Exp Brain Res (1999) 127:151–161

RESEARCH ARTICLE

Kathryn J. Jeffery · John M. O'Keefe

Learned interaction of visual and idiothetic cues in the control of place field orientation

Received: 23 November 1998 / Accepted: 13 March 1999

Abstract In a symmetrical environment (like a square box) hippocampal place cells use a mixture of visual and idiothetic (movement) information to tell them which way the environment is oriented. The present experiment tested the hypothesis that if the visual landmarks were mobile, place cells would learn to disregard these and rely on idiothetic cues instead. Place cells were recorded in a square box surrounded by circular black curtains. A cue card hung on the curtain behind one of the walls to break the fourfold symmetry. The relative influence of this card on the location of place fields was assessed each day by confining the rat on a rotating platter underneath an opaque cover, and then rotating the card and the platter by different amounts, to see whether subsequently recorded place fields had rotated with the card or with the rat. For some rats, these trials had been preceded by trials in which the card had been visibly moved from trial to trial, so that the rats had seen that it was mobile. Other rats received no prior visual information that the card was mobile. In the rats that had previously seen the card move, place fields initially rotated with the card but by the end of five sessions usually rotated with the rat instead. For rats that had never seen the card move, place fields always followed the card. Thus, the cells were able to "learn" that their preferred directional input, the card, was unreliable. A third group of rats, who were covered only for 30 s while the card was moved, showed mixed behaviour, suggesting a degradation of the idiothetic trace with time.

Key words Hippocampus · Place cells · Direction · Idiothetic cues · Visual cues · Landmark stability

K.J. Jeffery (✉) · J.M. O'Keefe
Department of Anatomy and Developmental Biology,
University College London, Gower Street,
London WC1E 6BT, UK
e-mail: kate@maze.ucl.ac.uk
Tel.: +44-171419-3393, Fax: +44-171391-1306

Introduction

Animals use, among other things, a mixture of visual and self-motion (idiothetic) information to determine their whereabouts. The neuronal basis of the "knowledge" of location appears to reside mainly in the hippocampus (O'Keefe and Nadel 1978). The present experiment set out to test the hypothesis that one of the visual inputs to the spatial representation, a directional input (telling the representation which way the environment is oriented), is not fixed but can be learned about.

In the rat, the spatial representation is instantiated in the firing of place cells (O'Keefe and Dostrovsky 1971), which are pyramidal cells in the hippocampus. Place cells fire whenever the animal is in a particular location in its environment, leading to the hypothesis that the hippocampus constructs a map of the environment which the animal can use to navigate (O'Keefe and Nadel 1978). The preferred firing location differs for each cell, and is called the cell's place field. The strongest determinant of whether a place cell fires is the distance of the rat from some (but not all) of the walls of the environment (O'Keefe and Burgess 1996). However, when the walls are physically identical, the cells also need information about direction to know which wall is which. It is this directional information that is the focus of the present experiment.

How does a place cell distinguish, for example, a north wall from a south wall? Previous experiments have shown that this directional information is provided by extra visual features of the environment outside of the immediate recording box. When these are unavailable, the cells use idiothetic cues (Sharp et al. 1995; Jeffery et al. 1997). These are cues, such as vestibular information (Stackman and Taube 1997) or motor efference signals, that are generated by the rat's own movements (Taube and Burton 1995). They can be used to update the place representation appropriately as the animal moves through the environment, even when the animal is temporarily deprived of external spatial information (Barlow 1964), such as in darkness (Markus et al. 1994; Quirk et

al. 1990). This periodic updating is known as dead reckoning or (more commonly) path integration, and it depends on continuous information about the animal's current position, current orientation, and both angular and linear motion.

Neither visual information nor path integration by itself is sufficient to tell place cells where to fire. Path integration is a cumulative process, and so any mistakes it makes also tend to accumulate over time. Visual information, on the other hand, depends on the stability of the visual environment. If the cues move around, then any aspects of the spatial representation that are anchored to them will also move around and so be useless to an animal that is trying to navigate (Biegler and Morris 1996). One way around these restrictions would be for visual and idiothetic cues to correct each other. For example, visual information could be used as it becomes available, to correct errors in the path integration process. In turn, idiothetic cues could be used as a reference to tell the spatial system whether a visual cue moved or not. If it did, the system could then learn to disregard it. It is the second of these two possibilities that concerns us here.

The hypothesis that idiothetic cues could be used to "ground" the spatial representation, and so reveal whether visual cues are mobile, was put forward by Knierim et al. (1995). They suggested that a rat enters a new environment with an internal, idiothetically based sense of direction, which is used as a baseline against which to test whether objects in the environment are stable enough to be used as landmarks. The present experiment set out explicitly to test this hypothesis. Visual and idiothetic directional information was placed in conflict to see which predominated in influencing the orientation of place fields. It was predicted that because of the repeated mismatch, the visual cue would fail to gain control over the orientation of place fields. Somewhat surprisingly, however, we found that visual information always predominated, unless the rats had prior experience that the cue card was unstable. An additional experiment suggested that this was because the rats were deprived of visual information for some time, while the conflict was being introduced, leading to a weakening of the idiothetic trace.

A preliminary subset of these data has been reported previously (Jeffery 1998).

Materials and methods

Subjects

Male Lister hooded rats (330–400 g) were housed singly in Perspex cages and maintained on a 12:12 h light:dark schedule with lights off at 3 p.m. Each rat was given sufficient food to maintain 90% of its free-feeding weight and allowed unlimited access to water.

Single unit recording

Prior to the start of the experiment the rats underwent surgical implantation of moveable microelectrodes with which to record multiple single neurons. The rats were anaesthetised using isoflurane and nitrous oxide and given an i.m. injection of buprenorphine (45 µg) for intra- and postoperative analgesia, and an s.c. injection of enrofloxacin (2.5 mg) as a prophylactic antibiotic. A 2-mm-diameter hole was drilled in the skull overlying the right dorsal hippocampus with a trephine bit, and two four-wire electrodes (tetrodes) were implanted into the overlying neocortex (bregma: –3.8 mm AP, 2.2 mm MIL, the deeper tetrode positioned 1.5 mm below brain surface). The tetrodes were separated by 300–500 µm so that while one tetrode was in a cell layer recording hippocampal cells, the other was in a cell-free zone acting as a reference. Each tetrode was made from four twisted strands of 17- or 25-µm-diameter HM-L-coated platinum-iridium wire (California Fine Wire) and the two tetrodes were held by a cannula which was attached to a microdrive, allowing the electrodes to be advanced through the brain in small steps. The assembly was held in place by means of jeweller's screws fixed to the skull, and dental acrylic. One of the screws was soldered to a gold Amphenol pin to enable the rat to be electrically grounded.

Recordings began 1 week after surgery. Each rat was connected to the recording equipment via lightweight hearing-aid wires and a socket which fitted onto the microdrive plug. The potentials recorded on each of the eight electrodes were passed through RC-coupled, unity-gain operational amplifiers mounted on the rat's head, and led to custom-built equipment (Gignomai Ltd.), where the signal was amplified (20000–40000 times) and bandpass filtered (500 Hz–9 kHz). Each of the four wires of one tetrode was recorded differentially with respect to one of the wires of the other.

Screening for hippocampal cells took place in a different room from the one in which the experiment was to be conducted, so as to minimise learning about the recording environment. The tetrodes were advanced by up to 200 µm daily in 50- to 100-µm steps, until hippocampal ripples appeared. Advancement then proceeded in 25-µm steps until complex spikes appeared. At this point, the rat was moved into the experimental room (see below) and the experiments begun. When place cells had been isolated their activity was recorded in trials lasting 4 min, while the rat chased rice grains in the apparatus described below. Unit activity was captured by monitoring each channel at 20-µs intervals and sampling 50 points/channel whenever the signal from any of the four electrodes exceeded a given threshold (a presumptive spike). Each spike event was stamped with the time since the start of the recording and the location of the animal. The data were stored on a hard disk and later transferred to a Sun Ultra workstation for analysis.

During the unit recording, the position of the rat was monitored by a video camera mounted directly above the apparatus. The location of the rat in the camera viewing area was converted into *x*–*y* coordinates by a TV-tracking system (Gignomai Ltd., UK), which detected a small DC light mounted on the recording cable near the rat's head. Every 20 ms the position of the rat was stored along with the unit data so that the whereabouts of the rat during activity of a given cell could subsequently be determined.

Apparatus

The experimental room was approximately 3×6 m in size, with a curtained-off area 2 m in diameter (Fig. 1) at one end. In the centre of this curtained area was a square box with sides 60 cm long and 25 cm high raised 30 cm from the floor. The long black curtains were completely closed and the main room lights were switched off so that external room cues were minimised. The centre of the box floor consisted of a circular platter 30 cm in diameter (i.e. half the width of the box) which could rotate independently of the box. This platter was driven by a motorised turntable situated just beneath the box, and could be turned at a slow angular velocity (0.13 rpm – intended to be undetectable by the rat). Thus, when the rat was confined to this part of the floor (see below), it could be rotated with respect to all other environmental cues except the platter itself.

Because the curtains provided no visual cues to indicate the orientation of the box, the environment possessed fourfold rota-

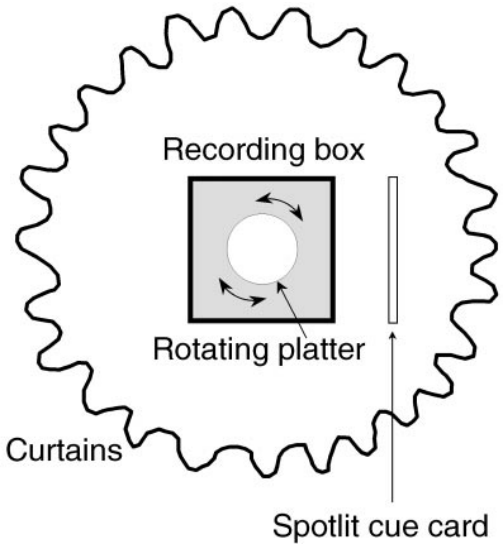

Fig. 1 The experimental apparatus. The black curtains formed an enclosure 2 m in diameter and 2 m high, in the centre of which a square box (60×60 cm) rested on 30-cm-high supports. The circular centrepiece of the floor of the box was attached to a rotating motor and could be turned at 0.13 rpm in either direction

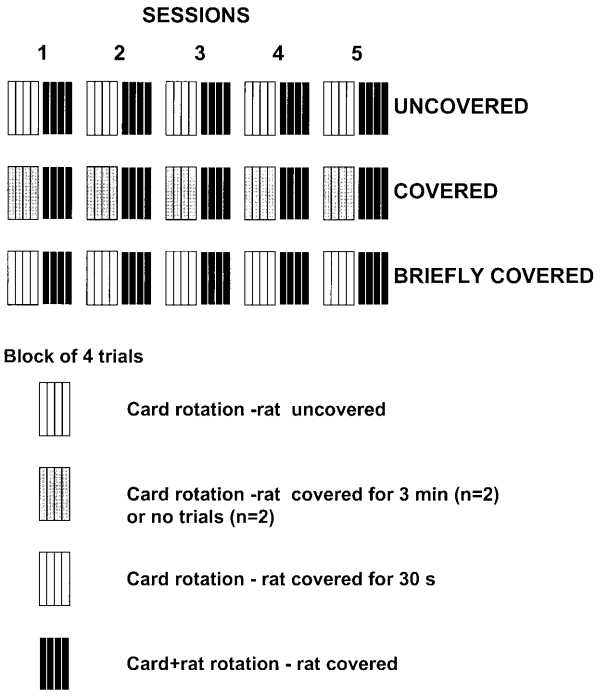

Fig. 2 The recording protocol. Each session began with a baseline recording trial (not shown). The latter half of each recording session consisted, for all rats, of four "conflict" trials during which the rat and the card were rotated by different multiples of 90°. The groups differed in their treatment during the four preceding training trials. Group Uncovered saw the card moved to a new location prior to the 4 min of place cell recording. Two rats from group Covered were confined under an opaque cover while the card was moved, while the other two rats received no training trials at all. Either way, these rats never saw the card move. Group Briefly Covered were covered for only 30 s while the card was moved. Five sessions were run altogether

tional symmetry. To break this symmetry, a large white cue card (1 m high × 76 cm wide) was suspended from the curtain rail just in front of the curtains, and behind one of the four walls of the box. Its position varied from trial to trial, and the principal variable of the experiment described below was whether its position was changed visibly (with the rat roaming the box freely and able to see the card move) or invisibly (with the rat confined by an opaque cover to the rotating platter).

Recording procedure

Screening for cells began at least 1 week after surgery, and the main part of the experiment usually began a few days later, when hippocampal place cells had been identified. At the start of each recording day, the rat was carried to the experimental room in an opaque box. Once inside the curtained arena the rat was lifted out, connected to the recording apparatus and placed in the square box. This was always done in the same part of the arena each day. At this time, the cue card was always situated in the east position, so that if the rat had been able to remain oriented on its way from the housing room, both the cue card and the start position would appear to be in a constant location at the start of the day. The experimenter always entered and left the environment through the same entrance. Because the rat's sense of direction was being manipulated throughout the experiment, this entrance would have appeared to the rat to vary randomly between north, south, east and west.

Since the purpose of the experiment was to investigate the relative influences of the card and the idiothetic cues on place fields, rotations of either the card alone, or both the card and the rat, were made as follows. The experimenter rotated the card by unhooking the card from its position behind one wall, carrying it to one of the other three positions and reattaching it. The corresponding spotlight was switched on and the previous one switched off. For rats that were not covered (see below), the card was visible throughout this procedure and so could be seen to move. For those trials in which the rat was also rotated, the animal was confined under an opaque cover to the rotating platter in the centre of the box. The cover was made of cardboard, of dimensions 18×23×32 cm high, and completely surrounded the rat except for a small hole in the top to allow the recording cable to pass through. The rat could thus see nothing but the inside of the cover while confined. The platter with the rat on it was then rotated slowly, at 0.13 rpm. The rotation took either 2 or 4 min to complete, depending on the angle of rotation (90° or 180°).

Using the above manipulations, the experiment proceeded as follows (Fig. 2). Each recording session began with a baseline trial, during which the rat foraged in the square box for 4 min while place cells were recorded. Following the baseline trial, all rats except two underwent four "training" trials during which they experienced one of three kinds of manipulation, described below. Ten minutes after each manipulation was completed, another 4-min trial was recorded.

Three experimental groups, of four rats each, were used in the experiment. The differences between the groups lay in their experience during the training trials. One group of rats (Group Uncovered) were allowed to roam freely in the box while the cue card was rotated from one position to another. These rats therefore were able to see that the card was mobile. The behaviour of place fields in these rats was compared with that of two other groups. In Group Covered, the rats either received no training trials at all (n=2), or they received trials similar to those of the Uncovered rat except that they were confined to the centre of the box under the cardboard cover for 3 min while the card was moved (n=2). These rats therefore never saw the card move. Based on preliminary results from these two groups, a third group was added. These rats were also exposed to training trials in which the card was moved from trial to trial, but in these rats their covering (and hence visual isolation) was made as brief as possible (the time taken to move the card, which was about 30 s; Group Briefly Covered). Thus,

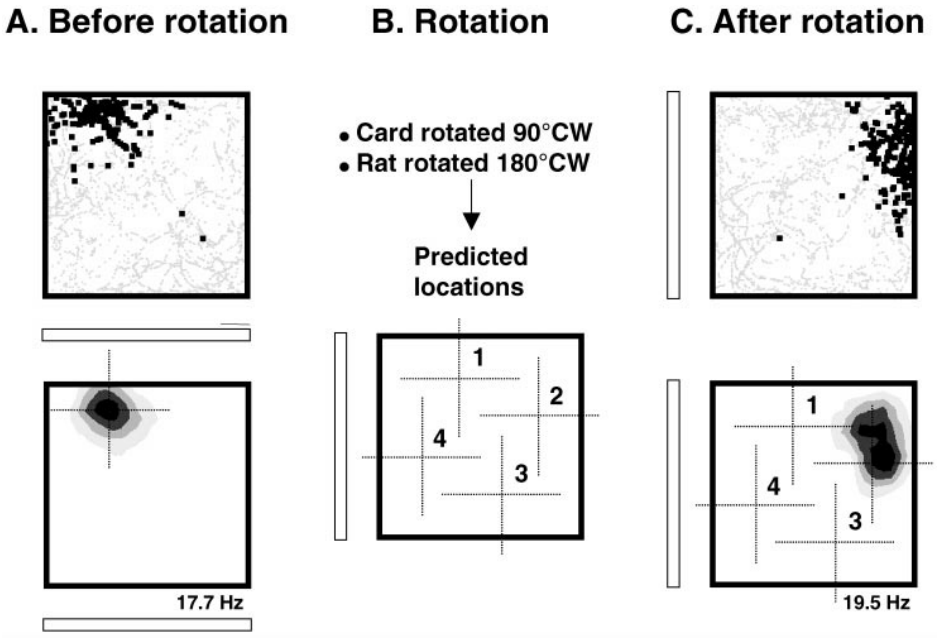

Fig. 3A–C An example of a single card+rat rotation, showing how the raw data were converted to an assessment of the place field's behaviour. **A** *Top panel* shows the square recording box, with the cue card (*white rectangle*) situated to the south of the box. The *stippled areas in the box* represent the positions of the rat during the 4-min recording trial. *Each small black square* represents one action potential of the cell, superimposed on the place where the rat was at the time. It can be seen that this particular cell mainly fired when the rat was in the northwest corner of the box. The *lower panel* shows the contour plot of the place field derived from the raw data (see text for method). The firing rate of the place cell at the field's peak (*located at the intersection of the dotted lines*) is shown *to the bottom right of the box*. **B** The rat and the card were then rotated by different amounts, the rat by 180° CW and the card by 90° CW. By rotating the previous peak location by multiples of 90°, four points were derived: (1) the predicted location of the field if it remained stationary with respect to the room, (2) the location if it rotated 90° CW (with the card), (3) the location if it rotated 180° (with the rat), and (4) the location if it rotated to the remaining corner. **C** Recording of the place field following that rotation. The field's new peak location was closest to point 2, the predicted location if it had rotated with the card. Thus, the field is adjudged to have followed the card

these rats resembled the two trained Covered rats except for the briefer duration of their isolation from the visual world when the card was moved. The purpose of this group was to see whether the differences between the Covered group and the Uncovered group might be explained by an effect of the duration of covering, rather than sight of the moving card per se.

A session consisted of a block of four card-only rotation training trials (except for the two non-trained Covered rats) followed by a block of four card+rat rotation test trials. In the card+rat rotation trials, given to all rats, the card and the idiothetic cues were both rotated, but by different amounts. These trials determined whether the card or the idiothetic cues now dominated. The rat was covered as described earlier and rotated by 90° or 180°. Meanwhile, the card was rotated by a different amount, also 90° or 180°, thus introducing a conflict between the two sources of directional information. After the rotation of the card and rat were complete, the rat was released and allowed to remain undisturbed in the environment for at least 10 min, before another 4-min foraging session was recorded to see what had happened to the place fields.

Each session contained a pseudorandom combination of rotations so that each card location, rat orientation, rotation and card/rat relationship was approximately equally experienced, evenly distributed across the experiment, and matched between animals. On a given recording day, between one and three sessions were run. If a session followed immediately from the preceding one, the last trial of that previous session was used as its baseline. Five sessions were run in total (see Fig. 2).

Data analysis

This was performed on a Sun Ultra workstation using proprietary software (Gignomai Ltd., UK). The collected waveforms were allocated to different cells by plotting the amplitude at the negative peak, the positive peak or the peak-to-peak amplitude of one electrode against that of each of the other three. The resulting clusters were separated by eye or by an automatic clustering algorithm. A place field was defined as a region of location-specific firing in which the peak rate after smoothing (see below) was greater than 1.0 Hz.

To determine the location and peak firing rate of each place field, a boxcar averaging process was used to convert the raw collection of spikes to a contour plot like that shown in Fig. 3. To achieve this, a 64×64 grid was placed on the camera viewing area (about 130×130 cm) and overlapping square bins of size 20×20 cm were placed around each grid point. For each bin (centred on a grid point) the firing rate for a given cell was determined by dividing the number of spikes it fired in that region by the amount of time the rat had spent there. Interpolation was then used to determine the firing rate for pixels between the grid points, and the resulting data were depicted as a grey-scale contour plot, autoscaled so that each grey gradation represented 20% of the peak firing rate. The point at the centre of the bin containing the highest firing rate was used as a measure of the location of the peak of a place field.

The following analysis was undertaken in order to determine whether a cell's field had rotated along with the card, the rat or neither. For each trial, the predicted position of the peak was calculated assuming the place field had done one of four things: remained stationary with respect to the room, rotated with the card, rotated with the rat, or rotated to some other position(s) (Fig. 3). This was done by subjecting the peak position from the previous trial to the various rotations to generate four predicted field positions. The Euclidean distance of the actual field position from the predicted positions was determined, and the field allocated the rotation corresponding to the smallest of these distances. The results produced by this analysis only disagreed with the experimenter's "eyeball"

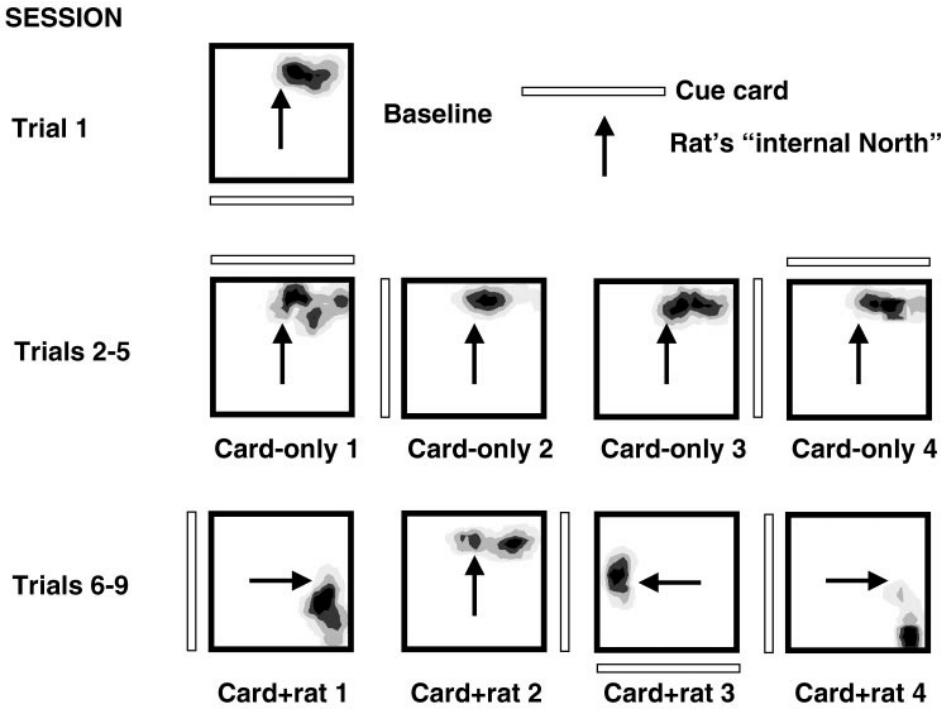

Fig. 4 Recording of a single place field for an entire session for an Uncovered rat. Trial 1 is a baseline run with the card (*white rectangle*) in the south – the field was located in the northeast corner. Trials 2–5 consisted of four card-only training trials, where the rat's internal direction sense (*arrows*) remained unrotated but the card was rotated to the north, west, east and then north again. Note that the place field did not rotate during these trials. Trials 6–9 consisted of four card+rat rotation test trials. The rat's internal direction sense was rotated 90° CW, 90° CCW, 90° CCW and 180° CW. The card was rotated 90° CW, 180°, 90° CW and 90° CW. It can be seen that on each trial, the place field's rotation corresponded to that of the rat and not the card. Thus, this cell had "learned" to use the rat's idiothetically based internal direction sense rather than the card as an orienting cue

assessment of the behaviour of the field on a small number of occasions (usually with a large field near the centre of the box).

For each session, the number of times the field followed the card, the rat or neither was expressed as a percentage of the total number of rotations of a given trial type (training or conflict).

Histology and cell localisation

After completion of recording each rat was killed with an overdose of sodium pentobarbital (Lethobarb, 10 mg) and perfused transcardially with saline followed by 4% paraformaldehyde. The brain was extracted and stored in 4% paraformaldehyde, and was later sliced coronally in frozen sections 40 µm thick, mounted and Nissl stained to allow visualisation of the electrode track. The location of each cell was estimated from the depth of the electrode at implantation plus the distance through which the microdrive had been advanced. This distance was superimposed on the electrode track obtained histologically.

Results

The procedure and resulting data of a typical card+rat rotation trial are shown in Fig. 3. It can be seen that with the cue card in the south position, the place cell was active whenever the rat was in the northwest corner of the box. The rat was then enclosed on the rotating platter, which was slowly rotated by 180°, over a period of 4 min. Meanwhile, the card was repositioned in the west. When the rat was released and a further trial recorded, the place field was now in the northeast: that is, it had rotated 90° CW from its previous position, thus following the card and not the idiothetic cues.

An example of the recording of a single place cell for a whole session is shown in Fig. 4. Because place fields that rotated (as opposed to switching off or on) always did so in unison, only the "best" cell (most clearly isolated and having the most stable field) from each session was selected to be representative of the orientation of the place representation and analysed in detail. Sometimes cells were recorded for the training or test part of a session only, and sometimes cells were recorded from more than one session. In total, 54 cells were recorded from the 12 rats, of which 46 were used in part or all of 1 session, 5 for between 1 and 2 sessions, 2 from 3 sessions and 1 for 5 sessions.

No differences were seen in the behaviour of cells from CA1 or CA3 and so these data are combined. For clarity, the results from group Covered and group Uncovered are presented first, and then compared with the results from the Briefly Covered group.

Occasionally it was not possible to run all four trials, and for one of the Uncovered rats, only four sessions were run. For these reasons, the numbers of trials presented below are not the same for all groups.

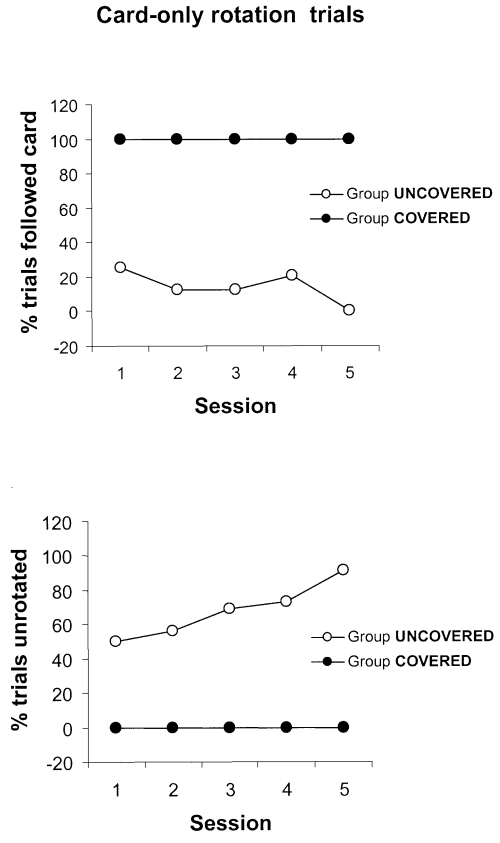

Fig. 5 Comparison between the behaviour of fields from the Covered and Uncovered rats during card-only rotations (the "training" trials). Only two rats from the Covered group underwent these trials. The *top graph* shows that for both Covered rats, the fields followed the card on 100% of trials. By contrast, for the four Uncovered rats, which saw the card move prior to each trial, the fields followed the card on only about 25% of trials. The *bottom graph* shows that, correspondingly, the fields in the Covered rats never remained unrotated whereas the fields in the Covered rats initially remained stationary on about 50% of trials, and by the end of five sessions were remaining stationary almost 100% of the time

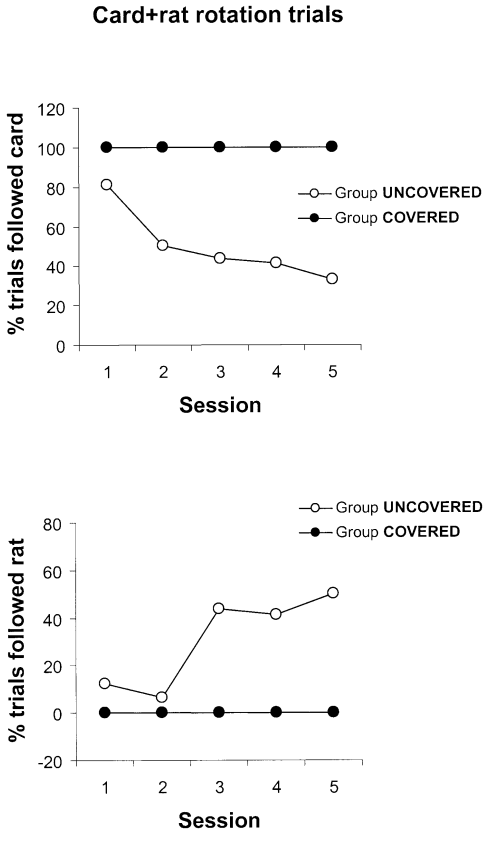

Fig. 6 Comparison between the behaviour of fields from the Covered and Uncovered rats during the card+rat rotation probe trials, where the card and the rat were rotated by different amounts. All rats underwent these trials. The *top graph* shows how often the fields followed the cue card: for the Uncovered rats this was 100% of the time, while for the Covered rats it began at about 80% of trials and by the end of five sessions it was only 40% of trials. The *bottom graph* shows how often the fields followed the rat's internal direction sense, and here the converse situation held: fields in the Uncovered rats never followed the rat while fields in the Covered rats initially followed it on about 12% of trials, and by the fifth session followed it on about 50% of trials

Card-only rotation trials (Covered vs Uncovered)

The card-only rotations (the training trials) occurred before the card+rat rotation trials for all four of the Uncovered rats and two of the Covered rats. The Uncovered rats saw the card move during these trials, while the Covered rats were enclosed for 3 min while the card was moved. Place cells were recorded after each trial to see whether the fields had rotated with the card, or remained stationary (and therefore stayed fixed relative to both the rat and the static background cues). It can be seen from Fig. 5 that the fields in the Covered rats followed the card on all the trials, but fields in the Uncovered rats followed it on very few of the trials. When the data were averaged across all 5 days, the place fields in the two Covered rats followed the card on all trials (40/40) while in the four Uncovered rats, they followed the card on only 15% of trials (11/75). Conversely, on the card-only rotations, fields in the Covered rats never remained stationary relative to the room, while those in the Uncovered rats remained stationary on 54 trials (72%). These differences were highly significant (χ^2=68.42, *df*=1, *P*<<0.0001).

These findings indicate that the place fields of rats that saw the card move were only weakly influenced by the card (despite its high salience), while those in rats that did not see it move were strongly influenced by it.

Card+rat rotation trials (Covered vs Uncovered)

Card+rat rotation trials followed immediately after the card-only rotation trials. Now, *all* the rats were hidden beneath the opaque cover while both the card and the rat were rotated. A difference between the two groups was evident from the first session and increased steadily across the 5 days (Fig. 6). For the Covered rats, place fields followed the card on every single trial across all

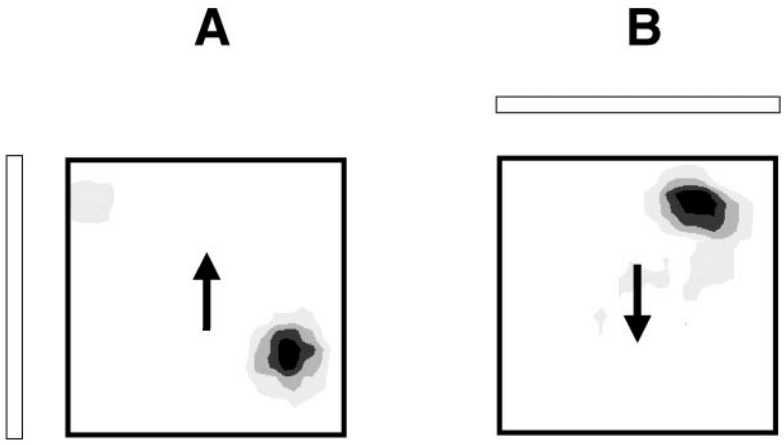

Fig. 7A, B Example of a field whose rotation equalled the sum of the rat's and the card's rotations. **A** Before the rotation, the field was in the southeast corner of the box. **B** The cue card was rotated 90° CW from the west to the north, while the rat was rotated 180° CCW. The field was now found in the northeast corner, corresponding to a rotation of 90° CCW

5 days. For the Uncovered rats, while the fields followed the card on the first session on 81% of trials (13/16), the card following showed a marked decline over the five sessions, and on the last session the fields followed the card on only 33% (4/12) of trials. Similarly, the tendency of the fields to follow the idiothetic cues was also markedly different between the two groups. For group Covered, the fields never followed the rat on any of the trials (since they had followed the card instead on all trials). For group Uncovered, the fields followed the rat on only 13% of trials (2/16) to begin with but by the fifth session this had increased to 50% (6/12) of trials. For the five sessions combined, these group differences were highly significant (χ^2=33.24, *df*=1, *P*<<0.0001). When the sessions were analysed separately, the value of χ^2 was seen to increase steadily, becoming significant by the third session.

Thus, although the card+rat rotation trials were conducted under identical conditions for the two groups, when the visual and idiothetic cues were placed in conflict the fields followed the visual cue in rats that had never seen the card move, but showed a strong tendency to follow the idiothetic cues instead, in rats that had seen it move.

Individual cells did not always rotate their fields following a rotation. Sometimes cells switched off or on following a manipulation, or shifted their fields slightly (often from a corner to further along a wall) or became unisolatable. An early (unquantified) impression was that fields showing unstable behaviour were most likely to be those with low rates and/or more dispersed fields, suggesting that perhaps they had weak inputs to begin with which were more prone to rearrangement.

Only in Uncovered rats was it the case that even the field selected for analysis (the "best" field) showed neither rat-following nor card-following behaviour. In 9 of the 13 trials, this "other" behaviour consisted of rotations to one of the other two locations within the box, which, interestingly, corresponded on all nine trials to the position that would be predicted if the rotation of the card was added to the rotation of the rat (Fig. 7). On the remaining trials, the fields either shifted along one of the walls of the box or split in two, resulting in the peak lying outside the region predicted by any of the rotations.

Despite these variations from pure rotational behaviour, no case was ever seen where a cell's field rotated by one amount while another cell's field rotated by a different amount. This suggests that the directional signal onto place cells is unitary (see "Discussion").

Comparison of card-only and card+rat rotation trials (Uncovered rats)

Card-following behaviour was compared in the Uncovered rats between trials that had been immediately preceded by visible card rotations (the card-only rotation trials) and those in which visible movement of the card had occurred tens of minutes or more beforehand (the card+rat rotation trials). The average card-following across the 5 days for the Uncovered rats in the card-only rotations was only 15% (11/75 trials), while for the same rats in the card+rat rotation trials it was considerably greater, at 52% (38/73 trials). A paired *t*-test was done to compare the percentage of trials in which fields followed the card during card-only rotations, with the percentage in the same rat that did so during the card-rat rotations. This showed these values to be significantly different [$t_{(3)}$=–3.53, *P*<0.05]. This finding shows that place fields that did not follow the card when it had been visibly moved would nevertheless sometimes follow it when it was then invisibly moved, showing that there was still an input from the card onto these cells, even when conditions were such that it could not override the other sources of information.

Briefly Covered rats

This group was introduced after preliminary results from the other two groups, in an attempt to determine why fields stopped following the cue card in the Uncovered rats. There are two obvious possibilities: one is that the movement of the card activated movement detectors somewhere in the visual system, indicating that the card was moving. The other is that because the rat was never isolated from the visual world while the card was moved,

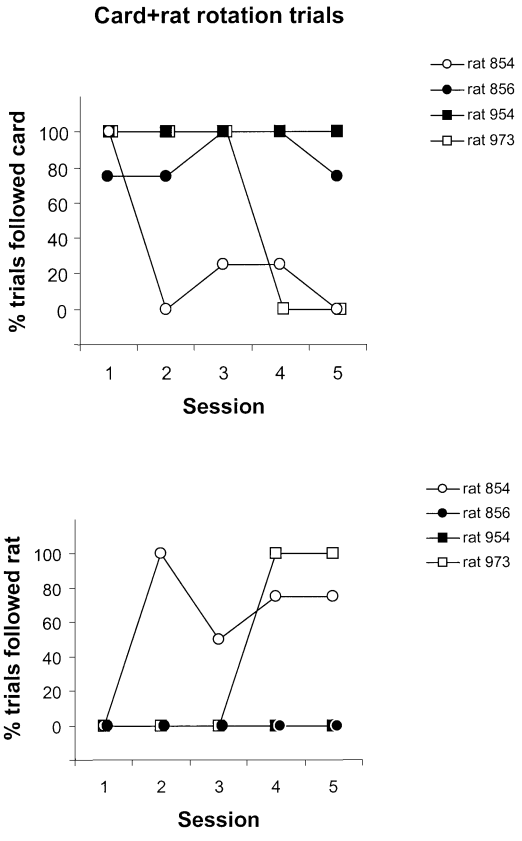

Fig. 8 Behaviour of fields from the four individual Briefly Covered rats, showing the dichotomy in behaviour: fields from rats 856 and 954 behaved like those in Covered rats, following the card most or all of the time (*top graph*) and the rat none of the time (*bottom graph*), while fields from rats 854 and 973 behaved like those in Uncovered rats, following the card initially but following the rat instead by the end of five sessions. Note that the transition from card-following to rat-following occurred abruptly in both rats, though on different sessions

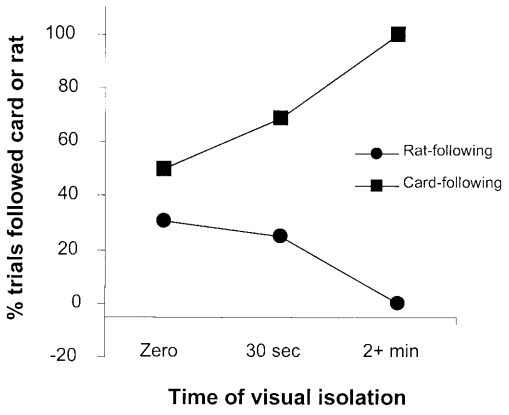

Fig. 9 Relationship between duration of covering and percentage of trials in which the fields followed the cue card (*squares*) or rat (*circles*). Note that in the "zero-time" condition, these rats had also seen the card actually move, which may or may not have contributed to the change in cue-following in these rats. Nevertheless, this graph suggests that prolonging the time of visual deprivation may increase the propensity of place fields to "trust" the card in preference to the internal direction sense

it was always fully oriented, and so able to detect that the card was moving with respect to other cues (such as idiothetic cues). By contrast, in the two Covered rats that received covered card-only rotations, and in all the rats during the card+rat rotation trials, the card was only seen again in its new position after the rats had been isolated from the visual world for between 2 and 4 min. Perhaps after several minutes of covering, the resulting conflict between the idiothetic cues and the visual cue was so weak that the influence of the cue card was not damaged by these trials.

Though these possibilities are not mutually exclusive, an attempt was made to separate them by adding a third group of rats (group Briefly Covered) which did not see the card move, but in which isolation from the visual world was made as short as possible so that the rats were close to fully oriented when they were released from the cover. Ideally, the time of isolation should have been close to zero, but in practice it took about 30 s to move the card.

The individual behaviours of fields from these four rats in the card+rat rotation test trials are shown in Fig. 8. It can be seen that two of the four rats behaved like Uncovered rats, in that their card-following decreased with time and fell to zero, while their rat-following increased markedly. The other two rats behaved like Covered rats in that their card-following remained robust and their rat-following remained at zero.

This finding leads to the tentative conclusion that the effect of covering the rats while the card was moved was time dependent, being strongest in rats that were isolated from the visual world for 2–4 min while the card was moved, moderate in rats isolated for only 30 s and non-existent in rats that were "hidden" for zero time (that is, that were uncovered while the card was moved; Fig. 9). Whether or not there is an additional contribution in the Uncovered rats from actually seeing the movement of the card cannot be determined at this stage.

Discussion

We found that when idiothetic cues were placed in conflict with a cue card, the card predominated in controlling the orientation of place fields, but only if the rats were always covered for several minutes while the card was moved. If the rats saw the card move for several trials each day, then place fields rapidly stopped following the card and started following the idiothetic cues instead. This switch of control from the visual to the idiothetic cues persisted even when these rats were covered during the card rotation preceding the "test" trials. This shows that although visual information was initially preferred, the idiothetic cues took control when the visual cue was seen to be mobile.

This raises the question of how place cells in rats that had seen the card move "knew" that it was unreliable. There seem to be two possibilities. The first is that the cells detected the card's actual movement, via activation

of visual motion detectors, or because the card shifted across a background constellation of static cues, sequentially revealing and obscuring features such as the folds of the curtains. The second is that they detected the resultant conflict between its new location and the idiothetic cues.

The latter alternative is the one proposed by Knierim et al. (1995) discussed in the "Introduction". At first sight it seems to be contradicted by the findings that place fields in the Covered group never stopped following the card, even though its location also conflicted with the idiothetic cues on every test trial. However, the results from the Briefly Covered group provide a possible resolution of this apparent paradox. In this group, whose rats were also isolated from the environment but for only 30 s, fields in two of the four rats behaved like those of the Uncovered rats: that is, they stopped following the card and started following the rat. Perhaps if it been possible to reduce this time to zero (that is, to move the card instantaneously), then all four rats might have behaved like the Uncovered rats. In other words, perhaps the relative preference of the cells for visual vs idiothetic cues is proportional, among other things, to the duration of visual deprivation (see Fig. 9).

Why should visual isolation cause place cells to prefer visual over idiothetic cues? One possibility is that the idiothetic directional representation needs to be continually updated by the visual cues (not necessarily the card) in order to remain at full strength (Goodridge and Taube 1995). If this is so, then in the card+rat rotation trials there might be attenuation of the idiothetic signal caused by the 2-4 min of isolation from the recording environment, in which case there would never really be a disagreement between the visual and idiothetic cues. It is as if the rats "forgot" which way the environment was oriented while they were enclosed in the box and hence saw no reason not to go on trusting the cue card after their release. In contrast, for the visible card-only rotations experienced by the Uncovered rats, visual information was always available and the idiothetic cues perhaps were always at full strength. Thus, when the card was moved on these trials, the idiothetic directional signal would conflict with that provided by the card and indicate that the card must have moved. Note that this explanation posits two different kinds of visual information: directional information, provided by salient polarising landmarks such as the card, and idiothetic-resetting information, which could theoretically originate from any stable environmental features (such as the corners of the box).

If the above explanation is true, it follows that the idiothetic directional signal becomes less reliable over a few minutes when it is not receiving periodic resetting information from the visual environment. This is supported by the results of several other experiments (Bures et al. 1998; Knierim et al. 1998) that show that the idiothetic cues lose their accuracy over time. Do they become weaker, are they perhaps flagged by some warning signal that indicates their reliability is to be questioned, or do they maintain strength but drift out of alignment with the visual environment? The results of the present experiment at first sight suggest that they either weaken or become labelled as unreliable, because if idiothetic cues remained at full strength during the covering procedure then they should still compete successfully with the cue card, even though they had drifted. However, because the recording environment in this experiment was a square and not a circle, any drift of the idiothetic "compass" would have caused it also to conflict with the box geometry when the rat was released (unless it happened to drift by some multiple of 90°). It might be that the combination of the box geometry and the cue card could override the idiothetic cues whereas the card alone could not.

Thus, although the present experiment suggests some sort of change in the internal compass following visual isolation, it cannot differentiate between drift of the orientation versus decay in the strength of the signal. It seems likely, though, that the signal drifts rather than attenuates. One way of resolving this question would be to repeat our experiment in a circular environment, in which the geometry of the recording chamber would never conflict with the idiothetic cues, no matter how far they might drift. In such an environment, then if the idiothetic cues drifted but did not weaken, the fields in Covered rats should behave like those in Uncovered rats.

Site of learning-related changes

Given that a visually unstable cue card became disconnected from the place cell representation, where in the brain did these experience-related changes in place fields originate? There are grounds for thinking that it is not in the hippocampus itself. This is because in rats where the fields changed their allegiance from the card to the idiothetic cues, no case was ever seen where one field rotated with the card while another rotated with the rat, even during the "changeover" from one preference to the other. If synapses onto place cells were mediating these changes, it might be expected that some cells would learn faster than others and so shift from the card to the rat earlier than others. Because this never happened, it seems likely that the directional signal onto place cells is homogeneous, and that the competition between the visual and idiothetic cues was resolved upstream of the place cells. However, perhaps the local network properties of the place cell representation could cause any "disagreement" about the orientation of the environment to be settled within the hippocampus itself. This seems less likely, but could be tested by inactivating the hippocampus while the card-movement experience was taking place. If the learning takes place outside the hippocampus, then such inactivation should not affect subsequent place field behaviour and the fields should switch to the rat, as usual. On the other hand, if it usually occurs on the place cells themselves, then if these are inactivated during the learning experience, they should subsequently

behave as though they had never seen the cue card move (assuming of course that such changes are activity dependent).

If the directional signal is resolved outside of the hippocampus, where is the likeliest place? It is thought that place cells receive their directional information from a system of head direction cells, which are located in various brain regions including the postsubiculum (Taube et al. 1990), anterior thalamus (Blair and Sharp 1995; Taube 1995) and posterior cortex (Chen et al. 1994). An elegant series of experiments has shown that head direction cells receive both idiothetic and visual inputs (Blair and Sharp 1996; Knierim et al. 1995; Taube and Burton 1995). Because the firing of anterior thalamic head direction cells predates that of postsubicular cells by about 25 ms (Blair et al. 1997), and because lesions of this structure abolish direction activity of head direction cells in the postsubiculum, it has been suggested that the anterior thalamus passes movement-related information on to the postsubicular cells, which are also receiving visual information (Goodridge and Taube 1997). It therefore seems plausible that the learning-related shift in influence from one information source to the other may take place on these postsubicular cells, in which case this would be a good place to look for visual/idiothetic competition. Again, this hypothesis could be tested with an inactivation experiment of the type described above.

Integration of visual and idiothetic cues

When the present results are taken in combination with those of previous studies, the following picture of visual/idiothetic interaction is emerging. It seems that the visual and idiothetic systems "help" each other, each correcting for the other's inadequacies.

First, the visual environment corrects errors in the idiothetic system so long as these are small, presumably because if there is only a small mismatch, it is more likely to be due to drift of the idiothetic signal because of accumulated errors than to a creeping shift in the location of a landmark. In support of this, if the rat is in the environment when the change is made, small (but not large) shifts of a cue card can "pull" place fields along with the card (Rotenberg and Muller 1997) and small (but not large) rotations of HD cell firing directions, made by rotating the rat in the dark, can be corrected by the visual cues when the lights are restored (Knierim et al. 1998). Second, the idiothetic cues take control when the mismatch is large, presumably because a large error occurring only a short time after the last visual updating is more likely to be due to shift of whatever object was the source of the visual stimulus, and less likely to be due to a sudden large mistake made by the path integration system. Evidence for this is the converse of that described above: with large shifts of a cue card, place fields remain unrotated (Rotenberg and Muller 1997) and with large rotations of the HD system made in the dark, visual cues do not restore its previous orientation (Knierim et al. 1998). This may not be the case, however, when the cue has been experienced to be stable for some time (Jeffery 1998).

The results of the present experiment add three further contributions to the picture. First, the ability of idiothetic cues to override the visual system in cases of large disagreements may depend on the strength ("confidence") of the idiothetic signal. We found that when animals were isolated from visual updating for 2 or 3 min while a visual/idiothetic conflict was introduced, the visual system then reliably predominated even for large conflicts of 180°. When animals were only isolated for 30 s while the conflict was introduced, then visual dominance was much less certain. The possible reasons for such a reduction in idiothetic influence were described earlier.

Second, the ability of the visual cues to override the idiothetic system depended on whether they had been seen to be mobile. If they had been seen to move, then they yielded control of place field orientation to the idiothetic cues, even if the rats had been isolated from visual updating information. This shows that the visual inputs are plastic, and also that the idiothetic signal was still present even after visual isolation. This provides the spatial system with a mechanism for learning to disregard mobile objects and avoiding trying to use them as landmarks.

Third, the visual and idiothetic cues may interact via summation of the strengths of their signals. Evidence for this is only scant at present, but is shown by the fact that in cases where place fields rotated but did not follow either the rat or the card, the location could be predicted by summing the rotation of the card *and* the rat. This leads to the hypothesis that when either the visual or idiothetic signal is strong, it overrides the other, but when they are matched in strength (as, for example, during weakening of the visual cue during instability learning) then the result is a compromise between the two signals.

It appears, then, that whether a visual or idiothetic signal predominates depends on a number of factors including the degree of conflict between them, the confidence of the idiothetic signal and the strength of the visual input.

To conclude, the results of the present experiment show that place cells can learn that a visual directional landmark is unstable, and to use the idiothetic cues instead. The likeliest indicator that the cue is unstable is the resulting mismatch between its direction and that indicated by the idiothetic cues. This mismatch appears to be attenuated by visual deprivation lasting more than a minute or so, suggesting that the idiothetic directional signal needs regular updating by visual cues to stop it either drifting or becoming weaker. The site of the learning about directional landmark stability is unknown, but is probably located outside the hippocampus itself.

Acknowledgements The authors are grateful to Stephen Burton, Jim Donnet, Dave Edwards and Clive Parker for technical assistance, and to Colin Lever for reviewing the manuscript. This work was supported by a Medical Research Council programme grant to John O'Keefe.

References

Barlow JS (1964) Inertial navigation as a basis for animal navigation. J Theor Biol 6:76–117

Biegler R, Morris RG (1996) Landmark stability: further studies pointing to a role in spatial learning. Q J Exp Psychol B 49:307–345

Blair HT, Sharp PE (1995) Anticipatory head direction signals in anterior thalamus: evidence for a thalamocortical circuit that integrates angular head motion to compute head direction. J Neurosci 15:6260–6270

Blair HT, Sharp PE (1996) Visual and vestibular influences on head-direction cells in the anterior thalamus of the rat. Behav Neurosci 110:643–660

Blair HT, Lipscomb BW, Sharp PE (1997) Anticipatory time intervals of head-direction cells in the anterior thalamus of the rat: implications for path integration in the head-direction circuit. J Neurophysiol 78:145–159

Bures J, Fenton AA, Kaminsky YU, Wesierska M, Zahalka A (1998) Rodent navigation after dissociation of the allocentric and idiothetic representations of space. Neuropharmacology 37:689–699

Chen LL, Lin LH, Green EJ, Bames CA, McNaughton BL (1994) Head-direction cells in the rat posterior cortex. I. Anatomical distribution and behavioral modulation. Exp Brain Res 101:8–23

Goodridge JP, Taube JS (1995) Preferential use of the landmark navigational system by head direction cells in rats. Behav Neurosci 109:49–61

Goodridge JP, Taube JS (1997) Interaction between the postsubiculum and anterior thalamus in the generation of head direction cell activity. J Neurosci 17:9315–9330

Jeffery KJ (1998) Learning of landmark stability and instability by hippocampal place cells. Neuropharmacology 37:677–687

Jeffery KJ, Donnett JG, Burgess N, O'Keefe JM (1997) Directional control of hippocampal place fields. Exp Brain Res 117:131–142

Knierim JJ, Kudrimoti HS, McNaughton BL (1995) Place cells, head direction cells, and the learning of landmark stability. J Neurosci 15:1648–1659

Knierim JJ, Kudrimoti HS, McNaughton BL (1998) Interactions between idiothetic cues and external landmarks in the control of place cells and head direction cells. J Neurophysiol 80:425–446

Markus EJ, Barnes CA, McNaughton BL, Gladden VL, Skaggs WE (1994) Spatial information content and reliability of hippocampal CA1 neurons: effects of visual input. Hippocampus 4:410–421

Muller RU, Kubie JL (1987) The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. J Neurosci 7:1951–1968

O'Keefe J, Burgess N (1996) Geometric determinants of the place fields of hippocampal neurons. Nature 381:425–428

O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res 34:171–175

O'Keefe J, Nadel L (1978) The hippocampus as a cognitive map. Clarendon Press, Oxford

Quirk GJ, Muller RU, Kubie JL (1990) The firing of hippocampal place cells in the dark depends on the rat's recent experience. J Neurosci 10:2008–2017

Rotenberg A, Muller RU (1997) Variable place-cell coupling to a continuously viewed stimulus: evidence that the hippocampus acts as a perceptual system. Philos Trans R Soc Lond B Biol Sci 352:1505–1513

Sharp PE, Blair HT, Etkin D, Tzanetos DB (1995) Influences of vestibular and visual motion information on the spatial firing patterns of hippocampal place cells. J Neurosci 15:173–189

Stackman RW, Taube JS (1997) Firing properties of head direction cells in the rat anterior thalamic nucleus: dependence on vestibular input. J Neurosci 17:4349–4358

Taube JS (1995) Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. J Neurosci 15:70–86

Taube JS, Burton HL (1995) Head direction cell activity monitored in a novel environment and during a cue conflict situation. J Neurophysiol 74:1953

Taube JS, Muller RU, Ranck JBJ (1990) Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. J Neurosci 10:420–435