Goal-directed shaping of cortical waves

Anton A Dogadov, Daniel E Shulz, Valérie Ego-Stengel, Isabelle Ferezou, Luc Estebanez

Goal-directed shaping of cortical waves

Anton A. Dogadov[#], Daniel E. Shulz, Valérie Ego-Stengel, Isabelle Ferezou*, Luc Estebanez*[#]

Université Paris-Saclay, CNRS, Institut des Neurosciences Paris-Saclay (NeuroPSI), 91400, Saclay, France.

*Co-last authors

[#]Correspondence: luc.estebanez@cnrs.fr, mail@dogadov.fr

ABSTRACT

At the surface of the cerebral cortex, mesoscopic scale brain dynamics are characterized by waves of synchronized neuronal activity. These waves have been shown to impact the processing of sensory information, but can they be actively shaped by the subject in a goal directed manner? To address this question, we designed a fast widefield optical brain-machine interface that enables real-time detection and reinforcement of individual traveling waves that display a specific displacement at the surface of the mouse somatosensory-motor cortex. Mice learned to generate such Conditioned Waves, which became progressively more stereotyped and were preceded by a characteristic suppression of cortical activity. Conditioned Waves were typically related to a limb movement that triggered somatosensory inputs. However, only the movements that occurred in conjunction with the identified suppressive cortical state could drive Conditioned Waves. This work demonstrates that propagating cortical waves can be subject to operant control. It provides evidence for the plasticity and functional relevance of large-scale cortical dynamics, and establishes a new paradigm for manipulating mesoscale brain activity.

INTRODUCTION

During awake, active behavior, the activity of cortical neurons is typically described as dominated by asynchronous, fast fluctuations in membrane potential (Steriade et al. 2001; Poulet and Petersen 2008). However, the advent of high-resolution recording methods applied to large cortical territories has revealed that such “active” cortical state is not uniform but rather shows patterns of synchronous firing resulting in a constant flow of waves travelling across the cortical network (Ma et al. 2016; Mohajerani et al. 2010; Luczak et al. 2009; Shahsavarani et al. 2023). These waves tend to move from sources, often located in the primary sensory areas, towards sinks — some in associative areas (Mohajerani et al. 2013; Liang et al. 2023). Their time averaging during trained behaviors reveals networks of sensory, motor and high-order cortical areas that activate in a highly structured fashion as the behavior unrolls (Allen et al. 2017; C. H. Chen et al. 2017; Gilad et al. 2018; Musall et al. 2019; Aggarwal et al. 2022).

Despite the prominence of these large-scale dynamics during active wakefulness, the effective role of individual wave events remains an open question. One hypothesis is that waves are largely stochastic events. In this view, their shape and displacement may passively follow the paths dictated by the network architecture (Ermentrout and Kleinfeld 2001; Mohajerani et al. 2013); they are not controlled by the subject and do not convey specific information, akin to spike timing in the rate coding interpretation of action potential firing.

Alternatively, individual cortical waves may actively contribute to information processing in the brain. Indeed, cortical state transitions (Crochet and Petersen 2006; Ferezou et al. 2006) and more generally waves (Arieli et al. 1996) directly impact the integration of sensory inputs, leading to either large or very limited evoked activity in response to the same input at the periphery. This contribution of individual waves to information processing is also supported by evidence that they can carry sensory information (Gonzales et al. 2025), and by experiments highlighting that individual cortical waves may be the source or substrate of illusory percepts (L. M. Chen et al. 2003; Jancke et al. 2004). However, so far, there has been no direct evidence that individual cortical waves are indeed a relevant functional component of brain activity that could be actively generated by the nervous system towards the resolution of a behavioral task.

The development of motor brain-machine machine interfacing has emerged as an efficient strategy to interrogate the agency of subjects on internal variables of their brain. In particular, neuronal operant control training has uncovered the ability of a subject to drive the activity of individual neurons (Arduin et al. 2013; Fetz 1969) as well as small clusters of neurons (Koralek et al. 2012; Clancy et al. 2014; Goueytes et al. 2022) and even pairs of distant neuronal clusters (Clancy and Mrsic-Flogel 2021). Here we ask whether cortical spatio-temporal dynamics can be orchestrated at the larger scale of cortical waves to meet behavioral needs. In particular, we tested if mice can learn via operant conditioning to generate individual waves with a specific displacement at the surface of the brain. To this end, we monitored in real-time the cortical waves moving across the primary somatosensory-motor cortical areas of awake, water-restricted, transgenic mice (Emx1-Cre x Ai-95) expressing the GCaMP6f calcium reporter across excitatory neurons. We tracked online the local maxima of the cortical waves, and we conditioned the delivery of water rewards to the occurrence of wave trajectories that moved through two adjacent 500 µm square patches at the surface of the cortex. Simultaneously, high-speed videography of the mouse body allowed monitoring the forelimb and hindlimb movements.

We found that a majority of the mice learned to increase the frequency of the Conditioned Waves as they learned to collect more water rewards. As they generated cortical waves, the mice deployed a strategy that combined limb movements — which generated excitatory inputs to the primary somatosensory cortex — with strong pre-movement cortical state modulations that shaped them into cortical waves that fulfilled our conditioned criteria for reward. This shows that mice can exert a degree of control on the properties of individual waves as they address behavioral challenges, and suggests that individual waves are a functionally relevant player of cortical information processing.

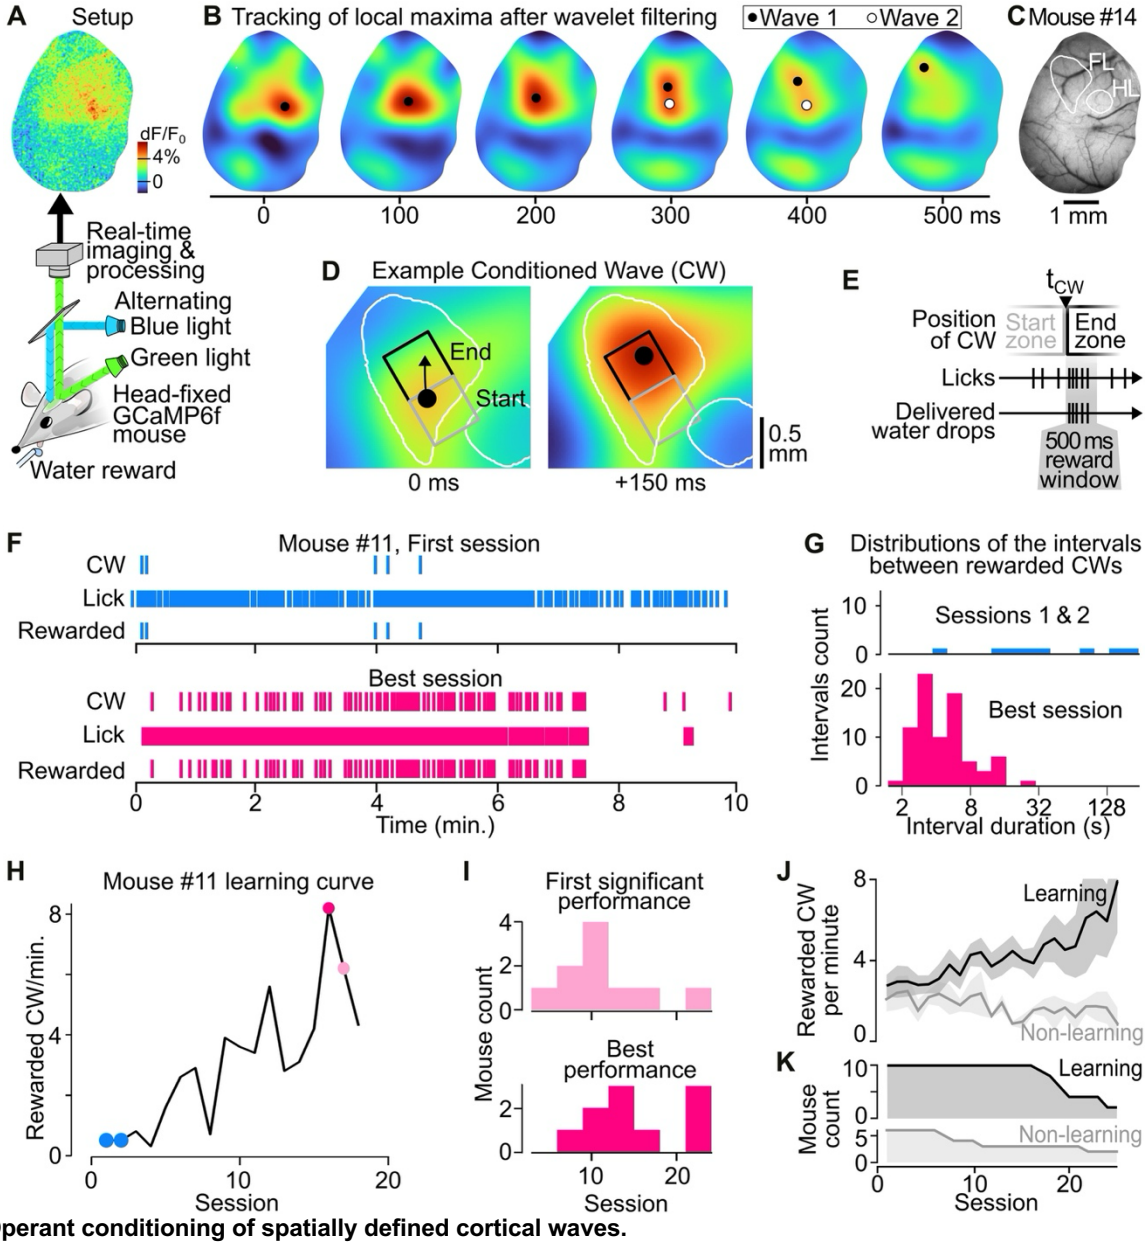

Figure 1. Operant conditioning of spatially defined cortical waves.

(**A**) A water-restricted GCaMP6f-expressing mouse with a chronic window above the somatosensory and motor cortices was head-fixed under a macroscope. Alternating imaging with blue/green illumination aimed to correct for intrinsic hemodynamic related signals in functional imaging. Water rewards were conditionally delivered from a spout upon licking.

(**B**) Example GCaMP6f fluorescence signals. Images were acquired at 100 Hz and processed in real-time. After hemodynamics correction (example in **A**), they were filtered with 2D wavelets (100 µm to 1 mm bandpass). The position of several local maxima corresponding to individual Waves could be simultaneously tracked online.

(**C**) Blood vessels at the surface of the cortex shown by green illumination in the same mouse as in **A**, **B**. Hindlimb (HL) and Forelimb (FL) touch representations were functionally mapped by imaging sensory-evoked signals under anesthesia.

(**D**) In the same mouse, example of a cortical Conditioned Wave (CW): as it travels across the cortical surface, it enters the Start and then the End zone. Black dot: position of the detected local maximum of the wave. Gray and Black squares: Start and End zones defined by the experimenter. White contour: FL and HL touch representations (see **C**).

(**E**) The time of a CW exiting the Start zone and entering the End zone is t_{CW}. t_{CW} marks the start of a 500 ms reward window, during which each lick on the spout triggers the release of a water droplet, collected by the tongue in the same lick movement.

(**F**) Example time sequence of CWs, Licks, and Rewarded CWs in mouse #11 during its first session (top, blue) and the session with the highest Rewarded CW frequency (best session, bottom, magenta).

(**G**) Corresponding histogram of the time intervals between consecutive Rewarded CWs during the first two sessions (top) versus the best session (bottom), shown on a logarithmic scale.

(**H**) Frequency of Rewarded CWs across the training sessions of the same mouse #11. Color code: outcome of a training-wide Kruskal-Wallis test followed by multiple comparison tests. Light magenta: Session where the intervals between consecutive Rewarded CWs were significantly shorter than during sessions 1 & 2. Saturated magenta: best session.

(**I**) Population distribution of the number of training sessions until the first session with a significant increase in performance (top) and until the best session (bottom).

(**J**) Average Rewarded CWs frequency across mice with a significant increase in performance (Learning, black line), versus no significant increase (Non-learning, gray line), Light background: SEM (n values shown in panel **L**).

(**K**) Evolution of the mouse count as training progressed.

RESULTS

Mice learned a cortical wave-dependent task

We have developed a novel widefield optical brain–machine interface to test whether mice can learn to generate individual mesoscopic waves with a specific displacement at the surface of the somatosensory-motor cortex in order to get rewards. Sixteen mice expressing GCaMP6f in excitatory neurons of the neocortex were implanted with a head fixation bar and a chronic optical window over the somatosensory and motor cortices (see Materials and Methods). After getting water-restricted mice used to remain head-fixed on a setup and freely obtain water by licking from a spout, we recorded their spontaneous cortical activity through widefield calcium imaging at 100 Hz (Figure 1A). To detect events of synchronous cortical activity (100 µm-1 mm spatial range), images were filtered with a 2D wavelet transform and the local maxima (x/y coordinates) of these events were tracked in real time (Figure 1B) at the surface of the somatosensory-motor cortex. Consecutive local maxima distant by less than 200 µm defined the trajectory of was we refer to here as a Wave (Figure 1C).

We selected two nearby 500 × 500 µm square zones (Start and End) within the optical window, and defined the Waves that were detected first in the Start, and then in the End zone, as Conditioned Waves (CW, Figure 1D). The time of first entrance in the End zone t_{CW} initiated a 500 ms opportunity window, during which water droplets were instantly delivered upon licking (Figure 1E). In order to facilitate the initial phase of the learning, we selected the location of the Start and End zones so that 1 to 5 Conditioned Waves/min occurred during spontaneous brain activity. Because of the diversity of local wave statistics observed across mice, the location of the Start and End zones varied within the somatosensory cortex, in the vicinity of the limb areas, that we identified by recording GCaMP6f signals evoked by passive limb stimulations under anesthesia (Figure 1C,D, Material and Methods).

We trained the mice daily, during sessions lasting 20 to 30 minutes. We quantified performance — defined as the frequency of Rewarded CWs — and carried out all subsequent analysis on the first 10 minutes, when the highest motivation was apparent (Matteucci et al. 2022). To assess the evolution of mice performance with training (example in Figure 1F), we computed for each session the distribution of the time intervals between Rewarded CWs, and compared it to sessions 1 and 2 combined (Figure 1G,H, Methods). Of the 16 trained mice, 10 increased significantly their performance compared to the first 2 sessions ("Learning" group, example outcome from Kruskal-Wallis test followed by multiple comparison tests shown in Figure 1H. Their first significant performance was achieved after 11.1 ± 5.4 (mean ± SD) training sessions (Figure 1I, top). Several of these mice kept progressing with additional training, and achieved their best performance after 15.0 ± 5.6 training sessions (Figure 1I bottom, Figure 1J). It should be noted that although not all the mice underwent the same number of training sessions, all Learning mice were trained at least during 18 sessions (Figure 1K).

A selective increase in frequency of Conditioned Waves

The increase in performance observed in the learning group may be based on evolution in cortical dynamics and/or licking behavior. Quantification of detected waves revealed a significant increase in the frequency of CW (Figure 2A) in the context of an overall increase in Wave frequency within the imaged window (Figure 2B). We asked if these learning-driven changes were caused by a homogeneous increase in Wave frequency across the cortical surface, or if instead the mice managed to generate a spatially focused increase of CW frequency. To probe this, we explored the spatial structure of the Wave trajectories, and more specifically the spatial density of their individual time points. With learning, the density increased mainly around the Start/End zones (Figure 2C-G).

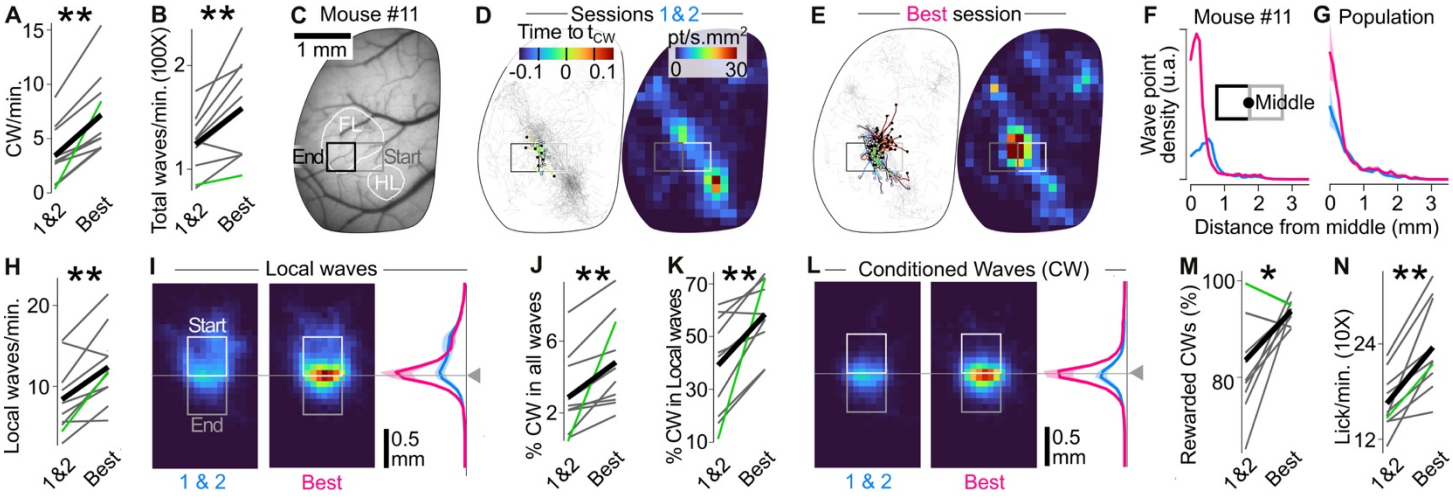

Figure 2. Spatial shaping of the cortical waves in Learning mice

(**A**) Frequency of Conditioned Waves (CW) during the first two sessions and during the session with the best performance. **: Wilcoxon p = 0.002. Gray: individual mice. Green: mouse #11. Black: population average.

(**B**) Same as **A** for the frequency of all waves detected on the imaged cortical surface. **: Wilcoxon p = 0.0098.

(**C**) Case study: Location of the start and end zones for example mouse #11. White delineations: primary FL and HL areas of the cortex identified by passive stimulations (see methods).

(**D**) Left: Trajectories of all waves detected during sessions 1 and 2 in mouse #11 (gray), overlaid with trajectories of CWs (colored). The changing color along CW trajectories indicates time relative to t_{CW}. White dot: first point of CW. Black dot: last. Right: point density of all Wave trajectories.

(**E**) Same as **D** for the best session.

(**F**) Radial distribution of Wave point density, computed for all Waves of mouse #11. First and second (blue) versus best session (magenta), centered on the middle point of the task (see insert schematic).

(**G**) Same as F at the population level in the Learning group. Light background: SEM.

(**H**) Evolution of the frequency of Local waves: all Waves that entered the Start zone, including CWs. **: Wilcoxon p = 0.0059.

(**I**) Population average of the density of the Local wave points at the surface of the cortex, around on the Start and End zones. Left: Sessions 1&2. Middle: Best session. Right: Marginal distribution of Local wave point density along the main axis of the Start/End zones (task main axis).

(**J**) Evolution of the proportion of CW among all recorded Waves. **: Wilcoxon p = 0.002.

(**K**) Same as **A** for the percentage of CW among Local waves. The proportion increases significantly over training. **: Wilcoxon p = 0.0039.

(**L**) Same as **I** for CW points.

(**M**) Evolution of the proportion of Conditioned Waves that were actually rewarded. *: Wilcoxon p = 0.027.

(**N**) Evolution of the frequency of licking. **: Wilcoxon p = 0.002.

The occurrence of Local Waves — defined as all waves that entered the Start zone, even if not Conditioned — increased significantly (Figure 2H). Following learning, we found that they spent most of their lifetime at the interface between the Start and End zones (Figure 2I). When we then focused on CWs, we found that their proportion increased both with respect to all Waves (Figure 2J) and with respect to the Local Waves (Figure 2K). Consistent with the general trend observed in Local Waves, learning led CWs to become spatially focused on the interface between the Start and End zones (Figure 2L).

Among CWs, a large and significantly increasing proportion were rewarded (Figure 2M). This came together with a strong increase in the frequency of licking (Figure 2N), suggesting that the mice were not limited in their ability to lick, and learned to increasingly systematically lick in order to capture any reward opportunity. This observation led us to focus the rest of our analysis on the timing of Conditioned Waves t_{CW}, which is directly related to the cortical dynamics, rather than on the timing of Rewarded CWs, as they are largely equivalent, and the later may be impacted by the potential stochasticity of the lick trains.

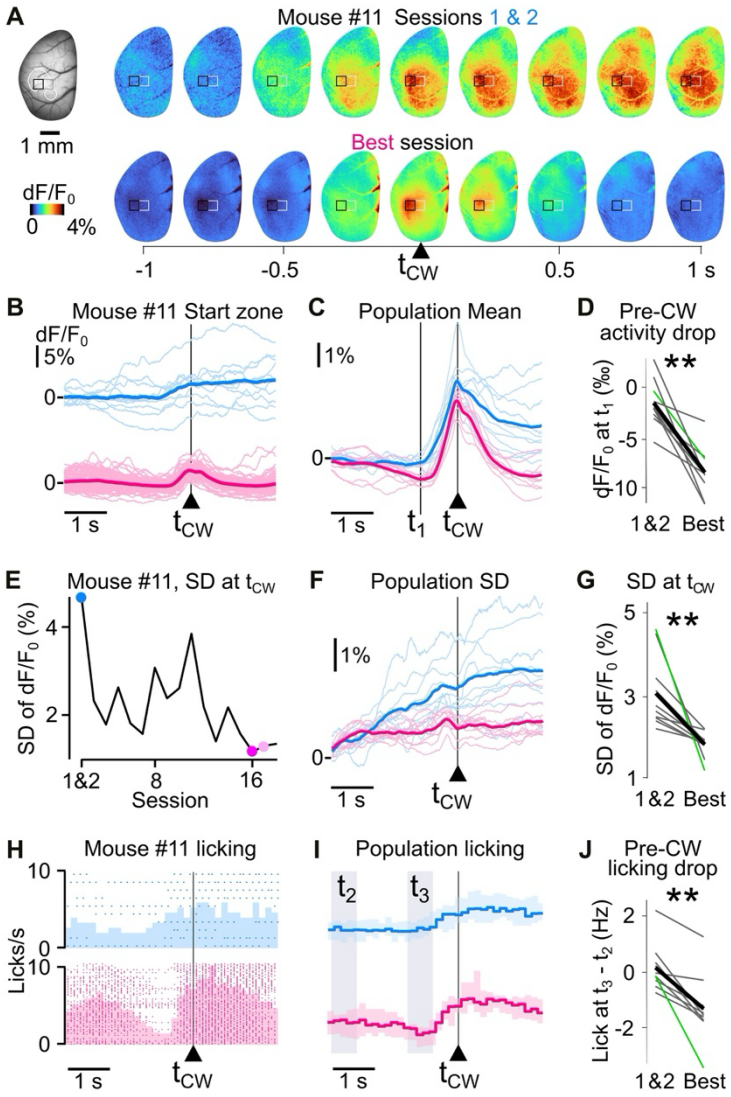

Figure 3. Spatiotemporal shaping of cortical dynamics by the operant conditioning of Wave trajectories

(A) Average GCaMP6f signal across the cortical window in mouse #11, aligned on t_{CW}, during sessions 1&2 (top) and the best performance session (bottom).

(**B**) Averaged calcium fluorescence signal in the 'Start' area aligned on t_{CW} from the same mouse as in **A**. Thin lines: individual CWs. Thick line: average across all CWs of the session. Top, blue: sessions 1&2. Bottom, magenta: best session.

(**C**) Averaged calcium fluorescence signal around t_{CW}. Thin lines: individual mice. Thick line: population average. Blue: sessions 1&2. Magenta: best session.

(**D**) Population analysis of the pre-CW suppression (quantified at the t_1 mark indicated in **C**). **: Wilcoxon p = 0.002. Gray: individual mice. Green: mouse #11. Black: population average.

(**E**) Evolution of the standard deviation of the 'Start' area calcium fluorescence signal across individual CWs, quantified at t_{CW}, across sessions for mouse #11. Blue dot: sessions 1 & 2. Magenta: sessions with significantly higher Rewarded CW frequency as in sessions 1 & 2. Dark magenta: best session.

(**F**) Population analysis of the standard deviation across individual CWs, aligned on t_{CW}. Thin lines: individual mice. Thick line: population average.

(**G**) Population analysis of the evolution of the pre-CW standard deviation, quantified at t1. **: Wilcoxon p = 0.002.

(**H**) PSTH around t_{CW} of the licking in the first two (top) versus best session (bottom) in mouse #11. The raster of licking times is shown on top of the PSTH.

(**I**) Average PSTH of the instantaneous licking frequency during the first and second (blue) versus the best session (magenta), aligned on t_{CW}. Light background: SEM.

(**J**) Change of lick rate quantified between t_2 and t_3 as defined in **I**. **: Wilcoxon p = 0.002.

Spatiotemporal shaping of cortical activity and licking around CWs

So far, we focused our analysis on the evolution of Waves trajectories that were specifically reinforced by operant conditioning. We next asked if the underlying mesoscale cortical dynamics was also shaped during this learning process. To this end, we analyzed the averaged, hemodynamics-corrected, non-filtered widefield calcium signals acquired around each t_{CW}. During the first two sessions (combined), the calcium activity around t_{CW} exhibited a broad spatial spread (example in Figure 3A,B top). In contrast, after training, the average calcium activity around CW time was more focused in space around the Start/End zones, and revealed a strong suppression before and after the CW (example in Figure 3A,B bottom). This drop of activity just before CW was significant at the population level (Figure 3C,D). In addition, we found that the inter-trial variability in calcium signal dynamics decreased with training (example in Figure 3E, population analysis in Figure 3F,G). This reduced variability suggests that mesoscale cortical dynamics — and in particular the pre-CW suppression — were operant-conditioned and were therefore required for the initiation of the cortical waves trajectory that were rewarded in our task.

Consistent with the large magnitude of this cortical suppression and its potential impact on cortically controlled behaviors, we found that licking was significantly down-regulated at the same time as mesoscale cortical activity (Figure 3H-J) after learning, which combined with the overall increased in licking that we described earlier (Figure 2N, also visible in Figure 3I).

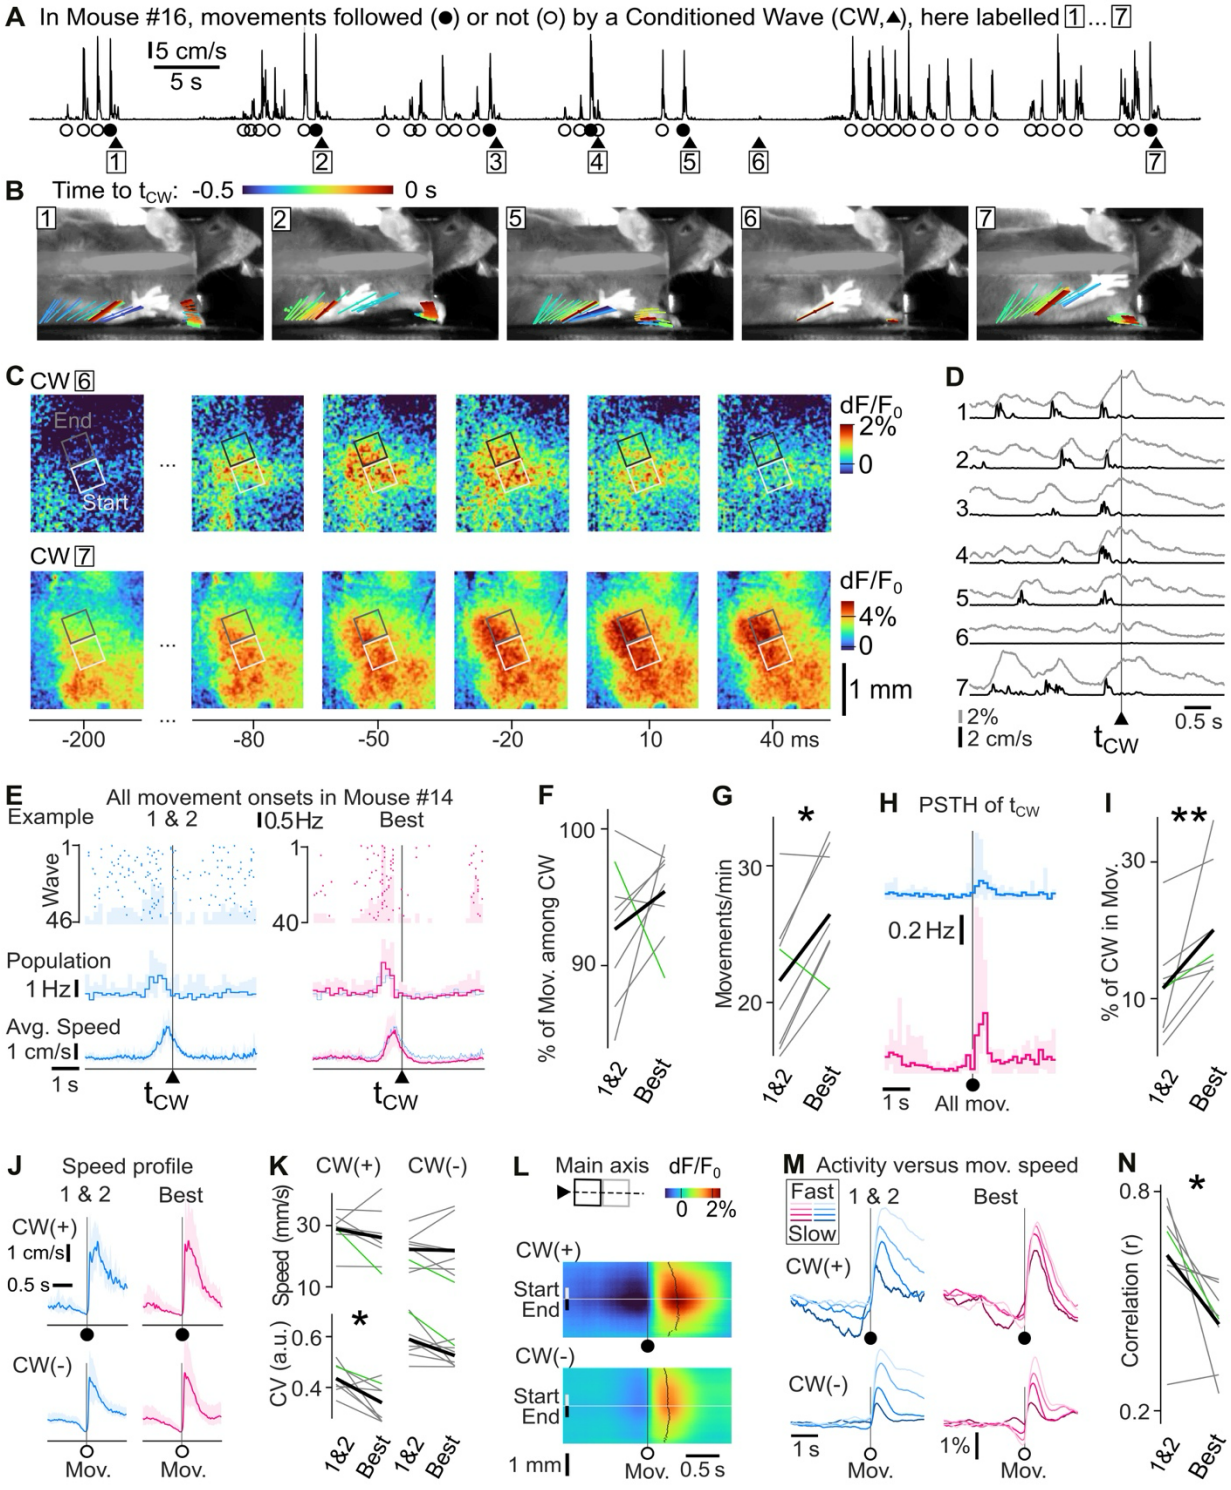

Figure 4. The interaction of forelimb movement inputs and cortical state-like dynamics underpins Conditioned Waves.
(**A**) Limbs speed (average of forelimb and hindlimb) of mouse #16, during 90 s of the best performance session. Triangle: CWs 1 to 7. Filled circles: movement associated with a CW (t_{CW} is between movement onset and offset +100 ms). Open circles: movement without a CW.
(**B**) Side view of mouse #16 behavior, overlaid with the movements of the right hindlimb and forelimb palms in a 0.5 s window prior to t_{CW} in 5 of the CWs shown in **A**.
(**C**) Snapshots of the calcium signal imaged around t_{CW} for CWs 2 and 6 (see **A** and **B**). White: Start zone. Black: End zone.
(**D**) Superimposed limb movements (black) and average calcium dynamics quantified over the Start zone (gray) around the 7 CWs shown in **A**.
(**E**) Distribution of the limb movements around t_{CW}, in the first two sessions (left, blue) and the best session (right, magenta). Top: scatter plot of detected limb movement onsets. Overlapping PSTH: same data. Middle: population median of the limb movement onsets PSTHs (n = 8). Light background: 5% and 95% percentiles. Bottom: Average limb speed. Light background: standard deviation.
(**F**) Evolution of the proportion of CWs that were preceded by a detected movement onset between the first two sessions, and the best session. Green line: mouse #16, thick line: population average.
(**G**) Evolution of the frequency of detected forelimb movements *: Wilcoxon p = 0.039.
(**H**) Population median of the time distribution of the t_{CW} around movement onsets. Light background: 5% and 95% percentiles (n = 8).

(**I**) Proportion of the movements that are immediately followed by a CW (latency ≤100 ms after movement offset). **: Wilcoxon p = 0.008.
(**J**) Population average speed profile of the detected movements in the first two (left, blue) and best session (right, magenta). n = 8. Top: movements associated with a CW (CW(+), filled circle: movement-onset Bottom: other movements (CW(-): open circle). Light background: standard deviation.
(**K**) Average and coefficient of variation (CV) of the peak speed of the CW(+) and CW(-) movements across learning (sessions 1 & 2 versus best session). The coefficient of variation decreased significantly for Conditioned Waves. *: Wilcoxon p = 0.023 (n = 8).
(**L**) Population average cortical dynamics around movement onsets across the task main axis, for CW(+) (top) and CW(-) movements (bottom), in the best session (n = 8). Black curve: time of the peak of the activity along the task main axis.
(**M**) Evolution of the population average calcium dynamics in the Start area, aligned on movement onsets, split by movement speed. Left: first two sessions. Right: best session.
(**N**) Population evolution with learning of the Pearson correlation coefficient between the peak activity in the Start zone (t_{CW} - 0.1 to t_{CW} + 0.5 s) and the average limb movement speed between t_{CW} - 1.5 s and t_{CW} + 0.5 s. *: Wilcoxon p = 0.016.

Peripheral inputs are integrated by cortical circuits to generate Conditioned Waves

Next, we asked what was the origin of the coordinated cortical activation that is generated by the mice as they solve the task. In the limb areas of the primary somatosensory areas where we positioned the Start and End areas, the major source of peripheral somatosensory inputs should be the limbs proprioceptive and touch inputs. Therefore, we hypothesized that during the task, the mice may perform limb movements to trigger cortical activity towards solving the task. To probe this hypothesis, we tracked the right forelimb and hindlimb position through all sessions (for 8 out of 10 'Learning mice', see Methods). We averaged the speed of the two limbs into an overall "limb speed" variable (example in Figure 4A) that we used to detect individual limb movements based on the thresholding of the speed (speed > 15 mm/s, see Methods). We often found — even during the best performance session of trained animals — repeated limb movements that persisted until one was immediately followed by a CW (movement examples in Figure 4B, associated cortical waves in Figure 4C, time courses in Figure 4D). Overall, although some CWs were not associated to any movement (wave 6 in Figure 4A-D), limb movements featured consistently before CWs and movement onsets were increasingly time locked with CWs (Figure 4E): At the population level, 92.4% of the CWs during the first two sessions were immediately preceded by limb movements, a proportion that (non-significantly) increased with learning (Figure 4F).

However, even if the frequency of the limb movements irrespective of Conditioned Waves increased significantly (Figure 4G) and CWs were increasingly aligned on movements (Figure 4H,I, see also Figure 4E), the proportion of all movements that were followed by a CW remained low (20%, Figure 4I). These frequent failure of limb movements to generate CW may be explained by inadequate or inconsistent movements. The comparison of the profile of the limb movements that led to CW (labelled CW(+) movements) versus non-Conditioned brain activity (CW(-)) did not reveal any clear distinctive feature. The average movement speed time course remained similar (Figure 4J), while we observed a specific and significant reduction of the variability of the CW(+) movements (Figure 4K). Overall, this suggests that the mice did stabilize and reinforce a subset of CW-leading movements, but failed to optimize it to improve decisively their efficiency in triggering CWs, with still on average ~80% of CW(-) movements at this end of the training effort.

We therefore hypothesize that instead of shaping their limb movements to generate peripheral inputs that would materialize in the cortex as a Conditioned Wave, the mice combined the rough mass cortical activation triggered by peripheral movements with an active sculpting of these inputs by cortical state dynamics.

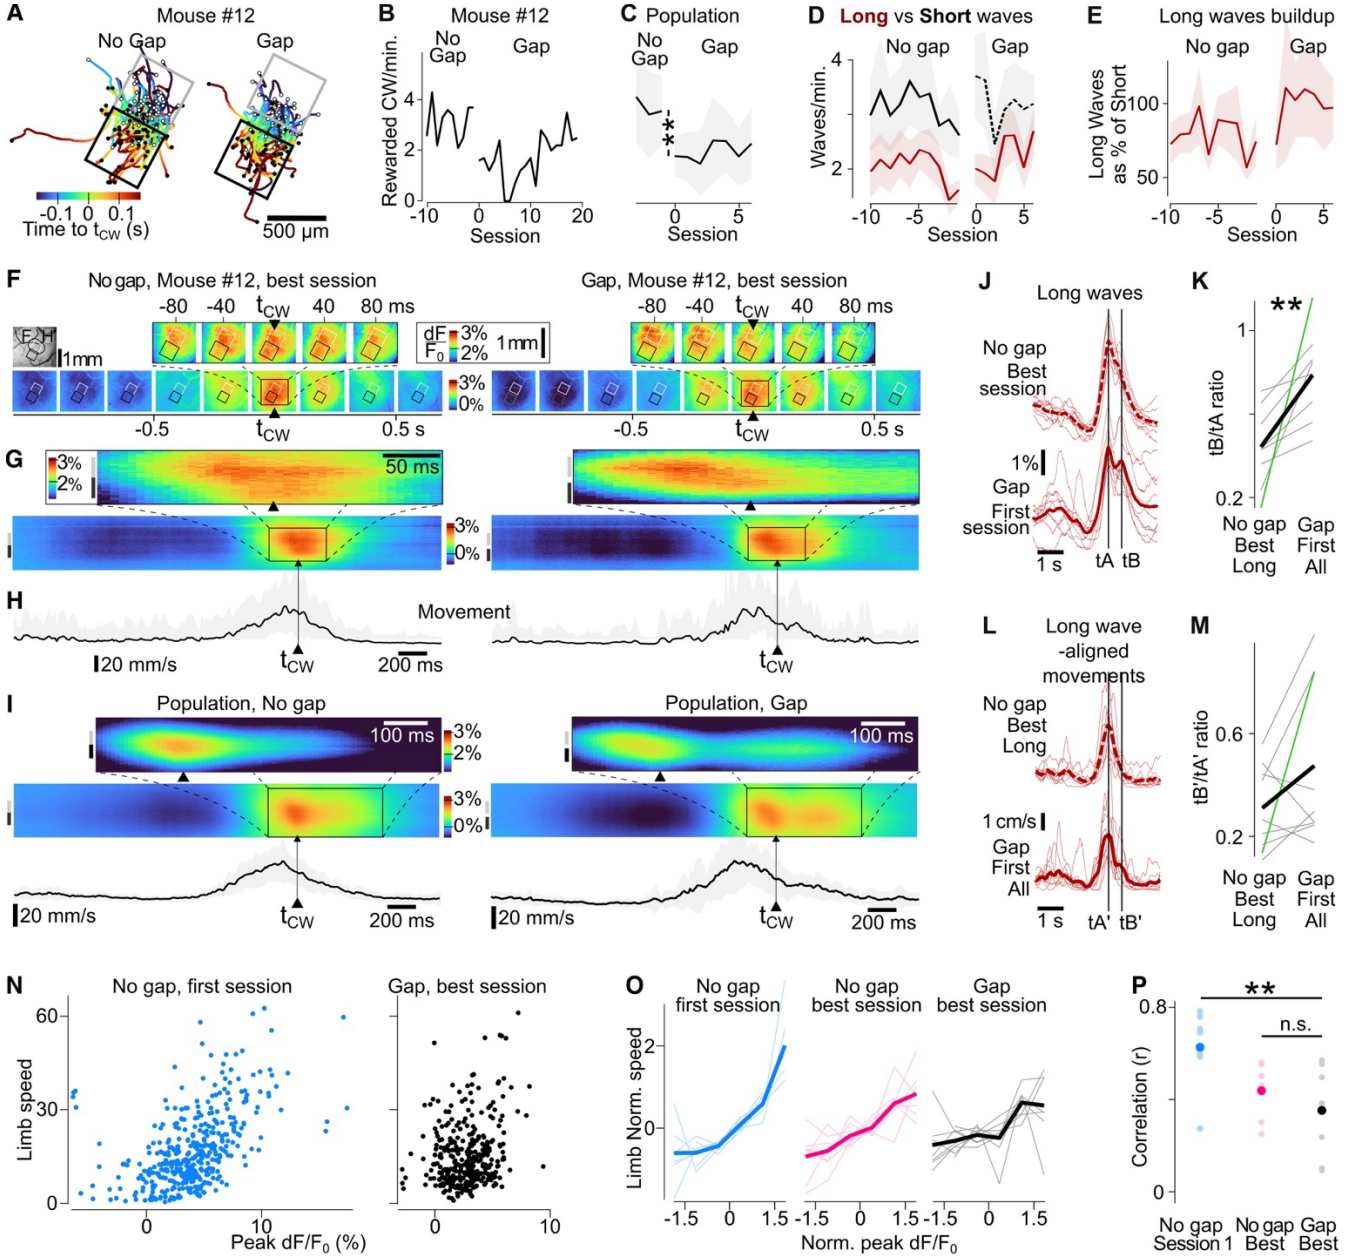

Figure 5. The operant conditioning of larger cortical wave travel triggers additional reshaping of cortical dynamics.

(**A**) Example trajectories of the detected Waves in the No Gap (left) versus the Gap condition (right) in (mouse #12, best session). Color code: time around t_{CW}.

(**B**) Example performance (Rewarded CW frequency) of mouse #12 around the transition from the No-Gap to the Gap condition.

(**C**) Same as **B** at the population level. Light Background: SEM. *: Wilcoxon p = 0.0080 (n= 9 mice).

(**D**) Evolution of the frequency of waves that would fulfil the criteria of the Gap condition (Long Waves, brown), and of waves that would fulfil the No Gap condition, but not the Gap condition (Short Waves, black). Continuous lines: fulfil the ongoing task condition. Light background: SEM.

(**E**) Ratio between the count of Short and Long Waves. Light background: SEM.

(**F**) Example in mouse #12 of the average of t_{CW}-aligned cortical dynamics during the initial, No Gap condition (left, best session in this condition) versus the Gap condition (right, best session in the condition).

(**G**) Time/Space view of the same data as in **F**. Data was projected on the task main axis.

(**H**) Average limb speed around t_{CW}, No-Gap (left) versus Gap (right) condition, in the same mouse example. Light background: SEM.

(**I**) Population analysis matching **G**-**H**.

(**J**) Population average time profile of the Calcium signal in the Start zone during Long Waves that match the Gap condition, in the best No Gap session (Top), versus the first Gap session (bottom). The waves are time-aligned on their maximum.

(**K**) Ratio extracted from the waves measured in **J** between the time of the wave peak (tA) and the time of the rebound observed in the Gap condition (tB). Green line: mouse #12. **: Wilcoxon p = 0.0040.

(**L**) Same as **J** for limb movements, also aligned on the peak of the cortical activity in the Start zone. tA' and tB' are shifted compared to tA and tB so that tA' is at the peak of the limb movements average.

(**M**) Same as **K** for limb movements. No significant difference.

(**N**) Peak value of the cortical activity versus average limb speed around t_{CW}. The data points of all mice are merged. Left (blue): No gap condition, first session. Right (black): Gap condition, best session.

(**O**) Average Limb speed observed before t_{CW} as a function of the average cortical activation in the Start zone at t_{CW}. Light colors: individual mice. Tick line: population averages.

(**P**) Pearson correlation between the CW-related cortical activity and movements, as reported in **N**,**O**. **: Mann-Whitney p = 0.002.

To probe this, we compared the cortical dynamics triggered by CW(+) versus CW(-) movements. Despite comparable movement dynamics in these two conditions, we found very different cortical dynamics. In particular, in trained mice, before CW(+) movements, we found a strong suppression of activity that was absent before CW(-) movements (Figure 4L).

To understand if this difference in cortical state before movement may have had an impact on the cortical dynamics triggered by the limb movements, we sorted them by limb speed. In the first training sessions, there was a direct link between limb speed and calcium wave amplitude, for both CW(+) and CW(-) movements (Figure 4M, left). However, in the best session, after training, this relation appeared largely abolished in CW(+) movements, but not in CW(-) movements (Figure 4M, right). Consistently, computing the linear regression between the speed of individual movements and amplitude of the corresponding calcium activation confirmed a correlation that was significantly (WIlcoxon $p = 0.016$) reduced with learning (Figure 4N).

Overall, these findings show that with learning, the production of Conditioned Waves remained tied to peripheral movements. However, we found that these peripheral inputs were shaped by pre-movement cortical state-like dynamics, making them independent from key limb movement characteristics such as speed.

Spatial separation of Start and End zones stretches cortical dynamics.

After training on the initial protocol during 21 sessions on average, we aimed to probe the ability of the mice to generate longer cortical waves. To this end, we transitioned the mice to a Gap condition, where the Start and End zones were separated by 130 µm. To solve this new task, the mice had to produce a Long Wave that crossed this entire space gap in order to obtain a reward (Figure 5A). We maintained the rest of the protocol unchanged. Nine mice were trained with this Gap protocol for an additional 13 sessions on average.

At the time of the transition from the No gap to the Gap protocol, the frequency of Rewarded CWs dropped significantly (Figure 5B,C). This was expected since through the No gap training, the frequency of Long Waves matching the requirement of the Gap protocol was lower than that of Short Waves whose trajectory only fulfill the No gap task requirements (Figure 5D). This lower frequency of Conditioned Waves persisted throughout the training in the Gap protocol. Still, we found that the ratio of Long Waves over Short Waves, which was stable in the No gap condition, increased, although not significantly, with training in the Gap condition (Figure 5E). Overall, this Gap protocol was more challenging than the No gap condition for the mice and they struggled to recover a large Rewarded CW frequency (Figure 5C), although some mice were successful at increasing performance in the Gap condition (Figure 5B).

The mice performed this new task by building on their ability to solve the initial No gap protocol, and further reshaped their cortical dynamics. In particular, they maintained the cortical hyperpolarization before wave onset (Figure 5F-I, No gap: left, versus Gap: right column), and generated longer duration CWs when compared to the No gap condition (Figure 5M,N). These extended CWs appeared to split into a first main wave followed by a significant rebound of activity (Figure 5O-P). Except in 3 mice, this second phase of the Gap CWs was not associated with a matching evolution of limb movements (Figure 5Q,R). To further probe this dissociation of the CW dynamics from the associated movements, we repeated our study of the link between limb movement speed and the amplitude of the cortical dynamics, which we found was already diminished by learning in the No gap condition (Figure 4L-N). A scatter plot of all mean limb speed around t_{CW} versus the peak calcium dynamics across the whole population showed that in contrast to the clear positive relationship observed in non-trained mice (Figure 5N,O, Pearson $r = 0.63$), the linear relationship was lost after the training, first in the No gap and then in the Gap protocol (Figure 5N,O, $r = 0.35$ in the best Gap session). Overall, this significant (Figure 5P) loss of the link between movement and

calcium activity in the primary forelimb cortex was even more visible in the No gap than in the Gap condition.

DISCUSSION

Using an optical brain-machine interface based on the real-time processing of wide field calcium signals, we have shown that mice can generate specific cortical waves of neuronal activity at the mesoscopic scale in order to solve an operant conditioning task. These waves were typically preceded by a limb movement that triggered a burst of cortical activity. However only movements occurring in conjunction with a preparatory suppressive cortical state drove Conditioned Waves. These findings suggest that mice can shape individual cortical waves of activity through unexpected mesoscale cortical processing.

A learning-based mesoscale brain-machine interface

Invasive brain-machine interfacing provides a unique opportunity to probe the information processing architecture of the brain, by extending the operant conditioning paradigm to internal variables of the brain. However, so far, this interfacing strategy based on learning has only been successful for small groups of neurons (Arduin et al. 2013; 2014; Clancy et al. 2014; Fetz 1969; Koralek et al. 2012; Neely et al. 2018; Prsa et al. 2017) or small (~0.1 mm^2) patches of cortical tissue trained through the real-time processing of local 1-photon calcium imaging (Clancy and Mrsic-Flogel 2021) or LFP recording (Engelhard et al. 2013; Shi et al. 2024; Chauvière and Singer 2019). In these experiments, when a group of neurons was conditioned to solve a task, a single neuron often took the lead in controlling the brain-machine interface (Abbasi et al. 2023). Here, in contrast, we conditioned the delivery of a reward to the coordinated activity of a large population of cortical neurons shaping into a traveling peak of mesoscopic activity. Therefore, our experiments open up the possibility to operant-condition the spatio-temporal features of neuronal activity at brain-wide, at a scale that is relevant to electrocorticogram measurements (ECoG), in addition to the decoding strategies that have dominated human application of ECoG, both towards the decoding of speech (Anumanchipalli et al. 2019), and body control restoration (Lorach et al. 2023). Our experiments also revealed the buildup of large-scale changes in activity across several millimeters of the imaged cortical area, both before and after the CW (Figure 3,4 and 5). We hypothesize that these unrewarded response of the cortex to our training are the sign of fundamental constraints on the plasticity of cortical mesoscale dynamics that will need to be further explored towards efficient mesoscale brain-machine interfacing.

A simple readout of complex mesoscale cortical dynamics

In the context of our brain-machine interface experiments, we aimed to train mice to control parameters of individual cortical mesoscale waves. These waves show complex and variable shapes that can be hard to capture through a compact set of descriptors. In fact, the identification of such descriptors has been a major effort in the field of optical imaging. One strategy has been to identify principal components of the spatial dimensions of the imaged cortical dynamics, and to project back the spontaneous spatiotemporal dynamics into the identified subspace (Mohajerani et al. 2013; Musall et al. 2019; Surinach et al. 2023). Another direction has been to compute the flow fields of wave speeds at the surface of the cortex, and focus on the singularities that structure these fields, and in particular sources and sinks (Liang et al. 2023; Mohajerani et al. 2013).

Given the challenge of operant conditioning, which requires a simple variable that could be computed in real-time and manipulated by the mouse, we selected a first order descriptor: the wave position. Several definitions of wave position are possible, including the position of the local maximum within wave territory, as well as the position of the wavefront. However, the wavelet filtering of the raw signal that is required to extract the dynamics of individual waves in real-time (Figure 1) precluded an accurate measurement of the wavefront. Therefore, we chose to track the

local maxima of the wave, as this measurement is immune to such filtering artefacts. This means that the constraints that we applied to the cortex were local to the ~500 µm area at the Start/End zones border, while the mice deployed mesoscale dynamics at a millimetric scale in order to solve the task (Figures 3-5). Future work on the design of a more spatially defined constraints may reveal additional modes of mesoscale operant conditioning.

Mesoscale cortical activations rely on peripheral movements

As the mice learned the task, Conditioned Waves were mainly generated in conjunction with limb movements (Figure 4F). We hypothesize that touch and proprioception generated from these movements, as well as motor commands in the nearby limb motor cortex, all contributed to the activation of the limb somatosensory cortex, which is the location where we have positioned our operant-conditioning windows. It is therefore likely that the mechanism used by the mice to generate the mass activation of cortical tissue required by the task relied on peripheral inputs, taking advantage of the cortical mesoscale activation that can be triggered by sensory inputs (Jancke et al. 2004; Ferezou et al. 2007; Afrashteh et al. 2021). This does not match the lack of movement reported in several brain-machine interfacing tasks that relied on the conditioning of few cortical neurons (Clancy et al. 2014; Clancy and Mrsic-Flogel 2021; Prsa et al. 2017). We hypothesize that this is because the bulk cortical activation that is required in our task can be hard to generate for the cortical circuitry in the absence of peripheral inputs, in contrast to the activation of a small subset of neurons, that may be achieved locally by the cortico-striatal circuitry (Koralek et al. 2012) without the need for peripheral inputs.

Interplay of peripheral inputs with cortical state.

In our experiment, only a small minority (20%) of the movement-triggered bulk cortical activation resulted in a reward (Figure 4I), and the movements that led to either successful or unsuccessful generation of Conditioned Waves were similar (Figure 4J,K). In contrast, the cortical dynamics were vastly different. In particular, we observed a pre-wave suppression of cortical activity that was specific to the Conditioned Waves, and became stronger with learning (Figure 4L,M, Figure 5F-I). Furthermore, while movement amplitude was correlated with the magnitude of associated cortical activation in naïve mice, we observed a specific suppression of this relationship for the movements preceding a Conditioned Wave throughout learning. (Figure 4M). These findings suggest that Conditioned Waves were generated by the conjunction of self-generated somatosensory inputs with internal cortical dynamics that were reinforced with operant conditioning. This internal processing was likely hard to pull out for the mice, as success rates remained low (despite learning, Figure 4I), and depended on the features of the cortical wave that had to be generated, as shown by the drop of performance in the Gap condition (Figure 5B). The rebound in activity observed after Conditioned Waves in the Gap task configuration (Figure 5J-M) suggests that orchestrating the neuronal dynamics required by the new rule demands an even more complex form of adaptation that calls for further investigation.

Relationship of Conditioned Waves with spontaneous waves

Several of our findings suggest that to achieve the behavioral task, the mice were able to control the timing and adjust parameters of existing wave patterns, but that they were not able to generate new wave patterns from scratch. In particular, we have observed that learning was facilitated when the selected Start and Finish zones were located within a cortical region that spontaneously features a high wave density (data not shown), in our case in the primary somatosensory and motor areas (Figure 1C). In addition, the Conditioned Waves retained the local orientation observed in the first sessions through the training sessions, even when this orientation was not optimal to achieve the task (Figure 2D,E). Finally, we have found that it was challenging to extend the distance traveled by the Conditioned Waves beyond their naturally occurring size (Figure 5B-E)

Therefore, Conditioned Waves in our experiments are likely to rely on a similar anatomical substrate as spontaneous cortical waves. Still, these waves did not match fully their spontaneous counterpart at they were preceded by a large suppression after training (Figure 3,4).

Suppression of the activity around Conditioned Waves

One of the most striking differences between Conditioned and Non-Conditioned Waves in our data is the strongly suppressed cortical state that precedes and follows Conditioned Waves. This suppression of ongoing activity is reminiscent of the quiet cortical state that has been well described in the somatosensory cortex (Petersen et al. 2003; Crochet and Petersen 2006; Ferezou et al. 2007). A transition to a quiet state in preparation for the generation of a cortical wave triggered by peripheral inputs would optimize its magnitude, as sensory inputs trigger stronger cortical activations in this context than during active cortical states (Fanselow and Nicolelis 1999; Petersen et al. 2003; Crochet and Petersen 2006; Ferezou et al. 2007). However, in several ways, the dynamics we observed around Conditioned Waves do not match the established features of the typically reported "quiet" state. In particular, we observed that after learning, the Conditioned Waves were shortened and followed by a second phase of hyperpolarization, while sensory inputs in the context of a quiet state should have transitioned the cortex to a long-lasting active state, rather than rapidly reverse to a quiet state. In addition, we found that after learning, the magnitude of the Rewarded Waves was not anymore modulated by the amplitude of the movement, in contrast to the established observations of cortical integration of inputs with cortical state (Petersen et al. 2003; Crochet 2011). Finally, we found that this transitory suppression of activity was time-locked to CW-related movement onsets (Figure 4L,M) and was accompanied with a structured modulation of the licking (Figure 3H-J).

This suggests that the pre-Conditioned Wave suppression is part of a voluntary action planning for the generation of a Conditioned Wave, rather than a transition to quiet state, which is generally associated with non-attentive, non-task engaged behavior. In particular, it could be related to an attentional mechanism, similar to the attention-driven spatial focus in the primary visual area, where non-task related activity is dampened, in contrast to task-relevant inputs (Brefczynski and DeYoe 1999; Dugué et al. 2020; Herrmann et al. 2010; Murray 2008). In visual tasks, such attentional effects have been linked to higher order cortical areas such as the frontal eye field, as well as to the pulvinar and thalamic nuclei.

Overall, we conclude from these observations that the mice generated a tailor-adjusted cortical processing state in order to generate Conditioned Waves from peripheral inputs. This remained challenging even after repeated training, as it was only successful in 20% of the limb movements (Figure 4I). Previous work on BMIs based on the readout of individual neurons have shown that the cortico-striatal loop is required to enable the learning of a new pattern of single neuron activity in order to solve a BMI task (Koralek et al. 2012; Neely et al. 2018). Further experimental work will be required to determine the actual subcortical circuitry and mechanisms that support the generation of Conditioned Waves in our experiments, including the suppression of cortical activity prior to Conditioned Waves.

MATERIALS AND METHODS

Animals

Experiments were performed on adult Emx1-Cre (B6.129S2-Emx1$^{tm1(cre)Krj}$/J; Jax # 5628) x Ai95 (B6;129S-Gt(ROSA)26Sor$^{tm95.1(CAG-GCaMP6f)Hze}$/J; Jax # 24105) transgenic mice expressing GCaMP6f in cortical excitatory neurons (Madisen et al., 2010). Male and female animals were used for the experiments (n = 16: 11 males, 5 females). Housing was enriched with a wheel, a tunnel, nesting material, and toys in a 12-hour light cycle, with food *ad libitum*. Protocols were in accordance with the French and European (2010/63/UE) legislations relative to the protection of animals used for experimental and other scientific purposes. All experimental procedures were approved by the

French Ministry of Education and Research, after consultation with the Ethical Committee #59 (authorization number: APAFIS#25932-2020060813556163 v7).

Surgery

Adult mice (at least 8 weeks of age) were anesthetized using isoflurane anesthesia (induction 3 to 4%, maintenance 1 to 1.5%). Suppression of paw withdrawal, whisker movement, and eye-blink reflexes was used to control for the depth of anesthesia. Mice received anti-inflammatory medication (Meloxicam (8mg/kg), subcutaneously), were placed on a heating blanket with the head stabilized by a custom-designed nose clamp. Their eyes were kept moist with Ocry-gel (TVM Lab). Following a local subcutaneous injection of lidocaine (4 mg/kg), the skin covering the skull was cut, and Povidone-iodine 5% was applied to treat the surgical wound. A drop of 3% hydrogen peroxide was applied to the skull surface, and the skull was scraped with a scalpel. An oval-shaped craniotomy with a 6 mm longer axis was performed above the left somatosensory and motor cortices. A 6 mm transparent window was glued in place using gel cyanoacrylic glue. A head fixation post was glued on the occipital bone. Dental cement was used to cover the skull and secure the head post and cranial window. Mice were monitored daily post-surgery until its weight returned to the pre-surgery level.

Calcium imaging and real-time wave tracking

An imaging and real-time image-processing setup (Figure 1A) was assembled (R&D Vision, France) from a JAI GO-2400M-PMCL camera mounted on a custom-built tandem-lens epifluorescence macroscope (Ratzlaff and Grinvald, 1991) that combined two photographic lenses in a face-to-face configuration: a 50 mm f/1.2 Nikkor lense (sample side), and a 12.5 – 75 mm f/1.2 Fujinon zoom lens (camera side). Recordings were performed at 200 Hz with two alternating light sources: green, mounted on the side of the optical path (Thorlabs M530L4 green LED light source with a Thorlabs SM1U25-A collimator and Thorlabs LEDD1B driver) and blue, mounted on the main optical path (SciMedia Ltd. LEX3-Blue LED light source with a Chroma ET480/40x filter). Separation of the blue source and camera light paths was achieved with a dichroic mirror (Semrock FF495-Di03 reflecting all wavelengths < 495 nm). A band-pass filter (Edmund Optics 525/45, OD6) restricted to green the spectrum of light that reached the camera. The blue light was used for GCaMP6f fluorescence excitation, while the green reflected light aimed to capture hemodynamic-related signals that contaminate the GCaMP6f fluorescence. To correct the GCaMP6f signal from these intrinsic contaminants, each image acquired under blue illumination was divided in real-time by the subsequent image taken under green illumination, generating a new sequence of frames at 100 Hz. To prevent high round-off error during division of integer-valued images, each blue image was multiplied by 64 before division.

After the correction for hemodynamics, fluorescence variation (dF/F_0) images were obtained relative to a baseline fluorescence image F_0, calculated from the 100 frames preceding the current frame. A 2D Mexican wavelet was applied to the data in the spatial dimension to retain components of cortical activity with sizes between 100 µm and 1 mm. After wavelet application, Otsu's method (Otsu, 1979) was used to create a binary mask corresponding to blobs of activity, and the local maximum within each blob, representing either a stationary or moving wave, was identified.

For each frame, all local maxima detected in that frame were compared with the last known coordinates of all tracked waves during the preceding ten frames. A local maximum detected in the current frame was considered part of an existing tracked wave if the distance between the current local maximum and the last known coordinate of the wave (within the preceding ten frames) was minimal and not greater than ten pixels (0.2 mm), and if the difference in relative fluorescence value was less than 0.01 relative units. If no such match was found, the detected maximum was considered a new wave to be tracked. Conversely, if no new match was found for a previously tracked wave for ten frames, that wave was considered to have ended.

To be considered Conditioned Wave (CW), a wave must have been detected first in the Start zone, in at least one frame, and then be tracked until it is detected in the End zone for at least two frames. Each new Conditioned Wave detected opened immediately a 500 ms opportunity window for reward

delivery. Each subsequent frame in which the wave continued to meet the Conditioned condition extended the opportunity window by 10 ms. All licks detected inside the opportunity window were rewarded. If at least one lick was detected, the CW was labeled as Rewarded. The latency between image acquisition and opportunity window opening in the case of Conditioned Wave detection was less than 10 ms, and the system was able to track up to 10 waves simultaneously. Multiple waves could be detected as Conditioned simultaneously. If another wave was detected during the opportunity window, their opportunity windows would overlap.

Cortical mapping

After the mouse fully recovered from surgery (minimum of 3 weeks), the right hindlimb and forelimb were stimulated with a rod actuated by a piezoelectric bender, to localize the hindlimb- and forelimb-related regions in the somatosensory cortices of the anesthetized mouse (isoflurane 0.75–1.5%). The widefield calcium signal from 45 stimulations was recorded and averaged. The one frame acquired 350 ms after stimulus onset was filtered using a 30x30 2D median filter. An extended-maxima transform (h = 0.01) was then applied to segment and identify the contour of the zone of evoked response.

Habituation protocol and setting of initial task parameters

Once the mouse was trained to stay head-fixed in the setup, it was placed on a water restriction protocol and trained to receive water from a lick-port. Spontaneous activity was collected to determine the placement of the Start and End zones so that 1-5 Conditioned Waves spontaneously occurred per minute. For 15 of the 16 mice, there was no gap between the Start and End zones in the initial training protocol. For mouse #13, a 130 µm gap was present between the zones since the start of the training.

Training Protocol

Water-restricted mice were placed in the setup, and an image of the brain was taken under blue illumination. Alignment of the Start and End zones between the current and previous sessions was performed using phase correlation registration. The training protocol continued for 10 – 30 minutes per session, with one session per day. During the training session, water droplets were delivered by a lick-port upon lick detection by a capacitive sensor within the reward opportunity window. During the session, we recorded the times of licks, opportunity window opening times (t_{CW}), reward delivery times, real-time detected wave trajectories, calcium signals and body movement (see below) synchronously. Only the first 10 minutes of session time were used for analysis, when the highest motivation was apparent.

Analysis of behavioral performance

In order to identify the specific sessions during which a mouse performance was significantly better than baseline, we computed for each session the distribution of the time intervals between rewarded waves (Figure 1G). A reference, *chance performance* distribution was defined by cumulating the intervals between Rewarded CWs observed during training sessions 1 and 2. A Kruskal-Wallis test was then applied to the set of distributions of intervals between Rewarded CWs computed across training sessions. If this first test identified that not all distributions of intervals between Rewarded CWs had the same distribution ($p < 0.05$), we then applied a Tukey-Kramer multiple comparison test to identify the sessions with a distribution significantly shorter (i.e., Rewarded CWs appeared more frequently) than that of sessions 1 & 2 ($p < 0.05$). Only time intervals longer than 500 ms were considered, corresponding to the duration of the reward opportunity window.

Recording and quantification of body movement

During the training and habituation sessions, body movement was recorded at 100 Hz using a DMK33UX273 camera with a Fujinon DF6HA-1B 1:1.2/6 mm lens, equipped with an EFFI-RLSW-00-050-IR infrared ring light. The camera was synchronized with the blue light source used for calcium imaging. The right-side (contralateral to the craniotomy) hindlimb and forelimb were tracked using DeepLabCut (version 2.1.8.2) employing a ResNet-101 convolutional neural network (Mathis et al. 2018). Two points were tracked per limb, at its proximal and distal aspects. For each frame,

the x/y coordinates of each limb were calculated as the average of the corresponding proximal and distal points. The pixel coordinates were converted into millimeters based on the headpost length located in the median plane, used as a reference. To reduce potential detection noise, a Kalman filter (acceleration variance 9×10^6 mm²/s⁴) was applied to each limb coordinate to estimate position and instantaneous velocity. A generic instantaneous "Limb speed" was computed as the average of the absolute speed of the hindlimb and forelimb. Velocity peaks above 15 mm/s were detected to identify individual movements. For each detected movement, we searched from the peak, in both time directions, until velocity fell below 4 mm/s. Two lines passing through the peak and these points defined the movement onset and offset as their crossings with zero.
For mice #5 and #6, body movements were not recorded during the first training session; therefore, they were excluded from the initial movement analysis (Figure 4) but not from movement analysis after opening the gap between Start and End zones (Figure 5).

Opening the gap between Start and End zones
9 of the 10 mice attributed to the Learning group were involved in the Gap condition (Mouse #16 was excluded). The gap opening was performed after an average of 21 sessions on the initial protocol. For 7 of the 9 mice, the End zone was shifted along the task main axis (Figure 4L) by 130 µm to create the gap. For mice #5 and #6, the 130 µm gap was added between the Start and End zones without moving the zone. For 7 of the 9 mice, the gap was introduced on the initial location of the Start and End zones. In addition, in two mice, the location of the Start and End zones was modified before opening the gap: in mouse #11, the Start and End zones were inverted for 14 sessions, and in mouse #14, both zones were rotated around their common center by 32° for 10 sessions.

Perfusion and Histology
The mice were anesthetized using isoflurane anesthesia (induction 3 to 4%, maintenance 1 to 1.5%) and placed on a stereotaxic frame. After removal of the implanted window, DiI fluorescent marker was injected at a depth of 1 mm at the corners of the Start and End zones. Following the administration of an overdose of pentobarbital (150 mg/kg), mice were perfused transcardically with saline followed by paraformaldehyde (4% in 0.1 M phosphate buffer). After overnight post-fixation of the brain in paraformaldehyde, 80 µm-thick coronal brain slices were exposed to 1:1000 DAPI for 5 minutes.

ACKNOWLEDGMENTS

Experimental assistance and technical expertise were provided by Aurélie Daret and Guillaume Hucher. We thank Thomas Deneux (NeuroPSI), Valentine Dhers and Olivier Lambert (R&D-Vision, France) for hardware and software development for the behavioral setup; Rodrigo Cofre and Alain Destexhe for their feedback and discussions throughout the project; Clément Picard, Rumeysa Can, and Konstantina Xanthopoulou Koukogia for their contribution to data analysis.

This work was funded by CNRS (CNRS 80|Prime, PRIME interdisciplinary label), Fondation pour la Recherche Médicale, Fondation 3DS, Agence Nationale pour la Recherche (ANR JCJC Mesobrain, ANR PRC Expect, Motorsense, Hermin, PROFouNd, PerBaCo), Lidex NeuroSaclay, Idex Brainscopes, iCODE, hCODE, and the H2020 RISE iNavigate project (number: 873178).

BIBLIOGRAPHY

Abbasi, Aamir, Henri Lassagne, Luc Estebanez, Dorian Goueytes, Daniel E Shulz, and Valerie Ego-Stengel. 2023. ‘Brain-Machine Interface Learning Is Facilitated by Specific Patterning of Distributed Cortical Feedback’. *SCIENCE ADVANCES*.

Allen, William E., Isaac V. Kauvar, Michael Z. Chen, Ethan B. Richman, Samuel J. Yang, Ken Chan, Viviana Gradinaru, Benjamin E. Deverman, Liqun Luo, and Karl Deisseroth. 2017. ‘Global Representations of Goal-Directed Behavior in Distinct Cell Types of Mouse Neocortex’. *Neuron* 94 (4): 891-907.e6. https://doi.org/10.1016/j.neuron.2017.04.017.

Anumanchipalli, Gopala K., Josh Chartier, and Edward F. Chang. 2019. 'Speech Synthesis from Neural Decoding of Spoken Sentences'. *Nature* 568 (7753): 493–98. https://doi.org/10.1038/s41586-019-1119-1.

Arduin, Pierre-Jean, Yves Frégnac, Daniel E. Shulz, and Valérie Ego-Stengel. 2013. '"Master" Neurons Induced by Operant Conditioning in Rat Motor Cortex during a Brain-Machine Interface Task'. *The Journal of Neuroscience* 33 (19): 8308–20. https://doi.org/10.1523/JNEUROSCI.2744-12.2013.

Arduin, Pierre-Jean, Yves Fregnac, Daniel E. Shulz, and Valerie Ego-Stengel. 2014. 'Bidirectional Control of a One-Dimensional Robotic Actuator by Operant Conditioning of a Single Unit in Rat Motor Cortex'. *Frontiers in Neuroscience* 8 (July). https://doi.org/10.3389/fnins.2014.00206.

Arieli, Amos, Alexander Sterkin, Amiram Grinvald, and Ad Aertsen. 1996. 'Explanation of the Large Variability in Evoked Cortical Responses' 273.

Brefczynski, Julie A., and Edgar A. DeYoe. 1999. 'A Physiological Correlate of the "spotlight" of Visual Attention'. *Nature Neuroscience* 2 (4): 370–74. https://doi.org/10.1038/7280.

Chen, Janice, Yuan Chang Leong, Christopher J Honey, Chung H Yong, Kenneth A Norman, and Uri Hasson. 2017. 'Shared Memories Reveal Shared Structure in Neural Activity across Individuals'. *Nature Neuroscience* 20 (1): 115–25. https://doi.org/10.1038/nn.4450.

Chen, Li M., Robert M. Friedman, and Anna W. Roe. 2003. 'Optical Imaging of a Tactile Illusion in Area 3b of the Primary Somatosensory Cortex'. *Science* 302 (5646): 881–85. https://doi.org/10.1126/science.1087846.

Clancy, Kelly B, Aaron C Koralek, Rui M Costa, Daniel E Feldman, and Jose M Carmena. 2014. 'Volitional Modulation of Optically Recorded Calcium Signals during Neuroprosthetic Learning'. *Nature Neuroscience* 17 (6): 807–9. https://doi.org/10.1038/nn.3712.

Clancy, Kelly B., and Thomas D. Mrsic-Flogel. 2021. 'The Sensory Representation of Causally Controlled Objects'. *Neuron* 109 (4): 677-689.e4. https://doi.org/10.1016/j.neuron.2020.12.001.

Crochet, Sylvain, and Carl C H Petersen. 2006. 'Correlating Whisker Behavior with Membrane Potential in Barrel Cortex of Awake Mice'. *Nature Neuroscience* 9 (5): 608–10. https://doi.org/10.1038/nn1690.

D'Aquila, Paolo S. 2024. 'Licking Microstructure in Response to Novel Rewards, Reward Devaluation and Dopamine Antagonists: Possible Role of D1 and D2 Medium Spiny Neurons in the Nucleus Accumbens'. *Neuroscience & Biobehavioral Reviews* 165 (October):105861. https://doi.org/10.1016/j.neubiorev.2024.105861.

Dugué, Laura, Elisha P. Merriam, David J. Heeger, and Marisa Carrasco. 2020. 'Differential Impact of Endogenous and Exogenous Attention on Activity in Human Visual Cortex'. *Scientific Reports* 10 (1): 21274. https://doi.org/10.1038/s41598-020-78172-x.

Ferezou, Isabelle, Sonia Bolea, and Carl C.H. Petersen. 2006. 'Visualizing the Cortical Representation of Whisker Touch: Voltage-Sensitive Dye Imaging in Freely Moving Mice'. *Neuron* 50 (4): 617–29. https://doi.org/10.1016/j.neuron.2006.03.043.

Fetz, Eberhard E. 1969. 'Operant Conditioning of Cortical Unit Activity'.

Gilad, Ariel, Yasir Gallero-Salas, Dominik Groos, and Fritjof Helmchen. 2018. 'Behavioral Strategy Determines Frontal or Posterior Location of Short-Term Memory in Neocortex'. *Neuron* 99 (4): 814-828.e7. https://doi.org/10.1016/j.neuron.2018.07.029.

Goueytes, Dorian, Henri Lassagne, Daniel E Shulz, Valérie Ego-Stengel, and Luc Estebanez. 2022. 'Learning in a Closed-Loop Brain-Machine Interface with Distributed Optogenetic Cortical

Feedback'. *Journal of Neural Engineering* 19 (6): 066045. https://doi.org/10.1088/1741-2552/acab87.

Grinvald, A, Ee Lieke, Rd Frostig, and R Hildesheim. 1994. 'Cortical Point-Spread Function and Long-Range Lateral Interactions Revealed by Real-Time Optical Imaging of Macaque Monkey Primary Visual Cortex'. *The Journal of Neuroscience* 14 (5): 2545–68. https://doi.org/10.1523/JNEUROSCI.14-05-02545.1994.

Herrmann, Katrin, Leila Montaser-Kouhsari, Marisa Carrasco, and David J Heeger. 2010. 'When Size Matters: Attention Affects Performance by Contrast or Response Gain'. *Nature Neuroscience* 13 (12): 1554–59. https://doi.org/10.1038/nn.2669.

Jancke, Dirk, Frédéric Chavane, Shmuel Naaman, and Amiram Grinvald. 2004. 'Imaging Cortical Correlates of Illusion in Early Visual Cortex'. *Nature* 428 (6981): 423–26. https://doi.org/10.1038/nature02396.

Jeon, Brian B., Thomas Fuchs, Steven M. Chase, and Sandra J. Kuhlman. 2022. 'Existing Function in Primary Visual Cortex Is Not Perturbed by New Skill Acquisition of a Non-Matched Sensory Task'. *Nature Communications* 13 (1): 3638. https://doi.org/10.1038/s41467-022-31440-y.

Jouvet, Michel. 1967. 'The States of Sleep'. *Scientific American* 216 (2): 62–72. https://doi.org/10.1038/scientificamerican0267-62.

Kenet, Tal, Dmitri Bibitchkov, Misha Tsodyks, Amiram Grinvald, and Amos Arieli. 2003. 'Spontaneously Emerging Cortical Representations of Visual Attributes'. *Nature* 425 (6961): 954–56. https://doi.org/10.1038/nature02078.

Koralek, Aaron C., Xin Jin, John D. Long Ii, Rui M. Costa, and Jose M. Carmena. 2012. 'Corticostriatal Plasticity Is Necessary for Learning Intentional Neuroprosthetic Skills'. *Nature* 483 (7389): 331–35. https://doi.org/10.1038/nature10845.

Liang, Yuqi, Junhao Liang, Chenchen Song, Mianxin Liu, Thomas Knöpfel, Pulin Gong, and Changsong Zhou. 2023. 'Complexity of Cortical Wave Patterns of the Wake Mouse Cortex'. *Nature Communications* 14 (1): 1434. https://doi.org/10.1038/s41467-023-37088-6.

Lorach, Henri, Andrea Galvez, Valeria Spagnolo, Felix Martel, Serpil Karakas, Nadine Intering, Molywan Vat, et al. 2023. 'Walking Naturally after Spinal Cord Injury Using a Brain–Spine Interface'. *Nature* 618 (7963): 126–33. https://doi.org/10.1038/s41586-023-06094-5.

Luczak, Artur, Peter Barthó, and Kenneth D. Harris. 2009. 'Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations'. *Neuron* 62 (3): 413–25. https://doi.org/10.1016/j.neuron.2009.03.014.

Ma, Ying, Mohammed A. Shaik, Mariel G. Kozberg, Sharon H. Kim, Jacob P. Portes, Dmitriy Timerman, and Elizabeth M. C. Hillman. 2016. 'Resting-State Hemodynamics Are Spatiotemporally Coupled to Synchronized and Symmetric Neural Activity in Excitatory Neurons'. *Proceedings of the National Academy of Sciences* 113 (52). https://doi.org/10.1073/pnas.1525369113.

Madisen, Linda, Theresa A Zwingman, Susan M Sunkin, Seung Wook Oh, Hatim A Zariwala, Hong Gu, Lydia L Ng, et al. 2010. 'A Robust and High-Throughput Cre Reporting and Characterization System for the Whole Mouse Brain'. *Nature Neuroscience* 13 (1): 133–40. https://doi.org/10.1038/nn.2467.

Matteucci, Giulio, Maëlle Guyoton, Johannes M. Mayrhofer, Matthieu Auffret, Georgios Foustoukos, Carl C.H. Petersen, and Sami El-Boustani. 2022. 'Cortical Sensory Processing across Motivational States during Goal-Directed Behavior'. *Neuron* 110 (24): 4176-4193.e10. https://doi.org/10.1016/j.neuron.2022.09.032.

Mohajerani, Majid H, Allen W Chan, Mostafa Mohsenvand, Jeffrey LeDue, Rui Liu, David A McVea, Jamie D Boyd, Yu Tian Wang, Mark Reimers, and Timothy H Murphy. 2013. 'Spontaneous Cortical Activity Alternates between Motifs Defined by Regional Axonal Projections'. *Nature Neuroscience* 16 (10): 1426–35. https://doi.org/10.1038/nn.3499.

Mohajerani, Majid H., David A. McVea, Matthew Fingas, and Timothy H. Murphy. 2010. 'Mirrored Bilateral Slow-Wave Cortical Activity within Local Circuits Revealed by Fast Bihemispheric Voltage-Sensitive Dye Imaging in Anesthetized and Awake Mice'. *The Journal of Neuroscience* 30 (10): 3745–51. https://doi.org/10.1523/JNEUROSCI.6437-09.2010.

Murray, S. O. 2008. 'The Effects of Spatial Attention in Early Human Visual Cortex Are Stimulus Independent'. *Journal of Vision* 8 (10): 2–2. https://doi.org/10.1167/8.10.2.

Musall, Simon, Matthew T. Kaufman, Ashley L. Juavinett, Steven Gluf, and Anne K. Churchland. 2019. 'Single-Trial Neural Dynamics Are Dominated by Richly Varied Movements'. *Nature Neuroscience* 22 (10): 1677–86. https://doi.org/10.1038/s41593-019-0502-4.

Neely, Ryan M., Aaron C. Koralek, Vivek R. Athalye, Rui M. Costa, and Jose M. Carmena. 2018. 'Volitional Modulation of Primary Visual Cortex Activity Requires the Basal Ganglia'. *Neuron* 97 (6): 1356-1368.e4. https://doi.org/10.1016/j.neuron.2018.01.051.

Okun, Michael, Amir Naim, and Ilan Lampl. 2010. 'The Subthreshold Relation between Cortical Local Field Potential and Neuronal Firing Unveiled by Intracellular Recordings in Awake Rats'. *The Journal of Neuroscience* 30 (12): 4440–48. https://doi.org/10.1523/JNEUROSCI.5062-09.2010.

Otsu, Nobuyuki. 1979. 'A Tlreshold Selection Method from Gray-Level Histograms' SMC-9 (1).

Poulet, James F. A., and Carl C. H. Petersen. 2008. 'Internal Brain State Regulates Membrane Potential Synchrony in Barrel Cortex of Behaving Mice'. *Nature* 454 (7206): 881–85. https://doi.org/10.1038/nature07150.

Prsa, Mario, Gregorio L. Galiñanes, and Daniel Huber. 2017. 'Rapid Integration of Artificial Sensory Feedback during Operant Conditioning of Motor Cortex Neurons'. *Neuron* 93 (4): 929-939.e6. https://doi.org/10.1016/j.neuron.2017.01.023.

Sakata, Shuzo, and Kenneth D. Harris. 2009. 'Laminar Structure of Spontaneous and Sensory-Evoked Population Activity in Auditory Cortex'. *Neuron* 64 (3): 404–18. https://doi.org/10.1016/j.neuron.2009.09.020.

Schiemann, Julia, Paolo Puggioni, Joshua Dacre, Miha Pelko, Aleksander Domanski, Mark C.W. van Rossum, and Ian Duguid. 2015. 'Cellular Mechanisms Underlying Behavioral State-Dependent Bidirectional Modulation of Motor Cortex Output'. *Cell Reports* 11 (8): 1319–30. https://doi.org/10.1016/j.celrep.2015.04.042.

Shahsavarani, Somayeh, David N. Thibodeaux, Weihao Xu, Sharon H. Kim, Fatema Lodgher, Chinwendu Nwokeabia, Morgan Cambareri, et al. 2023. 'Cortex-Wide Neural Dynamics Predict Behavioral States and Provide a Neural Basis for Resting-State Dynamic Functional Connectivity'. *Cell Reports* 42 (6): 112527. https://doi.org/10.1016/j.celrep.2023.112527.

Steriade, M., I. Timofeev, and F. Grenier. 2001. 'Natural Waking and Sleep States: A View From Inside Neocortical Neurons'. *Journal of Neurophysiology* 85 (5): 1969–85. https://doi.org/10.1152/jn.2001.85.5.1969.

Surinach, Daniel, Mathew Rynes, Kapil Saxena, Eunsong Ko, A. Redish, and Suhasa Kodandaramaiah. 2023. 'Strategy Dependent Recruitment of Distributed Cortical Circuits during Spatial Navigation'. https://doi.org/10.21203/rs.3.rs-2997927/v1.

Wang, Ziyue, Xiang Fei, Xiaotong Liu, Yanjie Wang, Yue Hu, Wanling Peng, Ying-wei Wang, Siyu Zhang, and Min Xu. 2022. 'REM Sleep Is Associated with Distinct Global Cortical Dynamics and

Controlled by Occipital Cortex'. *Nature Communications* 13 (1): 6896. https://doi.org/10.1038/s41467-022-34720-9.