Neural Correlates of Visual Localization and Perisaccadic Mislocalization

Bart Krekelberg,[1,*] Michael Kubischik,
Klaus-Peter Hoffmann, and Frank Bremmer[2]
Department of General Zoology and Neurobiology
Ruhr University Bochum
44780 Bochum
Germany

Summary

While reading this text, your eyes jump from word to word. Yet you are unaware of the motion this causes on your retina; the brain somehow compensates for these displacements and creates a stable percept of the world. This compensation is not perfect; perisaccadically, perceptual space is distorted. We show that this distortion can be traced to a representation of retinal position in the medial temporal and medial superior temporal areas. These cells accurately represent retinal position during fixation, but perisaccadically, the same cells distort the representation of space. The time course and magnitude of this distortion are similar to the mislocalization found psychophysically in humans. This challenges the assumption in many psychophysical studies that the perisaccadic retinal position signal is veridical.

Introduction

Spatial localization around the time of a saccade provides a rare glimpse of the visual mechanisms that continuously operate to provide us with a stable percept of the world. Matin and Pearce (1965) started a long series of experiments and controversies when they showed that a visual stimulus, briefly presented just before a saccade, is mislocalized in the direction of the saccade. Since this report, most authors have interpreted the mislocalization as a failure of the visual system to correctly remap its presaccadic coordinate system to the postsaccadic coordinate system (Bischof and Kramer, 1968; Dassonville et al., 1992; Honda, 1991; Mateeff, 1978; Matin and Pearce, 1965; Ross et al., 1997).

In line with this, several authors have suggested that perisaccadic mislocalization is related to the dynamics of receptive fields in various cortical areas. For instance, there is evidence of perisaccadic changes in the size and center of receptive fields in the lateral intraparietal area (Duhamel et al., 1992; Kubischik and Bremmer, 1999), the frontal eye field (Umeno and Goldberg, 1997), the superior colliculus (Walker et al., 1995), and area V4 (Tolias et al., 2001). Without an explicit assumption about how these neurons encode position, however, it is difficult to establish a firm link between these physio-

logical findings and the psychophysics of a distorted perception of space.

In other words, to find a neural correlate of mislocalization, one first has to find a neural correlate of localization. In our search for this correlate, we focused on the dorsal stream of the primate brain, as this is commonly thought to be most relevant to encoding *where* stimuli are (Mishkin et al., 1983). Specifically, we recorded from the medial temporal area (MT), the medial superior temporal area (MST), the ventral intraparietal area (VIP), and the lateral intraparietal area (LIP). Not only are these cells part of the putative where pathway, they also carry information on eye position and those in LIP show saccade-related activity: all properties that, intuitively, seem to be related to perisaccadic position perception.

If the neurons we recorded from are truly part of a where pathway, it should be possible for an ideal observer to interpret their spikes in terms of the spatial location of a stimulus. Our analysis finds an interpretation of a neuron's firing rates that an ideal observer can use to extract position information. This interpretation—called a codebook—explicitly links neural firing to the percept an ideal observer would have. Our working hypothesis is that mislocalization occurs around the time of saccades because the neurons that represent position during normal operation are in some, as yet unspecified, manner "disturbed" by the imminent saccade. A downstream area, modeled here by the ideal observer, is unaware of this disturbance and therefore derives an erroneous representation of position.

Results

In our experiments, monkeys faced a 60° by 60° tangent screen and fixated a small dot 10° off the vertical midline. The monkeys were trained to make an immediate saccade to the dot when it jumped to the position 10° on the other side of the midline (target position) (see Figure 1 and Experimental Procedures). There were three kinds of trials, which were randomly interleaved. In the first, the monkey fixated the starting position and a large luminous bar was flashed at one of six horizontal positions at least 300 ms before a saccade (preflashes). Then, after the monkey had made a saccade and fixated the target position for at least 300 ms, a second bar was flashed at the same position (postflashes). In the second kind of trial, only one bar was flashed in a time window of ±200 ms around the saccade (periflashes). All bar positions occurred equally often in the pre-, peri-, and postepochs. In the third kind of trial, no bar was ever flashed.

Eye Position

To assess position encoding before, during, and after saccades, it is clearly critical to have accurate knowledge of the eye position. Not only because we wish to assess localization at various times around the onset of a saccade, but also because mislocalization depends on the retinal position of a flashed stimulus (Bischof and

*Correspondence: bart@salk.edu
[1]Present address: The Salk Institute, Vision Center Laboratory, 10100 N. Torrey Pines Rd., La Jolla, California 92037.
[2]Present address: Department of Physics, Marburg University, Renthof 7, D-35032 Marburg, Germany.
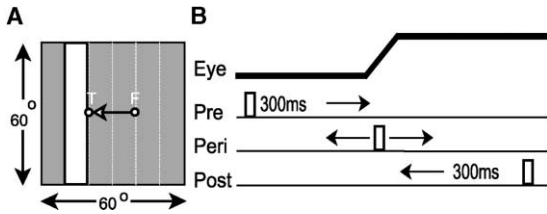
Figure 1. Experimental Paradigm

(A) The screen, flashed bars, fixation (F), and target (T) in our setup.
(B) In a single trial, bars are flashed while the monkey is fixating, long before (pre) and after (post) the saccade. Or they are flashed perisaccadically: within ±200 ms of saccade onset. In control trials, no bar was flashed.

Kramer, 1968; O'Regan, 1984). Hence, to be sure that the effects we find are not due to time-varying changes in eye position, we analyzed fixation accuracy and precision. Figure 2A shows the horizontal eye position from long before until long after the saccade, averaged over all trials in the data set. Clearly, fixation is accurate until saccade onset. Figure 2B compares the fixation accuracy among the pre-, peri-, and postepochs. For the periepoch, we included eye positions up until the onset of the saccade. Fixation is accurate in all three epochs. The jitter of the eye position (standard deviation of the eye position in a single epoch) is shown in Figure 2C and is also very similar in all three epochs. The difference in fixation among the three epochs, and even the spread in eye positions during the saccade, is smaller than either the width of the visual stimuli or the size of mislocalization effects we discuss below.

Position Decoding

For each stimulus, we determined the average response in a window from 50 to 250 ms post-stimulus onset. From these responses we calculated, for each neuron, the conditional probability of observing a particular firing rate given the presentation of a flash at a particular position. Using Bayes' rule this probability can be converted to the probability that an observed firing rate was caused by stimulation at a particular position. We call

such a Bayesian lookup table, linking firing rates to likely stimulated positions, a codebook. For each cell we determined such a codebook. Figure 3 shows an example for a single cell. We assumed that cells provide independent estimates, and we constructed a population codebook by simply combining the single-cell codebooks. In Experimental Procedures we discuss the precise parameterization of the codebooks: we chose the parameterization that most accurately encodes position during fixation. Hence, within the class of parameterizations we investigated, these are optimal for position encoding.

Localization during Fixation

To be a candidate for the encoding of position, a cortical area must consistently relate firing rate to position during fixation. In our analysis this means that we should determine how well the codebooks perform for pre- and postflashes. More specifically, it is the transsaccadic generalization of the codebook based on preflashes to the decoding of postflashes (and vice versa) that is relevant. To test this, we need to make an assumption about the coordinate system these cells encode information in: does the rate of a cell provide information on the position in the world or the position on the retina? Given that the posterior parietal cortex is involved in multiple coordinate transformations (for a review, see Snyder, 2000), there is no strong a priori reason to choose a specific coordinate system. Instead, we determined both the ability to encode world position and the ability to encode retinal position.

In both cases we determined a codebook based on the presaccadic rates and decoded the postsaccadic responses. The performance of the codebooks was determined with a bootstrap analysis. For each stimulated position x, we determined how often it was decoded to be at position y. For a correct decoding, x equals y. Figure 4A shows the average percentage correct of the codebooks if we assume that these cells encode world position. As there are six stimulated world positions, the chance level of performance is 17%. None of the codebooks perform significantly above chance. Figure 4B shows the performance of the same cells on the
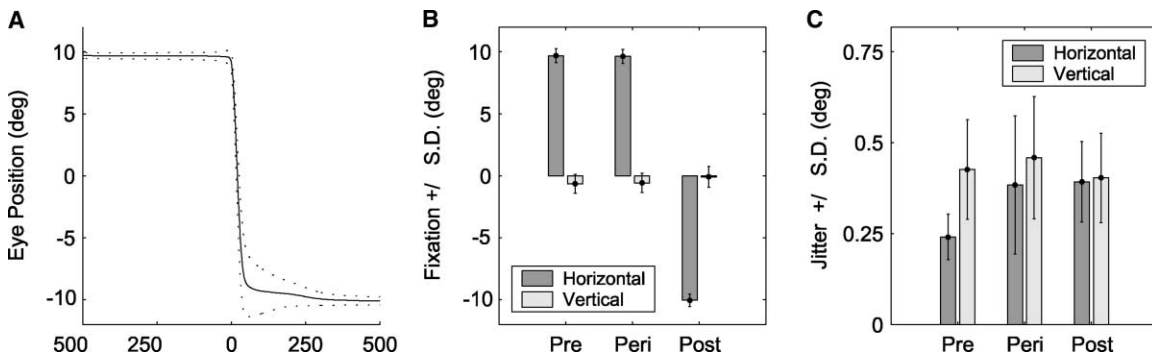

Figure 2. Eye Position

(A) The average eye position averaged over all 55,000 trials as a function of the time to the saccade. The dashed lines show the standard deviation.
(B) The average horizontal and vertical eye position in pre-, peri-, and postepochs.
(C) The average deviation (jitter) from the fixation point during an epoch. Error bars show standard deviations of the jitter over all trials. The periepoch averages only include data up to the start of the saccade.
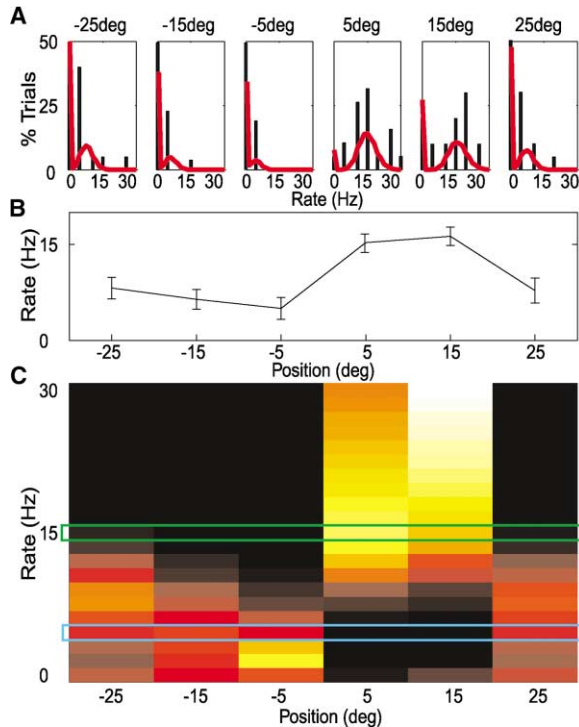
Figure 3. An Example Codebook for a Single Cell (C200-1)

(A) Histograms of the actually observed spike rates and (red line) the parametric description we used to describe this cell's conditional firing rate probability: P(r|x).

(B) This cell's conventional spatial tuning curve that relates position of a flash to an evoked mean firing rate.

(C) The codebook for this cell. The color codes the probability that a flash at a particular position evokes a given firing rate. Black is zero probability, white is the maximum probability. Decoding works as follows: if this cell fires at 15 Hz, two flash positions could have been stimulated (green box). The most probable stimulus is at 5°. This cell "votes" for position 5 with a large weight. A flash at 15° could also lead to a 15 Hz firing rate in this cell; this position gets a smaller vote. If this same cell fires at 5 Hz, there is much more uncertainty about the stimulus (blue box). A flash at −15° is the most probable stimulus, but flashes at −25, −5, or 25 could also evoke this firing rate. The cell votes for each position with a weight proportional to the probability at that position. The ambiguity that results from partial votes for multiple positions is resolved at the population level by combining the votes from all cells and choosing the position that received the largest sum of votes.

decoding of retinal position. As there are only four retinal positions that are stimulated by both pre- and post-flashes, the codebook is restricted to four positions and chance performance is 25%. The performance of the neurons in the superior temporal sulcus (MT and MST, analyzed together and referred to as STS) is greatly above chance. Areas VIP and LIP are above chance on only two out of four positions, hence the large standard deviations in their performance and their relatively poor performance overall (Figure 5 will discuss the details). We conclude that we have found a functioning fixation codebook in the STS, but not in VIP or LIP.

Next, we analyze the errors the fixation codebooks make. Our analysis determines a matrix that tabulates the percentage of trials in which position x was stimulated but the codebook decoded position y. Figure 4
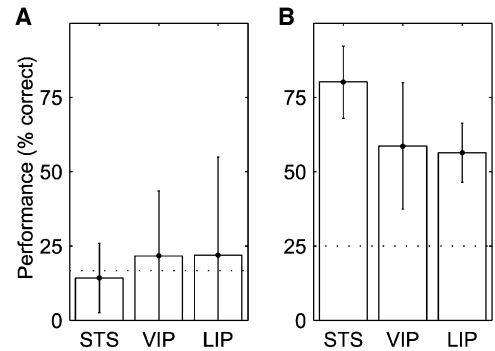

Figure 4. Fixation Codebook Performance

(A) Decoding in world coordinates.
(B) Decoding in retinal coordinates.
The dashed lines show chance performance (17% for six positions, 25% for four positions), error bars indicate the standard deviation over all positions.

only shows averages of the diagonal elements of this matrix. Figure 5 displays the complete matrix. Diagonal elements show the percentage of correct trials per position and as such gives a more detailed view of the extent to which the population is capable of encoding position in a particular coordinate system. Off-diagonal elements are trials in which the codebook gave the wrong answer. If a population actually encodes in eye-centered coordinates before a saccade but we decode it in world coordinates after a saccade, one would expect mislocalization errors the size of the saccade but in a direction against the saccade. This is clearly what is observed for both LIP and STS in Figure 5A. For instance, the world position +25 is erroneously decoded as world position +5: the difference between these positions is precisely the 20° leftward saccade. For VIP the results are more complicated; there is an indication of world coordinate encoding for some positions (+15), but other positions are clearly retinally encoded (+5).

Figure 5B shows the details of the performance on retinal position encoding. For VIP and LIP, the presence of off-diagonal elements in Figure 5B shows that the rates are not a good indicator of retinal position. In fact, the performance on only two out of four positions is significantly above chance. In the STS, however, the diagonal (correct) elements dominate and there is no particular bias to an off-diagonal element.

The analysis in Figure 5 contrasts retinal with world coordinates. It is known, however, that many areas in the posterior parietal cortex encode not only retinal position signals but also the position of the eye. Such multiplexing of information leads to a coordinate system intermediate between retinal and world coordinates. If a neuron has a strong eye position signal, the rates evoked by postflashes could be quite different from those evoked by preflashes, even when presented at the same retinal position. Hence, when tested in retinal coordinates, eye position signals could reduce the performance of a pre-codebook tested with postflashes. To test this possibility, we determined the performance of a codebook based on preflashes, tested with preflashes, and the performance of a codebook based on postflashes, tested with postflashes. Because these coding/
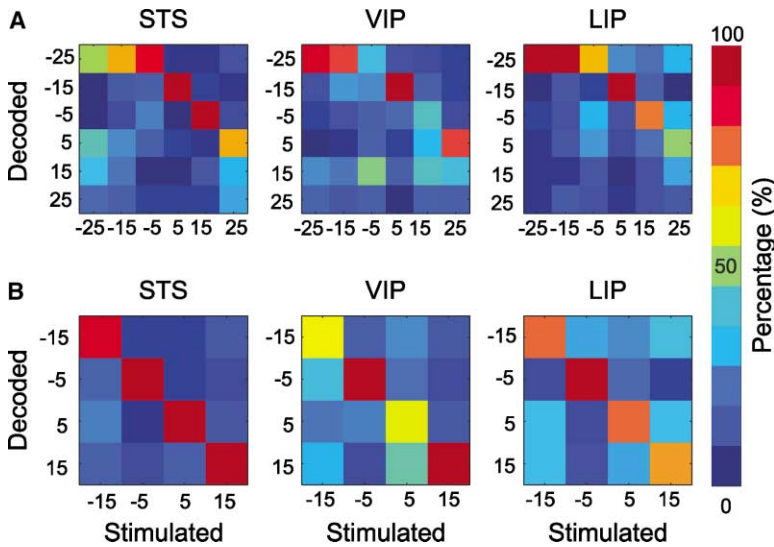
Figure 5. Fixation Codebook Performance and Errors

(A) World coordinates.
(B) Retinal coordinates.
The stimulated position is on the horizontal axis, the decoded position on the vertical axis. The color represents the percentage of trials. Diagonal elements are the correct decodings.

decoding procedures involve no change in eye position, they are immune to eye position effects. In each case, we randomly selected 75% of trials to create the codebook and used the remaining 25% to test the performance (crossvalidation; see Experimental Procedures). We compared the average performance of these same-interval codebooks with codebooks based on pre-flashes tested with postflashes. For a fair comparison the latter different-interval codebooks were also based on a random subset of 75% of the trials and tested with 25% of trials. This results in somewhat lower average level performance than that shown in Figure 4.

In the STS as well as LIP, the performance of the same-interval codebooks was approximately 10% better than the performance of the different-interval codebooks (STS, 65% improved to 75%; LIP, 44% improved to 53%). This difference, however, was not large enough to exclude the possibility that it is due to chance (p > 0.05). Hence, the effect of the eye position signals in these areas is either small enough for all cells to allow pure retinal-position encoding, or the effect is large in some cells but the effect is cancelled out at the population level (Bremmer et al., 1997).

In VIP the performance increased from 47% in the different-interval codebooks to 66% in the same-interval codebooks. This 19% increase is significantly above chance (p < 0.05). Hence, eye position signals play a more important role in area VIP than in the either LIP or the STS. This fits well with the finding that many cells in area VIP encode in a coordinate system that is intermediate between retinal and world centered (Duhamel et al., 1997). This also implies that the decoding performance in VIP could be improved if we knew just how each cell multiplexes retinal with eye position signals. For the cells at hand, however, we have no data that independently assess this and we are forced to use the next best coordinate system: retinal coordinates.

The analysis of fixation codebooks pools information over multiple cells and it is instructive to see the dependence of performance on population size. We sampled subpopulations of increasing size from our complete data set and repeated the above analysis to determine

the performance on retinal position encoding as a function of population size. Figure 6 shows that the performance of the STS population steadily increases from 46% for ten cells to the 80% of the complete data set of 125 cells. The performance of areas VIP and LIP shows a much shallower increase with the number of cells.

So far we have only considered stimuli that were temporally separated from saccade onset by at least 300 ms. We now address the question whether these populations of neurons encode the position of stimuli presented just before saccades. Figure 7A shows the average performance of codebooks based on preflashes on decoding periflashes in retinal coordinates. The performance is not above chance: this means that even the STS, which provides accurate position information during fixation, fails to provide accurate position information perisaccadically. This analysis was based on codebooks defined by preflashes. One could argue that the brain might switch to a different codebook specifically geared to perisaccadic position encoding. To investi-
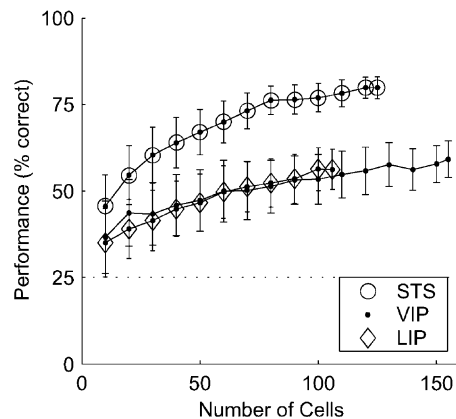

Figure 6. Population Size Effects

The average percentage of correctly decoded retinal positions as a function of the number of cells in the codebook. The dashed line shows chance performance. Error bars are the standard deviation of the performance over 50 populations of the same size.
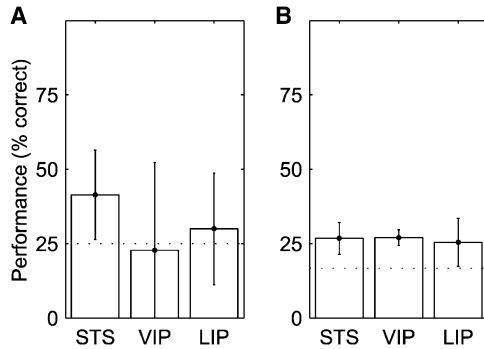
Figure 7. Fixation Codebook Performance for Perisaccadic Flashes

(A) Decoding based on a codebook defined by preflashes.
(B) Decoding based on a codebook defined by periflashes.
The dashed lines show chance performance (17% for six positions, 25% for four positions), error bars indicate the standard deviation over all positions.
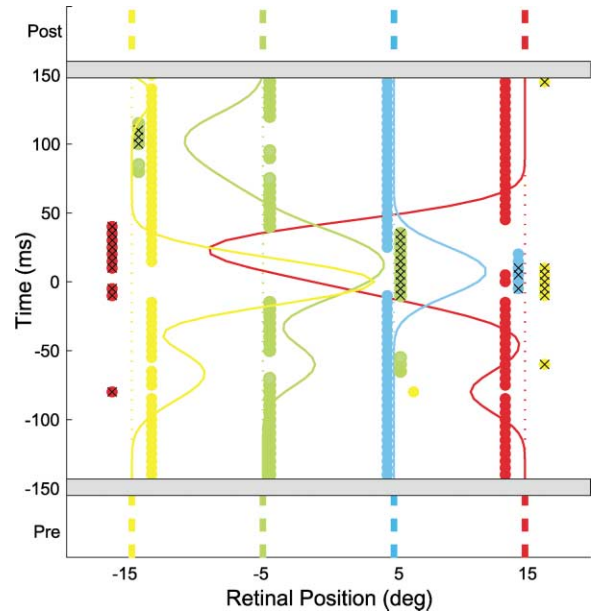

Figure 8. Perisaccadic Mislocalization in the STS

The dots are raw data points: each dot represents the decoded position at a give time. The color of the dot represents the stimulated position. To show multiple dots in a single decoded position, they straddle the actually decoded position. The crossed dots are significantly different from veridical (p < 0.05). The solid lines are smooth interpolations of these raw data points. The faint dotted curves indicate which position was actually stimulated. Time zero is saccade onset. The dashed curves show the decoded positions in the pre and post periods. Gray bars show where our analysis of the periepoch ends and the pre- and postepochs start.

gate this hypothesis, we determined a codebook based on flashes presented between 200 and 0 ms before saccade onset and, using a crossvalidation approach, decoded flashes from the same period (see Experimental Procedures). Figure 7B shows the average performance of these codebooks on all six retinal positions. Again, the performance is not significantly above chance. (This is also the case when position coding is restricted to the four parafoveal positions.)

For our VIP and LIP populations, we conclude that we have been unable to find a fixation codebook that gives adequate position information for all positions. This prevents us from using these populations in the further analysis. On the other hand, for the STS, we did find a codebook that consistently relates firing rate to retinal position during fixation. An ideal observer, or an area downstream from the STS, could use this codebook to interpret the firing rates of the STS in terms of stimulated retinal position. Figure 7A shows that perisaccadic position decoding based on area STS is much impaired, and moreover, Figure 7B shows that a downstream area can gain nothing by changing its codebook to one specifically geared to perisaccadic stimulation. It seems logical, therefore, to assume that such an area would keep on using the codebook that is known to work during fixation. The next section will investigate the errors of localization that such an area would make.

Perisaccadic Mislocalization

What happens if we decode the stimulus responses in the temporal vicinity of the onset of a saccade in terms of the STS codebook that works well during fixation? We constructed the optimal codebook based on all pre- and postsaccadic flashes and used this to decode the firing rates induced by perisaccadic flashes presented at various times with respect to saccade onset. The result of this analysis relates stimulated position to decoded position as a function of the time to the saccade. Figure 8 shows the results. In the pre and post periods, the figure shows the most frequently decoded position. For all actually stimulated positions this is the veridical position, which reinforces the results of Figure 4B. Start-

ing approximately 100 ms before the saccade, the representation of retinal position in the STS is strongly disturbed. Approximately 70 ms after saccade onset, the representation is accurate again. Many psychophysical studies describe precisely such temporal dynamics of mislocalization (for reviews, see Ross et al., 2001; Schlag and Schlag-Rey, 2002). Moreover, the average magnitude of the mislocalization in the STS is, during the saccade, on the order of 10°, which corresponds to half the saccade amplitude. This too is similar to what has been found in psychophysical studies in humans (Ross et al., 2001; Schlag and Schlag-Rey, 2002) and monkeys (Dassonville et al., 1992). Crucially, however, the mislocalization in the STS is entirely due to *retinal* position errors; nearly all psychophysical studies, on the other hand, report and interpret their data in terms of world coordinates. We will return to this issue in the Discussion.

Ideally, one would also like to compare the details of the mislocalization in the STS with those found in human psychophysics. Before we do this, however, some warnings are appropriate. First, our basic analysis is restricted to a much coarser resolution (10°) than most psychophysical studies: if a flash is mislocalized to position 15 in our analysis, this only means that this was a better estimate than the other three possible positions. The interpolated data curves in Figure 8 try to get around this restriction, but can only do so at the cost of further assumptions on the representation of position (see Ex-

perimental Procedures). Second, the psychophysical literature mostly discusses mislocalization in terms of flashes at particular positions in the world. It has been shown early on, however, that the retinal position of a flash is in fact the strongest determinant of its mislocalization (Bischof and Kramer, 1968; O'Regan, 1984) Later studies, however, often confounded retinal position with the time of the flash relative to saccade onset.

Figure 8 shows mislocalization in the direction of the saccade for retinal position 15 and mislocalization against the direction of the saccade for the other three positions. Such position-dependent changes in the direction of mislocalization have also been found psychophysically. The precise relationship between position and mislocalization, however, is a matter of ongoing debate. This relationship is strongly dependent on the precise experimental setup (Lappe et al., 2000) and, as is clear from the raw data in many studies, varies among subjects. Nevertheless, there are many clear examples of mislocalization very similar to what we find in the STS of these two monkeys in the studies of Bischof and Kramer (1968), O'Regan (1984), and Honda (1993). In the study of Ross et al. (1997) and Morrone et al. (1997), all flashes beyond the saccade target are mislocalized against the direction of the saccade, while flashes between fixation point and target are mislocalized in the direction of the saccade. This fits with the opposed mislocalizations of the yellow, green, and red curves in Figure 8, but not with the blue curve. Note however, that other studies (Honda, 1993) have found shifts against the direction of the saccade at this position.

Many psychophysical experiments find mislocalization in the perisaccadic period from 50 ms before the saccade to saccade onset. Because the eyes have not yet started to move, mislocalization in this period cannot be attributed to the retinal motion signals that arise for stimuli flashed after the eye starts to move. Figure 8 confirms that perisaccadic mislocalization in the STS is also found for flashes presented before saccade onset; the average mislocalization in the 50 ms before saccade onset is 5°. The retinal image motion of the flashed bars cannot explain such presaccadic mislocalization.

Saccade and Target Onset Responses

When the eye starts to move, many neural signals are likely to change: signals related either to saccade goal or to eye position, and also visual signals. The latter are present in any experiment that is not in absolute darkness. In our experiment they could include movement of the background and the saccade target. We use the control trials, in which the monkey made a saccade but no bars were flashed, to determine the (joint) influence of these changes on the firing of the neurons.

We analyzed responses in the same time window used in the analysis above: from 50 to 250 ms after an event. First, we determined the fraction of cells whose response to the appearance of the target was more than three standard deviations above or below the baseline firing rate. Only 6% of STS cells showed such a significant target response. This confirms that the saccade target itself is not a potent visual stimulus. We then determined the fraction of cells that responded significantly to saccade onset. Note that it is not possible in

our paradigm to distinguish pure saccade responses from saccade-related but visual responses (to the moving target or background). Twenty-four percent of STS cells responded significantly to saccade onset; in fact, their saccade response was enhanced by 20% compared to the flashed bar response. The average change in firing rate for all cells around saccade onset, however, is only 30% of the change in firing rate evoked by stimulation with the bright flashes. Hence, the response to the target alone and the population response to saccade onset (which includes the onset of motion of the target on the retina), are both small and it seems unlikely that they cause the mislocalization.

A stronger test of the influence of the target and saccade onset can be performed within our decoding framework. First, we decoded the firing rates evoked by the onset of the saccade target in terms of the codebook based on presaccadic flashes. This decoding is not accurate; there is a bias to the most peripheral positions. This merely means that the codebook defined by bright, large flashed bars does not transfer to the small red dot. We then decoded the response to saccade onset. If saccade onset events (including the target's motion) induced the mislocalization by themselves, one would expect that the decoding of these saccade onset rates is different from that of the target onset rates. This is not the case; the bias found for target responses remains. We conclude that even though there are visual and nonvisual events that change the firing in STS neurons perisaccadically in the absence of flashed bars, these changes by themselves are not enough to fully explain the mislocalization discussed in the previous section. An interaction of the strong visual stimulation by the flashes and the perisaccadic signal changes appears to be necessary.

Discussion

Our results show that areas MT and MST are capable of encoding the retinal position of large flashed bars during fixation but are thoroughly confused when these flashes are presented perisaccadically. For an ideal observer who interprets MT and MST output in terms of the firing rates that are evoked during fixation, this confusion leads to perisaccadic mislocalization effects that are similar to those found in psychophysical experiments with human subjects. This mislocalization in the STS, however, is found in a retinal frame of reference. This contradicts the usual interpretation of the psychophysical data that mislocalization occurs because veridical information in a retinal reference frame is translated into a world reference frame with the aid of an erroneous eye position signal. Instead, our data support the view that retinal position encoding itself is inaccurate in the temporal vicinity of a saccade.

The Where Pathway

For areas VIP and LIP, we could not find a reliable codebook for either retinal or world position. This does not imply that these areas are incapable of encoding position, just that our sample of cells cannot do this based on their mean firing rate and in pure retinal or pure world coordinates. In fact, for VIP we showed that a mixed

coordinate system could do better than a purely retinal coordinate system. This is clearly related to the head-centered, eye-centered, and intermediate receptive fields found in VIP (Duhamel et al., 1997). Without an independent assessment of a cell's reference frame, however, the decoding analysis cannot be done. It remains entirely possible, therefore, that in the correct frame of reference, these cells do encode position. Interestingly, given the fact that VIP receives a considerable part of its input from area MT (Maunsell and Van Essen, 1983), it would inherit the retinal mislocalization errors present in this input.

Our analysis of LIP showed that using a mixed frame of reference could not improve the performance on position coding. There are many possible explanations for this. The stimuli we used may not have been optimal for LIP, or their behavioral irrelevance may have contributed to poor encoding (Gottlieb et al., 1998). Moreover, because we searched for LIP cells by looking for saccade-related activity, our population of LIP cells may well be more closely involved in the planning of saccades than the spatial representation of the visual environment.

In MT and MST we found that position was encoded in retinal coordinates only. This is surprising because these areas have access to both retinal information and extraretinal eye position information (Bremmer et al., 1997). In principle, this combination of eye position information with retinal position information should be enough to provide head-centered positions (Zipser and Andersen, 1988). Eye position signals in MT and MST, however, have been determined in the absence of visual stimulation (Bremmer et al., 1997). Hence, our finding that MT and MST do not combine eye position information with retinal position information indicates that the retinal signal dominates and masks the eye position signal. This serves as a caveat that the presence of eye position signals in itself does not automatically guarantee the ability to encode in head coordinates.

Perisaccadic Mislocalization

Previous accounts of perisaccadic mislocalization have mainly focused on the recalibration of coordinate systems required to link pre- and postsaccadic retinal position to a position in the world. Such a recalibration involves the judicial combination of an eye position signal with a signal representing the retinal position of objects. The analysis of mislocalization in these accounts assumes that the retinal position signal is basically accurate, but that it is combined with an inaccurate (or damped) eye position signal (Dassonville et al., 1992). Our findings, however, show that the mislocalization can already be found in the retinal coordinate system of areas MT and MST. A downstream area that encodes in world coordinates but relies on MT and MST for its retinal (position) information would inherit this mislocalization.

Hence, in this view, perisaccadic mislocalization is not the result of an inaccurate eye-position signal, but rather the result of a disturbance of retinal position encoding. Our data do not address the question of what disturbs areas MT and MST perisaccadically, but we believe that there are two main possibilities, and we discuss them in turn.

Retinal Effects

As our data show that mislocalization can already be found in retinal coordinates, it is tempting to conclude that mislocalization must be due to retinal, not extraretinal, signals. During saccades, the visual environment sweeps over the retina at high speeds and this strong visual disturbance may interfere with processing and cause mislocalization. In fact, it has been shown that some errors of localization can also be found without saccades: when a sudden rapid background movement follows a flashed stimulus, the stimulus is mislocalized against the direction of the background motion (Mackay, 1970; Morrone et al., 1997; O'Regan, 1984). Other findings, however, speak against a purely retinal interpretation of mislocalization. For instance, mislocalization against the direction of simulated saccades has not been found with sudden background movements (Morrone et al., 1997).

Our data support the view that retinal factors alone cannot fully explain mislocalization. First, we found mislocalization even before the eyes had started to move, hence without the possible influence of retinal smear of the bars. Second, we could not find an equivalent disturbance in the firing of the cells when saccades were made in the absence of flashed bars (see Saccade and Target Onset Responses). Hence, whatever retinal stimulation there is besides the flashed bars, it in itself is not large enough to disturb MT and MST. It remains possible, though, that a nonlinear interaction of the saccade-induced visual stimulation (by the background, the saccade target, or other visual references) with the stimulation by a flashed bar causes the effect. If the visual system localizes objects relative to visual references in the background (Krekelberg and Lappe, 2000, 2001; Lappe et al., 2000; O'Regan, 1984), such an interaction may be expected.

Extraretinal Effects

In our opinion, there are two possible sources of extraretinal effects that could disturb position encoding in MT and MST perisaccadically. Even though these sources are conceptually very different, they may turn out to be closely related.

First, it is known that areas MT and MST carry information on the direction of impending saccades (Recanzone and Wurtz, 1999), as well as the current position of the eye (Bremmer et al., 1997). By their very nature, these signals change perisaccadically. Therefore, one possible interpretation of our findings is that these signals interfere with the encoding of retinal position. It should be stressed, however, that this interference is not the computation that determines world position from a combination of an eye position signal and the retinal position signal. If this computation took place in MT and MST, our analysis shown in Figure 4 would have found an encoding of world position with an accuracy similar to that of the encoding of retinal position.

Second, the disturbance may be related to saccadic suppression. This is the phenomenon that during saccades, the sensitivity of the visual system is reduced (Holt, 1903; Ross et al., 2001). This suppression is particularly strong for the kind of stimuli processed in the dorsal stream (Burr et al., 1994) and has been interpreted

as a mechanism to avoid processing the retinal smear induced by rapid eye movements. The time course of suppression is very similar to that of mislocalization (Diamond et al., 2000), and Burr et al. (1994) have suggested that, to achieve suppression, the dorsal stream partially shuts down around saccades. Thiele et al. (2002) recently linked this partial shut down to changes in the response of areas MT and MST. In this view, perisaccadic mislocalization comes about because the system tries to hide the retinal motion induced by saccades. By doing so, however, it interferes with other tasks that MT and MST may be involved in, such as encoding the position of objects on the retina.

Experimental Procedures

Recording

The monkey sat in a primate chair and made 20° saccades to a visual target for liquid reward. We recorded extracellularly from the anterior and posterior bank of the superior temporal sulcus and from the intraparietal sulcus from four hemispheres in two monkeys. Eye position was sampled at 500 Hz with a scleral eye coil system (Skalar, Delft) with an accuracy of 1 min arc. Saccade onset was determined offline by finding the sample after which the eye velocity exceeds 5% of the maximum velocity for at least three consecutive samples. Saccade latencies were restricted to lie in the range 80–300 ms. Trials in which no saccade could be detected that satisfied this criterion were discarded (5% of trials). Average saccade latency was 201 ± 11 ms, average duration 44 ± 6 ms.

Animal treatment, housing, surgical, and recording procedures were in accordance with EU guidelines on the use of animals in research (European Communities Council Directive 86/609/ECC). Details are discussed elsewhere (Bremmer et al., 1997). The analysis reported in this paper used 125 cells from MT and MST, defined as the direction-selective cells on the posterior and anterior bank of the STS, respectively. Responses to flashed bars from putative MT and MST cells were very similar, and we decided to treat these as a single population. We refer to this population as the STS population. In the intraparietal sulcus, we recorded from 158 cells with a clear preference for moving stimuli and no saccade-related activity. We identify these with area VIP. Finally, we recorded from 106 cells in the intraparietal sulcus with saccade-related activity. We will refer to these as neurons from area LIP. In the first animal the anatomical location of the cells has been confirmed histologically and shown to agree well with our physiological definition. All analyses reported here use these exact same cells, and we had no selection criteria for inclusion beyond the presence of a visual response. On average, we recorded 25 repetitions per bar position in each of the pre-, peri-, and postepochs. Hemispheres and saccade direction were mirror reversed to normalize all data to "leftward" saccades and recordings in a "left" hemisphere.

Stimulus

The luminance of the bars was 10 cd/m^2, duration 8 ms, and they were projected on a screen 48 cm in front of the monkey. The centers of the 10° wide, nonoverlapping bars always occupied the same six world positions: −25, −15, −5, 5, 15, and 25° from the midline. These positions were not adjusted to the cells' receptive fields. The fixation point (at 10°) and target (at −10°, each on the horizontal meridian) were identified by a small red dot (0.5° diameter). The fixation point disappeared when the target appeared.

Decoding

The recordings allowed us to estimate the conditional probability of observing an average firing rate (r) after a flash was presented at position x: P(r|x). For each cell we used Bayes' rule to construct a lookup table that inverts this relationship and relates the observed firing rate (r) to the posterior probability P that a flash was presented at position x:

$$P(x|r) = P(r|x) \times P(x)/P(r).$$

From P(x|r) we can determine the stimulus that is most likely to have caused an observed rate; this is called "decoding."

The prior distribution of the stimuli P(x) is the relative frequency of occurrence of a stimulus, which was a flat 1/6 in world coordinates, or 1/4 for the four retinal positions. Given the limited recording time available for a cell in an awake behaving monkey, estimating the complete distributions P(r|x) and P(r) is not feasible. We parameterized the conditional probability P(r|x) and the prior P(r) in terms of more easily estimable quantities such as the mean and standard deviation of the firing rate. Parameterization is necessary because we test the codebook with a different subset of trials from those used to setup the codebook; not all rates recorded in the test set need to have been recorded in the codebook set. Moreover, parameterization reduces the influence of outliers in the recorded firing rates.

To estimate the prior distribution of the firing rates, we determined the distribution of all observed firing rates in 200 ms windows in all cells and all trials. Importantly, this includes periods when a stimulus is in the cell's classical RF, periods when there is a stimulus outside the RF, periods when there is no stimulus present at all, as well as periods when the monkey makes a saccade. As such this is a description of the firing rates a cell may have during a typical day. For our data set, this distribution was well described by an exponential distribution with a decay parameter given by the standard deviation in the firing rate of the cells. We used this to model the prior distribution of the firing rates of each cell as P(r) ∼1/σ × exp(−r/σ), where σ is the standard deviation of the firing rates of that cell. Using such a prior effectively puts less weight on those trials in which few spikes were recorded. Given the noisy nature of cells' responses and electrophysiological recordings, this seems reasonable.

To determine the parameterization of the conditional probability P(r|x), we used an optimality criterion: we explored a number of possible parameterizations and chose the parameterization that led to the best performance on the encoding of position during fixation. We tested Gaussian, exponential, Poisson, and modified Poisson and Gaussian parameterizations. The modified Poisson distribution consistently (i.e., when decoding preflashes, postflashes in all three areas) led to the best results. This modified Poisson is a Poisson distribution based on the mean firing rate of the cell, but the probability of zero spikes is given by the experimental probability. This modification reflects the observation that we usually recorded more trials with zero spikes than predicted by a pure Poisson law. Figure 3A shows an example of an observed distribution of firing rates and the modified Poisson function we used to describe this experimental distribution.

The time window to determine the spike rate was determined by evaluating the encoding of position information for a range of windows. We settled on a window from 50 to 250 ms after stimulus onset; this window includes most visual responses that are typical for these areas (Schmolesky et al., 1998), and it led to the most accurate encoding of position information for pre- and postsaccadic flashes. We used the same window for all analyses.

To evaluate the fixation codebooks, we determined a codebook based on the preflashes and tested with postflashes. Assuming that cells provide independent evidence, we multiplied the posterior probabilities of the individual codebooks and determined which position was most likely to have been stimulated according to the population of cells. In a bootstrap validation approach, we repeated this decoding 1000 times, each time resampling the population of N cells and recalculating the population codebook. We used standard bootstrap 95% confidence limits to test statistical significance of bootstrap estimates.

We used crossvalidation to test the performance of a codebook for one particular epoch with flashes from that same epoch. A random selection of 75% of the trials was used to setup the codebook, and the remaining 25% of trials was used to test the performance. This was repeated in bootstrap fashion, each time resampling a new, randomly selected, subset of 75% of the trials. Performance measures were averaged over the bootstrap sets. The inevitably smaller number of test flashes in crossvalidation causes an increase in the variance of the bootstrap estimates. To compensate for this, we increased the number of bootstrap sets to 4000.

Perisaccadic Decoding

To decode the perisaccadic rates, we first constructed a codebook for each cell based on all its pre- and postsaccadic trials. We then

decoded the rates evoked by perisaccadic flashes. The retinal position of flashes during the saccade was determined by linearly interpolating the eye position signal between the fixation point and the target for each trial. To sum evidence over a large enough set of trials, we considered all flashes within a 25 ms window to be presented "at the same time." This 25 ms window was shifted from 150 ms before saccade onset to 150 ms after saccade onset in 5 ms steps. Each step results in a single raw data point in Figure 8. Here too we used a 1000 sample bootstrap resampling method to avoid sampling bias. Importantly, the average decoded position was not determined by a geometric mean over the bootstrap sets, as this would confound noise with a bias toward the fovea. Instead, we determined the most frequently decoded position over all bootstrap sets. This analysis results in the raw data points of Figure 8. To look at the position representation with a higher resolution than we actually sampled, we filtered these data points with a 25 ms Gaussian filter. This interpolation implicitly makes two assumptions. First, abrupt changes in the decoded position are due to our coarse sampling of the visual field, and the true nature of the representation of position is smooth. Second, it effectively assumes retinotopy in the representation of position: in other words, a vote for $+15°$ becomes a partial vote for $+10$, but not for -15.

References

Bischof, N., and Kramer, E. (1968). Untersuchungen und Überlegungen zur Richtungswahrnehmung bei willkürlichen sakkadischen Augenbewegungen. Psychol. Res. *32*, 185–218.

Bremmer, F., Ilg, U.J., Thiele, A., Distler, C., and Hoffmann, K.P. (1997). Eye position effects in monkey cortex. I. Visual and pursuit-related activity in extrastriate areas MT and MST. J. Neurophysiol. *77*, 944–961.

Burr, D.C., Morrone, M.C., and Ross, J. (1994). Selective suppression of the magnocellular visual pathway during saccadic eye movements. Nature *371*, 511–513.

Dassonville, P., Schlag, J., and Schlag-Rey, M. (1992). Oculomotor localization relies on a damped representation of saccadic eye displacement in human and nonhuman primates. Vis. Neurosci. *9*, 261–269.

Diamond, M.R., Ross, J., and Morrone, M.C. (2000). Extraretinal control of saccadic suppression. J. Neurosci. *20*, 3449–3455.

Duhamel, J.R., Colby, C.L., and Goldberg, M.E. (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. Science *255*, 90–92.

Duhamel, J.-R., Bremmer, F., Ben Hamed, S., and Graf, W. (1997). Spatial invariance of visual receptive fields in parietal cortex neurons. Nature *389*, 845–848.

Gottlieb, J.P., Kusunoki, M., and Goldberg, M.E. (1998). The representation of visual salience in monkey parietal cortex. Nature *391*, 481–484.

Holt, E.B. (1903). Eye movements and central anaesthesia. Psychol. Rev. *4*, 3–45.

Honda, H. (1991). The time courses of visual mislocalization and of extraretinal eye position signals at the time of vertical saccades. Vision Res. *31*, 1915–1921.

Honda, H. (1993). Saccade-contingent displacement of the apparent position of visual stimuli flashed on a dimly illuminated structured background. Vision Res. *33*, 709–716.

Krekelberg, B., and Lappe, M. (2000). A model of the perceived relative positions of moving objects based upon a slow averaging process. Vision Res. *40*, 201–215.

Krekelberg, B., and Lappe, M. (2001). Neuronal latencies and the position of moving objects. Trends Neurosci. *24*, 335–339.

Kubischik, M., and Bremmer, F. (1999). Peri-saccadic space representation in monkey inferior parietal cortex. Soc. Neurosci. *25*, 1164.

Lappe, M., Awater, H., and Krekelberg, B. (2000). Postsaccadic visual references generate presaccadic compression of space. Nature *403*, 892–895.

Mackay, D.M. (1970). Mislocation of test flashes during saccadic image displacements. Nature *227*, 731–733.

Mateeff, S. (1978). Saccadic eye movements and localization of visual stimuli. Percept. Psychophys. *24*, 215–224.

Matin, L., and Pearce, D.G. (1965). Visual perception of direction for stimuli flashed during voluntary saccadic eye movements. Science *148*, 1485–1488.

Maunsell, J.H.R., and Van Essen, D.C. (1983). The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. J. Neurosci. *3*, 2563–2586.

Mishkin, M., Ungerleider, L., and Macko, K. (1983). Object vision and spatial vision: two central pathways. Trends Neurosci. *6*, 414–417.

Morrone, M.C., Ross, J., and Burr, D.C. (1997). Apparent position of visual targets during real and simulated saccadic eye movements. J. Neurosci. *17*, 7941–7953.

O'Regan, J.K. (1984). Retinal versus extraretinal influences in flash localization during saccadic eye movements in the presence of a visible background. Percept. Psychophys. *36*, 1–14.

Recanzone, G.H., and Wurtz, R.H. (1999). Shift in smooth pursuit initiation and MT and MST neuronal activity under different stimulus conditions. J. Neurophysiol. *82*, 1710–1727.

Ross, J., Morrone, M.C., and Burr, D.C. (1997). Compression of visual space before saccades. Nature *386*, 598–601.

Ross, J., Morrone, M.C., Goldberg, M.E., and Burr, D.C. (2001). Changes in visual perception at the time of saccades. Trends Neurosci. *24*, 113–121.

Schlag, J., and Schlag-Rey, M. (2002). Through the eye, slowly: delays and localization errors in the visual system. Nat. Rev. Neurosci. *3*, 191–200.

Schmolesky, M.T., Wang, Y., Hanes, D.P., Thompson, K.G., Leutgeb, S., Schall, J.D., and Leventhal, A.G. (1998). Signal timing across the macaque visual system. J. Physiol. *79*, 3272–3278.

Snyder, L.H. (2000). Coordinate transformations for eye and arm movements in the brain. Curr. Opin. Neurobiol. *10*, 747–754.

Thiele, A., Henning, P., Kubischik, M., and Hoffmann, K.P. (2002). Neural mechanisms of saccadic suppression. Science *295*, 2460–2462.

Tolias, A.S., Moore, T., Smirnakis, S.M., Tehovnik, E.J., Siapas, A.G., and Schiller, P.H. (2001). Eye movements modulate visual receptive fields of V4 neurons. Neuron *29*, 757–767.

Umeno, M.M., and Goldberg, M.E. (1997). Spatial processing in the monkey frontal eye field. I. Predictive visual responses. J. Neurophysiol. *78*, 1373–1383.

Walker, M.F., Fitzgibbon, E.J., and Goldberg, M.E. (1995). Neurons in the monkey superior colliculus predict the visual result of impending saccadic eye movements. J. Neurophysiol. *73*, 1988–2003.

Zipser, D., and Andersen, R.A. (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. Nature *331*, 679–684.