Neural Correlates of Sleep, and Game-Playing in Bonnet Macaques

A Thesis

submitted to
Indian Institute of Science Education and Research Pune
in partial fulfilment of the requirements for the
BS-MS Dual Degree Programme

by

Jason Joby

Indian Institute of Science Education and Research Pune
Dr. Homi Bhabha Road,
Pashan, Pune 411008, INDIA.

April, 2024

Supervisor: Prof SP Arun

Certificate

This is to certify that this dissertation entitled "Neural Correlates of Sleep, and Game-Playing in Bonnet Macaques" towards the partial fulfilment of the BS-MS dual degree programme at the Indian Institute of Science Education and Research, Pune represents study/work carried out by Jason Joby at the Indian Institute of Science, Bangalore, under the supervision of Prof SP Arun, Professor, Centre for Neuroscience, Indian Institute of Science, Bangalore during the academic year 2023-2024.

Prof SP Arun

Jason Joby

Committee:

Prof SP Arun

Dr Collins Assisi

This thesis is dedicated to Biscu, Lily, Maria, Ancy, and Joby

Declaration

I hereby declare that the matter embodied in the report entitled "Neural Correlates of Sleep, and Game-Playing in Bonnet Macaques" are the results of the work carried out by me at the Centre for Neuroscience, Indian Institute of Science, Bangalore, under the supervision of Prof SP Arun and the same has not been submitted elsewhere for any other degree. Wherever others contribute, every effort is made to indicate this clearly, with due reference to the literature and acknowledgement of collaborative research and discussions. The work presented in Chapter 2 of this thesis was done with equal contribution from Sudhanshu Bharadwaj, BS student at the Undergraduate Programme, IISc.

Jason Joby

Acknowledgements

All the work presented in this thesis is only thanks to the guidance of my supervisor, Prof SP Arun. I am truly grateful for having a mentor who constantly goes above and beyond and always ensures that I am on the right path. He has helped me make significant progress in my understanding of research and my outlook towards it, how to approach scientific problems, and quite importantly, how to work efficiently.

The work presented in Chapter 1 was not possible without the help of the entire VisionLab, IISc, including several alumni members. Their suggestions and feedback were crucial to several analyses in my project, and the entire lab setup at the primate facility was built from the hard work of many current and ex-lab members, whom I look up to very much. I would like to thank my VisionLab colleagues: Debdyuti Bhadra, Eshita Bansal, Georgitta Valiyamattam, Jhilik Das, Kedarnath P, Prithu Purkait, Shubhankar Saha, Sini Simon, Sudhanshu Bharadwaj, Surbhi Munda, Tahiti Manna, Vaishnavi Adella, Thomas Cherian, Georgin Jacob and Chandrakant Harjpal for being great friends and labmates throughout this journey. The sleep recordings protocol itself was designed and setup by Surbhi Munda, and the work towards the neural data and CCTV video alignment was done in collaboration with her. Surbhi, Sini, Shubho, Deb, Jhilik and Thomas have been the most caring mentors I could ask for.

The entire work presented in Chapter 2 of this thesis has had equal contribution from Sudhanshu Bharadwaj. Working with him has been an absolute pleasure, and it truly showed me how far working together can take a project.

I would also like to thank my family, for their constant support and encouragement. They have always been there for me, and I am truly indebted to them. Their role in supporting me through this journey simply cannot be put into words.

I would like to thank all my friends, who have stood by me during the ups and downs of this time. They are too many to name, but to name a few (in no particular ranking or ordering): Pooja and Shruthi have just always surpassed my expectations; Parijat, Anirban, Reetish and Soumyodeep have been a second family; Kashika, Krishna, Arjun and Surya have always made sure I was cared for; Tanishi, Aniket, Avi and Digant are so close to me that they have contributed to core tenets of my personality through the years I have known them.

Abstract

In visual neuroscience, the traditional system of recording neural activity (a highly constrained and artificial environment) often falls short due to its limited capacity to capture natural behaviours. One of these natural behaviours (and an extremely prominent one) is sleep. Sleep is an essential behavioural state that is known to be crucial for the proper functioning of the brain and the body. Despite its importance, the neural correlates of sleep are not well understood, especially in nonhuman primates. In this thesis, we investigate the neural correlates of sleep in bonnet macaques, a species of nonhuman primates. We recorded neural activity from the inferotemporal (IT) region, the ventro-lateral prefrontal cortex (vlPFC), and the ventral premotor area (PMv) of the brain, while the monkeys slept naturally. We find that the firing rates of the neurons in the PMv region are significantly lower during sleep compared to wakefulness, while the firing rates in the IT region have a marked increase in their variance. We also find that the power in the delta band of the local field potentials (LFPs) is significantly higher during sleep compared to wakefulness, in all the regions. These results suggest that there are clear neural correlates of sleep that can be distinguished from those during wakefulness, and that these have similarities to those observed in humans and rhesus macaques. In the second part of the thesis, we present our training of the macaques towards a new task paradigm - that of using four direction buttons on the touchscreen to move a character around the screen, with the goal of obtaining "reward" objects. Overall, this thesis provides new insights into the neural correlates of sleep in bonnet macaques as well as showcases a paradigm for the training for playing games, and highlights the importance of studying natural behaviours in nonhuman primates.

Contents

List of Figures

Chapter 1

Neural correlates of sleep

1.1 Introduction

It has become increasingly understood, that adequate sleep is essential to a healthy physiology and proper bodily function. This behavioural state appears to be essential to animal life across a variety of animal taxa, wherein inadequate sleep has been shown to negatively impact development, cognitive function, and life span [1]. Despite this consistency in the absolute requirement of a sleep behavioural state across taxa, when it comes to duration of sleep as well as the timing, a large variation is found within the species among different animal phyla. Even amongst humans, sleep durations are known to vary widely, ranging from less than 5 hours to greater than 10 hours [2]. This variation in sleep duration and timing is also seen in other animals, and is known to be influenced by a variety of factors: in mammals, parameters such as brain size, diet, social hierarchy and BMI have been identified to affect total sleep times [3]. Sleep disorders affect a significant portion of the human population, and are known to be associated with a variety of health problems, including obesity, diabetes, cardiovascular disease, and even mortality [1]. More than 80 sleep disorders have been currently identified by the American Academy of Sleep Medicine - the most common of which are insomnia, narcolepsy, sleep apnea and restless leg syndrome. Studies in India for different sleep disorders have shown a high prevalence of insomnia (25%) and sleep apnea (37%), among others. College students were found to have an increased prevalence of insomnia compared to the general population: 35% [4]. Deprivation of sleep is also known

to be a significant factor in motor vehicular accidents and disasters [1].

Although human subjects are used widely in the study of sleep and sleep disorders, a deep study of brain function during any behavioural state requires invasive methods of investigation. However, this is particularly true for sleep, with its lack of overt behavioural activity (almost, by definition). The use of animal models is thus essential to understanding the neural mechanisms underlying sleep.

Animal models commonly used for sleep research include rodents and felines (carnivores). However, these have many drawbacks due to their fragmented sleep patterns, and the differences in their sleep architecture compared to humans: of the light-dark cycle, rodents sleep during the light phase, and carnivores are awake mainly at the transitions between the cycle's phases. Nonhuman primates, on the other hand, have sleep patterns that are more similar to humans and thus are better suited for sleep research, and several studies over the years have reported successful sleep recordings in large nonhuman primates [5]. The close phylogenetic relationship with humans also makes them almost a perfect model for filling the gap between rodent/feline and human sleep research. The macaques, and the rhesus monkey, in particular, have been used as a model for human sleep research, due to their similarities in sleep architecture, sleep stages, and sleep patterns. Importantly, the rhesus macaque has been the most used NHP model for the study of other diseases and disorders, including those of the nervous system, and thus has a wealth of information and characterization available.

However, for a long time, the use of nonhuman primates in sleep research was limited by the lack of suitable recording techniques, and the need for the animals to be restrained physically during the recordings or shifted to completely new and artificial environments. This has been a significant limitation, as it is known that the sleep architecture of nonhuman primates is affected heavily by the restraint [6], and thus the sleep patterns are markedly different from those during their natural state - the sleep patterns are fragmented and have prolonged bouts of wakefulness. Other than the physical recording setting, the recording techniques themselves have also evolved over the years. The use of fully implantable biotelemetry transmitters has allowed for continuous recordings over multiple days and has been a significant advancement in the field. EEG (electroencephalography) recordings have been the most common method of studying sleep [7], across all models, including nonhuman primates. However, recent advances in wireless biotelemetry devices, especially for electrocorticography (ECoG) and local field potential (LFP) recordings, have allowed for the study

of natural sleep at a much higher spatial resolution than what is possible with EEG.

1.2 Methodology

1.2.1 Neural recording setup

The experiments were performed on two male bonnet macaques (*Macaca radiata*), who were housed in a naturalistic environment[8]. Broadband neural signals were recorded at a 25kHz sampling rate by means of a wireless logger, which records from a total of 256 Pt/Id electrodes (Floating Microelectrode Arrays). In Monkey 1 (M1), 128 electrodes were placed in the inferotemporal (IT) region, 64 electrodes were placed in the ventro-lateral prefrontal cortex (vlPFC) and 64 in the ventral premotor area (PMv). In Monkey 2 (M2), 128 electrodes were placed in the IT region and 128 electrodes were placed in the PMv region.

1.2.2 Sleep recordings

The monkey was housed in an area of the naturalistic housing environment which has good visibility through the CCTV camera system. Neural data collection was wirelessly triggered and continued for a duration of approximately 2 hours, starting at approximately the time when the monkey naturally starts falling asleep (as noted via the CCTV). The start of the recording was accompanied by a series of flashes in a separate room (not visible to the monkey), which were used to synchronize the neural data with the CCTV video recording, post-hoc.

1.2.3 Data analysis

Data preprocessing

Neural data: The wideband data (recorded at a sampling rate of 25kHz) was low-passed to 500Hz and then downsampled to 1kHz, to generate LFP (Local field potential) data. Multi-unit activity (MUA) data was generated offline using the Plexon Offline Sorter™, by

first filtering with a low-cut Butterworth filter (4-pole) with a 250Hz cutoff. Waveform/spike sorting was then done following a (negative) amplitude threshold of 3.75 sigmas from the signal mean for each channel.

Video data: The video data was manually annotated from CCTV recordings of the monkey during a sleep recording session: "bouts" of sleep and wakefulness were marked based on visible behavioural markers such as a huddled position while sitting on the floor/log, with the head bent down. Wakefulness was marked by open eyes, raised head, movements around the room, scratching the self etc. There were no visible artificial or natural lights during this period, but the CCTV system was equipped with an infrared light source for night-time recordings (imperceivable by monkeys or human eyes). From these behavioural markers, we generate proxies of the sleep and wakefulness states, which we term as *HeadDown* and *HeadUp* bouts, respectively. It is to be noted that, although a *HeadUp* bout is necessarily a wakeful state, a *HeadDown* bout is not necessarily a sleep state, and could also be a relaxed wakeful state. It was ensured during annotation that at any point in time, the monkey was either in a *HeadDown* or *HeadUp* state, based on the behavioural markers.

Alignment: The neural data was aligned with the video data by using repeated flashes on a screen in a separate room. The flashes were visible in both the CCTV camera system, as well as on a separate camera system (WhiteMatter™ e3Vision) which is time-aligned with the neural data acquisition system (WhiteMatter™ eCube Server). By manual selection of the pixel area of the screen in both video systems, an intensity over time signal was generated for both, which was then aligned by thresholding for the intensity jump at the time of the first flash. This alignment (time difference) was then used to align the neural data with the video data.

Artifacts: The neural data was checked for artifacts, and the data was cleaned manually by assigning the time points of the "blips" as artifact time points. Later, based on the needs of the particular analysis, the data at those time points may be assigned NaN values. The characteristic of the artifacts/blips was that they were high amplitude square waves, which occur across all channels simultaneously every 2686 seconds, or 44.7mins, and last for exactly 1310ms, which makes for easy manual demarcation.

Bout lengths and counts

The lengths of the *HeadUp* bouts were calculated from the video data by finding the time difference between a *MoveHeadUp* event to the next *MoveHeadDown* event and vice versa for a *HeadDown* bout. Subsequently, a histogram of these bout lengths was made using a 20-second bin width, and both monkeys' histograms were plotted together. This was done separately, for both *HeadUp* and *HeadDown* bouts. (1.2.3 Data preprocessing – Video data). The fraction of total time spent in all the bouts less than 20 seconds was calculated as the sum of the durations of all the bouts less than 20 seconds, divided by the total duration of all the bouts.

1.2.4 Firing rates analysis

The firing rate was calculated from the MUA data per channel by binning the spike times into 1s bins (unless mentioned otherwise). The firing rate was then calculated as the number of spikes in a bin, divided by the bin width (1s). If any timebin contains an artifact timepoint, as defined in (1.2.3 Data preprocessing – Artifacts), the firing rate for that timebin was assigned NaN. Unless mentioned otherwise, the firing rates were z-scored (mean-subtracted and divided by the standard deviation) across all timebins for each channel, to make the firing rates comparable across channels.

Event-aligned firing rates

To visualize the firing rates of the channels around the time of the transition from *HeadDown* to *HeadUp* and vice versa, we align the firing rates of the channels to the time of the transition. Hence, time 0 in the plots was the time of the respective transition event. These transition events were labelled *MoveHeadUp* and *MoveHeadDown* events, and were defined as the time of the first frame of the video in which the monkey was seen to be transitioning from *HeadDown* to *HeadUp* and vice versa, based on the behavioural markers in (1.2.3 Data preprocessing – Video data).

To visualize the raw firing rates data, for each bout, we place colourmaps wherein each pixel's colour value represents the firing rate of a particular channel (along the rows), at

a particular timebin with respect to the transition event (along the columns). The bouts preceding and succeeding a particular transition event were plotted together with the same (artificial) y-displacement to aid visualization. For this particular analysis, a bin size of 0.5sec was used while computing firing rates, for a cleaner visualization. The events were sorted on the y-axis based on the duration of the bout just before the transition event.

We then calculate the average firing rates across the channels per region, average them across all these transition events and plot them to visualize the change in firing rates per region, around the time of the transition. Since while averaging across the bouts, the time duration of each of the bouts was not the same, the standard error of the mean (SEM) for each timebin was also calculated and plotted as a background shadow with lower opacity (each timebin's average across bouts may have a different number of contributing bouts - for example, as M1 only has 3 *HeadUp* bouts with a duration >60seconds, any timebin in the *MoveHeadDown*-aligned plots aligned at less than -60sec would not have more than 3 bouts contributing to the calculated average).

Firing rates comparison

To compare the average firing rates across the *HeadDown* and *HeadUp* states, there are two methods: average the firing rates over all the timebins which occur during all *HeadDown* or *HeadUp* bouts, or, average the firing rates over all the timebins which occur during each *HeadDown* or *HeadUp* bout, and then average these averages across all the bouts. The second method, however, assigns an equal weightage to all the bouts, which may be problematic as the bouts can be of very different lengths. The first method was visualized in the results section, via the plotting of the distributions of the average firing rates of the regions during the *HeadDown* and *HeadUp* states. Bin sizes for the distributions were 0.05. Dashed vertical lines present the averages from *HeadUp* and *HeadDown*. To perform statistical significance tests, a custom MATLAB function written at the VisionLab was used (`statcomparemean`), which performs the following process. An Anderson-Darling Normality Test was performed using MATLAB's `adtest` function to check whether the underlying distributions were normally distributed. If yes, a Two Sample t-Test (`ttest2`) was performed to check whether the means were significantly different. Otherwise, a Wilcoxon rank-sum test (`ranksum`) was performed, which tests the hypothesis that the "medians are equal". The p-value from the resulting test was compared to a critical value of 0.05 to compute

significance. (1.2.3 Data preprocessing – Neural data)

Correlation of firing rates with states

To check how the firing rates of the channels correlate with the *HeadDown* and *HeadUp* states, we construct a vector of 1s and 0s, wherein those timebins which were during *HeadUp* bouts are 1s and the 0s correspond to the timebins during *HeadDown* bouts. We then calculate the Pearson's linear correlation coefficient of the vector of firing rates of each channel with this vector.

Linear discriminant analysis-based classifier

We employed a linear discriminant analysis-based classifier (MATLAB's `classify`) to distinguish between *HeadUp* and *HeadDown* bouts based on the firing rates of all the channels during a given timebin. The classifier was trained using labelled firing rates data, where timebins during *HeadUp* bouts were assigned a label of 1 and *HeadDown* bouts a label of 0. We randomly selected an equal number of points from both *HeadUp* and *HeadDown* bouts for training. The number of points was calculated as the 1/3rd of the smaller set out of *HeadUp* and *HeadDown*, which in the case of both monkeys was the *HeadUp* set. We then evaluated the accuracy of the classifier on the held-out data to gauge its effectiveness in correctly classifying *HeadUp* and *HeadDown* timebins. To mitigate potential sampling bias, we performed 5-fold cross-validation, wherein we were able to verify that the accuracy of the classifier remained within a 3% range across the folds.

Cross-channel correlations and Principal component analysis

We calculated the cross-channel correlation matrix for the firing rates data during the *HeadDown* and *HeadUp* bouts. This revealed a high correlation between the channels of each region, which could lead to multicollinearity in the training data, and thus affect the performance of the classifier. To mitigate this, we performed a principal component analysis (PCA) on the MUA data, to reduce the dimensionality of the data and remove multicollinearity. We retained the principal components that explained 90% of the variance and used these

components for the classifier instead. All the remaining steps remain the same as that in (1.2.4). All subsequent analysis of the classifier was done using the model trained with the PCA-reduced data.

Classifier performance evaluation

To visualize the classifier's predictions, we plot the classifier's predicted probability values (for one class) for every timebin in the recording, and compare it with the actual labels, by underlaying a colour-coded background. We then create a confusion matrix, with the actual labels on one axis and the predicted labels on the other, to visualize the classifier's performance across the two classes - the percentages calculated were based on the total actual labels of each class. We then also average the classifier's probability values across all bouts, aligned to the transition events, to visualize the classifier's performance across changes in the state.

1.2.5 Local field potential analysis

Local field potential data was generated as specified in (1.2.3 Data preprocessing – Neural data). All analysis of the LFP presented in this study involved spectral analysis. This was done by using either MATLAB's `pwelch` function, or the Chronux [9, 10] package's `mtspecgramc` function: this will be explicitly mentioned in the individual sections.

Event-aligned spectrograms

The LFP data was first z-scored using MATLAB's `zscore` function, and to obtain clearer colourmaps, outliers (defined as >3 after z-scoring - meaning greater than 3 σ from the global mean in the original data) and artifacts (manually defined timepoints, (1.2.3 Data preprocessing – Artifacts)) were set to NaNs, before the spectral analysis. The overall layout follows the same philosophy as 1.2.4, wherein bouts surrounding a transition event were plotted together, with the same y-displacement. The spectrogram of the LFP's average across each region was then shown as a colourmap, with different pixel values showing the different powers of the frequencies. The spectrogram was made using the `mtspecgramc`

function, which uses a multi-taper spectrogram method, with a time-bandwidth product of 3, and 6 tapers. The frequency range for the spectrogram was 0-50Hz, and the time window used was 1s, with a 50% overlap in successive windows. Once the power values and the frequencies were obtained, the log of the power was plotted, against the log of the frequencies, as the final colourmap.

Power spectral densities

The power spectral densities (PSDs) of the LFP data (mean-subtracted) were calculated using the `mtspecgramc` function, by first computing the spectrogram over the entire duration of the recording, following which the log of the powers for the timebins which were during *HeadUp* and *HeadDown* bouts were averaged separately. The PSDs were then plotted for the *HeadUp* and *HeadDown* states, for each region. The frequency range for the PSDs was 0-50Hz, and the other parameters were also the same as that in (1.2.5).

Powers in canonical bands

Although there's plenty of variation across literature in the exact frequency ranges which define each canonical frequency band, for this study, they were defined corresponding to delta (1–4 Hz), theta (4–7 Hz), alpha (7–15 Hz), beta (15–31 Hz), and gamma (>31 Hz) bands, following suit from Xu et al., 2019 [11]. From the spectrogram, the average of the log of the powers of the frequencies within each of these bands was calculated, and then plotted against time, for each region. The backgrounds were coloured appropriately to visualize the change in the bands' powers with the change in state. Following this, the average powers across all *HeadDown* time points versus all *HeadUp* time points was computed, for each band separately.

1.3 Results

What is the behavioural pattern of sleep and wakefulness, around the time of onset of sleep?

As our sleep recordings were performed around the time of day when the monkeys start to go to sleep, this onset-of-sleep period is expectedly marked with multiple bouts of wakefulness and sleep interspersed with each other. This is visualized in Figure 1.1, which shows an overview of both monkeys' recordings from a behavioural standpoint. A majority of the time in both recordings is spent in the *HeadDown* state, but interspersed (but shorter) bouts of wakefulness/ *HeadUp* were observed. These *HeadUp* bouts occur throughout the recording and also imply that the maximum duration of a *HeadDown* bout that we can capture during a single recording is not very long.

The distribution of the durations of the *HeadDown* and *HeadUp* bouts, for both monkeys, are shown in Figure 1.2 A and Figure 1.2 B, respectively. A significant fraction of these bouts are of very short duration, in both monkeys and especially for *HeadUp*: 71% and 67% of the *HeadUp* bouts have a duration below 20 seconds, in M1 and M2 respectively. In *HeadDown* bouts, 28% and 32% are less than 20 seconds, for M1 and M2 respectively. The fraction of total time spent in all these bouts less than 20 seconds is only 2% in both M1 and M2.

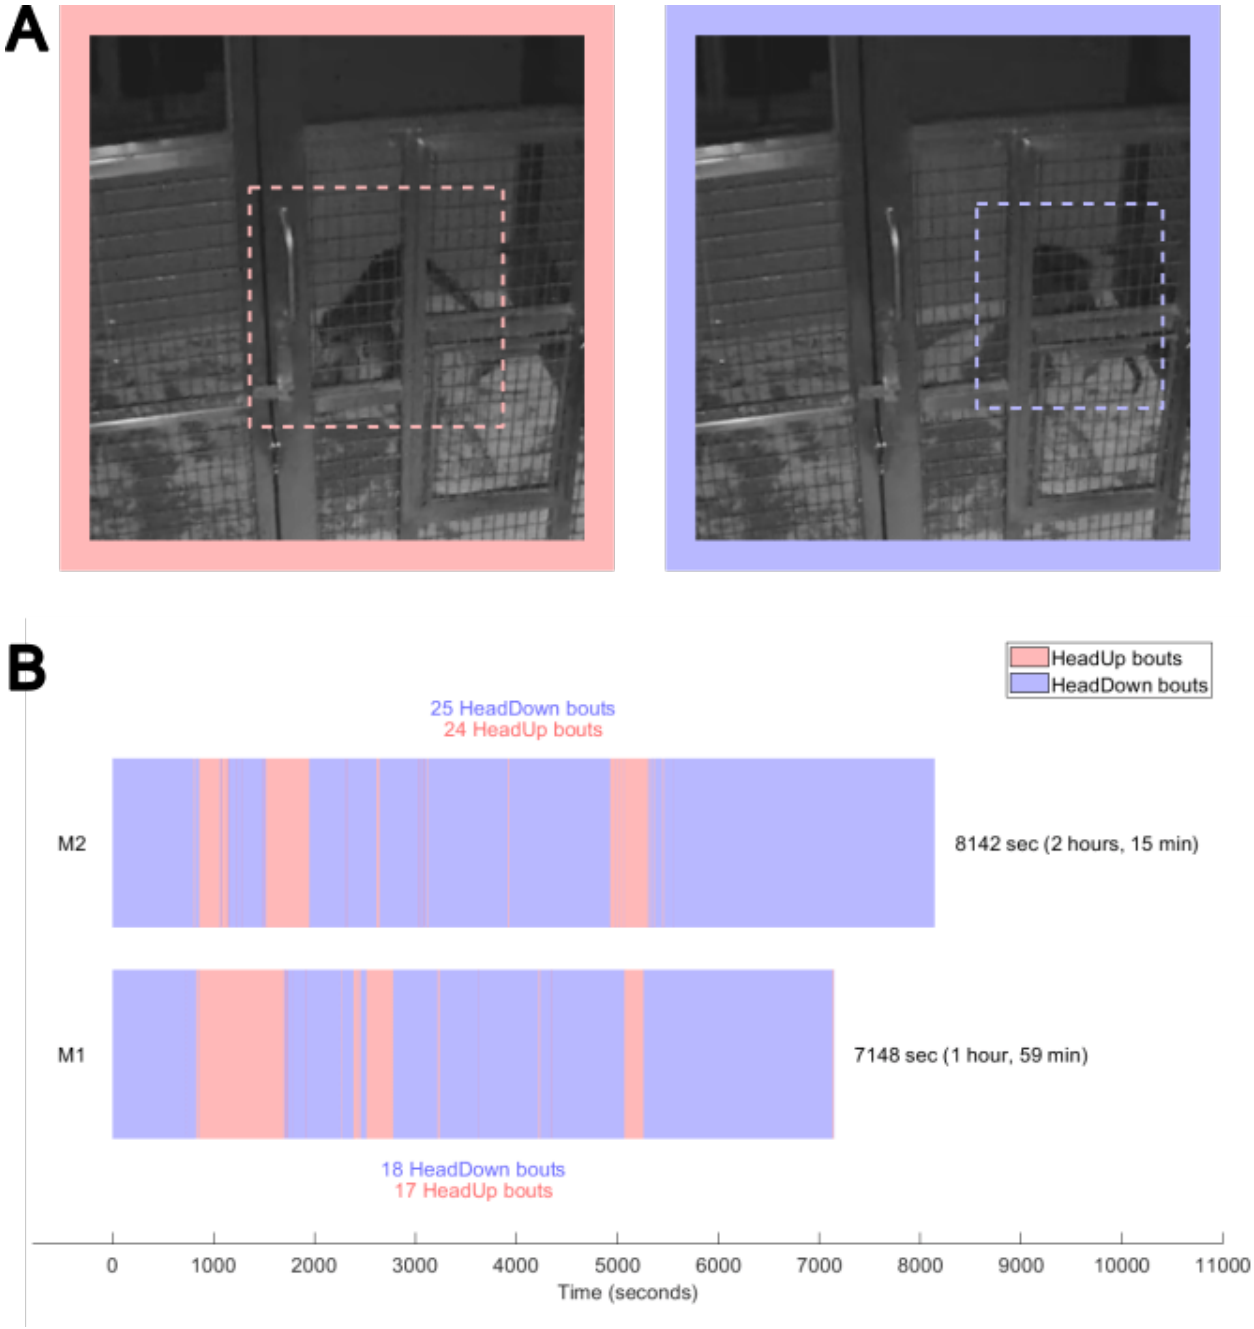

Figure 1.1: **Panel A** *Left* - Example CCTV video frame (cropped) of M2 during a *HeadUp* bout (active and moving around), *Right* - during a *HeadDown* bout (in a huddled up position with the head bent down). **Panel B** - Sleep recordings overview for M1 and M2. The x-axis is the time in seconds, while the coloured rectangles represent the *HeadDown* bouts (blue) and *HeadUp* bouts (pink), placed according to when these bouts occur during the recording.

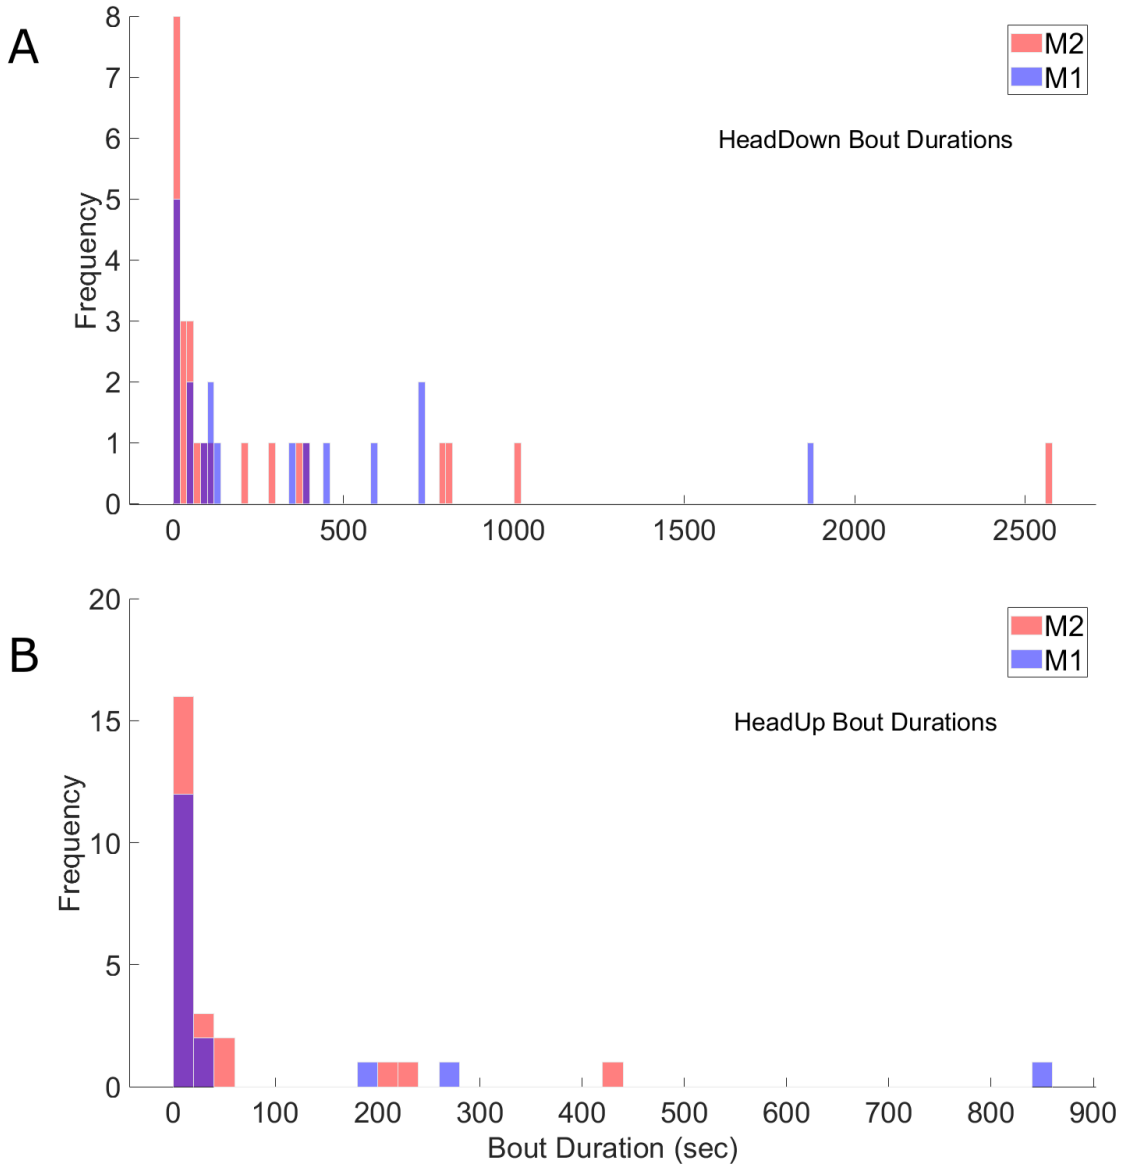

Figure 1.2: **Panel A** - Histogram of *HeadDown* bout durations for M1 and M2. **Panel B** - Histogram of *HeadUp* bout durations for M1 and M2. The x-axis is the duration of the bout in seconds, and the y-axis is the number of bouts of that duration. Bin sizes are of 20 seconds. Overlapping histogram areas are coloured in purple.

How do the firing rates change as the monkey changes between HeadUp and HeadDown states?

Every switch from a *HeadUp* bout to a *HeadDown* bout is marked as a *MoveHeadDown* event, and vice-versa for *MoveHeadUp* events. These events are the crucial (and only) behavioural markers of a possible state change. Event-aligned firing rate plots are used as a tool to visualize exactly how the different regions are with respect to their firing rates, before and after these two types of transitions: from *HeadUp* to *HeadDown* and from *HeadDown* to *HeadUp*. Figure 1.3 shows the event-aligned firing rates of the channels in the IT region around the time of the transition from *HeadUp* to *HeadDown*, which are *MoveHeadDown* events. The same plot is made for the PMv and vlPFC regions, and also for *MoveHeadUp* events instead. All other plots are available in the supplementary figures section.

Visually, it is easy to observe that in the IT region, the firing rates are more variable after a *MoveHeadDown* event, that is, during *HeadDown* bouts. For the other regions, it is not easy to visually discern a pattern, and thus we move on to the next section to quantify the differences in the firing rates during the two states.

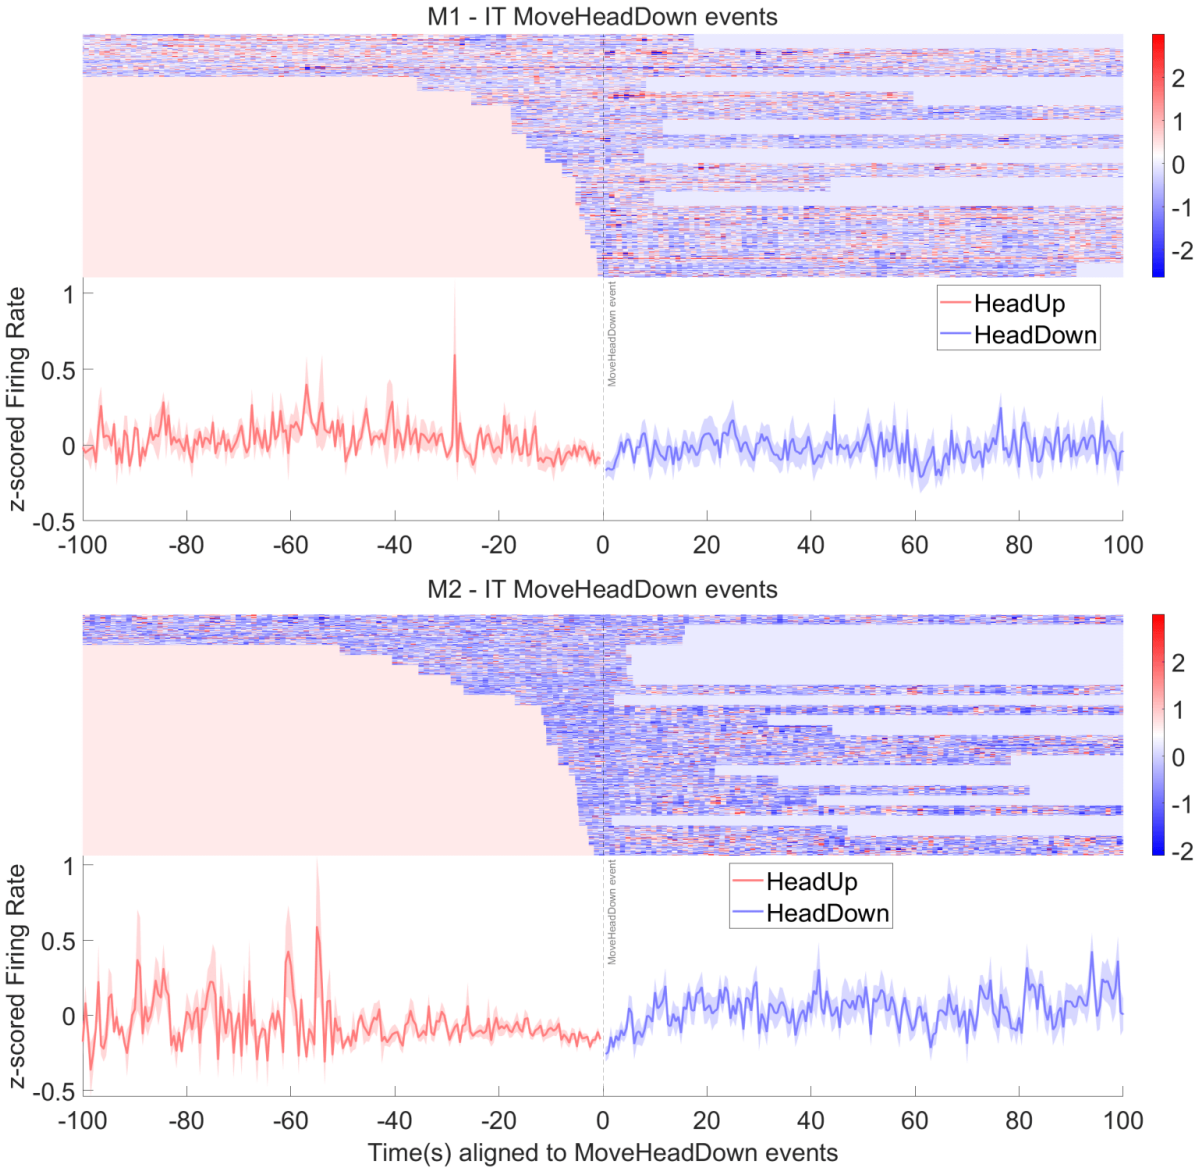

Figure 1.3: Event-aligned firing rates of the channels in IT region around the time of the transition from *HeadUp* (red) to *HeadDown* (blue), which are *MoveHeadDown* events. The x-axis is the time in seconds relative to the transition event, and the y-axis is the z-scored firing rate of the region-average, but subsequent events are displaced on the y-axis such that the earliest event is at the bottom. The dashed vertical line is the time of the transition event. Top is M1, bottom is M2.

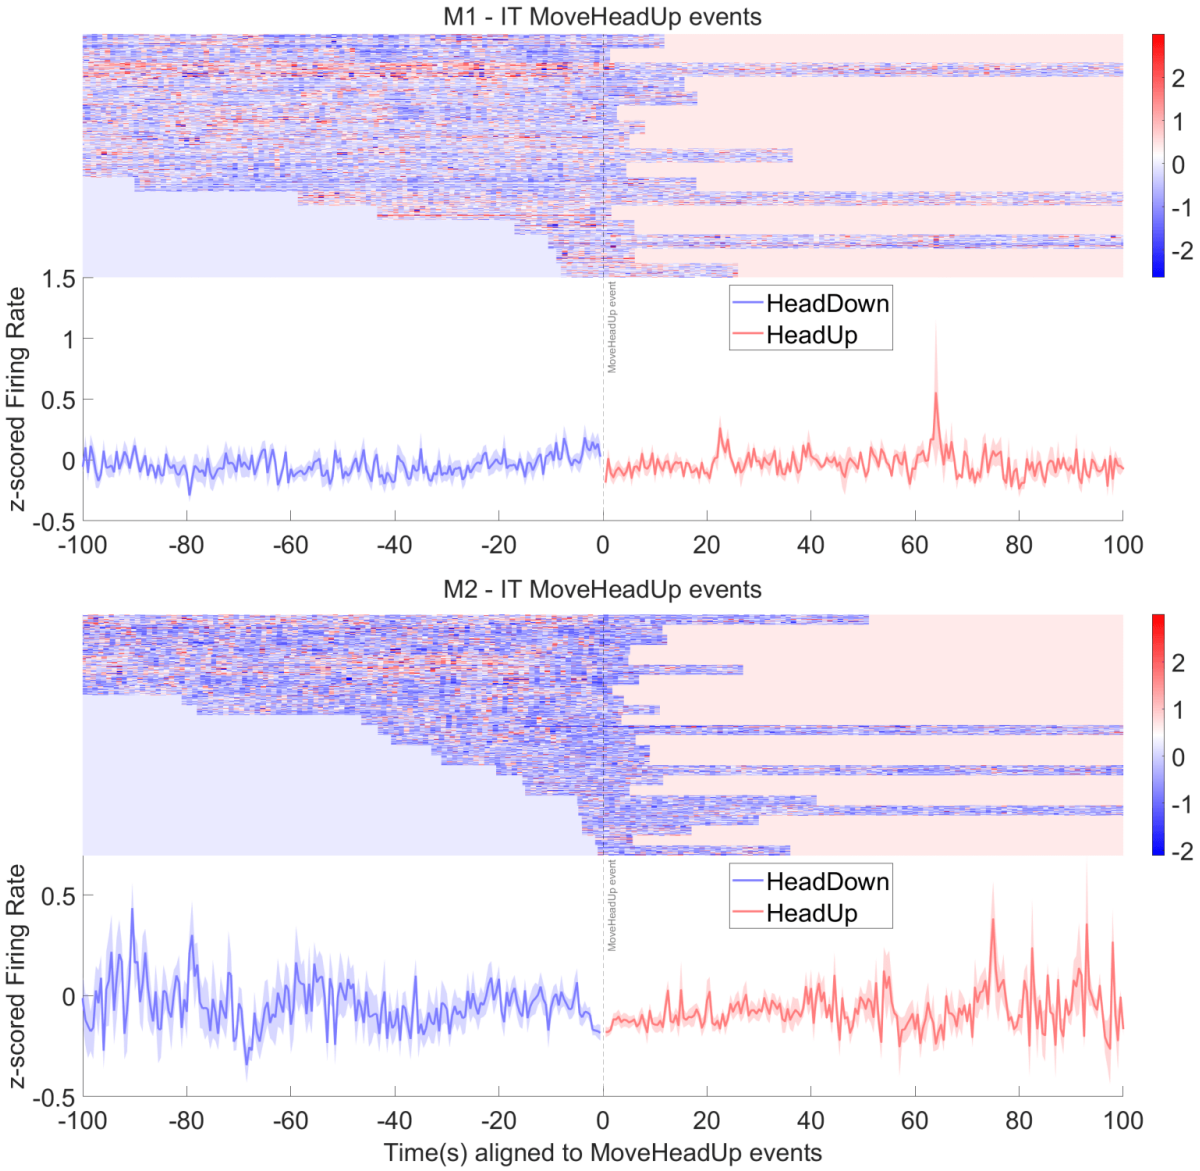

Figure 1.4: Event-aligned firing rates of the channels in IT region around the time of the transition from *HeadDown* (blue) to *HeadUp* (red), which are *MoveHeadUp* events. The x-axis is the time in seconds relative to the transition event, and the y-axis is the z-scored firing rate of the region-average, but subsequent events are displaced on the y-axis such that the earliest event is at the bottom. The dashed vertical line is the time of the transition event. Top is M1, bottom is M2.

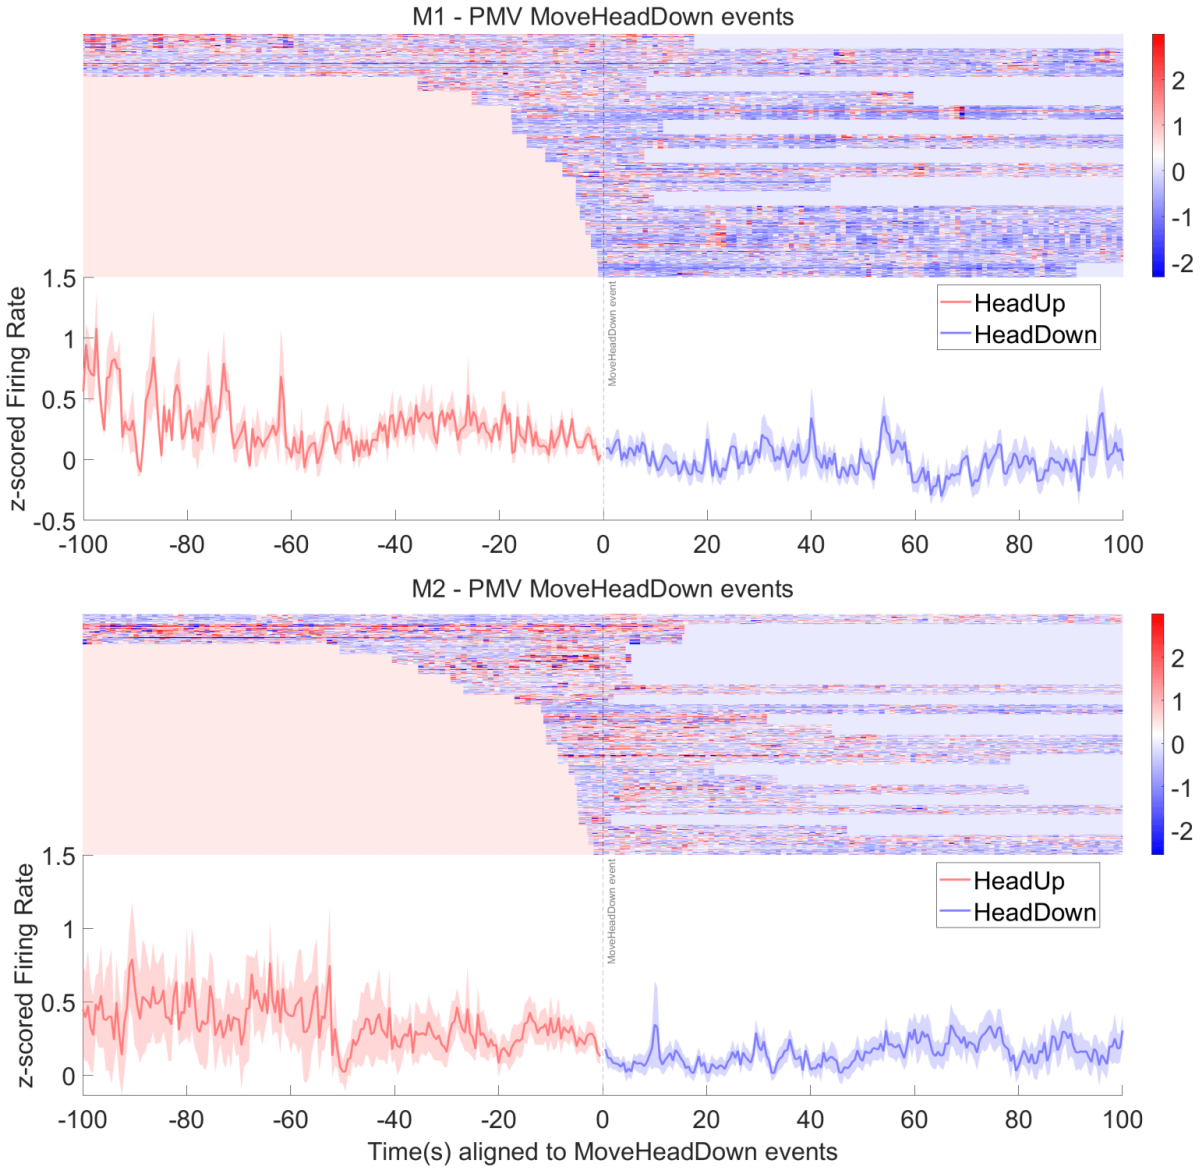

Figure 1.5: Event-aligned firing rates of the channels in PMv region around the time of the transition from *HeadUp* (red) to *HeadDown* (blue), which are *MoveHeadDown* events. The x-axis is the time in seconds relative to the transition event, and the y-axis is the z-scored firing rate of the region-average, but subsequent events are displaced on the y-axis such that the earliest event is at the bottom. The dashed vertical line is the time of the transition event. Top is M1, bottom is M2.

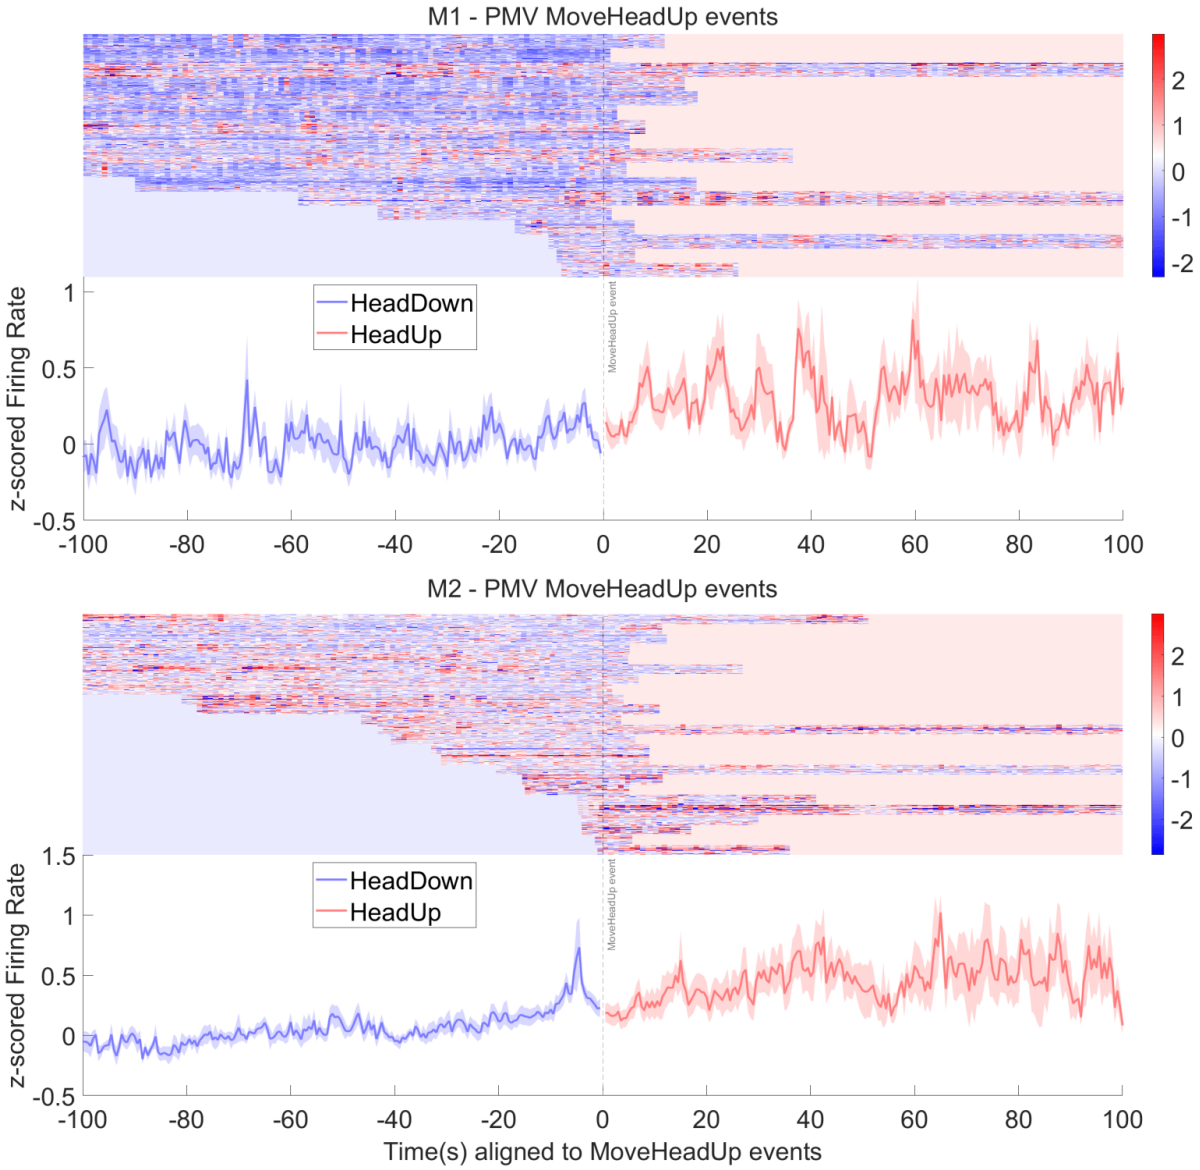

Figure 1.6: Event-aligned firing rates of the channels in PMv region around the time of the transition from *HeadDown* (blue) to *HeadUp* (red), which are *MoveHeadUp* events. The x-axis is the time in seconds relative to the transition event, and the y-axis is the z-scored firing rate of the region-average, but subsequent events are displaced on the y-axis such that the earliest event is at the bottom. The dashed vertical line is the time of the transition event. Top is M1, bottom is M2.

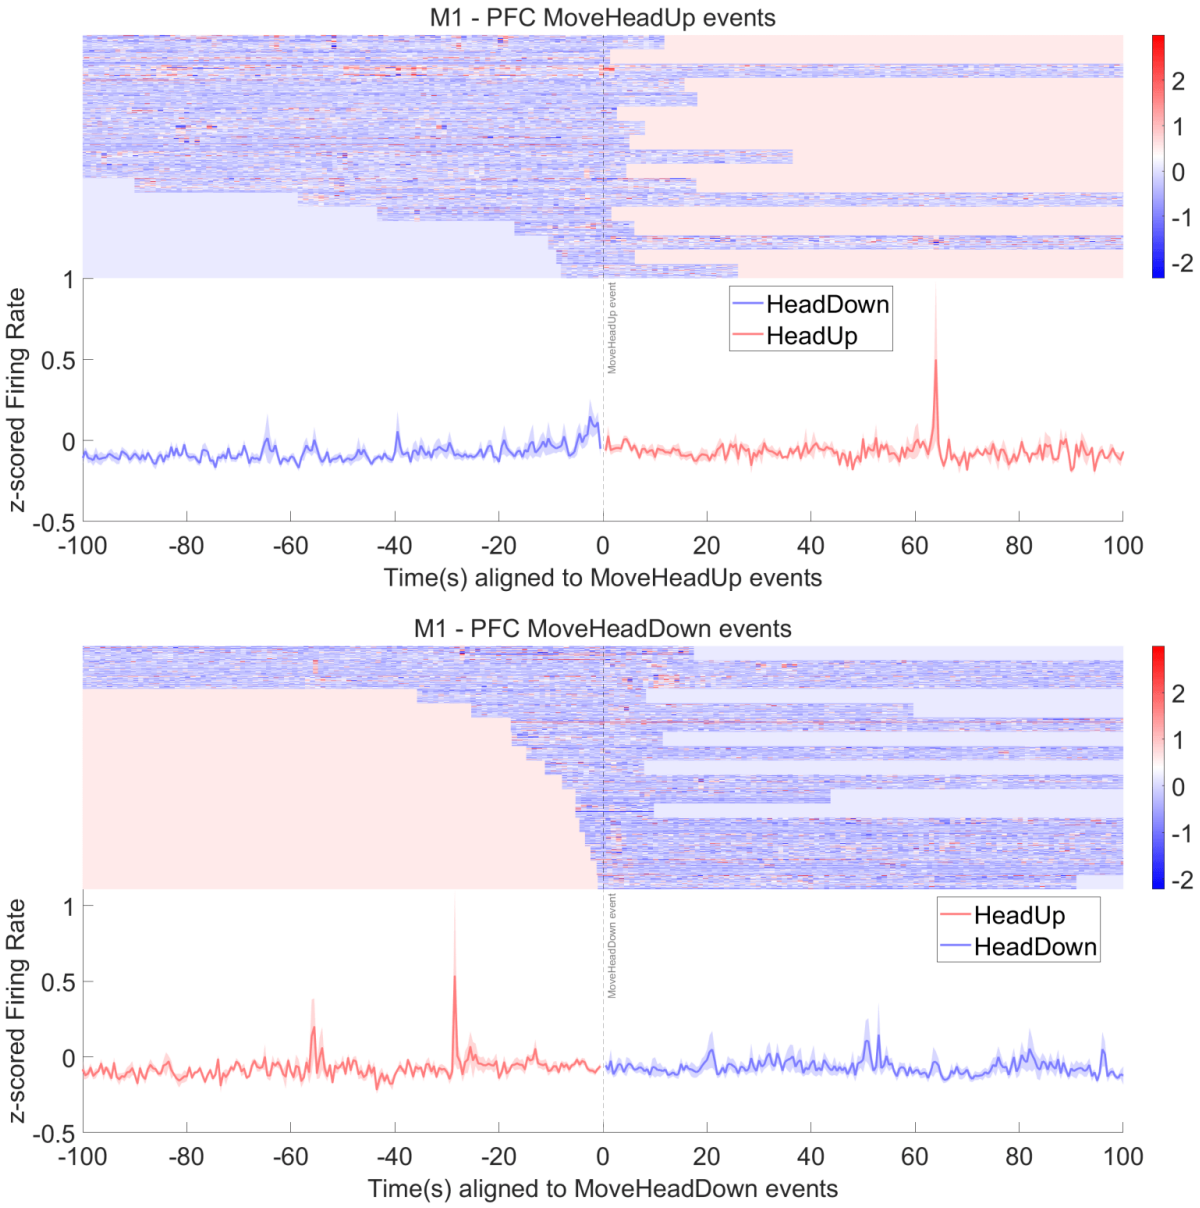

Figure 1.7: Event-aligned firing rates of the channels in vlPFC region for both *MoveHeadUp* (top) and *MoveHeadDown* (bottom) events. vlPFC region recordings are only available from M1. The x-axis is the time in seconds relative to the transition event, and the y-axis is the z-scored firing rate of the region-average, but subsequent events are displaced on the y-axis such that the earliest event is at the bottom. The dashed vertical line is the time of the transition event.

How do the firing rates during the two states compare, globally?

In the previous subsection, we see the firing rates of the regions, just around the time of the different transitions. In this section, we compare the average firing rates of the channels during the *HeadDown* and *HeadUp* states, across the whole recording. The distributions of (z-scored) firing rates of the channels during the *HeadDown* and *HeadUp* states are shown in Figure 1.8. The average firing rates (dashed lines in corresponding colours) of the channels during the *HeadDown* and *HeadUp* states are significantly different in all regions, as calculated using the Wilcoxon rank-sum test (1.2.4 Firing rates comparison – paragraph 2), but the difference in the averages themselves is more pronounced in the PMv region than in the IT and vlPFC regions. There's an overlap in the distributions of the firing rates during the two states across all the regions, but the way the shape of the distributions changes across the different states is consistent across both M1 and M2.

For the PMv regions, the *HeadDown* state has a lower average firing rate than the *HeadUp* state for both monkeys.
Means: M1 : -0.08 and 0.35, M2 : -0.06 and 0.38 for *HeadDown* and *HeadUp* respectively.

For the IT regions, the distribution of the firing rates during the *HeadDown* state is more spread out than that during the *HeadUp* state (has a larger variance during *HeadDown*). Variances: M1 : 0.18 and 0.03, M2 : 0.19 and 0.08 for *HeadDown* and *HeadUp* respectively. The mean of the distribution is shifted negatively for *HeadDown*, for M1, but positively for M2.

For the vlPFC region, the distribution of the firing rates is right-skewed during the *HeadDown* state compared to during the *HeadUp* state (and a positively shifted mean). Means: M1 : 0.02 and -0.07, for *HeadDown* and *HeadUp* respectively.

This section helps us concretely summarize the differences between the firing rates during the two states, which may have been difficult to make out visually from the analysis done in the last section.

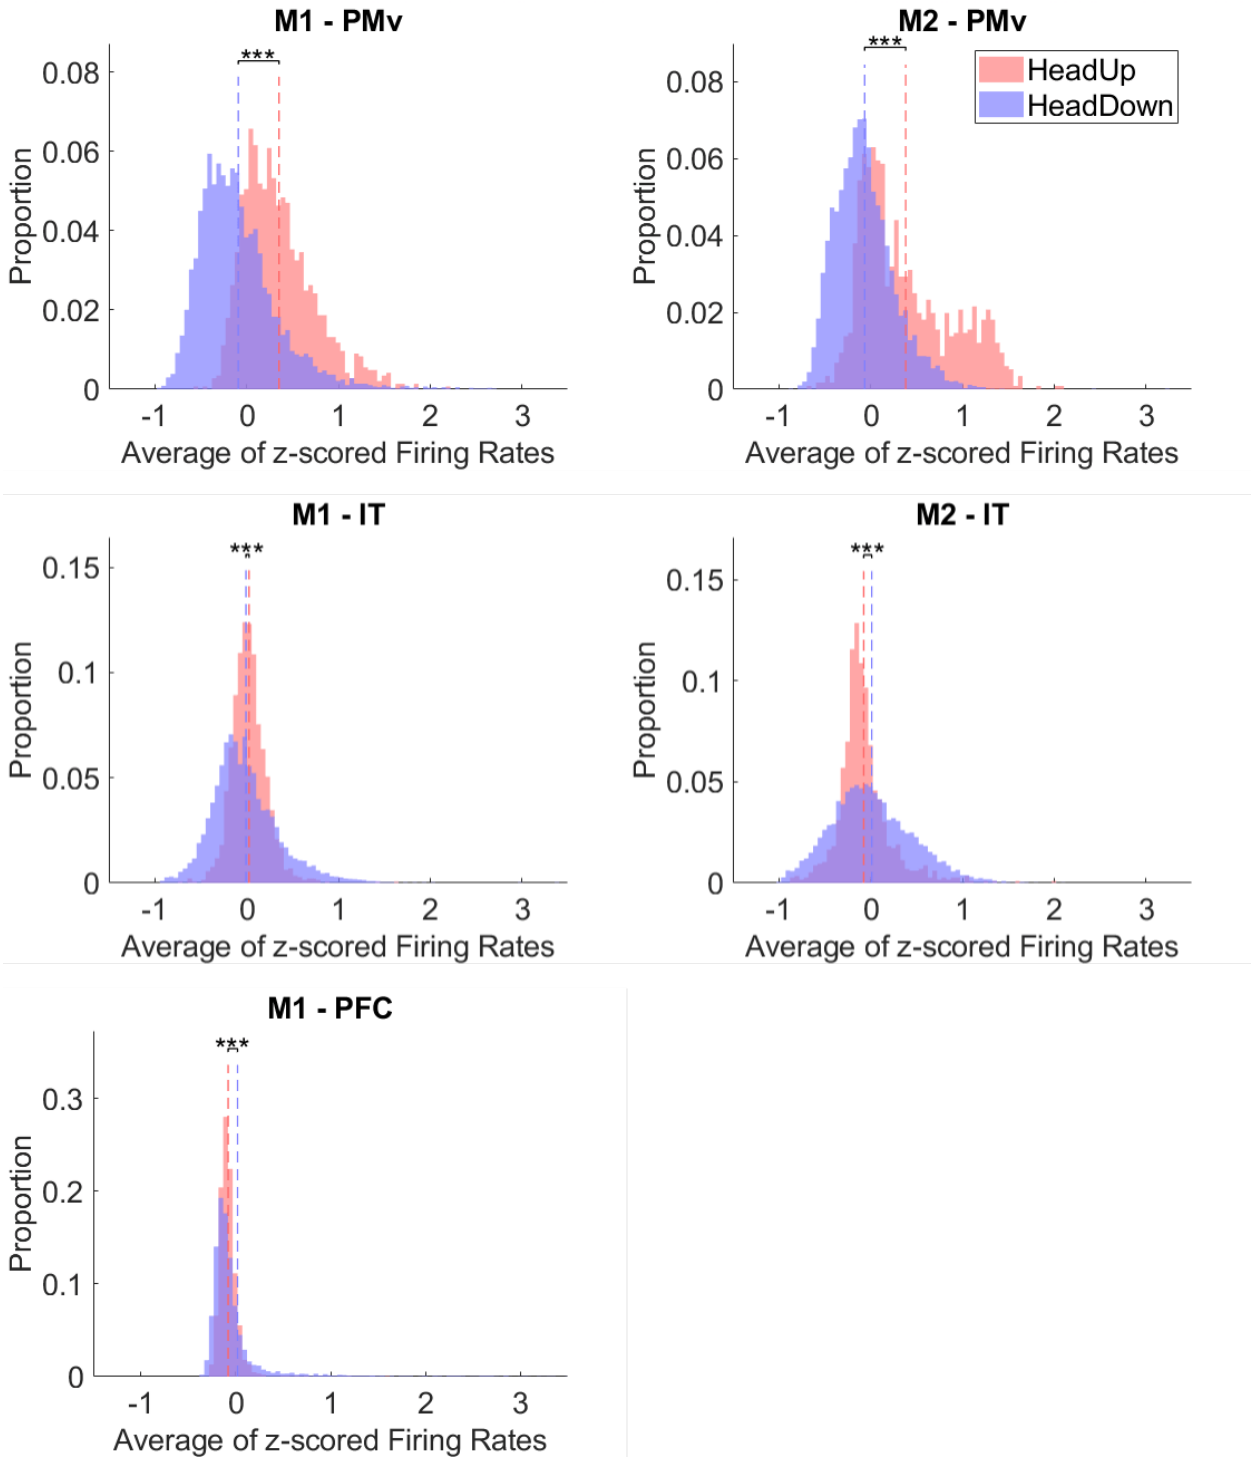

Figure 1.8: Firing rate distributions (normalized to probability/proportion) of the regions' average firing rates (z-scored) during the *HeadDown* and *HeadUp* states. The y-axis hence represents the number of timebins (1 second) that have that particular binned firing rate. The dashed vertical lines are the means of the corresponding distributions. The p-values obtained from the Wilcoxon rank-sum test of the *HeadUp* and *HeadDown* rates are shown as significance values:
*** corresponds to $p \leqslant 0.001$.

Are the firing rates of different channels modulated differently by the two states?

The last section provides an analysis that averages over all the channels in a region. This section provides a channel-wise look at the correlations between how the firing rates of each channel change as the states themselves change. A positive correlation here implies that as the state changes from 0 to 1 (*HeadDown* to *HeadUp*), the firing rates also increase. On the other hand, a negative correlation implies a reduction in firing rates as the state changes from *HeadDown* to *HeadUp*.

The correlation coefficients for each channel are shown in Figure 1.9, colour-coded for positive and negative values of the correlation. In the subsequent lines, we report the percentage of channels in each region that have positive correlations For both monkeys, the PMv region has a strong bias towards positive correlations (M1: 87.5%, M2: 93.7% positives). The IT region has a mix of positive and negative correlations (M1: 48.4%, M2: 35.9% positives). The vlPFC region has a strong bias towards negative correlations in M1 (23.4% positives).

Can we train a classifier to predict the state of the monkey, based on the firing rates?

To test whether the firing rates of the channels can be used to predict the state of the monkey, we trained a linear discriminant analysis-based classifier to distinguish between *HeadDown* and *HeadUp* bouts, based on PCA-reduced forms of the firing rates per channel. The PCA was done for dimension reduction, and the number of principal components retained was such that 90% of the variance was explained.

From this model, a probability for each of the classes is predicted for every timebin in the whole recording, and the probability for the *HeadUp* class (score) is plotted against time in Figure 1.10 (the probability of the other class, *HeadDown*, is always equal to one minus the probability of the *HeadUp* class). The background of the plot is coloured such that the time in which the monkey is in a *HeadUp* bout or a *HeadDown* bout is coloured differently from each other. We see a high matching of the predicted score, with the states. This visualization, actually allows us to see the evolution of the prediction within the bouts - for example, in M1, there are certain *HeadDown* bouts that seem to have a slowly increasing

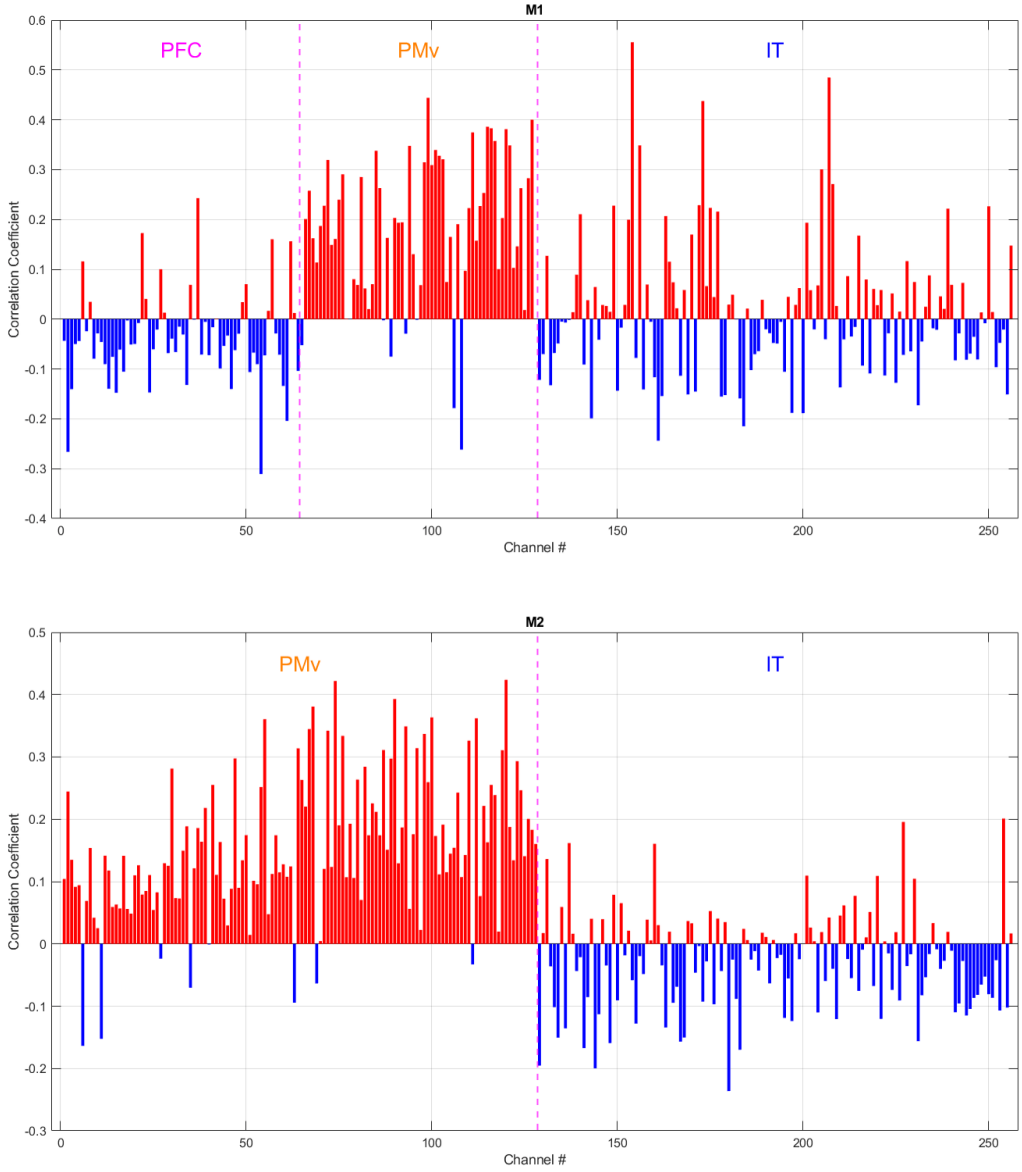

Figure 1.9: Correlation of the firing rates of the channels (over the entire recording) with the *HeadDown* and *HeadUp* states. The colour of the bar represents the sign of the correlation - red for positive correlation and blue for negative.

probability from near 0 towards 1, which is then followed by a waking up event and a *HeadUp* bout. On the other hand, for M2, although the classifier's scores match up well with the actual labels, no particular within-bout evolution can be observed, visually.

We calculate the confusion matrix based on the predicted classes for each timepoint (defined as the class with the higher probability score, in the classifier) (Figure 1.11 B), and also the accuracy by finding predicted labels, and comparing them with the actual labels. The accuracy of the classifier is 92.7% for M1 and 87.1% for M2. This tells us that that we are indeed able to predict the state of the monkey using a linear model on the firing rates of the channels.

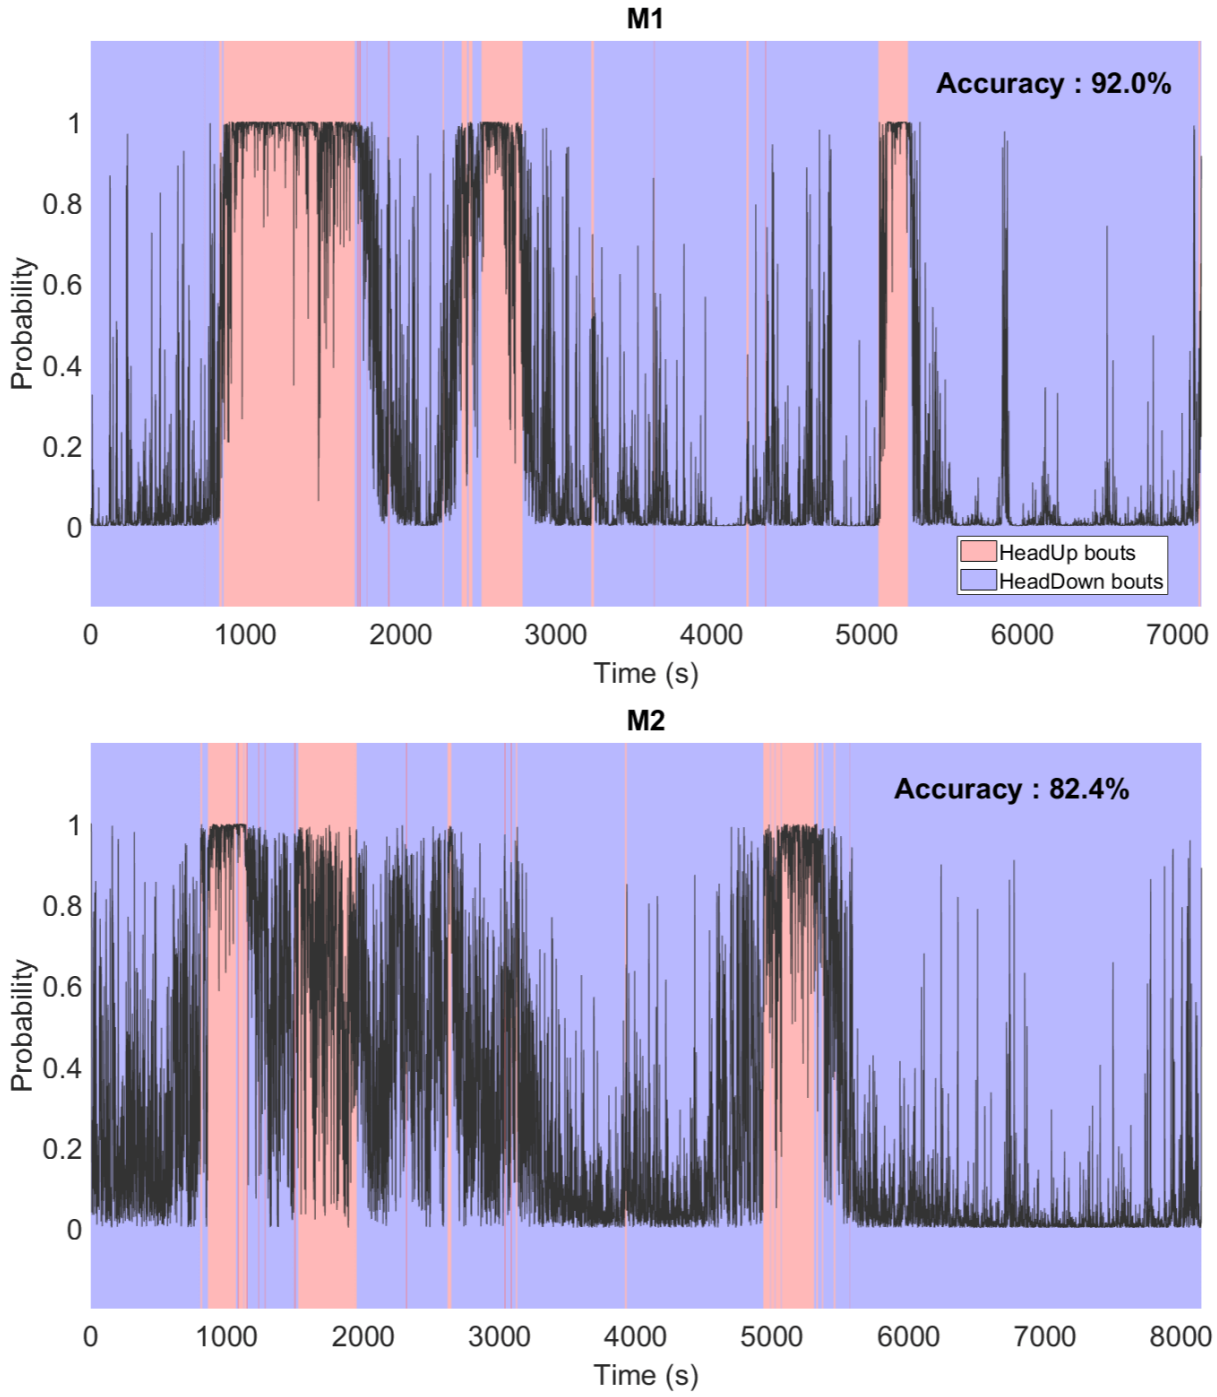

Figure 1.10: Predicted probability of the classifier for the *HeadUp* state, in every timebin in the recording, plotted against time. The background is coloured based on the actual state of the monkey at that time, with *HeadUp* in red and *HeadDown* in blue.

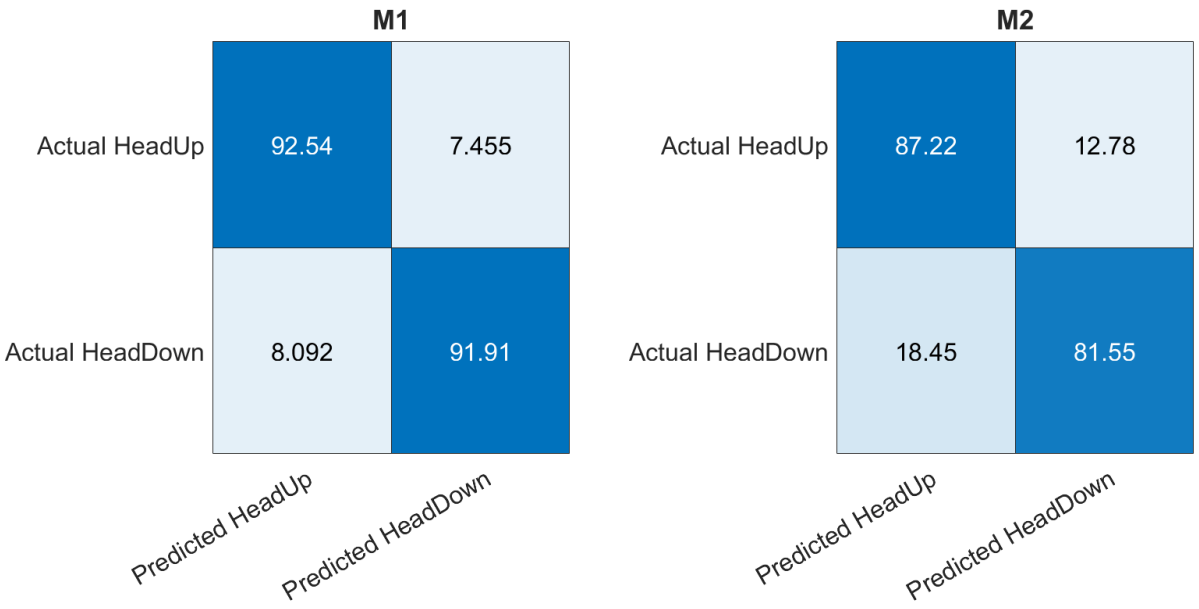

Figure 1.11: Confusion matrix of the classifiers for each of the monkeys. The percentages are calculated based on the total actual labels of each class (normalized by total of each row).

How do the spectral powers change between HeadUp and HeadDown states?

We begin our LFP analysis by computing spectrograms, and then aligning the different bouts with corresponding time points when our behavioural state changes are marked. The event-aligned spectrograms of the LFP data are shown in Figures 1.14 to 1.18. They are a tool, in a similar fashion as for the event-aligned firing rates, to be used to visualize the underlying data, and how it changes, before we move on to any further analysis. It gets harder to visually discriminate many features using the spectrogram plots given their high-resolution requirements, but are an important starting step in the analysis pipeline.

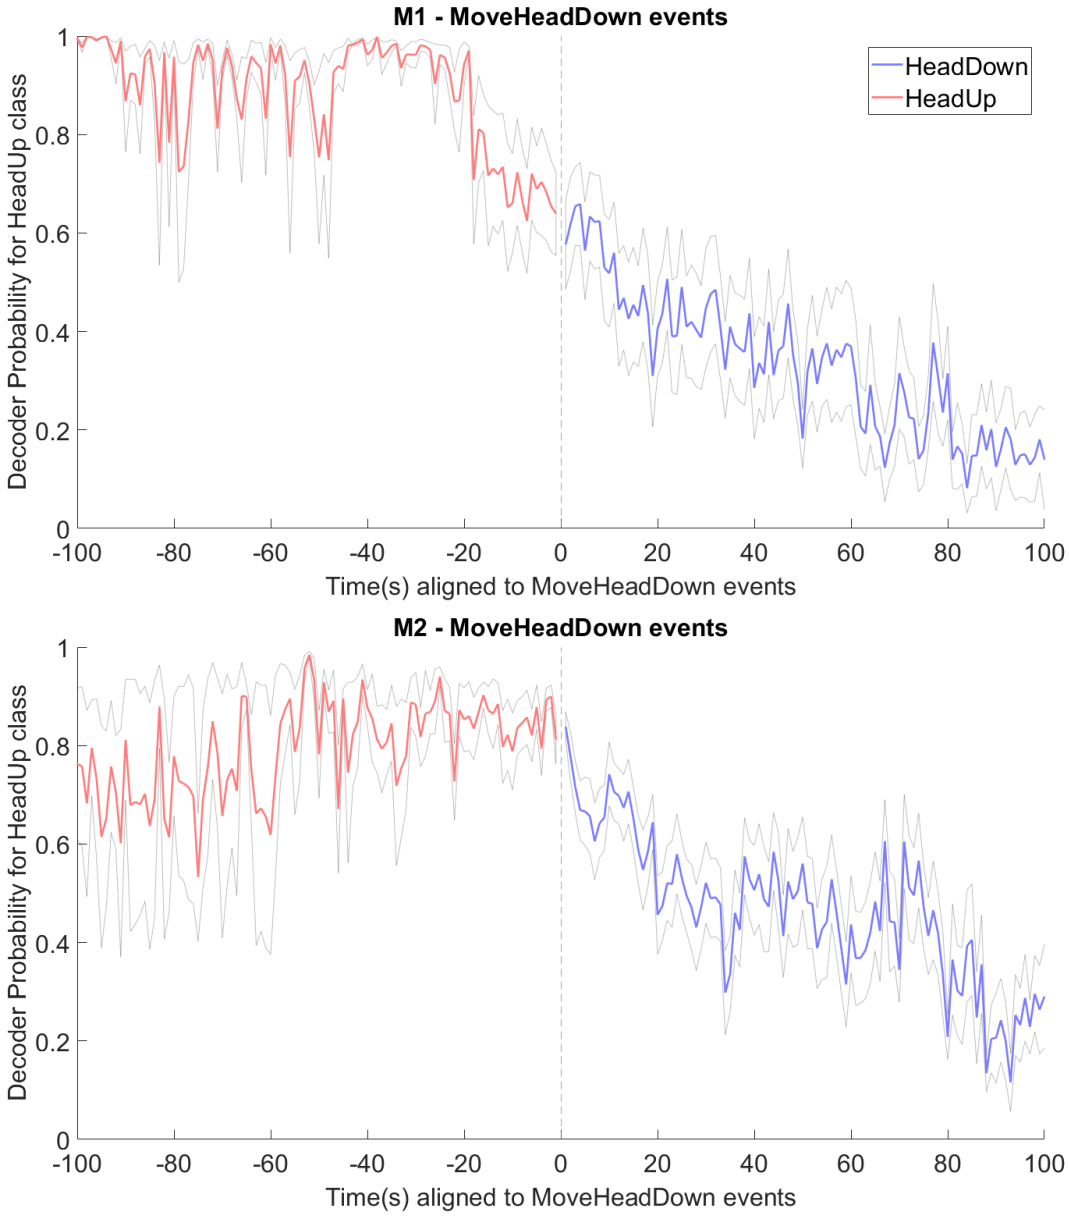

Figure 1.12: Classifier predicted scores surrounding the *MoveHeadDown* events, for M1 and M2. The x-axis is the time in seconds relative to the transition event, and the y-axis is the predicted probability of the classifier, for the class *HeadUp*. The dashed vertical line is the time of the transition event.

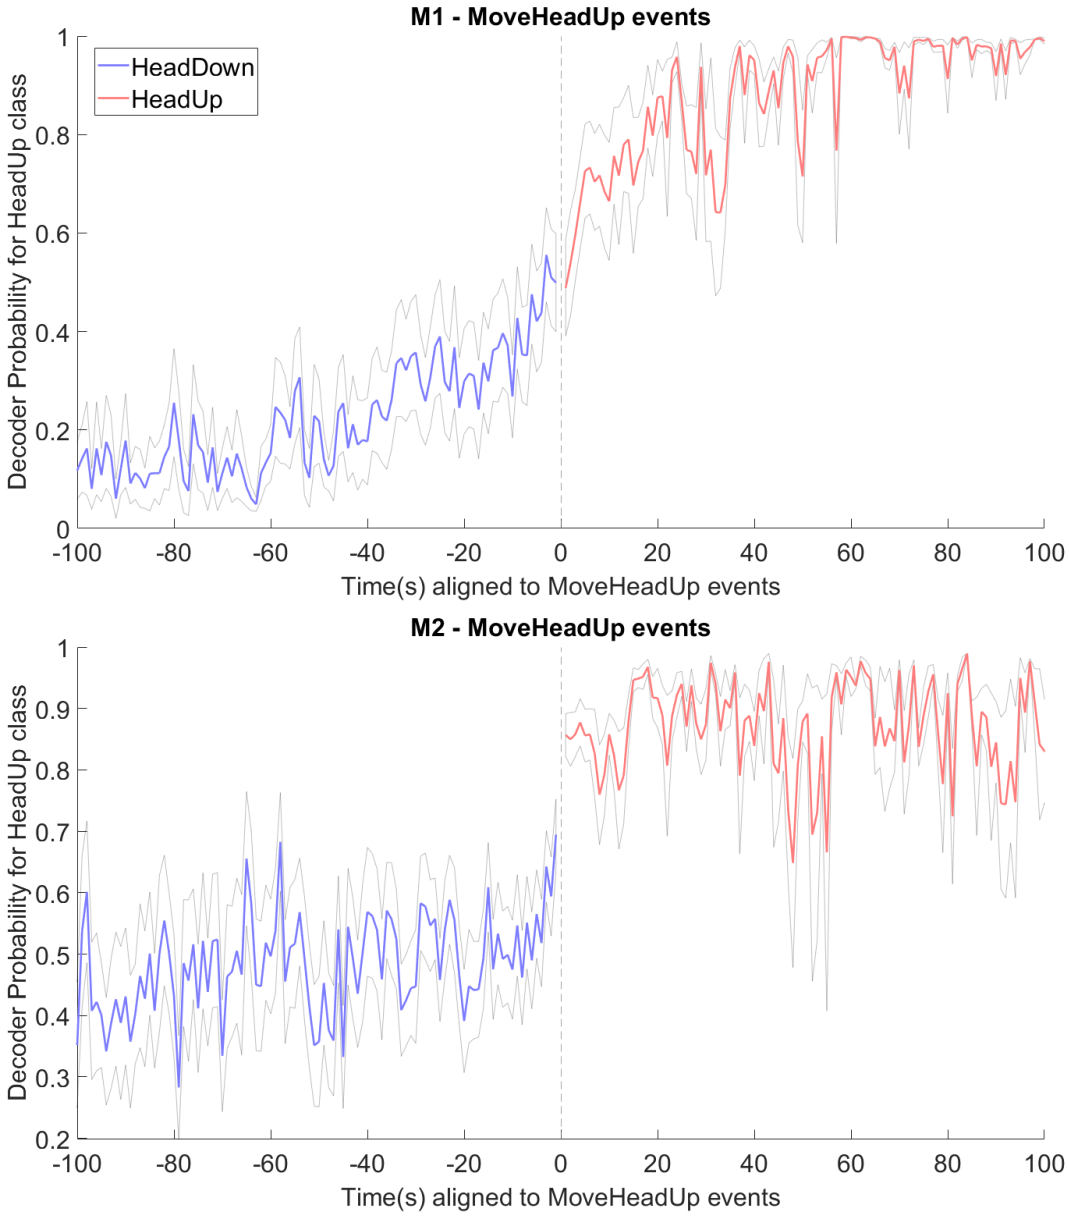

Figure 1.13: Classifier predicted scores surrounding the *MoveHeadUp* events, for M1 and M2. The x-axis is the time in seconds relative to the transition event, and the y-axis is the predicted probability of the classifier, for the class *HeadUp*. The dashed vertical line is the time of the transition event.

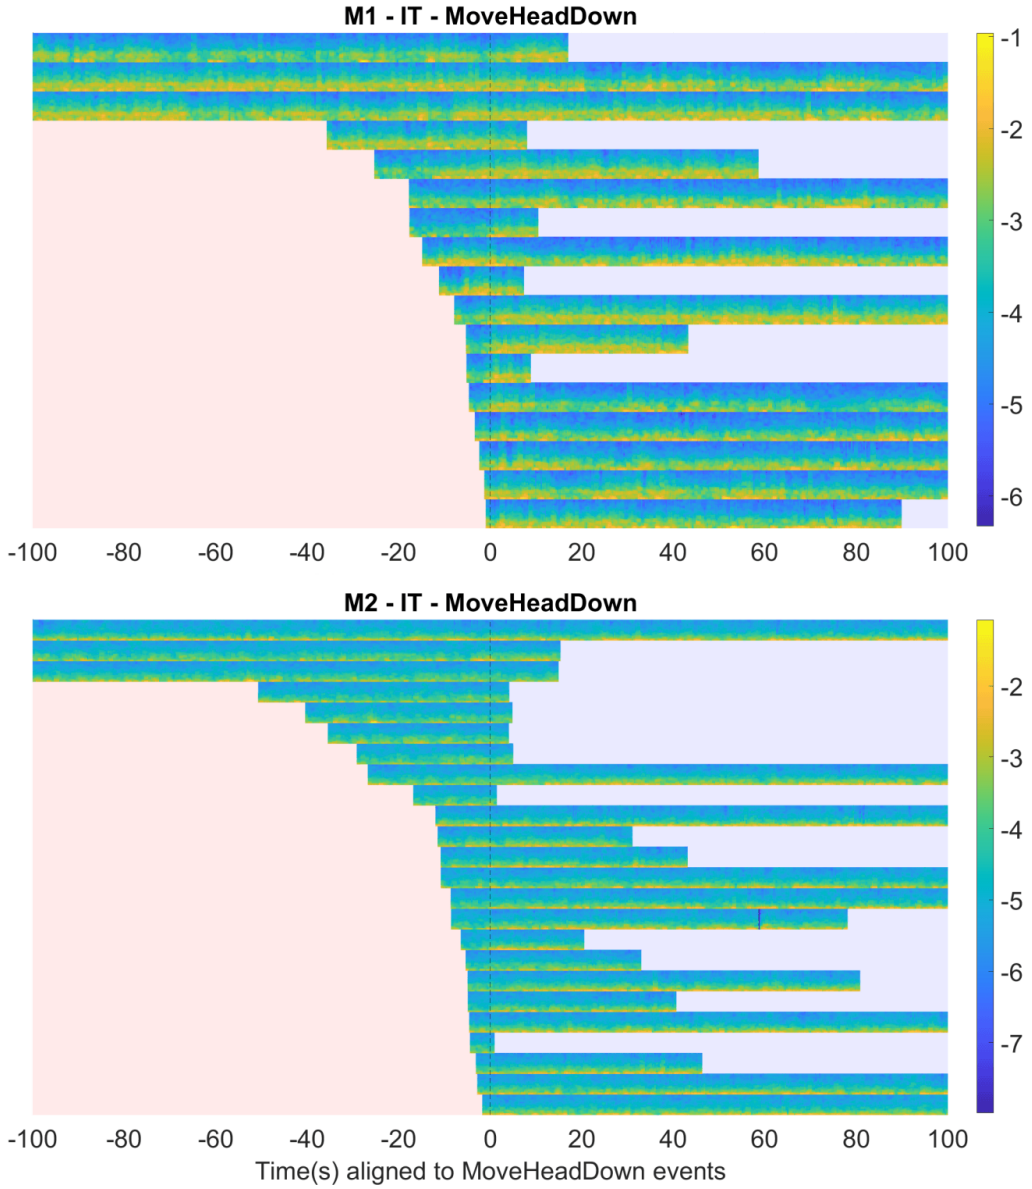

Figure 1.14: Spectrograms (IT region) of pairs of bouts aligned to the MoveHeadDown events. The colourmap represents the log of power over the log of frequencies (0-50Hz) across time in seconds, in this region. Time 0 is the time of the transition event. Top is M1, bottom is M2.

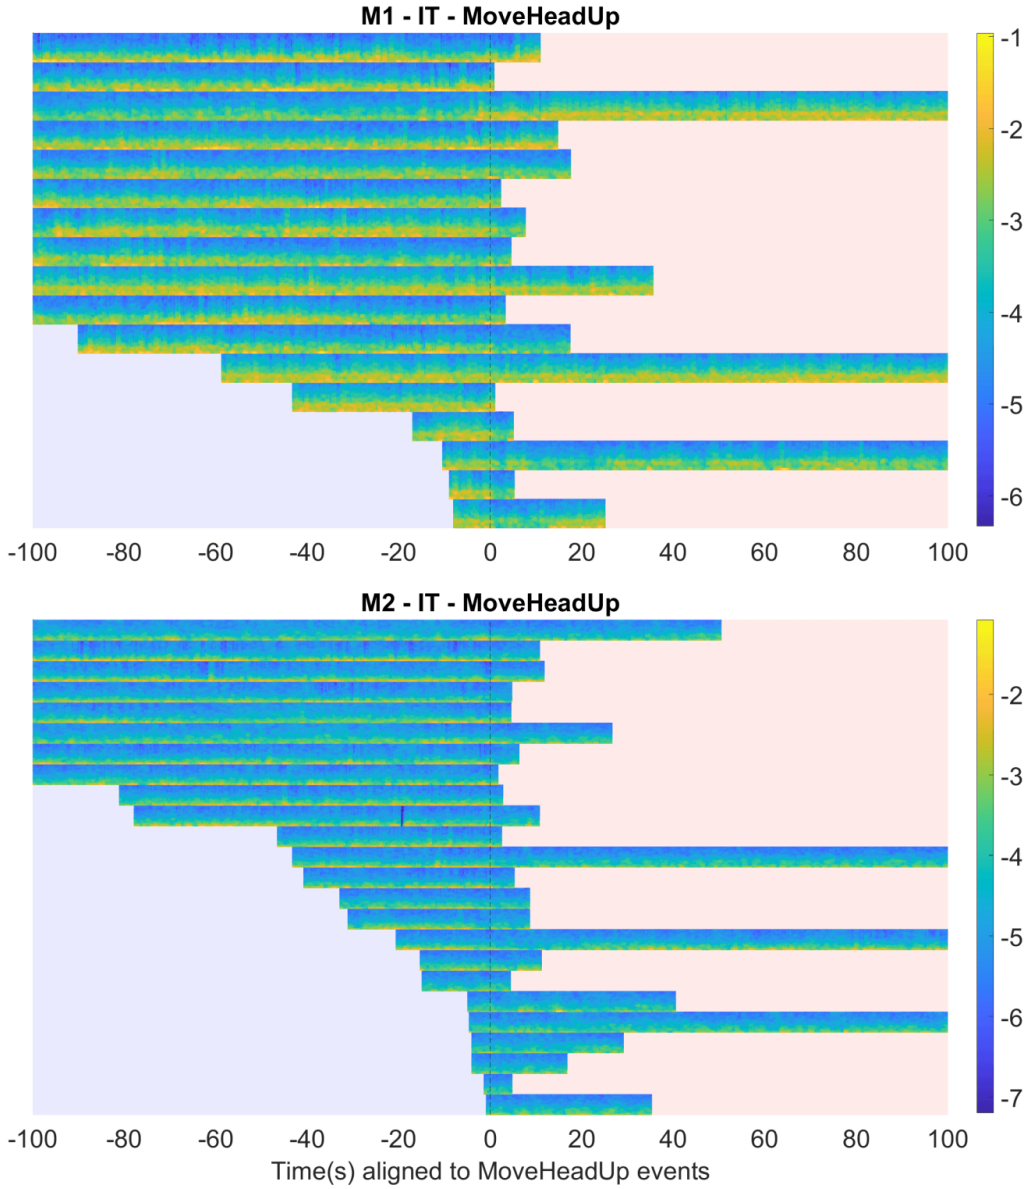

Figure 1.15: Spectrograms (IT region) of pairs of bouts aligned to the MoveHeadUp events. The colourmap represents the log of power over the log of frequencies (0-50Hz) across time in seconds, in this region. Time 0 is the time of the transition event. Top is M1, bottom is M2.

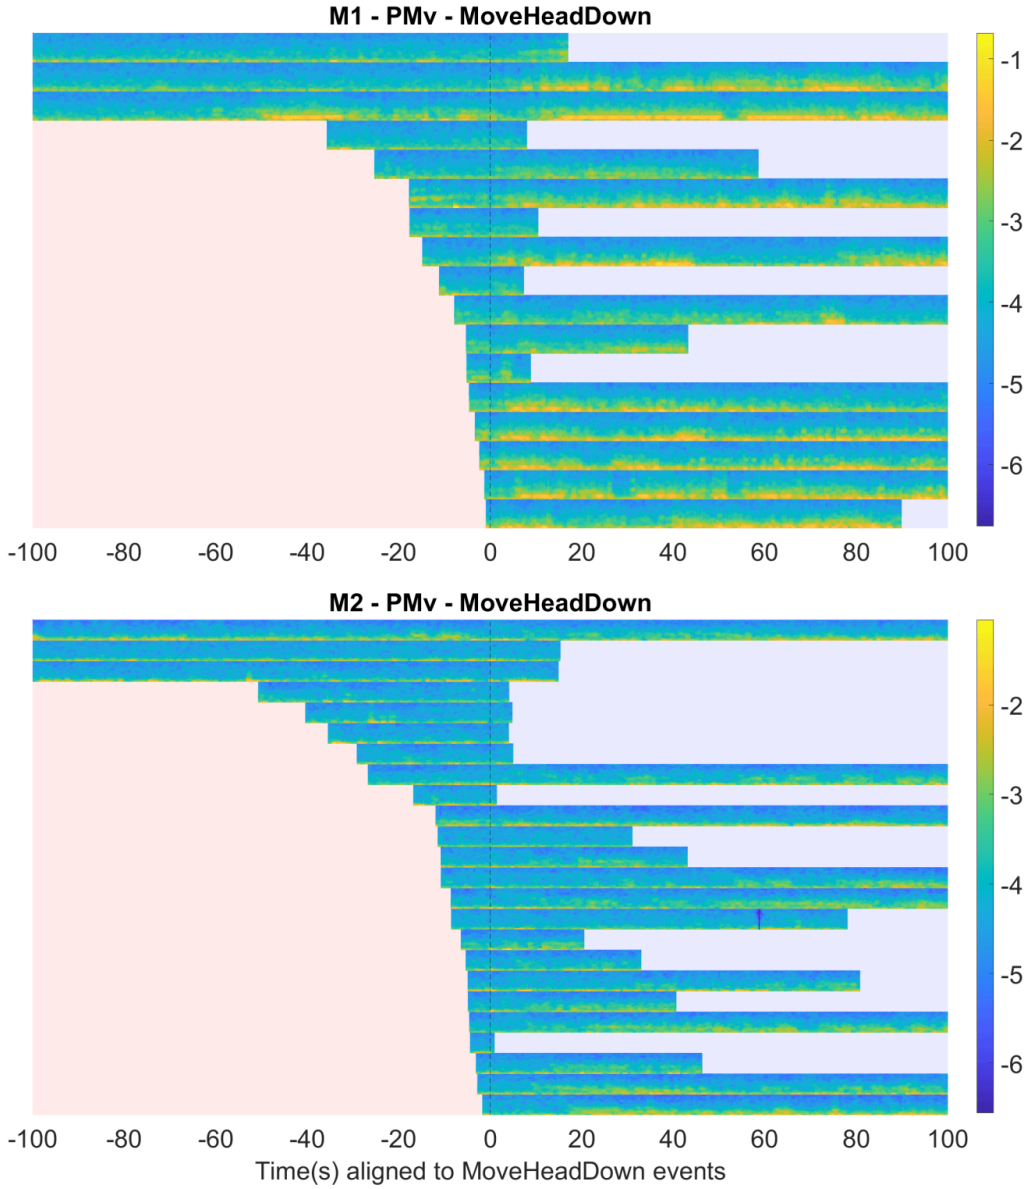

Figure 1.16: Spectrograms (PMv region) of pairs of bouts aligned to the MoveHeadDown events. The colourmap represents the log of power over the log of frequencies (0-50Hz) across time in seconds, in this region. Time 0 is the time of the transition event. Top is M1, bottom is M2.

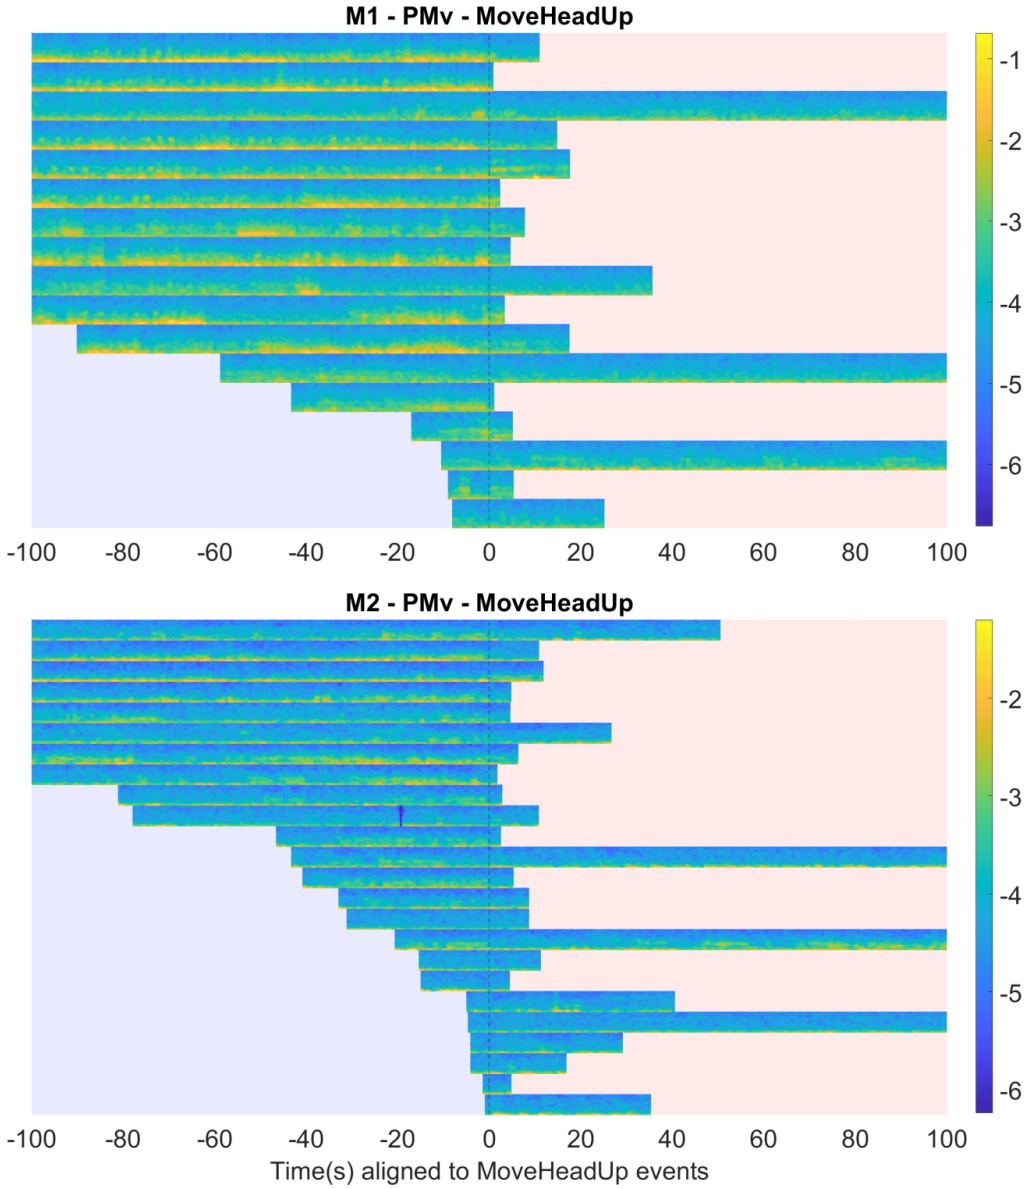

Figure 1.17: Spectrograms (PMv region) of pairs of bouts aligned to the MoveHeadUp events. The colourmap represents the log of power over the log of frequencies (0-50Hz) across time in seconds, in this region. Time 0 is the time of the transition event. Top is M1, bottom is M2.

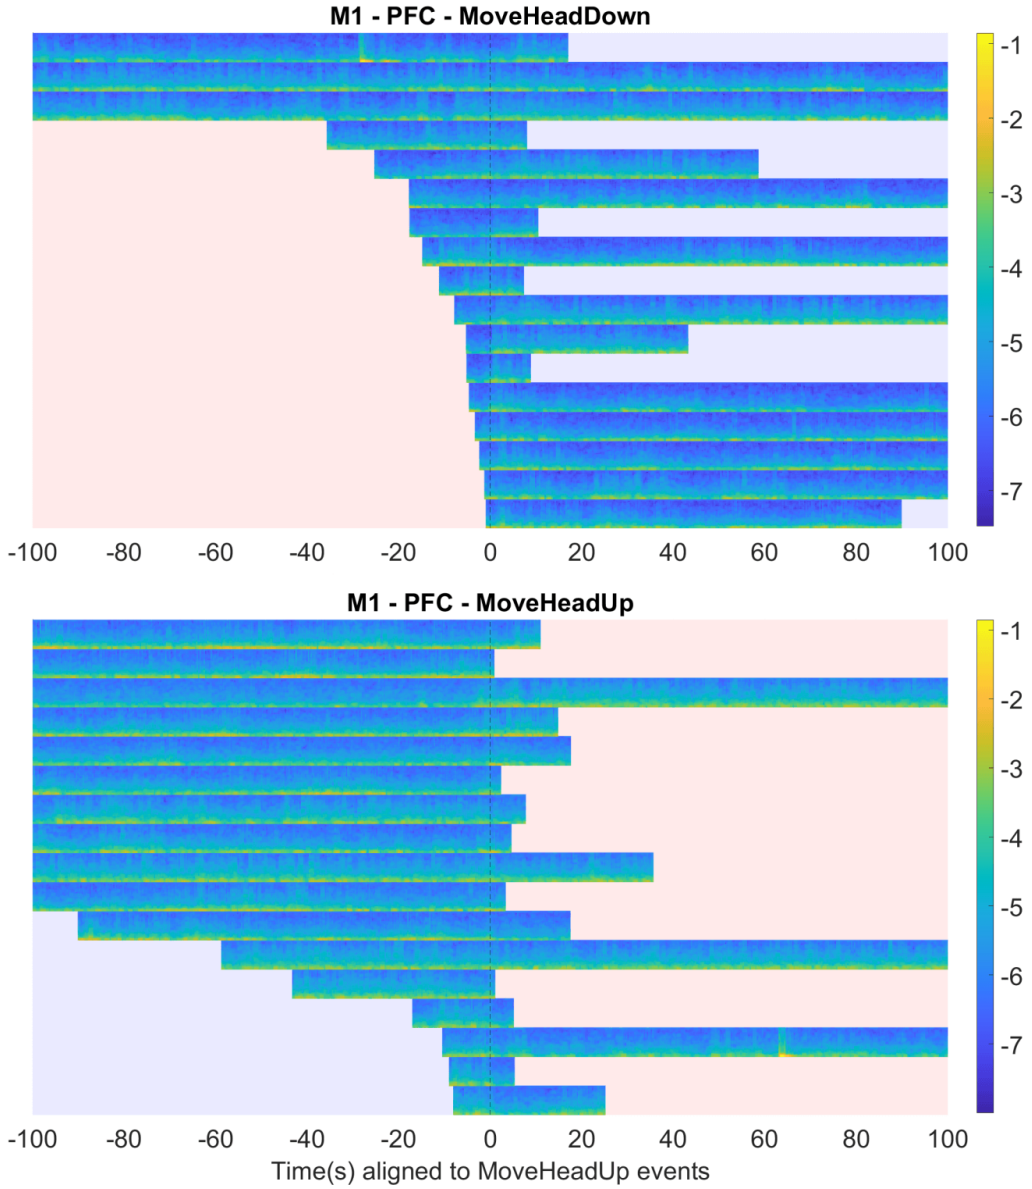

Figure 1.18: Spectrograms (vlPFC region) of pairs of bouts aligned to the MoveHeadDown events (top) and MoveHeadUp events (bottom). vlPFC region recordings are only available from M1. The colourmap represents the log of power over the log of frequencies (0-50Hz) across time in seconds, in this region. Time 0 is the time of the transition event.

How do the spectral powers during the two states compare, globally?

Following the last section, we move on to computing the power spectral density, which shows a global picture of the powers of all the different frequencies and allows us to compare them across the two states. The PSDs of the LFP data are shown in Figure 1.19. The plots across the regions and the monkeys show a bump in the power in the lowest frequencies (0-1Hz) during the *HeadDown* state. In 3/5 regions across the monkeys, there is an increase in the power in the lower frequencies during the *HeadDown* state, and either a flipped higher power for the *HeadUp* state in the higher frequencies, or the difference becomes negligible. However, other differences are not consistent across the monkeys or are not prominent.
It is to be noted that these graphs show/compare time-averages over the duration of *HeadUp* and *HeadDown* bouts, and thus do not capture the dynamics of the changes in the powers within the bouts.

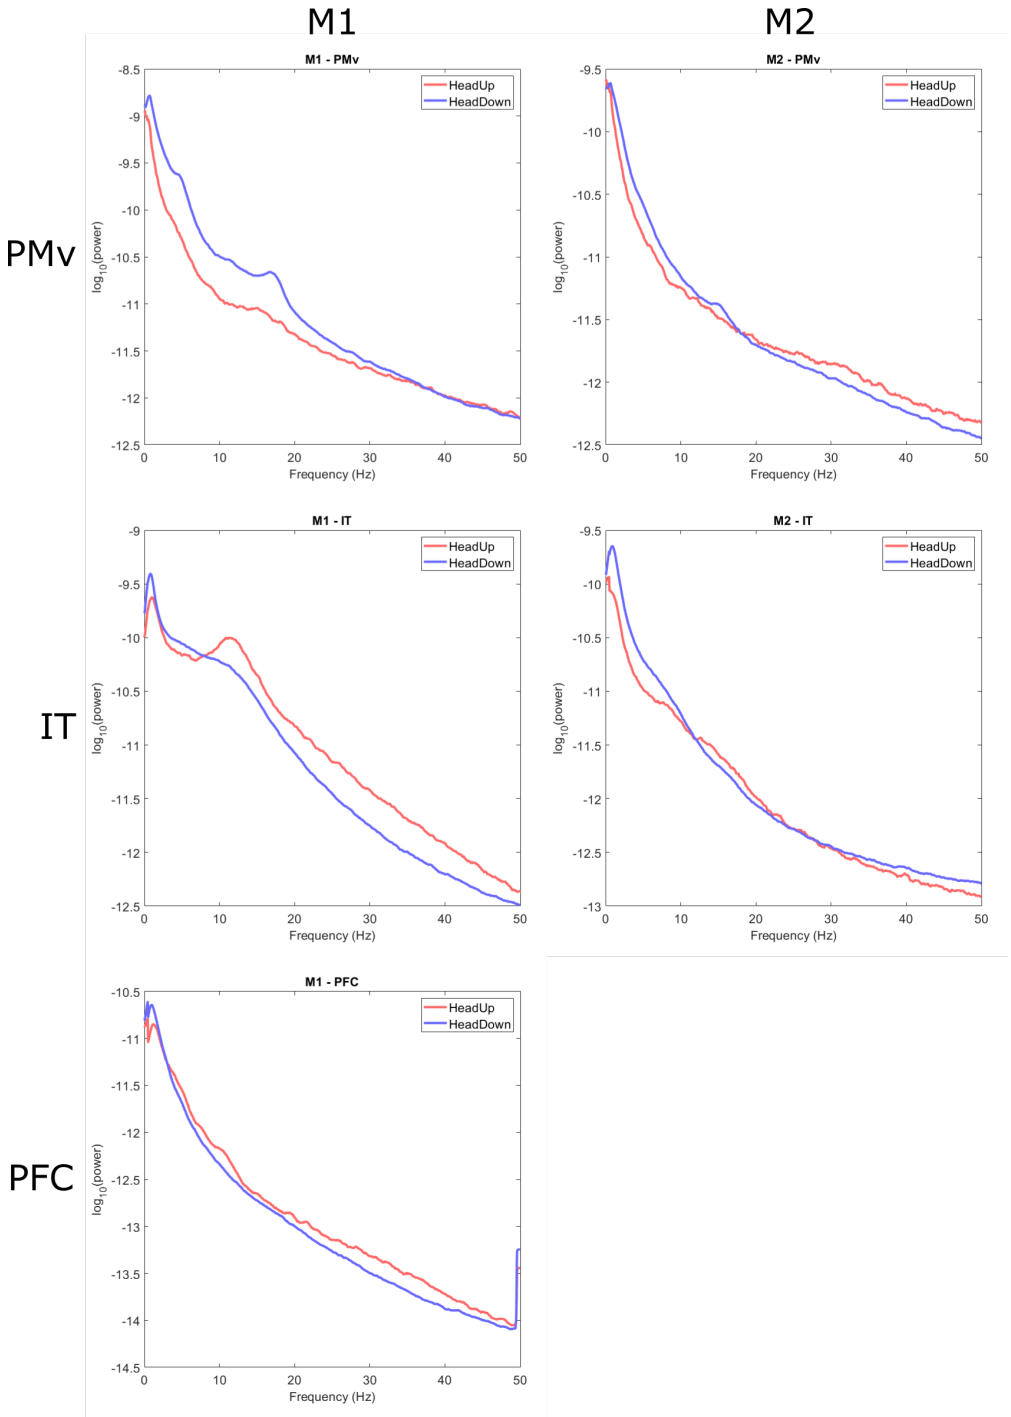

Figure 1.19: Power Spectral Density graphs across the *HeadDown* and *HeadUp* bouts, for each region. The x-axis is the frequency in Hz (0-50Hz), and the y-axis is the log of the power.

How do the average powers in the canonical bands change with the state of the monkey?

The average powers across all *HeadDown* time points versus all *HeadUp* time points is computed, and then, the difference between the *HeadDown* average and the *HeadUp* average (*HeadDown- HeadUp*) is found. This is done separately for each band, and the differences are plotted for each region in Figure 1.20. For all the regions, the delta band has a significantly higher power during the *HeadDown* state (with the exception of M1's vlPFC being higher, but not significantly). For both PMv and IT and across both monkeys, the theta band has a significantly higher power during the *HeadDown* state. On the other hand, the beta and gamma bands show a significantly higher power during the *HeadUp* state, in most cases. The alpha band does increase during *HeadDown* for both M1 and M2's PMv but is inconsistent across monkeys for IT.

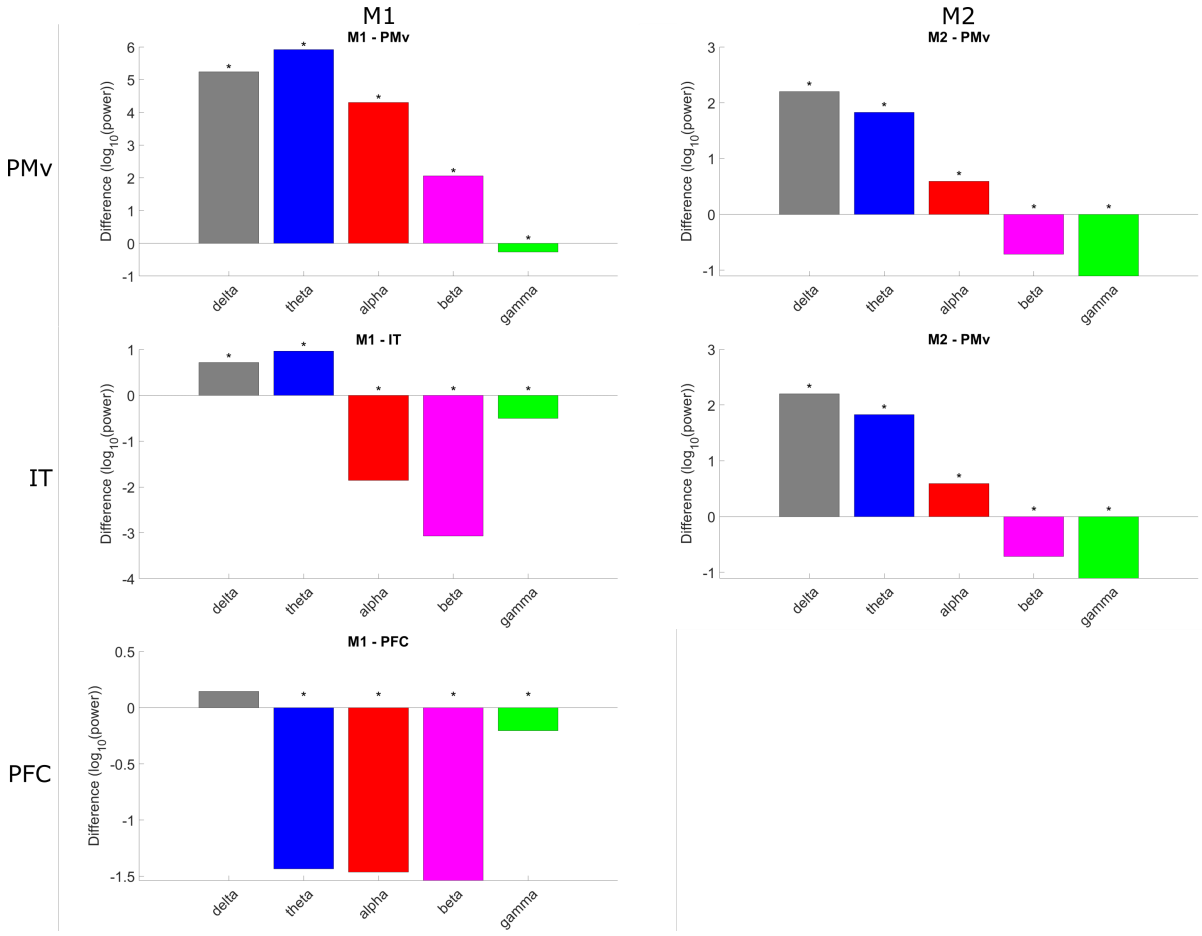

Figure 1.20: Difference in power of different bands during *HeadDown* minus during *HeadUp*, for each region. The x-axis is the frequency band, and the y-axis is the difference in the log of power (*HeadDown* minus *HeadUp*) across the entire duration of the two states. The * represents a significant difference in the power between the two states (p-value $<$ 0.05). The bars are coloured for different frequency bands - gray for delta (1–4 Hz), blue for theta (4–7 Hz), green for alpha (7–15 Hz), yellow for beta (15–31 Hz), and red for gamma (>31 Hz) band.

1.4 Discussion

We performed wireless intracortical recordings of 2 bonnet macaques, from the inferotemporal, ventro-lateral prefrontal cortex and the ventral premotor areas at the onset of sleep. Both monkeys had several bouts of wakefulness and sleep, which also gives us several transition events. This behaviour is in line with previous studies on the sleep cycles of (naturally sleeping) rhesus macaques, a very close relative of the bonnet macaque. The naturalistic housing environment and the wireless nature of the recordings allowed us to capture the neural data while the monkeys slept in their natural position and time of day, which sets it apart from a majority of sleep studies, which are performed in restrained lab settings. Behavioural analysis of sleep using video data, however, has been a well-established method [12].

The changes in neural activity were analyzed, both across the two types of bouts as well as surrounding the time of the transitions from one state to the other. The average firing rates in the different regions were found to be significantly different during the two states, with the PMv region showing the most pronounced difference. The IT region had a more variable firing rate during the *HeadDown* state, and the vlPFC region had a right-skewed distribution of the firing rates during the *HeadDown* state. Similar observations could be made visually, at the times of the transitions: increased firing rate in the PMv regions after a transition to the *HeadUp* state, and a more varied firing rate curve in the IT regions during the *HeadDown* state.

An analysis of the correlation of the firing rates of the channels with the *HeadDown* and *HeadUp* states found a large number of channels that had a positive correlation in the PMv region (implying a higher firing rate during the *HeadUp* state), and a mix of positive and negative correlations in the IT region. The vlPFC region had a strong bias towards negative correlations in M1, implying a lower firing rate during the *HeadUp* state in those channels.

A large majority of sleep studies, across animal models utilize EEG as the neural signal, which lacks the spatial resolution of methods such as ECoG or intracortical recordings [5]. Hence, this study provides a unique look at the neural data during sleep from the regions of IT, PMv and vlPFC, as no previous intracortical studies have been performed in these regions for macaques (the closest would be the study by Yin et al., 2014 [13], which was performed in the motor cortex (M1, leg region)).

A linear discriminant analysis-based classifier based on the firing rates was able to distinguish between the two states with high (92% - M1 and 87% - M2) accuracy. This is a significant result, as it shows that the neural activity (firing rates of the channels) can be used to predict the state of the monkey. There are several studies that have trained deep networks using EEG signals to predict the state of the animal, including the stage within sleep too [14, 15] - however, here we predict amongst the two states using MUA/spike data, and show that a simple linear model works well for this case.

The LFP data was analyzed using spectrograms and power spectral densities, and the powers in the canonical bands were found to be significantly different during the two states. The delta and theta bands had a higher power during the *HeadDown* state, and the beta and gamma bands had a higher power during the *HeadUp* state. This is in line with previous studies on the sleep cycles of rhesus macaques [11], which show a higher power in the lower frequencies during consolidated sleep, and a higher power in the higher frequencies during wakefulness.

However, within this study, there is much left to be explored in the spectral analysis. Some particular future work includes the use of the different spectral bands to classify the stages of sleep, such as REM and NREM sleep, and even the stages within NREM sleep, through the identification of spindles, K-complexes, etc. - this is how the majority of sleep studies are performed on EEG data, and this study provides a rare opportunity to do so with intracortical data. Further, the sleep recordings were performed on days on which the monkeys had performed a task through which they learnt the spatial location of some reward; a future study could involve the analysis of the neural data during sleep, to see if there is any replay of the task-related neural activity during sleep. This analysis typically involves the identification of specific neural activity "sequences" during the learning task, and then looking for those sequences in the sleep recording. Several previous studies have been done in other model organisms such as rats, but not as much in macaques [11].

Chapter 2

Button-controlled player movement in a maze-solving task

2.1 Introduction

Scientists have in the past trained monkeys to various extents, including having them play games [27, 26]. However, having access to brain data while they're doing so is much more rare, and such an opportunity is presented to us at our lab. The bonnet macaques at our lab have been well-trained to do certain touchscreen experiments, and have been religiously doing them for years, as part of various scientific studies. These experiments may vary in the underlying complexity of the task paradigm itself, but the macaques have shown that they can become extremely proficient at doing the tasks, with enough practice. With this in mind, we attempted to train the monkeys how to control the movement of a character on the screen by the use of 4 on-screen buttons, to move the character to certain reward objects as a part of a "maze" layout. A wide range of possible studies are opened up following such a training. In this chapter, I aim to outline the methods we used to train the monkeys to do such a task.

2.2 Methods

The (freely-moving) monkeys have been previously trained to interact with, and use, a touchscreen, to perform several tasks on them. The typical tasks are divided into a number of trials, and for performing a trial correctly, the monkey receives an amount of juice, delivered through a juice spout, as well as a "correct sound", which is played at the end of a trial. For an incorrect trial, no juice is delivered, and an "incorrect sound" is played. Some tasks are designed with the aim of collecting neural data during active fixation on a target area on the touchscreen, during which various stimuli are presented. Other tasks involve the use of memory and decision-making, such as a temporal same-different task, in which the monkey is shown two images in succession, has to decide if they were the same or different, and choose one out of two buttons on the screen based on this decision. The tasks are designed in MonkeyLogic [19], NIMH's open-source software for behavioural control and data acquisition, which is used to present stimuli on the screen, and record the monkey's responses.

2.2.1 Task design

The subtasks used during training may have several modifications, but the general task paradigm goes as follows: (Figure 2.1)

- A red "hold" button is presented, which must be touched for the trial to begin.
- The trial begins with a character (an image of a monkey, or human) being presented on the screen, in a "maze" layout. The maze also has reward objects (an image of a banana, or strawberry) placed at certain locations. The number and locations of rewards depend upon the particular level or subtask. A short delay period within which only the maze layout is present, follows this, before the next step.
- Four buttons, each with an arrow pointing in a different direction, are presented on the screen, beside the maze layout. The monkey must touch one of these buttons to move the character in the correct direction (the direction in which the reward is, relative to the character). Only one button can be touched at a time.

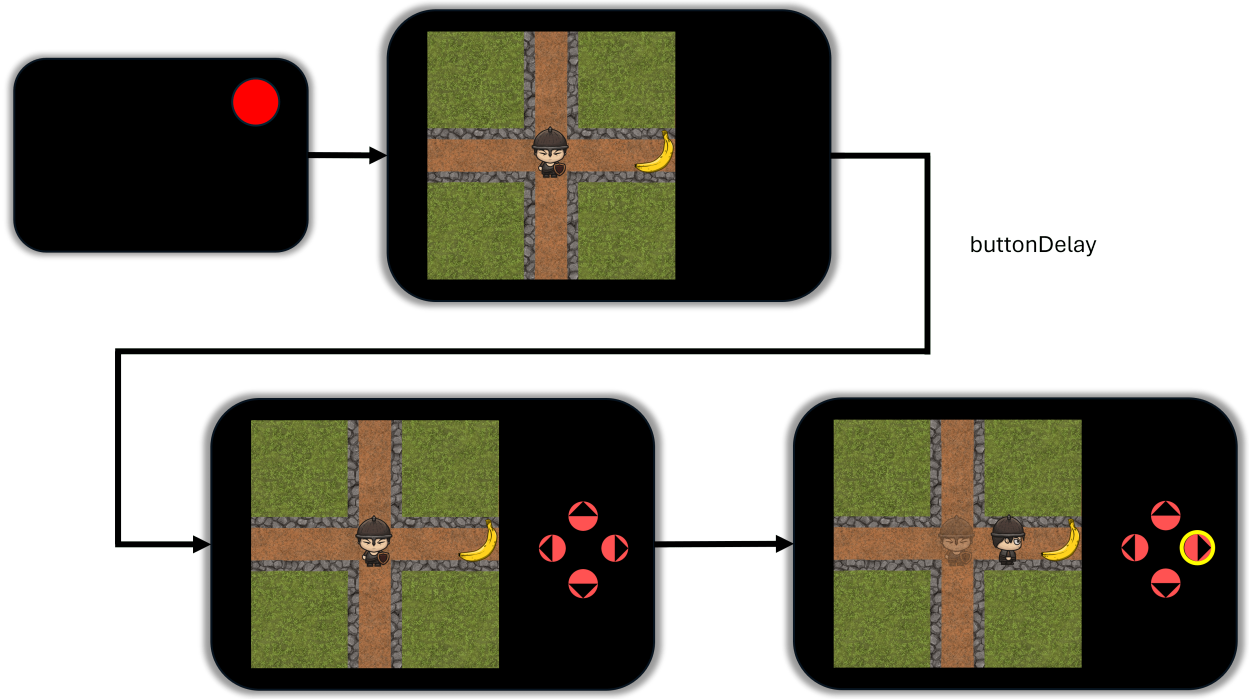

Figure 2.1: Overview of the task design. The red button is the hold button, which must be touched to start the trial. A maze layout is shown, with the character (here, a cartoon human) and reward objects (here, a banana). The four buttons are shown, which the monkey must touch to move the character. The arrows represent the flow of time in the task. In the last panel, the lower opacity character at the center is only representative, and not actually shown on the screen.

- The character moves in the direction of the button touched, and stops when it reaches a reward object. The trial ends, and the monkey receives juice and a correct sound.

- If the character reaches an end of the maze which doesn't have a reward object, the trial ends, and the monkey receives no juice and an incorrect sound. In another version of the task, the character is simply stopped, and the monkey can continue to control it, but is instead given a timeout period, after which the trial ends.

2.2.2 Task training

To train the monkeys to be able to do the task, we followed a step-by-step training procedure, which involved two "levels", and multiple sublevels/subtasks for each level. The first level

involves a plus or X maze, in which the character starts at the center, and then has to move to the end of one of the four arms of the maze. In each trial, the reward is placed in one of these 4 ends of the maze.

In level 2, the maze layout is made more complicated, with more arms (an H-like shape - Figure 2.6), and more reward objects per trial. The juice reward and corresponding correct sound are given for every reward object reached, at the time it is reached. Once a reward is "consumed", it disappears from the screen, and reaching the same location again will not be rewarded.

Level 1

The first level of training involved training the monkeys to move the character to the end of one of the four arms of the maze, based on the location of the reward object.

In general, a sublevel can be thought of as a single day of training, although it may exceed a day per level, depending on the performance of the monkey. The sublevels in this level were as follows:

- **Level 1, Sublevel 1** (Figure 2.2): The character starts at the center of the maze, and the reward is placed at the end of one of the arms. However, in one trial, a maximum of two buttons appear. The correct button, and the button opposite to the correct button (an incorrect button). As detailed in the figure, we progress through this level with a changing opacity of the incorrect button.

- **Level 1, Sublevel 2** (Figure 2.3): The character starts at the center of the maze, but the maze is now an X shape. Also, the causality of the task is now flipped: the character moves by itself, in the direction of the reward, and "consumes" the reward, but the monkey must touch the corresponding button (which is located right below, or next to, the reward's location), in order to get the juice reward and the correct sound. In this stage, the movement and reward location of the "right" and "down" trials are actually the same - however, the buttons are different.

- **Level 1, Sublevel 3** (Figure 2.4): This sublevel is intended to gradually bring the monkey to the plus maze layout. We slowly (across trials) move the reward objects

toward where they would have been, in the plus maze layout - correspondingly, the character also moves in a vector that reaches the shifted reward location. The button locations are kept the same as in the previous sublevel. By the end of this sublevel, the reward objects are at the correct locations for the plus maze layout, but the buttons are still at the edges of the maze, as in sublevel 2.

- **Level 1, Sublevel 4** (Figure 2.5): In this level, we move the buttons gradually (across trials) to the intended locations (grouped together, to the side of the maze, as can be seen in Figure 2.1, or Figure 2.5, bottom panel). Also, by the end of this level, all 4 buttons are shown at the same time, as in the previous levels the maximum number of buttons shown per trial was 2. At the end of this level, the monkey is trained to move the character to the end of one of the arms of the maze, based on the location of the reward object, using the buttons at the side of the maze.

Level 2

The level 2 training involves the larger, much more complex maze layout. The sublevels in this level are simply described as every possible starting position, which are the nodes at every end of the maze, and every joint of the arms of the maze, and for each start position, rewards are randomly placed at different nodes of this maze. The number of rewards varied from a minimum of 1 to a maximum of 3. The buttons were placed at the side of the maze, as in the previous level. The training ensures that the monkey completely understands the maze's layout and the restrictions in place for the movement of the character.

Figure 2.6 Shows one example start point and a possible reward layout, and a possible path that the monkey might take to reach all the rewards.

There are two major changes in the game's design in this level:

- At every node, the player artificially stops for a short duration (usually 300ms, but this was changed over the days of training)
- At every node, only the buttons which are valid moves, are shown at full opacity, while the other button(s) is(are) shown at 10% opacity. Two examples of such nodes and corresponding valid buttons are shown in Figure 2.6 B and Figure 2.6 C. Touching an

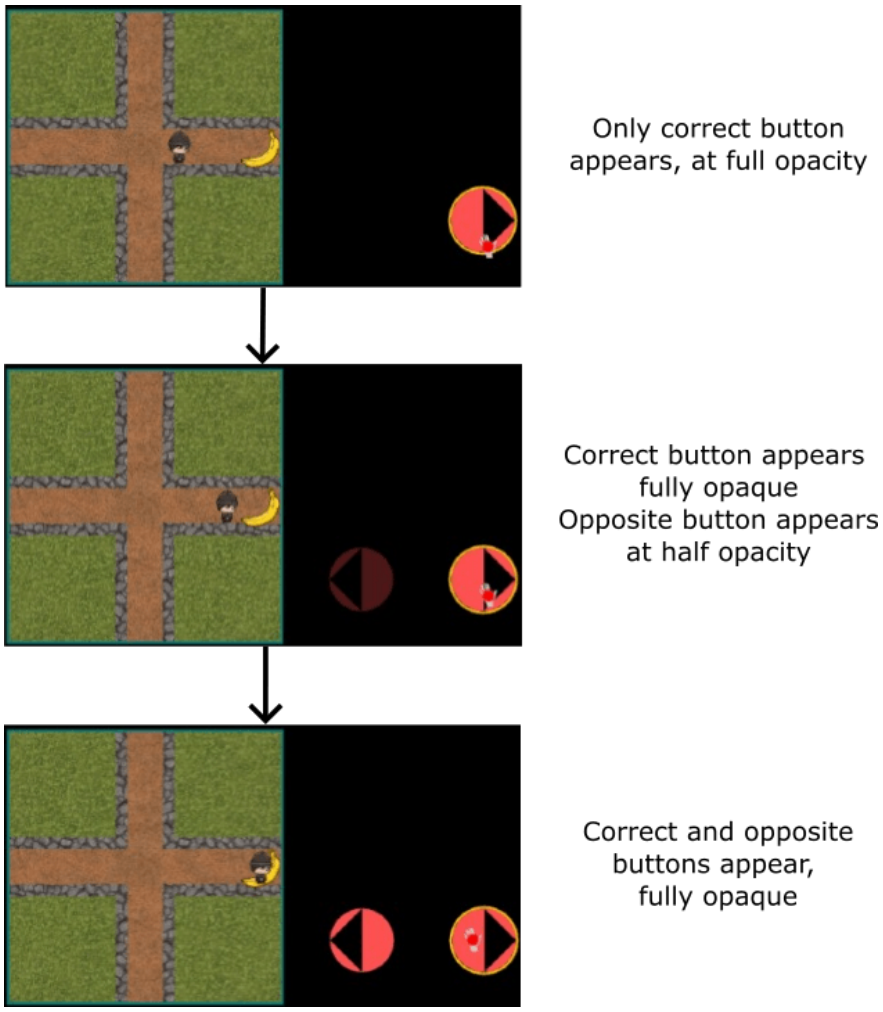

Figure 2.2: Level 1, Sublevel 1: Character at center, reward at end of arm. Two buttons: correct and incorrect, but the opacity of the incorrect button is slowly increased.

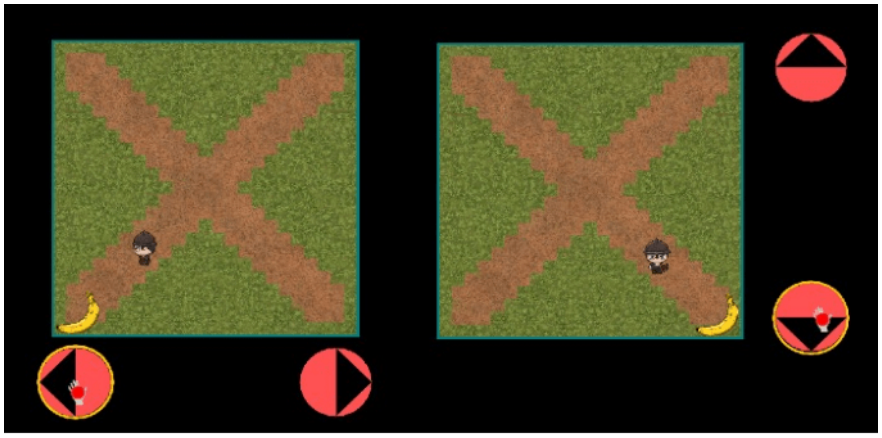

X maze,
Flipped Causality:
movement, followed by button press

Figure 2.3: Level 1, Sublevel 2: Character at center, but maze in X shape. Character moves to the reward always, but monkey needs to touch the correct button after the movement to get a reward.

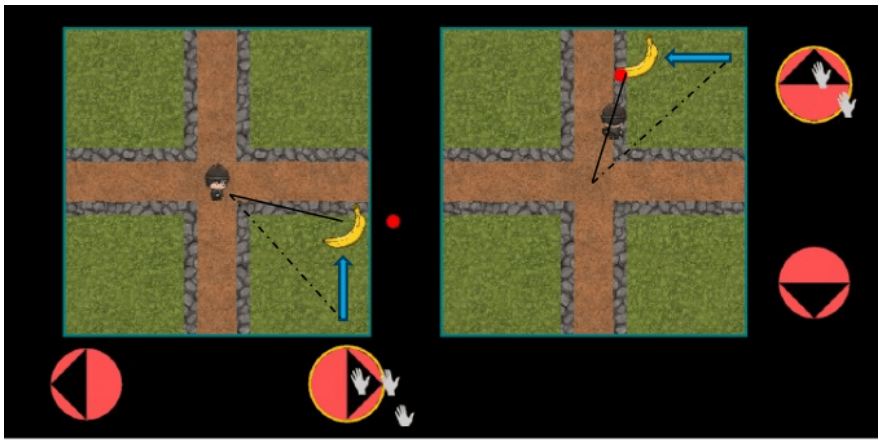

Flipped Causality
Gradually returned to + maze

Figure 2.4: Level 1, Sublevel 3: Gradual transition back to plus maze layout. Reward objects and character move accordingly. Same button locations as previous sublevel.

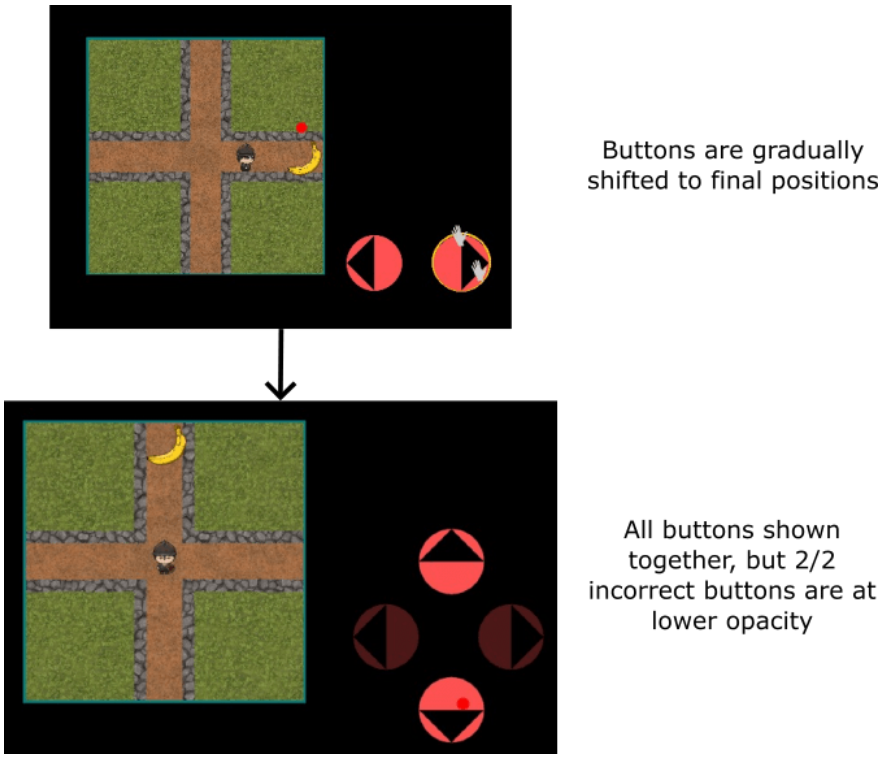

Figure 2.5: Level 1, Sublevel 4: Buttons are gradually moved to intended locations, partial hints can be given by lowering the opacity of 2 incorrect buttons. All 4 buttons are shown.

invalid button does not move the character, and the monkey must touch any of the other available buttons to move the character to the next node.

2.3 Results

We were able to successfully train M2 to navigate the Level2 task/maze completely. M1 has completed the Level1 task but is yet to complete the training for Level2. This measure of success is currently defined as the monkey being able to reach all the rewards in a single trial, without timing out the trial (which was set at 30 seconds). This timing of each trial requires the monkey to make the optimal choices at almost all the nodes. The monkey was able to do this consistently, across multiple days. However, a robust analysis of the performance of the monkeys, and the learning curve, is yet to be done, based on the behavioural data collected during the training.

2.4 Discussion

The training of the monkeys on this new task paradigm opens up the possibility of a wide range of studies, which can be performed on the monkeys. The task is designed to be complex enough to require the monkey to make decisions at every node, and yet simple enough to be able to be completed in a reasonable amount of time, along with more intermittent receipt of reward, as an incentive for good performance. An important distinction of this paradigm compared to others is that the monkey's actions/behaviour dynamically and continually affects the stimuli presented to it. In trials that require multiple movements, the monkey's previous actions in the trial can affect its performance through the rest of the trial - a mistake is required to be first corrected, as otherwise, it's more than likely that this mistake will get compounded over the trial. This is in contrast to many other tasks, in which the stimuli are pre-determined, and the monkey's actions are simply responses to these stimuli - one response determines the fate of that trial. A study currently being planned is to investigate the sense of "agency", which we humans clearly feel while controlling a character in a video game. Some other studies that are possible with this task are related to the neural correlates of path-planning, decision making etc.

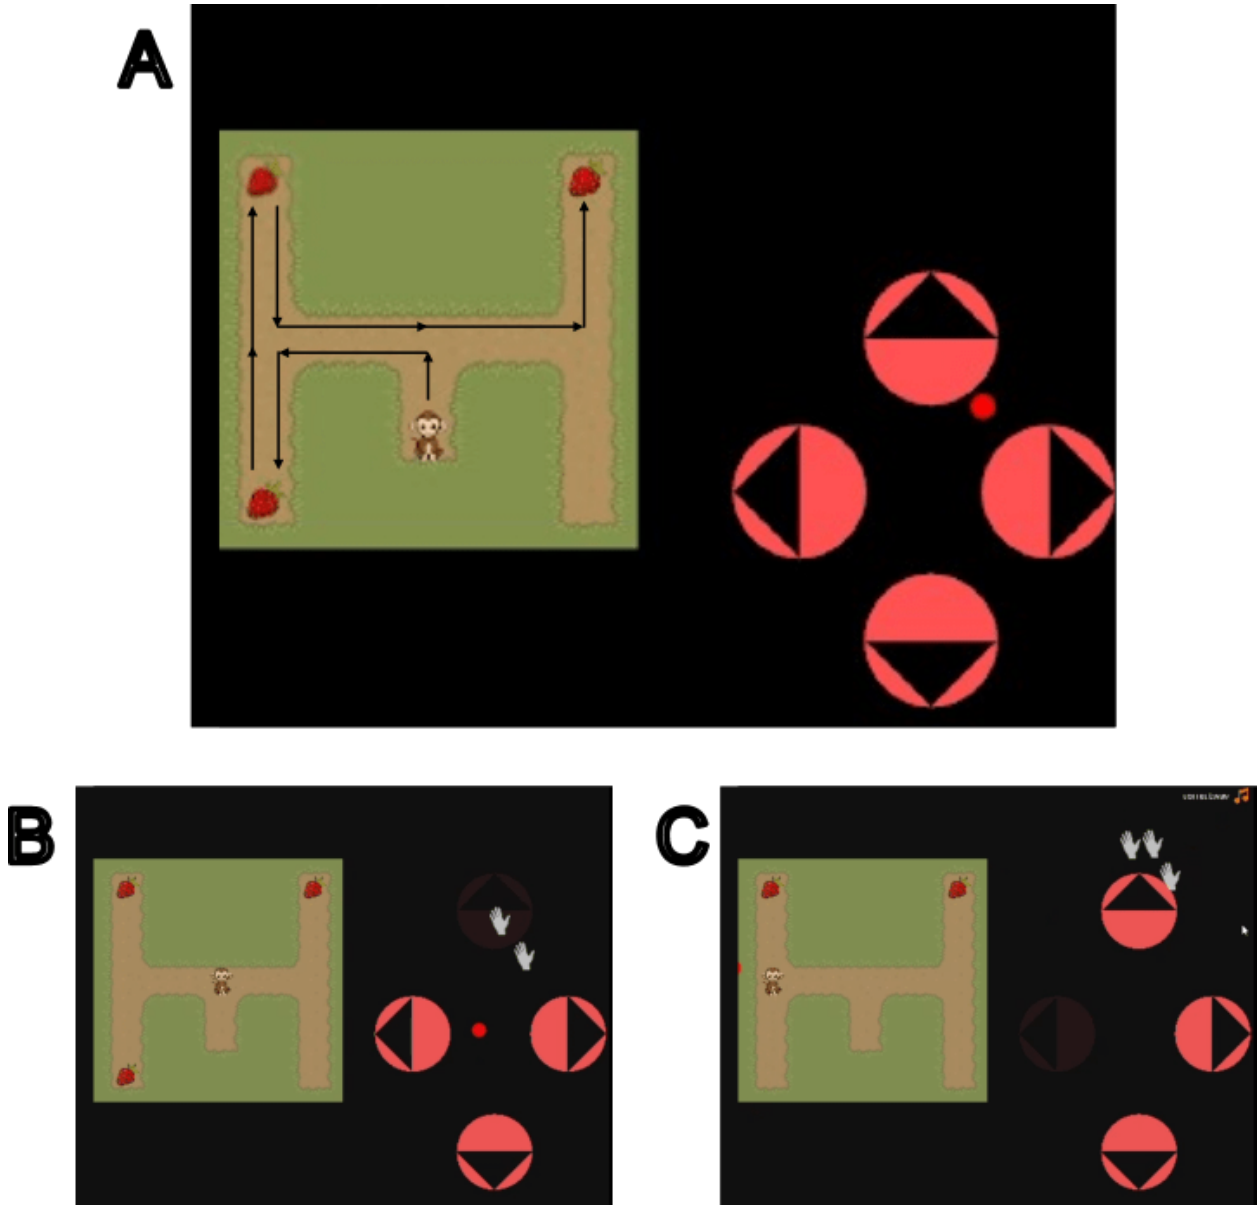

Figure 2.6: **Panel A** shows an example trial, with the H-like maze layout, the character at one of the possible nodes, and 3 rewards placed at different nodes. The black arrows show a possible path that can be taken to obtain all 3 rewards. **Panels B** and **C** show nodes at which the character might be during a trial, where its constrained such that it can only move in 3 directions, and the corresponding buttons show which moves are valid (fully opaque) and invalid (10% opacity).

Bibliography

[1] F. P. Cappuccio, L. D'Elia, P. Strazzullo, and M. A. Miller, "Sleep Duration and All-Cause Mortality: A Systematic Review and Meta-Analysis of Prospective Studies," *Sleep*, vol. 33, pp. 585–592, May 2010.

[2] E. Kronholm, M. Härmä, C. Hublin, A. R. Aro, and T. Partonen, "Self-reported sleep duration in Finnish general population," *Journal of Sleep Research*, vol. 15, pp. 276–290, Sept. 2006.

[3] A. C. Keene and E. R. Duboue, "The origins and evolution of sleep," *Journal of Experimental Biology*, vol. 221, p. jeb159533, June 2018.

[4] K. Datta, A. Bhutambare, and H. N. Mallick, "Systematic Review of Prevalence of Sleep Problems in India: A Wake- up Call for Promotion of Sleep Health," Dec. 2023.

[5] A. J. Jakobson, A. R. Laird, J. J. Maller, R. D. Conduit, and P. B. Fitzgerald, "Brain Activity in Sleep Compared to Wakefulness: A Meta-Analysis," *Journal of Behavioral and Brain Science*, vol. 02, no. 02, pp. 249–257, 2012.

[6] G. Gauquelin-Koch, J.-P. Blanquie, M. Viso, G. Florence, C. Milhaud, and C. Gharib, "Hormonal response to restraint in rhesus monkeys," *Journal of Medical Primatology*, vol. 25, pp. 387–396, Dec. 1996.

[7] I. G. Campbell, "EEG Recording and Analysis for Sleep Research," *Current Protocols in Neuroscience*, vol. 49, Oct. 2009.

[8] G. Jacob, H. Katti, T. Cherian, J. Das, KA. Zhivago, and SP. Arun, "A naturalistic environment to study visual cognition in unrestrained monkeys," *eLife*, vol. 10, p. e63816, Nov. 2021.

[9] "Chronux MATLAB Package." http://chronux.org/.

[10] P. Mitra and H. Bokil, *Observed Brain Dynamics*. Oxford University Press, Dec. 2007.

[11] W. Xu, F. de Carvalho, and A. Jackson, "Sequential Neural Activity in Primary Motor Cortex during Sleep," *Journal of Neuroscience*, vol. 39, pp. 3698–3712, May 2019.

[12] J. Muñoz-Delgado, G. Luna-Villegas, R. Mondragón-CeballoS, and A. Fernández-Guardiola, "Behavioral characterization of sleep in stumptail macaques (Macaca arctoides) in exterior captivity by means of high-sensitivity videorecording," *American Journal of Primatology*, vol. 36, no. 3, pp. 245–249, 1995.

[13] M. Yin, D. A. Borton, J. Komar, N. Agha, Y. Lu, H. Li, J. Laurens, Y. Lang, Q. Li, C. Bull, L. Larson, D. Rosler, E. Bezard, G. Courtine, and A. V. Nurmikko, "Wireless Neurosensor for Full-Spectrum Electrophysiology Recordings during Free Behavior," *Neuron*, vol. 84, pp. 1170–1182, Dec. 2014.

[14] V. Svetnik, T.-C. Wang, Y. Xu, B. J. Hansen, and S. V Fox, "A Deep Learning Approach for Automated Sleep-Wake Scoring in Pre-Clinical Animal Models," *Journal of Neuroscience Methods*, vol. 337, p. 108668, May 2020.

[15] J. G. Ellen and M. B. Dash, "An artificial neural network for automated behavioral state classification in rats," *PeerJ*, vol. 9, p. e12127, Sept. 2021.

[16] L. K. Barger, T. M. Hoban-Higgins, and C. A. Fuller, "Gender differences in the circadian rhythms of rhesus monkeys," *Physiology & Behavior*, vol. 101, pp. 595–600, Dec. 2010.

[17] M. A. Cruz-Aguilar, E. Hernández-Arteaga, M. Hernández-González, I. Ramírez-Salado, and M. A. Guevara, "Principal component analysis of electroencephalographic activity during sleep and wakefulness in the spider monkey (Ateles geoffroyi)," *American Journal of Primatology*, vol. 82, p. e23162, Aug. 2020.

[18] K.-C. Hsieh, E. L. Robinson, and C. A. Fuller, "Sleep Architecture in Unrestrained Rhesus Monkeys (Macaca mulatta) Synchronized to 24-Hour Light-Dark Cycles," *Sleep*, vol. 31, pp. 1239–1250, Sept. 2008.

[19] J. Hwang, A. R. Mitz, and E. A. Murray, "NIMH MonkeyLogic: Behavioral control and data acquisition in MATLAB," *Journal of Neuroscience Methods*, vol. 323, pp. 13–21, July 2019.

[20] M. J. Jorgensen, S. J. Suomi, and W. D. Hopkins, "Using a Computerized Testing System to Investigate the Preconceptual Self in Nonhuman Primates and Humans," in *Advances in Psychology*, vol. 112, pp. 243–256, Elsevier, 1995.

[21] T. Kaneko and M. Tomonaga, "The perception of self-agency in chimpanzees (*Pan troglodytes*)," *Proceedings of the Royal Society B: Biological Sciences*, vol. 278, pp. 3694–3702, Dec. 2011.

[22] K. Masuda and I. V. Zhdanova, "Intrinsic activity rhythms in Macaca mulatta: Their entrainment to light and melatonin," *Journal of Biological Rhythms*, vol. 25, pp. 361–371, Oct. 2010.

[23] Olga Bukhtiyarova, Olga Bukhtiyarova, Sylvain Chauvette, Sylvain Chauvette, Josée Seigneur, Josée Seigneur, Igor Timofeev, and I. Timofeev, "Brain States in Freely Behaving Marmosets," *Sleep*, May 2022.

[24] G. Santhanam, M. D. Linderman, V. Gilja, A. Afshar, S. I. Ryu, T. H. Meng, and K. V. Shenoy, "HermesB: A Continuous Neural Recording System for Freely Behaving Primates," *IEEE Transactions on Biomedical Engineering*, vol. 54, pp. 2037–2050, Nov. 2007.

[25] L. A. Toth and P. Bhargava, "Animal Models of Sleep Disorders," *Comparative Medicine*, vol. 63, pp. 91–104, Apr. 2013.

[26] D. A. Washburn and R. S. Astur, "Exploration of virtual mazes by rhesus monkeys (Macaca mulatta)," *Animal Cognition*, vol. 6, pp. 161–168, Sept. 2003.

[27] Q. Yang, Z. Lin, W. Zhang, J. Li, X. Chen, J. Zhang, and T. Yang, "Monkey plays Pac-Man with compositional strategies and hierarchical decision-making," *eLife*, vol. 11, p. e74500, Mar. 2022.

[28] A. Zhou, S. R. Santacruz, B. C. Johnson, G. Alexandrov, A. Moin, F. L. Burghardt, J. M. Rabaey, J. M. Carmena, and R. Muller, "A wireless and artefact-free 128-channel neuromodulation device for closed-loop stimulation and recording in non-human primates," *Nature Biomedical Engineering*, vol. 3, pp. 15–26, Jan. 2019.