# Head direction cells use a head-referenced dual-axis updating rule in 3D space

Williams, M[1,2], Street, JS[3], Burgess, N[3,4], Jeffery, KJ[5*].

1 Department of Clinical and Movement Neurosciences, Institute of Neurology, University College London

2 Institute of Behavioural Neuroscience, University College London, where the experimental work was done

3 Department of Clinical and Experimental Epilepsy, UCL Queen Square Institute of Neurology, University College London

4 UCL Institute of Cognitive Neuroscience, University College London

5 School of Psychology & Neuroscience, University of Glasgow

* Corresponding author

**Author contributions**

MW gained funding, designed the experiment, designed and built the apparatus, collected the data, conceptualised and performed analyses, edited the manuscript; JS provided advice and support, and edited the manuscript; NB provided advice and support, and edited the manuscript; KJ conceptualised the experiment, gained funding, conceptualised and performed analyses, and edited the manuscript.

**Acknowledgements**

The work was supported by a BBSRC training grant to MW (BB/M009513/1), a Wellcome Principal Research Fellowship to NB (202805/Z/16/Z) and a Wellcome Investigator Award to KJ (103896/Z/14/Z).

## Abstract

Head direction (HD) cell signalling, a compass signal for the brain's spatial map, may

support spatial orientation during 3D movement by using a 'dual axis' (DA) updating

rule that summates head rotation around a local axis and local-axis rotation around the

gravity axis. This theoretically enables continuous horizontal orientation by

compensating for errors that would accrue if rotations were tracked in the head-plane

only. To test for operation of this rule, we recorded HD cell activity from rats exploring a

hemisphere. We assessed HD cell tuning curves in either the standard horizontal

reference frame or in two reference frames governed by a DA rule, with the local axis

referenced to either the head (DAH) or the local surface (DAS). Tuning curves were

sharpest and tallest in the DAH reference frame, supporting operation of a DA rule and

suggesting that the relationship of the head to gravity is a component of the input to the

HD system.

**Main**

Spatial cognition in large-scale space requires a stable internal compass, implemented by the head direction (HD) neurons found in several brain regions. HD cells signal the azimuth (horizontal direction) of the nose with directionally oriented tuning curves, firing maximally when the nose aligns with the cell's preferred firing direction (PFD)[1]. For a given cell, firing falls off smoothly as the nose's angular distance from the PFD, called alpha, increases, falling to zero at approximately alpha = +/- 45 degrees. Network activity is organised as a hypothetical "ring attractor," [2–4] updating in response to the left-right head-turns that dynamically vary alpha. A difficulty with this notion is that during 3D movement, such as on a sphere (Fig. 1a), a ring attractor network with a purely yaw-based updating rule would accumulate so-called Berry-Hannay errors [5] due to occult changes in azimuth arising from 3D movements (Supp. Fig. 1a). However, HD cells do maintain a stable PFD on a sphere (Fig. 1b), indicating an ability to correct these errors. This might indicate pure reliance on external directional environmental cues, but such reliance would be vulnerable to losing sight of these cues. Previous work has suggested that the vestibular updating that HD cells use to supplement visual orienting on the horizontal plane could still work in 3D if HD cells could track rotations around two axes simultaneously [6,7]: head rotations around a local axis, and local-axis rotations around the gravity axis. These would be summed to yield a local PFD, adjusting with rotation of the head's dorso-ventral (DV) axis through space (Fig. 1c) such that on returning to a horizontal pose, the local and global PFDs become re-aligned. This hypothetical modification to the ring attractor architecture would allow effective navigation in 3D using only a 2D representational substrate, and even during loss of external vision, thus maintaining global orientation with minimal neural resource use.

Indirect support for use of this "dual axis" (DA) rule was recently reported by Page et al.[6] and Laurens and Angelaki[7], but these studies did not directly reveal its operation, and furthermore, did not address the question of whether the local axis is referenced to the environment surface (allocentric) or to the head itself (egocentric). Either is plausible, given the intersection of external and internal sensory inputs in this system. To

investigate, we recorded HD cells from postsubiculum (PoS) and dorsal anterior

thalamus (ADN) in rats exploring a hemispherical arena (see video[8]) , to test whether

activity is best described by a DA rule, and if so whether the local axis is allocentric or

egocentric. We refer to three rotational reference frames, each defined by an axis of

rotation and the corresponding orthogonal plane (Fig. 1d), as follows: the horizontal

plane plus gravity axis (GA), the plane tangent to the environment surface plus the

surface-normal axis (SA), and the transverse plane of the head plus the dorso-ventral

head axis (HA). The DA rule using the SA axis is referred to as DAS, and using the HA is

DAH. We thus investigated whether the PFDs of the neurons would be best described

by a purely horizontally referenced updating rule (GA), or better described by the DAS or

DAH (Fig. 1e), an effect that should be greatest at the azimuthal poles defined by the

PFD (Fig. 1f).

Figure 1 Dual axis updating of HD cells  (a) Rat exploring a sphere (still from video[a]). (b) Position

and spike data from a single sphere recording session. The path of the animal is shown in grey and

spikes are shown as dots color-coded by azimuth (see colorbar). Left-right: Path from side view,

example HD cell, and example non-HD cell. (c) Correcting for rotations of the head around the

gravity axis (See Supp. Fig. 1 for more details). By the DA rule, as the rat progresses around the

sphere its local reference frame rotates opposite to the rotation direction of travel to maintain

alignment with the global frame. (d) The three orientational reference frames (each described by

an axis and its plane) examined in this study. (e) Vector fields generated from the recording

session of a single HD cell, showing the horizontal pointing direction of the cell's PFD in three

different reference frames: note the divergence from parallel with the DA rules. (f) Theoretical

divergence of local PFDs from the global PFD, as a function of position on the sphere (hot colors =

greatest).

To investigate whether the DAS or DAH rule better describes HD firing in 3D than the GA

rule that applies in the horizontal plane, we recorded 2270 single units, using tetrodes,

in 10 rats from either anterodorsal nucleus of the thalamus (ADN, n = 5) or

postsubiculum (PoS, n = 5) while the animals freely foraged over the surface of a 1-m

diameter hemisphere. Position in 3D coordinates was determined  by tracking a head-

mounted LED with 5 cameras, and 3D orientation of their heads using an inertial

measurement unit (IMU) calibrated with a visual ground truth (Supp Fig. 2).

The animals did not sample all directions uniformly over the hemisphere, preventing

assessment of full tuning curves over the sphere surface. Instead, we adopted a model-

fitting approach that allowed us to infer the alpha at each moment in a local plane

relative to a local PFD. Local alpha was computed based on a broadly estimated global

PFD assessed in the horizontal reference frame (overhead view), together with

measured instantaneous head orientation and firing rate. We describe tuning curves by

a simplified version of the von Mises equation for HD cells (Equ. 1), relating firing rate to

alpha as follows:

$$FR = we^{(\kappa cos(\alpha))} \qquad \text{Equ. (1)}$$

Given a set of head directions and firing rates, this equation can be used in a search

process to discover the parameters of PFD, κ (kappa; tuning curve sharpness) and $w$

(weight; equivalent to peak) for a given cell's trial data. For all the recorded cells we first

applied this procedure to the global birds-eye-view data, using a measure of

instantaneous firing rate derived from a smoothed form of the spike train (Online

Methods; Supp. Fig. 3). We then selected the subset of HD cells (n = 436; Online

methods), based on κ > 0.5. This unbiased selection produced results that agreed well

with visual assessment of tuning curve quality, and the cells had typical tuning curve

properties in a standard horizontal reference frame (Supp. Table 1). Using the derived

model parameters we generated for each cell a model tuning curve, based on the

recorded firing rates, and generated an initial model fit to the tuning curves using

Pearson's correlation (R).

We then transformed the data by hypothesizing that cells are described best with

reference not to global PFD but to a local PFD, this being the global PFD modulated by a

DA rule. We inferred the local alpha at each moment by computing and summing two

rotations: (i) the yaw rotation of the head in either its own plane (DAH) or the tangent

plane of the sphere (DAS) relative to vertical (the GA, projected onto the plane), and (ii)

the yaw rotation of either the SA or HA in the horizontal plane, relative to the point on

the circle antipodal to the global PFD (that is, PFD + $\pi$ radians). This was chosen as the

reference point because it is where alpha is known, from prior work[9], to be zero when

the nose points directly upwards. As a null control, we computed alphas in the

egocentric reference plane of the rat's head (a single-axis rule ignoring gravity).

Using the resultant set of local alphas we then used the model search process as

previously to re-derive new $\kappa$ and $w$ model parameters, re-binned and re-plotted the

data to create new tuning curves, and fitted the new model to these. An illustration of

this procedure (see Online Methods for details) is shown for the example cell in Fig. 2a.

The tuning curves have been created using alphas from the four transformations. With

the two DA transformations the tuning is sharpened, both visibly and as reflected in the

peak-rate and $\kappa$ values. This effect is more pronounced for DAH than DAS.

**Figure 2 Sharpening of tuning curves with change in reference frame.** (a) Directional firing rate histograms of a single ADN cell computed in a standard gravity-axis reference frame (GA), surface-axis referenced dual-axis frame (DAS), head-referenced dual-axis frame (DAH) or head-referenced only frame (HA). The DAH tuning curve was tallest (Peak rate) and sharpest (kappa) with the greatest directionality (Rayleigh vector) and best von Mises model fit (red line). (b) Tuning curve parameters for all HD cells (n = 436; ** = p < 0.01, *** = p < 0.001). (c) Pairwise comparisons of the three tuning curve parameters for GA vs. DAH, color-coded by brain region, showing the distribution weighted towards DAH, above the y = x line of equality (black solid line).

We used linear mixed-effects models to examine the effect of model type (GA, DAS, DAH) on four response variables: peak firing rate, Rayleigh vector length, κ, and model fit. All models were fitted using maximum likelihood estimation, with model type as a fixed effect and random intercepts for rat and for cell nested within rat. Analyses were conducted using MATLAB's `fitlme` function. There was a significant effect of model type for all four variables (Supp. Table 2) with DAH significantly surpassing both other models for all four variables. DAS was next, surpassing GA for all variables except κ. The

pairwise comparisons for the GA vs. other models illustrate the robustness of the

effect across the cell population (Fig. 3a) and the superiority of DAH over DAS (Fig. 3b;

see Supp. Fig. 4 for comparisons between all three models). We compared the model fit

quality by using Pearson's R for visualisation (Fig. 3 b and c) and the AIC and BIC (Online

Methods) for quantification (Fig. 3d), finding that DAH consistently outperformed the

other models (Supp. Table 3).

**Figure 3 Model comparisons for the 436 HD cells.** (a) Pairwise comparisons as for Fig. 2,
showing the Pearson's R values for the GA model compared with each of the others. (b) DAS
compared with DAH showing a better fit for DAH. (c) Raw Pearson's R for the models ** = p <
0.01, *** = p < 0.001).(d) Count of the best-fit model using each of the criteria: log likelihood
(logL), Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), and Pearson's
correlation (R). DAH was the best model for all four.

Operation of a DA rule requires detection of the tilt of the head relative to gravity, raising

the possibility to find cells tuned to this parameter in our data set. This could have been

evident as cells tuned to a particular elevation, or cells with tuning curves in the HA

reference frame. We first looked at firing rate as a function of pure elevation (angle of

the nose above or below horizontal). Behavioural sampling of elevation angles was

restricted, (Supp. Fig. 5), with angles distributed around a peak of -0.52 radians,

consistent with the observed preference of the animals for circumferential and

downwards travel. However we found no clear evidence for cells tuned to elevation: a

plot of firing rates vs. elevation showed increased variance at the higher elevations

(Supp. Fig. 5b) but this was likely due to the dwell-time undersampling; neither visual

inspection nor statistical analysis (not reported) found evidence of a systematic

elevation-rate relationship.

We then looked at the subset of cells that were best described by the HA

transformation, which assessed angular deviation from the GA axis projected onto the

head plane independently of HA azimuth.  We identified those from the whole set of

2270 cells having  $\kappa$ > 0.5 in this transformation (indicating at least some directional

tuning), and from these, selected the subset with peak rates and R-vector values that

exceeded those for the DAH transformation, meaning that the HA was the better

descriptor. Visual inspection of these (n = 40; Supp. Fig. 6) revealed tuning curves lying

to either side of vertical, but none actually *on* the vertical (alpha = 0). These cells could

potentially serve as a mediating signal for azimuth-independent local axis rotation.

Note that there were very few of them however (1.8%).

It is possible that local axis azimuth could also independently be coded by location-

specific firing (place fields) on the sphere, which is plausible in light of past reports that

neurons in PoS showed conjunctive place x direction coding [10]. We thus looked at the

spatial distribution of spikes over the hemisphere for all 2270 cells, We identified a

small no. (52/2770) of neurons that showed clear place fields by eye: 7 from PoS and 45

from ADN. However, the angular location of the place fields closely accorded with the

direction of the cell's PFD, and when we computed conjunctive position x direction rate

maps the "place fields" were clearly seen to spuriously result from directional

undersampling of the areas devoid of spikes. Overall, therefore, we found no clear

evidence of positional encoding by these neurons.

Our results from the HD cells provide strong support for the dual-axis hypothesis,

because modulating the inferred PFDs by the DAH rule very reliably sharpened tuning

curves relative to the baseline curves. When comparing the two variants of the rule,

surface-referenced and head-referenced, the head-referenced fit was superior. Thus,

the cells are continuously integrating rotations of the head in the head plane with

rotations of the head axis around the gravity axis. Note that "rotation" does not

necessarily imply a dynamic process: static signals could be used by the brain to

compute instantaneous local alpha, provided there is a static sensory cue to global

direction (such as a landmark) and to the direction of gravity (from, presumably, the

vestibular otolith organ).

We can thus re-formulate the HD cell von Mises function to account for the second

axis, as follows:

212
$$FR = we^{(\kappa \cos(\alpha_{HH} + \alpha_{HG} + \pi))}$$
Equ. (2)

where $\alpha_{HH}$ is the alpha of the head rotation in the head plane around the HA (referenced

to the projection of the GA onto the head plane) and $\alpha_{HG}$ is rotation of the HA in the

horizontal plane around the GA. Although we have here described the head rotations in

two separate reference frames, mathematically the two summed rotations equate to a

rotation around a single tilted axis; hence "tilted azimuth" in Laurens and Angelaki's

formulation. We prefer to describe the axes separately since the two axes (head axis

and gravity) are natural organising reference frames for the brain, and this

nomenclature also reinforces the notion that activity can shift around the ring attractor

if the head rotates in its own plane (as we assumed previously) *or* if that plane rotates

around the GA, or both. The added role of the GA azimuth requires extension of

standard ring attractor models of HD updating to incorporate gravity. How the brain

processes these reference frames, and how and where they interact, remains to be

elucidated.

## Online methods

## Experimental procedures

### Animals and Care

Subjects were 10 adult male Lister-Hooded rats, obtained from Charles River, housed
in a 11hr/11hr day/night cycle, with 1hr of dawn and 1hr of dusk. Pre-surgery, animals
were housed in groups of 2-6 in a large cage (dimensions ≈2m x 2m x 2m). They were
provided with apparatus to encourage 3D movement, including boxes, ledges, ropes,
wooden balls, and spherical and volumetric climbing apparatus. Post-implantation
they were housed singly but in clear-topped cages so they could see each other. All
experiments were performed under UCL's establishment licence and policies, in
accordance with the Animal (Scientific Procedures) Act (1986).

### Surgery

Surgery was performed once animals reached 350g. Animals were anesthetized with
isoflurane and implanted with four tetrodes attached to an Axona poorlady microdrive
using standard stereotaxic procedures. Coordinates used for ADN were (in mm for AP,
ML and DV respectively) bregma -1.8, 1.3, -4.0, and for PoS -7.5, 3.3, -2.0. After the
tetrodes had reached the appropriate coordinates, the outer cannula of the drive was
lowered to the brain surface, sealed with petroleum jelly, and fixed in place with dental
cement. Once the drive was fully secure, the wire to the ground pin was soldered to the
ground wire of the drive, and the animal was removed from the stereotaxic frame. Post-
operatively, animals recovered in a heated recovery box until mobile. Analgesia was
given for 3 days in the form of meloxicam (0.2mg/Kg) in condensed milk, and wet high-
protein mash. Animals were closely monitored for 7 days, after which experiments
began.

*Apparatus*

The screening arena consisted of a 1.2m x 1.2m open topped wooden box, with 60cm

walls painted grey, a black antistatic polyvinyl floor, and a white cue card on one wall.

The arena was set in a room with ample distal cues, including the recording PC, and the

experimenter themselves.

In the experimental room, a 1m diameter spherical recording arena was situated in the

centre of the room, surrounded by clear distal landmark cues on three walls, with the

neural recording setup, PC, and experimenter behind a curtain on one side of the room

to avoid distracting the rat. The sphere was constructed from paper mâché with the

bottom 1/4 cut off to form a flat base. It was painted with non-reflective grey paint, and

covered with a plastic mesh to allow animals to grip the surface. A 15cm ledge around

the bottom prevented animals falling to the floor, and the whole apparatus was placed

on a 40cm high trolley to prevent escape. Since the animals never strayed below the

equator, we refer to it as the hemisphere. The cable of the neural recording system was

suspended with elasticated thread with a system of pulleys to increase travel distance

and allow the animal to move over the hemisphere unimpeded. Recordings were 20-30

minutes, depending on animal willingness to forage.

*Recording Protocol*

For hemisphere sessions, the headstage was plugged into the microdrive attached to

the rat's head, while the rat was distracted with a cereal reward. Once plugged in, the

rat was placed on the hemisphere, where he foraged for malt paste until satisfied;

usually up to 30 minutes. Malt paste was manually reapplied approximately every 5

minutes, primarily around the sides of the hemisphere to motivate exploration on the

steeper surface. Except for when re-baiting the hemisphere the experimenter remained

at the PC, behind a curtain, in order to monitor orientation, position, and neural

recording quality. If the wire to the headstage became twisted, malt paste was used as

bait to persuade the rat to turn back the other direction to untwist the cable. After

recording concluded, the rat was picked up, given another reward, the headstage

disconnected.

*Screening and recording*

Animals were screened for HD cells daily. Recordings were performed using an Axona

recording system. Tetrodes were advanced by 50µm each day after recording.

Screening sessions lasted 6 minutes each, after which the experimenter performed a

brief assessment of cell clusters and HDC presence in TINT (Axona). If HD cells were

present, the animal was moved on to the main recording session.

*Positional Tracking*

On the hemisphere, 3D position was recorded at 50Hz using a 5-camera setup (see[11]. A

1Hz TTL pulse from the recording system drove an LED in the frame of each camera,

which was used to synchronize the 5 analogue video streams offline. Camera feeds

were digitized and adjusted for lens distortion, and 3D reconstructions of the position of

a LED on the headstage were made using the MATLAB function triangulateMultiview.

*Orientation Tracking*

3D Head orientation was recorded on the hemisphere using a FSM300 (Hillcrest) inertial

measurement unit (IMU). The IMU was mounted to the headstage such that it was level

with the head (as defined by lambda and bregma), with the Y axis aligned with the nose,

the X axis in the direction of the right ear, and the Z axis aligned with the dorsoventral

axis of the animal. The same 1Hz TTL pulse used to synchronise the cameras was used

for the IMU. The output of the IMU includes an accuracy metric, 0 being lowest and 3

highest. Data for which the accuracy metric was below 3 (meaning estimated error > 2°)

was not used in analysis – these data typically amounted to 0-5% of a session's data.

A Bland-Altman analysis of a session in which a headset with both LEDs and IMU was

recorded, while being moved by hand, showed a mean difference between the two

recording methods of 0.88° (Supp. Fig. 1).

By recording HDCs and azimuthal nose direction in two sessions, one using the

traditional LED method, one using the IMU, we confirmed that the IMU was able to

provide accurate enough readings to identify HDCs (Supp fig). Simultaneous recording

of the orientation of the IMU and LED apparatus, rotated by the experimenter,

confirmed the accuracy of the IMU (Supp. Fig. 1).

*Histology*

At the end of recordings, once no new HDCs had been recorded for 7 days, animals

were culled and brains extracted. Anaesthesia was induced with 5% isoflurane, then

animals were given an intraperitoneal (IP) overdose of pento- barbital (Pentoject). While

the pentobarbital took effect, lesions were made at the tips of two tetrodes to aid

visualisation post-mortem. After breathing and reflexes had ceased, animals were

perfused with phosphate buffered saline (PBS), then 10% formalin in order to fix brain

tissue. The brain was carefully extracted from the skull, placed in 10% formalin and

refrigerated for at least 3 days. Brains were sliced into coronal sections of width 50µm,

and mounted on glass slides. After drying, sections were stained with cresyl violet,

washed with water and ethanol, and coverslipped. Sections were imaged using a bright-

field microscope (Leica DMi8), and tetrode placement was assessed manually by the

experimenter using a rat brain atlas (Paxinos and Watson, 2013).

**Neural data analysis**

*Clustering Single Units*

To identify putative individual neurons, single unit data were clustered offline using Tint

(Axona Ltd), which performs automated clustering based on principal component

analysis (PCA) and an expectation-maximisation algorithm (Klusta-kwik,[12]). The

resultant clusters were then manually curated. Clusters were excluded based on

violation of a 2ms refractory period, or if they contained < 100 spikes. Once HD cells

were identified (see below), the data were further refined to remove possible duplicates

by identifying cells recorded in the same session from the same tetrode with PFDs of an

angular separation < $\pi/6$ radians, and keeping on the highest-rate cell. This selection

procedure resulted in a final set of 436 HD cells.

*Directional and spatial analysis*

To determine the firing directions of neurons in global space we projected the 3D head

direction data onto the horizontal plane, binned the firing rate data into 6-deg

directional bins and divided the head-direction dwell time in each bin. HD cells were

identified based on the model-fitting procedure described below, with a κ (tuning curve

sharpness) criterion of 0.5. The resultant cells had tuning curves that looked on visual

assessment to be typical HD cells.

We ran a spatial analysis in MATLAB to look for place cells or place x direction

("topodirectional") cells as have been reported previously in PoS[10]. Two-dimensional

rate maps were made for each recorded cell as follows: the environment was

discretized into 50 × 50 cm spatial bins, and occupancy-normalized firing rates were

computed by dividing the number of spikes per bin by the time spent in that bin,

excluding bins with less than 100 ms occupancy. The resulting rate maps were

smoothed using a 2D Gaussian kernel (σ = 5 cm) applied to both spike and occupancy

maps prior to division. A circular arena mask was applied by calculating the maximum

distance from the centre of mass of the trajectory and discarding bins beyond this

radius. Peak firing rate was defined as the maximum value in the rate map. Spatial

information content (bits/spike) was calculated following the Skaggs formulation, using

occupancy-normalized firing rates and bin probabilities:

$$I = \sum_i p_i \frac{\lambda_i}{\bar{\lambda}} \log_2 \left( \frac{\lambda_i}{\bar{\lambda}} \right)$$

where $p_i$ is the occupancy probability of bin $i$ and $\lambda_i$ is the mean firing rate in bin $i$

*Model-fitting analysis*

Rats moved relatively freely over the hemisphere (see video) but spent less time on the

steeper parts. They often tended to circle around the hemisphere, sometimes shuffling

around the steeper parts sideways, or else walked down the hemisphere as far as they

dared before backing upwards again. This behaviour restricted the full sampling of

directions at each location, particularly nearer the equator. This led to our developing

an alternative analysis method based on model-fitting to the firing rate profile, as
follows.

The FR-direction relationship of a tuning curve is described by a modified von Mises
distribution (below).

$$FR = we^{(\kappa cos(\alpha))} \quad \text{Equ. (1)}$$

where *FR* is the firing rate, *w* (weight) is a gain factor, κ (kappa) modulates tuning curve
sharpness and α (alpha) is the angular distance of the head from the cell's preferred
firing direction, both measured in the plane of interest. We used this relationship to
examine the directional firing of neurons over the hemisphere as follows:

1. We used instantaneous firing rates instead of spike counts as these provide a
truer picture of the activity of a cell at a given moment. These were derived from
individual spike trains at 50 Hz. The elephant electrophysiology toolbox for
Python was used [13]; specifically the function
*elephant.statistics.instantaneous_rate,* which implements the Gaussian kernel
method described in Shimazaki and Shinomoto [14]. Optimisation of kernel width,
as described in that work, resulted in an extremely large kernel width of 1630ms.
Lower kernel widths were judged to be more appropriate, with 200ms decided
on as a reasonable compromise between detail and smoothing (Supp. Fig. 2).
2. We next used the recorded data, with directions referenced to the horizontal
plane (i.e., the azimuths), to estimate the values of w and k for a given cell, using
custom Matlab code (fit_tuning_curves.m). This was done by an iterative model-
fitting process, mfit_optimize,[15] which searched the parameter space for the
best fit. It uses gradient descent to find the best combination of all three
parameters, bounding *α* within the range *-π - π*, W within the range 0.01 to 100,
and κ within the range 0.5-20. We applied this model fit process to the binned
tuning curves rather than the raw data because it produced much better fits to
the tuning curves and was also very much faster.
3. We based our HD cell selection criterion not on model fit quality, which was
often degraded by noise in the data even for obvious HD cells, but rather κ, for

which selection above a threshold of 0.5 (the starting value) produced good

results according to visual inspection, with evident peaked directional tuning,

and Rayleigh vectors almost all > 0.2. This procedure yielded 537 cells. We then

identified co-recorded cells that were detected on the same tetrode and had

PFDs with $< \pi/6$ angular separation, and removed all but the best (highest peak

rate) on the basis that they could be the same cell. The final HD cell data set was

thus n = 436.

4.  We then used the derived parameters to generate a synthetic tuning curve,

which was compared with the recorded data using Pearson's R.

5.  The alphas were then transformed according to the four models (details below)

using custom-written DAS_transform.m.

6.  Using the adjusted alphas, new tuning curve histograms were generated and

fitted with von Mises functions using mfit_optimize again using

fit_new_all_cells_tuning_curves.m to produce new tuning curve histogram

parameters κ, peak rate, and Rayleigh vector.

7.  Alphas were transformed under GA/DAS/DAH/HA, binned (60 bins, $-\pi$-\pi$-\pi$ to

$\pi$\pi$\pi$), and fitted with von Mises functions to obtain κ, w, peak rate, and Rayleigh

vector. For model comparison we evaluated the Poisson log-likelihood $\log L$ of

the raw sample-level data given each model's fitted curve, and computed AIC

($2k$ - $2\log L$) and BIC ($k\ln n$ - $2\log L$) with $k \in \{2,3\}$ and $n$ equal to the number of

samples ($dt$ = 1/50 s). Lower AIC/BIC indicate better fit; Pearson's R is reported

for interpretability.

*Transformations*

The four transformations in step 4 were as follows:

GA rule (gravity axis only): This involved projecting the nose vector onto the horizontal

plane and computing azimuth (head direction in the horizontal plane) in the standard

way.

DAS rule (dual axis; surface and gravity): this modulated the alphas according to how

much the surface normal axis of the environment local to the rat was rotated around

the GA, with the reference (zero) point being in the opposite direction from the cell's

PFD:

$$FR = we^{(\kappa \cos(\alpha_{HS} + \alpha_{SG} + \pi))} \qquad \text{Equ. (3)}$$

where $\alpha_{HS}$ is the head rotation around the SA (referenced to the projection of the GA

onto the tangent plane) and $\alpha_{SG}$ is rotation of the SA in the horizontal plane around the

GA. The $\pi$ term is added because the zero reference for calculating head $\alpha$ is taken to be

when maximal firing (local PFD) coincides with the rat facing upwards, which occurs at

the opposite side of the sphere from the cell's global PFD.

DAH rule (dual axis; head and gravity): this modulated the alphas according to how

much the DV axis of the head was rotated around the GA, with, again, the reference

(zero) point being in the opposite direction from the cell's PFD. The modified von Mises

equation describing the head-referenced dual-axis rule (DAH) is:

$$FR = we^{(\kappa \cos(\alpha_{HH} + \alpha_{HG} + \pi))} \qquad \text{Equ. (2)}$$

Where $\alpha_{HH}$ is the alpha of the head rotation in the head plane around the HA (referenced

to the projection of the GA onto the head plane) and $\alpha_{HG}$ is rotation of the HA in the

horizontal plane around the GA.

HA rule (head axis only): This ignores the tilt of the environment surface or the head and

simply updates according to left-right turns around the head axis, ignoring the head axis

azimuth. We expected this to produce poor model fits and this manipulation was

included as a null control.

*Elevation analysis*

We defined elevation as the angle of the nose vector relative to the horizontal plane.

Since HD cells sensitive to this parameter have been reported previously (ref; called

"pitch" in that experiment) we investigated the distribution of elevation angles, from -

$\pi/2$ (nose pointing down) to $+\pi/2$ (nose pointing up).

 **Supplementary Tables**

| | All cells | | | HD cells | | | t | p |
|---|---|---|---|---|---|---|---|---|
| | **All** | **PoS** | **ADN** | **All** | **PoS** | **ADN** | | |
| n | 2270 | 630 | 1640 | 436 | 141 | 295 | | |
| **Standard GA reference frame** | | | | | | | *t* | *p* |
| Pk rate (Hz) | 3.82 | 2.99 | 4.13 | 6.51 | 3.62 | 7.89 | -6.59 | $6.22 \times 10^{-11}$ |
| R vector | 0.18 | 0.19 | 0.17 | 0.47 | 0.50 | 0.45 | 2.85 | $4.6 \times 10^{-3}$ |
| kappa | - | - | - | 2.15 | 3.06 | 1.73 | 5.37 | $1.26 \times 10^{-3}$ |
| **DAS transformed** | | | | | | | | |
| Pk rate (Hz) | 4.06 | 3.13 | 4.42 | 7.08 | 3.74 | 8.68 | -6.84 | $1.32 \times 10^{-11}$ |
| R vector | 0.18 | 0.19 | 0.17 | 0.49 | 0.52 | 0.48 | 2.39 | $8.53 \times 10^{-3}$ |
| kappa | 0.91 | 1.19 | 0.81 | 2.30 | 3.16 | 1.88 | 5.29 | $9.96 \times 10^{-3}$ |
| **DAH transformed** | | | | | | | | |
| Pk rate (Hz) | 4.28 | 3.24 | 4.68 | 8.10 | 4.24 | 9.94 | -7.03 | $4.08 \times 10^{-12}$ |
| R vector | 0.19 | 0.20 | 0.19 | 0.54 | 0.56 | 0.53 | 1.51 | NS |
| kappa | 1.20 | 1.61 | 1.05 | 3.57 | 4.88 | 2.94 | 5.82 | $5.78 \times 10^{-5}$ |
| **HA transformed** | | | | | | | | |
| Pk rate (Hz) | 3.78 | 2.96 | 4.10 | 4.61 | 2.14 | 5.79 | -7.94 | $8.60 \times 10^{-8}$ |
| R vector | 0.11 | 0.11 | 0.10 | 0.18 | 0.19 | 0.18 | 1.01 | NS |
| kappa | 0.58 | 0.55 | 0.59 | 0.68 | 0.58 | 0.73 | -1.04 | NS |

Supplementary table 1 Supplementary Table 1: Basic firing properties in standard horizontal coordinates and under the three transformations, compared between cell types and brain regions

 **Supplementary table 2**

| Variable | Model | Mean | Comparison | $t$ | $p$ | $d$ |
|---|---|---|---|---|---|---|
| Peak rate | GA | 5.59 | GA vs DAS | 6.00 | .0000 | 0.41 |
| | DAS | 6.17 | GA vs DAH | `6.57 | .0000 | 1.76 |
| | DAH | 7.18 | DAS vs DAH | 10.58 | .0000 | 1.57 |
| R. vector | GA | 0.48 | GA vs DAS | 8.09 | .0000 | 0.55 |
| | DAS | 0.50 | GA vs DAH | 24.45 | .0000 | 1.66 |
| | DAH | 0.55 | DAS vs DAH | 16.37 | .0000 | 1.11 |
| kappa | GA | 2.77 | GA vs DAS | 1.71 | .0878 NS | 0.12 |
| | DAS | 2.91 | GA vs DAH | 17.02 | .0000 | 1.15 |
| | DAH | 4.19 | DAS vs DAH | 15.31 | .0000 | 1.04 |

476 Supplementary Table 2: Tuning curve param
490 eters under the different reference frame transformations, and pairwise comparisons

492 **Supplementary table 3**

| Variable | Model | Best-model count |
|---|---|---|
| AIC | GA | 108 |
| | DAS | 112 |
| | DAH | 199 |
| | HA | 17 |
| BIC | GA | 108 |
| | DAS | 112 |
| | DAH | 199 |
| | HA | 17 |
| logL | GA | 108 |
| | DAS | 112 |
| | DAH | 199 |
| | HA | 17 |
| R | GA | 22 |
| | DAS | 32 |
| | DAH | 382 |
| | HA | 0 |

**Supplementary Table 3:** Model comparisons: Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), log likelihood (logL) and Pearson's correlation coefficient (R).

     **Supplementary Figures**

**Supplementary Figure 1: Berry-Hannay errors following exploration in 3D**. (a) A rat
travels over a sphere, carrying its internal, local reference frame (small colored arrows) with it. By
the end of the journey, having returned back to a horizontal alignment, its local reference frame
has dissociated from the global one (large colored arrows). (b) The coordinate reference frames
from (a) expressed relative to the rat (egocentrically). The first turn (iii) rotates the rat in its internal
coordinate frame (it is now facing the "yellow" direction; inner arrows) but the local frame
remains aligned to the global frame (outer arrows). The second turn (v) both rotates the rat within
its own frame *and* the local frame in the global frame, the result being a dissociation. The
dissociation occurs because in 3D, rotation can indirectly occur in a given plane by rotations
occurring in the orthogonal planes, which a simple planar ring attractor would not detect.

**Supplementary Figure 2: IMU and classic LED-based head direction recording methods show a high degree of agreement**. Bland-Altman analysis of paired recording with LEDs and IMU: blue points are individual readings at 50Hz. Left shows LED and IMU readings plotted against one another, with line of best fit, $R^2$ value, and number of samples. Right shows mean difference between the two sets of readings.

**Supplementary Figure 3: Instantaneous firing rate calculation** (a) Left: cluster space for a

hemisphere session, showing a HD cell in red, other units in blue. Right, waveforms for the same cell,

highest amplitude waveform in red. Plots show waveform means and standard deviations. Peak, width,

and signal to noise ratio (SNR) also shown. (b) Gaussian smoothing to generate the FR profile. Example

spiketrain (red vertical lines) and estimated instantaneous firing rates calculated using gaussian

kernels (Shimazaki and Shinomoto, 2010). Various sigma values (kernel widths) are shown in different

colours. Sigma value used in final analysis was 200ms (purple)

**Supplementary Figure 4: Pairwise comparisons of the four major tuning curve parameters** for individual neurons (n = 436). The dotted black line shows the linear regression relative to the y = x line (solid black line).

551  **Supplementary Figure 5: Calculation of elevation tuning.** (a) Distribution of nose elevation
552  angles. (b) Firing rates relative to the cell's mean, for all 2270 cells as a function of nose elevation
553  (mean +/- s.e.m)

**Supplementary Figure 6:** Tuning curves for the 40 cells showing best directionality for the HA reference frame. Each subplot shows (L:R) GA, DAS, DAH and HA transformations (the latter in orange). These cells could potentially provide a possible mediating signal for the local axis rotation (independent of axis azimuth).

556
557
558
559

## References

**References**

1.    Taube, J. S. Head direction cells and the neurophysiological basis for a sense of direction.

*Prog Neurobiol* **55**, 225–256 (1998).

2.    Skaggs, W. E., Knierim, J. J., Kudrimoti, H. S. & McNaughton, B. L. A model of the neural

basis of the rat's sense of direction. *Adv Neural Inf Process Syst* **7**, 173–180 (1995).

3.    Zhang, K. Representation of spatial orientation by the intrinsic dynamics of the head-

direction cell ensemble: a theory. *J Neurosci* **16**, 2112–2126 (1996).

4.    Khona, M. & Fiete, I. R. Attractor and integrator networks in the brain. *Nat. Rev. Neurosci.*

**23**, 744–766 (2022).

5.    Jeffery, K. J., Wilson, J. J., Casali, G. & Hayman, R. M. Neural encoding of large-scale three-

dimensional space-properties and constraints. *Front Psychol* **6**, 927 (2015).

6.    Page, H. J. I., Wilson, J. J. & Jeffery, K. J. A dual-axis rotation rule for updating the head

direction cell reference frame during movement in three dimensions. *J. Neurophysiol.* **119**,

192–208 (2018).

7.    Laurens, J. & Angelaki, D. E. The brain compass: A perspective on how self-motion updates

the head direction cell attractor. *Neuron* **97**, 275–289 (2018).

8.    Williams, M. & Jeffery, K. Video: rat on sphere. (2025).

9.    Stackman, R. W., Tullman, M. L. & Taube, J. S. Maintenance of rat head direction cell firing

during locomotion in the vertical plane. *J. Neurophysiol.* **83**, 393–405 (2000).

10.   Cacucci, F., Lever, C., Wills, T. J., Burgess, N. & O'Keefe, J. Theta-modulated place-by-

direction cells in the hippocampal formation in the rat. *J Neurosci* **24**, 8265–8277 (2004).

11.   Grieves, R. M. *et al.* The place-cell representation of volumetric space in rats. *Nat.*

*Commun.* **11**, 789 (2020).

12.   Harris, K. D., Kadir, S. N., & Goodman, D. F. M. Masked KlustaKwik. https://klusta-

team.github.com/klustakwik. (2013).

13.   Denker, M., Kern, M. and Richter, F. Elephant 1.1.0.

https://doi.org/10.5281/zenodo.10955610 (2024) doi:10.5281/zenodo.10955610.

14.  Shimazaki, H. & Shinomoto, S. Kernel bandwidth optimization in spike rate estimation. *J*

*Comput Neurosci* **29**, 171–182 (2010).

15.  Gershman, S. 'MFIT'. https://github.com/sjgershm/mfit (2022).

591