

Aurélié Cordier

**Master Nanotech - MNIS
Academic year 2019/2020**

In Kreiman Laboratory, Harvard Medical School, Boston (US)

Recognition of minimal images in the human brain

from 1/03/20 to 28/08/20

Confidential

Under the supervision of :

- **Antonino CASILE, PhD** (Antonino_Casile@hms.harvard.edu)
and **Gabriel KREIMAN, PhD** (Gabriel.Kreiman@childrens.harvard.edu)
- **Phelma Tutor: Prof. Panagiota MORFOULI** (Panagiota.Morfouli@grenoble-inp.fr)
- **EPFL Tutor: Prof. Silvestro MICERA** (Silvestro.Micera@epfl.ch)

**Ecole nationale
supérieure de physique,
électronique, matériaux**

Phelma

Bât. Grenoble INP - Minatec
3 Parvis Louis Néel - CS 50257
F-38016 Grenoble Cedex 01

Tél +33 (0)4 56 52 91 00
Fax +33 (0)4 56 52 91 03

<http://phelma.grenoble-inp.fr>

**Ecole Polytechnique Fédérale
de Lausanne**
Section of Electrical and
Electronic Engineering

EPFL
Route Cantonale,
1015 Lausanne, Switzerland

Tél +41 (0)21 693 46 18

<https://www.epfl.ch>

ACKNOWLEDGMENTS

First of all, I would like to express my gratitude towards my advisor Antonino Casile who offered me the special opportunity of working with his team. His supervision, although the initial project could not be completed, was very helpful and reassuring.

Special thanks go to Gabriel Kreiman, who kindly hosted me in his lab and offered me guidance and advice all along this project. Our weekly virtual meetings, along with lab meetings, virtual seminars and talks taught me a lot of interesting things and helped me make the best of this very valuable experience. The entire group at Kreiman lab was always welcoming and I am grateful for our virtual interactions that provided me with all the help when I needed it.

I would also like to thank my Phelma advisor, Panagiota Morfouli, for her supervision throughout these six months. I am thankful for the entire Phelma team including my masters coordinator, Liliana Prejbeanu, that helped make this project happen in spite of a very special academic year.

Finally, I would like to thank my EPFL advisor, Silvestro Micera, who kindly agreed to oversee my thesis and made possible this opportunity to experience research work here in Boston.

TABLE OF CONTENTS

<i>Aknowledgments</i>	<i>i</i>
<i>Abstract</i>	<i>iii</i>
<i>Résumé</i>	<i>iv</i>
<i>Astratto</i>	<i>v</i>
<i>Glossary</i>	<i>vi</i>
<i>Introduction</i>	<i>1</i>
<i>Motivation and framework</i>	<i>2</i>
<i>Background</i>	<i>4</i>
The human visual system	4
Neurophysiological recordings	7
<i>Material and methods</i>	<i>9</i>
Subjects	9
Intracranial Field Potential recordings	9
Stimulus presentation and task	11
Image selection	11
Pre-processing of the data	14
Further processing	14
<i>Data analyses</i>	<i>15</i>
Relabeling	15
Visual responsivity	15
Latency measurements	17
Comparisons	18
<i>Results</i>	<i>21</i>
Behavioral recognition	21
Neural responses and visual responsivity	23
Visual selectivity	24
Timings	25
Comparisons	27
<i>Discussion</i>	<i>29</i>
<i>Summary</i>	<i>30</i>
<i>References</i>	<i>32</i>

ABSTRACT

Visual object recognition is performed without effort by humans even though it requires a series of complex computations which are, for now, not well understood. This study introduces the concept of minimal images, smallest configurations where an image is recognizable to the human vision, to study the processes by which the brain uses visual features to carry out computations underlying visual recognition. Crucially, the role of these visual features is revealed at the minimal level and a tiny change in the image configuration is enough to completely lose recognition. A neurophysiological experiment was conducted with twelve subjects implanted with intracranial electrodes. Visual representations elicited by minimal as well as sub-minimal images could be observed, and category-selective responses could be discriminated. Although the two image conditions did not result in distinguishable neural features, the results seem to endorse previous observations regarding behavior and sensitivity to perceptual discrimination.

RÉSUMÉ

La reconnaissance d'objets est conduite sans efforts par le cerveau humain, bien que cela requiert une série de calculs complexes qui ne sont, jusqu'à aujourd'hui, que très peu connus. Cette étude introduit le concept d'image minimale, plus petite configuration pour laquelle une image est reconnaissable par le système visuel, pour étudier les procédés par lesquels le cerveau utilise des caractéristiques visuelles pour résoudre la reconnaissance visuelle d'un objet. Le rôle de ces caractéristiques visuelles est spécialement révélé au niveau "minimal", et un changement infime dans la configuration de l'image a pour conséquence la perte immédiate de reconnaissance. Une expérience a été conduite sur douze sujets implantés avec des électrodes intracrâniennes. Des réponses visuelles induites par les images minimales et sous-minimales ont pu être observées au niveau de populations neuronales, ainsi que des réponses sélectives à différentes catégories d'images. Bien que les deux configurations d'images n'ont pas résulté en des caractéristiques neuronales distinguables, les résultats de l'étude semblent confirmer certaines observations précédentes à propos de tendances comportementales et de la sensibilité à la discrimination perceptuelle.

ASTRATTO

Il riconoscimento degli oggetti è condotto senza sforzo dal cervello umano, anche se questo richiede una serie di calcoli complessi che, per ora, non sono ben compresi. Questo studio introduce il concetto di immagine minimale, la configurazione più piccola per la quale un'immagine è riconoscibile dal sistema visivo, per studiare i processi con cui il cervello utilizza le caratteristiche visive per risolvere il riconoscimento visivo di un oggetto. Il ruolo di queste caratteristiche visive è rivelato a livello "minimo" e un piccolo cambiamento nella configurazione dell'immagine comporta una perdita immediata del riconoscimento. È stato condotto un esperimento su dodici soggetti impiantati con elettrodi intracranici. Le risposte visive indotte da immagini minime e sub-minime potrebbero essere osservate a livello di popolazioni neurali, nonché risposte selettive a diverse categorie di immagini. Sebbene le due configurazioni dell'immagine non si traducono in caratteristiche neurali distinguibili, i risultati dello studio sembrano confermare alcune osservazioni precedenti sulle tendenze comportamentali e sulla sensibilità alla discriminazione percettiva.

GLOSSARY

MIRC : Minimally Recognizable Configuration

subMIRC : **Sub**-Minimal Recognizable Configuration

EEG : Electroencephalography

sEEG : Stereo-Electroencephalography

ECoG : Electrocorticography (also iEEG: intracranial EEG)

LFP : Local Field Potentials

fMRI : functional Magnetic Resonance Imaging

MEG : Magnetoencephalography

INTRODUCTION

The placement took place in the Kreiman laboratory affiliated with Boston Children’s Hospital and Harvard Medical School in Boston (USA) and lasted 24 weeks. The Kreiman laboratory is a leading team of experts in the field of computational neuroscience and neurophysiology. The main focus of their research lies in visual recognition, learning and memory. They are specifically interested in understanding how neural circuits carry out computations in our brains and further developing biologically-inspired artificial intelligence. Combining behavioral studies, invasive neurophysiological recordings and computational models, their research leads to many publications on topics such as visual search, pattern completion, real-life memories as well as building models and machines that are able to see and interpret the world like humans.

The purpose of this project is to study visual recognition of reduced images at a physiological level in the human brain. The work is based on a study carried out by Ullman et al. in 2016 where they introduce the concept of so-called “minimal images” to investigate the differences between humans and computational models regarding image recognition [1]. The project focuses on the recognition of these minimal images at a physiological level, analyzing data collected by invasive neurophysiological recordings (ECoG and sEEG) in 12 subjects that performed a variation of the task by Ullman et al.

MOTIVATION AND FRAMEWORK

Following decades of evidence in neuroscience, a fundamental question that remains is understanding how neural activity gives rise to perception and behavior in the physical world. The visual system turns out to be an excellent model for this purpose. The visual system is defined by the part of the central nervous system required for visual perception, including receiving, processing and interpreting visual information to build representations of our environment. Physically, this consists of the eye and retina, fibers that convey visual information, the superior colliculus and parts of the cerebral cortex [2]. The fact that around half of the non-human primate neocortex is involved at some stage in analyzing visual information tells a lot about the computational complexity of processes underlying vision throughout evolution and many questions remain yet to answer [3]. However, as the processing hierarchy for visual information appears to be quite similar to the general functional structure of the brain, an understanding of the visual system would provide crucial information on the mechanisms and processes by which the brain carries out computations inducing perception and behavior in general.

A major field of investigation in visual neuroscience is seeking to define cortical areas, building blocks of the cerebral cortex, and the way they are inter-related to create information pathways that constitute the basis of visual perception in humans and animal models. Such knowledge is critical to help people suffering from visual impairments due to cerebral cortex damage, among other things [4]. Gaining knowledge about the anatomy of visual processing stages in the brain would also allow to build more robust and biologically-relevant technologies such as artificial intelligence, sensors, cameras etc. In summary, understanding the brain computations and being able to extend and generalize such algorithms to new systems could highlight new ways of repairing broken neuronal circuits as well as to augment normally functioning circuits [5].

An interesting angle to study the structure of the visual system is through object recognition. Among all the functions of vision, recognition probably is one of the most crucial. Object recognition is essential to most tasks our brains carry out every day: from identifying objects and people faces to recognizing letters while reading, walking in the street or driving a car. We are able to recognize complex shapes in a very short amount of time, ranging around 150 ms [6][7].

Our ability to discriminate among different object categories, faces and scenes, or to distinguish among similar objects is called visual selectivity and is a key attribute of vision. Another important property of object recognition relates to the robustness to object transformations (e.g. changes in size, position, rotation, illumination). Both aspects are remarkably well handled by the human brain and we are mostly able to recognize objects even though their projections on the retina is never twice exactly the same. This robustness to transformations is, however, one of the main challenges for computational models of vision.

Although such transformations usually keep the general content of the image, single pixels can almost completely change and most computational models are not currently able to categorize objects as well

as humans do [4]. This is one of the questions computer-based visual recognition aims to address: what are the visual features and representations used by the brain (and not replicated in machines) that are critical for recognition? In recent work, neural network models of visual recognition, including biological and deep neural networks, have shown a lot of progress and have started to equalize human performance in some tasks. These models are typically trained on image samples to learn and extract features and representations and use them to categorize objects. What remains unclear, however, is whether these representations and learning processes are similar to those used by the human visual system. By introducing minimal recognizable images, it can be shown that human vision uses features and processes that are not currently used by models and that are critical for recognition [1].

The concept of minimal images was introduced by Ullman et al. in 2016 [1]. The main idea is to reduce images in size and resolution until a point where they become unrecognizable. The behavioral study carried out by Ullman et al. shows that there seem to be a sharp transition where small changes to an image make it become unrecognizable to the human vision. This phenomenon cannot be accounted for by the best computational models of vision, suggesting that the human visual system uses features and processes that current models do not use and that are critical for the task.

In this work, we want to investigate whether neural responses to minimal images can help uncover some of the mechanisms that differ between humans and machines. First, it will be useful to search for visually responsive signals and assess where and when visual responses to stimuli occur. Then we want to ask whether visual responses differ between the so-called MIRC (minimally recognizable configurations) and the so-called subMIRC (slightly modified versions of the MIRC images that are unrecognizable). A full understanding of the use of visual information and features would help shed light on the cortical mechanisms underlying visual recognition, thus possibly enhancing the ability of current models to learn from visual experience and to deal with detailed image interpretation [1].

BACKGROUND

THE HUMAN VISUAL SYSTEM

The visual system is known as the part of the central nervous system that is required for visual perception, including receiving, processing and interpreting visual information to build representations of the environment. Physically, this consists of the sensory organ that is the eye, but also fibers that convey visual information, the superior colliculus and parts of the cerebral cortex known as the visual cortex [2]. Fig. 1 (a) represents the main pathway followed by visual information from the eye to the visual cortex.

Visual pathway

Vision starts with photons impinging on photoreceptors in the retina. The light coming to the outer surface of the eye is successively reflected and refracted through the cornea-pupil-lens system which acts as a microscope and inverts the image printed on the retina with respect to the visual field.

Photoreceptors in the retina then act as transducers to convert visual information into chemical components for the body. Two different types of receptors are distinguished: rods and cones. Rods are mainly associated with movement information, including depth and slight differences in brightness. Cones are associated with the perception of colors and fine shape details [8][9]. Information flows through the retinal circuit from photoreceptors to retinal ganglion cells that produce the retinal output.

The neural signals processed by the retina are then projected via the axons of the retinal ganglion cells through the optic nerves, dividing and partially crossing over in the optic chiasm. Thus some of the visual information is directly sent towards a part of the thalamus referred to as the lateral geniculate nucleus (LGN) [10], while the remaining part of the information follows an indirect path through the superior colliculus in the midbrain where eye movements such as saccades and other motor responses are processed [11][12].

The optic chiasm, located at the base of the thalamus, is the stage where the information coming from the right (or left) visual field of both eyes is mixed and sent to the LGN in the opposite brain hemisphere. The LGN acts as a connection between the optic nerve and the primary visual cortex (V1) [13]. In the human brain, it is composed of six layers of cells called magnocellular and parvocellular cells, which purpose is to start processing rods and cones information respectively [14]. Neurons in the LGN then carry visual information to the visual cortex for further processing.

V1 transmits information to two primary extrastriate regions through a ventral stream and a dorsal stream [18][4]. This is shown in Fig. 1 (b). The dorsal stream goes through areas such as V2 and medial temporal area (MT/V5) as well as medial superior temporal. It is associated with spatial information processing, including motion, position and depth perception and is referred to as the “where/how pathway”. The ventral stream goes through visual areas V2 and V4 to the inferior temporal cortex (ITC). It is associated with object perception and representations including colors and shapes, and is often referred to as the “what pathway”, hence most interesting for object recognition studies [19].

As neural signals go further into areas of the visual cortex, more associative, complex processes take place. The final psychological and perceptual experience of vision includes some aspects of memory, expectations and predictions subserved by brain areas that are apparently unrelated to vision [10]. Two of these important areas are the prefrontal cortex and the medial temporal lobe, including the hippocampus and surrounding structures. They are connected to the last visual stages of both ventral and dorsal streams [4]. The prefrontal cortex plays a role in the moment-to-moment and task-dependent interpretation of visual input to make cognitive decisions, and the hippocampus and surrounding areas play a role in the consolidation of information through long-term memories [4][20][21][22].

Internal connections

The visual system is an extremely complex organization comprising billions of interconnected neurons. Recent work combined with decades of evidence have helped uncover features about the anatomy of interconnections in the visual cortex, but the main part of the mystery remains.

As mentioned previously, cortical areas come with a six-layers structure. Inputs and outputs of each visual area share patterns of connectivity of two types that can be distinguished: ascending “feedforward” (or bottom-up) pathways where visual information first undergoes computations in lower layers; or descending “feedback” (or “recurrent”, top-down) pathways where information originates in the upper cortical layers, higher stages of visual processing [3][5].

Although both bottom-up and top-down modulations are likely to play a role along the way of visual computations underlying object recognition, the relative contribution of bottom-up compared to top-down mechanisms in different aspects of visual object recognition is unclear. For instance, rapid categorization tasks and certain transformations of isolated objects such as scale or position changes seem to be described using purely bottom-up computational models [6][7]. However, more complex tasks like the ones by Tang et al. or Epshtein et al. involving recognition of objects from partial information seem to be too difficult of a problem for purely feedforward model architectures and may involve non-negligible contributions from recurrent and/or horizontal internal connections [23][24][25].

One way to gain knowledge into processes taking place in the visual cortex is to use timing analyses. Previous work showed that latencies in neural responses could correlate with a need for additional, possibly recurrent computations when a recognition task is complicated by some factor [23][25]. As argued by S. Ullman and his research team, top-down modulations are also likely to be involved in the visual processing of minimal images. This is endorsed by the fact that the purely feedforward computational models that they use fail to replicate human recognition performance [1]. Thus, an interesting question to ask with minimal images would be whether there are differences in latencies in the neural responses that correlate with recognition, because this could indicate that we should investigate the role of recurrent computations in the recognition process.

NEUROPHYSIOLOGICAL RECORDINGS

In order to better understand the functionality of neuronal circuits underlying visual information processing, it is necessary to examine the brain activity in response to visual stimuli at a high spatial and temporal resolution. The patients taking part in this experiment suffer from a severe form of

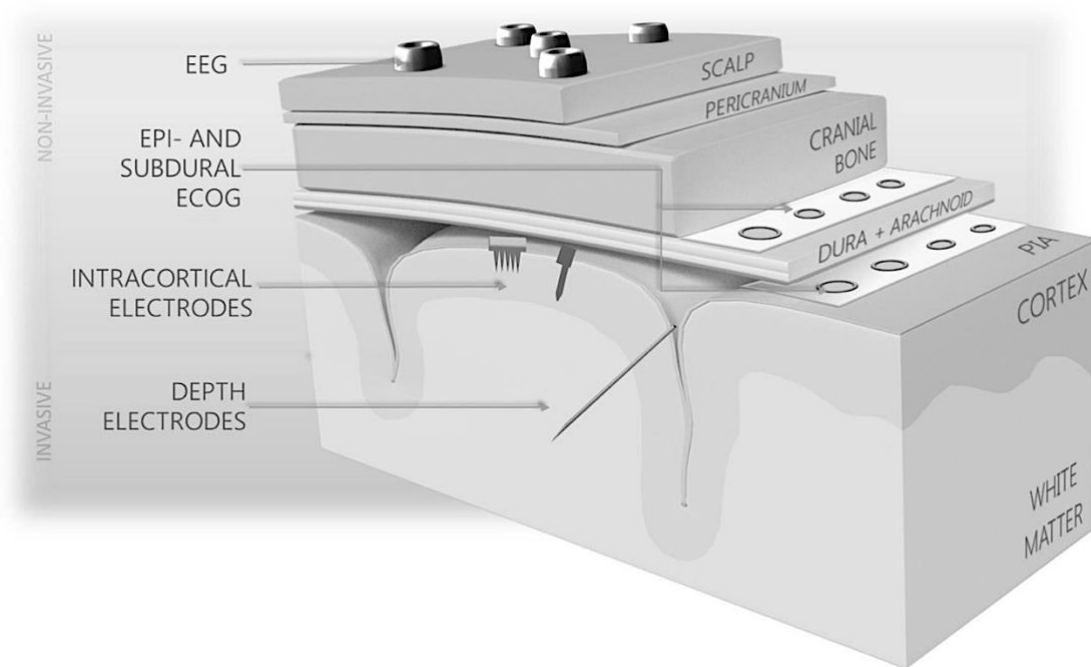


Fig. 2: Schematic depicting the various types of brain recordings (EEG, ECoG, LFP, single neuron) and their location with respect to the brain. Less invasive systems provide recordings of lower resolution compared to intracortically-implanted electrodes. [28]

epilepsy that cannot be cured by medication. As a treatment for the seizures, they have to be implanted with electrodes into their brains to help identify epileptogenic zones (i.e. brain areas generating epileptic seizures that should eventually be surgically removed to get freedom from the seizures [26]). Several recording techniques allow to achieve this. Non-invasive techniques like scalp electroencephalography (EEG) typically have a poor resolution in detecting epileptic foci due to the high resistance underlying tissues (skin, skull) blocking the electrical signal [27]. Hence other recording methods are used, allowing to implant electrodes inside the skull as pictured in Fig. 2 [28]: electrocorticography (ECoG, or intracranial electroencephalography, iEEG) and stereo-electroencephalography (sEEG, or depth electrodes)[29]. The choice depends on the clinical need of patients. In this experiment, both types of intracranial electrodes are involved.

ECoG is a type of neurophysiological recording that uses electrodes placed directly below the skull on the brain exposed surface. sEEG recordings typically involve placing electrodes at a depth inside a given brain area, as opposed to placing electrodes over the same area in ECoG recordings. Compared to conventional EEG and other non-invasive neuroimaging methods such as functional magnetic resonance imaging (fMRI) and magnetoencephalography (MEG), ECoG and sEEG recordings have the clear advantages of high spatial resolution (millimeter scale) and high temporal resolution (millisecond scale). A further asset of these intracranial recordings lies in the fact that they are not susceptible to artifact contamination from eye movements or blinks, which are known to be very detrimental to the quality of scalp EEG [30]. On the flipside, the procedure is highly invasive, especially considering ECoG electrode implantation requires a full craniotomy surgery. Thus, the implants are fully dependent on clinical needs and it should be noted that electrodes may not be associated in any way with behaviors underlying visual processing, depending on their location in the brain.

Other recording techniques such as local field potentials (LFP) recordings and single neuron recordings have been shown to provide neural responses with similar or better features [31][32]. Namely, microwire electrodes implanted extracellularly inside the cortex allow to monitor the extracellular voltage at a millisecond temporal resolution and neuronal spatial resolution. They are more selective compared to ECoG [33]. However, they are not extensively used for medical purposes in epilepsy treatments, and little remains known about the activity of single neurons in the human cortex.

MATERIAL AND METHODS

SUBJECTS

Subjects of the neurophysiological experiment were 12 patients (5 male, 11 to 43 years old) with pharmacological intractable epilepsy. They were admitted either into Boston Children's Hospital (CHB) or Johns Hopkins Hospital (JH) to localize their seizure foci for potential surgical resection.

INTRACRANIAL FIELD POTENTIAL RECORDINGS

The subjects were implanted with intracranial electrodes (Ad-Tech, Racine, WI, USA) that were arranged into either sEEG electrodes or ECoG grids and strips. All the data was collected during periods without any seizure events. Depth electrodes contained from 6 to 10 recording sites. Each subdural grid or strip contained from 4 to 64 recording sites and their layout was such that the electrode centers had 1 cm separation with one another. Each recording site was 2 mm in diameter. The number of recording sites per subject ranged from 80 to 223, for a total of 1712 electrodes. Throughout the text the recorded signal is referred to as "intracranial field potential" (IFP). This nomenclature aims to distinguish this particular type of electrophysiological recording from scalp EEG recordings or LFPs.

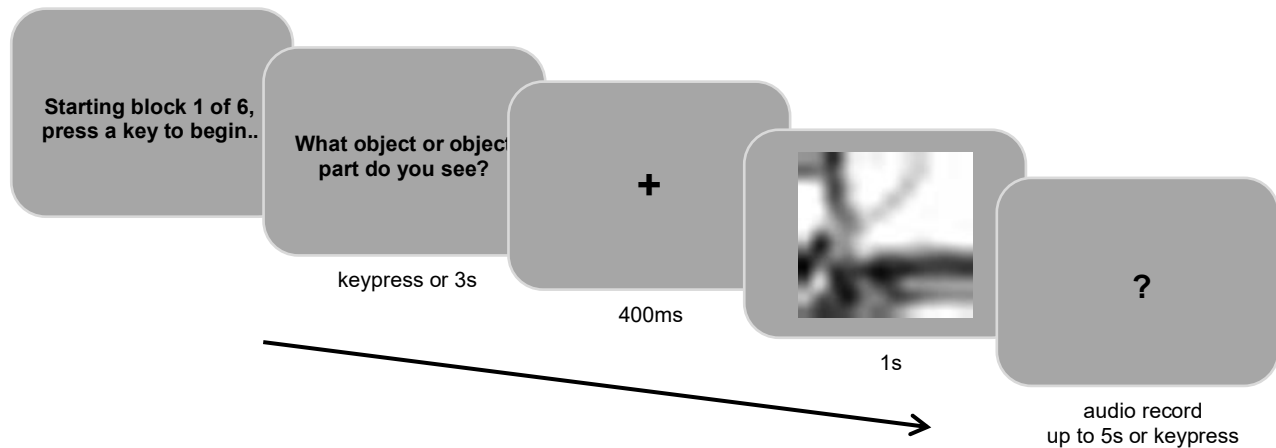


Fig. 3: Schematic depicting the timeline of the experiment. Subjects are first presented with a fixation cross for a delay period of 400 ms before stimulus presentation (1 s). They are asked to report the stimulus content in an audio recording which is then analyzed for correctness of recognition.

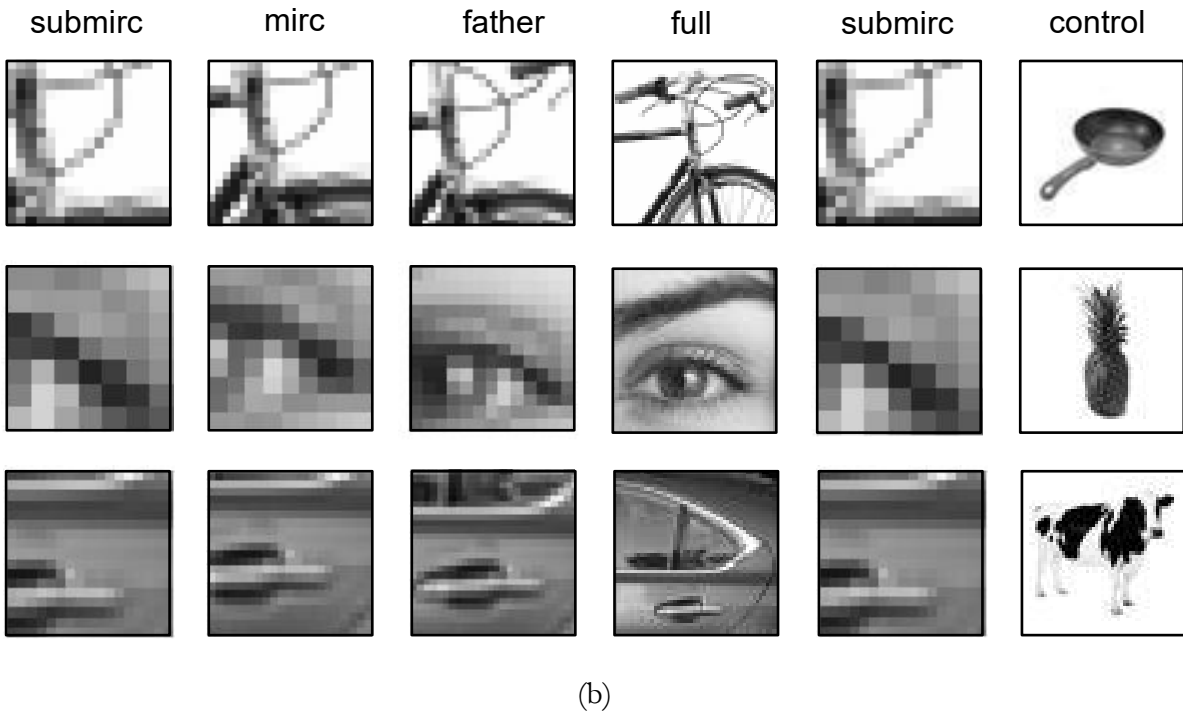
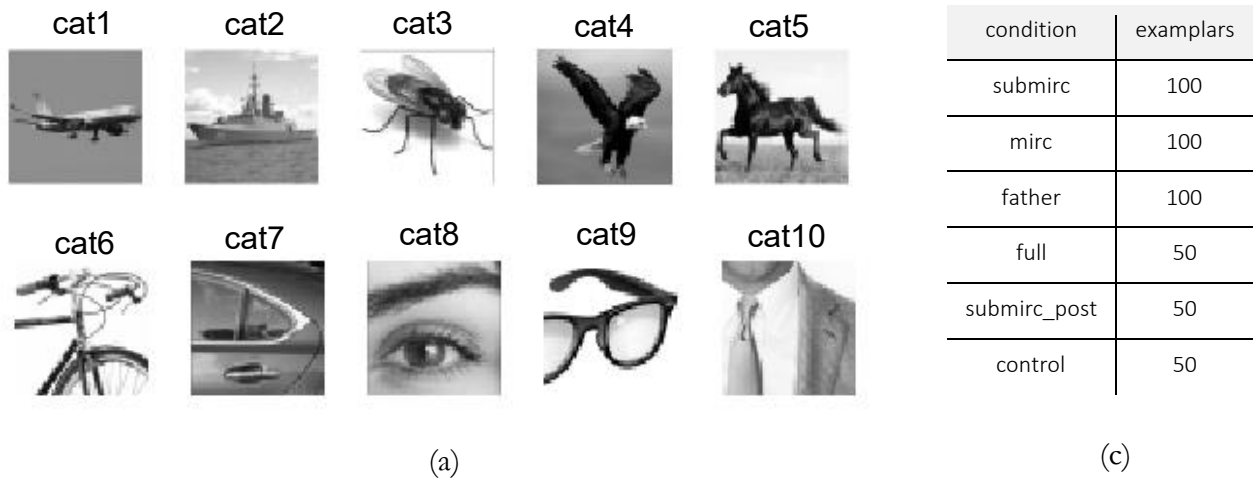


Fig. 4: Images presented during the experiment. (a) 10 image categories were shown to subjects. (b) Images were presented in six different conditions. They were presented in a specific order to avoid artifacts due to unwanted perceptual discrimination. Exemplars were shown in the order represented in the figure, from left to right. Control images were interspaced between the trials. (c) Images were not represented with the same number of exemplars.

STIMULUS PRESENTATION AND TASK

After being introduced to the task, subjects are first presented with a blank screen with a centered fixation cross for a delay period of 400 ms. This aims to provide a visual aid for subjects to keep their gaze steady and cue them to the area where the stimuli will be presented, so as to remove artifacts related to unwanted eye movements that could propagate to later stages of the visual system. After the fixation cross, an image from the dataset appears at the center of the screen for a duration of 1 s. Finally, subjects are presented with a blank screen and asked to say out loud what they recognized on the image in an audio recording. The task timeline is summarized in Fig.3. After each trial, no correct/incorrect feedback was provided.

IMAGE SELECTION

Subjects were presented with a diverse set of visual stimuli belonging to 10 different categories (bike, car, eagle, eye, eyeglasses, fly, horse, plane, ship, suit) within 4 image conditions (subMIRC, MIRC, father, full) which are shown in Fig.4 (a) and (b). Images of all conditions are presented in grayscale. SubMIRC, MIRC and father images are all descendants of the full image of the same category. An additional image condition was added, consisting of images of various objects unrelated to the 10 image categories used in the experiment. These are highly recognizable images used only for control purposes throughout the experiment. These control images have a resolution of 100×100 pixels and are presented in grayscale.

The 10 full images used in the experiment and their descendants are the same images than the ones used by Ullman et al. in their MIRC's project [1]. They are obtained through a psychophysics experiment on Amazon's online Mechanical Turk platform. The full images have a 50×50 pixels resolution. Descendants are obtained by gradually reducing the full image size or resolution by steps of 20%. At each step, a single image patch from each of the 10 images is taken and presented to observers. If a patch is recognizable (i.e. more than 50% of observers correctly recognized the image), then five descendants are taken from the patch. One descendant is obtained by reducing the resolution of the patch and the four others are obtained by cropping 20% of the patch on one corner. This is pictured in Fig.5 (a). For instance, the 50×50 original full images produces four cropped images with size 40×40 , and one 40×40 copy of the original with reduced resolution. The process was reiterated until finding an image patch where none of the five descendants reach a recognition rate of 50%. This patch was then labelled as MIRC, and its descendants were labelled as subMIRC's (Fig.5 (a), (b) and (c)). For presentation purposes, all patches were rescaled to 100×100 pixels by image interpolation so that the size of the presented image was homogeneous across trials without adding or losing information.

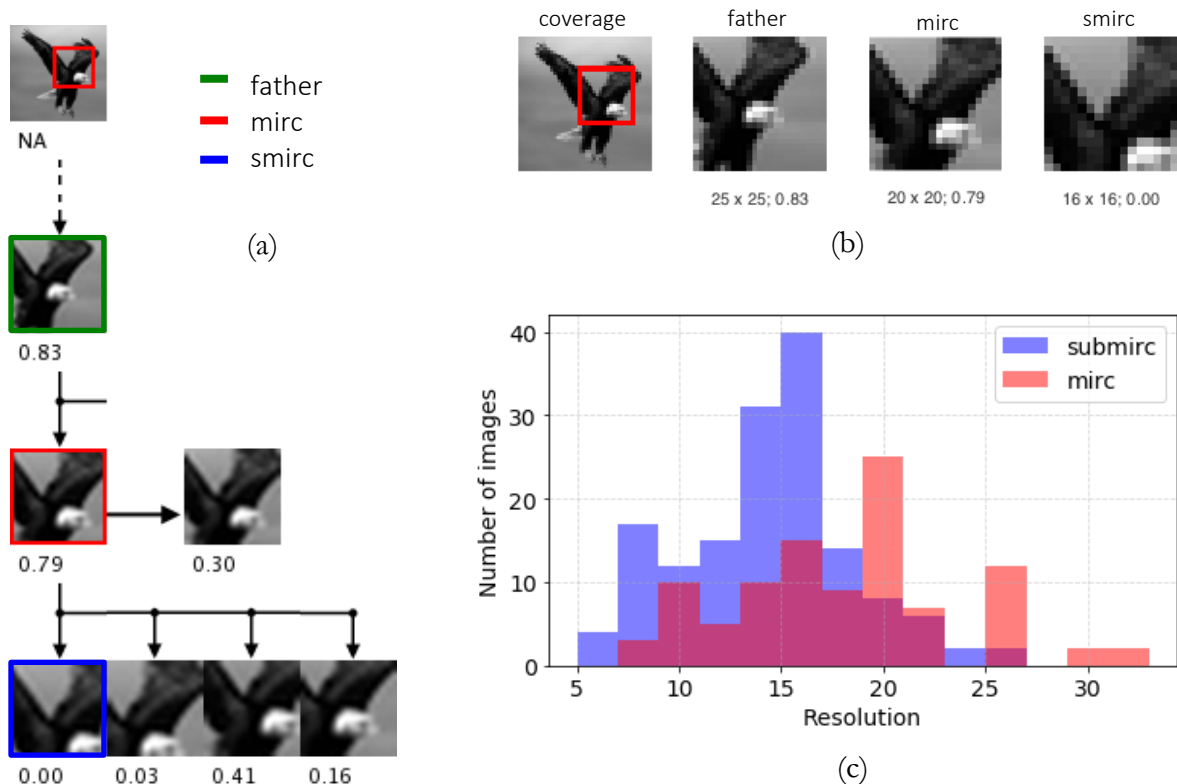


Fig. 5: MIRC and subMIRC images for testing. (a) Reduced images were obtained through a psychophysics experiment by reducing original images either in lowering resolution or cropping one corner. MIRCs correspond to patches where the five descendants are unrecognized and the subsequent descendants are labelled as subMIRCs. The recognition rate is indicated under the images. (b) Father, MIRC and subMIRC images obtained from a specific patch and used for testing. Resolution and recognition rate are indicated for each image. It can be noticed that the small change in the image from MIRC to subMIRC induces a strong drop in recognition to the subMIRC.s in the psychophysics experiment. (c) Distribution of resolution across MIRCs and subMIRCs for testing.

Compared to the experiment by Ullman et al., an additional image condition (father) was added. This condition is a less reduced version of the MIRC (as shown Fig.5 (a)) and aims to account for cases where subjects might not recognize the MIRC condition. Because our number of subjects is highly reduced compared to the psychophysics experiment used to get MIRCs and subMIRCs, we do not want to “lose” data for one patient if his recognition threshold happens to be different than expected (i.e. in between the father and the mirc). Adding this new condition is a way to capture information for these particular cases.

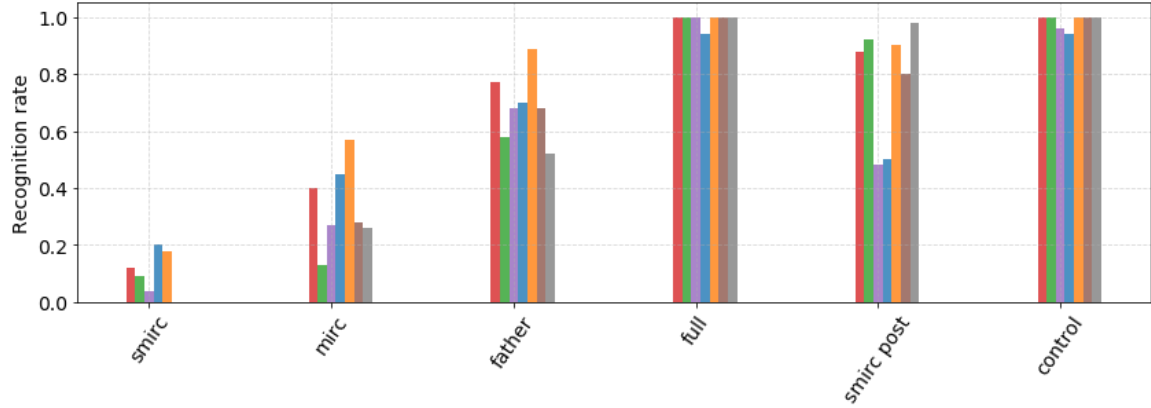


Fig. 6: Distribution of the recognition rate across image conditions in a preliminary behavioral study carried out with 7 non-epileptic subjects. Images were shown in blocks of two categories as in the final experiment. Compared to the final experiment where exemplars from the two categories are presented in a random order, here the exemplars of condition 1 of category 1 are all presented in a row, followed by exemplars of condition 1 of category 2, then condition 2 etc. Thus 20 exemplars of subMIRCs in the same category are presented in a row. The figure shows that this does not seem to affect the recognition performance.

Again because the number of epileptic subjects is obviously limited, special attention has to be granted to the design of the task. One primary difference with the task by Ullman et al. is that in our experiment, more than one image exemplar will be shown to each subject. Each subject will even see all conditions and all categories, and they will see them several times. A primary concern that was raised during the early phases of the study was that showing several same subMIRCs images in a row to a subject would make them more likely to be recognized. A preliminary behavioral study on non-epileptic subjects eventually removed that concern (see Fig. 6). However, images still have to be presented to subjects in a very specific order during the neurophysiological experiment to account for visual recognition.

Trials are organized in blocks where the succession of images is as pictured in Fig. 4 (b). Because the subMIRC condition is supposed to be the hardest to recognize, it is presented first. Then follow the MIRC, father and full image conditions. Control images are interspaced between the trials for control purposes. After full images, subjects are presented with a second occurrence of subMIRC images, which are exactly the same exemplars as the ones that were shown before. This aims to study whether the second occurrence induces some kind of perceptual discrimination due to the fact that subjects just saw the full image corresponding to the same category.

Two versions of the experiment have to be distinguished. For one subject, subMIRC images from all 10 categories were first presented on the screen, then all MIRCs from the 10 categories, then all fathers etc. This structure of the experiment had to be changed for psychological reasons. Since 100 subMIRC images (see Fig.4 (c)) and 100 MIRC images were first shown to the patient, there was a risk that none

of the images would be recognizable and the task would be considered “too hard” for the patient who might give up on trying to recognize subsequent stimuli. For this reason, all other subjects were presented with five blocks of images composed of only two categories at a time. For instance, they would be shown first 20 subMIRC exemplars belonging to bike and eagle categories, then 20 MIRC exemplars of bikes and eagles, then 20 father images, 10 full images and 10 other subMIRCs. The following block would consist of the same succession of images in e.g. horse and ship categories.

PRE-PROCESSING OF THE DATA

The signal was recorded from each electrode with sampling rates of 1000 Hz or 2000 Hz depending on subjects, before being amplified and notch-filtered to remove noise at 60 Hz. In two subjects, eye positions were recorded simultaneously with the physiological recordings, but this data was not analyzed at this stage of the study.

FURTHER PROCESSING

All the operations made on the data and data analyzes have been carried out with Python (Jupyter notebooks, Anaconda, Inc) using the scikit-learn toolbox in addition to standard libraries (Scipy, Numpy) and toolboxes. Matlab (R2020, Mathworks) was also used in certain cases, along with the signal processing toolbox.

DATA ANALYSES

RELABELING

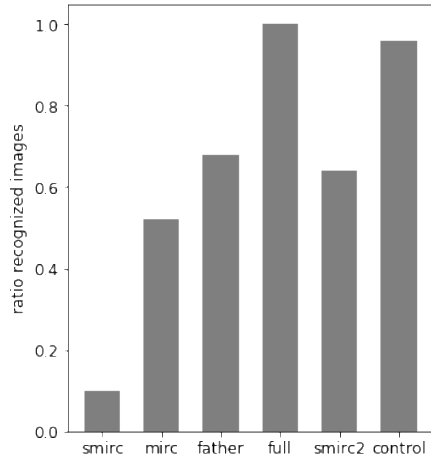
Because the study aims to investigate what happens around the threshold where an image becomes unrecognizable, one of the first steps to take with the data is to relabel the trials around this threshold. Among subjects and image categories, the same image condition does not always lead to a similar recognition rate, which can induce noise in the data regarding the phenomena under study. It is also important to put aside of the study cases where subMIRC images are recognized, as well as cases where father images are unrecognized. The method chosen to achieve optimal relabeling was to compute recognition rates (i.e. proportion of correct trials) in all conditions corresponding to an image category. Fig. 7 (a) and (b) show an example subject where, even though overall MIRC and father images seem to be recognized (a threshold of 50% has been used as a criterion for the recognition rate), the behavioral responses differ a lot across categories. SubMIRC, MIRC and father images in eagle, plane and ship categories are always unrecognized. Conversely, images in the car category are always recognized, even subMIRCs. In the bike category, MIRCs are not recognized but father images are. In this case, the category father images have been relabeled as MIRCs and MIRCs have been relabeled as subMIRCs. An example of relabels is presented in Fig. 7 (c)

VISUAL RESPONSIVITY

A first interesting question to answer is whether there are electrodes, among the dataset, that show visual responsivity related to stimulus presentation. As mentioned in the previous section, the dataset is composed of a total of 1712 electrodes, most of them most likely located in areas unrelated to the visual system. Hence discriminating visually responsive electrodes would allow to reduce the dimensionality of the data and still keep relevant information.

Previous studies have demonstrated that neural responses in the visual cortex usually occur in the window between 50 ms and 300 ms post stimulus onset [7][6]. The risk in considering neural responses occurring after this time period is that information contained in the data would be unrelated to the presentation of the stimulus, or even the experiment. Hence, this is the time window that will be considered to assess responsivity.

The method used involves comparing the IFP magnitude to the average baseline response in the period from -200 ms to 50 ms post stimulus onset. Comparisons are made at every time point in the 50 ms to 300 ms post stimulus onset time window using a paired t-test. This parametric test is commonly used to compare the means of two groups against a null-hypothesis of equal distributions

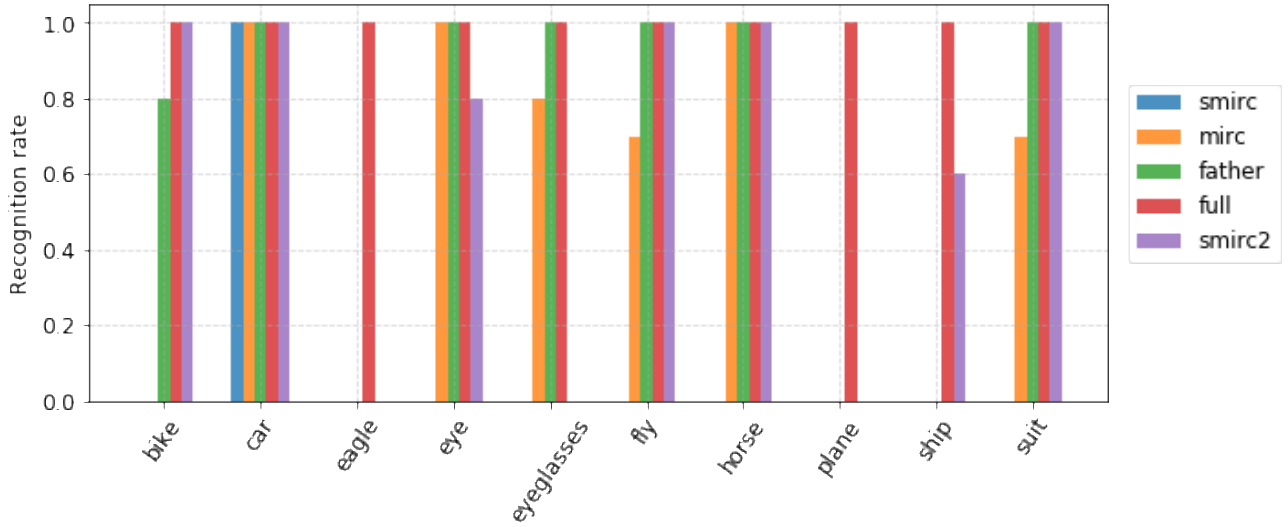


(a)

	smirc	mirc	father	full	smirc_post
bike	0	1	2	4	8
cardoor	9	9	9	9	9
eagle	10	10	10	4	10
eye	1	2	3	4	5
fly	1	2	3	4	5
glasses	1	2	3	4	5
horse	1	2	3	4	5
plane	10	10	10	4	10
ship	10	10	10	4	10
suit	1	2	3	4	5

0: sub-sMIRC; 1: sMIRC; 2: MIRC; 3: father; 4: full; 5: sMIRC_post; 8: sub-sMIRC_post; 9: allrecognized; 10: allunrecognized

(c)



(b)

Fig. 7: Example relabeling procedure for one subject of the experiment. (a) Behavioral recognition performance across image conditions. Full and control images show very good performance as expected. SubMIRCs, MIRCs and father images show increasing performance. subMIRCs_post are recognized much more often than subMIRCs. (b) Recognition rate by image category. Eagle, plane and ship categories are not recognized by this subject. The car category is too easily recognized and no recognition threshold can be observed. In the bike category, the threshold turns out to be a MIRC/father transition. (c) New labels per category and condition. The remaining part of the study will focus on conditions 1,2,4 and 5 in particular. Labels 6 and 7 correspond to control images, which do not require relabeling.

with the following assumptions: the two samples must be related and following a normal distribution. They should also be free from any outliers. The first two hypotheses seem reasonable considering the data under study. To remove potential outlier trials from the study, each trial was compared against the mean distribution of the IFP response magnitude across all trials. A trial was considered an outlier when its magnitude was higher than the mean and four standard deviations of the IFP magnitude distribution. The remaining data passed the criterion given in equation (1).

$$mag_{trial} < \overline{mag} + 4 \cdot smag \quad (1)$$

To assess whether an electrode was visually responsive, it was chosen to only consider trials corresponding to full and control images, as they are supposed to be best recognized and thus might elicit the most significant responses. It has to be noted that we distinguished two types of “responsivity” in this study. Hypothetically, brain areas might either respond to the presentation of a stimulus regardless of the image condition (essentially, the onset of pixels on the screen induces a visual response), or they might respond according to the intensity of the internal representations elicited by the image and thus respond better to images that are easily recognizable. As we are interested in capturing both cases, full and control images were the most relevant to take into account in the comparison.

The p-value is computed following equation (2) where μ_0 is the null hypothesis for the comparison between the two samples ($\mu_0 = 0$ in this case) and n is the size of sample x .

$$p = \frac{\bar{x} - \mu_0}{s_x / \sqrt{n}} \quad (2)$$

An electrode is considered visually responsive if the comparison rejects the null hypothesis (p-value below 0.01) for more than 35 consecutive milliseconds. In addition, the IFP response magnitude has to be higher than 70 μV to pass the criteria. These two thresholds have been obtained through trial and error. The method leads to a false discovery rate of 0%, as computed over 100 iterations of randomly shuffling the baseline and epoch responses.

In such visually responsive electrodes, responsivity was then assessed at a per-condition level, computing the p-values corresponding to the trials in a given image condition. If a condition is significantly different from its baseline response for more than 35 consecutive milliseconds, then it is responsive. With this criterion, an electrode that is considered visually responsive is not necessarily responsive in all the image conditions.

LATENCY MEASUREMENTS

To characterize the latency of neural responses, a first attempt was made considering the moment where the maximum IFP amplitude occurs in the time period between 50 ms and 300 ms post stimulus onset. This method turned out to be highly unprecise in some cases, hence it was chosen to turn to a different criterion. After discriminating visually responsive electrodes in a first step, the latency has been defined as the first time point where the neural response is considered significantly different from the baseline response according to the t-test p-value criterion. The latency was then defined at a per-condition level, where all the trials in a given condition were averaged.

COMPARISONS

To compare the different conditions at a neural level, several types of analyzes have been run. This has been done in an iterative process, trying to always extract the best possible information for the data at hand.

Parametric tests

Statistical comparisons of the neural responses in different conditions were made first. The analyses include a one-way analysis of variance (ANOVA) to compare the variance related to each of the image conditions, and a t-test to run pairwise comparisons.

Conditions were considered significantly different from one another when the p-value reached a value below 0.01. Comparisons were made at a time point level as explained in the previous section on visual responsivity.

Latencies

Latencies between two different conditions were compared computing the corresponding value of the cumulative binomial distribution function. Latencies were assessed over all subjects and visually responsive electrodes and arranged in pairs of conditions. The proportion of values tending towards each of the two conditions was computed. The value of the cumulative binomial distribution then accounts for the probability of observing the same number or fewer successes over the same number of experiments, when a single experiment has a probability p of success (here $p = 0.5$). Equations (3) and (4) provide the general formula.

$$F(x) = P(X \leq x) = \sum_{m=0}^x f(m) = f(0) + f(1) + \dots + f(x) \quad (3)$$

$$\text{where} \quad f(k) = \binom{n}{k} p^k (1-p)^{n-k} \quad (4)$$

Classification

To discriminate further the neural responses corresponding to conditions, it was used a machine learning approach where classifiers were trained on the neural data before being tested. As labels are available at a trial level, the analysis falls into supervised learning techniques. The upside of this approach compared to other techniques is that it allows to combine several features of the data to try and differentiate the conditions.

To run the analysis, all the visually responsive electrodes are combined together in a so-called “pseudo-population” of electrodes (i.e. across subjects, measures were not recorded simultaneously). Data is randomly sampled so as to obtain an identical number of trials per condition across subjects. The data is then split in two separate datasets: a training set which allows to train the machine learning classifier to discriminate conditions, and a testing set on which the performance of the model is then evaluated. This separation of the data is crucial so that the model does not “see” test data before the actual test, which would induce an overestimation of the classifier performance, also referred to as “overfitting” in machine learning terms.

The classifier performance accounts for the accuracy of its prediction on the test data and is computed as an averaged score over a certain number of train-test splits, also referred to as the “cross-validation” step in machine learning terms. This is to ensure the prediction is not biased towards a particular result, due to the random character of the train-test split and the small number of trials per condition resulting from the sub-sampling over subjects.

In this study, a support vector machine was used as the classifier. This type of model is commonly used in classification analyses of neural data and is known to result in good generalization performance with such type of dataset.

The training data accounted for 70% of the dataset, and test data accounted for the remaining 30%. A 5-fold stratified shuffle split was used as cross-validation. Binary comparisons were run across conditions, and all features were computed as an averaged value over 50 random sub-sampling iterations. For each comparison, the chance level was computed by testing the model on data with randomly shuffled labels and averaging the score over 100 iterations.

Several features have been computed and tested with the model. Only features allowing to get the best performance of the model were then kept.

IFP magnitude. The magnitude of the IFP response was computed in the time window from 50 ms to 500 ms post stimulus onset.

Time of maximum magnitude. The time at which occurs the maximum IFP amplitude of the signal in the time window from 50 ms to 500 ms post stimulus onset.

Sliding windows. The amplitude of the IFP response averaged in sliding time windows in bins of 20 ms. This was also done in the time window from 50 ms to 500 ms post stimulus onset.

Principal components analysis (PCA). The principal components of the data, according to a principal component analysis. This analysis allows to reduce the dimensionality of the data by projecting it to a lower dimensional space and keeping the relevant information. Five principal components were used in this analysis.

Frequency analysis. The instantaneous power was computed from the time series. Previous research has shown that electrical activity in the brain could be analyzed from the power in different frequency bands [34]. One of them in particular is highly interesting regarding the analysis of neural signals in the visual cortex: the gamma frequency band [35]. Definitions of frequency bands vary across the literature. The interval 30Hz-100Hz was picked for the gamma band, as it is defined in several studies [34][36][37].

The signal is first bandpass-filtered using a Butterworth filter of order 5 in the desired frequency band. The complex representation of the signal is then obtained via the Hilbert transform and the instantaneous power is obtained through getting the magnitude.

To evaluate the influence of a feature on the classification performance of the model, we proceeded to look at the distribution of each feature across the dataset. When the distribution of the feature allows to distinguish conditions, then it is relevant to use it for the task. In most analyses, the features used were the magnitude and the timings of maximums and minimums.

RESULTS

BEHAVIORAL RECOGNITION

All subjects taking part in the neurophysiological experiment were first evaluated on their recognition performance at a behavioral level. Fig. 8 shows the distribution of these evaluations across subjects and different image conditions. It can be noticed immediately the good recognition performance at full and control images, which was certainly expected. Cases where subjects gave incorrect feedback

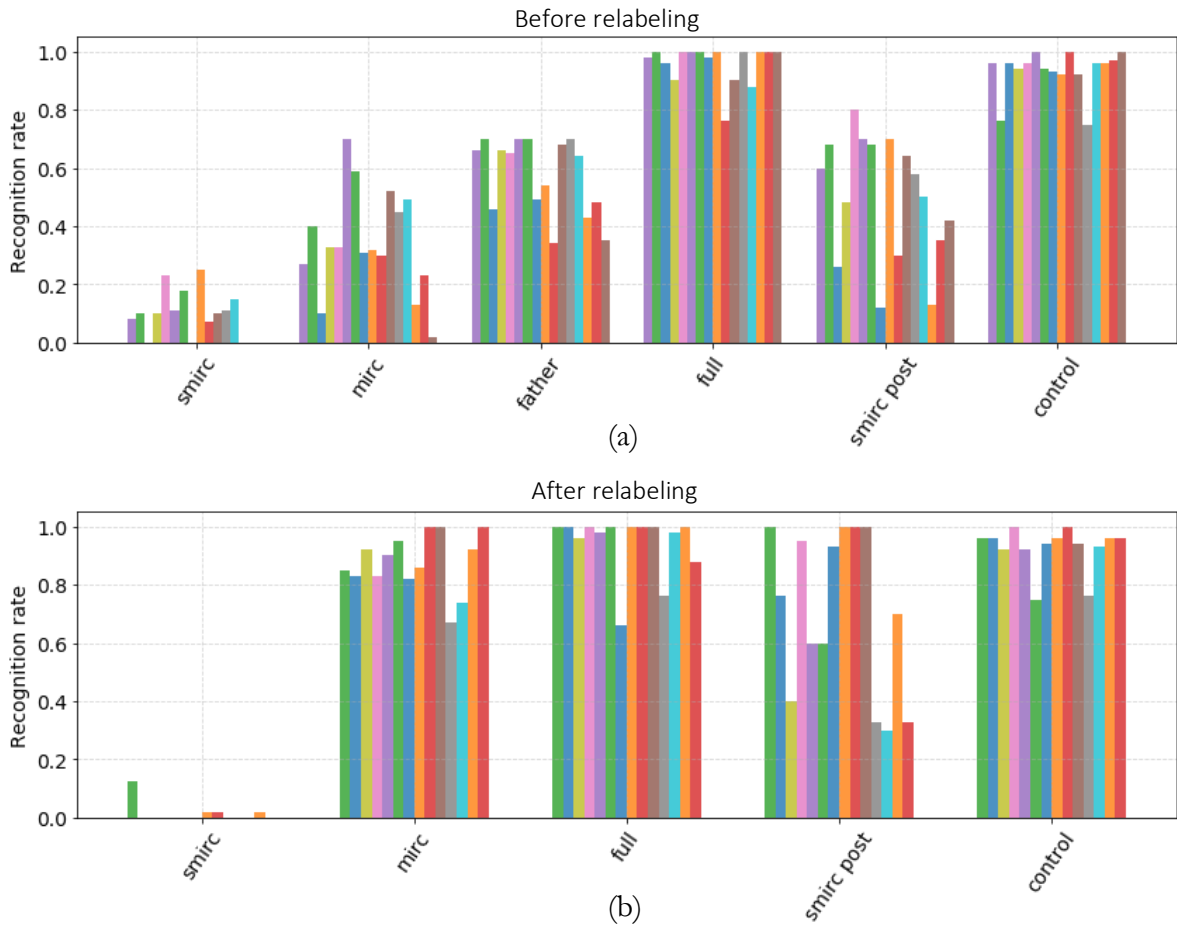


Fig. 8: Distribution of the behavioral recognition to the different image conditions across all subjects taking part in the study. Each color accounts for one subject. (a) Before relabeling. It can be observed that most subjects do not pass the 50% threshold of recognition for MIRC images. This highlights the relevance of father images and the need to differentiate recognition between image categories. (b) After relabeling. A strong recognition drop occurs between the subMIRC images and MIRC images. Besides, subjects do seem to recognize much better the second occurrence of the sMIRC images at a behavioral level.

to full and control images occurred mostly due to memory issues and lack of attention probably due to fatigue (e.g. subject did not get the time to see the image or did not get the time to reply). The epileptic patients involved in the study indeed are under heavy medication and they are also recovering surgery.

Overall, the recognition to father images is higher than to MIRC images, which is in turn higher than subMIRC images. However, it would have been reasonable to expect that MIRC images in Fig. 8 (a) would be better recognized compared to subMIRC images. One result of Ullman et al. in their MIRC study was observing a sharp drop in recognition from MIRC images to subMIRC images across most observers taking part in their task. They suggest that the transition occur for the same images regardless of individual experience and thus different subjects would share similar visual representations about the images [1]. Our behavioral data shows that not only the recognition threshold varies across subjects but, even

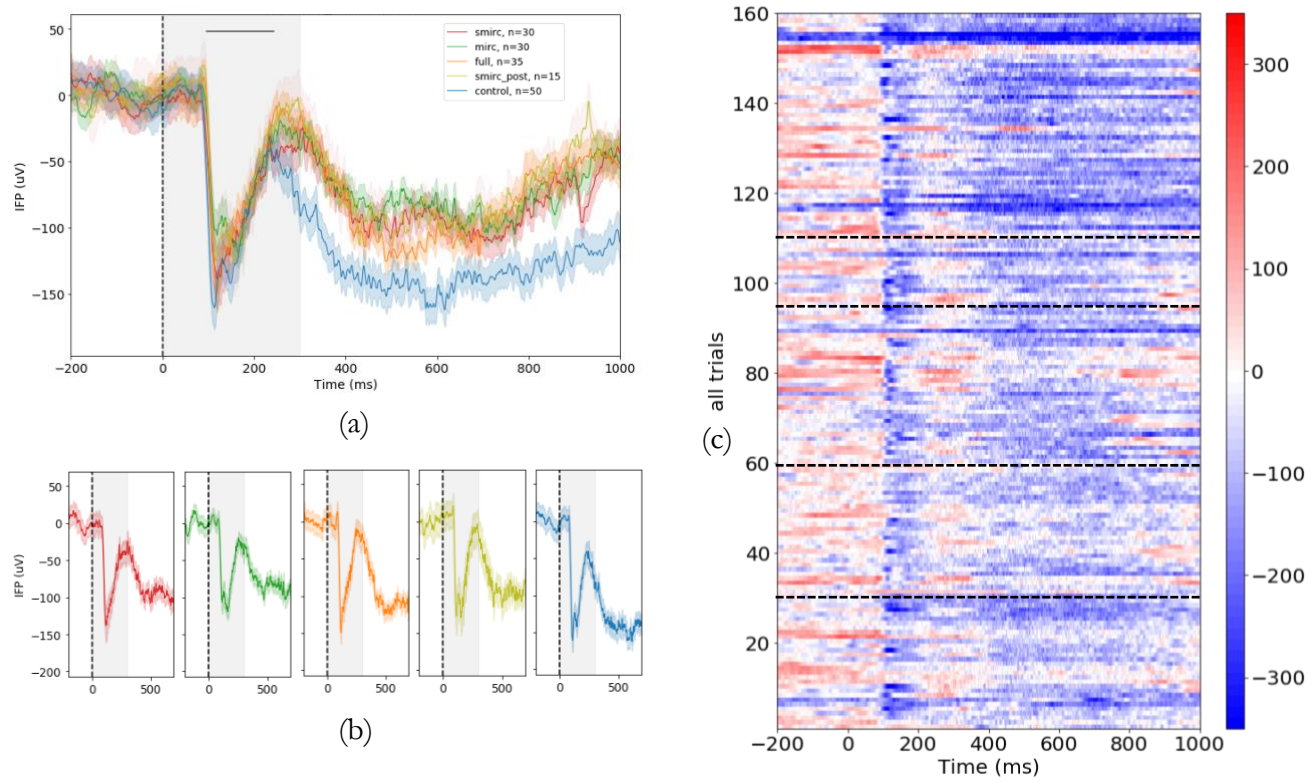


Fig. 9: Example neural response from an electrode in the left temporal cortex. Time zero indicates the onset of stimulus presentation on the computer screen. (a) IFP response averaged by image condition. The horizontal black line indicates the time period when the IFP response significantly differs from the baseline response (t-test, $p < 0.01$). (b) IFP response averaged by condition. The trials have been previously relabeled. (c) IFP response across all trials. Each line accounts for a trial. On the ordinate axis in increasing order are subMIRC, MIRC, full, subMIRC_post and control trials. The color indicates the intensity of the response in microvolts. Visual responsivity can be observed around 100 ms post stimulus onset.

more importantly, it varies significantly across image categories. Our data is, however, much more reduced in number of subjects, which could explain getting high variance in behavioral recognition compared to Ullman’s work.

Fig. 8 (b) shows the same result after a relabeling procedure was applied to the data. A sharp drop in recognition (0.86 ± 0.08) can now be observed in between subMIRC and MIRC images, confirming the result by Ullman et al. Besides, the second occurrence of subMIRCs also show much higher recognition performance (difference of 0.65 ± 0.3) than the first occurrence, suggesting that subjects might perceptually discriminate these stimuli on internal representations built with MIRC and full images that they saw just before.

NEURAL RESPONSES AND VISUAL RESPONSIVITY

Neural responses were recorded across 1712 electrodes in 12 subjects, and it could be observed responsivity to visual stimuli in some of them. An example electrode is shown in Fig. 9, where a visual response to the stimuli occur around 100 ms. Another visually responsive electrode is presented in Fig. 10.

56 electrodes were evaluated visually responsive in the time interval from 50 ms to 300 ms post stimulus onset, according to the data analysis described in the previous section. Among these electrodes, 51 were specifically responsive to the full image condition, 50 were responsive to the MIRC condition, 36 were responsive to the subMIRC condition and 22 were responsive to the subMIRC_post condition. All electrodes were specifically responsive to control images, which are the easiest to recognize. This result shows that images that are more recognizable seem to elicit higher neural responses on average across subjects. This agrees with a result by Lerner et al. where they show that greater brain activation occurs in informative image fragments over non-informative fragments [38]. Even more so in our study, nearly as many electrodes are responsive to MIRCs and full images, but a drop in the number of responsive electrodes occurs for the subMIRC condition. This reminds of the drop in recognition rates that could be observed in the behavioral part of the study, and that has been mentioned previously.

The low number of electrodes responsive to the second occurrence of the subMIRCs can be explained by the fact that this condition usually has only half the number of trials compared to other conditions, or even less depending on the relabeling. This results in electrodes being rejected of the analysis for this condition, because it is not represented by enough trials to be considered relevant (fewer than 10 trials).

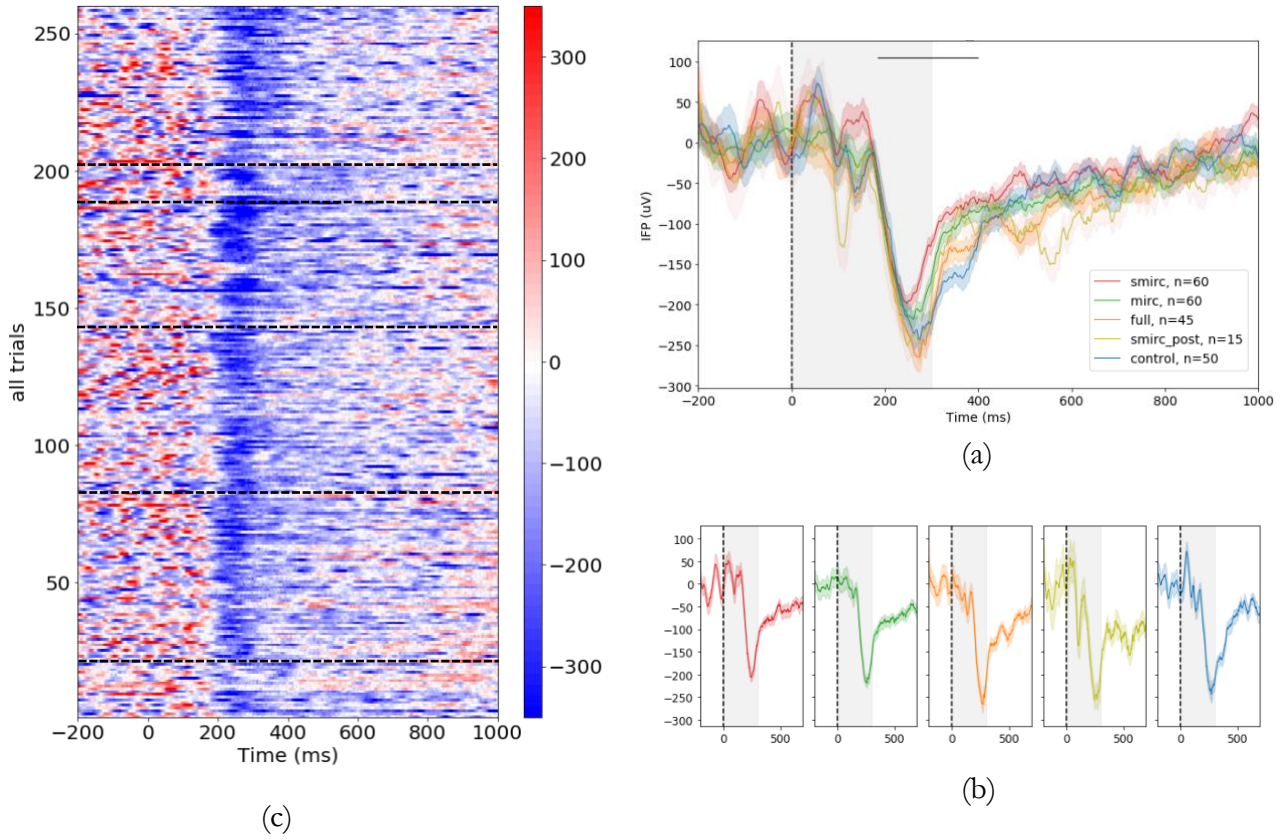


Fig. 10: Example neural response from an electrode in the temporal cortex. It can be observed that the visual response seems to be shifted with respect to the example electrode presented in Fig. 9 (c) In increasing order on the ordinate axis: sub-sMIRC, subMIRC, MIRC, Full, subMIRC_post and Control.

VISUAL SELECTIVITY

Although this is not the main focus of the study, it was interesting to investigate whether some electrodes were selective to some image categories in particular. In Fig. 11 an electrode in the occipital lobe is pictured where some image categories seem to induce a higher visual response compared to others. This result seem to correlate with a recent fMRI study on minimal images, where it was shown that minimal images were able to elicit category-selective responses in higher visual areas such as the lateral occipital cortex, selective to objects [39]. However, the localization of electrodes has not been thoroughly evaluated in our study, thus it is not possible to say with sufficient accuracy where the electrode in Fig. 11 is actually implanted.

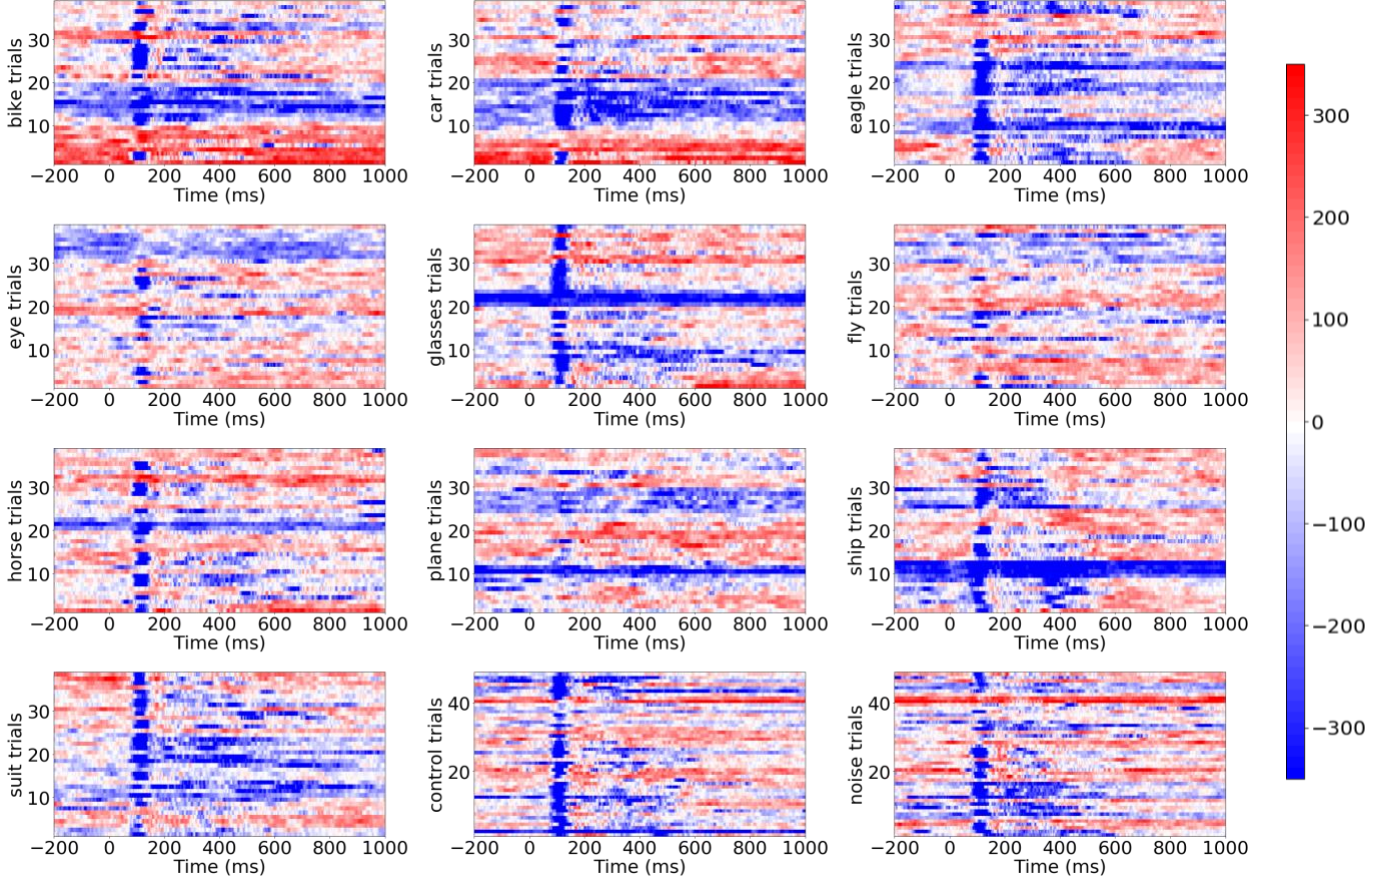


Fig. 11: Neural response from an electrode in the occipital lobe. The raster plots picture every trial as a line, and the intensity of the response corresponds to the color scale in microvolts on the right. It can be observed that the visual response appears in different intensities depending on the image category. Plane and fly trials are weakly represented compared to glasses and eagle categories, although plane and fly categories were recognized by the subject at a behavioral level.

TIMINGS

Differences in latencies could be observed in particular electrodes such as the ones pictured in Fig. 13 (a) and (b). To try and account for these differences, latency measurements were presented in pairwise comparisons between image conditions. This is shown in Fig. 12. The analysis shows a clear difference for the second occurrence of subMIRCs in particular, which seem to have higher latencies compared to subMIRC, full and control images. This could be explained by the fact that this condition calls on short-term memory and perceptual representations from the previous full image. Pattern completion skills could also be involved, that have been shown to correlate with higher latencies [23]. This result endorses the behavioral result that shows much higher performance for these subMIRC_post compared to subMIRCs.

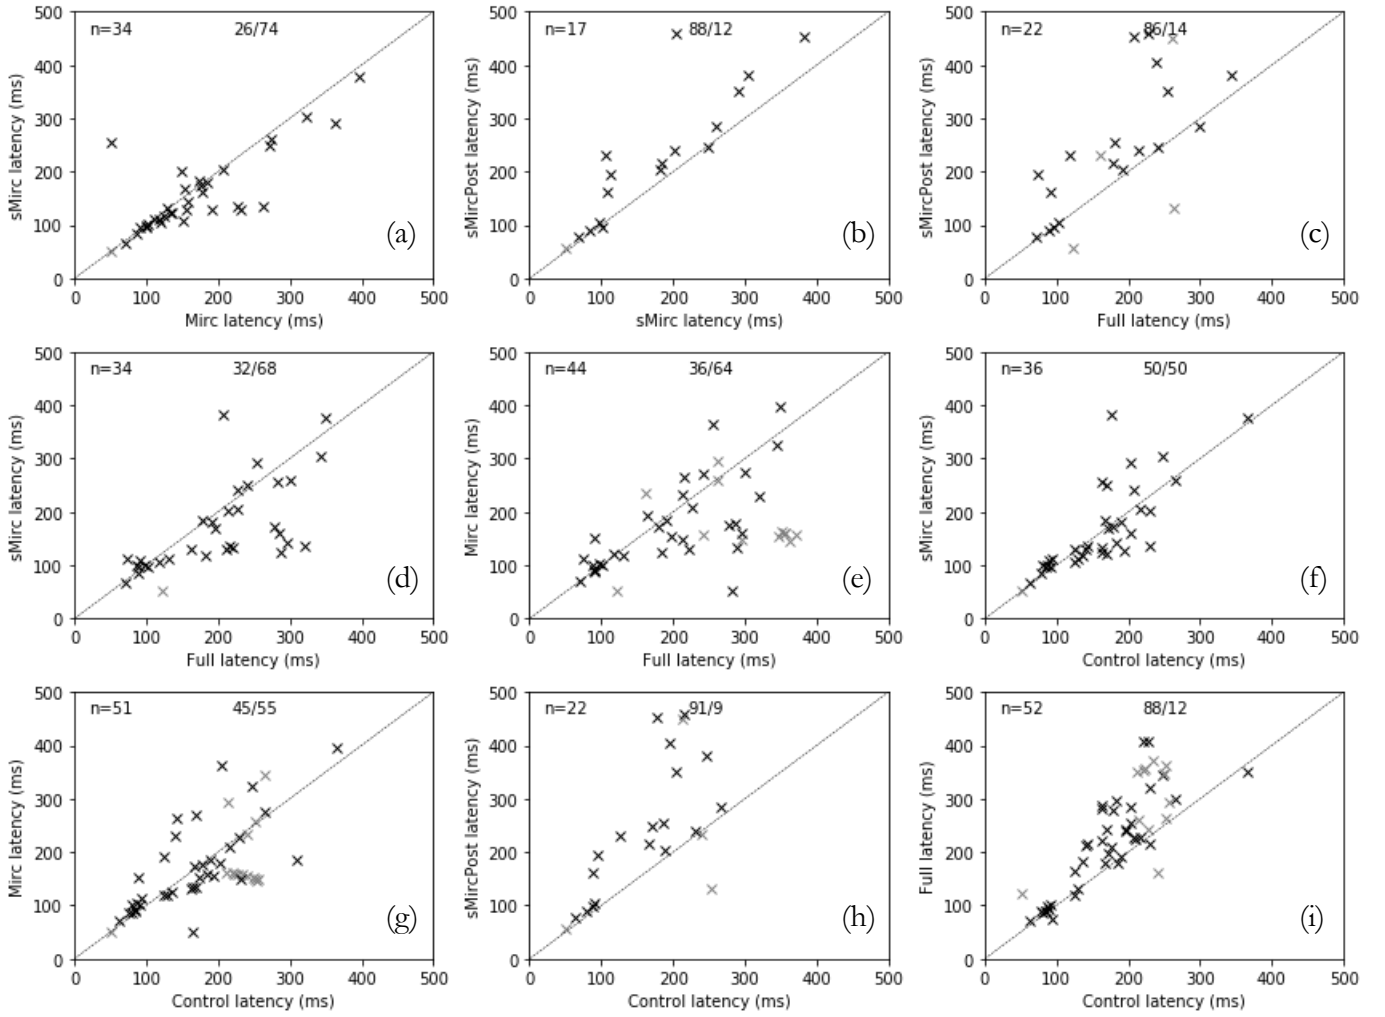


Fig. 12: Pairwise comparisons of latencies. Each plot represents a comparison and each point on a plot accounts for an electrode among the 56 visually responsive electrodes. On the top left corner of each plot is given the number of electrodes visually responsive to both conditions represented in the plot. On the top center is given the percentage of points respectively above and below the diagonal. It can be observed that the subMIRC_post condition has significantly longer timings for all conditions it is compared to (plots (b), (c) and (h)).

The latency results hardly show any other tendency on other conditions and this is probably a sign that the metric could be further optimized. Other than that, latencies might not characterize the different conditions as well as expected, and it would be useful to include other features of the data to try drawing a clearer comparison between the conditions.

COMPARISONS

A first step was to be able to visually compare the different conditions. Fig. 13 (a) and (b) show two example electrodes where differences can be observed between subMIRCs and MIRCs and subMIRC_post, both in timing and in amplitude.

Across the dataset, neural responses differ a lot and finding a good metric to accurately compare the conditions has been a real challenge. In order to include more features of the data into the analysis, it was adopted a machine learning approach where a classifier was trained on data from visually responsive electrodes to try and categorize the different conditions. The results of pairwise comparisons are presented in Fig. 14 (a) and (b). The features used for the classifier were the magnitude of the neural response along with the timing of both the maximum and minimum amplitude.

Fig. 14 (a) and (a) show the classification score is, except for one comparison, significantly above chance. Thus, the model is able to discriminate pairs of conditions when compared to one another based on the given features.

Interestingly, the classification performance for the subMIRC/MIRC comparison is very low above chance level. SubMIRCs were expected to elicit lower responses compared to MIRCs, especially given the fact that the visual responsivity analysis led to evaluate quite fewer subMIRC-responsive electrodes compared to MIRC-responsive electrodes.

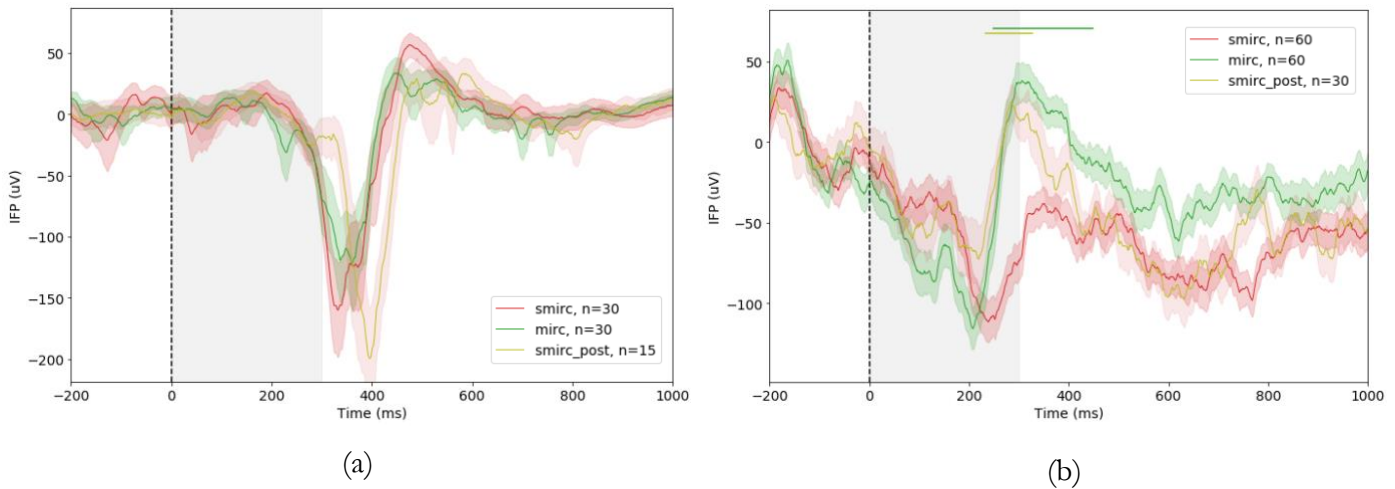


Fig. 13: Visual representation of the comparison between subMIRC, MIRC and subMIRC_post images in two example electrodes. (a) An electrode in the left temporal cortex, where a difference in timing can be observed between subMIRC_post and the two other conditions (b) An electrode located in the mesial temporal lobe, where the subMIRC condition seem to differ in amplitude compared to the other conditions. The horizontal bars show the time period when the subMIRCs are significantly different from the MIRCs (green) and subMIRC_post (beige) (t-test, $p < 0.01$).

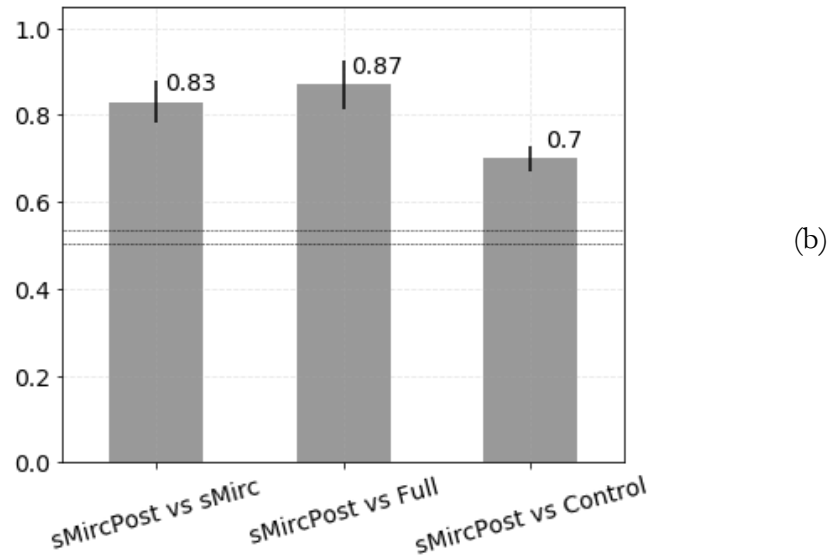
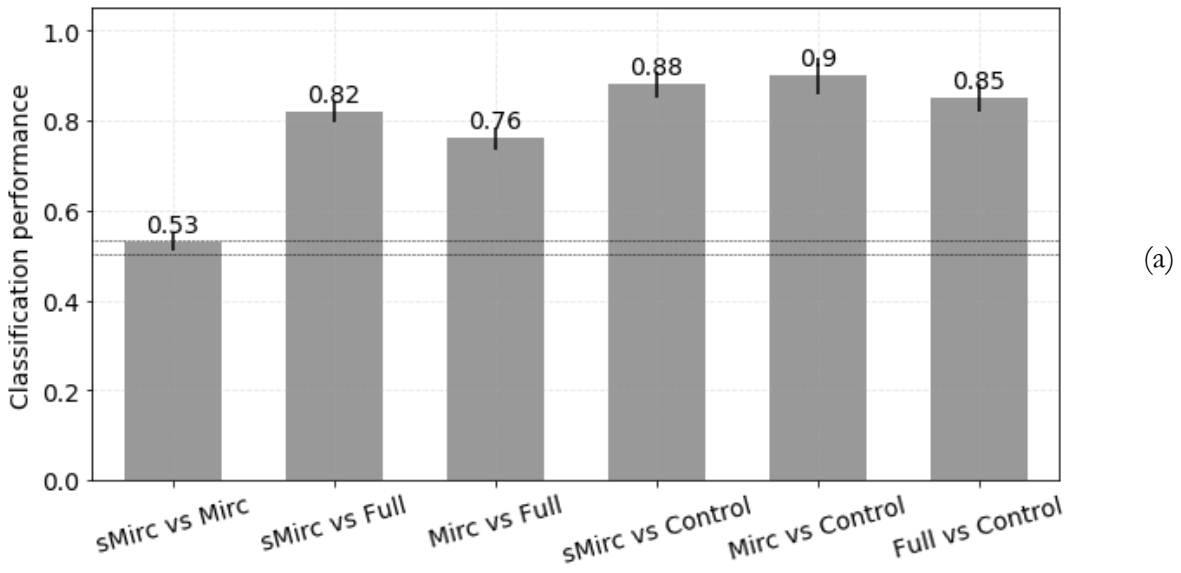


Fig. 14: Classification performance results. (a) Results of the classification task (Data analyses section). The lower dashed line represents the chance level obtained by randomly shuffling the labels over 100 iterations. The upper dashed line accounts for the statistical significance of the result, defined as 3 standard deviations above chance score. Data from all subjects were taken into consideration in this computation. (b) Comparisons with subMIRC_post images. This computation was done removing 3 subjects from the study, due to the low number of trials remaining for subMIRC_post images after relabeling.

DISCUSSION

Overall, the results obtained are consistent with previous analyzes on minimal images [1][39]. The behavioral data allowed to observe a sharp threshold in recognition at the MIRC level, confirming the result by Ullman et al. (2016) and introducing once again this type of image as an unique configuration requiring each and every feature available to be effectively used for recognition.

The fact that the second occurrence of subMIRC image shows both a much better recognition performance than the first occurrence at a behavioral level, and higher latencies, seem to be consistent with the hypothesis that such condition induces perceptual discrimination calling on different computations maybe related to short-term memory and pattern completion skills. It could be argued that such computations would involve feedback modulations in the visual cortex, similarly to the task by Tang et al. [23]. It would then be interesting to study whether this observation can be accounted for by using computational models of vision.

An important step to take further the analysis would be to properly run the localization of electrodes in subjects. This part of the study was put aside for logistic reasons. The brain locations provided in this report are based on predicted surgery plans and notes from doctors, but a more accurate account of electrode implantation could be made by looking at subjects CT and MRI images.

Considering visual selectivity, it would be interesting to observe whether the example electrode from Fig. 11 belongs to a brain area that is known as “category-selective” such as the lateral occipital cortex, as well as whether other electrodes can be found category-selective in the same area.

Comparing image conditions in groups of electrodes in a similar localization could be a good way to grasp the specific processes occurring in different areas. For example, MIRCs could be expected to elicit higher responses compared with subMIRCs in high-level areas engaged in visual recognition, but not necessarily in the early visual cortex that is less involved in pattern recognition [39]. This might be part of the explanation for the poor recognition performance on subMIRCs and MIRCs in the machine learning analysis (Fig. 14 (a)), if most visual responses do not occur in high-level visual areas across the dataset.

Further analysis could also be run with the data to improve our understanding of the processes underlying recognition. Another interesting way to compare conditions is known as “demixed” principal component analysis and has been introduced in a paper by Kobak et al. as a way to account for simultaneous representations of different elements of a task [40]. This analysis reduces dimensionality using the dependance of the data on external parameters such as stimuli and decisions. The output is then a “mixed” representation over a time and stimulus components.

Also, a lead to improve the accuracy of latency measurements and statistical comparison tests would be to try reducing the sensitivity to type I errors (false positive) using multi-comparison correction methods [41][42].

SUMMARY

Twelve patients with pharmacological intractable epilepsy completed a neurophysiological experiment aiming to study object recognition at a minimal image level. Minimal images are unique configurations where visual features are crucial for recognition and the smallest change to the image can make it unrecognizable. Results show that such a small change in size or resolution induces a sudden drop in recognition performance ($86\% \pm 8\%$) at a behavioral level. At a neural level, visual responses do not seem to be specific enough to draw a distinction between minimal and sub-minimal images. Interestingly, neural responses to sub-minimal images seem to match perceptual discrimination when images are presented a short time after an easily recognizable version of the image. Higher latencies are also involved with these types of sub-minimal images, suggesting bottom-up modulations might take place in the computational account for recognition.

RÉSUMÉ

Douze patients présentant de l'épilepsie intractable pharmacologiquement ont accompli une expérience neurophysiologique visant à étudier la reconnaissance d'objet à un niveau d'images minimales. Les images minimales sont des configurations uniques où les caractéristiques visuelles sont cruciales pour la reconnaissance et le plus petit changement fait à une image peut la rendre irreconnaissable. Les résultats montrent qu'un petit changement de taille ou de résolution entraîne une baisse soudaine des performances de reconnaissance ($86 \% \pm 8 \%$) au niveau comportemental. Au niveau neuronal, les réponses visuelles ne semblent pas être suffisamment spécifiques pour faire une distinction entre les images minimales et sous-minimales. Fait intéressant, les réponses neuronales à des images sous-minimales semblent correspondre à de la discrimination perceptuelle lorsque les images sont présentées peu de temps après une version facilement reconnaissable de l'image. Des latences plus élevées sont également impliquées dans ces types d'images sous-minimales, ce qui suggère que des modulations ascendantes pourraient avoir lieu dans le procédé de calcul pour la reconnaissance visuelle à ces images.

REFERENCES

- [1] S. Ullman, L. Assif, E. Fetaya, and D. Harari, “Atoms of recognition in human and computer vision,” *Proc. Natl. Acad. Sci.*, vol. 113, no. 10, pp. 2744–2749, Mar. 2016, doi: 10.1073/PNAS.1513198113.
- [2] “The visual system.” <https://www.nature.com/subjects/visual-system>.
- [3] D. J. Felleman and D. C. Van Essen, “Distributed hierarchical processing in the primate cerebral cortex,” *Cereb. Cortex*, vol. 1, no. 1, pp. 1–47, 1991, doi: 10.1093/cercor/1.1.1.
- [4] J. Blumberg and G. Kreiman, “How cortical neurons help us see: Visual recognition in the human brain,” *J. Clin. Invest.*, vol. 120, no. 9, pp. 3054–3063, 2010, doi: 10.1172/JCI42161.
- [5] James J. DiCarlo, D. Zoccolan, and N. C. Rust, “How does the brain solve visual object recognition?,” *Bone*, vol. 23, no. 1, pp. 1–7, 2014, doi: 10.1038/jid.2014.371.
- [6] T. Serre, G. Kreiman, M. Kouh, C. Cadieu, U. Knoblich, and T. Poggio, “A quantitative theory of immediate visual recognition,” *Prog. Brain Res.*, vol. 165, 2007, doi: 10.1016/S0079-6123(06)65004-8.
- [7] H. Liu, Y. Agam, J. R. Madsen, and G. Kreiman, “Timing, Timing, Timing: Fast Decoding of Object Information from Intracranial Field Potentials in Human Visual Cortex,” *Neuron*, vol. 62, no. 2, pp. 281–290, 2009, doi: 10.1016/j.neuron.2009.02.025.
- [8] D. Atchison and G. Smith, *Optics of The Human Eye*. Butterworth Heinemann, Oxford, United Kingdom, 2000.
- [9] P. G. J. Barten, *Contrast Sensitivity of the Human Eye and Its Effects on Image Quality*. SPIE Press, 1999.
- [10] C. Van De Pol, “Basic Anatomy and Physiology of the Human,” *Basic Anat. Hum. Vis. Syst.*, no. mm, p. chapter 6 pg238-246, 2014, [Online]. Available: http://www.usaarl.army.mil/publications/HMD_Book09/files/Section 13 - Chapter 6 Anatomy and Structure of the Eye.pdf.
- [11] R. H. Wurtz and J. E. Albano, “Visual-motor function of the superior colliculus,” 1980. Accessed: Aug. 27, 2020. [Online]. Available: www.annualreviews.org.
- [12] J. M. Sprague and T. H. Meikle, “The role of the superior colliculus in visually guided behavior,” *Exp. Neurol.*, vol. 11, no. 1, pp. 115–146, Jan. 1965, doi: 10.1016/0014-4886(65)90026-9.

- [13] J. D. Haynes, R. Deichmann, and G. Rees, “Eye-specific effects of binocular rivalry in the human lateral geniculate nucleus,” *Nature*, vol. 438, no. 7067, pp. 496–499, Nov. 2005, doi: 10.1038/nature04169.
- [14] P. H. Schiller and J. G. Malpeli, “Functional specificity of lateral geniculate nucleus laminae of the rhesus monkey,” *J. Neurophysiol.*, vol. 41, no. 3, pp. 788–797, 1978, doi: 10.1152/jn.1978.41.3.788.
- [15] R. J. Douglas and K. A. C. Martin, “Neuronal circuits of the neocortex,” *Annu. Rev. Neurosci.*, vol. 27, pp. 419–451, 2004, doi: 10.1146/annurev.neuro.27.070203.144152.
- [16] R. B. H. Tootell *et al.*, “Functional analysis of primary visual cortex (V1) in humans.,” in *National Academy of Sciences (NAS Colloquium) Neuroimaging of Human Brain Function.*, 1998, doi: 10.17226/6236.
- [17] I. Rentzeperis, A. R. Nikolaev, D. C. Kiper, and C. van Leeuwen, “Distributed processing of color and form in the visual cortex,” *Front. Psychol.*, vol. 5, no. AUG, pp. 1–14, 2014, doi: 10.3389/fpsyg.2014.00932.
- [18] J. V. Haxby *et al.*, “Dissociation of object and spatial visual processing pathways in human extrastriate cortex,” *Proc. Natl. Acad. Sci. U. S. A.*, vol. 88, no. 5, pp. 1621–1625, 1991, doi: 10.1073/pnas.88.5.1621.
- [19] S. H. Creem and D. R. Proffitt, “Defining the cortical visual systems: “What”, “Where”, and “How”,” *Acta Psychol. (Amst)*, vol. 107, no. 1–3, pp. 43–68, 2001, doi: 10.1016/S0001-6918(01)00021-X.
- [20] E. K. Miller, “The Prefrontal Cortex and Cognitive Control,” vol. 1, no. October, pp. 137–154, 2007, doi: 10.1007/978-3-540-45702-2_10.
- [21] L. R. Squire, C. E. L. Stark, and R. E. Clark, “The medial temporal lobe,” *Annu. Rev. Neurosci.*, vol. 27, pp. 279–306, 2004, doi: 10.1146/annurev.neuro.27.070203.144130.
- [22] H. Eichenbaum, “A Cortical – Hippocampal System for declarative memory,” *Nat. Rev. Neurosci.*, vol. 1, no. October, pp. 41–50, 2000.
- [23] H. Tang *et al.*, “Spatiotemporal Dynamics Underlying Object Completion in Human Ventral Visual Cortex,” *Neuron*, vol. 83, no. 3, pp. 736–748, 2014, doi: 10.1016/j.neuron.2014.06.017.
- [24] H. Tang *et al.*, “Recurrent computations for visual pattern completion,” *Proc. Natl. Acad. Sci. U. S. A.*, vol. 115, no. 35, pp. 8835–8840, 2018, doi: 10.1073/pnas.1719397115.
- [25] B. Epshtein, I. Lifshitz, and S. Ullman, “Image interpretation by a single bottom-up top-down cycle,” *Proc. Natl. Acad. Sci.*, vol. 105, no. 38, pp. 14298–14303, Sep. 2008, doi:

10.1073/PNAS.0800968105.

- [26] A. M. Siegel, “Presurgical evaluation and surgical treatment of medically refractory epilepsy,” *Neurosurg. Rev.*, vol. 27, no. 1, pp. 1–18, 2004, doi: 10.1007/s10143-003-0305-6.
- [27] T. Yang, S. Hakimian, and T. H. Schwartz, “Intraoperative ElectroCorticoGraphy (ECog): indications, techniques, and utility in epilepsy surgery,” *Epileptic Disord.*, vol. 16, no. 3, pp. 271–279, Sep. 2014, doi: 10.1684/epd.2014.0675.
- [28] K. M. Szostak, L. Grand, and T. G. Constandinou, “Neural interfaces for intracortical recording: Requirements, fabrication methods, and characteristics,” *Front. Neurosci.*, vol. 11, no. DEC, 2017, doi: 10.3389/fnins.2017.00665.
- [29] E. S. Hughes, W. D. Gaillard, G. M. Jacyna, and S. J. Schiff, “Feature-based detection of epileptic seizures using ECoG recordings and a PAM model,” in *Annual International Conference of the IEEE Engineering in Medicine and Biology - Proceedings*, 1995, vol. 17, no. 2, pp. 933–934, doi: 10.1109/iembs.1995.579355.
- [30] T. Ball, M. Kern, I. Mutschler, A. Aertsen, and A. Schulze-Bonhage, “Signal quality of simultaneously recorded invasive and non-invasive EEG,” *Neuroimage*, vol. 46, no. 3, pp. 708–716, Jul. 2009, doi: 10.1016/j.neuroimage.2009.02.028.
- [31] X. Liu *et al.*, “Epileptic seizure detection with the local field potential of anterior thalamic of rats aiming at real time application,” in *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, 2011, pp. 6781–6784, doi: 10.1109/IEMBS.2011.6091672.
- [32] C. Im and J.-M. Seo, “A Review of Electrodes for the Electrical Brain Signal Recording,” *Biomed Eng Lett*, vol. 6, pp. 104–112, 2016, doi: 10.1007/s13534-016-0235-1.
- [33] I. Fried, K. A. MacDonald, and C. L. Wilson, “Single neuron activity in human hippocampus and amygdala during recognition of faces and objects,” *Neuron*, vol. 18, no. 5, pp. 753–765, May 1997, doi: 10.1016/S0896-6273(00)80315-3.
- [34] M. J. Jutras and E. A. Buffalo, “Synchronous neural activity and memory formation,” *Current Opinion in Neurobiology*, vol. 20, no. 2. Elsevier Current Trends, pp. 150–155, Apr. 01, 2010, doi: 10.1016/j.conb.2010.02.006.
- [35] E. Spaak, M. Bonnefond, A. Maier, D. A. Leopold, and O. Jensen, “Layer-specific entrainment of gamma-band neural activity by the alpha rhythm in monkey visual cortex,” *Curr. Biol.*, vol. 22, no. 24, pp. 2313–2318, 2012, doi: 10.1016/j.cub.2012.10.020.
- [36] M. Ramirez-Elias, V. Arce-Guevara, M. O. Mendez, and A. Alba, “Evaluation of EEG signal

- during the A-phases of the cycling alternating pattern using principal component analysis,” Feb. 2016, pp. 1–3, doi: 10.1109/embc.2015.7395938.
- [37] N. Axmacher, F. Mormann, G. Fernández, M. X. Cohen, C. E. Elger, and J. Fell, “Sustained neural activity patterns during working memory in the human medial temporal lobe,” *J. Neurosci.*, vol. 27, no. 29, pp. 7807–7816, Jul. 2007, doi: 10.1523/JNEUROSCI.0962-07.2007.
 - [38] Y. Lerner, B. Epshtein, S. Ullman, and R. Malach, “Class information predicts activation by object fragments in human object areas,” *J. Cogn. Neurosci.*, vol. 20, no. 7, pp. 1189–1206, Jul. 2008, doi: 10.1162/jocn.2008.20082.
 - [39] Y. Holzinger, S. Ullman, D. Harari, M. Behrmann, and G. Avidan, “Minimal Recognizable Configurations Elicit Category-selective Responses in Higher Order Visual Cortex,” *Washington, DC APA, Guidel. Dev. Panel Treat. Posttraumatic Stress Disord. Adults.*, 2019, doi: 10.1162/jocn.
 - [40] D. Kobak *et al.*, “Demixed principal component analysis of neural population data,” *Elife*, vol. 5, Apr. 2016, doi: 10.7554/eLife.10989.
 - [41] H. Abdi, “The Bonferonni and Šidák Corrections for Multiple Comparisons.” Accessed: Aug. 27, 2020. [Online]. Available: <http://www.utd.edu/>.
 - [42] J. Lawrence, “Familywise and per-family error rates of multiple comparison procedures,” *Stat. Med.*, vol. 38, no. 19, p. sim.8190, May 2019, doi: 10.1002/sim.8190.