

Habit learning is associated with major shifts in frequencies of oscillatory activity and synchronized spike firing in striatum

Mark W. Howe¹, Hisham E. Atallah¹, Andrew McCool, Daniel J. Gibson, and Ann M. Graybiel²

McGovern Institute for Brain Research, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139

Contributed by Ann M. Graybiel, August 15, 2011 (sent for review July 15, 2011)

Rhythmic brain activity is thought to reflect, and to help organize, spike activity in populations of neurons during on-going behavior. We report that during learning, a major transition occurs in task-related oscillatory activity in the ventromedial striatum, a striatal region related to motivation-dependent learning. Early on as rats learned a T-maze task, bursts of 70- to 90-Hz high- γ activity were prominent during T-maze runs, but these gradually receded as bursts of 15- to 28-Hz β -band activity became pronounced. Populations of simultaneously recorded neurons synchronized their spike firing similarly during both the high- γ -band and β -band bursts. Thus, the structure of spike firing was reorganized during learning in relation to different rhythms. Spiking was concentrated around the troughs of the β -oscillations for fast-spiking interneurons and around the peaks for projection neurons, indicating alternating periods of firing at different frequencies as learning progressed. Spike-field synchrony was primarily local during high- γ -bursts but was widespread during β -bursts. The learning-related shift in the probability of high- γ and β -bursting thus could reflect a transition from a mainly focal rhythmic inhibition during early phases of learning to a more distributed mode of rhythmic inhibition as learning continues and behavior becomes habitual. These dynamics could underlie changing functions of the ventromedial striatum during habit formation. More generally, our findings suggest that coordinated changes in the spatiotemporal relationships of local field potential oscillations and spike activity could be hallmarks of the learning process.

basal ganglia | beta rhythm | gamma rhythm | neural circuit reorganization | network dynamics

Rhythmic field-potential activities in particular frequency bands are known to characterize different on-going behavioral states: low frequency (α -band) rhythms distinguish sleep from waking, higher (β -band) frequencies distinguish movement from rest, and high (γ -band) frequencies distinguish attentive from inattentive states. These activities in different frequency bands can be coordinated by phase or amplitude coupling, and especially when synchronous, such oscillations are proposed to represent on-line processing modes that can facilitate interactions across brain regions and binding of neural ensembles, and that can provide a temporal structure for neural activity related to memory processing and the creation or maintenance of brain states (1–4).

It is well documented that spike activity can be organized relative to particular rhythms detected in local field potentials (LFPs), and can even be dependent on such rhythms (5, 6), yet little is known about how such oscillatory activity changes across learning, a dynamic state in which spike activity is reorganized. We explored the possibility that changes in the oscillatory structure of a network might reorganize local circuit activity as a behavior is learned and stamped in as habitual. We took as our starting point that learning-related changes in both θ -band and γ -band activity have been documented in animals learning associative and navigational tasks (7–11), and we capitalized on the fact that a dynamic progression of activity is thought to occur in striatal circuits during learning, with the ventromedial striatum

being required for the initial acquisition of reward-motivated behaviors (12, 13). We trained and then overtrained rats on an associative T-maze task, and during the entire learning period we recorded spike and LFP activity chronically in the ventromedial striatum. We found that as learning proceeds, there is a marked transition in the dominant mode of on-going oscillatory activity from primarily locally synchronous high- γ to widespread β -activity, especially near trial end. Remarkably, populations of both interneurons and projection neurons transiently synchronized their spiking to brief bursts of γ - and β -activity, and did so inversely, with interneurons spiking near oscillation troughs and projection neurons spiking near oscillation peaks. These findings suggest that there is a change in the spatiotemporal dynamics of on-going oscillatory and spike activity in the ventromedial striatum as learning progresses, with transient bouts of local, γ -synchronized inhibition occurring frequently at task completion early in learning, and transient bouts of global, β -synchronized inhibition taking over as a successful behavior becomes habitual. These coupled oscillatory and spiking patterns may underlie the crystallization of striatal circuit changes that occurs as initial exploratory activity yields to habitual behavior through learning.

Results

β -band (15–28 Hz) LFP Oscillations in the Ventromedial Striatum Increase Transiently at Goal-Reaching During Habitual Performance of an Associative T-Maze Task. We recorded spike and LFP activity bilaterally in the ventromedial part of the striatum (here called VMS) as seven rats were trained to perform the T-maze task shown in Fig. 1A (14). After turning left or right according to auditory instruction cues, the rats received a chocolate reward at the end of the indicated end-arm (see *Methods*). Tetrodes were confirmed to be in the VMS, mainly within the nucleus accumbens core (Fig. 1A and Fig. S1).

Task-related modulation of oscillatory activity occurred in multiple frequency ranges in the session-averaged LFPs recorded during training (Fig. 1B and C). Particularly prominent was a transient increase in power in the β -band (15–28 Hz) that occurred directly after goal-reaching (Fig. 1B). The increase was strong on correctly performed trials but weak on incorrect trials (Fig. 1B and C); it did not exhibit a clear relationship to either running speed or acceleration (Fig. 1C).

To investigate this β -band activity, we examined raw and band-pass filtered LFP traces from single trials. Remarkably, the extended period of increased β -power visible in the session-averaged spectrograms after goal-reaching did not represent continuous

Author contributions: M.W.H., H.E.A., and A.M.G. designed research; H.E.A. performed research; M.W.H., H.E.A., A.M., D.J.G., and A.M.G. analyzed data; and M.W.H. and A.M.G. wrote the paper.

The authors declare no conflict of interest.

¹M.W.H. and H.E.A. contributed equally to this work.

²To whom correspondence should be addressed. E-mail: graybiel@mit.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1113158108/-DCSupplemental.

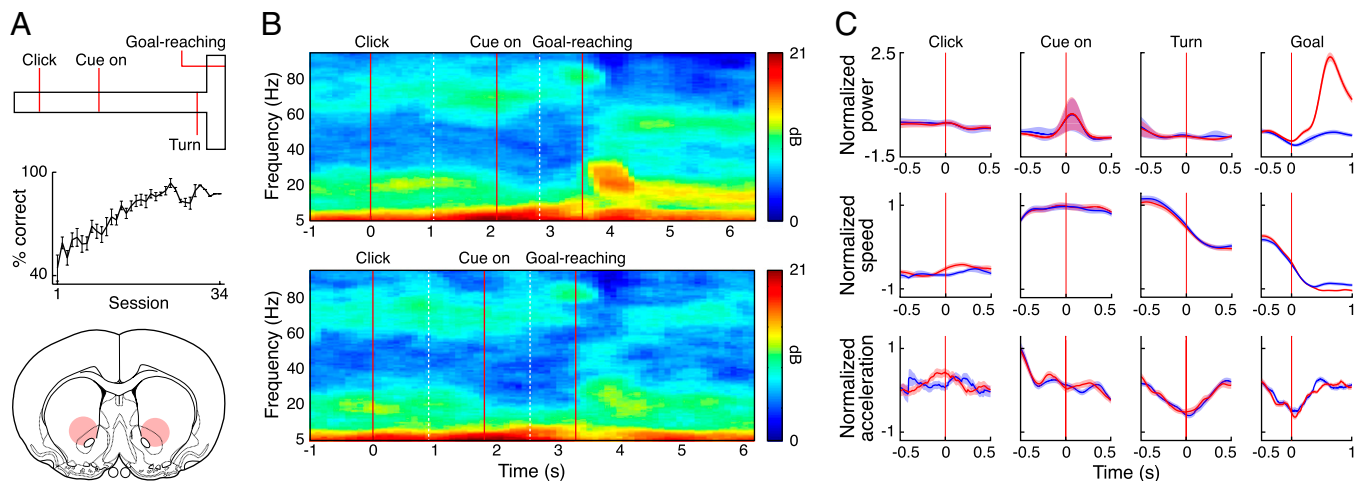


Fig. 1. β -Power in ventral striatum increases after goal-reaching on correct trials. (A) (Top) The T-maze task. (Middle) Average learning curve for all animals ($n = 7$). (Bottom) Region sampled by tetrode recordings (red shading). (B) Trial-averaged spectrograms for correct trials ($n = 22$ trials, Upper) and for incorrect trials ($n = 18$ trials, Lower) for a representative 40-trial session. White dashed lines indicate edges of concatenated peri-event windows. (C) (Top) The z-score normalized β -power (15–28 Hz) averaged across all sessions and recording tetrodes for correct trials (red) and incorrect trials (blue). (Middle) Average z-score normalized run-speed. (Bottom) Average z-score normalized acceleration. Shading, SEMs computed across sessions.

β -band activity, but instead, corresponded to brief bursts of high-amplitude β -activity, lasting about two to four cycles (~ 100 – 200 ms), which occurred at slightly different times in different trials.

To determine whether this transient β -burst activity was related to reward delivery on the correct runs, we performed two control experiments in which we manipulated the primary reward. In one ($n =$ two rats over 13 sessions, with chocolate milk reward), we delayed the chocolate milk delivery by 2 s (Fig. 2A). The increase in β -power after goal-reaching did not shift forward in time, but instead, continued to occur just after goal-reaching (Fig. 2A, Lower). However, β -power in incorrect trials, nearly absent before the delay sessions, became stronger during the trials with reward delay (Fig. 2A).

In the second control, after 10 d of overtraining on the task, we suddenly changed the identity of the primary reward at one end-arm of the T-maze from chocolate milk to strawberry milk reward ($n =$ one rat over six sessions) (Fig. 2B and C). Despite the fact that the rat had been preexposed to strawberry milk in his home cage and drank it readily, there was a large ($P < 0.001$, two-tailed t test) decrease in postgoal β -power on correct trials for both end-arms (Fig. 2B and C). This decrease was significantly stronger for the end-arm baited with strawberry milk than for the chocolate milk end-arm ($P < 0.01$, two-tailed t test) (Fig. 2C) and could not be accounted for by changes in the run speed of the animal (Fig. 2C). After 5 consecutive days of exposure to the new reward, the β -power had rebounded nearly to preswitch levels (Fig. 2B and C). Collectively, these results indicate that

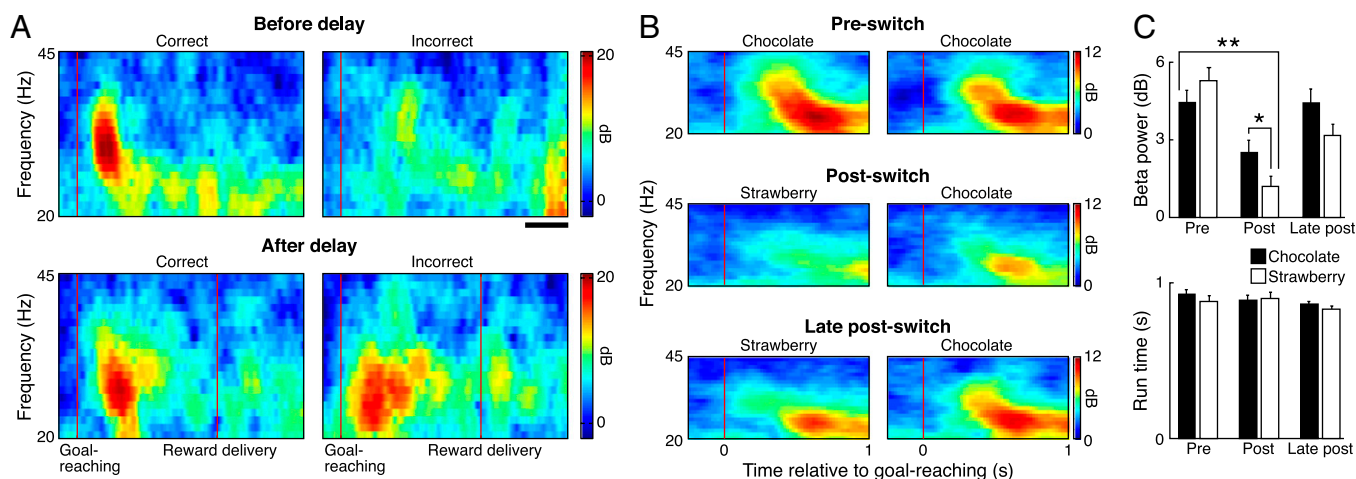


Fig. 2. Increases in β -power after goal-reaching are not dependent on the presence of primary reward, but are influenced by changes in the primary reward. (A) Reward-delay control. Session-averaged β -power for correct (Left) and incorrect (Right) trials during a standard 40-trial session (Upper) and during a control session in which reward was delayed for 2 s after goal-reaching (Lower). (Scale bar, 0.5 s) (B) Reward-switch control. (Top) β -Power averaged across correct trials on the left and right sides of the maze (Left and Right columns, respectively) during two sessions averaged just before switching the primary reward. (Middle) Average β -power after goal-reaching for two sessions averaged immediately after primary reward on the left arm of the maze was switched from chocolate milk to strawberry milk. (Bottom) Average β -power for two sessions averaged 5 d after the switch to the strawberry milk as reward. (C) Average β -power during the period from 0 to 0.8 s after goal-reaching (Upper) and average run times from turn offset to goal-reaching (Lower) for the sessions shown in B. Error bars represent SEM. * $P < 0.01$, ** $P < 0.001$.

postgoal β -power is not tied to the receipt of a particular primary reward at a particular time, but could reflect an internal state that is modifiable by unexpected changes in the primary reward.

During Behavioral Learning, High- γ LFP Oscillations in the VMS Weaken as β -Band LFP Oscillations Strengthen. We analyzed the entire set of data recorded during training to determine whether the patterns of oscillatory LFP activity changed across learning (Fig. 3). Both β -band and γ -band trial-averaged activities were composed, when examined trial-by-trial, of transient bursts of oscillations lasting one to several cycles (Fig. 3 *A, B, E*, and *F*). The occurrence of trial-end β -activity was highly experience-dependent and was accompanied by inversely changing high γ -band activity in the LFPs (Fig. 3 *C* and *D*, and Fig. S2). High- γ power was strong both before and after goal-reaching during initial training and the first days of overtraining, but diminished significantly as performance improved and became more restricted in duration (Fig. 3 *C* and *G*) ($R = -0.28, P < 0.001$ before, $R = -0.17, P = 0.05$ after). In sharp contrast, the rise in β -band power after goal-reaching was scarcely detectable, even on correct tri-

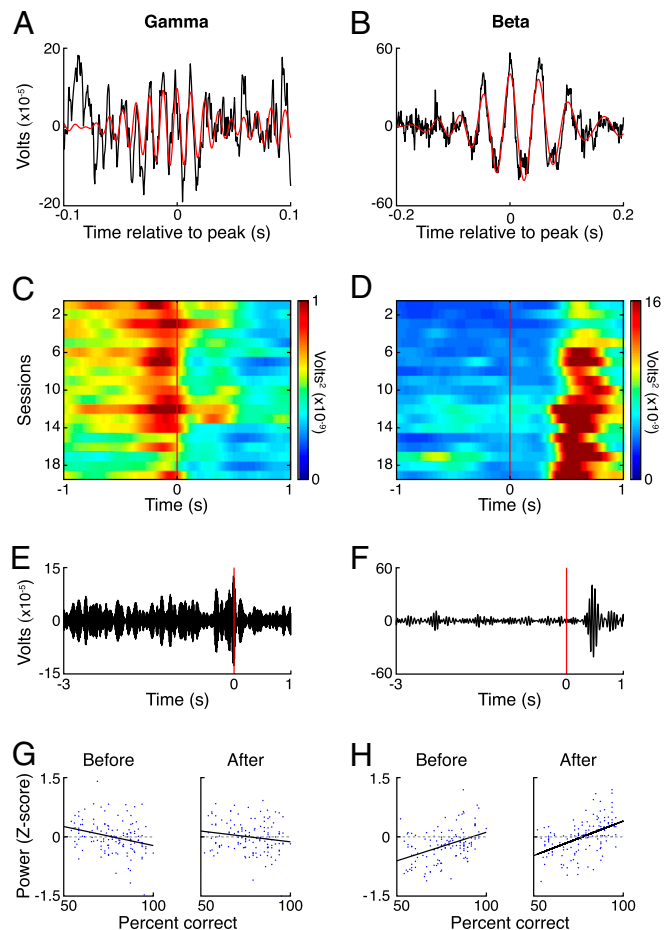


Fig. 3. β -Power during task performance increases with learning, whereas γ -power decreases with learning. (*A* and *B*) Representative examples of single high-amplitude γ - (*A*) and β - (*B*) bursts. The raw LFP signal is shown in black, and the band-pass filtered signal is overlaid in red. (*C* and *D*) γ - (*C*) and β - (*D*) power around goal reaching for all sessions run by one representative rat. Each row represents the average power over 18 randomly selected correct trials. (*E* and *F*) Representative band-pass filtered traces from single trials showing task-related emergence of high-amplitude γ - (*E*) and β - (*F*) bursts. (*G* and *H*) Mean z-score normalized γ - (*G*) and β - (*H*) power before goal-reaching (*Left*) and after goal-reaching (*Right*) as a function of percent correct for all sessions run by all rats ($n = 7$) combined.

als, at the beginning of training (Fig. 3*D*). This β -band activity did not become prominent during the correct runs until about the time the rats reached the learning criterion, and it then continued to increase throughout the overtraining period (Fig. 3 *D* and *H*) ($R = 0.53, P < 0.002$). β -Band power before goal-reaching was much weaker than that after goal-reaching (Figs. 1*B* and 3*D*) but also was positively correlated with percent-correct performance (Fig. 3*H*) ($R = 0.42, P < 0.0002$). These findings demonstrate that high- γ power during T-maze performance significantly decreases with learning in the VMS, but simultaneously recorded β -power increases with learning, especially just after goal-reaching.

Spike Firing of Medium Spiny Neurons and Fast-Spiking Interneurons Are Synchronized with High-Amplitude β - and γ -LFP Oscillations but Have Opposite Phase Relationships. We found that the spike activities of both putative fast-spiking interneurons (FSIs) and putative projection neurons recorded (identified as in *Methods* and Fig. S3) were modulated during identified high γ - or β -bursts. Nearly 50% of the FSIs ($n = 163$ of 332) were significantly modulated (Rayleigh's test, $P < 0.05$) during high-amplitude β -bursts, and similar numbers ($n = 161$ of 332) were modulated around high-amplitude γ -bursts (Rayleigh's test, $P < 0.05$). Just over 50% of the β -modulated FSIs (98 of 163) were modulated by both β - and high- γ rhythms. Remarkably, as a population, the spike firing of the FSIs was transiently synchronized to peak near the troughs of both the high-amplitude γ -oscillations and the high-amplitude β -oscillations (Fig. 4*A* and *B*). Given the transition from high- γ to β -band prominence in the postgoal oscillatory bursts, these findings indicate that, as a population, the FSIs tended to fire near the phase troughs of high- γ oscillations early in learning but near the phase troughs of β -oscillations late in learning.

In the dorsal striatum, neurons exhibiting high firing rates have been found to inhibit projection neurons *in vitro* (15), but few such direct inhibitory interactions have been reported *in vivo* (16). We reasoned that if the FSIs that we recorded inhibited

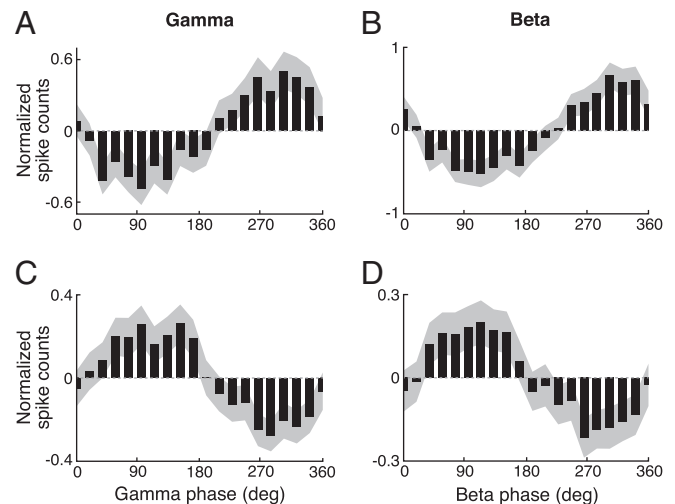


Fig. 4. Spikes of putative FSIs and medium spiny neurons are inversely but coordinately modulated during high-amplitude β - and γ -oscillations. (*A*) Mean z-score normalized spike-phase histograms averaged for all significantly modulated FSIs ($n = 161$) around periods of high-amplitude γ -oscillations. Zero phase is defined as the positive-going zero crossing. (*B*) Spike-phase histograms constructed as those in *A* for significantly β -modulated FSIs ($n = 163$) around periods of high-amplitude β -oscillations. (*C*) Spike phase histograms for all γ -modulated projection neurons ($n = 488$) around periods of high-amplitude γ -oscillations. (*D*) Spike-phase histograms as in *C* for all β -modulated projection neurons ($n = 783$) around high-amplitude β . Shaded lines in *A* to *D*, $2 \times$ SEM.

projection neurons, then the spikes of the projection neurons might also be synchronized to the β -burst oscillations, but with opposite phase. This finding was true for many of the projection neurons. Over a quarter of these neurons showed spike-field modulation (about 28%, $n = 783$ of 2,796 units) during β -bursts and about 17% (488 of 2,796) during γ -bursts, and the phase distribution of their spikes was opposite to that of modulated FSI spiking during β -bursts: the projection neurons fired near the positive peaks in the LFP (Fig. 4 C and D), regardless of the oscillatory frequency. This antiphase relationship could reflect intermittent inhibition of projection neurons by FSIs.

Antiphase Spiking of FSIs Is Synchronized Locally During γ -Bursting but Globally During β -Bursting. We next analyzed the spatial distribution of synchronization patterns of the β - and γ -bursts recorded on different tetrodes in the VMS. β -Bursts were highly synchronous across the recording sites (Fig. 5 A and B). Synchrony was much stronger during high-amplitude β -bursts than during random periods outside of β -bursts (Fig. 5B). In contrast, during high-amplitude γ -bursts, synchrony across tetrodes spaced through the VMS was comparatively weak (Fig. 5 C and D). Thus, the γ -activity that appeared early during behavioral learning appeared to be predominantly local, whereas the β -activity that appeared later during learning appeared to be more widely distributed across the VMS territory explored.

This spatial disparity in synchrony between β - and γ -bursting also held for spiking (Fig. 6 A and B). Modulated FSIs displayed a robust oscillation when aligned to peaks of both β - and γ -bursts on the same tetrode from which the spikes were recorded. When spikes were aligned to peaks of γ -bursts recorded on other spatially separated tetrodes, these synchronous oscillations were significantly reduced around the oscillation peaks (Fig. 6A) ($P < 0.0001$, two-tailed t test on baseline-subtracted means around peak). By contrast, the spike-field synchrony remained robust when we aligned the spikes of β -modulated units to peaks of β -bursts

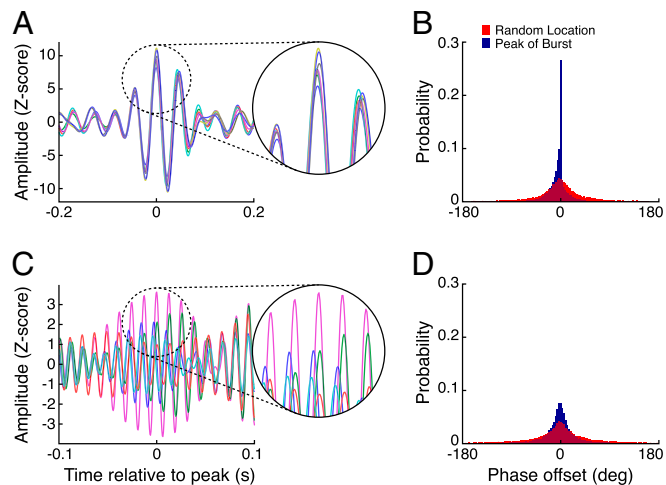


Fig. 5. Bursts of high-amplitude β -oscillations in the LFP are globally synchronized, whereas bursts of high- γ oscillations are more spatially localized. (A) Overlaid β -band-pass filtered traces from eight LFPs simultaneously recorded on spatially separated tetrodes in the VMS around a high-amplitude β -burst in a single trial. Traces are centered on the peak of the highest amplitude signal. *Insets* show tight alignment of the peaks on different channels. (B) Histogram of phase offsets on all LFP channels relative to the peak of the highest-amplitude LFP trace (reference phase = 90°) during all identified β -bursts (> 2.5 SD above mean, blue) and during an equal number of randomly selected nonburst periods during the task (red). Phase differences around 0 indicate strong synchrony across electrodes. (C and D) Equivalent plots as A and B for γ -oscillations.

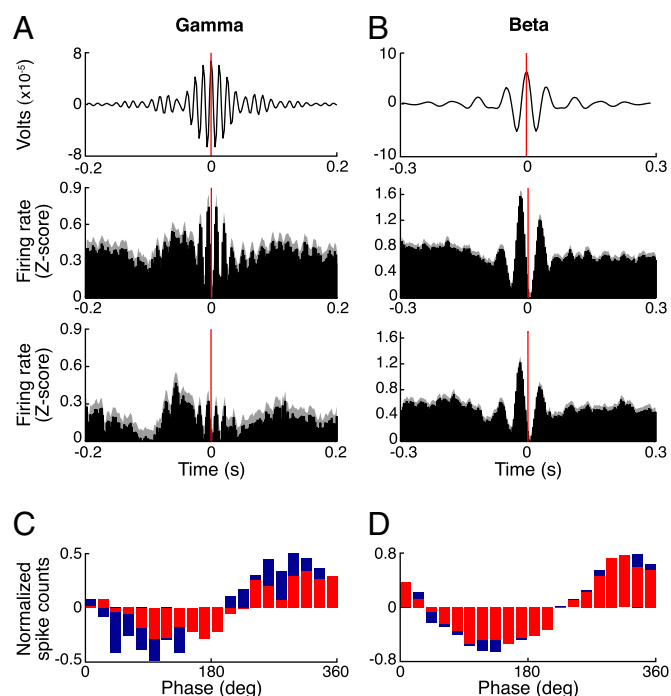


Fig. 6. FSIs synchronize globally during high-amplitude β -bursts, but are more weakly spatially synchronized during γ -bursts. (A) (Top) Average band-pass filtered γ -trace used to construct histograms below. (Middle) Normalized spike histograms of all γ -modulated FSIs aligned on the peaks of all high-amplitude γ -bursts on the local tetrode on which the spikes were recorded. (Bottom) Normalized spike histogram as above aligned to peaks of high-amplitude γ -bursts recorded on a randomly selected tetrode different from the one the spikes were recorded. Shaded bars, SEM. (B) Normalized firing-rate histograms, as in A, for high-amplitude β -bursts. (C and D) Spike-phase histograms, as in Fig. 4, constructed around γ - (C) and β - (D) bursts recorded on the same electrode as the spiking (blue) and randomly selected tetrodes spatially separated from the tetrode the unit was recorded on during local burst times (red).

on nonlocal tetrodes (Fig. 6B) ($P = 0.48$, two-tailed t test). These findings could help to account for why some studies have reported the presence of FSI synchrony, but others have not (17, 18). Global synchrony of FSIs may occur only transiently and rarely, prominently during the 100- to 200-ms periods in which β -bursts occur.

To quantify this spatial effect on spike-field synchronization, we constructed spike-phase histograms, as illustrated in Fig. 4, but we plotted the spikes relative to nonlocal band-pass filtered LFPs during the high-amplitude β - and γ -events (Fig. 6 C and D). Nonlocal spike-phase relationships during high-amplitude γ -oscillations were significantly weaker than local spike-phase relationships (Fig. 6C) ($P < 0.01$, Kruskal-Wallis test on the absolute values of mean phase distributions). However, such reductions did not occur for the spikes synchronized to the high-amplitude β -oscillations recorded (Fig. 6D) ($P = 0.57$, Kruskal-Wallis test); they showed a nearly equivalent magnitude of spike-phase modulation for local and nonlocal LFPs. Finally, spike-spike coherence analysis of pairs of simultaneously recorded FSIs demonstrated a significant coherence peak around the β -frequency during high-amplitude β -bursting (Fig. S4). These findings suggest that the brief bursts of β -band oscillation represent periods of global synchrony in the VMS network, whereas high-amplitude γ -oscillations primarily represent local synchrony.

Discussion

Our findings demonstrate that as animals learn a simple T-maze task, the temporal and spatial organization of LFP oscillations in

the VMS undergoes a major transition in which bursts of oscillatory activity in the high- γ range occurring near the completion of successful maze runs are supplanted by β -band bursts. Populations of striatal interneurons and projection neurons synchronize to opposite phases of both of these rhythms, with the consequence that over the course of learning, spike-field synchrony also shifts frequency from γ - to β -bands. As our evidence favors spatial heterogeneity of γ -burst synchronies early in learning, and more distributed and uniform β -burst synchronies late in learning, our findings suggest that both the spatial and temporal structure of neural processing in the ventral striatum profoundly changes during learning.

β -Band Activity in the VMS Occurs in Transient Bursts During Which Spiking of Fast-Firing Interneurons and Projection Neurons Transiently Become Synchronized with Opposite Phase Alignment. Our findings establish that β -band oscillations in the VMS occur predominantly in brief, approximately 100- to 200-ms bursts. We found no evidence for prolonged periods of β -oscillations when we analyzed the data trial by trial, consonant with findings by Dejean et al. (19). Remarkably, the β -bursts that we observed directly corresponded to transient episodes in which the spikes of striatal FSIs and projection neurons were synchronized to the β -band oscillations. Spiking of the FSIs was concentrated near the troughs of the β -bursts, whereas spiking of the projection neurons was strongest near the peaks. Because FSIs are known to make monosynaptic inhibitory connections with projection neurons, but no equivalent functional connection has been found from projection neurons to interneurons, at least in the dorsal striatum (20), it is likely that the alternating periods of firing at the β -frequency are generated by an interneuron-driven suppression of projection neurons. Thus, the LFP β -bursts, lasting one to a few cycles of 15- to 28-Hz oscillation, likely correspond to fleeting episodes in which the main output neurons of the VMS are rhythmically and synchronously suppressed but populations of inhibitory FSIs are synchronously excited.

The synchronization of striatal interneurons and projection neurons associated with β -bursts was experience-dependent. The bursts of β -band LFP oscillations developed as behavioral learning of the task advanced. Moreover, the β -bursts occurred with highest probability and amplitude at task-end, and occurred only on correctly performed rewarded trials. Reward likelihood was thus clearly critical for the postgoal β -bursts. The nature of these reward effects was not simple, however. The presence and identity of the reward modulated the strength of the bursts, but the β -bursts were not tightly time-locked to reward delivery itself, and they were not time-locked to the licking or consumption of the reward. It seems more likely that the bursts, and the accompanying synchronizations of spiking, mark the expected task-end when the animals have learned the task well and are in a habitual mode of successful performance.

These characteristics suggest that the prominent β -bursts that occur at task-end represent brief neural events in which striatal networks can be reconfigured and reset to reflect successful completion of a learned behavior. This idea could be seen as fitting the proposal that β -band oscillations at trial end serve a function in maintaining the status quo (4). β -bursts and associated spike synchronization did occur at other task-times, however. The timing of these bursts was more variable than that of the task-end β -bursts (Fig. S5). We emphasize here that it could be critical to the function of these β -bursts that they are associated with transient and coordinated synchronization of different sets of striatal network neurons: one set—the FSIs—synchronizing near the troughs of the β -oscillations, and another set—the output neurons—synchronizing near their peaks.

γ -Band Activity and β -Band Activity Are Inversely Modulated Through the Course of Learning. Our findings confirm those of van der

Meer and Redish (8) that high- γ oscillations decrease as learning proceeds. By making detailed comparisons of the high- γ bursts and β -band bursts that we recorded around goal-reaching, we found that there was an inverse relation between the two over the course of learning. γ -Power was strongest early in learning and became weaker during overtraining, whereas β -power was relatively weak early on but became progressively stronger with training. Strikingly, the spikes of the FSIs and projection neurons had the same inverse phase relations to the γ -band oscillations as they had to the β -band oscillations: the FSIs spiked near the phase troughs of γ -oscillations, and the projection neurons near the phase peaks. Frequency modulation was not exclusive for individual cells as some were modulated by both β - and γ -rhythms. Nevertheless, the relationship of the neuronal spiking to the two different oscillations suggests that there was a major switch from spike-timing regulation during high- γ bursts early in training to spike-timing regulation during β -bursts later in training, after behavioral accuracy had reached asymptote.

During Habit Learning, a Shift Occurs in Network Processing from Local to Global Spike-LFP Synchronization. We found a marked difference between the spatial patterns of synchrony for the high- γ and β -band oscillations. The early-appearing γ -rhythms were synchronized relatively weakly across recording sites, in agreement with observations by Kalenscher et al. (21), whereas the late-appearing β -oscillations were broadly synchronized across tetrode recording sites spread across the VMS. This difference in spatial structure of the synchrony patterns could have major implications for the accompanying neural processing. As shown in the schematic model of Fig. 7, the spatiotemporal patterning suggested by our findings is one in which transient, local γ -synchronized spike-field patterns are prominent during early learning, whereas transient bouts of β -synchronized spike-field synchrony, strongest at task-end, are widely coordinated across the VMS late in learning.

What could be the mechanistic significance of this transition from local γ -synchronization to more global β -synchronization of spike firing as a result of learning? One possibility, raised by

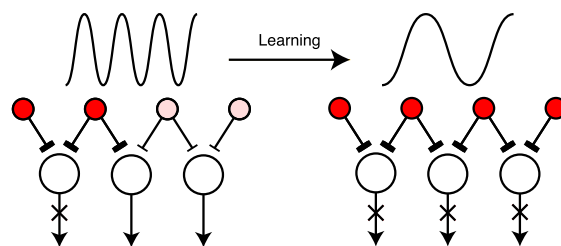


Fig. 7. Model for changes in spatiotemporal dynamics of oscillatory LFP activity and spiking in the VMS during habit learning. The left column represents activity patterns early in learning, when high-amplitude γ -bursts in LFPs recorded in the VMS are prominent at task-end. The right column represents LFP activity patterns after overtraining and habit formation, when β -bursts have become strong at task-end. Large open circles represent projection neurons, and smaller shaded circles represent inhibitory interneurons. The intensity of shading represents the probability of firing at a given time point during a high-amplitude γ - (Left) or β - (Right) event. A projection neuron is inhibited if it receives strong (dark red, dark projection lines) inhibition from two converging interneurons. Early in training, during high- γ bursts, interneurons synchronize with each other only locally in small groups, and thus provide only local, nonconvergent inhibition to small populations of projection neurons. Late in training, during high-amplitude β -bursts, interneurons synchronize their firing across wider spatial areas and across a longer cycle period, and as a consequence, project powerful convergent inhibition to a larger population of projection neurons. The result of this β -range synchrony is a widespread inhibition of ventral striatal output during habit formation.

previous studies (1–4), is that oscillatory activities facilitate, in frequency-dependent ways, communication between interconnected brain regions. By this view, the shift from high- γ to β -burst synchronization could reflect a shift in the functional connectivity of ventral striatal networks during learning.

Another possibility is that the oscillations represent a particular state of microcircuit operation within the striatal network itself. The shift in balance between high- γ and β during learning appears to reflect a shift from a spatially asynchronous mode of transient FSI firing to a mode in which FSIs transiently become globally synchronized. By this view, in the early, exploratory phase of learning, when γ -bursts at task-end are strong, different subsets of projection neurons might independently communicate with local interneurons. This type of network environment could favor plasticity and flexibility. When, later in learning, bursts of β -oscillations emerge at task-end, they could promote a more homogenous network structure and a more nearly fixed pattern of output as behavior becomes habitual.

A third possibility is that the widespread FSI synchrony during β -bursts, by allowing a longer intraburst time for spatial and temporal summation of FSI spikes than that during γ -bursts—and thus a larger pool of interneurons firing together—could promote more effective inhibition of projection neurons through enhanced spatial and temporal summation of the inhibitory postsynaptic currents (Fig. 7). FSIs, which synchronize strongly as a population during β -bursts, fire on average at around the β -frequency (17). A powerful common input to a large population of FSIs, as has been found for some cortico-striatal projections (22, 23), could thus produce a global resetting of FSI firing and transiently synchronize spiking of the entire population at around their intrinsic mean firing frequency. Thus, based on our findings, we predict

that, as learning proceeds, the output of the ventral striatum could become progressively dampened at trial end.

Much evidence suggests that as habits and procedures are acquired, activity in the VMS region, including the nucleus accumbens core, is essential early in the learning process (12, 13). The findings we report here suggest that oscillatory dynamics in the VMS may reflect a process by which initial local network adjustments give way to a more widespread resetting of striatal networks as learning proceeds.

Methods

Procedures were approved by the Massachusetts Institute of Technology Committee on Animal Care and accorded with the National Research Council's Guide for the Care and Use of Laboratory Animals. Male Long Evans rats ($n = 7$) were implanted with 12-tetrode headstages (six per hemisphere) attached to independently movable microdrives (14). Tetrodes were lowered (6–7 mm, about 2 wk) to target sites in the VMS (A/P +1.5 mm, M/L \pm 2.1 mm). Rats were then trained on a T-maze task (40 trials per session). The rats were instructed to turn left or right in response to a tone to receive a chocolate reward at the end of the indicated end-arm. Rats were trained to a criterion of 72.5% correct trials ($P < 0.01$, χ^2 test) and then overtrained for at least 10 consecutive days. Spikes and LFPs were recorded throughout training. LFP spectral analysis was conducted using Chronux algorithms (<http://chronux.org>) and in-house MATLAB software. Detailed methods are available in *SI Methods*.

ACKNOWLEDGMENTS. We thank Christine Keller-McGandy, Alex McWhinnie, and Henry F. Hall for their help, Drs. Richard Courtemanche and Jason Ritt for their comments on the manuscript, and Dr. Joseph Feingold for sharing valuable information about patterns of task-related β -band bursting in macaque monkeys. This work was funded by National Institutes of Health Grant R01 MH060379 and a Mark Gorenberg graduate student fellowship.

- Buzsaki G (2006) *Rhythms of the Brain* (Oxford Univ Press, New York).
- Kopell N, Whittington MA, Kramer MA (2011) Neuronal assembly dynamics in the beta1 frequency range permits short-term memory. *Proc Natl Acad Sci USA* 108:3779–3784.
- Wang XJ (2010) Neurophysiological and computational principles of cortical rhythms in cognition. *Physiol Rev* 90:1195–1268.
- Engel AK, Fries P (2010) Beta-band oscillations—Signalling the status quo? *Curr Opin Neurobiol* 20:156–165.
- Brandon MP, et al. (2011) Reduction of theta rhythm dissociates grid cell spatial periodicity from directional tuning. *Science* 332:595–599.
- Koenig J, Linder AN, Leutgeb JK, Leutgeb S (2011) The spatial periodicity of grid cells is not sustained during reduced theta oscillations. *Science* 332:592–595.
- van der Meer MA, et al. (2010) Integrating early results on ventral striatal gamma oscillations in the rat. *Front Neurosci* 4:300.
- van der Meer MA, Redish AD (2009) 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.
- DeCoteau WE, et al. (2007) Learning-related coordination of striatal and hippocampal theta rhythms during acquisition of a procedural maze task. *Proc Natl Acad Sci USA* 104:5644–5649.
- Popescu AT, Popa D, Paré D (2009) Coherent gamma oscillations couple the amygdala and striatum during learning. *Nat Neurosci* 12:801–807.
- van Wingerden M, Vinck M, Lankelma J, Pennartz CM (2010) Theta-band phase locking of orbitofrontal neurons during reward expectancy. *J Neurosci* 30:7078–7087.
- Atallah HE, Lopez-Paniagua D, Rudy JW, O'Reilly RC (2007) Separate neural substrates for skill learning and performance in the ventral and dorsal striatum. *Nat Neurosci* 10:126–131.
- Hernandez PJ, Sadeghian K, Kelley AE (2002) Early consolidation of instrumental learning requires protein synthesis in the nucleus accumbens. *Nat Neurosci* 5:1327–1331.
- Barnes TD, Kubota Y, Hu D, Jin DZ, Graybiel AM (2005) Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature* 437:1158–1161.
- Koós T, Tepper JM (1999) Inhibitory control of neostriatal projection neurons by GABAergic interneurons. *Nat Neurosci* 2:467–472.
- Lansink CS, Goltstein PM, Lankelma JV, Pennartz CM (2010) Fast-spiking interneurons of the rat ventral striatum: Temporal coordination of activity with principal cells and responsiveness to reward. *Eur J Neurosci* 32:494–508.
- Berke JD (2008) Uncoordinated firing rate changes of striatal fast-spiking interneurons during behavioral task performance. *J Neurosci* 28:10075–10080.
- Gage GJ, Stoetznner CR, Wiltschko AB, Berke JD (2010) Selective activation of striatal fast-spiking interneurons during choice execution. *Neuron* 67:466–479.
- Dejean C, et al. (2011) Power fluctuations in beta and gamma frequencies in rat globus pallidus: Association with specific phases of slow oscillations and differential modulation by dopamine D1 and D2 receptors. *J Neurosci* 31:6098–6107.
- Chuhma N, Tanaka KF, Hen R, Rayport S (2011) Functional connectome of the striatal medium spiny neuron. *J Neurosci* 31:1183–1192.
- Kalenscher T, Lansink CS, Lankelma JV, Pennartz CM (2010) Reward-associated gamma oscillations in ventral striatum are regionally differentiated and modulate local firing activity. *J Neurophysiol* 103:1658–1672.
- Parthasarathy HB, Graybiel AM (1997) Cortically driven immediate-early gene expression reflects modular influence of sensorimotor cortex on identified striatal neurons in the squirrel monkey. *J Neurosci* 17:2477–2491.
- Kita T, Kita H, Kitai ST (1985) Local stimulation induced GABAergic response in rat striatal slice preparations: intracellular recordings on QX-314 injected neurons. *Brain Res* 360:304–310.