Operant conditioning of cortical waves through a brain-machine interface

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

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

Correspondance: luc.estebanez@cnrs.fr

ABSTRACT

At the surface of the cerebral cortex, activity dynamics measured at a mesoscopic (about 0.1 to 1 mm in mouse cortex) scale are characterized by both spontaneous and behavior-related dynamical waves of synchronized neuronal activity. These waves are thought to participate in information propagation and processing, but remain poorly understood.

To assess if such mesoscopic coordinated neuronal dynamics can be controlled in a goal-directed manner, we implemented a task in which mice are trained to generate waves with specific trajectories in order to obtain rewards. We tracked propagating waves at the surface of the somatosensory-motor cortex of head-fixed mice in real time by means of wide field calcium imaging, and conditioned the delivery of rewards to the detection of wave trajectories responding to specific spatiotemporal criteria.

We found that the majority of the trained mice significantly increased their performance, mainly by increasing the frequency of rewardable waves. As the mice learned to achieve this task, we observed changes in the spatiotemporal patterns of the cortical waves. By revealing that, upon learning, neuronal activity can be shaped at the mesoscopic scale to generate specific waves patterns, our work opens up new perspectives for brain-machine interfacing.

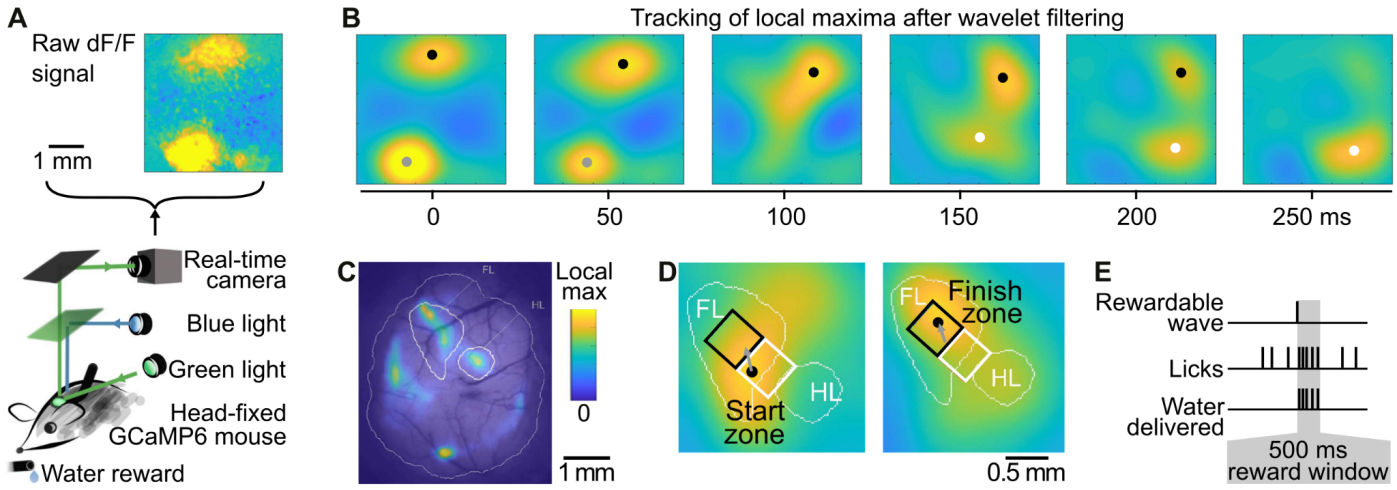

Figure 1. Setup for tracking mesoscale waves in real time, and operant conditioning task.
(A) General schematics of the setup. A water-restricted GCaMP6f-expressing mouse with a chronic window above the somatosensory and motor cortices was head-fixed under a macroscope. Note that the delivery of water reward is conditional to licking.
(B) Example of recorded cortical activity. GCaMP6f fluorescence signal imaged at 100 Hz was processed in real-time. After hemodynamic artifact removal, the images were filtered with a 2D wavelet transform (100 µm to 1 mm bandpass), and the position of local maxima corresponding to each wave was tracked (filled circles).
(C) Density of the local maxima of waves detected during one session in a naïve mouse (#11). White outlines indicate the forelimb (FL, left) and hindlimb (HL, right) cortical representations defined by functional intrinsic imaging.
(D) Example cortical wave traveling across the cortical surface. Black dot: position of the local maximum of the wave. White and Black squares: Start and Finish zones defined by the experimenter. The maximum of the filtered wave traveled from the Start to the Finish zone, with a time interval of 200 ms between the two frames.
(E) Waves that were detected in the “Start” and then in the “Finish” zone triggered the opening of a 500 ms reward opportunity window, during which the mouse could trigger delivery of water rewards from the lick-port by licking.

INTRODUCTION

One of the most striking characteristics of cortical dynamics is the existence of distinct activity states, separated by rather abrupt transitions from one state to the next. These different activity regimes, which were first described in sleep (Jouvet 1967; Steriade, Timofeev, and Grenier 2001), can dominate the membrane potential fluctuations of the neurons during wake, both in motor cortices (Schiemann et al. 2015) and in sensory cortices (Crochet and Petersen 2006; Luczak, Barthó, and Harris 2009; Okun, Naim, and Lampl 2010; Poulet and Petersen 2008; Sakata and Harris 2009).

Imaging of these dynamics at the scale of cortical areas — the *mesoscale* — using membrane voltage reporters (Arieli et al. 1996; Grinvald et al. 1994) and calcium reporters of spiking activity (Allen et al. 2017; Gilad et al. 2018), revealed that the awake brain active state decomposes into a constant flow of waves of neuronal depolarization (Ma et al. 2016; Mohajerani et al. 2010). These waves move from sources, often located in the primary sensory areas, towards sinks — some in associative areas (Mohajerani et al. 2013, Liang et al., 2023).

Despite the prominence of these dynamics at the surface of the cortex during behavior, their function remains unclear. Time averaging of cortical waves during trained behaviors reveals networks of sensory, motor and high-order cortical areas that activate in a highly structured fashion as the behavior unrolls (J. Chen et al. 2017; Gilad et al. 2018; Musall et al. 2019; Shahsavarani et al. 2023). However, the 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 is not controlled by the subject and does 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, individual cortical state transitions (Crochet and Petersen 2006; Ferezou, Bolea, and Petersen 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 to information processing of individual waves is also supported by evidence that they can carry sensory information, as indicated by the discovery of spontaneous cortical waves of activity that adhere to the topographical representations of sensory inputs during anesthesia (Kenet et al. 2003; Mohajerani et al. 2013), and by experiments highlighting that individual cortical waves may be the source or substrate of illusory percepts (Chen, Friedman, and Roe 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, which could be actively controlled 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 for interrogating the agency of a subject on internal variables of its brain. In particular, neuronal operant control training has uncovered the ability of a subject to control 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 asked whether mice could control some of the properties of cortical waves. More precisely, we tested if they can learn — via operant conditioning — to generate individual waves with a specific trajectory at the surface of the brain. We further asked if this may have an impact on the broader cortical dynamics around the conditioned wave. To meet these aims, 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.

We found that a majority of the mice learned to increase the frequency of the rewardable waves, and licked during the reward windows in order to collect water. While increasing their performance, the mice reshaped their cortical activity by reducing the variability of the rewarded waves, as well as reducing overall cortical activity across the sensori-motor cortex before and after the rewarded waves.

RESULTS

Mice learned to significantly increase the frequency of 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 the mice were habituated to remain head-fixed on a setup and lick from a lick-port, we recorded the spontaneous cortical activity of naïve mice through widefield calcium imaging at 100 Hz (Figure 1A). Local maxima were detected in the signal after filtering with a wavelet transform, and their x-y-coordinates were tracked in real time (Figure 1B).

We observed that the density of local maxima in naïve mice was not spatially homogeneous, with a higher density located in the somatosensory and motor cortical areas corresponding to the contralateral hindlimb and forelimb (Figure 1C). For each mouse, we selected two nearby 500 × 500 µm square zones located in the somatosensory-motor cortex, such that the frequency of waves detected in the first zone, called “Start”, and moving to the second zone, called “Finish,” was around 1-5 waves per minute (Figure 1D). These zones could be located in different cortical regions for different mice.

We conditioned the mice to create such waves in order to obtain a water reward. As soon as a candidate wave, here labeled “rewardable wave”, was detected, a 500 ms opportunity window opened, allowing the mouse to obtain a water reward if licking (Figure 1E). Licking outside of the opportunity window did not lead to a reward. All the analyses were carried out in the first ten minutes of each session, a period during which we observed consistently high thirst motivation (Matteucci et al. 2022).

We defined the mouse performance during the session as the number of rewarded waves per minute during the first ten minutes of the session. To probe the evolution of mice performance with training (Figure 2A), we focused on the distribution of the time intervals between rewarded waves (Figure 2B). We identified mice that learned the task by looking for the distributions of inter-reward intervals across sessions that were significantly skewed towards short values (i.e., rewarded waves appeared more frequently), versus the reference distribution obtained in sessions 1 and 2. This was achieved with a Kruskal-Wallis test followed by multiple comparison tests to identify all sessions where the performance was significantly higher (see methods, Figure 2C).

Of the 16 mice involved in the training protocol, 10 showed a significant increase in performance ("Learning" group, increase in performance shown in Figure 2D-F), while the remaining mice were classified as the "Not Learning" group (6 mice). In the Learning group, the first session meeting the learning criterion was achieved after 11.1 ± 5.4 (mean ± STD) consecutive training sessions (Figure 2E, top). Several of these mice kept progressing with additional training (Figure 2F) and on average achieved their best performance after 15.0 ± 5.6 consecutive training sessions (Figure 2E, bottom). It should be noted that although for various practical reasons not all the mice underwent the same number of training sessions, all Learning mice remained in training for at least 18 sessions (Figure 2G).

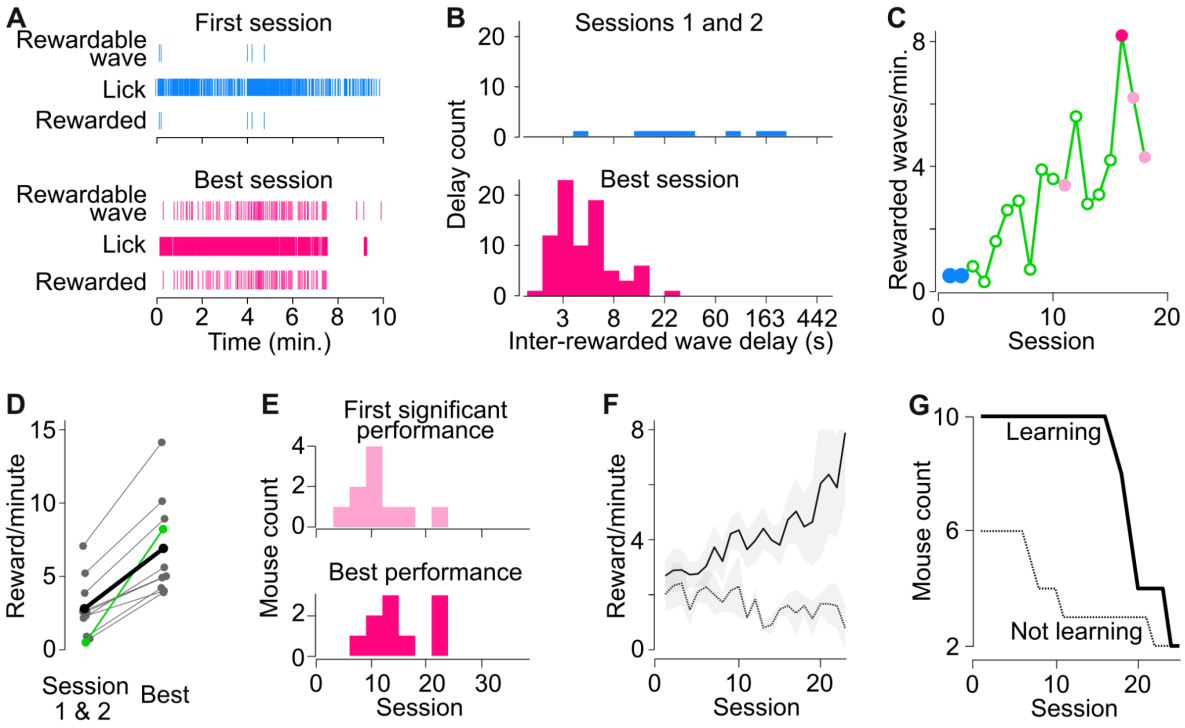

Figure 2. Mice increase the reward rate with learning
(A) Time sequence of the rewardable waves, licks, and rewarded waves for mouse #11 during the first session (top, blue) versus its best session (bottom, magenta).
(B) Histogram of the time intervals between consecutive rewarded waves for mouse #11 during the first two sessions (top) versus the session with the highest reward rate (bottom), shown on a logarithmic scale.
(C) Frequency of rewarded waves across the training sessions of mouse #11. Saturated magenta: best performance (same as session in panel A). Light magenta: Sessions where the inter-rewardable wave intervals were significantly shorter than during the combined sessions 1 and 2 (in blue).
(D) Frequency of rewarded waves from session 1 & 2, versus the session with the best performance plotted for the Learning mice. Light gray: individual sessions. Green: case study (mouse #11). Black: population average.
(E) Population distribution of the number of training sessions until the first session with a significant increase in performance (top) and until the best performance (bottom).
(F) Average reward frequency across mice that showed a significant increase in performance (Learning group, solid line, N = 10), versus no significant increase (Not learning group, dotted line, N = 6), Light background: SEM.
(G) Evolution of mice withdrawn from the training protocol across training sessions. In total, 10 mice were involved in the *Learning* group, and 6 in the *Not Learning* group.

Performance in the learning group increases with the frequency of rewardable waves

The increase in performance in the Learning group could be due to several factors. First, the mice could generate more waves of activity traveling from the Start to the Finish zone, thereby opening more windows of reward opportunity. By comparing the occurrence of rewardable waves during the first two sessions (combined) versus the session with the best performance, we observed a significant increase in frequency (Figure 3A). This increase in rewardable waves was accompanied by an overall increase in the frequency of waves across the imaging field (Figure 3B), as well as a specific increase in the proportion of waves that matched the operant conditioning constraints among all waves (Figure 3C). Similarly, when focusing on local waves (defined as all waves that were detected in the Start zone, irrespective of their trajectory), we found that their frequency increased significantly with training (Figure 3D), but again, the proportion of waves that matched the operant conditioning constraints among these waves increased significantly (Figure 3E).

To further examine this buildup of rewardable waves with learning, we mapped the density of all wave maxima at the surface of the cortex (Figure 3F). We found that the density of waves increased with training, mainly in an area less than 0.5 mm away from the border between the Start and Finish zones (Figure 3G,H). This increase of wave density at the Start/Finish border was particularly visible when limiting the mapping to Local waves (Figure 3I) or only to rewardable waves (Figure 3J).

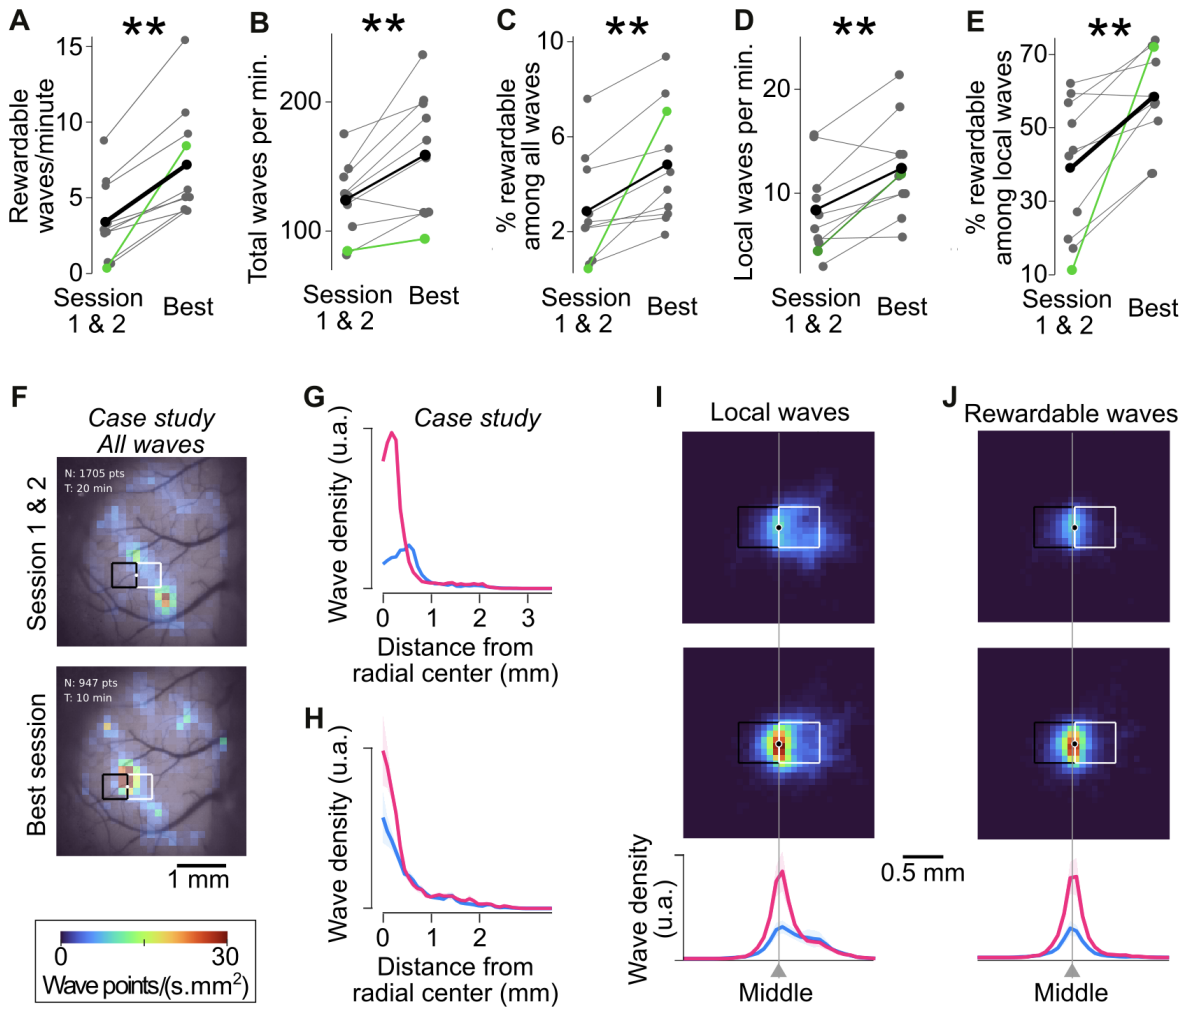

Figure 3. Mice increase task performance by differentially increasing the frequency of rewardable waves
(**A**) Frequency of rewardable waves during the first two sessions and during the session with the best performance. **: Wilcoxon p = 0.002. Light gray: individual mice. Green: case study (mouse #11). Black: population average.
(**B**) Same as **A** for the frequency of all waves detected on the imaged cortical surface. **: Wilcoxon p = 0.0098.
(**C**) Same as **A** for the percentage of all waves that were rewardable. The proportion increases significantly over training. **: Wilcoxon p = 0.002.
(**D**) Same as **A** for the frequency of "Local waves", which are all waves detected in the Start zone, including the rewardable waves. **: Wilcoxon p = 0.0059.
(**E**) Same as **A** for the percentage of the Local waves that were rewardable. The proportion increases significantly over training. **: Wilcoxon p = 0.0039.
(**F**) Case study: in mouse #11, density of maxima of all cortical waves tracked by our setup, on the first sessions (top) and on the best performance session (bottom).
(**G**) Radial distribution of wave maxima density, computed for all waves of mouse #11. First and second (blue) versus the best session (magenta). The center point used for the analysis is positioned at the middle of the border between the start and finish zones.
(**H**) Same as G at the population level in the Learning group. Light background: SEM.
(**I**) Population average of the density of the local wave maxima at the surface of the cortex in the Learning group (N = 10). All sessions are re-aligned on the Start and Finish zones. Top: Sessions 1&2. Middle: Best session. Bottom: marginal distribution of wave density along the main axis of the Start/Finish zones.
(**J**) Same as G, limited to rewardable waves.

Overall, we found that in a context of an overall increase in wave frequency, the mice differentially increased the frequency of rewardable waves, and in particular waves that are focused on the task-defined Start/Finish border.

A second possibility for the mice to improve their behavioral performance is to coordinate their licking with the onset of rewardable waves, i.e., with the time window when reward becomes available. To probe this, we segmented the quantification of licking into bursts, which are stereotyped lick sequences (D'Aquila et al. 2024, see Methods). Across training, we did not observe a systematic change in the overall lick burst rate (Figure 4A) — although the frequency of individual licks

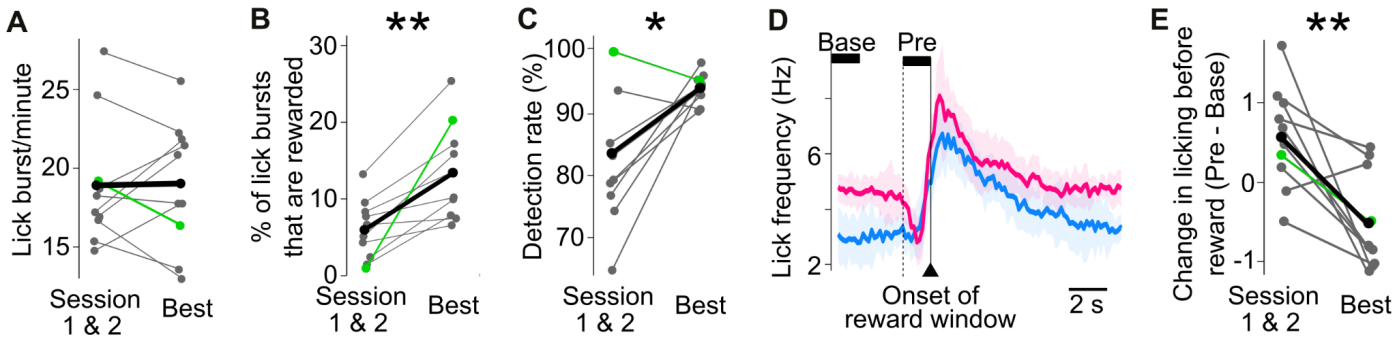

Figure 4. Evolution of licking strategies during learning

(**A**) Frequency of licking, measured as the frequency of the onset of the bursts of licks. The change during learning is not significant at population level.

(**B**) Percentage of the lick bursts that include licks that were rewarded. The proportion increased significantly. **: Wilcoxon p = 0.002.

(**C**) Percentage of the rewardable waves that were actually rewarded. The increase in this detection rate is significant. *: Wilcoxon p = 0.027.

(**D**) Average time course of the instantaneous licking frequency during the first and second (blue) versus the best session (magenta), at the start of the opening of reward windows. Light background: SEM.

(**E**) Change of licking rate in the Pre-Rewarded wave window (labeled Pre in D), normalized by the baseline rate from the Base window (see D). Over training, the licking began to drop just before the occurrence of a rewarded wave window, compared to baseline. **: Wilcoxon p = 0.004.

increased significantly (Wilcoxon p = 0.002) see baseline licking rate in Figure 4D. However, the proportion of rewarded lick bursts — in which at least one lick occurred during the opportunity window — increased significantly (Figure 4B). Simultaneously, the proportion of rewardable waves that were actually rewarded increased significantly with training sessions (Figure 4C). Overall, during training, the mice learned to produce licking patterns that were more efficient in collecting water rewards during the open reward time windows.

This led us to ask if the instantaneous relationship between licking and rewardable waves was also shaped by learning. To this aim, we computed the temporal distribution of licks around rewarded wave onset (Figure 4D). We found that following training, the licking rate dropped significantly below baseline in the 1.5 s time window preceding the occurrence of a rewardable wave (Figure 4E). We hypothesize that this may mark a preparatory time when the mouse stops licking as it focuses on the production of a rewardable wave and accompanying licks.

Diversity of the strategies to solve the wave trajectory task.

We described earlier (Figure 3F-J) the learning-related increase in wave frequency at the Start/Finish border. We asked if the learning was also impacting the structure of individual waves. To this aim, we explored the evolution of descriptors of individual waves, namely the distance covered by the wave travel between the first and last detected local maxima pertaining to the wave; the variability of the orientation of the wave travel (color-coded on wave trajectories in Figure 5A,B), and the ratio between the count of rewardable waves, versus the number of inverse waves that entered the Finish before the Start area.

We found a broad diversity in the evolutions of the wave descriptors during training, but no systematic learning-related change. Still, the properties of wave paths were dominated by two stereotypes (Figure 5C,D). On one side, we found mice that generated small size waves that were not consistent in their orientation from one rewardable wave to the next. These mice generated almost as many waves in the rewardable direction as in the inverse direction (example mouse #8 in Figure 5A). This strategy possibly corresponds to the creation of “standing waves” at the borderline between the “Start” and “Finish” zones, which minimized the travel required to be detected as Rewardable.

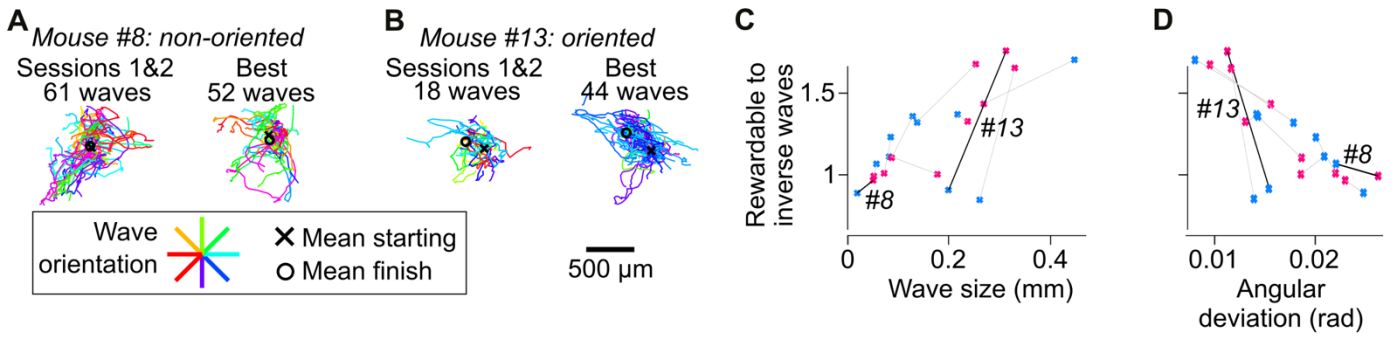

Figure 5. Shaping of cortical waves trajectories during learning.
(**A**) All trajectories of rewardable waves, color-coded by wave orientation (defined as the angle of the segment connecting the first and last maxima of the wave) of mouse #8, *non-oriented*. All waves from session 1&2 and best sessions are shown. Even after training, multiple wave orientations are apparent.
(**B**) Mouse #13 is *oriented*. After training, one wave orientation has become dominant.
(**C**) Population scatter plot showing in session 1&2 (blue) and best session (magenta) the distance of the waves from start to finish (x axis), versus the ratio of the count of rewardable waves and the count of waves that crossed the start and finish windows in the reverse direction (y axis). Examples #8 and #13 are highlighted by a black line.
(**D**) Same as C, with the angular deviation in the y axis. Angular deviation is the average angular distance from the orientation of individual waves to the average rewardable wave.

On the other side, we observed mice that generated waves covering larger distances, and which were clearly oriented. These mice selectively generated waves in the rewarded direction, much more than in the inverse direction (example mouse #13 in Figure 5B). This strategy corresponds to the generation of traveling waves covering longer distances and more specifically oriented to meet the task constraints. Between these two extrema, we found that most mice generated waves with intermediate characteristics (Figure 5C,D).

Spatiotemporal shaping of cortical activity related to rewarded waves

So far, we have focused our analysis on the displacements of the local wave maxima with respect to the “Start” and “Finish” zones. Here we hypothesize that the promotion of specific wave trajectories and the simultaneous licking that we observed requires a larger-scale shaping of the cortical activity, beyond the immediate cortical location of the Start/Finish zones. To probe this hypothesis, we analyzed the hemodynamics-corrected widefield calcium signals — without the wavelet processing — corresponding to rewarded waves. During the first two sessions (combined), the calcium activity around the reward exhibited a broad spatial spread (Figure 6A). In contrast, during the highest performance session, the average calcium activity around reward time was more focused in space around the “Start” and “Finish” zones (Figure 6B) and revealed a strong suppression of the activity before and after reward time. This drop in cortical activity before and after the rewarded wave was also visible at the population level (Figure 6C).

Beyond the average dynamics, we found that the variability of the dynamics evolved with training, and decreased from the first to the last training sessions (example in Figure 6D-E, population analysis in Figure 6F). These findings suggest that beyond the travel of the local maxima of the waves, we achieved an operant conditioning of the cortical activity at the mesoscale.

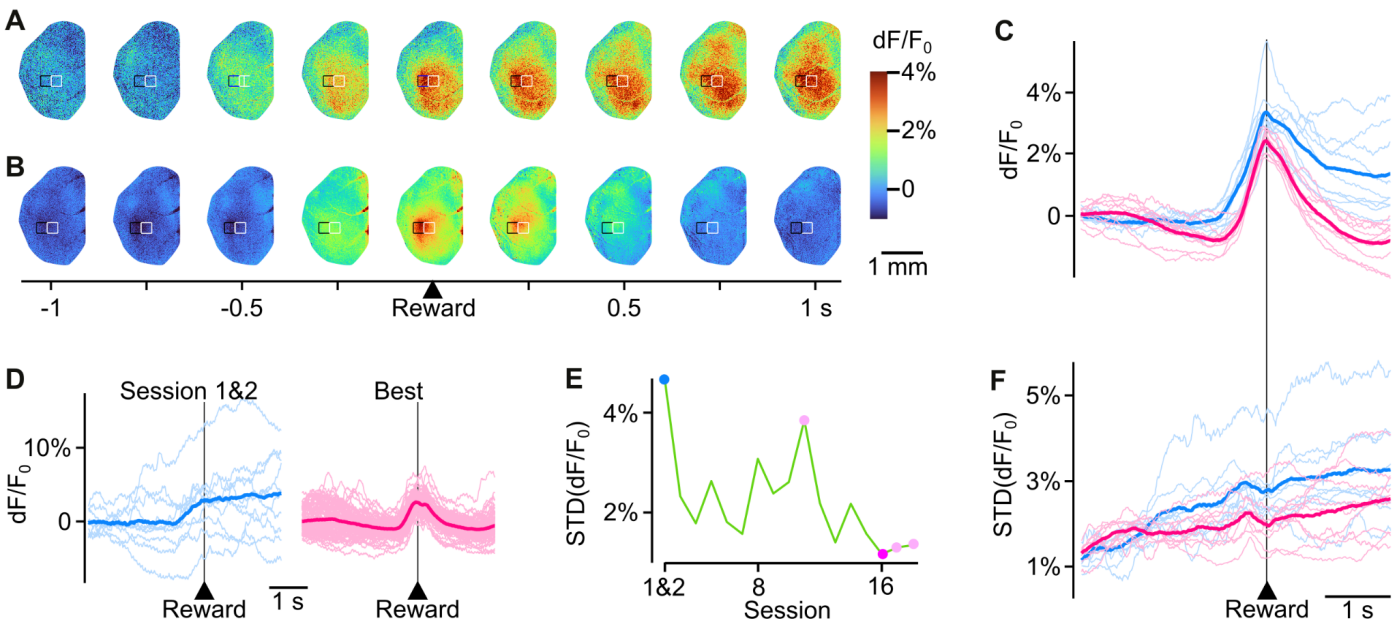

Figure 6. Spatiotemporal shaping of cortical activity beyond the wave local maxima
(**A**) Average GCaMP6 signal across the cortical window, aligned on reward time, during sessions 1&2 (mouse # 11).
(**B**) Same as A for the best performance session.
(**C**) Average reward-aligned calcium fluorescence signal across all mice (*Learning* group, N = 10) around reward times. Thin lines: individual mice. Thick line: population average. Blue: sessions 1&2. Magenta: best session.
(**D**) Average calcium fluorescence signal in the 'Start' area around reward times. Individual trials (thin lines) from the same mouse as in A. Thick line: average. Left: Sessions 1&2. Right: Best session.
(**E**) Evolution of the standard deviation of the calcium fluorescence signal across sessions. Blue: sessions 1&2. Light Magenta: sessions with significantly higher reward frequency as in sessions 1&2. Dark magenta: best session.
(**F**) Same as C for the standard deviation of the signal.

DISCUSSION

Using an optical brain-machine interface (BMI) based on real-time processing of wide field calcium signals, we have found that cortical mesoscale wave activity can be operantly controlled by mice in the context of a goal-directed task.

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, Galiñanes, and Huber 2017) or small patches of cortical tissue (Clancy and Mrsic-Flogel 2021). In a previous study, when we conditioned a group of neurons to solve a task, a single neuron most often took the lead in controlling the BMI (Abbasi et al. 2023, see also Jeon et al. 2022). Here, in contrast, we conditioned the delivery of a reward to the coordinated activity of a large population of cortical neurons that is needed for the emergence of a mesoscopic wave. Beyond this, we have shown that this cortical wave task leads to large scale changes in activity across the trained cortical hemisphere, both before and after the rewarded wave (Figure 6A-C). Therefore, our experiments open up the possibility to operantly condition the spatio-temporal features of brain activity at larger scales. This could include the shaping of electrocorticogram (ECoG)-scale dynamics, in addition to the decoding strategies that have dominated human application of ECoG, both towards the decoding of speech (Anumanchipalli, Chartier, and Chang 2019), and body control restoration (Lorach et al. 2023).

Decomposition of complex mesoscale cortical dynamics into movements of local waves maxima

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 manipulated by the mouse, we sought to compute in real time a low-dimensionality descriptor of each individual wave, instead of these established complex descriptors. Specifically, we selected the first order descriptor, that is, the wave position. Several definitions of the 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.

Mechanisms underlying the generation of rewardable waves

The rewardable waves recorded during the first sessions of training in naïve animals may emerge from very different mechanisms than the ones produced in expert animals after learning. Indeed, the rewarded waves recorded in trained mice were preceded and followed by a strong suppression of cortical activity, which was not visible during the first two training sessions (Figure 6).

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, in our case in the primary somatosensory and motor areas (Figure 1C). In addition, the rewarded cortical 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. Finally, we have found that it was challenging to extend the distance traveled by the rewarded waves beyond the size of physiological waves (data not shown).

We conclude that rewarded waves in our experiments are likely to rely on a similar anatomical substrate as spontaneous cortical waves. In addition, we hypothesize that the suppression of activity that we observed around the time of the reward after training may be due 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.

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 therefore be required to determine the actual subcortical circuitry and mechanisms that support the generation of rewardable waves, including the suppression of cortical activity prior to and after reward time.

MATERIALS AND METHODS

Animals

Experiments were performed on adult Emx1-Cre (Jax # 5628) x Ai95 (Jax # 24105) transgenic mice expressing GCaMP6f in a majority of cortical excitatory neurons (Madisen et al., 2010). Male and female animals were used for the experiments (n=16). Housing was enriched with a wheel, a tunnel, nesting material, and toys in a normal 12-hr 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 comply with the ARRIVE guidelines, and 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. A heating blanket maintained the body temperature measured rectally at 37°C. The eyes were kept moist with Ocry-gel (TVM Lab) and the head was stabilized with a custom-designed nose clamp. Following a subcutaneous injection of lidocaine (4 mg/kg), the skin covering the skull was cut, and diluted Betadine (1:2) 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 using a drill above the left somatosensory and motor cortices. A 6 mm transparent window was glued in place using gel cyanoacrylic glue. A head fixation bar was glued with liquid cyanoacrylic glue above the midline of the brain. The surgical wound edges were glued to the skull with cyanoacrylic glue, and dental cement was used to cover the skull. A subcutaneous injection of Meloxicam (8 mg/kg) was administered in the gluteal region. The mouse was monitored post-surgery until its weight returned to the pre-surgery level.

Cortical mapping

After the mouse fully recovered from surgery (minimum of 3 weeks), the right hindlimb and forelimb were stimulated with a piezoelectric element to localize the hindlimb- and forelimb-related regions in the somatosensory and motor cortices of the anesthetized mouse (isoflurane 0.75–1.5%). The average widefield calcium signal from 45 stimulations was recorded. The averaged image corresponding to 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.

Calcium imaging and real-time wave tracking

A real-time setup was designed and assembled in collaboration with R&D Vision (Saint-Maur-des-Fossés, France). A JAI GO-2400M-PMCL camera with a 12.5 – 75mm f/1.2 zoom lens and an Edmund Optics 525/45 50 mm diameter OD6 mounted filter was used to record widefield calcium signals. Recordings were performed at 200 Hz with two alternating light sources: green (Thorlabs M530L4 green LED light source with a Thorlabs SM1U25-A collimator and Thorlabs LEDD1B driver) and blue (Brainvision LEX3-Blue LED light source with a Chroma ET480/40x 25 mm diameter mounted filter). The green light source was used to capture hemodynamic signals, and the blue light source was used for GCaMP6f excitation. To minimize hemodynamic effects, each image taken under blue illumination was divided by the subsequent image taken under green illumination in real-time, 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, relative fluorescence variation (dF/F0) images were obtained relative to a baseline fluorescence image, 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 rewardable, 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 Finish zone in at least the current and the previous frame. Each new rewardable wave opened immediately a 500 ms opportunity window for reward delivery. Each subsequent frame in which the wave continued to meet the rewardable 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 wave was labeled as "rewarded." The latency between image acquisition and opportunity window opening in the case of rewardable 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 rewardable simultaneously. If another wave was detected during the opportunity window, their opportunity windows would overlap.

Habituation Protocol and Initial Data Collection

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 "Finish" zones.

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 "Finish" 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, mice collected rewards via a lick-port, controlled by a signal sent to an electromagnetic valve when a lick was detected during the reward opportunity window. The lick signal was detected using a capacitive sensor. During the session, we recorded the times of licks, opportunity window opening times, reward delivery times, real-time detected wave trajectories, and calcium signals. Body movement was recorded synchronously. Only the first 10 minutes of session time were used for analysis.

Analysis of behavioral data

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

To analyze licking, we segmented the lick time series into bursts. We defined bursts as a series of licks that were taking place less than 300 ms from each other, as in our previous works (Goueytes et al. 2022, Abbasi et al. 2023).

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, the DiI marker was injected at a depth of 1 mm at the corners of the "Start" and "Finish" zones. Following the administration of an overdose of pentobarbital (150 mg/kg), mice were perfused transcardially with saline followed by paraformaldehyde (4% in 0.1 M phosphate buffer). After an 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), 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.