J Neurophysiol 103: 1658–1672, 2010.
First published January 20, 2010; doi:10.1152/jn.00432.2009.

Reward-Associated Gamma Oscillations in Ventral Striatum Are Regionally Differentiated and Modulate Local Firing Activity

Tobias Kalenscher,[1,*] Carien S. Lansink,[1,2,*] Jan V. Lankelma,[1] and Cyriel M. A. Pennartz[1]

[1]*Department of Cognitive and Systems Neuroscience, Swammerdam Institute for Life Sciences, University of Amsterdam, Amsterdam, The Netherlands; and* [2]*Department of Anatomy and Neurobiology, University of Maryland School of Medicine, Baltimore, Maryland*

Submitted 19 May 2009; accepted in final form 16 January 2010

Kalenscher T, Lansink CS, Lankelma JV, Pennartz CMA. Reward-associated gamma oscillations in ventral striatum are regionally differentiated and modulate local firing activity. *J Neurophysiol* 103: 1658–1672, 2010. First published January 20, 2010; doi:10.1152/jn.00432.2009. Oscillations of local field potentials (LFPs) in the gamma range are found in many brain regions and are supposed to support the temporal organization of cognitive, perceptual, and motor functions. Even though gamma oscillations have also been observed in ventral striatum, one of the brain's most important structures for motivated behavior and reward processing, their specific function during ongoing behavior is unknown. Using a movable tetrode array, we recorded LFPs and activity of neural ensembles in the ventral striatum of rats performing a reward-collection task. Rats were running along a triangle track and in each round collected one of three different types of rewards. The gamma power of LFPs on subsets of tetrodes was modulated by reward-site visits, discriminated between reward types, between baitedness of reward locations and was different before versus after arrival at a reward site. Many single units in ventral striatum phase-locked their discharge pattern to the gamma oscillations of the LFPs. Phase-locking occurred more often in reward-related than in reward-unrelated neurons and LFPs. A substantial number of simultaneously recorded LFPs correlated poorly with each other in terms of gamma rhythmicity, indicating that the expression of gamma activity was heterogeneous and regionally differentiated. The orchestration of LFPs and single-unit activity by way of gamma rhythmicity sheds light on the functional architecture of the ventral striatum and the temporal coordination of ventral striatal activity for modulating downstream areas and regulating synaptic plasticity.

INTRODUCTION

Oscillations of synchronously discharging neuronal cell populations in the gamma range (30–100 Hz) have been observed in many cortical and subcortical structures and in several species. They are believed to underlie a wide range of perceptual, motor, and cognitive functions (Engel et al. 2001; Fries et al. 1997; Masimore et al. 2005; Montgomery and Buzsáki 2007; Singer 1993; van der Meer and Redish 2009b) and have been hypothesized to contribute to sensory awareness and consciousness (Engel and Singer 2001; Singer and Rose 1997; Srinivasan et al. 1999; Tononi et al. 1998). Recent studies demonstrated that gamma-band activity in local field potentials (LFPs) is also prominent in the rat and human striatum (Berke 2009; Berke et al. 2004; Cohen et al. 2009; DeCoteau et al. 2007; Masimore et al. 2005; van der Meer and Redish 2009b). In the dorsal striatum of rats, gamma oscillations were reported to partially correlate with movement initiation (Masimore et al. 2005). In humans, gamma-band activity in ventral striatum (VS) is connected to processing information about monetary gains and losses (Cohen et al. 2009).

The VS is one of the principal target areas of the mesolimbic dopamine projection and is a main input structure of the basal ganglia (Mogenson et al. 1980; Pennartz et al. 1994; Voorn et al. 2004). The basal ganglia play an important role in initiating and coordinating motor actions (Alexander et al. 1986, 1990; Graybiel 1994) and the VS has been implicated in motivation and invigoration of goal-directed behaviors (Berridge and Robinson 1998; McBride et al. 1999; Robbins and Everitt 1996; Salamone et al. 2001; Setlow et al. 2003), drug addiction (Bardo 1998; Koob and LeMoal 1997; Robinson and Berridge 1993; Self and Nestler 1995; Wise 1996, 1998), off-line memory processing (Hernandez et al. 2002; Lansink et al. 2008; Pennartz et al. 2004), pavlovian and instrumental conditioning and reward prediction (Dalley et al. 2005; Knutson and Cooper 2005; McClure et al. 2003; O'Doherty et al. 2004; Robinson and Carelli 2008; Schultz et al. 2000), and decision making (Acheson et al. 2006; Cardinal et al. 2001; Cohen et al. 2009; Wickens et al. 2007). Single neurons in VS respond to rewards or reward-predicting cues (Apicella et al. 1991; Lansink et al. 2008; Robinson and Carelli 2008; Roitman et al. 2005; Schultz et al. 1992; Tran et al. 2005; Wan and Peoples 2006; Yanagihara et al. 2001) and to attributes affecting the value of expected rewards, such as reward magnitude and delay (Izawa et al. 2005; Singh et al. 2008).

To account for the wide diversity of functions in which the VS is involved, we have proposed that the VS is composed of functionally distinct ensembles of neurons (Pennartz et al. 1994). However, notwithstanding the long-known anatomical compartmentation of the VS (Voorn et al. 1989) and recent physiological indications for reactivating ventral striatal cell groups (Lansink et al. 2008), evidence for a functional heterogeneity in the ventral striatal microarchitecture beyond the core/shell divide remains scarce (Robinson and Carelli 2008). Moreover, despite the well-documented role of the VS in reward-related behavior, the function of rhythmic gamma-band activity recorded from this structure remains largely unexplored (but see Cohen et al. 2009; Masimore et al. 2005; van der Meer and Redish 2009b). More specifically, it is of great interest to study whether gamma rhythmicity is expressed uniformly across the VS or is heterogeneous and regionally differentiated, depending on ongoing events and behaviors. Moreover, it is unclear whether VS gamma activity should be characterized as a "carrier rhythm" to which neurons expressing specific information become phase-

* These authors contributed equally to this work.
Address for reprint requests and other correspondence: T. Kalenscher, Department of Cognitive and Systems Neuroscience, Swammerdam Institute for Life Sciences, University of Amsterdam, Science Park 904, 1098 XH, Amsterdam, The Netherlands (E-mail: T.Kalenscher@uva.nl).

locked or whether any informational content can be attributed to gamma oscillations themselves.

Here, we aimed to characterize the function of gamma oscillations in VS during reward-seeking behavior. Rats performed a reward-collection task in which they searched for and consumed three different types of rewards, one per trial, on a triangular track. Because of the importance of VS for motivated learning, reward seeking, and decision making, we hypothesized that ventral striatal gamma oscillations convey not only motor-related but also reward-related information relevant to perform the task. Moreover, we examined whether single-unit activity may phase-lock to gamma-band oscillations in LFP and whether the spike phase modulation could be linked to reward properties of the LFP and single-unit data. Finally, we tested whether gamma oscillations were expressed homogeneously throughout the VS or whether there was evidence for a heterogeneous, differentiated rhythmicity. A preliminary account of this study previously appeared in abstract form (Kalenscher et al. 2008).

METHODS

Subjects

Four male Wistar rats (375–425 g; Harlan Nederland, Horst, The Netherlands) were individually housed under a 12-h/12-h alternating light–dark cycle with light onset at 8:00 am. All experiments were conducted in the animal's day period. On training and recording days, access to water was limited to 2 h following training or recording, although food was available without restriction. All experimental procedures were in accordance with national guidelines on animal experimentation and conformed with the guiding principles in the care and use of animals of the American Physiological Society.

Pretraining

Before surgery, rats were pretrained on a linear track [185 × 10 cm (length × width), 40 cm elevated from the floor] to shuttle back and forth for reward available at both ends of the track. During pretraining, rats were familiarized with three kinds of rewards, which differed in both taste and texture [sucrose solution (10%), vanilla dessert, and chocolate mousse]. In each of 10–12 pretraining sessions, two of three reward types were provided to the rat. Each type was assigned to a specific well location throughout the session but type–location combinations differed between sessions. The wells were baited following a partial reinforcement schedule in which the probability that the rat would find a well baited was 50%.

Task

During recordings, rats underwent a protocol in daily sessions that consisted of a rest period (20–60 min), followed by a phase of reward-searching behavior on a triangular track (20 min) and concluded by a second period of rest (60–120 min). Neural activity during rest periods was not of interest for the current study and is thus reported elsewhere (Lansink et al. 2008, 2009). The track was triangular in shape (equilateral sides, 90 cm; width, 10 cm; cf. Fig. 3) and was novel to the rats at the first recording session. On the track, the rats were required to run in a clockwise direction to detect and consume rewards available at reward wells positioned in the center of each arm. The three types of rewards used during pretraining were also used as rewards in the triangular track task. On each lap, the track was baited with only one of the three reward types at one reward location. The assignment of reward type to reward location was fixed throughout all sessions. Rewards were distributed according to a pseudorandom schedule in which, in a given number of reward-site passages, the three ports were baited and unbaited an equal number of times. Moreover, all possible combinations of two subsequent trials were represented equally (possible combinations: two successive reward sites were both baited; two successive reward sites were unbaited; one site was baited and two were unbaited; one site was unbaited and two were baited).

Surgery and recordings

Rats were implanted with a multielectrode microdrive containing seven to eight individually movable tetrodes directed unilaterally to the VS of the right hemisphere (1.8 mm anterior and 1.4–3 mm lateral to Bregma; Paxinos and Watson 1996) and additional tetrodes to the hippocampus (reported elsewhere; cf. Lansink et al. 2009). The reference electrodes used in the current analysis were placed in the corpus callosum and near the hippocampal fissure. A skull screw located in the caudal part of the parietal skull bone contralateral to the drive location served as ground. Only LFPs that were recorded from different tetrodes were taken into account in our analysis. According to the exit grid of the microdrive, the distance between tetrodes was at least about 330 μm and maximally about 1 mm, not taking into account the difference in depth of the tetrodes. Spike trains from individual cells, LFPs, and position of the rat's head on the track were recorded using a 64-channel Cheetah recording system (Neuralynx, Bozeman, MT). When signals exceeded a manually preset voltage threshold, waveforms were sampled at 32 kHz for 1 ms (filter settings: 600–6,000 Hz). LFPs were continuously sampled at 1,990 Hz and band-pass filtered between 1 and 475 Hz.

Data analysis

LFP ANALYSIS. To assess event-related variations in gamma power, we computed spatial power maps. The power maps contained the average gamma power associated with locations on the triangle track (bin size: 3 pixels, or 7.5 mm, according to the resolution of the videotracker; we report mean values averaged across 3 × 3-pixel grids). Power maps were constructed for general reward-site visits (all site visits were pooled across visits when rewards were present and absent), baited reward-site visits (only when rewards were present; visits when rewards were absent were left out from the analysis), unbaited reward-site visits (only reward-absent visits were included), and for differences in gamma power (baited minus unbaited reward-site visits). The power maps in Fig. 3 (see RESULTS) were computed by transforming raw LFPs using the Morlet wavelet (with the Matlab Wavelet toolbox; The MathWorks, Natick, MA) and taking the mean power of the coefficients for the gamma band (30–100 Hz; coefficients were computed for the 30- to 40-, 40- to 50-, 50- to 60-, . . . , 90- to 100-Hz bands; bin size: 3 pixels). Gamma power was obtained by squaring the filtered data, with each value averaged by the number of events. We opted for the 30- to 100-Hz range instead of smaller bands to maximize the gamma range covered.

For statistical analysis, we quantified gamma power before and after baited and unbaited reward-site visits. Gamma power was obtained by squaring the LFP signal after filtering the raw LFP data (frequency range: 30–100 Hz; sample rate: 1,990 Hz) with a yulewalk filter for frequencies $<$50 Hz (Yulewalk order 50) and a Chebyshev filter for frequencies $>$50 Hz (Chebyshev order 6, ripple in band-pass 0.1 dB). The thus obtained gamma power time series was then smoothed with a zero-phase moving average filter (window width: 100 samples). To test for task-related differences, gamma power was averaged across samples in a window of 4 s, centered on the arrival at a baited or unbaited reward site. Main effects and the interaction of "types of reward site," "baitedness," and "time relative to arrival" (gamma power in the interval 2–0 s before arrival at a reward site vs. gamma power in the interval 0–2 s after arrival at a reward site) on gamma power were assessed using a within-subject ANOVA. Re-

ward-site arrivals were signaled by the crossings of off-line installed "virtual photobeams" (cf. Lansink et al. 2008), positioned right before the point at which the rat reached each reward well. Differences in gamma power were considered significant when $P < 0.01$.

SPIKE SORTING. In addition to LFP analysis, we examined single-unit activity. Spikes from neurons were separated from those emitted by other neurons recorded on the same tetrode by grouping spikes with similar distributions of waveform properties across the four channels of a tetrode using standard automated and manual clustering methods (i.e., BubbleClust and MClust, respectively). BubbleClust groups spikes based on nearest-neighbor distances, clustering spikes that are close to each other, given features of the waveform such as peak-amplitude or area under the curve, and principal components of a spike on each tetrode channel. MClust facilitates manual selection of clusters by allowing users to limit cluster membership based on boundaries drawn on two-dimensional plots of the waveform features. Clusters of spikes were attributed to a single unit on the basis of waveform characteristics and when they exhibited <0.1% of spike intervals within a 2-ms period in their interspike interval histograms.

REWARD-RELATED FIRING PATTERNS. To identify reward-related units, perievent time histograms were constructed for the rewarded and nonrewarded conditions for each reward site. The histograms were aligned to reward-site arrivals. Reward-related responses were assessed within a period of 1 s before and 1 s after arrival at a reward site. Spike counts were binned in 250-ms intervals. The eight bins comprising the reward period were each compared with three bins taken from the corner passage opposite to the well under scrutiny within the same lap (Wilcoxon's matched-pairs signed-rank test, $P < 0.01$; cf. Lansink et al. 2008). We verified that the firing in the control period of three bins was not marked by specific deviations from the firing in all intermediate segments between corners and reward sites using perievent time histograms and plots of the spatial distribution of firing rates. Units were categorized as significantly reward related when one or more bins in the reward period were significantly different from each of the three reference bins. For further details, see Lansink et al. (2008).

PHASE-LOCKING. Bursts of spike activity in striatum can be time-locked to the cycle of one or more of the oscillatory bands in LFPs (Berke et al. 2004). Here, we tested whether firing patterns of single units in ventral striatum were phase-locked to local gamma oscillations. Spikes were binned (bin size: 10°) and aligned according to the phase of the gamma cycle. The phase was obtained by Hilbert-transforming high-gamma power periods in the filtered LFP signal. We defined periods with high gamma power as those periods where the amplitude of the gamma-filtered signal (range: 30–100 Hz) exceeded a 1% confidence interval around the mean. The LFP signals between the amplitude peaks and troughs identified in this way were grouped together into one high-gamma episode whenever the peak-to-trough interval was <90 ms. Episodes with low-gamma amplitudes (not exceeding the confidence interval) were not considered in this analysis. Significant time-locking was computed for every LFP–neuron combination using the Rayleigh test ($P < 0.001$).

ANALYSIS OF MOVEMENT PARAMETERS. Using an array of light-emitting diodes on the headstage, a video tracking system extracted the rat's position on the track at 60 Hz with a resolution of 2.5 mm/pixel. The behavior of the rat was also stored on videotape. The rat's position in time was estimated by low-pass filtering and zero-phase moving average filtering the digital video tracking signal to remove artifacts attributed to rapid head jerking movement, grooming, and so on. Movement velocity was estimated by downsampling the filtered position signals by a factor 10 and dividing the distance traveled between two sample points by the travel time. Acceleration was computed as the first derivative of the velocity vector.

HISTOLOGY. The final position of the tetrodes was marked by passing a 25-μA current lasting 10 s through one lead of each tetrode to produce a small lesion. The next day, animals were transcardially perfused with a 0.9% NaCl solution followed by 4% paraformaldehyde in phosphate-buffered saline (0.1 M, pH 7.4) before the brains were removed. Coronal brain sections (40 μm) were cut on a Vibratome and Nissl-stained for verification of tetrode tracks and endpoints. All of the tetrode endings were in the VS approximately between 2.2 and 1.2 anterior to Bregma and between 1.6 and 3.0 laterally, compared with an atlas of the rat brain (Paxinos and Watson 1996; cf. Fig. 1). To estimate the number of recordings originating from the core and shell subdivisions of the VS, we first assessed the endpoints of the individual tetrodes in the histological sections and converted them to coordinates according to the atlas (Fig. 1). We then estimated the approximate depths of the tetrodes in each session by subtracting the average travel distance from the tetrode endpoint.

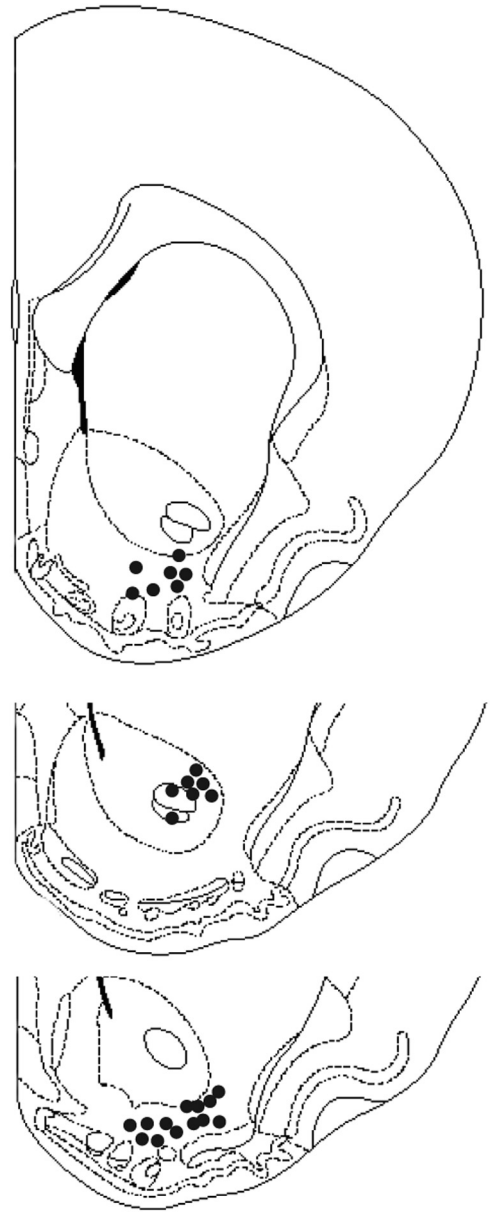

FIG. 1. Histological assessment of endpoint locations. The tips of the tetrodes were in the right ventral striatum approximately between 2.2 and 1.2 anterior to bregma and between 1.6 and 3.0 lateral to midline with respect to the coordinates in Paxinos and Watson (1996).

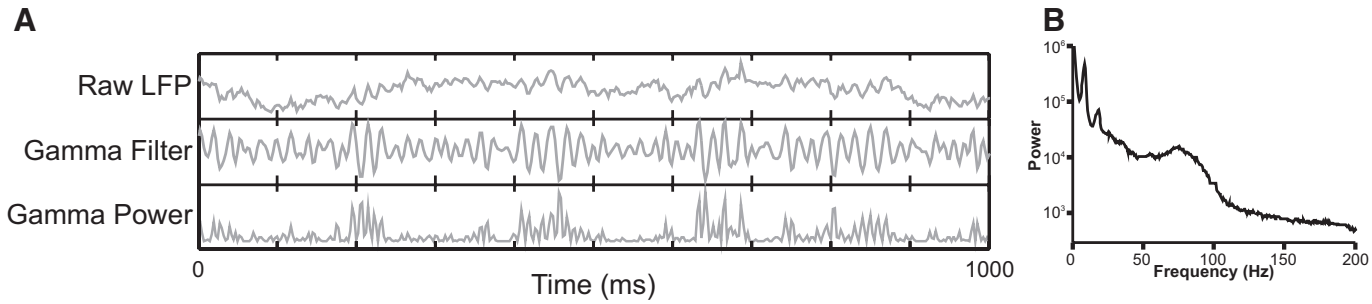

FIG. 2. Gamma oscillations in local field potential (LFP) recordings. *A*: 1,000-ms episode of raw LFP (*top trace*), the same episode filtered for gamma bands (*middle trace*), and gamma power (*bottom trace*). This example shows that short and distinct bursts of gamma activity occurred frequently. *B*: exemplar power spectrum. In most recordings, power in the gamma band was markedly elevated during task performance. The *y*-axis is log-scaled.

Thirty sessions yielded recordings from 150 locations from the core and shell regions of nucleus accumbens. Ensembles from most sessions were likely to contain both core and shell recordings. Six recording sessions, yielding 46 LFPs, were identified as containing core-only recordings. A more precise anatomical localization is currently not possible with the available tetrode array technique.

RESULTS

Behavior and LFP recordings

We recorded neural activity in the VS of four rats performing the reward-collection task on the triangle track described earlier. The rats ran on average 38 (±4.2 SE) laps per session at an average running speed of 46.8 pixels/s (±2.3; ~117 ± 5.8 mm/s). The median traveling time between reward sites was 6.5 s (25th percentile: 2.5 s; 75th percentile: 13.9 s; range: 0.9–158.2 s). We report data from 150 LFP recordings and 761 single units that were obtained in 30 training sessions with, on average, 5 LFP recordings per session (±0.4) and 25.4 single units per session (±2.14).

Transient gamma episodes in LFP and power spectrum analysis

Figure 2*A* shows an exemplary period of raw LFPs during task performance and the same LFP sequence filtered for gamma bands and gamma power. This figure illustrates a phenomenon that was characteristic for most of the recordings: LFPs exhibited prominent brief, transient bursts of gamma activity that usually lasted about 30 to 200 ms. A power spectrum analysis revealed that gamma activity (~50–100 Hz) was pronounced during task performance (Fig. 2*B*), suggesting that gamma frequencies may be related to aspects of the task and/or movement parameters.

Gamma oscillations are task related

Figure 3 shows examples of LFP power maps of a single typical session. Figure 3*A* displays an LFP recording in which gamma power significantly discriminated between general reward-site visits (visits of baited and unbaited sites pooled). Compared with the left and right reward sites (sucrose solution, left; vanilla dessert, right), gamma power was noticeably ele-

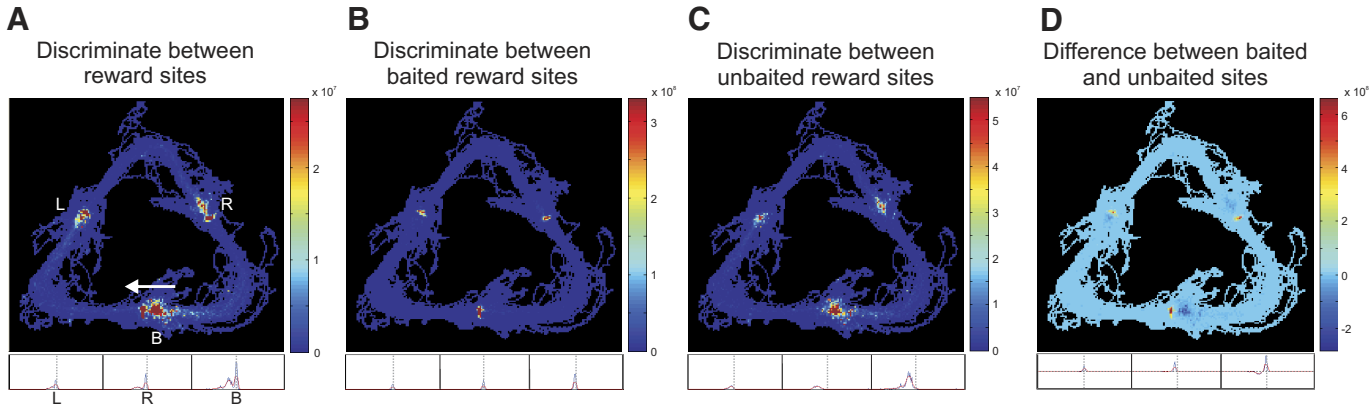

FIG. 3. Examples of power heat maps of task-related gamma activity. The maps show the LFP's mean gamma power on the triangular track. The color bars at the side of each map indicate gamma power in arbitrary units. Units were arbitrary because of the various amplification factors and filter parameters during LFP processing, but are commensurate with μV^2. The plots below the heat maps represent gamma power peristimulus time histograms (PSTHs) 1 s before to 1 s after arrival at a reward site (as indicated by the gray vertical dotted lines). Rats ran on the track in a clockwise direction, as indicated by the white arrow in *A*. Reward sites were positioned at the center of each arm. The bottom reward site (indicated by letter B in *A*) contained chocolate mousse, the left reward site (L) sucrose solution, and the right reward site (R) vanilla dessert. *A*: gamma power discriminated between general reward-site visits (baited and unbaited pooled together) and was lowest before and during visits of the left reward site (sucrose). *B*: gamma power before/during/after baited reward-site visits (visits of unbaited sites were discarded from this analysis). Gamma activity discriminated between reward types. In this example, it was highest for chocolate mousse (bottom reward site). *C*: gamma power in relation to empty reward sites (visits of baited sites were discarded from the analysis). Gamma activity discriminated between reward locations and was higher during visits of the bottom reward site compared with the other 2 sites. *D*: difference in gamma power between baited and unbaited reward sites. The difference in gamma power was often negative before reward-site visits, but became positive during the visit (see text for details).

vated shortly before or on arrival at the bottom reward site (chocolate mousse). Figure 3*B* shows the mean gamma power, averaged across visits when the reward sites were baited, with all unbaited visits removed from this analysis. This panel also shows that gamma power discriminated between combinations of reward site and reward type: compared with visits of reward sites baited with sucrose (left) and vanilla (right), gamma activity was higher when visiting the reward site baited with chocolate mousse (bottom). Gamma oscillations shown in Fig. 3*C* discriminated between unbaited reward sites: bottom reward-site arrivals were associated with clearly higher gamma power than that of arrivals at the left and right reward sites. In the session displayed in Fig. 3*D*, gamma power discriminated between baited and unbaited reward-site visits; the panel shows the difference in gamma power in baited minus unbaited trials.

To assess whether gamma oscillations in the LFPs were task related, we tested for differences in gamma power using a within-subject ANOVA, with the factors "reward site" and "baitedness" (that is, baited vs. unbaited; see METHODS). Furthermore, it is possible that gamma power was different when the rat approached a reward site than when it resided at or left a reward site. We therefore included "time relative to arrival" as a further factor in the ANOVA, with the levels "before arrival" (2 to 0 s before arrival) and "after arrival" (0 to 2 s after arrival).

The ANOVA revealed that, out of a total of 150 LFPs, gamma power in 47 LFPs (31.3%) significantly differentiated between reward sites in general (all $P < 0.01$, all $F > 5.2$), gamma power in 4 LFPs (2.7%) significantly differed between baited and unbaited reward sites (all $P < 0.01$, all $F > 9.1$), and there was a significant main effect of time relative to arrival on gamma power in 48 LFP recordings (32%; all $P < 0.01$, all $F > 9.1$). Furthermore, there were significant interactions between reward site and baitedness in 20 LFPs (13.3%; all $P < 0.01$, all $F > 6.8$), between reward site and time relative to arrival (before/after) in 29 LFP recordings (19.3%; all $P < 0.01$, all $F > 5.7$), and between baitedness and time relative to arrival in 46 LFPs (30.7%, all $P < 0.01$, all $F > 8.1$).

Out of the 47 LFPs with a main effect of reward site on gamma power, the majority (44 LFPs, 93.6%) had the highest gamma power in association with reward site 3 (chocolate) visits, 3 LFPs (6.4%) had the highest gamma power in association with reward site 1 visits (sucrose), and no LFP had maximum gamma power with reward site 2 visits (vanilla); it is possible that this pattern was related to the rats' reward preferences, but future research needs to determine this possibility. Of the 4 LFPs whose gamma power significantly discriminated between baited and unbaited reward sites, gamma power in one LFP (25%) was higher in baited than that in unbaited reward sites and it was higher in unbaited than that in baited reward sites in 3 LFPs (75%). Figure 3*D* suggests that the difference in gamma power between baited minus unbaited visits was negative before, but positive during or after arrival (compare bottom reward site). This was also found in the three other sessions showing a significant difference in gamma power between baited and unbaited sites. It is tempting to speculate that this difference reflected the rats' reward expectation and prediction error, but future research needs to confirm this hypothesis. In addition, in all 48 LFPs with a significant main effect of time relative to arrival, gamma power was higher before than that after arrival at a reward site. Of the 47 LFP recordings that had a main effect of reward-site visits, 25 LFPs (53.2%) also had higher gamma power before than that after reward-site arrival. Compare Table 1 for a summary of the results.

In the following, we will refer to LFPs as *task-related* when their gamma rhythmicity either discriminates between reward-site visits (either on account of spatial or motivational attributes), baitedness, time relative to arrival, or when there was a significant interaction between any of these factors. In total there were 82 LFPs (54.7%) with significantly task-related gamma oscillations (33 LFPs in rat 32, 6 in rat 41, 26 in rat 71, and 17 in rat 83). None of the LFPs with reward site–sensitive or baitedness-sensitive gamma oscillations stemmed from the 46 LFP traces in rat 41 that contained core-only recordings. Thus it is likely that the majority of the task-related patterns in LFPs arose from shell or rostral pole, but not core of VS, although this issue deserves further study.

The gamma power estimate on which this analysis was based was obtained by averaging gamma power across time points in a 4-s window, centered around reward-site arrivals (± 2 s before/after arrival at a site). Additional analyses on average gamma power obtained from a 2-s window (± 1 s before/after arrival at a site) yielded similar results. The time windows encompassed different behavioral events: approaching a reward site, halting, reward consumption, and also reward-site departure. However, these events were common to all three reward sites. Because gamma rhythmicity differentiated between reward-site visits, it was probably not related to them.

For the analysis of unbaited reward-site visits reported earlier, we considered all trials in which a given reward site was unbaited and averaged gamma power across segments of ± 2 s before/after arrival at the site (or ± 1 s, respectively; see preceding text). It is possible that changes in gamma power following preceding visits of a baited reward site may have contaminated the signal of the unbaited reward sites. Even though we cannot fully rule out this hypothesis, we still consider it unlikely. First, due to our reward schedule, an empty reward site followed a baited reward site in only 50% of the trials; in the other 50%, the previous reward site was likewise empty. Thus any spillover effect of the previously baited reward site on gamma power should be heavily diluted. Second, spillover effects from visits of the previous sites should be less pronounced the more time has passed since the visit of the previously baited site. However, the replication of our analysis with smaller segments of ± 1 s around arrivals at the unbaited reward sites yielded essentially identical results (note that the smaller the segments, the more time has passed since the visit of the preceding site).

The European power network operates with an AC at a rate of 50 Hz. It is thus possible that the present results were influenced by 50-Hz noise. To test this possibility, we repeated the ANOVA after applying a notch filter to our data that removed 50-Hz oscillations. This analysis revealed that gamma power in a slightly smaller number of notch-filtered LFPs significantly differentiated between reward sites (27 LFPs, 18%, all $P < 0.01$, all $F > 5.6$), but a higher number of LFPs differentiated between baited and unbaited reward sites (7 LFPs, 4.7%, all $P < 0.01$, all $F > 8.5$) and time relative to arrival (51 LFPs, 34%, all $P < 0.01$, all $F > 9.4$). The general decrease in the number of significantly reward-related LFPs is not surprising, given that the notch filter removes a presumably task-relevant

TABLE 1. *Number of LFPs with significantly reward related gamma power*

Power	Reward Site	Baitedness	Before vs. After Arrival at Reward Site
30–100 Hz	Total: 47 (31.3%)	Total: 4 (2.7%)	Total: 48 (32%)
	Reward site 1: 3 (2%)	Baited > unbaited: 1 (0.7%)	Before > after: 48 (32%)
	Reward site 2: 0	Unbaited > baited: 3 (2%)	After > before: 0
	Reward site 3: 44 (29.3%)		
30–100 Hz, with 50-Hz notch filter	Total: 27 (18%)	Total: 7 (4.7%)	Total: 51 (34%)
	Reward site 1: 6 (4%)	Baited > unbaited: 2 (1.3%)	Before > after: 50 (33.3%)
	Reward site 2: 5 (3.3%)	Unbaited > baited: 5 (3.3%)	After > before: 1 (0.7%)
	Reward site 3: 16 (10.7%)		
45–55 Hz	Total: 57 (38%)	Total: 19 (12.7%)	Total: 25 (16.7%)
	Reward site 1: 6 (4%)	Baited > unbaited: 17 (11.3%)	Before > after: 11 (7.3%)
	Reward site 2: 1 (1.3%)	Unbaited > baited: 2 (1.3%)	After > before: 14 (9.3%)
	Reward site 3: 49 (32.7%)		
70–85 Hz	Total: 8 (5.3%)	Total: 25 (16.7%)	Total: 64 (42.7%)
	Reward site 1: 7 (4.7%)	Baited > unbaited: 0	Before > after: 64 (42.7%)
	Reward site 2: 0	Unbaited > baited: 25 (16.7%)	After > before: 0
	Reward site 3: 1 (0.7%)		

Power	Interaction: [Reward Site] × [Baitedness]	Interaction: [Reward Site] × [Time Relative to Arrival]	Interaction: [Baitedness] × [Time Relative to Arrival]	Triple Interaction: [Reward Site] × [Baitedness] × [Time Relative to Arrival]	Any Effect (General Task-Relatedness)
30–100 Hz	20 (13.3%)	29 (19.3%)	46 (30.7%)	11 (7.3%)	82 (54.7%)
30–100 Hz, with 50-Hz notch filter	17 (11.3%)	14 (9.3%)	47 (31.3%)	2 (1.3%)	81 (54%)
45–55 Hz	27 (18%)	41 (27.3%)	21 (14%)	19 (12.7%)	75 (50%)
70–85 Hz	15 (10%)	16 (10.7%)	61 (40.7%)	5 (3.3%)	84 (56%)

The table shows the number and the proportion (percentage) of LFPs whose gamma power was significantly modulated by reward site visits, baitedness, time relative to reward site arrival (before vs. after arrival), and/or any interaction term between these factors. The indicated percentages of LFPs are relative to the total number of recorded LFPs ($n = 150 = 100\%$). Proportion indications in the main text may differ from the percentages indicated in this table, depending on the reference group used to compute the percentage values. Significance was assumed when $P < 0.01$.

bandwidth section from the signal (van der Meer and Redish 2009b). Importantly, the fact that a substantial number of LFPs still showed task-related gamma oscillations after removing all 50-Hz activity suggests that the main effects reported in this study cannot be attributed to AC noise. All numbers of LFPs with significantly modulated gamma power are shown in Table 1.

Task correlates of different gamma bands

A number of groups (Berke 2009; Berke et al. 2004; Sharott et al. 2009; van der Meer and Redish 2009b) have found task-related differences in striatal gamma power in the low (~50-Hz) and high (~80-Hz) ranges. For example, van der Meer and Redish (2009b) reported that 80-Hz oscillations increase in power as animals approach a reward site, but they did not find this effect in 50-Hz oscillations. Conversely, on departure from a reward site, Berke (2009) found a decrease in power in the 50-Hz range and an increase in the 80-Hz range. We tested whether we would also find a functional distribution across gamma bands in the present study. We repeated our analysis with LFP signals filtered in the 45- to 55-Hz range (in the following referred to as 50-Hz range) and in the 70- to 85-Hz range (in the following referred to as 80-Hz range).

In the 50-Hz range, we found that gamma power in 57 LFPs (38%; instead of 47 LFPs when considering the entire bandwidth of 30–100 Hz) significantly differentiated between reward-site visits (all $P < 0.01$, all F values >5.8; see Fig. 4*A* for an example), 19 LFPs (12.7%) differentiated between baited and unbaited reward sites (all $P < 0.01$, all F values >8.3), and there was a significant interaction between reward-site visits and baitedness in 27 LFPs (18%; all $P < 0.01$, all F values >6.3). Gamma power in 25 LFPs (16.7%) was significantly different before versus after arrival at a reward site (all $P < 0.01$, all $F > 8.2$). Of these, 11 LFPs had higher gamma power before arrival at a reward site, compared with that during or after arrival, and 14 LFPs showed the opposite pattern.

Conversely, in the 80-Hz range, we found that gamma power in only 8 LFPs (5.3%) significantly differentiated between reward-site visits (all $P < 0.01$, all $F > 10.3$), 25 LFPs (16.7%) differentiated between baited and unbaited reward sites (all $P < 0.01$, all $F > 8.6$), and there was a significant interaction between reward-site visits and baitedness in 15 LFPs (10%; all $P < 0.01$, all $F > 5.4$). Importantly, gamma power in 64 LFPs (42.7%) was significantly different before versus after arrival at a reward site (all $P < 0.01$, all $F > 9.1$; see Fig. 4*B* for example). All 64 LFPs had a higher gamma power before compared with that during or after reward-site arrival. All results are summarized in Table 1.

Chi-square (χ^2) tests revealed that a significantly higher number of LFPs had reward site–related gamma oscillations in

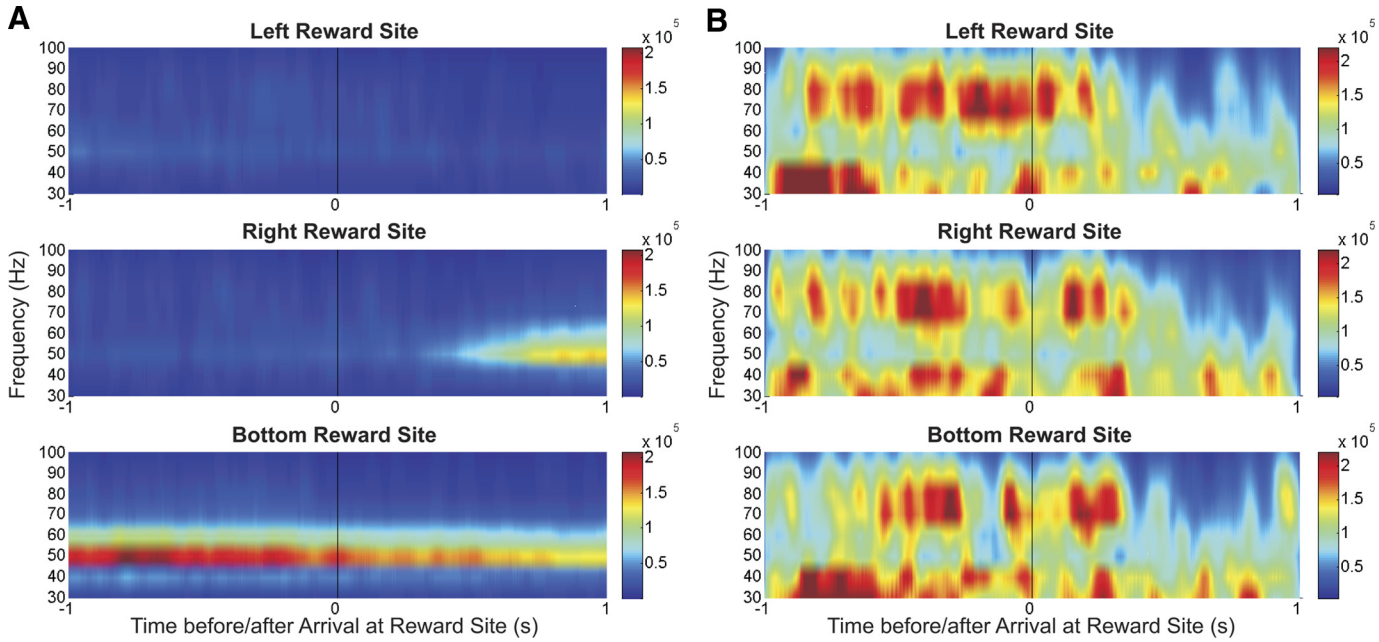

FIG. 4. Functional dissociation of different gamma bands. *A*: PSTH 1 s before until 1 s after arrival at a reward site, as indicated by the vertical black line in the center of the graphs) across the 30- to 100-Hz bandwidth (*y*-axis). The graphs show the PSTHs for the left, right, and bottom reward sites (*top* to *bottom*). They show data from a single exemplar session in which gamma oscillations in the 50-Hz range differentiated between reward sites. *B*: PSTH from another exemplar session in which gamma oscillations in the 80-Hz range were higher before than after arrival at a reward site.

the 50- than that in the 80-Hz range ($\chi^2 = 47.2$, $P < 0.001$), but gamma power discriminated between reward-site arrival versus departure in significantly more LFPs in the 80- compared with the 50-Hz range ($\chi^2 = 24.3$, $P < 0.001$). There was no significant difference in the number of LFPs with baitedness-sensitive gamma power in the 50- and 80-Hz ranges ($\chi^2 = 0.96$, $P = 0.33$).

In summary, linking to previous work, our data also suggest a functional differentiation in gamma bands: whereas gamma power was more often related to reward locations in the 50- than that in the 80-Hz range, we found that more LFPs had different gamma power depending on the timing relative to reward-site arrival in the 80- compared with that in the 50-Hz range. More specifically, gamma power in the 80-Hz range dramatically dropped after the arrival at a reward site relative to the approach phase.

Gamma power correlates weakly with movement parameters

The behavioral events in the present task were correlated with movement parameters. For example, when rats encountered an unbaited reward site, they halted shortly and then moved on directly to the next baited site, but they stopped longer at baited reward sites for reward consumption. Thus behavioral events, such as reward detection and discrimination, were potentially confounded by differences in movement speed and acceleration.

Indeed, a repeated-measures ANOVA across all rats revealed a significant main effect of reward site and baitedness on velocity in the 2 s before arrival at a reward site (all $F > 100$, all $P < 0.001$) and a significant interaction between reward site and baitedness [$F(2,720) = 46.7$, $P < 0.001$]. In general, velocity was higher when reward sites were baited [$F(1,360) = 298.3$, $P < 0.001$] and velocity was highest when approaching the reward site baited with sucrose solution (all $F > 75$, all $P < 0.001$). Despite the significance, the mean velocity preceding the different reward sites did not differ very much [sucrose: 84.8 ± 1.1 pixels/s (mean ± SE); vanilla: 63.3 ± 0.94; chocolate: 69.7 ± 1]. An ANOVA on the effects of reward site and baitedness on acceleration yielded comparable results.

We performed a linear regression analysis to test whether gamma power was correlated with movement velocity or acceleration. We considered correlations significant when $P < 0.01$. Gamma power in 79 of a total of 150 LFP recordings (52.7%) significantly correlated with velocity (average correlation coefficient 0.11 ± 0.008 SE; range: −0.05 to 0.28) and gamma power in 53 LFP recordings (35.3%) significantly correlated with acceleration (average correlation coefficient −0.03 ± 0.005 SE; range: −0.14 to 0.17). Gamma power in 50 LFPs with velocity-related gamma oscillations also showed significant acceleration-related gamma oscillations. The acceleration measure comprised positive acceleration (increase in running speed) and negative acceleration (decrease in running speed). Further regression analyses performed on positive and negative accelerations separately revealed that 37 LFP recordings (24.7%) significantly correlated with positive acceleration and 20 LFP recordings (13.3%) significantly correlated with negative acceleration.

All correlations were very weak, despite their statistical significance (significance can presumably be attributed to the very large number of data points). Figure 5*A* shows the LFP with the highest correlation (out of all LFPs) between gamma power and velocity ($r = 0.28$) and Fig. 5*B* shows the LFP whose gamma power correlated the most strongly with acceleration ($r = 0.17$, different session). These figures show that a clear and distinct relationship between movement parameters and power was not evident in any of the plots. Figure 5, *C* and *D* shows the distributions of correlation coefficients between gamma power of all LFPs with velocity (Fig. 5*C*) and accel-

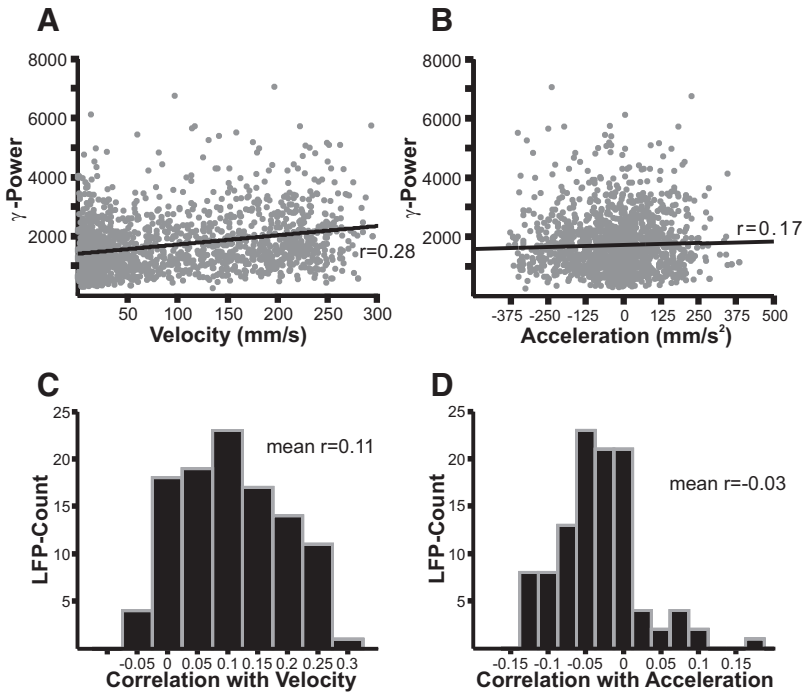

FIG. 5. Correlation between gamma power and movement parameters. *A* and *B*: this figure shows the 2 LFP recordings with the highest correlation between gamma power and running velocity (*A*, $r = 0.28$) and gamma power and acceleration (*B*, $r = 0.17$). *B* includes both positive and negative accelerations. A dissociation of positive and negative accelerations did not yield any better correlations with gamma power (see main text). Despite the relation between gamma power and movement parameters, a regression analysis revealed that task-related gamma modulation (see Fig. 3) remained significant even when including velocity and acceleration in the regression model. *C* and *D*: histograms showing the distribution of correlation coefficients of gamma power in all LFPs with velocity (*C*) and acceleration (*D*).

eration (Fig. 5*D*). Figure 5*D* suggests that gamma power was more often negatively than positively correlated with acceleration. However, because of the general weakness of the correlations, we emphasize the lack of strong relationships between gamma power and movement parameters.

Despite the weakness of the correlations, to rule out possible confounds, it is imperative to disambiguate the influence of movement parameters and task-related factors on gamma power modulation. To this end, we regressed gamma power against the continuous variables "velocity" and "acceleration" and the categorical variables "reward site" and "baitedness," using multiple linear regression analysis. The regression analysis was restricted to data covered by a 4-s window, centered on the time points of reward-site arrivals. We found that, even when including velocity and acceleration in the regression model, the factor "reward site" remained significant in 56 ($P < 0.01$) of all 150 LFP recordings (37.3%) and the factor "baitedness" in 8 recordings (5.3%, all $P < 0.01$; see Supplemental Fig. S1 for histograms of the regression coefficient distributions).[1] The subtle discrepancy in the number of significantly task-related gamma oscillations between the regression analysis and the ANOVA can be attributed to differences in the statistical methods and power. This analysis indicates that gamma power was presumably modulated by task-related factors independently from velocity or acceleration effects.

Single units phase-lock to gamma rhythms

A total of 761 well-isolated single units were recorded over 30 sessions. Of these neurons, 66 (8.7%) showed significant reward-related firing rates whose activity changed prior to or during reward-site visits (all $P < 0.01$, Wilcoxon test). We refer to these 66 neurons as having reward-related activity. For a more complete coverage of single-unit activity, refer to Lansink et al. (2008).

[1] The online version of this article contains supplemental data.

Activity bursts of single ventral striatal neurons were often rhythmic and regular. We tested whether their discharge patterns were locked to a particular gamma range of the local LFPs. For the purpose of this analysis, we considered all periods where the amplitude of the gamma-filtered LFP signal (range: 30–100 Hz) exceeded the 1% confidence interval.

Figure 6*A* shows a typical gamma-modulated neuron. The neuron's firing rate covaried with the cycle of the gamma oscillation, with peak firing rates slightly lagging after 0° of the gamma cycle (peaks in the phase cycle; note that the polarity of the LFP signal depends on the recording settings and is thus arbitrary, albeit identical for all recordings). Figure 6*B* shows another phase-locked neuron that had a preferred phase of about 180°. To assess how many neurons phase-locked to the gamma cycle, we tested whether single neurons had a significantly nonrandom phase distribution using Rayleigh's *r* test (significance threshold $P < 0.001$; we refer to single units as phase-locking if the Rayleigh statistic reaches significance, independent of the neuron's preferred gamma phase). In total, 37 of 761 neurons (4.9%) significantly phase-locked to one or more LFPs in the gamma range and 38 of 150 LFP recordings (25.3%) had one or more phase-locked single units. Note that we included all neurons in our analysis, including those neurons that fired very little or not at all during task performance (cluster cutting sessions included rest periods before and after the task that are not further covered here; cf. Lansink et al. 2008). The estimate of the percentage of phase-locked neurons and LFP–neuron combinations (see following text) may have been higher had we restricted our analysis to the more active neurons. To estimate the distribution of preferred phases, we computed a circular phase histogram by extracting the circular mean over phase for each significantly phase-locked neuron (Fig. 6*C*). This figure shows that the excitable phases in the gamma cycle were distributed, but the majority of neurons had a preferred phase ranging between 120 and 210°. It is of note

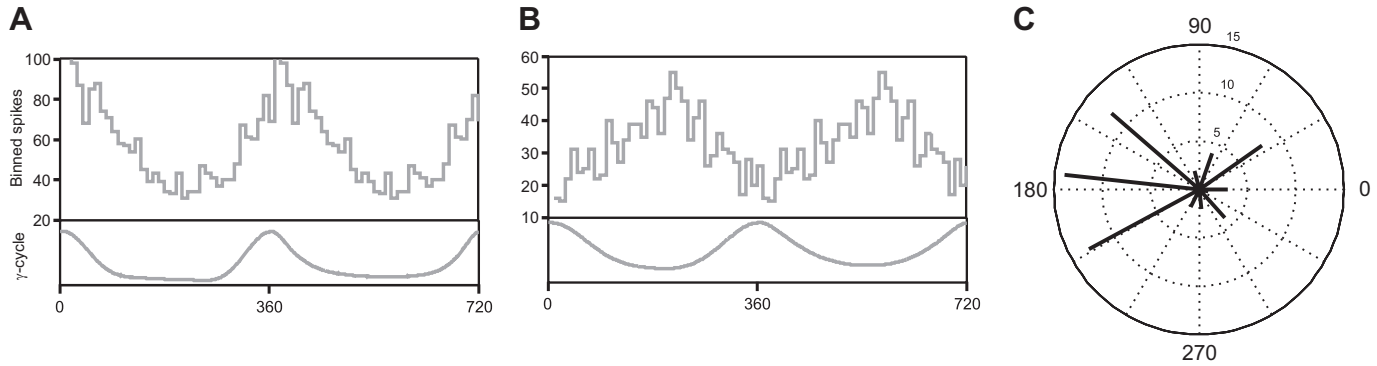

FIG. 6. Phase-locking of single units. *A*: example of a phase-locked unit. The graph shows the activity of a single neuron as a function of the gamma cycle of the LFP. The top part of the figure displays the average spike counts per bin (bin size: 10°) and the bottom part displays the gamma cycle. This neuron had a preferred phase of about 0° of the gamma cycle, i.e., the discharge probability was highest at 0° of the gamma cycle. Note that the polarity of the LFP signal is arbitrary and depends on the recording settings. Thus a preferred phase of 0° may also be interpreted as 180°. *B*: another phase-locked neuron with a preferred phase of about 180°. *C*: polar phase histogram showing the distribution of preferred phases across all neurons with a significantly nonrandom phase distribution, as determined with the Rayleigh statistic (all $P < 0.001$). The numbers next to the concentric circle indicate LFP counts. Most phase-locking neurons had a preferred phase around 180°.

that considerably more units and LFPs showed significant phase-locking when the analysis was not restricted to high-gamma power segments due to more spikes entering the analysis and thus higher statistical power.

Thirteen single units qualified as putative fast-spiking interneurons because of their high average firing rate of >8 Hz and particular waveform characteristics (small peak-to-valley width, valley shape, high decay rate of spike valley; Taverna et al. 2007). We assumed that the remaining cells were medium-sized spiny neurons by a large majority. After removing these 13 putative interneurons from the data set, we found less, but still substantial, phase-locking: 25 presumed medium-sized spiny neurons (3.3%) significantly phase-locked to one or more LFPs in the gamma range and 18 LFP recordings (12%) had one or more phase-locked medium-sized spiny neurons. Moreover, 12 of the 13 putative interneurons (92%) significantly phase-locked to one or more LFPs in the gamma range. In sum, we show that there was not only a significant relationship between gamma-band activity and firing of presumed interneurons, but also putative medium-sized spiny neurons.

Next, we tested whether there was a relationship between phase-locking and reward relatedness of the LFPs and single units. We found that 29 (35.4%) of the 82 LFPs with task-related gamma oscillations had one or more units that significantly phase-locked to their gamma rhythm. In contrast, only 9 (13.2%) of the 68 LFPs with task-unrelated gamma oscillations had one or more phase-locked neurons. Chi-square tests revealed that significantly more LFPs had gamma phase-locked neurons when their gamma oscillations were task related rather than task unrelated ($P < 0.001$, $\chi^2 = 13.3$). In addition, 9 of the 66 single units with task-related activation (13.6%) phase-locked to gamma rhythms of one or more LFPs, but only 28 of the 695 single units without task-related activity (4%) phase-locked to gamma oscillations in any LFP. The proportion of phase-locking units was significantly higher in the task-related group than that in the task-unrelated group ($P < 0.001$, $\chi^2 = 14.1$). In sum, these data suggest that phase-locking often went together with task relatedness of the LFP gamma oscillations and single-unit activity. The results are summarized in Table 2.

Gamma rhythmicity is regionally differentiated within VS

Next, we asked whether gamma rhythmicity was uniformly expressed across the VS or whether it was regionally differentiated. We found that, occasionally, the gamma power time series of some LFPs were remarkably different from those of the other LFPs of the same recording session. An example of a recording session that contained similar and dissimilar gamma power time series is shown in Fig. 7*A*. To quantify this observation, we calculated linear correlations between the gamma power time series in all LFPs acquired in a given recording session. This was done for every session separately. To rule out that variations in gamma power could be explained by differences in the reference signal, we restricted this analysis to 143 LFP traces that were referenced against the same electrode (placed in corpus callosum). Figure 7*B* displays the distribution of correlation coefficients across all sessions. The histogram reveals that, even though gamma rhythmicity in the majority of LFP traces correlated highly, a large number of correlations was remarkably small: 65 of 337 correlations (19.3%) had a correlation coefficient that was < mean ± 2SE and 26 correlations (7.7%) were even <0.1.

To test whether the distribution of correlation coefficients was significantly nonrandom, we calculated the correlation coefficients across time-shifted gamma power time series. To generate these series, for every session, we shifted the second LFP recording by 2 s, the third LFP recording by 4 s, the fourth LFP recording by 6 s, and so on, and then calculated linear correlation coefficients across those LFP traces. A Kolmogorov–Smirnov (KS) test revealed that the distribution of time-shifted correlation coefficients, which peaked close to zero, was significantly different from the original distribution (KS = 0.99, $P < 10^{-130}$).

It is possible that gamma oscillations sensitive to a particular task event are more prone to correlate than gamma oscillations that are sensitive to a different or no task event. To test this possibility, we examined whether LFP pairs with highly correlating gamma time series had a different likelihood to have behavioral correlates than LFP pairs with poorly correlating time series. An LFP pair was classified as poorly correlating if the correlation coefficient of their gamma time series was

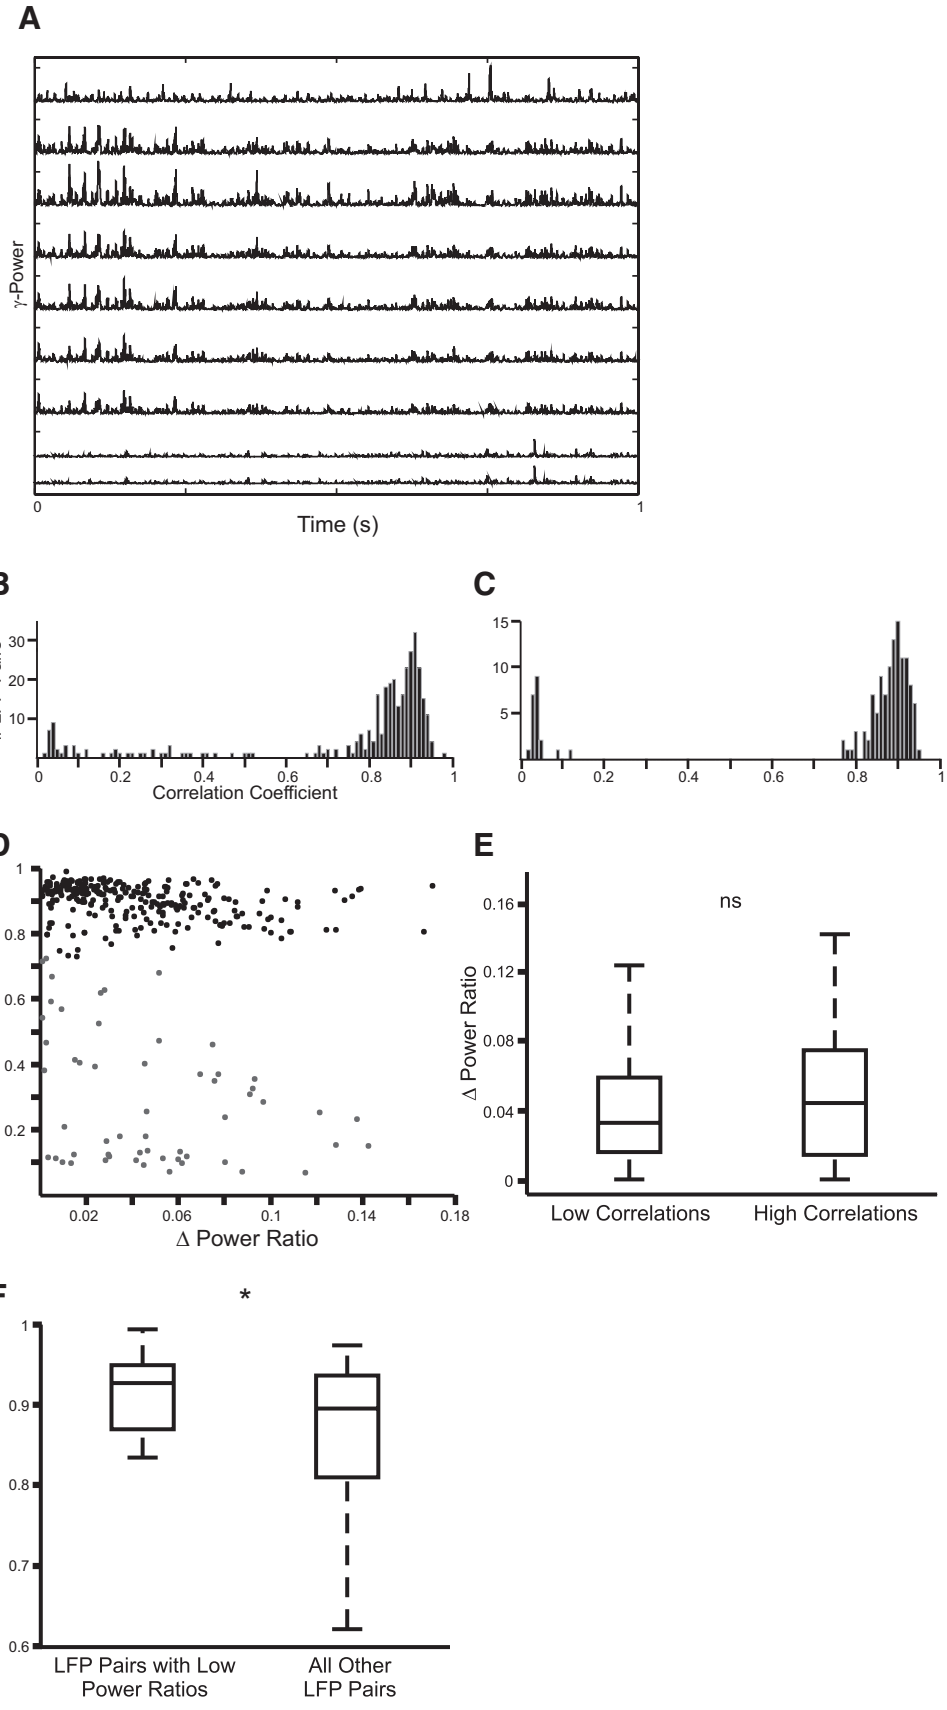

FIG. 7. Regional differentiation of gamma rhythms: within-session correlations of gamma power time series. All LFPs were referenced against the same electrode. *A*: example of a recording session in which some gamma power time series were highly correlated (*traces 2–7* from *top*); others were poorly correlated (*trace 1* and *traces 8 and 9* from *top*). *B*: distribution of correlation coefficients between gamma power time series. Correlation coefficients were computed across all possible pairwise LFP combinations within a recording session. The histogram shows the distribution of those correlations across all sessions. Even though the majority of gamma power series were significantly correlated, a large number of correlations was remarkably poor. *C*: distribution of correlations between gamma power time series, restricted to the data from core-only recordings. *D*: scatterplot mapping the correlation coefficients of LFP pairs against the difference in their gamma power ratios. Dots contributing to the low-correlation group in *E* are indicated in gray; dots contributing to the high-correlation group are indicated in black. If low correlations between 2 gamma power time series were the consequence of differences in the net gamma power, one would expect a clustering of dots in the lower right quadrant (high differences in gamma power would correspond to poor correlations) and upper left quadrant (low differences in gamma power would correspond to high correlations). The absence of such clustering suggests that low correlations were unrelated to discrepancies in the net gamma power difference between 2 LFPs. *E*: box plots of the difference in power ratios between LFPs in a pair (*y*-axis). All LFP pairs were grouped into a low-correlation group (correlation coefficient below the lower confidence interval limit; gray dots in *D*) and high-correlation group (correlation coefficients within the confidence interval, or above the upper limit; black dots in *D*). The horizontal line in the boxes correspond to the median difference in power ratios, the boxes encompass the 25th and 75th percentiles, and the whiskers cover the extent of the rest of the data. There was no significant difference in the distribution of power ratio differences. *F*: LFP pairs were grouped into the category "pairs where both LFPs had a low power ratio" and "pairs where only one or no LFP had a low power ratio." The correlation coefficients between these groups were significantly differently distributed and LFPs with low power ratios had a slightly higher median correlation, although the difference was presumably due to a difference in the distribution variance. $^*P < 0.05$ (Kolmogorov–Smirnov test).

below the confidence interval limit (mean − 2SE). A χ^2 test revealed that there was a significant difference in the number of cases in which neither of the two LFPs was task related between poorly and highly correlating LFP pairs (poor: 10 of 48 LFPs, 20.8%; high: 125 of 277 LFPs, 45.1%; $\chi^2 = 10.7$, $P < 0.001$). In contrast, there were not significantly more cases in which both LFPs were task related in poorly correlating LFP pairs (26 of 48 LFPs, 54.2%) compared with highly correlating LFP pairs (109 of 271 LFPs, 40.2%; $\chi^2 = 3.3$, $P = 0.07$). Interestingly, compared with poorly correlating LFP pairs, highly correlating LFPs had a significantly smaller number of cases in which only one LFP was task related, but the other was not (poor: 12 of 48 LFPs, 25%; high: 37 of 271 LFP, 13.7%; $\chi^2 = 4$, $P < 0.05$). This analysis yields the conclusion that a large proportion of highly correlated LFPs was not related to the task, underscoring that a correlate to particular task events was not necessary for gamma-band activity to synchronize. Next, we conducted an in-depth analysis in which we tested whether poorly correlating gamma series had identical reward-site relatedness or whether their gamma patterns were related to different reward sites. Thirty-nine of 49 LFP pairs (79.6%) whose gamma rhythmicity correlated poorly had the highest gamma power at the same reward site. When restricting the data set to those LFPs whose gamma power was significantly modulated by reward site, we found that all 21 (100%) LFP pairs with poorly correlated gamma power had the highest gamma power at the same reward site. This analysis further supports the conclusion that task relatedness was not necessary or sufficient to synchronize gamma oscillations. Even though caution is required when drawing conclusions, we consider the general dissociation between reward relatedness and gamma power correlation as further support for the hypothesis that information is represented by microanatomically and spatially distinct ensembles (cf. Pennartz et al. 1994).

As mentioned earlier, tetrode array recordings have some inherent uncertainty with regard to the exact microanatomical location of LFP recordings. It is possible that the nonuniformity in LFP traces may stem from the fact that different LFP signals were made up by different anatomical sources. For example, within a given session, some LFP signals may be more strongly influenced by activity in the core of VS, whereas other LFP signals may be more strongly influenced by VS shell activity. To address the possibility of homogeneity within one VS macrocompartment, we repeated the correlation analysis, this time restricted to the 46 LFP recordings obtained from the six sessions that contained core-only recordings (see HISTOLOGY). The results revealed an almost bimodal distribution of high and low correlations (Fig. 7*C*). Twenty-one of 137 LFP pairings (15.3%) had a correlation coefficient that was < mean − 2SE and 20 (14.6%) LFP pairs even had a correlation <0.1. Again, the distribution of correlation coefficients across time-shifted time series was significantly different from the original distribution (KS = 1, $P < 10^{-51}$). Thus many, but not all, of the LFPs with $r < 0.1$ shown in Fig. 7*B* stemmed from recordings in the core region of VS (cf. Fig. 7*C*).

Furthermore, it is possible that the low correlations between two gamma power series may be attributable to differences in the overall power. If, in a pair of two LFPs, the power spectrum reveals a high peak in the gamma range in one LFP, but not the other, then the correlation between the gamma power time series is likely to be low. To address this possibility, we estimated the relative gamma power for each LFP (core and shell recordings) and tested whether two gamma power time series that correlated poorly had a high difference in their gamma power. Relative gamma power was quantified as the gamma:delta (γ:δ) power ratio, which is a measure of the power in the gamma range relative to the power in the delta range

$$PR_{LFP} = \frac{\log (P_\gamma)}{\log (P_\delta)} \quad (1)$$

where PR_{LFP} indicates the γ:δ gamma power ratio for a given LFP, $\log (P_\gamma)$ indicates the log of the gamma power in the power spectrum and $\log (P_\delta)$ indicates the log of the delta power. We chose the power at 80 and 2 Hz to be representative for gamma power and delta power, respectively, because most power spectra peaked around these frequencies (Fig. 2*B*). If large differences in the gamma power ratio between two LFP gamma power series explained low correlations, then the differences in gamma power ratios in low-correlating LFP pairs should be higher than the gamma power ratios in high-correlating LFP pairs. We compared the differences in the gamma power ratios in the poorly correlating LFP pairs (those gamma power time series whose correlation coefficients were below the lower confidence interval limit) with the power ratios of those LFP pairs not exceeding the confidence interval. A Kolmogorov–Smirnov (KS)[2] test across the differences in the power ratios of high- and low-correlating LFPs did not reach significance ($P = 0.12$; Fig. 7*E*; see Fig. 7*D* for a scatterplot of the LFP pairs' correlation coefficients against their power ratios). Thus we had no evidence to assume that the power ratio differences between high- and low-correlating gamma power time series came from different distributions or could be attributed to strong differences in gamma power.

It is equally conceivable that poor correlations between gamma power time series are driven by generally low gamma powers: If the gamma powers in the power spectra of both LFPs in a pair are low, then their correlations are likely to be poor. We tested whether the gamma power time series of LFP pairs in which the power ratios were below the lower limit of the confidence interval correlated significantly more weakly than did LFP pairs in which one or both gamma power series were within the confidence interval. A KS test reached significance ($P < 0.05$), however, the effect was in the opposite direction as expected. The median correlation between the LFP pairs with low power ratios was slightly higher ($r = 0.93$) than the median correlation in the remaining LFP pairs ($r = 0.9$; Fig. 7*F*). Moreover, Fig. 7*F* suggests that the difference in the distribution of correlation coefficients was mainly due to differences in the variance of the distributions. Independent of what causes the difference in distributions, this analysis suggests that low power ratios in LFP pairs were not sufficient to explain low correlations. Put together, we conclude that the heterogeneity in gamma power time series across different tetrodes was not related to differences in the net gamma power, but was the consequence of locally differentiated sources.

[2] Note that a regression analysis was not possible because the data were distributed heteroskedastically due to their clustered, almost bimodal nature.

DISCUSSION

The VS holds a very prominent position in the brain's reward and motivation systems (Berridge and Robinson 1998; Heimer and Wilson 1975; Lansink et al. 2008; Mogenson et al. 1980; Pennartz et al. 1994; Robbins and Everitt 1996; Salamone et al. 2001; Schultz et al. 1992; Voorn et al. 2004). Here, we show that its functional role in reward-related information processing is reflected by variations in the gamma oscillations in LFPs. In line with previous reports (Berke et al. 2004; Masimore et al. 2005; van der Meer and Redish 2009b), we found frequent, distinct, and short episodes of gamma oscillations in ventral striatal LFP recordings that occurred during task performance. Importantly, gamma oscillations carried task-related information. In about 31% of our recordings, gamma activity discriminated between different reward-site approaches and visits, independent of whether a reward was present or absent. In roughly 3% of our recordings, gamma oscillations encoded whether a reward was present or absent, about 13% of our recordings showed a significant interaction between reward site and baitedness, and 32% of LFPs had higher gamma power before than that after arrival at a reward site. In total, gamma oscillations in almost 56% of our recordings were related to one or more task events. The fact that gamma power correlated with distinct events during task performance suggests that gamma oscillations did not simply reflect the general state of the animal, but encoded aspects of goal-directed behavior and reward. Thus we argue that informational content can be attributed to gamma oscillations.

In addition, we found a functional differentiation in the 50- versus 80-Hz gamma bands: gamma power was more often reward-site related in the 50- than in the 80-Hz range, but was more likely to differentiate between approach and arrival at a reward site in the 80-Hz range. More specifically, gamma power in the 80-Hz range was markedly higher during the approach phase toward a site compared with the halt at or departure from a site. This finding is consistent with that of van der Meer and Redish (2009b) who also showed an increase in power in the VS in the 80-Hz range during approach toward a salient site in a T-maze. In contrast, Berke (2009) found that ventral striatal gamma power in the 80-Hz range increased following, but not preceding, reward. In addition, van der Meer and Redish (2009b) reported that departure from a reward site was accompanied by an increase in power in the 50-Hz, not the 80-Hz, range. It is presently unclear why the results diverge. One possibility is that the different tasks imposed different cognitive and spatial demands on the rats and raised differences in expectational and decisional patterns. Whereas our task was relatively simple and did not require a high degree of cognitive sophistication, the task by van der Meer and Redish (2009b) consisted of a modified continuous T-maze involving a sequence of T-choices and Berke's task (Berke 2009) involved navigation in a four-arm radial maze also involving decision making. Thus these tasks presumably imposed higher working-memory and decision-making demands on the rats than did our task.

We further report that the firing patterns of almost 5% of the recorded single neurons in VS were time-locked to gamma rhythms in LFPs. LFP–unit combinations were significantly more likely to be entrained if the LFP gamma signal or the single unit, or both, were reward related. It is tempting to speculate that reward relatedness of the gamma rhythm facilitated entrainment of single units. Thus in addition to attributing informational context to VS gamma activity, gamma oscillations could additionally be characterized as a "carrier rhythm" to which neurons expressing specific information became preferably phase-locked.

Further, we found that about one fifth of gamma power time series correlated poorly with the other time series of the same session, suggesting that gamma rhythmicity was regionally differentiated within the VS, and even within the nucleus accumbens core as one of its macrocompartments (Voorn et al. 1989; Zaborsky et al. 1985).

Although VS ensembles may process contextual aspects of the task in addition to reward-related information, it is unlikely that changes in gamma activity in the present study reflected spatial instead of reward information exclusively, at least for a subset of our LFP recordings, because >13% of LFPs showed differential gamma activity to baited versus unbaited reward sites (either main effect for baitedness or interaction between reward-site visit and baitedness; an even higher number of LFPs, 31%, showed a significant interaction between baitedness and time relative to arrival at a reward site). If these recordings merely reflected the location of the reward site with respect to an internal or external coordinate system, gamma power should be independent of whether the food tray was full or empty. In addition, even though gamma power correlated weakly, but significantly, with locomotor movement in a number of recordings, the effects of locomotion and reward on gamma power variance were dissociable. Our analysis showed that, in most cases, gamma power covaried with reward information independently of movement parameters. Moreover, it is imaginable that gamma rhythmicity was related to halting at a reward site for reward consumption or to other consumption-specific motor actions. However, halting and reward consumption were common to all baited reward sites, yet gamma power discriminated between the sites. This suggests that gamma rhythmicity was not a function of consumption-related actions, but of cognitive aspects of reward-seeking behavior instead (cf. van der Meer and Redish 2009a,b).

It is conceivable that the oscillation pattern in the LFPs originates from volume conductance from sources outside the VS. If the recorded LFPs were irrelevant to local processing in the VS, then there should be no temporal coordination between local LFP rhythmicity and single-unit activity. However, inconsistent with this presumption, we found that almost 5% of the recorded VS single neurons fired preferentially at a particular phase of the gamma cycle (at $P < 0.001$), indicating that gamma oscillations had local relevance for neural activity and were not merely the consequence of volume conduction from different sources. Finally, we take the fact that many LFPs were regionally differentiated as further evidence against volume conduction, since LFPs should be commonly modulated otherwise. Discharge rate in reward-related single neurons and gamma power in task-related LFPs both, by definition, correlated with reward-related events. Thus there is a potential caveat that the strength of the observed single unit–LFP correlation depends on factors such as mean firing rate and gamma power due to reward-related events as common modulators. However, Table 2 shows that still a substantial number of neurons was phase-locked, but not reward related. This implies that reward relatedness per se was not necessary for phase-

TABLE 2. *Relationship between phase-locking and reward relatedness of the LFPs and single units*

	LFP Task-Related	LFP Not Task-Related	Unit Task-Related	Unit Not Task-Related
Phase-locked	29 (35.4%)	9 (13.2%)	9 (13.6%)	28 (4%)
Not phase-locked	53 (64.6%)	59 (86.8%)	57 (86.4%)	667 (96%)

This table shows the absolute number and proportion of LFPs and single units that were (or were not) phase-locked and task-related. An LFP was classified as phase-locked if it had at least one single unit that significantly phase-locked to its gamma oscillations. A single unit was classified as phase-locked if it significantly phase-locked to at least one LFP. LFPs and single units were classified as task-related if they showed significant modulation of gamma power or discharge rate, respectively, in relation to any of the task events. There was a significant difference in the proportion of phase-locked and not-phase-locked LFPs between task-related and -unrelated LFPs. The same was found for the single units (all $P < 0.001$).

locking. We conclude that it is unlikely that our results are purely attributable to simultaneous fluctuations in discharge rate and gamma power.

The majority, although not all, of significantly task-related gamma oscillations had higher gamma power, particularly in the 80-Hz band, during the approach phase to a reward site compared with during or after a visit of a reward site. Also, gamma power in the 80-Hz range in a very high number of LFPs (41%) showed an interaction between approach and baitedness, suggesting that 80-Hz oscillations were not only higher before than after arrival at a site, but also discriminated whether a site was baited. It is possible that the pseudorandom nature of our task schedule made upcoming rewards partially predictable (there was always only one reward site baited per full lap) and thus allowed the animals to develop specific expectations about imminent reward-site visits. In line with the well-documented role of ventral striatum in reward expectation (Cardinal et al. 2002; Ernst et al. 2004; Hassani et al. 2001; Khamassi et al. 2008; Schultz et al. 1992), it is thus possible that ventral striatal gamma rhythms, in particular in the 80-Hz range, reflected the reward-anticipatory state of the animals, possibly in combination with the act of approach and action value (Lauwereyns et al. 2002; Samejima et al. 2005). Alternatively, the sensory aspects of a reward, such as the odor or the vision of nearby food, may also account for the reward relatedness of gamma oscillations. Even though we cannot entirely rule out the latter possibility, we consider it unlikely to be the exclusive explanation because gamma power also discriminated between empty reward sites in which no primary sensory information was available (Fig. 3*C*). Future research needs to more clearly identify the precise functional significance of gamma rhythmicity in reward search.

One of the main findings of this study concerns the functional heterogeneity of gamma oscillations recorded on different electrodes placed throughout the VS. From the present study and an overview of anatomical, physiological, and pharmacological studies (e.g., Cardinal et al. 2002; Koya et al. 2009; Pennartz et al. 1994), it becomes reasonable to assume that the VS does not behave as a monolithic structure, but is likely to consist of functionally distinct ensembles that may be organized spatially. We previously hypothesized that the core and shell of the nucleus accumbens—the main functional entities in VS—do *not* mediate only one functional component, or subroutine, each, but instead consist of a collection of microanatomically and functionally distinct patches or neural ensembles that deal with different computational aspects of locomotion and various innate or learned behaviors (Pennartz et al. 1994). If this was true, one would expect to find a fine-grained differentiation in the manifestation of physiological parameters within compartments of the VS. In support of this hypothesis, we report a regional differentiation in the expression of gamma rhythmicity in many local LFPs, even when recorded from the core alone. Our findings are also consistent with recent evidence suggesting that different subsets of VS neurons encode distinct features during operant behavior (Robinson and Carelli 2008) and decision making (Izawa et al. 2005).

Put together, our findings suggest that, even within one particular substructure of VS, functional computations and their physiological expressions were highly differentiated. Although it is tempting to speculate that the anatomical compartmentation of the VS in patches and matrix components having specific input and output relationships may underlie the regional differentiation in gamma activity, more research must be done to elucidate the underlying spatial organization of this phenomenon. For example, it is elusive whether the VS's microcompartmentation follows an anatomical gradient, e.g., from medial to lateral (Voorn et al. 1989, 2004). Moreover, the VS has repeatedly been implicated in attributing incentive salience to cues associated with primary rewards and drugs of abuse (Berridge and Robinson 1998, 2003; Robinson and Berridge 1993; Robinson and Carelli 2008). A starting point to refine the regional and functional differentiation in VS would be to conduct LFP recordings in tasks in which rats learn the secondary reinforcing properties of cues or actions.

Optimal foraging behavior requires the ability to associate rewards with reward-predicting cues and spatial locations ("what and where") and to make decisions about which foraging option to pursue based on which cue promises the most valuable and most easily and quickly reachable rewards (Kalenscher and Pennartz 2008). Even though the VS has been implicated in associating rewards and reward features with reward-predicting cues (Cardinal et al. 2002; Floresco et al. 1997; Ito et al. 2008; Tran et al. 2002), it is unknown where the spatial and reward-feature information originates from. Research on spatial navigation suggests that the hippocampus contains cells that are selectively active for particular locations in space (O'Keefe and Dostrovsky 1971; Wilson and McNaughton 1993), represent a general map for locations and events in space (McNaughton et al. 2006; Moser et al. 2008; O'Keefe and Nadel 1978; Wood et al. 1999), and encode information about environmental changes, events, and sequences (Leutgeb et al. 2005; Pastalkova et al. 2008). Moreover, gamma oscillations are also found in hippocampus and are believed to play a role in organizing spatiotemporal information, such as sequences of positions in space (Senior et al. 2008). The hippocampus is therefore a prime candidate to provide the spatial and contextual information to the VS for the formation of reward associations. In support of this hypothesis, recent evidence suggests that hippocampus and VS closely communicate in reward-foraging

tasks (Lansink et al. 2009; Tabuchi et al. 2000). Gamma oscillations may serve as a temporal structure to temporally align spike trains across brain areas and thus facilitate integration of hippocampal and other limbic inputs into VS firing patterns and Hebbian synaptic modification. The idea that gamma power conveys contextual information by imposing a temporal structure on ventral striatal neurons is supported by our findings that LFPs with reward-related gamma oscillations were more likely to have phase-locked single units than reward-unrelated LFPs and that neurons with reward-related activity were more likely to phase-lock to LFPs than reward-unrelated neurons. Thus we conjecture that location- and reward-related phasic gamma bursts in hippocampus and VS may promote spatial and contextual learning phenomena, such as contextual conditioning and conditioned place preference.

In conclusion, regardless of the precise functions of gamma activity, it is remarkable that large-volume neuronal signals that appear to be differentiated across the VS carry task-specific information. Its massive nature can be inferred to have a major impact on target structures, such as ventral pallidum, ventral tegmental area, and lateral hypothalamus. Thus task-specific and regionally differentiated neural oscillations in the gamma range may constitute a key neural mechanism in promoting efficient foraging behavior in animals.

ACKNOWLEDGMENTS

We thank D. Redish and P. Lipa for providing the software used for cluster cutting .

GRANTS

This work was supported by Human Frontier Science Program Grant RGP0127/2001 to C.M.A. Pennartz, Netherlands Organization for Scientific Research Grants VICI 918.46.609 to C.M.A. Pennartz and VENI 016.081.144 to T. Kalenscher, and SenterNovem Grant BSIK-03053 to C.M.A. Pennartz.

REFERENCES

Acheson A, Farrar AM, Patak M, Hausknecht KA, Kieres AK, Choi S, de Wit H, Richards JB. Nucleus accumbens lesions decrease sensitivity to rapid changes in the delay to reinforcement. *Behav Brain Res* 173: 217–228, 2006.

Alexander GE, Crutcher MD, DeLong MR. Basal ganglia-thalamocortical circuits: parallel substrates for motor, oculomotor, "prefrontal" and "limbic" functions. *Prog Brain Res* 85: 119–146, 1990.

Alexander GE, DeLong MR, Strick PL. Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu Rev Neurosci* 9: 357–381, 1986.

Apicella P, Ljungberg T, Scarnati E, Schultz W. Responses to reward in monkey dorsal and ventral striatum. *Exp Brain Res* 85: 491–500, 1991.

Bardo MT. Neuropharmacological mechanisms of drug reward: beyond dopamine in the nucleus accumbens. *Crit Rev Neurobiol* 12: 37–67, 1998.

Berke JD. Fast oscillations in cortical-striatal networks switch frequency following rewarding events and stimulant drugs. *Eur J Neurosci* 30: 848–859, 2009.

Berke JD, Okatan M, Skurski J, Eichenbaum HB. Oscillatory entrainment of striatal neurons in freely moving rats. *Neuron* 43: 883–896, 2004.

Berridge KC, Robinson TE. What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? *Brain Res Brain Res Rev* 28: 309–369, 1998.

Berridge KC, Robinson TE. Parsing reward. *Trends Neurosci* 26: 507–513, 2003.

Cardinal RN, Parkinson JA, Hall J, Everitt BJ. Emotion and motivation: the role of the amygdala, ventral striatum, and prefrontal cortex. *Neurosci Biobehav Rev* 26: 321–352, 2002.

Cardinal RN, Pennicott DR, Sugathapala CL, Robbins TW, Everitt BJ. Impulsive choice induced in rats by lesions of the nucleus accumbens core. *Science* 292: 2499–2501, 2001.

Cohen MX, Axmacher N, Lenartz D, Elger CE, Sturm V, Schlaepfer TE. Good vibrations: cross-frequency coupling in the human nucleus accumbens during reward processing. *J Cogn Neurosci* 21: 875–889, 2009.

Dalley JW, Laane K, Theobald DE, Armstrong HC, Corlett PR, Chudasama Y, Robbins TW. Time-limited modulation of appetitive Pavlovian memory by D1 and NMDA receptors in the nucleus accumbens. *Proc Natl Acad Sci USA* 102: 6189–6194, 2005.

DeCoteau WE, Thorn C, Gibson DJ, Courtemanche R, Mitra P, Kubota Y, Graybiel AM. Oscillations of local field potentials in the rat dorsal striatum during spontaneous and instructed behaviors. *J Neurophysiol* 97: 3800–3805, 2007.

Engel AK, Fries P, Singer W. Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci* 2: 704–716, 2001.

Engel AK, Singer W. Temporal binding and the neural correlates of sensory awareness. *Trends Cogn Sci* 5: 16–25, 2001.

Ernst M, Nelson EE, McClure EB, Monk CS, Munson S, Eshel N, Zarahn E, Leibenluft E, Zametkin A, Towbin K, Blair J, Charney D, Pine DS. Choice selection and reward anticipation: an fMRI study. *Neuropsychologia* 42: 1585–1597, 2004.

Floresco SB, Seamans JK, Phillips AG. Selective roles for hippocampal, prefrontal cortical, and ventral striatal circuits in radial-arm maze tasks with or without a delay. *J Neurosci* 17: 1880–1890, 1997.

Fries P, Roelfsema PR, Engel AK, Konig P, Singer W. Synchronization of oscillatory responses in visual cortex correlates with perception in interocular rivalry. *Proc Natl Acad Sci USA* 94: 12699–12704, 1997.

Graybiel AM, Aosaki T, Flaherty AW, Kimura M. The basal ganglia and adaptive motor control. *Science* 265: 1826–1831, 1994.

Hassani OK, Cromwell HC, Schultz W. Influence of expectation of different rewards on behavior-related neuronal activity in the striatum. *JNeurophysiol* 85: 2477–2489, 2001.

Heimer L, Wilson RD. The subcortical projections of the allocortex: similarities in the neural associations of the hippocampus, the piriform cortex, and the neocortex. In: *Perspectives in Neurobiology Golgi Centennial Symposium*, edited by Santini M. New York: Raven Press, 1975, p. 177–193.

Hernandez PJ, Sadeghian K, Kelley AE. Early consolidation of instrumental learning requires protein synthesis in the nucleus accumbens. *Nat Neurosci* 5: 1327–1331, 2002.

Ito R, Robbins TW, Pennartz CM, Everitt BJ. Functional interaction between the hippocampus and nucleus accumbens shell is necessary for the acquisition of appetitive spatial context conditioning. *J Neurosci* 28: 6950–6959, 2008.

Izawa E, Aoki N, Matsushima T. Neural correlates of the proximity and quantity of anticipated food rewards in the ventral striatum of domestic chicks. *Eur J Neurosci* 22: 1502–1512, 2005.

Kalenscher T, Lansink CS, Lankelma JV, Pennartz CMA. Temporal organization of ventral striatal ensemble activity: reward selectivity and phase-locking to striatal gamma oscillations. *Soc Neurosci Abstr* 490.12/TT57, 2008.

Kalenscher T, Pennartz CMA. Is a bird in the hand worth two in the future? The neuroeconomics of intertemporal decision-making. *Prog Neurobiol* 84: 284–315, 2008.

Khamassi M, Mulder AB, Tabuchi E, Douchamps V, Wiener SI. Anticipatory reward signals in ventral striatal neurons of behaving rats. *Eur J Neurosci* 28: 1849–1866, 2008.

Knutson B, Cooper JC. Functional magnetic resonance imaging of reward prediction. *Curr Opin Neurol* 18: 411–417, 2005.

Koob GF, Le Moal M. Drug abuse: hedonic homeostatic dysregulation. *Science* 278: 52–58, 1997.

Koya E, Golden SA, Harvey BK, Guez-Barber DH, Berkow A, Simmons DE, Bossert JM, Nair SG, Uejima JL, Marin MT, Mitchell TB, Farquhar D, Ghosh SC, Mattson BJ, Hope BT. Targeted disruption of cocaine-activated nucleus accumbens neurons prevents context-specific sensitization. *Nat Neurosci* 12: 1069–1073, 2009.

Lansink CS, Goltstein PM, Lankelma JV, Joosten RN, McNaughton BL, Pennartz CMA. Preferential reactivation of motivationally relevant information in the ventral striatum. *J Neurosci* 28: 6372–6382, 2008.

Lansink CS, Goltstein PM, Lankelma JV, McNaughton BL, Pennartz CMA. Hippocampus leads ventral striatum in replay of place-reward information. *PLoS Biol* 7: e1000173, 2009.

Lauwereyns J, Watanabe K, Coe B, Hikosaka O. A neural correlate of response bias in monkey caudate nucleus. *Nature* 418: 413–417, 2002.

Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser MB. Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* 309: 619–623, 2005.

Masimore B, Schmitzer-Torbert NC, Kakalios J, Redish AD. Transient striatal gamma local field potentials signal movement initiation in rats. *Neuroreport* 16: 2021–2024, 2005.

McBride WJ, Murphy JM, Ikemoto S. Localization of brain reinforcement mechanisms: intracranial self-administration and intracranial place-conditioning studies. *Behav Brain Res* 101: 129–152, 1999.

McClure SM, Berns GS, Montague PR. Temporal prediction errors in a passive learning task activate human striatum. *Neuron* 38: 339–346, 2003.

McNaughton BL, Battaglia FP, Jensen O, Moser EI, Moser MB. Path integration and the neural basis of the "cognitive map." *Nat Rev Neurosci* 7: 663–678, 2006.

Mogenson GJ, Jones DL, Yim CY. From motivation to action: functional interface between the limbic system and the motor system. *Prog Neurobiol* 14: 69–97, 1980.

Montgomery SM, Buzsáki G. Gamma oscillations dynamically couple hippocampal CA3 and CA1 regions during memory task performance. *Proc Natl Acad Sci USA* 104: 14495–14500, 2007.

Moser EI, Kropff E, Moser MB. Place cells, grid cells, and the brain's spatial representation system. *Annu Rev Neurosci* 31: 69–89, 2008.

O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* 304: 452–454, 2004.

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

O'Keefe J, Nadel L. *The Hippocampus as a Cognitive Map*. Oxford, UK: Oxford Univ. Press, 1978.

Pastalkova E, Itskov V, Amarasingham A, Buzsáki G. Internally generated cell assembly sequences in the rat hippocampus. *Science* 321: 1322–1327, 2008.

Paxinos G, Watson C. *The Rat Brain in Stereotactic Coordinates.* New York: Academic Press, 1996.

Pennartz CM, Groenewegen HJ, Lopes da Silva FH. The nucleus accumbens as a complex of functionally distinct neuronal ensembles: an integration of behavioural, electrophysiological and anatomical data. *Prog Neurobiol* 42: 719–761, 1994.

Pennartz CM, Lee E, Verheul J, Lipa P, Barnes CA, McNaughton BL. The ventral striatum in off-line processing: ensemble reactivation during sleep and modulation by hippocampal ripples. *J Neurosci* 24: 6446–6456, 2004.

Robbins TW, Everitt BJ. Neurobehavioural mechanisms of reward and motivation. *Curr Opin Neurobiol* 6: 228–236, 1996.

Robinson DL, Carelli RM. Distinct subsets of nucleus accumbens neurons encode operant responding for ethanol versus water. *Eur J Neurosci* 28: 1887–1894, 2008.

Robinson TE, Berridge KC. The neural basis of drug craving: an incentive-sensitization theory of addiction. *Brain Res Brain Res Rev* 18: 247–291, 1993.

Roitman MF, Wheeler RA, Carelli RM. Nucleus accumbens neurons are innately tuned for rewarding and aversive taste stimuli, encode their predictors, and are linked to motor output. *Neuron* 45: 587–597, 2005.

Salamone JD, Wisniecki A, Carlson BB, Correa M. Nucleus accumbens dopamine depletions make animals highly sensitive to high fixed ratio requirements but do not impair primary food reinforcement. *Neuroscience* 105: 863–870, 2001.

Samejima K, Ueda Y, Doya K, Kimura M. Representation of action-specific reward values in the striatum. *Science* 310: 1337–1340, 2005.

Schultz W, Apicella P, Scarnati E, Ljungberg T. Neuronal activity in monkey ventral striatum related to the expectation of reward. *J Neurosci* 12: 4595–4610, 1992.

Schultz W, Tremblay L, Hollerman JR. Reward processing in primate orbitofrontal cortex and basal ganglia. *Cereb Cortex* 10: 272–284, 2000.

Self DW, Nestler EJ. Molecular mechanisms of drug reinforcement and addiction. *Annu Rev Neurosci* 18: 463–495, 1995.

Senior TJ, Huxter JR, Allen K, O'Neill J, Csicsvari J. Gamma oscillatory firing reveals distinct populations of pyramidal cells in the CA1 region of the hippocampus. *J Neurosci* 28: 2274–2286, 2008.

Setlow B, Schoenbaum G, Gallagher M. Neural encoding in ventral striatum during olfactory discrimination learning. *Neuron* 38: 625–636, 2003.

Sharott A, Moll CK, Engler G, Denker M, Grun S, Engel AK. Different subtypes of striatal neurons are selectively modulated by cortical oscillations. *J Neurosci* 29: 4571–4585, 2009.

Singer W. Synchronization of cortical activity and its putative role in information processing and learning. *Annu Rev Physiol* 55: 349–374, 1993.

Singer W. Consciousness and the structure of neuronal representations. *Philos Trans R Soc Lond B Biol Sci* 353: 1829–1840, 1998.

Singh T, Brown PL, Mullins SE, Schoenbaum G, Roesch M. Decision-related activity in ventral striatum reflects value and direction. *Soc Neurosci Abstr* 98.11, 2008.

Srinivasan R, Russell DP, Edelman GM, Tononi G. Increased synchronization of neuromagnetic responses during conscious perception. *J Neurosci* 19: 5435–5448, 1999.

Tabuchi ET, Mulder AB, Wiener SI. Position and behavioral modulation of synchronization of hippocampal and accumbens neuronal discharges in freely moving rats. *Hippocampus* 10: 717–728, 2000.

Taverna S, Canciani B, Pennartz CM. Membrane properties and synaptic connectivity of fast-spiking interneurons in rat ventral striatum. *Brain Res* 1152: 49–56, 2007.

Tononi G, Srinivasan R, Russell DP, Edelman GM. Investigating neural correlates of conscious perception by frequency-tagged neuromagnetic responses. *Proc Natl Acad Sci USA* 95: 3198–3203, 1998.

Tran AH, Tamura R, Uwano T, Kobayashi T, Katsuki M, Ono T. Dopamine D1 receptors involved in locomotor activity and accumbens neural responses to prediction of reward associated with place. *Proc Natl Acad Sci USA* 102: 2117–2122, 2005.

van der Meer MA, Redish AD. Covert expectation-of-reward in rat ventral striatum at decision points. *Front Integr Neurosci* 3: 1, 2009a.

van der Meer MA, Redish AD. Low and high gamma oscillations in rat ventral striatum have distinct relationships to behavior, reward, and spiking activity on a learned spatial decision task. *Front Integr Neurosci* 3: 9, 2009b.

Voorn P, Gerfen CR, Groenewegen HJ. Compartmental organization of the ventral striatum of the rat: immunohistochemical distribution of enkephalin, substance P, dopamine, and calcium-binding protein. *J Comp Neurol* 289: 189–201, 1989.

Voorn P, Vanderschuren LJ, Groenewegen HJ, Robbins TW, Pennartz CM. Putting a spin on the dorsal-ventral divide of the striatum. *Trends Neurosci* 27: 468–474, 2004.

Wan X, Peoples LL. Firing patterns of accumbal neurons during a pavlovian-conditioned approach task. *J Neurophysiol* 96: 652–660, 2006.

Wickens JR, Budd CS, Hyland BI, Arbuthnott GW. Striatal contributions to reward and decision making: making sense of regional variations in a reiterated processing matrix. *Ann NY Acad Sci* 1104: 192–212, 2007.

Wilson MA, McNaughton BL. Dynamics of the hippocampal ensemble code for space. *Science* 261: 1055–1058, 1993.

Wise RA. Addictive drugs and brain stimulation reward. *Annu Rev Neurosci* 19: 319–340, 1996.

Wise RA. Drug-activation of brain reward pathways. *Drug Alcohol Depend* 51: 13–22, 1998.

Wood ER, Dudchenko PA, Eichenbaum H. The global record of memory in hippocampal neuronal activity. *Nature* 397: 613–616, 1999.

Yanagihara S, Izawa E, Koga K, Matsushima T. Reward-related neuronal activities in basal ganglia of domestic chicks. *Neuroreport* 12: 1431–1435, 2001.

Zaborszky L, Alheid GF, Beinfeld MC, Eiden LE, Heimer L, Palkovits M. Cholecystokinin innervation of the ventral striatum: a morphological and radioimmunological study. *Neuroscience* 14: 427–453, 1985.