Master thesis on Cognitive Systems and Interactive Media
Universitat Pompeu Fabra

Assessing spatial orientation in the human hippocampus during virtual navigation

Marina Estévez Almenzar

Supervisor: Diogo Santos-Pata

July 2020

Universitat Pompeu Fabra
Barcelona

Master thesis on Cognitive Systems and Interactive Media
Universitat Pompeu Fabra

Assessing spatial orientation in the human hippocampus during virtual navigation

Marina Estévez Almenzar

Supervisor: Diogo Santos-Pata

July 2020

Contents

1 Introduction **1**

1.1 Hypothesis 2

2 Methods **3**

2.1 General Methodology Applied 3

2.1.1 LFP dataset 4

2.1.2 Experimental dataset 4

2.2 Experimental set-up 5

2.3 Procedures used to obtain data and results 5

2.3.1 Procedures used to obtain data 5

2.3.2 Procedures used to obtain results 7

3 Results and analysis **11**

3.1 Channels' contribution on orientation decoding 12

3.2 Frequencies' contribution on orientation decoding 13

4 Discussion and Conclusion **15**

List of Figures **16**

List of Tables **17**

Bibliography **18**

A Experimental dataset **20**

Abstract

Neural decoding uses electrophysiological signals recorded from the brain to find identifications between these signals and behavioral variables in the outside world. These decoding predictions can be used to better understand how areas of the brain relate to the observable states. In essence, neural decoding is a regression (or classification) problem relating neural signals to particular variables [1]. In this project, we attempt to identify human hippocampal local field potential (LFP) with spatial orientation during virtual navigation, guided by some important results from the literature such as the apparent relation between LFP and information about speed and location in rats [2], the identification between entorhinal cortex and information about head direction, self-location, and speed in rats [3], [4], the relation between thalamus and subiculum with 3D head direction [5] or the association of hippocampal theta oscillations to mnemonic processes [6]. To address this proposal, we use an adaptation of the convolutional neural network proposed by Frey et al. [7] that can be found here. Our results suggest that spatial orientation is encoded in human hippocampal LFP during virtual navigation.

Keywords: Hippocampus; Spatial orientation; Local field potential; Neural networks

Chapter 1

Introduction

Deciphering the neural code, despite being a complex task, has been longly studied as a central aim of Neuroscience. The understanding of the neural representation through electrophysiological signals and their linked behaviors has reached success with significant studies such as the cat experiment of Hubel and Wiesel (Hubel and Wiesel, 1959, [8]) and the experiment with rats with injured hippocampus by O'Keefe and Dostrovsky (O'Keefe and Dostrovsky, 1971, [9]). Although these two experiments inaugurated a list of successful researches in the field, the progress has been slow.

One of the reasons is that neural code is commonly sparse: one single region of the brain may represent several variables at the same time, in different extents, by the activation of different neural populations in this region. The behavioral and the sensory stimuli data linked to these noisy neural signals could be considered to be low dimensional, but what it is really complex is to find the correspondence between the neural representations of sensory features and its resulting observable states. Throughout these years, and yet recently, this identification has been carried out by the recognition capacity of one or several experimenters, guided by existing knowledge. Researchers attempted to identify which factors can be mapped, and some important results were achieved by using this method, such as the identification between local field potential (LFP) and information about speed and location in rats (McFarland, 1975 [2]). More recently, entorhinal cortex electrophysiological recording and the encoded information about head direction, self-location, and movement speed have also been identified in rats (Sargolini et al., 2016 [3]; Hafting et al., 2017 [4]).

Kim and Maguire, in 2018, by using functional magnetic resonance imaging (fMRI) multivoxel pattern similarity analysis, found evidence that the thalamus, particularly the anterior portion, and the subiculum encoded the horizontal component of 3D head direction (azimuth) [5].

In 2008, Lisman et al. reviewed evidence that theta and gamma oscillations work together to form a neural code that provides a way for multiple neural ensembles to represent an ordered sequence of items. In the hippocampus, this coding scheme is utilized during the phase precession, a phenomenon that can be interpreted as the recall of sequences of items (places) from long-term memory. The same coding scheme may be used in certain cortical regions to encode multi-item short-term memory [6].

1.1 Hypothesis

We have reviewed several identifications between hippocampal neural code and observable behavioral states related to spatial location, particularly to orientation. However, the identification between human hippocampal oscillations and spatial orientation is still unclear. Here we hypothesize that spatial orientation is encoded in human hippocampal local field potential.

Chapter 2

Methods

2.1 General Methodology Applied

The participant is a woman. She is 37 years old. She is a drug resistant epilepsy patient. She has 10 electrodes (a total of 125 contact points) place inside the skull. No more information about the participant was given to us.

The participant was asked to explore a virtual environment and find an object target. The experiment was split into 3 phases, every of which was composed of several trials. Every trials lasted between 5 and 37 seconds, depending on the difficulty level of the environment and on how well the participant's performance was. The difficulty of the environment depended on the number of clues included in it. A clue can be either local or global. A local clue in an specific trial is another object different to the target object, also located in the environment, that could help the participant to either orientating herself or remembering the location of the target. These local clues are different and are randomly located for every trial. A global clue is a bigger and static object such as a building, set at the same location for every trial. The possible target objects, as well as the possible local clues, were common objects such as plants, food, animals...

An image of the target object was presented to the participant at the beginning of every trial. There was a total of three different virtual environments and six different target objects.

During phase 1 all the objects were visible for every trial, so the participant could explore the environment and learn the objects location. This phase lasted a total of 24 trials (4 trials per target). During phase 2 the experimenters attempted to test whether the participant was able to recall the location of the objects. The target object wasn't visible in this phase, and the participant was asked for moving to the location of that specific target object and to press a button when she considered she was on top of the its location. This phase lasted a total of 48 trials (8 trials per target). During phase 3 all the local cues were removed from the environment. The aim of this phase was the same as in the previous phase, but with the difficulty of not having clues in the environment. This phase lasted 48 trials, and the search procedure remained as in the previous phases.

The experimenters provided us with two datasets: the local field potential (LFP) dataset and the experimental dataset.

2.1.1 LFP dataset

The LFP dataset (see Appendix B) contains the name, coordinates, group and ASEG region for every contact point (a total of 125 contact points), for every of which we were given the recording of the electrophysiological signal. These signals were recorded with a sampling rate of 512 Hz.

2.1.2 Experimental dataset

The experimental data contains the information about the virtual position and orientation recorded during the task. We also were given data about specific information about the trials, such as times and the kind of target object involved in every trial (see Appendix A). This experimental data was recorded with a frequency rate lower than the LFP's frequency, so we needed to do an interpolation between both signals, as we will see in section 2.3.1.

2.2 Experimental set-up

Electrophysiological recordings were made using a standard clinical EEG system (XLTEK, subsidiary of Natus Medical) at a sampling rate of 512 Hz. A unilateral implantation was performed in the patient, using from 10 intracranial electrodes (Dixi Medical, Besançon, France, diameter: 0.8 mm, 7 to 15 contacts, 2 mm length, with a space between contact center of 1.5 mm) stereotactically inserted using robotic guidance (ROSA, Medtech Surgical , Inc., Montpellier, France). Electrodes were implanted by the clinical team at Hospital del Mar (Barcelona, Spain). Targeted regions were established for clinical reasons to localize the epileptogenic area, including the left hemisphere hippocampus in the case of our participant. Electrodes were localized using postimplanation structural T1-weighted MRI scans.

2.3 Procedures used to obtain data and results

From now on, by data we will understand the input of the algorithm that we will be using to obtain the results (see section 2.3.2), and by results we will understand the output of this algorithm (see Chapter 3).

2.3.1 Procedures used to obtain data

Signal pre-processing

We first computed a zero-phase band-pass filter over the raw LFP channels (see Figure 1) to minimize the influence of the frequencies outside the range from 2 to 200 Hz. Then we notched and bipolarized the signals by subtracting the associated white-matter signal in order to get rid of the background noise (see Figure 2). We then interpolated the LFP signal and the experimental data for every trial (see Figure 3) in order to establish an only common sampling rate of 512 Hz.

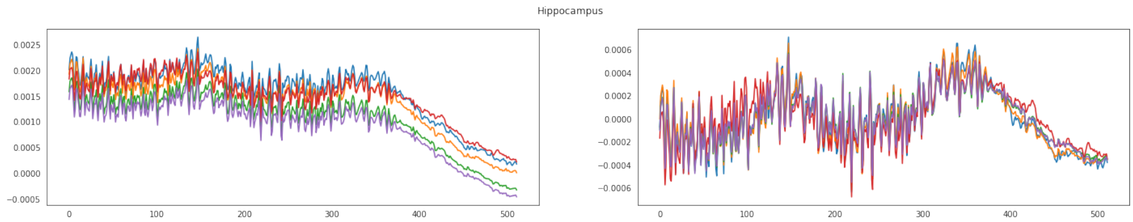

Figure 1: Zero-phase band-pass filter

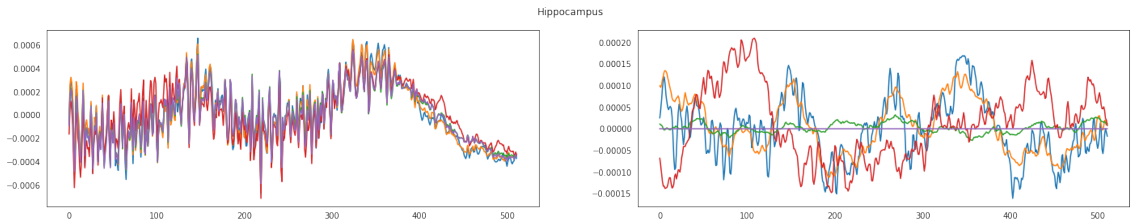

Figure 2: Bipolarizing signal after notched it

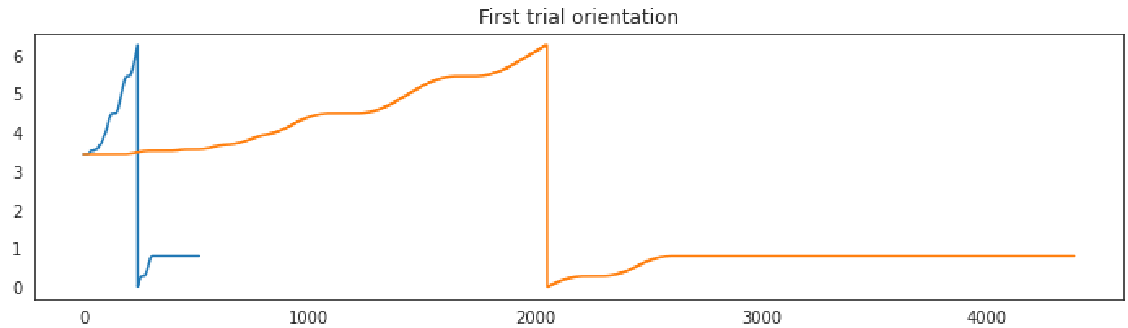

(a) Orientation signal corresponding to first trial before (in blue) and after (in orange) interpolation

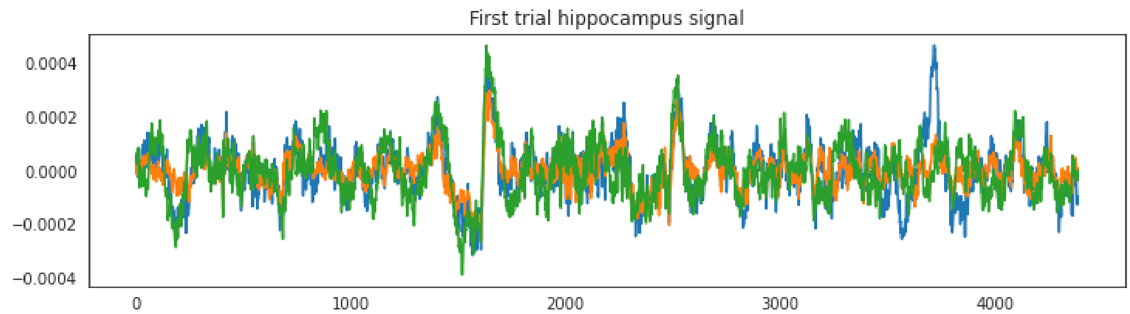

(b) Hippocampus signal corresponding to first trial

Figure 3: Interpolation process: Orientation signal before and after interpolating it with its corresponding electrophysiological signal

2.3.2 Procedures used to obtain results

To justify the procedures that we have used to obtain our results, we need to remember the point we were mentioning in Section 1, when we talked about the common methodology used to find identifications between neural code and observable behavioral states. This traditional methodology is based on the recognition capacity of one or several experimenters who, guided by existing knowledge, attempted to identify which neural factors can be mapped. This methodology necessarily favors the identification of the simplest and less noisy neural representations, which could make us think that some difficult to identify but still essential neural data is being either deprecated or not even considered.

Artificial Neural Networks (ANNs) suggest an alternative way of dealing with this task. These algorithms have proved in several cases that human performance on tasks similar to our task, such as image classification, can be exceeded (Krizhevsky et al., 2012 [10]; Mnih et al., 2015 [11]). They have also been successfully applied to the decoding process of behavioral states from observing neural firing rates, reaching very good results (Glaser et al., 2017 [1]; Tampuu et al., 2018 [12]). However, this accuracy is not the only (and maybe not the most important) advantage of ANNs: these networks also make some assumptions about the structure and relevance of the input data that evolve over the training steps of the algorithm. Then, in addition to using the trained network to make predictions based on new input, the results obtained can be analyzed to identify which elements of the input are more significant for the task itself [7].

Frey et al. [7], to test this proposal, developed a convolutional neural network able to take wide-band neural data as input, and offer predicted continuous regression variables as output. Their model differs from conventional decoding methods which typically use Bayesian estimators (Zhang et al., 1998 [13]) together with highly processed neural data. Bayesian decoding has as an important component a prior accurate calculation of expectations based on the way in which the neural data varies with the decoded variable, requiring considerable knowledge about the neural signal

[7]. In contrast, Frey et al. worked on a flexible and multi-purposes approach that requires few assumptions. The authors trained this model with electrophysiological recordings made from the CA1 pyramidal cell layer in freely foraging rodents. The network successfully recognized place cells (O'Keefe and Dostrovsky, 1971, [9]) and that theta-band oscillations in the LFP were informative about running speed (McFarland, 1975 [2]).

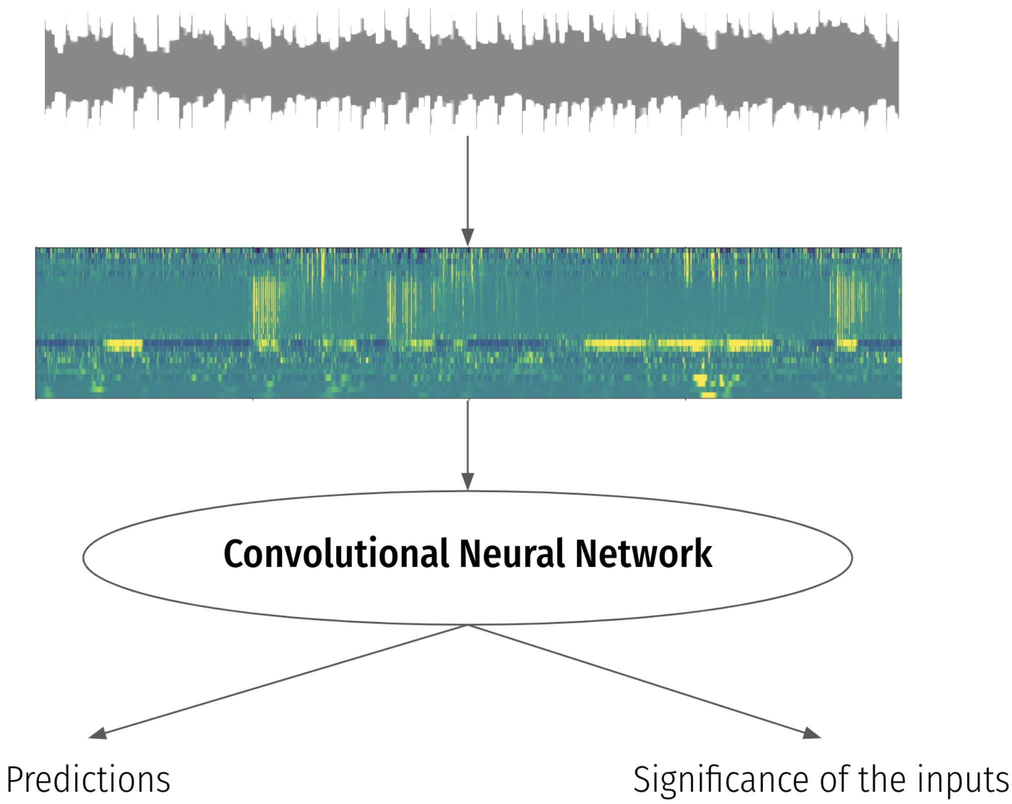

Figure 4: Diagram of the network with its inputs and outputs

We obtained our results from the execution of an adaptation of the Convolutional Neural Network proposed by Frey et al. in 2019[1]. This adaptation can be found here. We adapted it to deal with human neural data and test its accuracy, as well as to analyze the result obtained in order to find identifications between human neural

[1]https://github.com/CYHSM/DeepInsight

representations of sensory features and behaviors. We also implemented the option of visualizing the comparison of channels influence, as it is done in the original code for the visualization of frequencies comparison. As you can see in Figure 4, the network takes a three-dimensional wavelet transformed signal as input, together with the associated behavioral states, that is, the desired output (in our case, the orientation data). After training the network we obtain a model capable of making predictions about orientation based on new electrophysiological data and, what is more important for our project, a model capable of giving us information that, after being analysed, can tell us about the significance of different aspects of our input signal.

The convolutional neural network

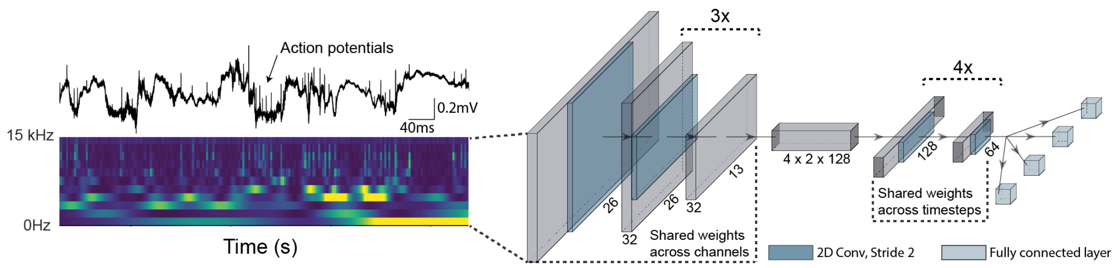

Figure 5: Network architecture. Frey et al., 2019 [7]

The model takes a three-dimensional wavelet transformed signal as input and uses convolutional layers connected to a regression head to decode continuous behaviour. It uses a kernel size of 3 throughout the model and keep the number of filters constant at 64 for the first 8 layers while sharing the weights over the channel axis, while then doubling them for each following layer. For downsampling the input it uses a stride of 2, intermixed between the time and frequency dimension (see Figure 5). It uses 2D-convolutions on a three-dimensional input to lower computational load and GPU memory consumption but also to improve generalization. To achieve this, it shares the weights across the channel dimension for the first 8 convolutions and across the time dimension for the last 6 convolutional layers. Sharing the weights across channels prevents overfitting to any channel specific feature, making sure there exists a global representation e.g. of spikes or other prominent oscillations in the local field potential [7].

Selection of significant contact points

We selected a total of 9 contact points out of 125 provided by the experimenters. From this 125 contact points, 4 belong to the left hippocampus, 1 belong to the left amygdala, 35 belong to left cerebral cortex, 1 belong to the choroid plexus and the rest belong to white matter. We finally selected 3 out of the 4 available contact points from the hippocampus after checking that one of the signals was too uniform compared to the rest of hippocampal signals (see Figure 1 right, green signal). We also selected all the available contact points from amygdala and choroid plexus, and 4 random contact points out of the 35 available ones from the cerebral cortex in order to balance the comparison between cerebral cortex and hippocampus. We selected the contact point from choroid plexus as a control variable since we did not expect to obtain any contribution to orientation decoding from this region.

We run the network in two different ways with two different purposes: we first used the network to obtain results with regard to the contribution of every channel (signals coming from the different contact points) on the decoding of orientation. We selected the most significant set of channels belonging to the same brain region and we run the network again with this selection as a new input in order to obtain results with regard to the contribution of every specific frequency of the signal on the decoding of orientation.

Chapter 3

Results and analysis

Although our intention is not to look for any significant conclusion from observing the behavioral data obtained from the participant performance during trials, it is interesting to have a look at them in order to envision a general idea about the participant's performance in terms of, for instance, time and capacity of exploration. In Figure 6a we can observe the evolution of the participant's position throughout the trials. From Figure 6b we can state that most of the trials lasted less than 20 seconds, and from 6c we can state that early trials took longer than late trials.

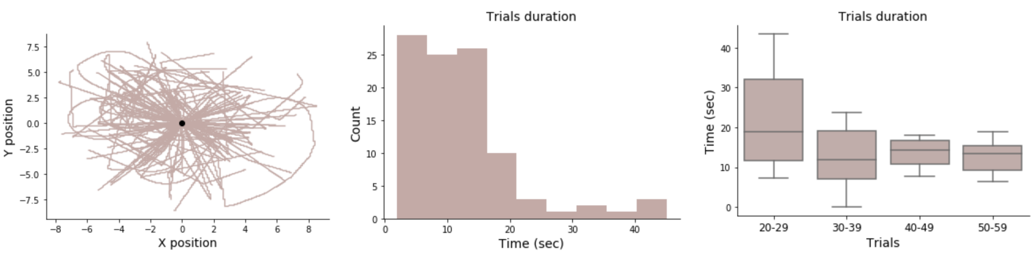

(a) Position of the participant over a number of trials.

(b) Number of trials by trial duration.

(c) Average time by trial split into groups from early to late trials.

Figure 6: Behavioral data showing the exploration capacity of the participant, the time she took for the trials.

3.1 Channels' contribution on orientation decoding

If we observe Figure 7 we can suggest that hippocampus and amygdala regions have a higher influence on decoding of orientation than cerebral cortex and choroid plexus. This last region, as we were expecting, has no any influence on the decoding of orientation process. To test the significance of the hippocampus's performance on the decoding of orientation we decided to compare it with the performance of the cerebral cortex, which also appears as an influential region but to a lesser extent. We carried out a Mann-Whitney test comparing the mean of the influences for every of these regions. As you can see in Figure 8, the mean of the influence of signals that come from hippocampus is significantly higher than the mean of the influence of signals that come from cerebral cortex. This, together with the fact that the mean for the hippocampus has been computed with just three signals and the mean for the cerebral cortex has been computed with four signals, suggests that hippocampal LFP encodes spatial orientation, which supports our hypothesis.

This does not mean that either the cerebral cortex or the amygdala is not influential on the encoding of orientation; this just means that the hippocampus influences it in a significantly higher way. The same analysis could be done comparing hippocampus and amygdala, but it would only make sense in case more amygdala contact points were available to establish a balanced comparison.

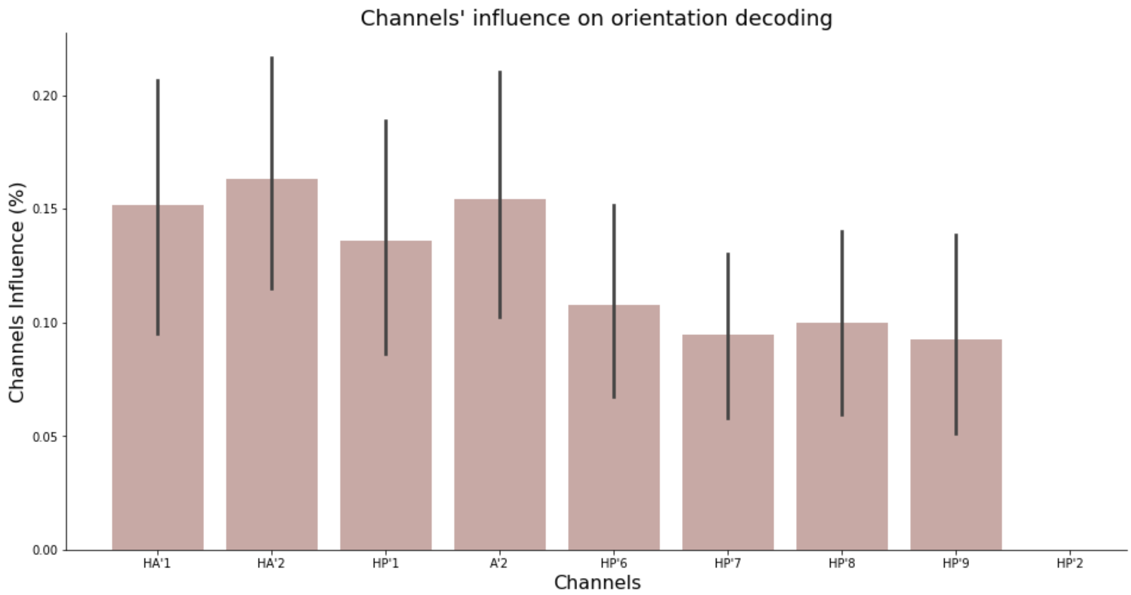

Figure 7: Channels' contribution on the decoding of orientation process. First three channels correspond to the hippocampus, the fourth to the amygdala, the next four channels correspond to the cerebral cortex and the last one correspond to the choroid plexus.

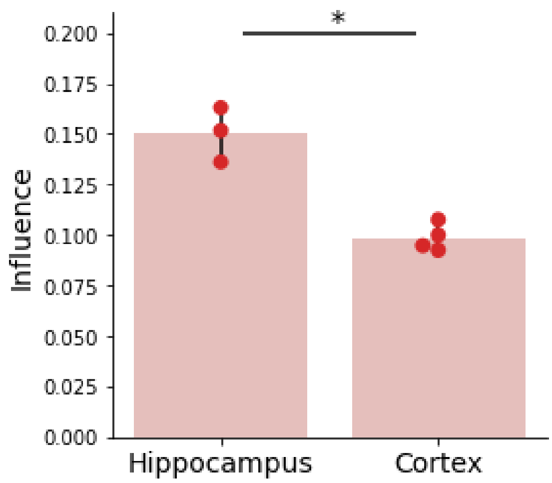

Figure 8: Mann Whitney U test comparing hippocampus and cerebral cortex associated influence on the decoding of orientation (statistic = 0.0, p-value = 0.0259).

3.2 Frequencies' contribution on orientation decoding

After deducing that the role of the hippocampus is significantly influential on the decoding of orientation, as we explained before, we can now use the same Convolutional Neural Network in order to obtain results in relation to how every specific frequency involved in hippocampal LFP is influencing the decoding of orientation. As you can see in Figure 9, most influential frequencies are the ones going from 4.0 to 8.0 Hz, approximately, that correspond to theta oscillations. This is in line with Lisman et al., who suggested that theta oscillations from hippocampus were involved in mnemonic processes such as integration of spatio-temporal information. We carried out another Mann Whitney U test in order to check this interpretation. As we can see in Figure 10, the mean of the influence of hippocampal theta oscillations is significantly higher than the mean of the influence of the others oscillations, including beta, delta, and gamma oscillations.

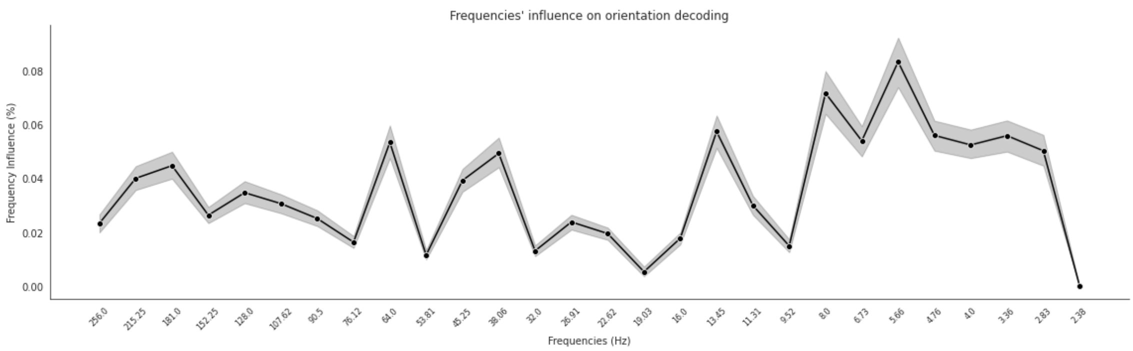

Figure 9: Frequencies' contribution on the decoding of orientation process

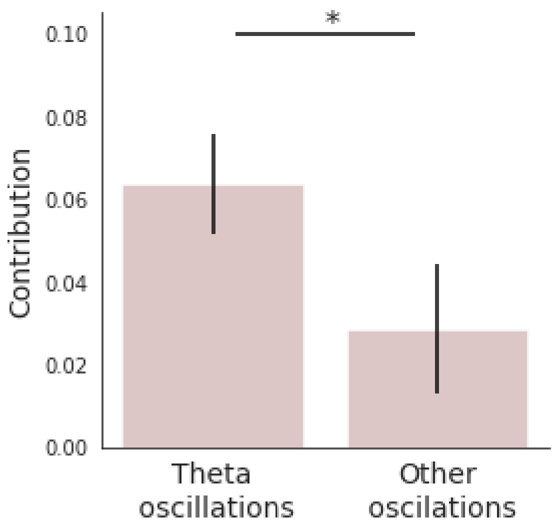

Figure 10: Mann Whitney U test comparing the influence of theta oscillations from hippocampus and the influence of the rest of oscillations on the decoding of orientation (statistic = 4.0, p-value = 0.0008).

Chapter 4

Discussion and Conclusion

We have obtained significant results that suggest that spatial orientation is encoded in human hippocampal local field potential, as we hypothesized. However, more extensive studies should be carried out to support our conclusion. It is essential to start by collecting a wider amount of electrophysiological data in order to elaborate a dataset with ideally a larger amount of signals coming from hippocampus and amygdala. This would allow us to improve the accuracy of our results about hippocampus and also to extend our work to the amygdala, which also appeared as a promising source of good results, as we observed in Figure 7. Also, the possibility of having access to more participant would allow us to compare the results obtained between participants, which would result in more accurate and detailed results.

We also state as a future step the improvement of the network either by adjusting more parameters or increasing the number of epochs in order to extend our results to the possibility of obtaining good predictions in terms of position and orientation when using new electrophysiological signals as input.

List of Figures

1 Zero-phase band-pass filter 6
2 Bipolarizing signal after notched it 6
3 Interpolation process: Orientation signal before and after interpolating it with its corresponding electrophysiological signal 6
4 Diagram of the network with its inputs and outputs 8
5 Network architecture. Frey et al., 2019 [7] 9
6 Behavioral data showing the exploration capacity of the participant, the time she took for the trials. 11
7 Channels' contribution on the decoding of orientation process. First three channels correspond to the hippocampus, the fourth to the amygdala, the next four channels correspond to the cerebral cortex and the last one correspond to the choroid plexus. 13
8 Mann Whitney U test comparing hippocampus and cerebral cortex associated influence on the decoding of orientation (statistic = 0.0, p-value = 0.0259). 13
9 Frequencies' contribution on the decoding of orientation process . . . 14
10 Mann Whitney U test comparing the influence of theta oscillations from hippocampus and the influence of the rest of oscillations on the decoding of orientation (statistic = 4.0, p-value = 0.0008). 14

List of Tables

1 Extract of the experimental dataset. Time in seconds, *Activation* in seconds, *TTL start markers* in samples, and *Game started* in samples 21

2 Extract of the LFP dataset 22

Bibliography

[1] Glaser, J. I., Chowdhury, R. H., Perich, M. G., Miller, L. E. & Kording, K. P. Machine learning for neural decoding. *arXiv preprint arXiv:1708.00909* (2017).

[2] McFarland, W. L., Teitelbaum, H. & Hedges, E. K. Relationship between hippocampal theta activity and running speed in the rat. *Journal of comparative and physiological psychology* **88**, 324 (1975).

[3] Sargolini, F. *et al.* Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science* **312**, 758–762 (2006).

[4] Hafting, T., Fyhn, M., Molden, S., Moser, M.-B. & Moser, E. I. Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**, 801–806 (2005).

[5] Kim, M. & Maguire, E. A. Encoding of 3d head direction information in the human brain. *Hippocampus* **29**, 619–629 (2019).

[6] Lisman, J. & Buzsáki, G. A neural coding scheme formed by the combined function of gamma and theta oscillations. *Schizophrenia bulletin* **34**, 974–980 (2008).

[7] Frey, M. *et al.* Deepinsight: a general framework for interpreting wide-band neural activity. *bioRxiv* 871848 (2019).

[8] Hubel, D. H. & Wiesel, T. N. Receptive fields of single neurones in the cat's striate cortex. *The Journal of physiology* **148**, 574 (1959).

[9] O'Keefe, J. & Dostrovsky, J. The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain research* (1971).

[10] Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105 (2012).

[11] Mnih, V. *et al.* Human-level control through deep reinforcement learning. *nature* **518**, 529–533 (2015).

[12] Tampuu, A., Matiisen, T., Ólafsdóttir, H. F., Barry, C. & Vicente, R. Efficient neural decoding of self-location with a deep recurrent network. *PLoS computational biology* **15**, e1006822 (2019).

[13] Zhang, K., Ginzburg, I., McNaughton, B. L. & Sejnowski, T. J. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *Journal of neurophysiology* **79**, 1017–1044 (1998).

Appendix A

Experimental dataset

Table 1: Extract of the experimental dataset. Time in seconds, *Activation* in seconds, *TTL start markers* in samples, and *Game started* in samples

Trial	**Time**	**Target type**	**Target ID**	**Activation of movement**	**TTL start markers**	**Game started**
0	22.46254	Food	3	25.48104	416784	418329
1	38.10352	Crystal	2	41.10556	429369	430906
2	62.66492	Object	5	65.68336	439724	441269
3	82.89115	Plant	3	85.90936	447810	449355
4	98.68200	Music	2	101.69970	463554	465099
...	...	...	...	...	...	...
117	2514.64400	Crystal	2	2517.66300	1690601	1692146
118	2525.89900	Food	3	2528.90000	1698652	1700188
119	2541.60700	Plant	3	2544.62500	1703715	1705260

Appendix B

LFP dataset

Table 2: Extract of the LFP dataset

name	x	y	z	group	ASEG region
A'2	-25.228381	16.345159	-3.865484	A'	18
A'3	-28.456763	17.690319	-3.730968	A'	3
A'4	-31.685146	19.035478	-3.596452	A'	2
A'5	-34.913528	20.380637	-3.461936	A'	2
...	...	...	...	...	...
HA'1	-22.999999	5	-3	HA	17
HA'2	-26.422901	4.543613	-2.429516	HA'	17
HA'3	-29.845803	4.087226	-1.859032	HA	17
HA'4	-33.268705	3.630839	-1.288549	HA	2
...	...	...	...	...	...
HP'1	-20.999999	-18	4	HP	17
HP'2	-24.426656	-18.685332	3.804191	HP'	31
HP'3	-27.853313	-19.370663	3.608382	HP	2
HP'6	-38.133285	-21.426658	3.020955	HP	3
HP'7	-41.559942	-22.111989	2.825146	HP	3
HP'8	-44.986599	-22.79732	2.629337	HP	3
...	...	...	...	...	...