

Neuron, Volume 109

Supplemental information

**Frequency of theta rhythm is controlled
by acceleration, but not speed, in running rats**

Emilio Kropff, James E. Carmichael, Edvard I. Moser, and May-Britt Moser

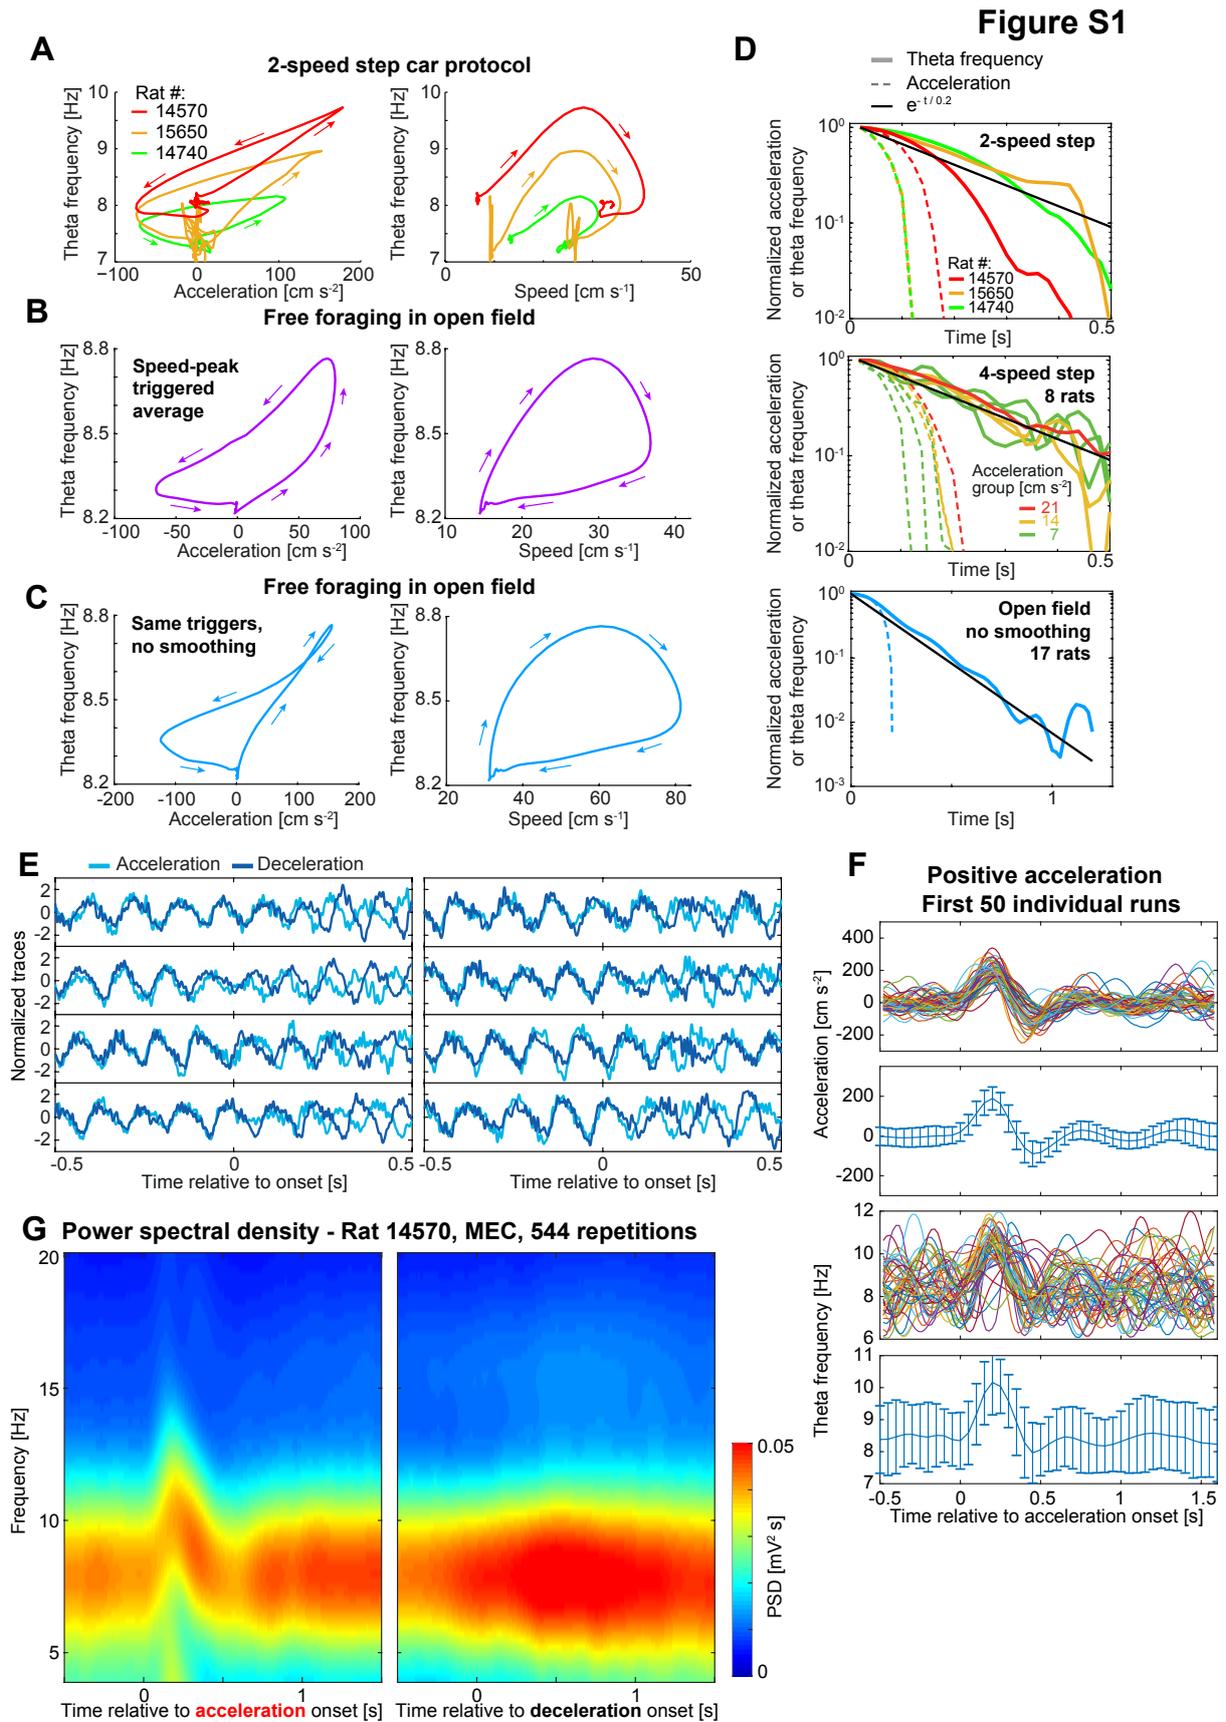


Figure S1. Instantaneous rise and slow decay of theta frequency during positive acceleration peaks (related to Figure 1).

(A) Curves of theta frequency vs. acceleration (left) or speed (right) for the data in Figure 1B (red; rat 14570) and for 2 other rats trained on the 2-speed step protocol (orange and green; rats 15650 and 14740). Every curve is the average over positive acceleration events pooling together all sessions. The speed step sizes were 30 cm s^{-1} (rat 14570), 20 cm s^{-1} (rat 15650) and 14 cm s^{-1} (rat 14740). Arrows indicate the direction of time flow. Note that all 3 rats developed the strategy of accelerating slightly more than required, and adjusted to the car speed by applying a correction in the form of a small deceleration trough immediately following the acceleration peak. This small deceleration was not part of the car protocol. While theta frequency and acceleration peak together in the top plot, the decay of theta frequency after the peak is slightly slower than that of acceleration, creating a hysteresis effect. Note in the bottom plot that speed cannot explain the dynamics of theta frequency during this kind of event. While theta frequency increases with speed during the first half of the event, it then decreases back to baseline as animals reach maximum speed. (B) Similar curves for speed-peak triggered averages of open field data (Figure 5E, top; pool of all open field data). Hysteresis is also present here. Note that, in contrast with (A), the peak of theta frequency occurs slightly after the peak of acceleration, resulting in a smooth-edged curve. There is, however, an important difference in the treatment of both datasets. While in (A) data was averaged across trials with no smoothing, data in (B), as for all other open field analyses, had no identical trials to average across, so smoothing along the temporal dimension was necessary to avoid high noise in the estimates of instantaneous speed, acceleration and frequency (Figure S5; STAR Methods). (C) To account for effects of smoothing, the plots in (B) are repeated here using the same windows for averaging (Figure 5E, top) but without smoothing acceleration and frequency. Interestingly, the frequency and acceleration peaks synchronize and the rounded shape disappears, while the slow decay in theta frequency creating hysteresis remains for lower values of theta frequency. This shows that, while the small misalignment between acceleration and frequency peaks observed in (B) (and also in Figure 5E) is an artifact of smoothing along the temporal dimension, the slow decay of theta frequency after peaking is not. Note that the curves of theta frequency vs. speed are qualitatively similar throughout (A) to (C), because, unlike acceleration, the relationship between speed and theta frequency is not affected by temporal precision. (D) Delayed decay of theta frequency compared to acceleration. Acceleration is normalized by its positive peaking value and drawn until it becomes negative. Theta frequency is normalized to be 0 at its minimum and 1 at its peaking value. The decay dynamics toward 0 (acceleration) or minimum (theta frequency) is plotted in logarithmic scale. Top: 2-speed step data shown in (A). Center: 4-speed step data shown in Figure 1D (only positive acceleration). Bottom: unsmoothed peak-triggered averages of open field data shown in (C). In all cases the decay of acceleration is faster than that of theta frequency. In most cases, the decay of theta frequency can be modelled by an exponential with a time constant of 0.2 s (black line). (E) As Figure 1C but for examples of individual runs (one for every row/column combination) with accelerating (light blue) or decelerating (dark blue) trajectories. (F) Instantaneous acceleration (top) and theta frequency (bottom) for the first 50 individual acceleration episodes in Figure 1B. Each panel shows individual runs plotted with different colours (top) and distribution (bottom; mean \pm s.d.). (G) Average spectrograms using the Fast Fourier Transform (MATLAB function `spectrogram()`) for the data in Figure 1B corresponding to acceleration (left) or deceleration (right) events. Although the temporal and frequency resolution of this method are

poorer than the one used in Figure 1B, the increase in theta frequency with positive acceleration and the lack of modulation by negative acceleration or speed are clearly observed.

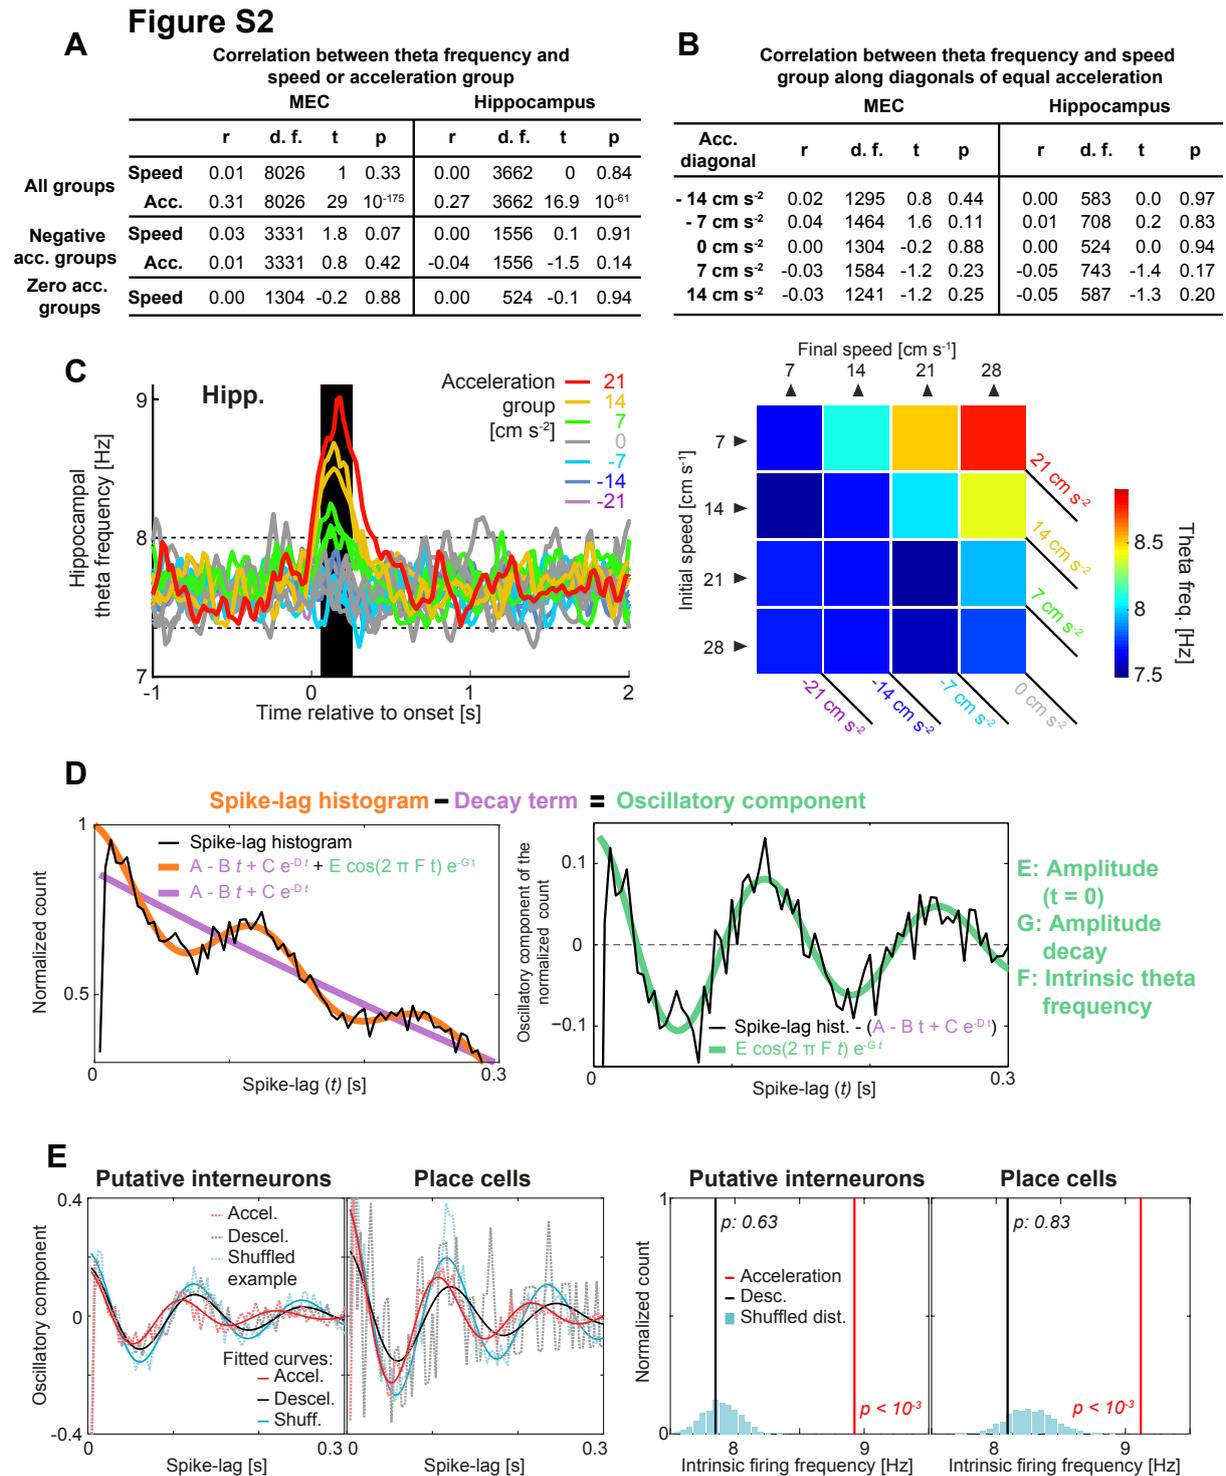


Figure S2. LFP and rhythmic spiking theta frequency in recordings from MEC and hippocampus (related to Figure 1).

(A) Table showing Pearson correlations between theta frequency on one side and speed or acceleration group on the other (statistics and p-values correspond to two-tailed t-tests for the null hypothesis of zero regression), based on pooled data of acceleration events in the 4-speed protocol. As in Figure 1E, theta frequency for each event is averaged over the whole acceleration segment (60 to 260 ms after the onset; black area in Figure 1D). Events are

organized according to the initial and final speed of the transition (Figure 1E). The acceleration of each event over a 1 s window is defined as the difference between final and initial speed, so that descending diagonals in Figure 1E define ‘acceleration groups’, with equal acceleration regardless of speed (see panel C, right). Similarly, the speed of each event is defined as the average between initial and final speed, so that ascending diagonals in Figure 1E form groups of equal speed regardless of acceleration. (B) Pearson correlation between theta frequency and speed along all diagonals of equal acceleration in Figure 1E. (C) Same plots as for MEC in Figure 1D, bottom (left) and Figure 1E (right) but here for hippocampal data. Diagonals of invariant acceleration (colour code as in Figures 1D and 1F) are indicated. (D) Method for estimating intrinsic theta frequency from spike-lag histograms. Population spike-lag histogram as in Figure 2A, obtained using data extracted exclusively from windows of 0.5 s after the onset of each acceleration or deceleration event and averaging over all selected neurons. Without normalization (introduced to improve visualization) it would represent an estimation of the probability of a cell in the population to spike at time t given that it has spiked at time 0 (STAR Methods). The spike-lag histogram shows an oscillation at theta frequency (intrinsic firing frequency) mounted on top of a decaying baseline. Left: the histogram (black) is fitted (orange) with a function containing a decay term (purple) and an oscillatory term (green). There are 7 fitting parameters, 4 of them describing the decay and 3 describing the oscillatory modulation. A and B describe a linear decay in the spike-lag histogram while C and D describe an exponential decay. In general terms the decay is a mixture of the natural decay due to the neural activity pattern (Climer et al., 2015) and an additional component due to the reduced size of the windows from which data is extracted (since windows are 0.5 s long, all spike-lag histograms necessarily reach a value of 0 at time 0.5 s). Since data in different cases are better fit by either a linear or an exponential decay, we chose to include both terms in the fit. In this way, a single equation is used to fit all data throughout the paper. Among parameters describing the oscillatory modulation, E describes the initial amplitude of the oscillatory component at time $t = 0$, G describes the exponential decay of the amplitude and F is the quantification of intrinsic theta frequency. Right: the decaying term can be subtracted from either the actual spike-lag histogram or the fit to obtain the oscillatory component. This has the only purpose of visually enhancing the demonstration of theta modulation. Note that, due to the limited amount of data inside small acceleration windows, we always apply this method to the pooled data of all available neurons within a given condition. For this reason, we get a single value of intrinsic frequency for each condition, rather than a distribution of values. (E) Application of method in D on hippocampal data. Plots as for MEC data in Figures 2A (left subpanel) and 2B (right subpanel), but here for hippocampal putative interneurons and place cells (as indicated) recorded in the 2-speed step sessions of Figure 1B. In both cases a significant increase in intrinsic firing frequency for acceleration (but not deceleration) episodes was observed, as in data from MEC. Putative interneurons: similarly, positive acceleration events were associated with a high intrinsic theta frequency (8.9 Hz, right-tail $p < 10^{-3}$) but not negative acceleration events (7.85 Hz, $p = 0.63$). Place cells: positive acceleration events (red) were associated with a high intrinsic theta frequency (9.12 Hz, right-tail $p < 10^{-3}$ by comparison with 1000 shuffles) but not negative acceleration events (black, 8.09 Hz, $p = 0.83$).

Figure S3

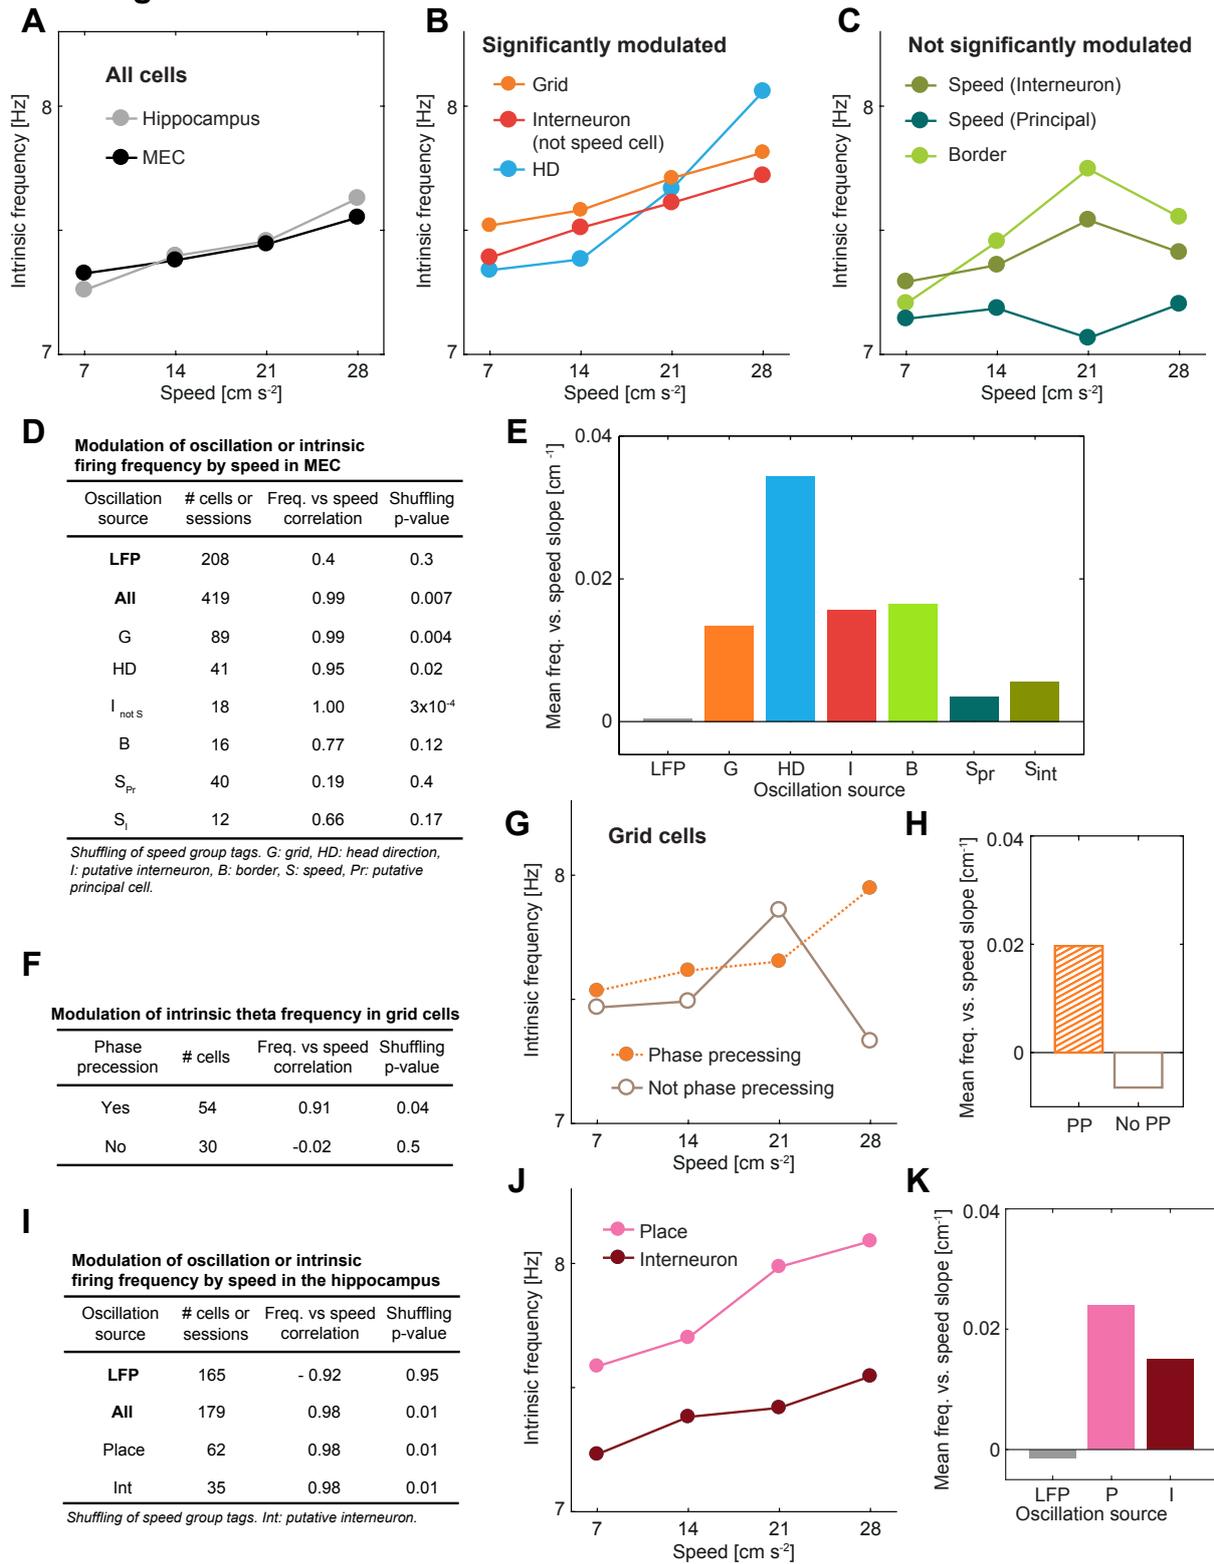


Figure S3. Intrinsic firing frequency is mildly modulated by running speed in some cell types, as expected when the cells display phase precession (related to Figure 2).

Phase precessing cells in the hippocampus (O'Keefe and Recce, 1993) and entorhinal cortex (Hafting et al., 2008) encode the position of the animal within a spatial field in their theta firing phase. As rats advance, the firing phase of place cells and some grid cells precedes along the

theta cycle, giving rise to a firing period slightly shorter than the LFP theta period and thus to an intrinsic firing frequency slightly higher than LFP theta frequency (Jeewajee et al., 2008). Since firing phase codes for position, faster runs imply stronger precession, and thus intrinsic firing frequency increases with running speed (Geisler et al., 2007; Jeewajee et al., 2008; O'Keefe and Recce, 1993). Here we assessed the relationship between intrinsic firing frequency and speed for different cell types in data from the 4-speed step protocol. To quantify intrinsic firing frequency we used the same method as in Figures 2 and S2, but instead of using windows of 0.5 s after acceleration onsets, data for spike-lag histograms were extracted from windows of variable length (typically > 1 s) corresponding to uniform motion segments, excluding 1 s after every acceleration or deceleration onset. (A) Intrinsic firing frequency as a function of speed group for all cells recorded in the 4-speed step protocol in the MEC (black) and the hippocampus (grey). A linear dependence with speed is observed (“All” in panels D and I), with small overall variations compatible with the observations in Figure 2C. (B) and (C) Similar to (A) but grouping neurons of the MEC by functional cell type (assessed from a supplementary open field session; STAR Methods). Cell types with significant and non-significant modulation by speed are plotted in (B) and (C), respectively. (D) Quantification and significance of the modulation of intrinsic theta frequency by running speed in the MEC, sorted by cell types. The significance was obtained by direct comparison with 10000 shuffles of speed group tags, which indicate for every segment of trajectory the corresponding constant running speed (7, 14, 21 or 28 cm s^{-1} ; criterion: Pearson correlation between frequency and speed group above the 95% of the shuffled distribution, i.e. right tail $p < 0.05$). The table also shows the quantification of the speed modulation of LFP theta oscillations (Figure 1G). Column 1: oscillation source (LFP or intrinsic firing cell type). Column 2: Number of sessions (for cell types, 1 session per cell). Column 3: correlation between the 4 values of fitted intrinsic frequency and the corresponding speed groups (7, 14, 21 and 28 cm s^{-1}). Note that for every cell type a single fit is performed with the pool of all the data falling into each speed group category. This means that the Pearson correlation is performed over only 4 data points, easily achieving extreme values. For this reason, significance was assessed using a high number (10000) of shuffles. Column 4: right-tailed p-values obtained as the fraction of 10000 speed-group tag shuffles with higher correlation than the non-shuffled value. G: grid cells, HD: head direction cells, I: non speed-modulated putative fast spiking cells, B: border cells, S_{pr} : putative principal speed cells, S_{int} : putative fast spiking speed cells. (E) Magnitude of the speed modulation of theta frequency for each oscillation source, measured as the mean difference in theta frequency between adjacent speed groups. (F) As (E) but grouping grid cells according to whether or not they exhibited phase precession (Hafting et al., 2008). Phase precession for every cell was assessed as a significant correlation ($p < 0.05$) between theta spiking phase and within-field firing position in the car experiment, pooling together spikes from all spatial fields. We observed no significant speed modulation of intrinsic firing frequency for non-phase precessing grid cells (Pearson correlation: - 0.02, p : 0.5 by comparison with 10000 shuffles of speed group tags). (G) and (H) As (A) and (D) respectively, but applied to phase-precessing (PP) and non phase-precessing (No PP) grid cells. (I), (J) and (K) As (D), (B) and (E), respectively, but for hippocampal LFP and functional cell types. For consistency, place cells, like other cell types, were classified according to their behaviour in a separate open field session, but since active cells in the open field were sometimes silent in the car session, an

additional condition of a mean firing rate of at least 0.2 Hz in the car session was required for inclusion in the analysis. Both cell types (place cells and putative fast spiking cells) exhibited a significant speed modulation.

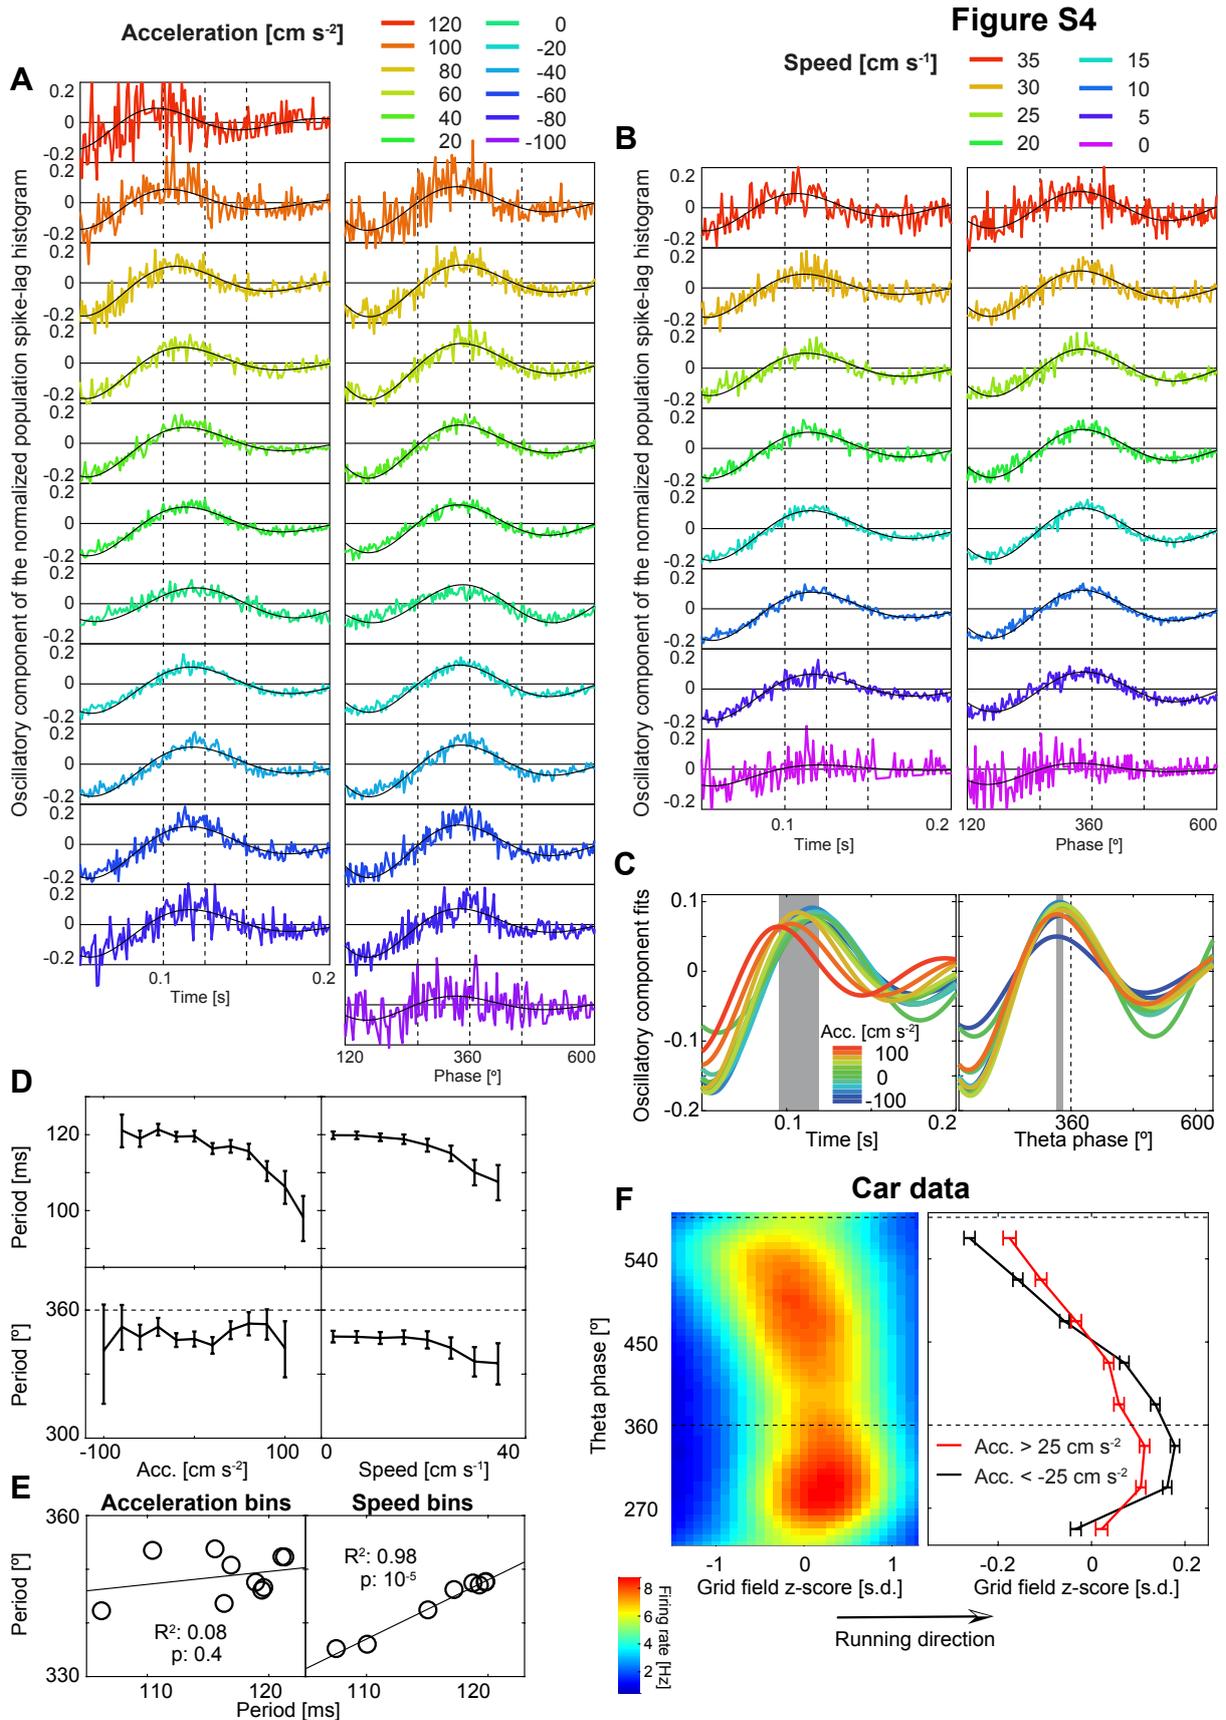


Figure S4. Grid cell phase precession explains modulation of intrinsic firing frequency by speed but not by acceleration (related to Figure 2).

We focused on grid cells to understand if the modulation of intrinsic theta frequency by speed interacts with the modulation by acceleration, and to what extent both types of modulation can be explained by the phenomenon of phase precession. We examined spike-lag histograms of the pool of all grid cells recorded in open field experiments (518 cells). Spike-lag histograms (as in Figures 2 and S2D) were built by using either the timestamps of spikes or the unwrapped LFP theta phase of the spikes. (A) Zoom of the first peak in the oscillatory component (as in in Figure 2A, right) constructed from grid cell spikes occurring at different accelerations. Acceleration is colour coded and shown in successive rows. Left: temporal spike-lag histograms. Right: LFP theta phase spike-lag histograms. Note that theta phase is unwrapped, so that it is a monotonically increasing variable that extends beyond 360° . Fits with a decaying oscillatory function (Figure S2D) are shown in black. Note that the first peak in the temporal oscillatory component moves to the left as acceleration increases but the peak in the LFP theta phase oscillatory component is acceleration-invariant. Only curves that could be properly fitted by the oscillatory decaying function are shown (goodness of fit determined by $R^2 > 0.8$). (B) As (A) but dividing the data into ranges of speed rather than acceleration. (C) Summary of fitted curves in (A), with the same colour-code for acceleration. The grey area shows the range in time (left) or LFP theta phase (right) where maxima occurred. Note that the position of maxima vary in the temporal domain (left) but are rather phase locked at a value slightly below 360° (dashed line) in LFP phase (right). (D) Fitted intrinsic theta period of temporal (top) and LFP theta phase (bottom) spike-lag histograms for data divided into acceleration bins (left) and speed bins (right), corresponding to data in (A) and (B) respectively. Period was calculated as the inverse of the fitted frequency (Figure S2D), and is plotted here instead of frequency to make LFP phase plots more intuitive. Theta-related peaks occurred earlier in time for increasing positive acceleration, consistent with higher intrinsic firing frequency (Pearson correlation between fitted period and acceleration bins: $r = -0.87$, $t(9) = -5.3$, $p: 5 \cdot 10^{-4}$), but at an acceleration-invariant LFP theta phase (mean: 349° , s.d.: 4.5° ; Pearson correlation with acceleration bins: $r = 0.16$, $t(9) = 0.48$, $p: 0.64$). (E) Intrinsic firing period in LFP theta phase as a function of the intrinsic firing period in time across acceleration bins (left) or speed bins (right). Consistent with the phenomenon of phase precession, 98% of the temporal variability across speed bins can be explained by an advance in spiking phase relative to LFP theta (Pearson correlation: $r = 0.99$, $t(6) = 16.8$, $p: 3 \cdot 10^{-6}$). In contrast, no such relationship is observed across acceleration bins (Pearson correlation: $r = 0.28$, $t(8) = 0.82$, $p: 0.44$). These results suggest that phase precession and acceleration are two dissociated sources of intrinsic frequency modulation, acceleration being the strongest one. In this light, ideas of theta frequency conveying a speed code that can be transformed into position need to be revised to include the effects of acceleration, or favored by models that do not rely on such a relationship to account for path integration (McNaughton et al., 2006) or phase precession (Mehta et al., 2002; Tsodyks et al., 1996). (F) Left: Phase precession plot, presented as a colour coded distribution of spikes in theta phase and space (represented by grid field z-score; rats run from negative to positive values) for the pool of all grid cells and fields recorded in the 2-speed step and 4-speed step car experiments. The dashed lines indicate the range of theta phases used to quantify phase precession, where phase varies monotonically with position. Right: Spikes with high positive (red) or negative (black) acceleration values (threshold: 25 cm s^{-2}) are selected. Phase precession plots for each condition are shown as a distribution (mean \pm s.e.m.) of z-score

for different theta phase ranges, after subtracting from the pool of all spikes in each condition its mean z-score to improve visual comparison. As expected from phase precession taking place in the form of faster but equally sized phase steps, positive acceleration data exhibited a steeper relationship between theta phase and position. To quantify the difference, we modelled the theta phase of spikes (inside the range between dashed lines) as a linear function of z-score and acceleration condition (sign). We found an interaction between these predictors, indicating a significantly steeper phase precession for positive acceleration (z-score coefficient: -0.075, $t(89963)$: -16.2, p : 3×10^{-59} ; z-score*positive_acceleration_condition coefficient: -0.083, $t(89963)$: -12.1, p : 2×10^{-33}).

Figure S5

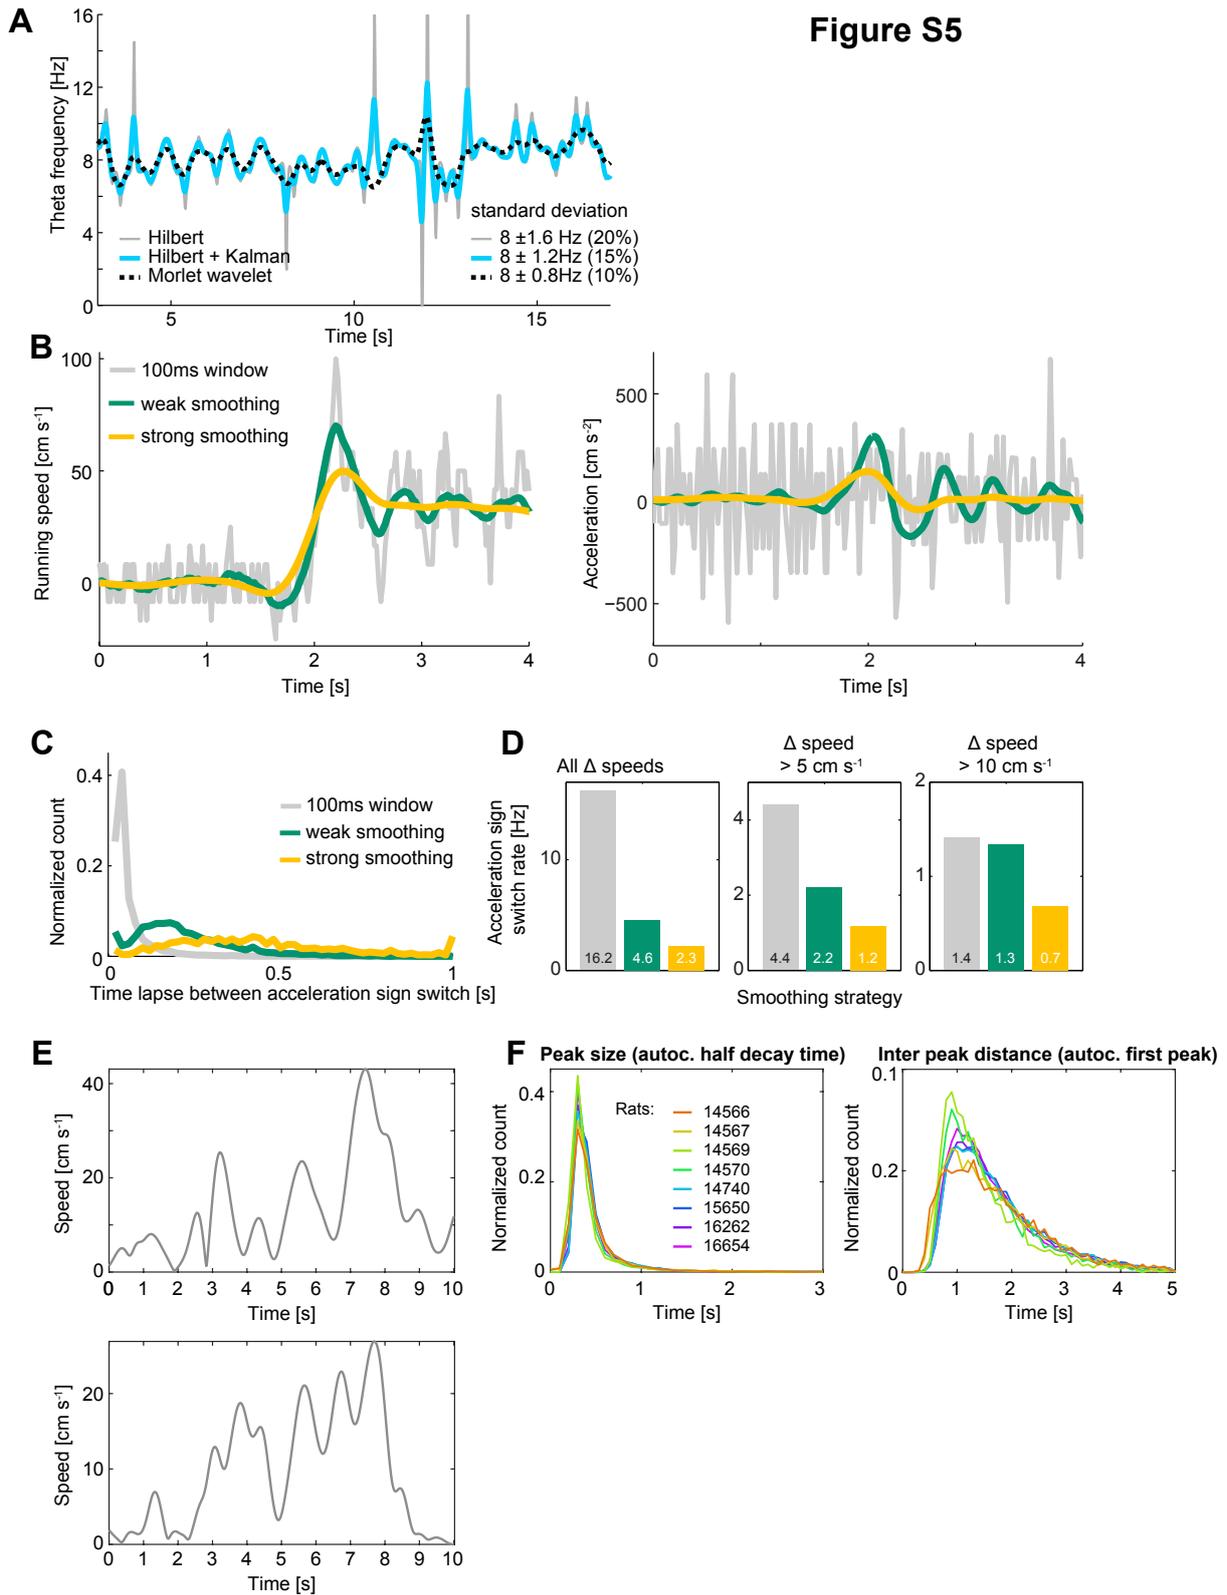


Figure S5. Estimations of instantaneous theta frequency and derivatives of position for open field data (related to Figure 3).

In contrast with car data, open field analyses cannot profit from averaging over identical trials. Instead, since all analyses are based on instantaneous measures of theta frequency, speed and

acceleration, temporal smoothing becomes necessary. (A) Representative example of different strategies for estimating theta frequency applied to 20 s of open field data, depicting a compromise between the detection of fast theta frequency variations on one hand and the estimation of unrealistic values on the other. The theta frequency component of the LFP is obtained through bandpass filtering (6-12 Hz). In the first strategy, the Hilbert transform is applied to obtain the instantaneous phase. The instantaneous slope of the unwrapped phase (i.e. the frequency of the oscillation) is estimated from the phase difference between neighboring timestamps (grey). The Hilbert transform assumes a single spectral component in the oscillation and the frequency derived from it in this direct way has strong deviations from the theta band when noise or secondary frequency components are not negligible. In order to avoid this, the second strategy makes use of a Kalman filter and smoother to obtain the instantaneous theta frequency as the derivative of the phase (light blue). This method still captures fast changes in theta frequency but avoids strong deviations from the theta band. Thus, it is used for open field analyses throughout the paper. An alternative strategy based on Morlet Wavelets rather than the Hilbert transform is also shown (black dashed line). It yields similar results but only captures relatively slow variations of the frequency and is therefore not used. (B) A similar trade-off in smoothing parameters is used for analyzing position and its derivatives. Smoothing should ideally reduce noisy variations in speed and acceleration, but still capture instantaneous variations related to body movements. We compared different strategies for smoothing speed (left) and acceleration (right) on a well-established body movement: a single acceleration event from 0 to 36 cm s^{-1} in the car, taking place roughly at $t = 2 \text{ s}$. The first strategy consists of estimating the x and y components of the velocity and acceleration vectors using the differences in position between timestamps 100 ms apart (grey). The result, as expected, is noisy, with higher noise levels for acceleration, since it is a higher order derivative of position. Alternative strategies consisted of using a Kalman filter and smoother with either weak (dark green) or strong (orange) smoothing. Note that the weak smoothing strategy reveals fast oscillations after the acceleration onset, possibly caused by head movements rather than body movements. Similar oscillations in instantaneous acceleration are observed, although in a much weaker form, when cross trial averaging is used instead of temporal smoothing (Figure 1B, top). If these oscillations corresponded to head swings, they would be greatly overestimated since the tracking utilizes LEDs some centimeters away from the head. In a rotational movement, speed and acceleration are proportional to the distance from the axis of rotation, so they would be several-fold higher measured at the level of the LEDs compared to, for example, at the level of the rat vestibular system. In contrast, body movement or translational acceleration would not suffer from this caveat, only present in the estimation of speed and acceleration during rotational movements. The fact that the observed oscillations do not seem to have a clear theta frequency counterpart (Figure 1B), although elsewhere acceleration modulates theta frequency, argues in favour of this interpretation. Taking this into account, we chose to be conservative and used the strong smoothing for open field analyses throughout the paper. This is a trade-off that allows for noise filtering while still capturing sudden instantaneous body accelerations such as the step-like transition of speed in (B). The same strong smoothing parameters were used in a previous article (Kropff et al., 2015). (C) To further demonstrate the convenience of our choice for smoothing position and its derivatives, we plot a histogram of the time lapse between switches of acceleration sign during an open field session recorded on the same day

as in (B), for the same 3 smoothing strategies (same colour code). Note that for the 100 ms window or the weak smoothing strategies, a significant number of switches took place after a single timestamp (20 ms), which corresponds to noise rather than body dynamics. (D) Based on (C), we calculated the acceleration sign switch rate during the whole session for all switches (left), for switches involving a speed difference of at least 5 cm s^{-1} (center) or 10 cm s^{-1} (right). In each case the mean rate is indicated. (E) Two representative examples of 10 s windows of free foraging behaviour. Speed is organized into concatenated peaks rather than alternating uniform motion levels. (F) The pool of all free foraging data for 8 rats (colour code) was divided into 10 s windows as in (E). For each window, the temporal autocorrelogram in bins of 20 ms was obtained. Left: For each animal, distribution of the half decay time, indicating the typical temporal span of speed peaks. Overall mean and s.d.: $0.46 \pm 0.3 \text{ s}$. Right: For each animal, distribution of the location of the first peak in the autocorrelogram, indicating the typical distance between consecutive peaks. Overall mean and s.d.: $1.69 \pm 0.8 \text{ s}$.

Figure S6

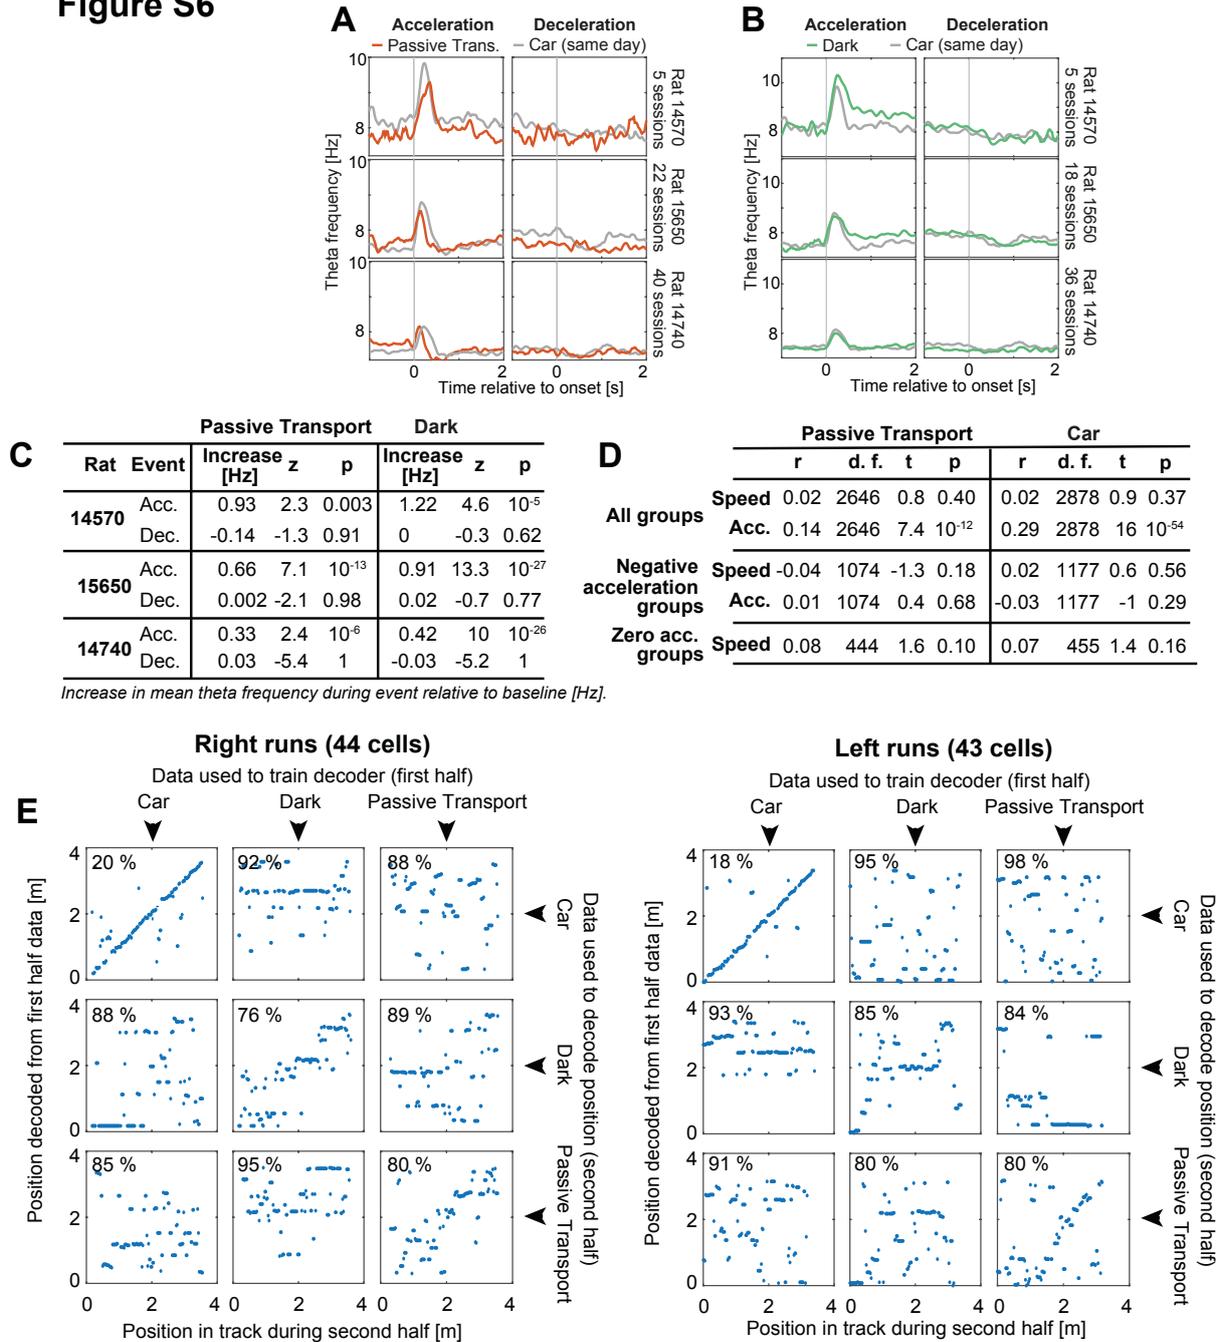


Figure S6. Visual or self-motion cue deprivation in the car does not impair modulation of theta frequency by acceleration but produces global remapping of grid cells (related to Figure 4).

Recent evidence points to a dissociation between the modulation by kinematic variables of neural activity and LFP in MEC (Hinman et al., 2016; Winter et al., 2015), as well as between LFP and the phase or rate codes of hippocampal place cells (Aghajan et al., 2015; Ravassard et al., 2013). We here provide further evidence in support of this idea, showing manipulations that preserve the modulation of theta frequency by acceleration but make grid cells remap into representations of space with lower stability and spatial information. (A) Mean theta frequency relative to acceleration (left) or deceleration (right) onsets for the Passive Transport (red) and

same-day Car (grey) sessions in Figure 4A. Rows correspond to each of the 3 rats trained in the 2-speed step protocol (step size for each rat, from top to bottom: 30, 20 and 14 cm s⁻¹). (B) Mean theta frequency relative to onset of acceleration (left) or deceleration (right) in Dark (green) and same-day Car (grey) sessions of Figure 4A. (C) Statistics for data in A and B indicating a significant increase in theta frequency during acceleration events (defined as in Figures 1D, 1E and S2) but not during deceleration events. The increase is measured relative to a shuffled baseline obtained from constant running episodes before and after the events. The mean increase (in Hz) with respect to baseline and the statistics (Mann-Whitney U test z and right-tail p value) are shown for Passive Transport (left) and Dark (right) sessions. (D) Statistics of correlations between theta frequency and speed or acceleration groups (as in Figure S2C) for Passive Transport vs. same-day Car data in sessions recorded using the 4-speed protocol (Figure 4B). (E) Passive Transport and Dark experiments in the bottomless car produce global remapping of grid cells, with new maps exhibiting lower stability and spatial information (Figure 4E). These measures are indicative of a poor spatial code, and should ultimately impair the decoding of position out of grid cell firing. In order to understand if this happened as a consequence of the remapping produced by the Dark and Passive Transport conditions, we studied the decoding of position in each condition and across conditions. For this purpose, rat 14740 was selected because it was the one with the highest number of grid cells recorded on days that included all 3 variants of a car protocol (2-speed step protocol in Car, Dark and Passive Transport conditions). Right (right panel) and left (left panel) runs were treated independently. Since the car ensures exact repetitions of the protocol, cells from all recording days were pooled together, resulting in 44 right run cells and 43 left run cells (directional grid cells were included only in the direction in which they were active). The first half of the data in each session was used to train the decoder, while the second half was used to test it. Population activity vectors for the first and the second halves of each session were constructed for all 2.5 cm bins along the track. For any given bin, the training population vector describing the activity of the network in that position was correlated with all testing population vectors, and the decoded position was defined by the bin with the testing population vector that maximized correlation. Decoding error (Figure 4E, bottom) was computed as the percentage of bins with a distance between decoded and actual position larger than 10 cm. The two main panels in this figure are composed of 9 subpanels each, corresponding to all combinations of 3 sets of training data (Car, Dark, Passive Transport) and 3 sets of testing data, as indicated. Every subpanel shows decoded position vs. actual position for all bins in the track. The decoding error is indicated in the top left corner. For both left and right runs, only the data from the regular car sessions could be decoded with an error lower than 20%, while all other combinations, including those in the diagonal (where training and testing data were extracted from the same sessions), had errors above 75%. In sum, although theta frequency modulation by acceleration was preserved in the Dark and Passive Transport conditions, grid cells remapped acquiring spatial representations with low stability and spatial information, from which position could be decoded with substantially lower accuracy. This supports the idea that LFP oscillations and neural activity in the MEC are dissociated pathways for the representation of kinematic information (Hinman et al., 2016; Winter et al., 2015). One of the studies observing such a dissociation (Winter et al., 2015) reported a disruption of the speed-theta frequency relationship in a Passive Transport condition that differed from ours in several

respects, e.g. non-identical spatial trajectories were compared, but the study did not investigate the relationship between acceleration and theta frequency.

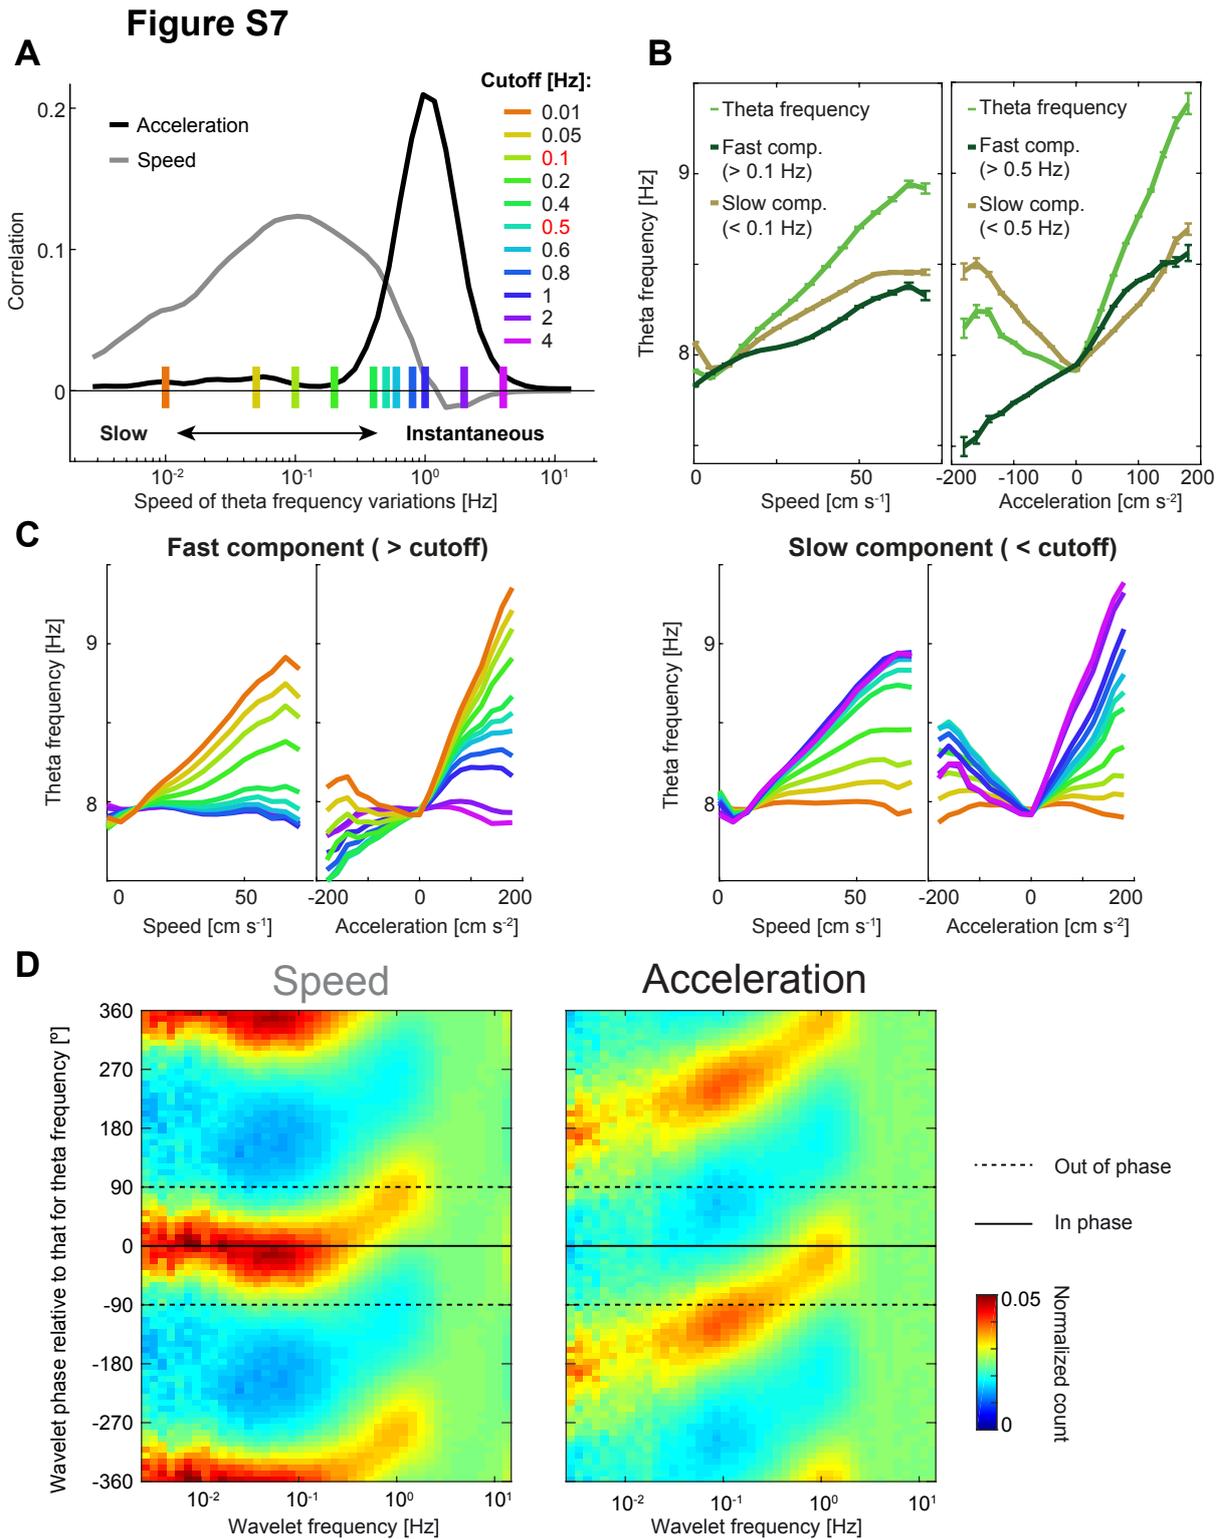


Figure S7. Different temporal dynamic range for correlations between theta frequency and speed or acceleration (related to Figure 6).

(A-C) Figure 6E repeated with a range of alternative cutoff values. (A) The bottom panel of Figure 6D is reproduced here adding with colour code the range of explored cutoff values (cutoff values used in Figure 6E, 0.1 Hz and 0.5 Hz, are indicated in red font). (B) Plots as in Figure 6E but with 0.1 Hz cutoff for speed and 0.5 Hz cutoff for acceleration. (C) Plots of fast

(left) and slow (right) varying components for a range of cutoff values, colour-coded as in (A). Inside each subpanel, the relationship between components of theta frequency and speed (left) or acceleration (right) is shown. Put together, these plots show that *ad hoc* cutoff frequencies used in Figure 6 are pivotal for the relationship between variables. Applying a cutoff value of 0.5 Hz (or higher) gives rise to an apparent correlation between speed and theta frequency. For acceleration the critical cutoff value is 0.1 Hz. For this value, the high-pass component of acceleration recovers the perfectly asymmetric relationship with theta frequency observed in the car (Figure 1F). The low-pass component, instead, exhibits a symmetric modulation of theta frequency, which, like the modulation by speed, has a spurious origin. Both modulations are observed in Figure 6A, despite the fact that theta frequency is modelled to respond exclusively to positive acceleration (model M3). (D) Full distributions used to draw Figure 6D, top. Every column in each colour plot shows, for a given Morlet wavelet frequency, the distribution of differences between the phase of the theta frequency component and the phase of the speed (left) or acceleration (right) components (2 complete phase difference cycles are shown to facilitate visualization). Warm colors indicate synchronization with a given phase difference. Solid lines indicate perfect synchronization (0° phase difference) while dashed lines indicate that Morlet wavelet components are out of phase ($\pm 90^\circ$ phase difference).