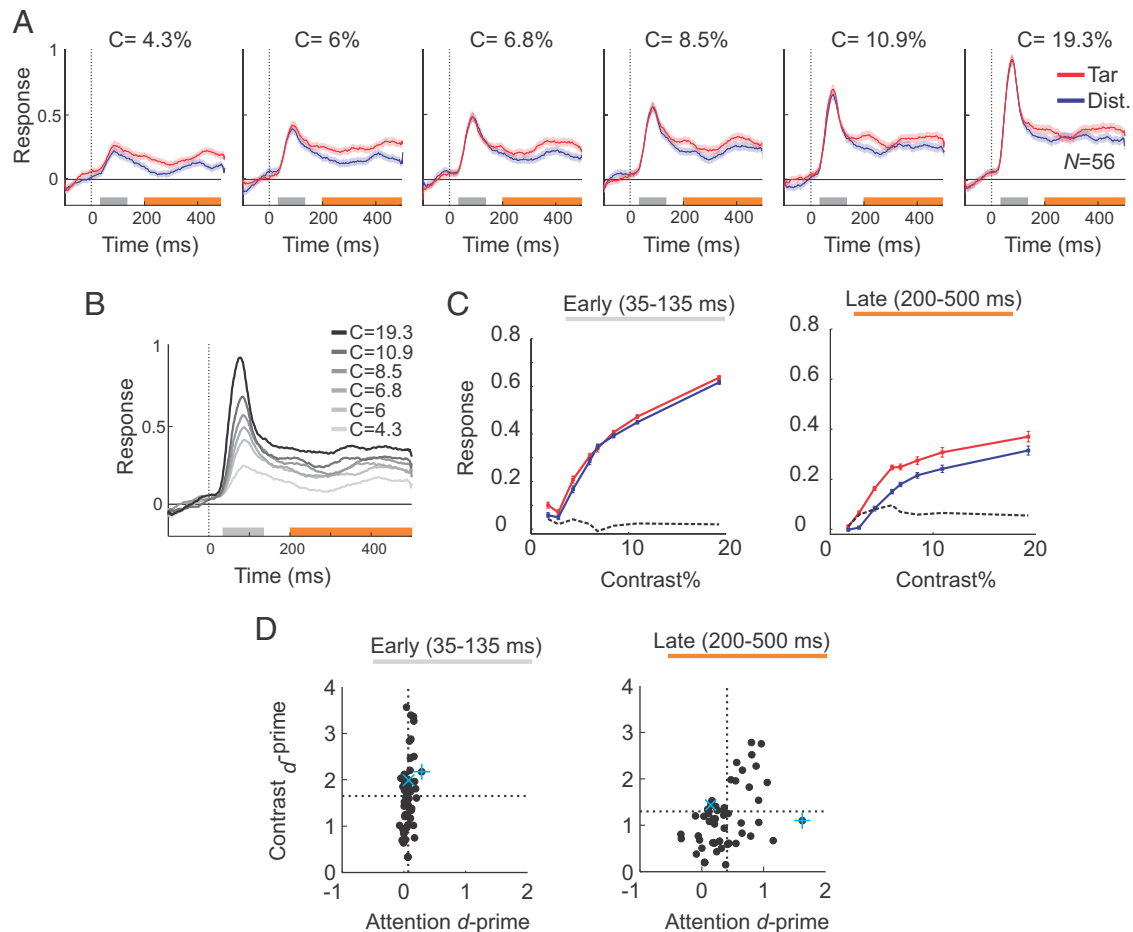
and  $\mu_i$  and  $\sigma_i$  represent the mean and the SD of the neuronal responses across  $n_i$  trials for the two stimuli. To compute attention  $d'$ , we first calculated the  $d'$  for the discrimination between the target and the distracter curve for every contrast, and then averaged these values across contrasts to obtain a single measure of the attention  $d'$ . Similarly, the contrast  $d'$  was computed for the discrimination between the responses to two contrast levels (4.3 and 19.3%), averaged across the two attention conditions.

**Decoding analysis.** We used the SVM<sup>light</sup> implementation (Joachims et al., 1999) of support vector machines (SVMs) (Vapnik, 1995). The input to the SVM classifiers were  $N$   $p$ -dimensional vectors with the neuronal responses,  $r = \{r_1, r_2, \dots, r_N\}$ , where  $N$  is the number of trials and  $p$  is the number of recording sites. The SVM classifier finds the weights that best separate the responses evoked by different contrasts or different attention conditions using a linear discriminant function. To evaluate the accuracy of classification, we used the leave-one-trial-out cross-validation method. In this method, the neuronal responses of all trials but one were used to derive the SVM, which was then used to classify the probe trial that was kept separate. This process was repeated for all the trials, and the reported accuracies represent the percentage of the trials that were correctly classified. To evaluate the performance of the contrast decoder, we trained the SVM with the data of two contrast levels (4.3 and 19.3%) and classified the single trial responses of any pair of contrasts. We tested whether the classification accuracies were significantly different from a baseline classification accuracy assessed with the leave-one-trial-out cross-validation method after randomly assigning class labels (shuffling). We obtained two distributions of the cross-validated classification accuracies for the original and shuffled data by bootstrapping ( $N = 100$ ); i.e., if there were  $t$  trials, we randomly selected  $t$  trials with replacement and repeated all calculations. We considered the classification accuracies to be significant if the 95% confidence intervals of these two distributions did not overlap.

## Results

### Behavioral performance in curve-tracing task

We used a curve-tracing task to control the locus of selective attention. The monkeys had to trace a target curve connected to the fixation point and to ignore a distracter curve that was not connected to the fixation point. At the end of the trial, the animals made a saccade to a larger circle at the end of the target curve. Previous studies demonstrated that human observers direct selective attention to the target curve during this task (Scholte et al., 2001; Houtkamp et al., 2003). Figure 2B shows the performance of the two monkeys in the curve-tracing task (Fig. 2A) as a function of stimulus contrast. Accuracy was high for the higher contrasts and came close to chance levels at the lowest contrasts. At these lower contrast levels, the two curves were difficult to discriminate from the gray background, and it is conceivable that the monkeys directed their attention to the wrong curve on a fraction of trials. Here we will therefore focus on the data from the stimuli for which the monkeys' performance was  $\geq 75\%$  correct (luminance contrast of  $\geq 4.3\%$ ).



**Figure 3.** V1 population responses in the curve-tracing task. **A**, Neuronal responses evoked by the target (red) and distracter (blue) curve at six contrast levels, averaged across all the recording sites ( $N = 56$ ). For our analysis, we measured the neuronal response magnitude in two time windows: an early window that contained the response transient (peak window, 35–135 ms, gray bar) and a later window when the neuronal activity reached a sustained level (200–500 ms, orange bar). **B**, Neuronal responses evoked by curves of different contrasts, averaged across the two attention conditions. **C**, Average contrast response function in the early (35–135 ms) and late (200–500 ms) time windows, for responses evoked by the target (red) or distracter curve (blue). The dashed black lines represent the difference between responses evoked by the target and the distracter curve (blue traces subtracted from red traces). **D**, Attention  $d'$  (abscissa) and contrast  $d'$  (ordinate) across recording sites in the early (left) and late (right) response windows. The dashed lines show the mean  $d'$ .

### Effects of attention and luminance contrast on the V1 population response

We recorded multiunit spiking activity (Supèr and Roelfsema, 2005) at 56 recording sites in area V1 of the two monkeys (17 sites in monkey G, 39 sites in monkey A) (Fig. 2C) in the same sessions in which we obtained the behavioral data described above. The use of chronically implanted electrode arrays permitted simultaneous recording of spikes of all the neurons. The receptive fields of the V1 always fell either on the target or on the distracter curve, and we ensured that the contour element that connected the fixation point to one of the curves was located outside the receptive fields of the neurons so that differences between the conditions could be attributed to attention shifts (Roelfsema et al., 1998).

Figure 3A shows population responses evoked by the target and distracter curve at various levels of luminance contrast. When a segment of one of the curves appeared in the receptive field, the neurons first exhibited a transient response, which was followed by a more sustained response caused by the continuous receptive field stimulation during the 500 ms fixation delay. Attention did not have a strong effect on the transient response, whereas the sustained response evoked by the target curve was stronger than the response evoked by the distracter curve, as

described previously (Roelfsema et al., 1998). To examine the encoding of stimulus contrast, we averaged responses across the target and the distracter curve at the various contrast levels (Fig. 3B). The strength of the population responses increased monotonically with the luminance contrast in both windows, although the responses evoked by different contrasts were more distinct during the peak response than in the later phase of sustained activity.

We next examined the interaction between attention and contrast, in two time windows. The first time window was centered on the transient response (Fig. 3A, 35–135 ms after stimulus onset, gray bar) and was chosen because previous studies demonstrated that attention has little effect during this phase of the neuronal responses in the curve-tracing task (Roelfsema et al., 1998, 2004). The second time window included the sustained response phase (200–500 ms, orange bar in Fig. 3A,B) when the effects of attention are more pronounced. We note that these two time windows were non-overlapping and had different lengths. Figure 3C shows the average contrast response function of the neurons in the early (Fig. 3C, left) and the late window (Fig. 3C, right). The dashed black lines in Figure 3C show the difference between the responses to the target and the distracter curves as an indicator of the strength of attentional modulation. It can be seen

that the effect of attention is greatest in the later phase of the neuronal responses. Importantly, the strength of modulation in this late phase appeared to be largely independent of the stimulus contrast, especially for stimulus contrasts  $>4\%$  (Fig. 3C, right, dashed black curve). To test the statistical significance of these effects, we applied an ANOVA on the neuronal responses in the late window with attention and contrast (6 levels; i.e., contrasts  $>4\%$ ) as factors. As expected, we observed significant main effects of attention ( $F_{(1,660)} = 82, p < 10^{-3}$ ) and contrast ( $F_{(5,660)} = 59, p < 10^{-3}$ ). Importantly, the interaction between attention and stimulus contrast was not significant ( $F_{(5,660)} = 0.62, p = 0.68$ ), indicating that the effect of attention was relatively constant across the various stimulus contrasts. This result is in line with a recent study that demonstrated additive effects of attention and contrast on response strength in area V1 (Thiele et al., 2009). Such an additive effect could be beneficial for the simultaneous decoding of the attended stimulus and stimulus contrast from the neuronal responses (Fig. 1) (see Simultaneous decoding of the attended stimulus and contrast on single trials).

#### Attention and contrast tuning at individual recording sites

The decoding scheme illustrated in Figure 1 relies on differences between tuning of neurons to attention and contrast. We therefore examined how well cells at individual recording sites discriminated between the attention conditions and the luminance contrasts. We computed attention and contrast  $d'$ -primes ( $d'$ ) in the early and late time windows (35–135 and 200–500 ms, respectively) to quantify how well individual recording sites discriminated between attention and contrast conditions on single trials (see Materials and Methods). Figure 3D (left) shows the distribution of attention and contrast  $d'$  across the population of V1 recording sites ( $N = 56$ ) in the early time window. As expected, the transient neuronal responses discriminated well between contrasts (mean  $d' = 1.65$ ; comparison between contrasts of 4.3 and 19.3%) but poorly between attention conditions (mean  $d' = 0.07$ ). In the late time window, the neurons discriminated less well between contrasts (mean  $d' = 1.29$ ; Mann–Whitney  $U$  test,  $p < 0.01$ ) but the discrimination between target and distracter curve improved (mean  $d' = 0.41$ , Mann–Whitney  $U$  test,  $p < 10^{-10}$ ) (Fig. 3D, right). There also were a number of recording sites with attention  $d'$  close to zero in the late window, indicating that these neurons were hardly influenced by the shifts of attention. We considered the possibility that the absence of an attentional effect was caused by a weaker response or a lower SNR of the MUA recordings at some of the recording sites. To examine this possibility, we computed the SNR as the ratio of the peak response to the SD of the spontaneous activity and compared the SNR between recording sites with the highest and the lowest attention  $d'$  values (lowest and highest 33% of sites,  $N = 19$ ). The difference between the SNR values in these two groups of recording sites was not significant ( $p > 0.1$ , Mann–Whitney  $U$  test), which indicates that the absence of attentional effects at weakly modulating recording sites was not caused by a small SNR. Moreover, the heterogeneity in the effects of attention were reproducible, because recording sites without an effect of attention on one day also tended to lack the attentional effect on other days. The attention  $d'$  values were significantly correlated across sites between recordings sessions on subsequent days (monkey A,  $\rho = 0.77$ ; monkey G,  $\rho = 0.88$ ; both  $p < 10^{-5}$ ) and even between sessions  $>6$  months apart (monkey A,  $\rho = 0.70$ ; monkey G,  $\rho = 0.75$ ; both  $p < 0.005$ ). Furthermore, we considered the possibility that the heterogeneity of the effect of attention might be related to the location of the receptive fields along the curve. To investigate

this possibility, we measured the correlation between the attention  $d'$  and the eccentricity of the receptive fields. The lack of a significant correlation ( $\rho = 0.2, p > 0.1$ ) indicates that the strength of the attentional effect does not depend strongly on RF location in accordance with previous findings (Roelfsema et al., 1998; Khayat et al., 2006). Together, these results indicate that the variable strength of the attentional effect across recording sites reflects a genuine heterogeneity across cells.

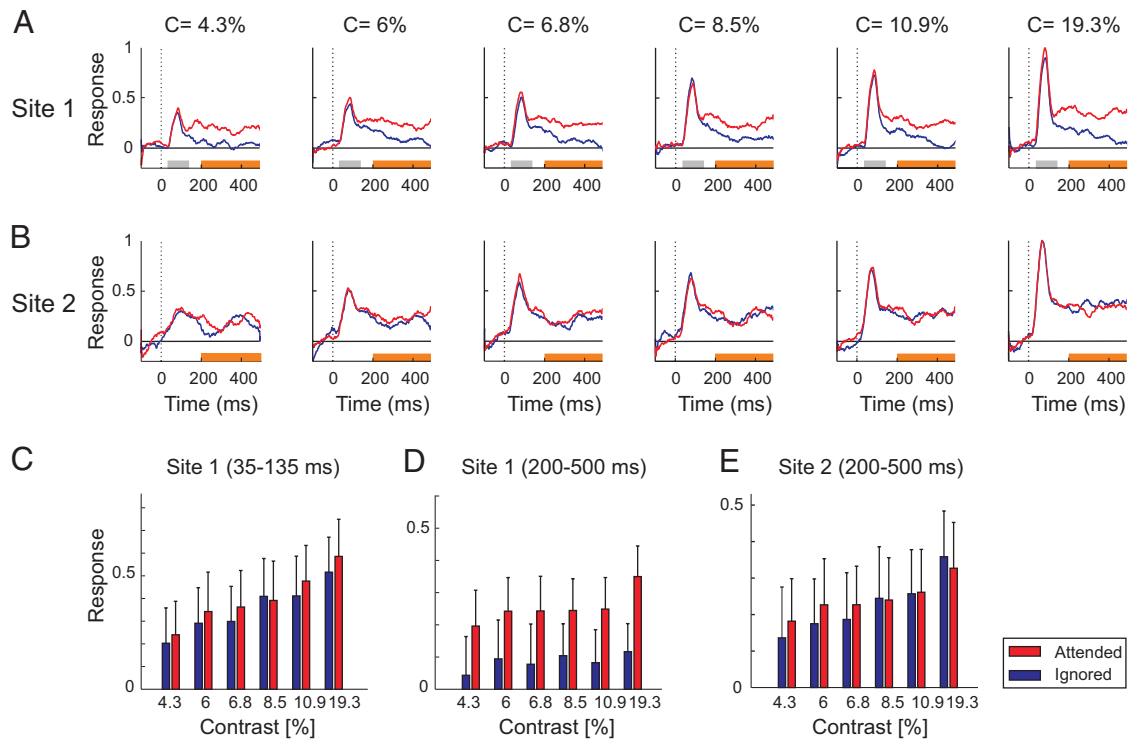
We here exploited the heterogeneity in the contrast and attention tuning across neurons and time windows to decode the attended stimulus and contrast at the same time, using two methods. First, we combined the neuronal responses from the two time windows. Figure 4A illustrates this approach for an example recording site. Neurons at this recording site had weak attentional modulation in the early window ( $d' = 0.08$ ) and were strongly modulated by stimulus contrast (with a  $d'$  of 2.0) (Fig. 4C). The same neurons had stronger attentional modulation in the late window ( $d' = 1.6$ ), whereas the contrast  $d'$  decreased ( $d' = 1.1$ ) (Fig. 4D). Thus, the attended stimulus and the luminance contrast could be decoded in different time windows for this example recording site.

A second possibility to decode the attended stimulus and contrast separately is to combine late window activity of neurons with strong and weak attentional modulation. Figure 4B shows the neuronal responses at a second recording site that was recorded at the same time as the recording site of Figure 4A and was only weakly modulated by attention (attention  $d' = 0.14$ ). The neurons at this second recording site discriminated relatively well between contrasts in the late time window (contrast  $d' = 1.44$ ) (Fig. 4E). We conclude that relatively pure attention and contrast signals also exist in the sustained response phase of different neurons. The examples of Figure 4 were chosen to illustrate the idea of using different time windows and different neurons, and we will now generalize these ideas to simultaneous recordings from a larger population of recording sites.

#### Simultaneous decoding of the attended stimulus and contrast on single trials

We wished to determine the linear combination of responses across a larger number of recording sites that would best decode the attended stimulus and the stimulus contrast, an approach that may benefit from the additive effects of contrast and attention on response strength (Fig. 3C). To this aim, we constructed an “attention decoder” and a “contrast decoder” SVM (see Materials and Methods). The advantage of support vector machines for decoding is that they optimize a linear model for classification that gives rise to good generalization and avoids overfitting of the data (Vapnik et al., 1996; Muñoz and Moguerza, 2006). Moreover, SVMs do not require an assumption about the statistical distribution of the noise across trials. As input to the decoders, we provided the single-trial activity of all V1 recording sites with receptive fields at nearby locations on the same curve (17 in monkey G and 39 in monkey A) (Fig. 2C) in the two time windows. The attention decoder was trained with the stimuli of all the contrasts and had to classify whether the receptive fields fell on the target or the distracter curve regardless of stimulus contrast. The contrast decoder was trained by using the trials of two contrast levels, 4.3 and 19.3%, in both attention conditions.

The attention support vector machine was defined by two weights,  $w_{i,early}^A$  and  $w_{i,late}^A$  for every recording site. The neuronal activity in the two windows  $r_{i,early}$  and  $r_{i,late}$  is multiplied by the weights and then summed across windows and recording sites, and if this sum exceeds a threshold,  $\theta_{Attention}$ , then the decoder



**Figure 4.** Variation in the effects of attention and contrast across two example recording sites that were recorded at the same time. **A**, Responses of neurons at recording site 1 (Fig. 3D, blue plus) exhibited strong attentional modulation. Responses evoked by the target curve are shown in red and responses evoked by the distracter in blue. **B**, Neuronal responses at a recording site with weak attentional modulation (Fig. 3D, blue cross). **C, D**, Magnitude of responses at recording site 1 in the early (**C**) and late (**D**) window evoked by the target (red bars) and distracter curve (blue bars). **E**, Neuronal activity evoked at recording site 2 in the sustained response window.

indicates that the curve falling in the RF is attended, while it indicates that the curve is not attended otherwise (Fig. 1):

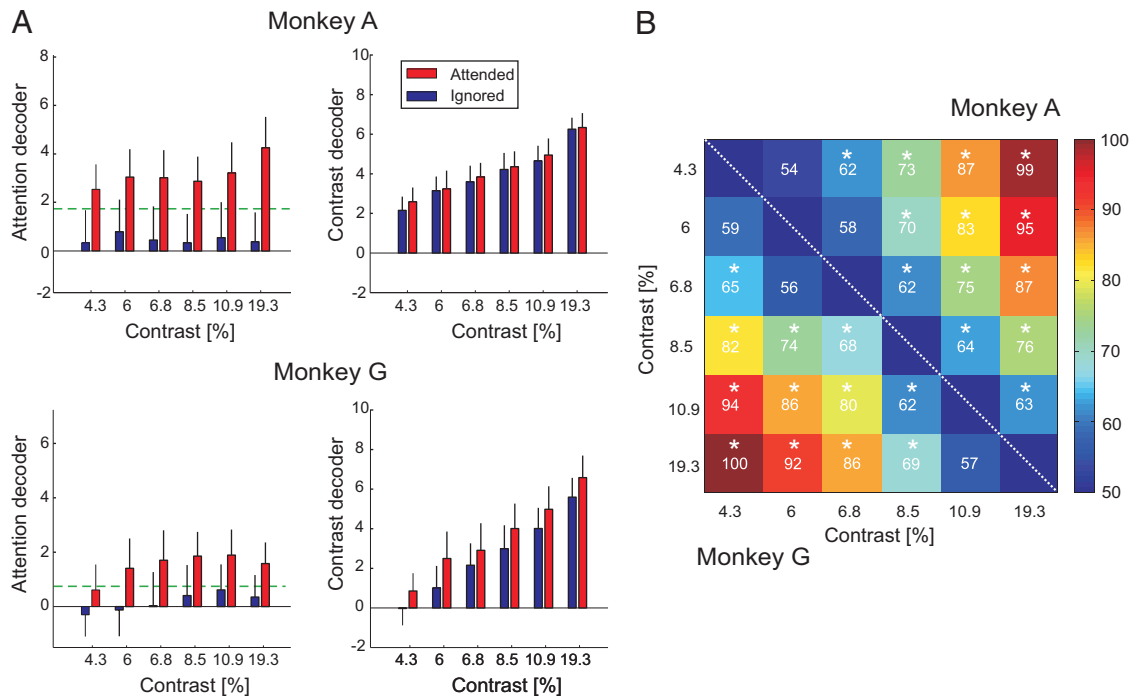
$$\begin{cases} \sum_i w_{i, \text{early}}^A r_{i, \text{early}} + \sum_i w_{i, \text{late}}^A r_{i, \text{late}} > \theta_{\text{Attention}} \Rightarrow \text{Attended} \\ \sum_i w_{i, \text{early}}^A r_{i, \text{early}} + \sum_i w_{i, \text{late}}^A r_{i, \text{late}} < \theta_{\text{Attention}} \Rightarrow \text{Not\_Attended} \end{cases} \quad (3)$$

To train the attention decoder, we used the neuronal responses evoked by the six contrasts  $>4\%$ , because the accuracy of the monkeys was  $<75\%$  for the two lowermost contrasts, and we cannot be confident about where they directed their attention (Fig. 2B) on those low-contrast trials. All trials with the RF on the target curve fell into one class, regardless of contrast, and all trials with the RF on the distracter curve fell into the other class. Figure 5A (left) shows the performance of the attention decoder in the two animals. In monkey A, the mean output of the decoder was stronger when the curve in the receptive fields was attended than when it was not, for every contrast. Thus, a linear combination of the activity of a population of V1 neurons can distinguish between attended and nonattended contours despite large variations in luminance contrast. In monkey G, the attention decoder also permitted classification with a single threshold, except for the responses evoked by the stimulus of 4.3% contrast, where the output of the attended condition was just below the  $\theta_{\text{Attention}}$ . We do not know whether this difference between the results in the two animals was caused by the smaller number of recorded neurons in monkey G or by his poorer performance at this low luminance contrast (Fig. 2). We next investigated how much information was present in the population of neurons about the locus of attention on a single trial by applying the leave-one-trial-out cross-validation method (see Materials and

Methods). The leave-one-trial-out accuracy for the attention decoder was 81% in monkey A and 72% in monkey G, showing that it is indeed possible to isolate the effects of attention by a linear combination of the activity of a population of neurons in area V1.

We next tested how well the performance of the attention decoder generalizes to data not used for training. To this end, we carried out two separate analyses. First, we investigated whether the attention SVM that was trained with the six highest contrasts generalized to the two lower contrasts that were not used for training (because the monkeys' performance was  $<75\%$ ). On correct trials with the two lowest contrasts, the output of the attention decoder was higher for the target curve than for the distracter curve (supplemental Fig. S3, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material), which implies that the decoder indeed generalized to the lower contrasts (supplemental Fig. 3, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). Interestingly, the output of the attention decoder was inverted on incorrect trials, as the higher output was now elicited by the distracter curve. These results suggest that the monkeys attended the wrong curve on error trials and imply that the decoder predicts if the monkey is going to make an error or not (Roelfsema and Spekreijse, 2001).

Second, we investigated how well the performance of an attention decoder that is trained with only one contrast level generalizes to other contrasts. Supplemental Figure 4 (available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material) illustrates that the decoders were most accurate if they were tested with the contrast used for training, generalized reasonably well to neighboring contrasts, and performed poorly for contrasts further apart. We therefore conclude that the linear weights that can separate the attention signal from the effects of contrast require a large enough range of contrasts in the training data.



**Figure 5.** Output of the attention and contrast decoders. **A**, The output of the attention decoder is shown in the left column for the two monkeys, and the output of the contrast decoder in the right column. The ordinates are in arbitrary units. Red bars show responses evoked by the attended stimuli, and blue bars responses evoked by nonattended stimuli as a function of luminance contrast. Error bars denote the SD of the decoder output across trials. The green dashed line depicts the classification threshold of the attention decoder ( $\theta_{Attention}$ ). **B**, Matrix of classification performance for all pairwise contrast comparisons. The values above the diagonal are from monkey A, and the values below the diagonal are from monkey G. For every pair of contrasts, we tested whether the contrast decoder classifies significantly better than chance level. To this end, class labels were shuffled and the 95% confidence interval of the performance of the decoder for shuffled data was calculated by bootstrapping ( $N = 100$ ). The stars show pairs of contrasts where the performance of the decoder was significantly better than chance ( $p < 0.05$ ).

In the Introduction, we suggested that a relatively pure attention signal might be extracted from area V1 by subtracting neuronal responses that are influenced by contrast but not by attention from neuronal activity that is modulated by both contrast and attention (Fig. 1). This could be achieved, for example, by subtracting activity in the early window from activity in the late window. In accordance with this view, the weights  $w_{i,early}^A$  that the SVM assigned to the early window were negative, on average, whereas the weights  $w_{i,late}^A$  in the late window were positive (supplemental Fig. 5A, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). Moreover, we observed a strong positive correlation between  $w_{i,late}^A$  and the attention  $d'$  of the neurons in the late window ( $\rho = 0.92, p < 10^{-10}$ ), as well as a negative correlation between  $w_{i,early}^A$  and the contrast  $d'$  in the early window ( $\rho = -0.62, p < 5.10 \cdot 10^{-7}$ ), indicating that the attention SVM assigned larger positive weights to recording sites that were more sensitive to the locus of attention and larger negative weights to sites that coded luminance contrast reliably (supplemental Fig. 5C,E, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). In other words, the attention decoder indeed subtracted responses strongly tuned to contrast from responses influenced by attention. We noted, however, that the correlation between the contrast  $d'$  and the weight of the attention decoder was also positive in the late window (supplemental Fig. 5E, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material), and that this seems inconsistent with the proposed decoding scheme. We suspected that this positive correlation was caused by variations in the level of sustained activity in the late window. Neurons with a weak sustained response are expected to have a small attention  $d'$  and also a small contrast  $d'$  in the late window. We therefore computed the partial correlation between contrast  $d'$  and the weights

of the attention decoder, factoring out the correlation with the attention  $d'$ . The resulting partial correlation between the contrast  $d'$  and the weights of the attention decoder was not significant, whereas the negative correlation between contrast tuning in the first window and the weights for the attention decoder remained significant after partialization ( $p < 0.01$ ). Similarly, the significant correlation between the attention  $d'$  in the late window and the attention decoder weights was also significant after partialization to the contrast  $d'$  in the late window ( $p < 10^{-10}$ ). These results, together, indicate that the attention decoder subtracted responses tuned to contrast from the responses tuned to attention to obtain a relatively pure attention signal, in accordance with the scheme of Figure 1.

### Decoding of contrast on single trials

The contrast SVM had to decode contrast while ignoring the effects of attention shifts on neuronal activity in area V1. We defined the contrast decoder equivalently to the attention decoder, with weights  $w_{i,early}^C$  and  $w_{i,late}^C$  for the two response windows of every recording site  $i$  and a single threshold,  $\theta_{Contrast}$ . The SVM was trained with only two contrasts, 4.3 and 19.3%, using trials of both attention conditions, but it was tested on the neuronal responses evoked by all contrast levels. Figure 5A (right) shows that the output of the contrast decoder increased monotonically with the contrast of the stimulus in both monkeys. Decoding of contrast worked well in monkey G, although there was a small residual influence of attention on the output of the decoder, and it was excellent in monkey A where the decoder output hardly depended on attention.

Figure 5B illustrates how well the contrast SVM trained with only two contrasts distinguished between any pair of contrasts on a single

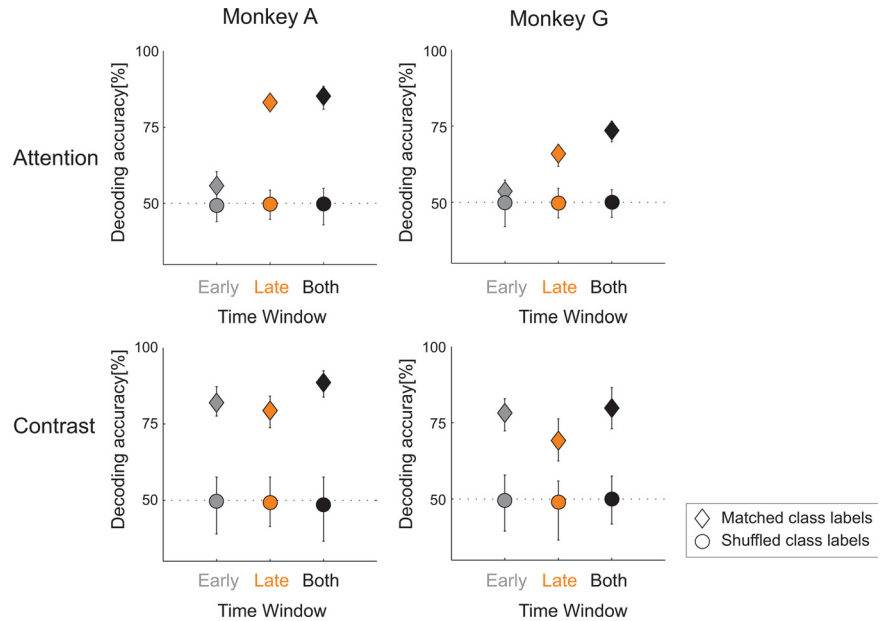
trial by determining the percentage of correctly classified trials (see Materials and Methods). The decoder worked above chance level for a number of the most difficult discriminations (neighboring contrast levels), whereas the trial-by-trial discrimination became more accurate for larger contrast differences. An analysis of the weights of the contrast SVM in the two time windows revealed that the weights were positive, because neurons were sensitive to contrast during their entire response, whereas weights in the early window were stronger than the weights in the late window (supplemental Fig. 5B, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material).

We next investigated the relation between contrast decoder weights of individual recording sites and their contrast and attention  $d'$  values (supplemental Fig. 5D,F, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). As expected, we observed a strong correlation between contrast  $d'$  values and decoder weights in both time windows. There was no significant correlation between attention  $d'$  values in the early window and decoder weights, but a significant positive correlation existed in the late window. This positive correlation may have been caused by the variability in the sustained response level, as was described for the attention decoder above. Indeed, when the correlation between the attention and contrast  $d'$  in the late window was factored out by a partial correlation analysis, the positive correlation between the weights of the contrast decoder and the attention  $d'$  changed into a negative correlation, but this correlation was not significant ( $\rho = -0.16, p > 0.2$ ) (supplemental Fig. 5F, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). In contrast, the significant correlations between contrast  $d'$  and the weights of the contrast decoder remained significant after partialization to the attention  $d'$  values (both  $p < 10^{-7}$ ) (supplemental Fig. 5D, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). These results indicate that the contrast decoder mainly used information from recording sites strongly tuned to contrast, which is also in accordance with the scheme of Figure 1.

It is likely that the approach to train the classifier with two very different contrasts and using the weights to also classify the intermediate contrasts underestimates the information present in the neuronal activity about nearby contrast levels. Indeed, a multi-class decoder that was designed according to a scheme proposed by Kamitani and Tong (2005) was even better able to discriminate between nearby contrast levels (supplemental Fig. 6, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material).

### Decoding of the attended stimulus and contrast in one time window

The above analysis combined neuronal responses in two time windows. A neuronal implementation of this decoding method would have to store traces of the activity in the first window for comparison with the neuronal responses in the second window. We next asked how well the stimulus that is attended and stimulus contrast can be decoded if neuronal activity in only one of these windows is available. We determined the accuracy of decoding in each of the two windows separately and compared the



**Figure 6.** Comparison of classification performances of the attention (top row) and contrast decoder (bottom row) from neuronal responses in the early, late, and both time windows. Dashed lines show the baseline classification level when the class labels were randomly assigned to the responses (shuffling). The error bars denote the 95% confidence intervals.

results to the decoding accuracy when the information in both windows was used (Fig. 6). In this analysis, we chose two neighboring contrast levels for the contrast decoder that were relatively difficult to discriminate (10.9 and 19.3%), and compared decoding performance to a baseline performance that could be achieved when we randomly assigned trials to two arbitrary categories (shuffling) (Fig. 6, circles), thereby ensuring that the classification accuracies are related to the attention conditions and contrasts, and not to regularities in the data due to factors beyond our immediate experimental control. As expected, we obtained the best decoding results by combining the responses in both windows. In the early window, decoding of contrast was good, but decoding of the attended stimulus was close to the baseline, because neuronal responses were hardly influenced by attention during this phase of the response (Fig. 3). In the late window, however, decoding of the attended stimulus was quite good (83.1% in monkey A and 65.9% in monkey G), and the contrast decoder could still distinguish the responses evoked by stimuli with nearby contrasts relatively well on single trials (79.4% in monkey A and 69.1% in monkey G). These results indicate that both the attended stimulus and the stimulus contrast can be decoded during the delayed response phase by linear combinations of the responses of neurons that are weakly and strongly influenced by attention.

### Discussion

Here we have demonstrated that it is possible to decode the locus of attention as well as the luminance contrast of a stimulus by combining information from a population of neurons in area V1. We extracted a relatively pure contrast signal by focusing on the neuronal activity during the first transient response, or on the sustained response phase of neurons that are only weakly influenced by attention shifts. Furthermore, we found that the combined effect of attention and contrast is well described by an additive interaction (Thiele et al., 2009), so that it is possible to isolate a relatively pure attention signal by subtracting neuronal responses that were weakly modulated by attention from re-



sponses that were modulated more strongly (Fig. 1). Thus, it is possible to decode the attended stimulus as well as the stimulus contrast with a reasonable accuracy from a relatively small sample of the neuronal activity in area V1, on a single trial.

Most previous studies used unambiguous cues to direct the monkey's attention (e.g., Reynolds et al., 2000; Williford and Maunsell, 2006). In our task, the attentional cue (the segment connected to the fixation point) became harder to see for low-contrast curves so that the animals may have directed their attention to the wrong curve on some trials. On these error trials, the output of the attention decoder was stronger for the distracter curve than for the target curve, as if the monkey directed his attention to the distracter (supplemental Fig. S3, available at [www.jneurosci.org](http://www.jneurosci.org) as supplemental material). These results are in line with a previous study showing that the representation of a distracter curve in area V1 becomes stronger than the representation of the target curve if the monkey makes a tracing error (Roelfsema and Spekreijse, 2001).

### Decoding accuracy

These data illustrate the virtues of techniques for simultaneous recording of neuronal activity at multiple cortical locations. The single-trial responses of cortical neurons are noisy, and the usual approach to remove this noise is to average neuronal activity across trials. However, neurons in higher cortical areas have to decode the attended stimulus and stimulus contrast by combining the responses of different neurons rather than different trials of the same cell (Deadwyler and Hampson, 1997; Nicolelis et al., 1997; Churchland et al., 2007). The responses of different V1 neurons exhibit correlated fluctuations, and these so-called noise correlations can reduce the amount of information that is carried about the visual stimulus on a single trial (Shadlen et al., 1996; Abbott and Dayan, 1999). Conventional methods that allow recording from one neuron at a time cannot estimate the noise correlation, and therefore do not provide realistic estimates of the amount of information present in the responses of a population of neurons. To test the effect of the noise correlations on decoding accuracy, we also emulated the situation where responses are recorded for one recording site at a time. We removed the noise correlations by shuffling trials of different recording sites (this is equivalent to combining recordings from different sessions) and recomputed the accuracy of the attention decoder. The accuracy of the attention decoder increased from 82 to 92% in monkey A and from 75 to 85% in monkey G. This result shows that the noise correlations indeed deteriorate decoding of responses evoked by the same curve (Poort and Roelfsema, 2009). Thus, simultaneous recordings offer a more realistic impression of the amount of information that visual neurons carry about the locus of attention and stimulus contrast.

The linear multivariate classification method used by us (SVM) automatically takes the redundant information present in the correlations between neuronal responses into account. Decoding of stimulus information using SVMs is now used commonly in fMRI research (Haynes and Rees, 2005; Kamitani and Tong, 2005; Kamitani and Tong, 2006; Haynes et al., 2007; Li et al., 2007; Mourão-Miranda et al., 2007). These decoding techniques can also be used for single-unit and multiunit data (Stark and Abeles, 2007), and here we applied them for the simultaneous decoding of the attended stimulus and stimulus contrast in the visual cortex. Due to the high signal-to-noise ratio and temporal resolution of multiunit activity compared with the fMRI signal, we could decode the attended stimulus and luminance

contrast at the same time, and on a finer time scale (time windows <400 ms) than can be achieved with fMRI.

There are a number of reasons why it is likely that our results provide a lower estimate of the amount of information that is present in a small population of neurons. First, we showed recently that noise correlations in area V1 are caused by relatively global trial-to-trial fluctuations in the overall level of neuronal activity (Chen et al., 2006; Poort and Roelfsema, 2009). These noise correlations reduce the amount of information about the locus of attention when combining neurons with receptive fields on the same curve, as in the present study. However, the difference between the response strengths evoked by the target and distracter curve is only little influenced by these global activity fluctuations, and it is possible to decode the locus of attention with an accuracy of ~90% on a single trial with an average of only four recording sites, provided that some of the receptive fields fall on different curves (Poort and Roelfsema, 2009). It is therefore likely that the accuracy of attention decoding in the present study would have increased had we recorded from neurons with receptive fields on different curves. Two recent studies demonstrated that attention also improves the quality of the sensory representation in area V4 by suppressing noise correlations (Cohen and Maunsell, 2009; Mitchell et al., 2009). This effect was not observed in area V1 in the curve-tracing task, however, because the average noise correlation did not differ between attended and nonattended curves (Roelfsema et al., 2004).

A second reason why we may have underestimated the amount of information is that we used two large time windows, one for the transient response and the other for the sustained response phase. A fine-grained analysis of the temporal profile of the neuronal responses might have increased the amount of information about attention and contrast. Such an analysis would become sensitive to the latency of neuronal responses, and response latency in areas V1 (Gawne et al., 1996), V4 (Williford and Maunsell, 2006), and MT (Thiele et al., 1999) depends on stimulus contrast. Information about response latency on individual trials might therefore have improved contrast decoding.

We note, however, that the integration of activity from different time windows requires storage of traces of the neuronal activity in previous time windows. This storage could occur, for example, in areas of frontal cortex where neurons code the traces of previously presented sensory stimuli (Miller and Cohen, 2001; Hernández et al., 2002). It is unclear, however, whether the information that the V1 peak responses carry about stimulus contrast is stored in this manner, and we therefore also decoded stimulus contrast using the sustained response phase only. Accuracy of decoding in the late time window was only slightly worse compared with the condition where both windows were used. This is because many V1 neurons retain their sensitivity to stimulus contrast in the later response phase.

### Implications for models of visual attention

The efficacy of our decoding algorithm relies on the heterogeneity of the effects of attention and contrast on the activity of neurons across recording sites. Neurons at some of the recording sites were mainly affected by attention during their sustained response, whereas neurons at other recording sites were only affected by contrast and hardly influenced by attention shifts, and the neurons at yet other sites were modulated by both attention and contrast (Fig. 4). This heterogeneity is not unexpected and has been observed in many, if not all, previous studies on the representation of attended versus unattended stimuli (Spitzer et al., 1988; Treue and Maunsell, 1996; Roelfsema et al., 1998, 2004;

Reynolds et al., 1999; Martínez-Trujillo and Treue, 2002; Roberts et al., 2007) as well as in studies on the coding of luminance contrast (Albrecht and Hamilton, 1982; Sclar et al., 1990; Thiele et al., 2000).

Previous studies suggested that the effect of attention on the neuronal responses resembles an increase in the effective contrast of the stimulus driving the cell (Reynolds et al., 2000; Treue, 2004), and a recent fMRI study observed similar effects of frontal-eye-fields microstimulation on the visual cortex (Ekstrom et al., 2009). The present results do not support the strictest interpretation of these “contrast-gain” models of attention. If the only effect of attention were an apparent increase of luminance contrast, then it would not have been possible to decode both factors. Nevertheless, our results are compatible with previous findings that the interaction between attention and contrast is well described by a contrast gain model for some cells (Reynolds et al., 2000; Treue, 2004), whereas other neurons conform to a response gain model or an additive model (Williford and Maunsell, 2006; Thiele et al., 2009). Our results indicate that this heterogeneity, and in particular the presence of neurons little influenced by attention, is beneficial for the simultaneous decoding of the attended stimulus and luminance contrast.

The existence of distinct representations of attention and stimulus contrast in the visual cortex has a number of advantages. First, it permits a veridical representation of the luminance contrast of visual stimuli regardless of attention shifts. Attention improves contrast sensitivity (Huang and Dobkins, 2005), and a number of psychophysical studies demonstrated that attention can increase perceived contrast (Carrasco et al., 2004; Ling and Carrasco, 2006; Störmer et al., 2009), an effect that may be related to the residual effects of attention on our contrast decoder (Fig. 5). However, other studies have observed only weak effects of attention on perceived contrast (Prinzmetal et al., 1997; Schneider, 2006; Schneider and Komlos, 2008; Palmer and Moore, 2009). Our results show that attention need not alter perceived contrast because there are also V1 neurons that are mainly affected by luminance contrast and are not influenced by attention shifts. The second advantage is that shifts of visual attention need not be dominated by high-contrast stimuli. If the codes of attention and contrast differ, attention can also be directed to low-contrast stimuli, which is in accordance with human psychophysics (Pashler et al., 2004; Einhäuser et al., 2008). In the contour grouping task of the present study, attention is directed to all contour elements of the target curve so that they are grouped in perception (Roelfsema, 2006). The separation of the neuronal codes for attention and contrast would permit the grouping of contour elements with varying contrasts if they belong to the same curve. Neurons with weak attentional modulation can code the contrast of the contour elements veridically, whereas the neurons with stronger modulation can label the set of contour elements that belong to the target curve.

## References

- Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11:91–101.
- Albrecht DG, Hamilton DB (1982) Striate cortex of monkey and cat: contrast response function. *J Neurophysiol* 48:217–237.
- Boynton GM (2009) A framework for describing the effects of attention on visual responses. *Vision Res* 49:1129–1143.
- Buracas GT, Boynton GM (2007) The effect of spatial attention on contrast response functions in human visual cortex. *J Neurosci* 27:93–97.
- Carrasco M, Ling S, Read S (2004) Attention alters appearance. *Nat Neurosci* 7:308–313.
- Chen Y, Geisler WS, Seidemann E (2006) Optimal decoding of correlated neural population responses in the primate visual cortex. *Nat Neurosci* 9:1412–1420.
- Churchland MM, Yu BM, Sahani M, Shenoy KV (2007) Techniques for extracting single-trial activity patterns from large-scale neural recordings. *Curr Opin Neurobiol* 17:609–618.
- Cohen MR, Maunsell JH (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci* 12:1594–1600.
- Deadwyler SA, Hampson RE (1997) The significance of neural ensemble codes during behavior and cognition. *Annu Rev Neurosci* 20:217–244.
- Einhäuser W, Rutishauser U, Koch C (2008) Task-demands can immediately reverse the effects of sensory-driven saliency in complex visual stimuli. *J Vis* 8:2.1–2.19.
- Ekstrom LB, Roelfsema PR, Arsenault JT, Kolster H, Vanduffel W (2009) Modulation of the contrast response function by electrical microstimulation of the macaque frontal eye field. *J Neurosci* 29:10683–10694.
- Gawne TJ, Kjaer TW, Richmond BJ (1996) Latency: another potential code for feature binding in striate cortex. *J Neurophysiol* 76:1356–1360.
- Haynes JD, Rees G (2005) Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8:686–691.
- Haynes JD, Sakai K, Rees G, Gilbert S, Frith C, Passingham RE (2007) Reading hidden intentions in the human brain. *Curr Biol* 17:323–328.
- Hernández A, Zainos A, Romo R (2002) Temporal evolution of a decision-making process in medial premotor cortex. *Neuron* 33:959–972.
- Houtkamp R, Spekreijse H, Roelfsema PR (2003) A gradual spread of attention during mental curve tracing. *Percept Psychophys* 65:1136–1144.
- Huang L, Dobkins KR (2005) Attentional effects on contrast discrimination in humans: evidence for both contrast gain and response gain. *Vision Res* 45:1201–1212.
- Hubel DH, Wiesel TN (1968) Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 195:215–243.
- Joachims T, Schölkopf B, Smola A (1999) *Advances in kernel methods—support vector learning*. Cambridge, MA: MIT.
- Kamitani Y, Tong F (2005) Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8:679–685.
- Kamitani Y, Tong F (2006) Decoding seen and attended motion directions from activity in the human visual cortex. *Curr Biol* 16:1096–1102.
- Khayat PS, Spekreijse H, Roelfsema PR (2006) Attention lights up new object representations before the old ones fade away. *J Neurosci* 26:138–142.
- Li S, Ostwald D, Giese M, Kourtzi Z (2007) Flexible coding for categorical decisions in the human brain. *J Neurosci* 27:12321–12330.
- Ling S, Carrasco M (2006) When sustained attention impairs perception. *Nat Neurosci* 9:1243–1245.
- Martínez-Trujillo J, Treue S (2002) Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron* 35:365–370.
- Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24:167–202.
- Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63:879–888.
- Mourão-Miranda J, Friston KJ, Brammer M (2007) Dynamic discrimination analysis: a spatial-temporal SVM. *Neuroimage* 36:88–99.
- Muñoz A, Moguerza JM (2006) Estimation of high-density regions using one-class neighbor machines. *IEEE Trans Pattern Anal Mach Intell* 28:476–480.
- Nicolelis MA, Ghazanfar AA, Faggin BM, Votaw S, Oliveira LM (1997) Reconstructing the engram: simultaneous, multisite, many single neuron recordings. *Neuron* 18:529–537.
- Palmer J, Moore CM (2009) Using a filtering task to measure the spatial extent of selective attention. *Vision Res* 49:1045–1064.
- Pashler H, Dobkins K, Huang L (2004) Is contrast just another feature for visual selective attention? *Vision Res* 44:1403–1410.
- Poort J, Roelfsema PR (2009) Noise correlations have little influence on the coding of selective attention in area V1. *Cereb Cortex* 19:543–553.
- Prinzmetal W, Nwachuku I, Bodanski L, Blumenfeld L, Shimizu N (1997) The phenomenology of attention: 2. Brightness and contrast. *Conscious Cogn* 6:372–412.
- Reynolds JH, Heeger DJ (2009) The normalization model of attention. *Neuron* 61:168–185.
- Reynolds JH, Chelazzi L, Desimone R (1999) Competitive mechanisms subserve attention in macaque areas V2 and V4. *J Neurosci* 19:1736–1753.
- Reynolds JH, Pasternak T, Desimone R (2000) Attention increases sensitivity of V4 neurons. *Neuron* 26:703–714.

- Roberts M, Delicato LS, Herrero J, Gieselmann MA, Thiele A (2007) Attention alters spatial integration in macaque V1 in an eccentricity dependent manner. *Nat Neurosci* 10:1483–1491.
- Roelfsema PR (2006) Cortical algorithms for perceptual grouping. *Annu Rev Neurosci* 29:203–227.
- Roelfsema PR, Spekreijse H (2001) The representation of erroneously perceived stimuli in the primary visual cortex. *Neuron* 31:853–863.
- Roelfsema PR, Lamme VA, Spekreijse H (1998) Object-based attention in the primary visual cortex of the macaque monkey. *Nature* 395:376–381.
- Roelfsema PR, Lamme VA, Spekreijse H (2004) Synchrony and covariation of firing rates in the primary visual cortex during contour grouping. *Nat Neurosci* 7:982–991.
- Roelfsema PR, Tolboom M, Khayat PS (2007) Different processing phases for features, figures, and selective attention in the primary visual cortex. *Neuron* 56:785–792.
- Schneider KA (2006) Does attention alter appearance? *Percept Psychophys* 68:800–814.
- Schneider KA, Komlos M (2008) Attention biases decisions but does not alter appearance. *J Vis* 8:3.1–3.10.
- Scholte HS, Spekreijse H, Roelfsema PR (2001) The spatial profile of visual attention in mental curve tracing. *Vision Res* 41:2569–2580.
- Sclar G, Maunsell JH, Lennie P (1990) Coding of image contrast in central visual pathways of the macaque monkey. *Vision Res* 30:1–10.
- Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci* 16:1486–1510.
- Spitzer H, Desimone R, Moran J (1988) Increased attention enhances both behavioral and neuronal performance. *Science* 240:338–340.
- Stark E, Abeles M (2007) Predicting movement from multiunit activity. *J Neurosci* 27:8387–8394.
- Störmer VS, McDonald JJ, Hillyard SA (2009) Cross-modal cueing of attention alters appearance and early cortical processing of visual stimuli. *Proc Natl Acad Sci U S A* 106:22456–22461.
- Supèr H, Roelfsema PR (2005) Chronic multiunit recordings in behaving animals: advantages and limitations. *Prog Brain Res* 147:263–282.
- Thiele A, Distler C, Hoffmann KP (1999) Decision-related activity in the macaque dorsal visual pathway. *Eur J Neurosci* 11:2044–2058.
- Thiele A, Dobkins KR, Albright TD (2000) Neural correlates of contrast detection at threshold. *Neuron* 26:715–724.
- Thiele A, Pooremaeli A, Delicato LS, Herrero JL, Roelfsema PR (2009) Additive effects of attention and stimulus contrast in primary visual cortex. *Cereb Cortex* 19:2970–2981.
- Treue S (2004) Perceptual enhancement of contrast by attention. *Trends Cogn Sci* 8:435–437.
- Treue S, Maunsell JH (1996) Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* 382:539–541.
- Vapnik V (1995) *The nature of statistical learning theory*. New York: Springer.
- Vapnik V, Golowich S, Smola A (1996) Support vector method for function approximation, regression estimation, and signal processing. *Adv Neural Inf Process Syst* 9:281–287.
- Vidyasagar TR (1998) Gating of neuronal responses in macaque primary visual cortex by an attentional spotlight. *Neuroreport* 9:1947–1952.
- Williford T, Maunsell JH (2006) Effects of spatial attention on contrast response functions in macaque area V4. *J Neurophysiol* 96:40–54.