Probabilistic decision-making in rodent foraging behaviour: contributions of distinct prefrontal cortex regions

Beatriz Simões Godinho
beatriz.godinho@tecnico.ulisboa.pt
beatriz.godinho@neuro.fchampalimaud.org

Instituto Superior Técnico, Lisboa, Portugal

November 2017

Abstract

Neurological disorders are considered the main cause of death and disability worldwide [1] in part due to the lack of information and understanding of how exactly the brain works. Progress made over the last decades revealed the importance of decision-making behaviour impairment in several mental disorders [2], and thus comprehending this behaviour is one of the chief goals of cognitive neuroscience. The study of the decision-making process in neuroscience has supported itself in foraging behaviour tasks to mimic the real-world decision process [3, 4], as it requires value-based strategies to evaluate and decide in a world containing various possible options [5]. In this work, mice performed a recent foraging task – flipping task – to test for the contributions of distinct prefrontal cortex (PFC) regions to the process of decision-making. The Anterior Cingulate Cortex (ACC) and the Orbitofrontal Cortex (OFC), which are two anatomically interconnected areas of the PFC, were inactivated through GABAergic interneurons' optogenetics stimulation in a VGat-ChR2 transgenic mouse line. The ACC and the OFC are associated with distinct yet sometimes overlapping roles in value-based decision-making tasks. However, it was possible to verify that both areas impact differently in the same behaviour. Evidences suggest ACC to be an important modulator in evidence accumulation, whereas OFC impairment translates in the disruption of the behavioural strategy. The flipping task proved to be efficiently designed to differentiate distinct brain regions functions and observe their impact in behavioural adjustments based on reward probability and travel cost variations.

Keywords: Value-based decision-making, Foraging, Flipping Task, VGat-ChR2 mouse line, Optogenetics, prefrontal cortex (PFC), Anterior Cingulate Cortex (ACC), Orbitofrontal Cortex (OFC)

1. Introduction

1.1. Decision-making and Value-based Tasks

The process of decision-making can be decomposed in four general steps: the recognition of the situation - or state -, the evaluation of action candidates - or options -, the selection of the option in reference to one's needs and the re-evaluation of the chosen action, based on the outcome. Individual goals reflect the decision maker's physiological and economical needs, but also external factors impact in the decision process. The environmental stochastic dynamics limit strongly our predictability of the future state, and the further ahead one tries to predict these hidden-states, the more difficult it becomes to succeed in this prediction [6].

One of the types of decison-making we can investigate - and accordingly design tasks to study it - is the value-based decision-making. Value-based tasks consider a complex set of situations where there is no single correct action and variables of the environment cannot be controlled by the subject [7]. This class of tasks is also associated with the psychological state of "uncertainty", since the underlying state of the environment is unknown to the decision maker [8].

In this light, the subject must adopt a strategy aimed at the maximization of the reward intake and for that several reward- and internal-related variables need to be evaluated [9]. For instance, these variables include the intention to make an action, the likelihood of that action resulting in a reward, the integration of perceptual signals that inform if that action will be rewarded, the probability and magnitude of reward associated with that specific action, the memory of the past action- and stimuli-outcome history, or some combination of all these factors [10, 11, 12].

In other words, a value-based decision-making task requires the subject to consider the different attributes of all available options, assess the value of each of those attributes and combine all these attributes into one coherent value representation that allows comparison with the other possible options [9]. This description comprises the main elements of *foraging* [3, 4], reason why this is one of the most used model-behaviours for studying this process [13].

Over the past decades several studies gave rise to the general consensus that reward magnitude and value attribution is represented in a small number of well-identified brain regions, interestingly, brain areas implicated in reward-seeking processes like *foraging* – the Basal Ganglia, the Anterior Cingulate and Orbitofrontal Cortices [14, 15, 16] and eventually the Lateral Intraparietal Area (LIP) and Dorsolateral Prefrontal Cortex (dlPFC) areas [11, 12].

1.2. ACC in the decision-making processes

In foraging tasks, animals must decide whether it is worth acting at all or evaluate whether it is worth continuing to engage in the current behaviour or to explore the existing alternatives. This decision-making pattern is recurrently

associated to ACC [17, 18].

Manipulations of the ACC affect subject's ability to initiate actions in the first place [19], their interpretation of action associated benefits and costs [20] and promotes action switching in relation to their values or to explore the alternative choice [21, 22, 23]. Evidence accumulation before switching was found to be analogous to the neural integrate-to-threshold mechanism, where higher gains in neuronal responses lead to a faster decision threshold achievement and increased travel-associated costs decrease neuronal gains rate to reach decision threshold [24].

In sum, we can easily find ACC associated with economic-value decisions, where it is implied in outcome monitoring and comparison as well as action associated values updating [25] for behaviour adjustment towards exploration contrarily to exploitation [17]. However, it might also reflect choice difficulty and alert/control engagement against default activity [26, 27]. Malfunctions in ACC are often associated with clinical diagnoses of depression, addiction, obsessive-compulsive disorder among others [28].

1.3. OFC in the decision-making processes

Discoveries in the OFC are often demonstrated experimentally either in reversal learning tasks or using reinforcer devaluation strategies [29, 30]. Results reveal that the acquisition of new associations is not affected by OFC lesions, but only for updating previously to lesion learned associations [31, 32].

Thinking of value-based decisions, OFC is linked to the process of value assignment, specifically coding for the different types of assignment related to a particular reward-type [33]. It is also considered a critical structure when decisions depend on the association of stimuli with reinforcement [34] and is identified as a distinct source of information integrator region that can provide outcome predictions based on confidence monitoring processes [35].

Generally, OFC is described as a powerful integrator based on perceptual evidence and memory that can have a crucial influence in value-based decision processes and learning [31, 36]. Also, because decision confidence is important for computing this value of decision outcome, it is believed that this region can monitor this same confidence level[35, 37]. Damage in human OFC has been described has enabling people to measure the consequences of their actions, such as in anti-social personality disorder [38]. Hyperactivation of the OFC combines in obsessive-compulsive behaviour and action-overthinking.

1.4. Optogenetic brain manipulation

In 1979, Francis Crick noticed one of the major necessities of neuroscience: to have a "method by which all neurons of just one type could be inactivated, leaving the others more or less unaltered" [39, p. 122]. Because of their lack of specificity for cell type and lack of temporal precision on the timescale of neural coding and signalling, respectively, electrodes and drugs were excluded for this effect over the years. At the same time, several techniques and scientific findings aligned on what ended up being called "optogenetics" [40, 41].

Optogenetics is a combination of genetic and optical methods that allow to achieve gain or loss of function of well-defined events in specific cells of living tissue, that not just neurons. It essentially lies on three core features: microbial opsins, the optogenetics actuators, neural engineering, for cell type-specific manipulation and neuron illuminaton, for spatiotemporally-resolved photostimulation [42].

1.4.1 Optogenetics stimulation in ChR2-positive GABAergic interneurons

Although different subcellular compartments of the cortical pyramidal (Pyr) cells are inhibited by distinct GABAergic neurons, the experimentally used GABAergic antagonists have too general action [43]. Thus, to overcome this weakness and enable selective manipulation of specific interneurons types for inhibition transmission, Atallah and his colleagues directly manipulated the activity of a genetically identified type of inhibitory interneurons to affect Pyr cells in response to visual stimuli. This was done targeting parvalbumin (PV)-expressing interneurons expressing Arch-GFP and ChR2-tdTomato, to respectively suppress and increase cell activity [44].

Relying on optogenetics, they found that PV cells perfom a linear operation on the response of Pyr cells in layer ⅔ of mouse V1. Simply said, PV cell spiking activation with ChR2 proteins linearly decreased Pyr cells activity, and the opposite for the Arch protein. In ChR2, maximum suppression of the Pyr cells means approximately zero spikes [44].

The described inhibition mechanism is the same used in the VGat-ChR2 trasngenic mouse line.

2. Materials and Methods

All experiments and procedures were reviewed and performed in accordance with the European Union Directive 2010/63/EU and the Champalimaud Centre for the Unknown Ethics Committee guidelines, and approved by the Portuguese Veterinary General Board, Direção-Geral de Alimentação e Veterinária (DGAV), approvals 0420/000/000/2011 and 0421/000/000/2016.

2.1. Animal Subjects

Twenty adult male mice from the in-house breeding colony were used, 10 for the ACC inactivation experiment and 10 to the OFC inactivation experiment. Because of optogenetic control, we used a mouse-line expressing an improved light-sensitive Channelrhodopsin-2/EYFP fusion protein (mhChR2::YFP).

Each inactivation experiment was performed on 10 VGAT-ChR2-EYFP line 8 transgenic mice [B6.Cg-Tg(Slc32a1-COP4*H134R/EYFP)8Gfng/J], from Jackson Laboratory [45]. For the ACC experiment, 7 heterozygous (Het) and 3 Control animals were used, while in the OFC we used 6 Het and 4 Control animals.

All animals were between ages 2 to 4 months old by the date of surgery, after which all animals were individually housed and kept under a normal 12 h light/dark cycle. Mice motivation was obtained via water deprivation having *ad libitum* food, both during training and testing. In the course of the experiments, animals' body weight was always kept higher than 85% of the initial weight.

2.2. ACC and OFC Surgeries

All surgical procedures for fibre implantation were carried out under aseptic conditions. Antibiotic (Enrofloxaxin, $5\,\mathrm{mg\,kg^{-1}}$) and analgesic (Buprenorphine, $0.1\,\mathrm{mg\,kg^{-1}}$) drugs were administered and dehydration was prevented by regular saline injection. Animals' eyes were covered

and protected by application of optical gel (Dexamytrex, $0.3\,\mathrm{mg\,g^{-1}}$ dexamethasone +$3\,\mathrm{mg\,kg^{-1}}$ of gentamicin). Inhalation of isoflurane (Vetflurane®, Virbac) was used as anaesthesia at 3 % during induction and 1 % to 2 % during surgery, with a 1.5 % mixture with O_2 and a flow rate of $0.8\,\mathrm{L\,min^{-1}}$.

During procedure, adjustments of isoflurane were made according to paw withdrawal reflex, and a body temperature of 37.5 °C was maintained. After craniotomy, zirconium fibre-optic cannulae (MFC _200/230-0.48 _3mm _ZF1.25(G)-FLT) were stereotactically (KOPF NEUROSTAR Stereotaxic) implanted according to ACC and OFC procedures, using a mouse brain atlas [46]. Fixation to the skull was done with dental cement (Super-Bond C&B).

ACC Procedure cannulae were implanted with a ±16 ° angle in each hemisphere targeting ACC brain region (anterior/posterior (AP): +1.90, media/lateral (ML): ±0.50, dorsal/ventral (DV): +1.75 from skull surface).

OFC Procedure cannulae were implanted vertically in each hemisphere targeting right above OFC brain region (AP: +2.90, ML: ±1.25, DV: +1.80 from skull surface).

2.3. Water Deprivation

In this work, animals were water deprived for three days prior to training, where day one was the removal of the water bottle and in the next two days they received controlled amounts of water: 800 µL and 600 µL, respectively. After training initiation, water availability was restricted to the behavioural sessions. On resting days, the mice always received 800 µL and 600 µL of water, on day one and on day two, respectively.

2.4. The Flipping Task

The flipping task, inspired in the switching task [47], consisted of tracking a single hidden variable in two mutually exclusive states - the left and right reward states - which corresponds to two distinct reward ports. At a certain point, poking on only one of the two reward ports is potentially rewarded - only one of the ports is on the high-state of reward.

The goal of the subject is to predict in which state it is - decision report is made by staying or switching ports. This prediction is based on stochastic evidence. Each port in the high-state gives 4 µL size reward with probability $\mathrm{P_{REW}}$. Furthermore, ports can change states with probability $\mathrm{P_{TRS}}$ – the port that was in the high-state alters to the low state, means $\mathrm{P_{REW}}$ equals zero, and the port that was in the low-state starts being potentially rewarded.

A sensor inside each port detects the animal's nose and validates the action as a poke when it stays in the port for at least 100 ms uninterruptedly. A series of consecutive pokes on one port was defined as a streak or a trial. Each session had a fixed duration of 25 min.

2.5. Optimal Performance

Being a foraging decision-making task, the aim is to optimize the reward rate. For that, subjects need to infer the hidden-state of the environment - are rewards on the left port or on the right port? -, and subsequently decide on their action – stay on the same port or switch. This statistical inference process is based on a series of rewards and omissions – evidences -, that give the subject information to calculate the instantaneous probability at a given time. Additionally, it requires establishing an appropriate threshold value for each situation that reflects the changing point - reward probability is so low that I am willing to switch ports.

But the way we interpret these signs is the key to correctly solve this problem. On one hand, rewards always indicate that the port is on the high-state, since they can only be obtained in the correct port. For this, rewards should always reset our beliefs about the state, independent of previous outcomes or any manipulations, like statistics or action cost. On the other hand, omissions can signal two very distinct states – (1) an unlucky omission poke, even though the port is on the high-state or (2) an omission poke at a depleted port when states have transitioned. For this, obtained omissions should be the factor impacting the instantaneous probability calculation, and subsequent behaviour adjustment to different environment manipulations.

Equation 1 describes the probability ratio PR of being in a low-state versus a high-state condition as a function of the number of consecutive omissions after the last reward.

$$PR = \frac{P(Low \mid n-1)}{P(High \mid n-1} = \frac{P_{TRS}(1-P_{REW})}{1-(1-P_{TRS})(1-P_{REW})}\left[\frac{1}{(1-P_{TRS})(1-P_{REW})}\right]^{n} - \frac{P_{TRS}(1-P_{REW})}{1-(1-P_{TRS})(1-P_{REW})} \quad (1)$$

where P_{TRS} is the port state transitioning probability, the P_{REW} is the reward probability and n is the number of consecutive omissions after the last reward. We can also calculate the instantaneous reward probability (P_{INST}) for the next poke at the same port as a function of the number of consecutive omissions after the last reward (n), knowing that the last poke was rewarded [Equation 2].

$$(P_{INST})_{n-1} = \big(1-P_{TRS}\big)^{n+1}\big(1-P_{REW}\big)^{n}P_{REW} \quad (2)$$

Note that the exponents of both terms are different as the last rewarded poke before a series of omissions could have already induced a state change.

Imperatively, the inferential process relies on the likelihood of omission occurrence while in the high-state, which is given by the probability of an omission at the correct port $(1-P_{REW})$ and the probability of still being at the correct port $(1-P_{TRS})$.

It should be mentioned that the likelihood ratio described in Equation 1 grows exponentially with the number of consecutive omissions after the last reward at a rate equal to:

$$r_{low/high} = -\log(1-P_{TRS}) - \log(1-P_{REW}) \quad (3)$$

and the instantaneous reward probability (Equation 2) declines exponentially at the same rate. In essence, both described parameters are expected to impact probability estimations on an equal basis.

2.6. Experimental Conditions

2.6.1 Statistics manipulation

The set of reward statistics was chosen for various reasons. Firstly, because behaviour is normatively affected by omissions with the exponential rate mentioned above (Equation 3), we expect to notice behavioural changes across the

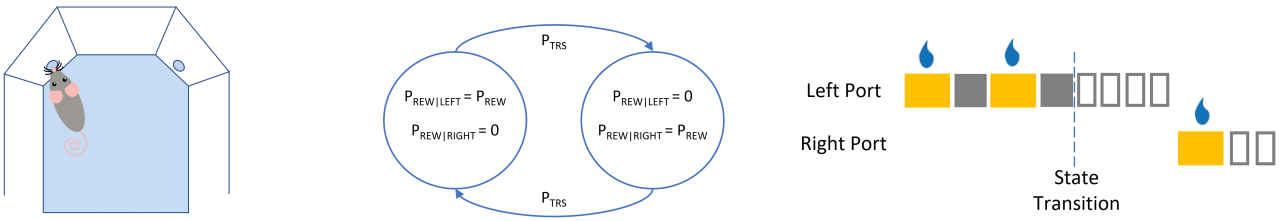

(a) (b) (c)

Figure 1: The flipping task. (a) Scheme of the experimental box used in the flipping task. A hexagonal black chamber with two reward ports was used in this task. (b) The subject must infer the hidden-state of the task in the two mutually exclusive ports. At a certain point in time, only one of the port is in the high state with reward probability P_{REW}. The port state can switch from high to low (no reward) with probability P_{TRS}, time at which the high-state transitions to the other port. (c) Typical trial of the flipping task, where the yellow squares represent rewarded pokes, grey squares represent non-rewarded pokes while the port is still in the high-state and white squares represent non-rewarded pokes when the port has already transitioned to the low-state.

entire dynamic range of P_{REW} and P_{TRS}, especially for higher values. Secondly, as P_{INST} is mathematical equally affected by changes in P_{REW} and P_{TRS}, these two variables were modified orthogonally to allow parameter comparison when it comes to behavioural adjustment. More, P_{REW} should be enough to guarantee subjects engagement during the task, as well as P_{TRS} should be fixed at somewhat high values (>20%) to produce an ample-enough dataset by keeping the trials (or streaks) relatively short.

Considering this, the three chosen statistics combinations are presented in Table 1.

Table 1: **Training an testing statistics used in the flipping task for reward probability manipulation.** Reward manipulation statistics (Stats) chosen for this work's protocols of the flipping task. Each statistic, or protocol, is composed of a reward probability P_{REW} and a transitioning state probability P_{TRS}. The 90-30 and 30-30 protocols were used for training purposes. The 30-30 and the 90-90 protocols were used for testing purposes.

Stats	P_{REW}	P_{TRS}
90 - 30	90%	30%
30 - 30	30%	30%
90 - 90	90%	90%

2.6.2 Travel distance manipulation

As seen before [24], evidence accumulation towards patch switching decision is also influenced by the distance between the patches: increasing this distance reflects in a greater effort (or cost) of travelling, delaying the decision of leaving the initial patch. Hence, one other interesting variable to manipulate in this experiment would be the distance that the mice had to travel between the two ports. Such condition was achieved by placing a 12 cm long transparent barrier between the ports, forcing the mice to go around it when switching sides.

2.6.3 Training and Testing Calendar

For the training and testing stages, the mice were divided into two blinds (two groups of 5 mice containing the same Het/Control ratio).

The mice were first trained in the flipping task for nine consecutive days, divided in two phases with two different statistics: phase one consisted of the first four days with $P_{REW} = 90\%$ and $P_{TRS} = 30\%$ and phase two consisted of the next five days with $P_{REW} = 30\%$ and $P_{TRS} = 30\%$. Because the latter involved an additional level of difficulty, which led to being more difficult to perform optimally, it required an extra day of training. Also, the last two days of training were done with the fibres connected to the cannulae and the masking light ON, so the mice would adapt to these new characteristics of the environment previous to actual laser stimulation. The size of reward during training was 4 µL and the time of a session varied between 25 min to 35 min, with a minimum reward consumption of 300 µL. After the training phase followed two days of resting after which the experimental tests started.

For the tests, the mice performed three distinct experimental blocks of eight days each, where the first two days were executed without photostimulation, in contrast with the next six days. In between every eight-day block, the mice experienced two resting days. Statistics $P_{REW} = 30\%$, $P_{TRS} = 30\%$ was assumed as the baseline statistics – second block of testing -, and so the referred manipulations of the statistics and the travel cost between ports were introduced in the first and third blocks of the experiment – the first block consisted of the baseline statistic with the barrier implemented and the third block consisted of the $P_{REW} = 90\%$, $P_{TRS} = 90\%$ statistics.

The last two blocks were counterbalanced, meaning the other blind of animals started with the $P_{REW} = 90\%$, $P_{TRS} = 90\%$ statistics as block one and barrier implementation as block three.

2.6.4 Box Design

The experimental setup consisted of two black hexagonal chambers (approximately $15 \times 15 \times 15$ cm) with two reward ports [Figure 1(a)]. Approximately 4 µL of water were delivered to the spout if a poke was determined to be rewarded. Overhead, eight blue light emitting diodes (LEDs) were used to mask photostimulation – "masking light" - and an overhead video camera (PlayStation Eye, Sony, JP) was placed to monitor behavioural performance during sessions (the videos were not taken for analysis purposes). Interfacing with the camera was done using Bonsai framework [48]. Behavioural boxes control was assured by a custom written script running on an Arduino® (Arduino Mega 2560, Arduino LLC) and human interfacing was made through a Python (Python Software Foundation. Python Language Reference, version 2.7 and 3.5.)

custom-written code.

2.6.5 Photostimulation

A 200 mW DPSS laser (Dhom Canada Co. Ltd) was used to deliver blue light (473 nm) to the setup. The laser output was controlled using an acousto-optic modulator (AOM) of an Arduino board. To lead the light to the implanted fibre-optic cannulae, were used optic fibres patch cords (200 µm diameter, NA 0.48, Doric Lenses Inc) and a split fibre patch cord to enable the ACC and OFC bilateral photostimulation. The power at the tip of the fibres was calibrated to approximately 3.6 mW, considering a maximum deviation power value of 10%. Trials were photostimulated with a 50% probability during the sessions, with pulses of 10 ms being delivered at 75 Hz for selected brain areas' excitation [44].

2.7. Histology

Cannula implantation site was confirmed via Magnetic Resonance Imaging (MRI) technique. For that, a cryogen-free permanent 1 Tesla Bruker ICON™ magnet was used, with a gradient strength of 450 mT m^{-1}. Due to the presence of head implants, a RF rat body coil was used. It was equipped with a water heating system in-line with a Thermo Scientific™ SAHARA PPO S5P Heated Bath Circulator set at 51 °C to maintain the animal at approximately 37.5 °C.

Animals were kept anaesthetized during the imaging by isoflurane (Vetflurane®, Virbac) inhalation: 4 % for induction and 2 % for the entire exam duration (1.5 % mixture with O_2 and a flow rate of 0.8 L min^{-1}).

The magnet was operated using MRI software package ParaVision® 6.0.1 pl3_ ICON, with a T2 RARE program set for collection of both coronal and transverse sectional images of the brain. Regular histology will also be performed.

2.8. Data Analysis

Behavioural data analysis was performed using custom-written scripts in Julia-0.6.0. An interval from 10 to 60 streaks was chosen to be ideal for data analysis of the behavioural results, once it represents the time when the animals are engaged in performing the task.

Results were represented as mean±s.e.m and statistical significance was accepted for $p < 0,05$. The statistical analysis was done in Julia-0.6.0 with Mixed-Models existing package, applying a linear mixed-effects model (lmm) for random effects that follow a Gaussian distribution for both unconditional and conditional distributions [49].

3. Results and Discussion

The results from the analysis of ACC and OFC inactivation behavioural experiments are showed and discussed in this chapter.

3.1. Data Validation

As mentioned in section 2.6.3, the testing phase consists of three blocks of eight days where the animals execute the task for the three different chosen protocols, that differ in the reward probability (P_{REW}), transition-state probability (P_{TRS}) and travel cost between the two ports (presence or absence of a barrier). These protocols are the (P_{REW} = 90%, P_{TRS} = 90%), the (P_{REW} = 30%, P_{TRS} = 30%) and the (P_{REW} = 30%, P_{TRS} = 30% with the addition of a 12 cm long barrier between the two ports). From now on, these protocols will be referred to as 90-90, 30-30 and 30-30 + Barrier for the sake of abbreviation.

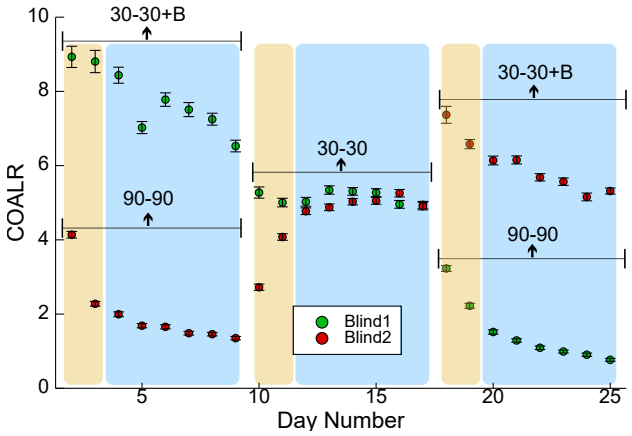

(a) ACC

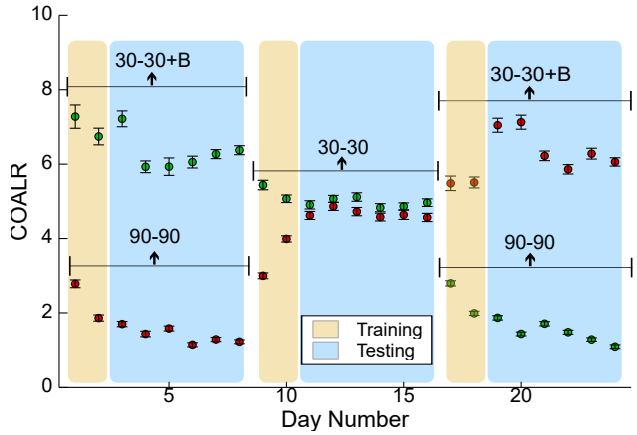

(b) OFC

Figure 2: General mouse behaviour while performing the flipping task in the ACC and OFC inactivation experiments. The consecutive omissions after the last reward (COALR) are plotted as a function of the days of the experiment for the two blinds, across mice for all sessions. The plateaus correspond to the different protocols adopted for the task, as indicated in the Figure. In light yellow, non-stimulated sessions (considered adaptation sessions) are indicated and each correspond to two days when the mice go through behavioural adaptation to the new protocol, to the fibre connection to the implanted cannulae and to the masking light of the Light-Emitting Diode (LED) set placed on the box's ceiling. In light blue, stimulated sessions (or testing sessions) are indicated and each correspond to six days of experimental testing of a certain protocols under optogenetic stimulation. n = 10 mice for each experiment.

In Figure 2 we can observe the general mouse behaviour while performing the testing stage of the flipping task, for both ACC [see Figure 2(a)] and OFC [see Figure 2(b)] inactivation experiments, divided in two batches for counterbalancing the protocols. In this study of mouse decision-making process, we do not want to consider the learning stage associated to this process. Hence, to analyse the obtained dataset from both inactivation experiments, we will only consider the steady-state behaviour sessions, i.e. stimulated sessions of the testing phase of the different protocols.

3.2. Mice used in the ACC and OFC inactivation experiments adjust their behaviour to the tested protocols

Mice are able to adjust their behaviour to the different protocols tested, albeit reward and state transition probabilities and travel cost manipulations. The same adaptation occurred for the animals that underwent fibre implantation surgery in the ACC and the OFC, independently of the optogenetic stimulation. Considering both surgery groups, and solely the unstimulated trials, we investigated whether the animals showed adaptation to the distinct protocols testing for a significant difference in the number of consecutive omissions after the last reward (COARL) among protocols using a linear mixed model $\left[lmm\,(COALR \sim\ 1 + Protocols + (1 \mid Mouse), p - value = \text{1e}{-99}\right]$.

3.3. Anterior Cingulate Cortex inactivation modulates evidence accumulation in value-based decision-making behaviour in mouse rodent model

3.3.1 ACC inactivation increases slows down evidence accumulation in mouse value-based decision-making behaviour

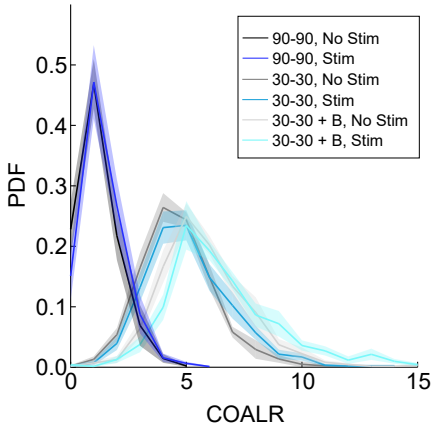

(a) Het.

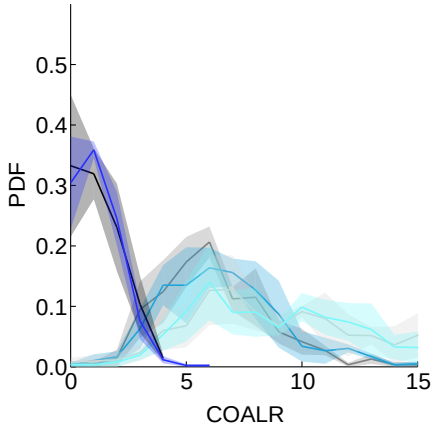

(b) Control.

Figure 3: ACC inactivation increases mouse patience in the flipping task. Probability Density Function (PDF) of the consecutive omissions after the last reward for the heterozygous (Het) (a) and Control (b) animals performing the flipping task during ACC inactivation sessions. Stimulated trials are represented in blue shades and unstimulated trials in black shades. The presented error bars correspond to s.e.m across mice (n = 7 Het + 3 Control).

A difference was observed in Het animals behaviour between stimulated and non-stimulated trials. Stimulation in Het animals increase COALR as shown by a shift to the right in the COALR distribution [see Figure 3(a)]. This genotype and stimulation interaction was significant for every protocol $\left(90-90: p-value = \text{1e}{-4}; 30-30: \text{1e}{-3}; 30-30+Barrier: p-value < \text{1e}{-4}\right)$.

In any of the protocols, the effect was proved not to be significant for stimulation independent of the genotype, therefore a simple artefact due to the laser light cannot explain the data.

The fact that animals were able to adjust their behaviour to the different protocols corroborates the idea that the ACC is up to a certain extent responsible for the integration of the action-outcome value representations – engage value, search value and costs [34, 50]. In the flipping task, the reward size is kept constant across protocols, so the benefits weight is constant [51]. The ACC must then use the weight of the cost associated with leaving a certain port – travel cost – and the value of leaving it too early - an error trial.

Then, it seems that the harder the protocol (i.e. the decision), the more control is necessary in the ACC to decide to switch to the alternative port. This idea resonates with the theory that ACC considers the perception of the difficulty associated to each protocol of the task and that foraging and difficulty values have been mistaken among them [26].

Overall, we can propose that the ACC slows down the decision process independently of the environmental characteristics, and favours the default behaviour, which is staying engaged at the same port.

3.3.2 ACC inactivation maintains mouse hidden-state notion in value-based decision-making behaviour

The hidden-state notion of the mice is not disrupted with ACC inactivation in this task. The animals can effectively reset the negative accumulation process every time they receive a reward (see Figure 4). The observation that the mean COALR in stimulated trial is independent (flat slope) of the position of the last reward in Figure 4(a) proves that the accumulation process is altered quantitatively and not qualitatively.

Likewise, we can conclude that the ACC does not disrupt the mouse behaviour. It acts as a modulator of the evidence accumulation process in the value-based decision-making behaviour.

3.3.3 Scale invariance of the data is preserved albeit ACC inactivation

According to Weber's law, we expect the scaling of the data to also be conserved in the stimulated trials of the flipping task, that reflects in a linear correlation between the standard deviation an the mean values of COALR.

In Figure 5 we can observe this representation for the Het animals in the stimulated and unstimulated trials. The linear variation of Standard Deviation (STD) with the Mean is true, except for the 90-90 data, where we can identify a small deviation. This has been seen before for Poisson distribution models due to discretization of the collected dataframes [52]. Nonetheless, we can still observe the linear pattern correlating STD and Mean values of COALR in this task, and for the 90-90 statistics the identified deviation is coherently present for both stimulated and non-stimulated trials.

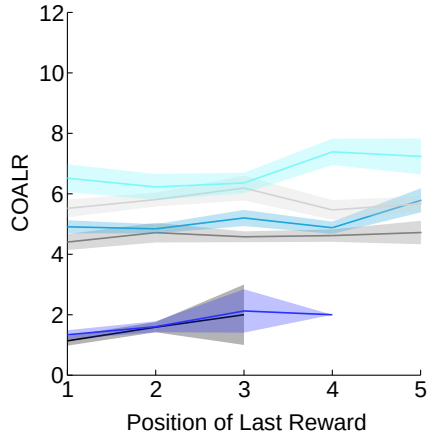

(a) Het.

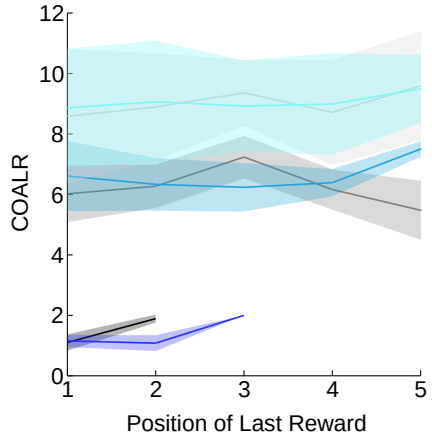

(b) Control.

Figure 4: ACC inactivation maintains mouse notion of the hidden-state. Representation of COALR as a function of the position of the last reward for the Het (a) and Control (b) mouse groups. Stimulated trials are represented in blue shades and unstimulated trials in black shades. The presented error bars correspond to s.e.m across mice (n = 7 Het + 3 Control).

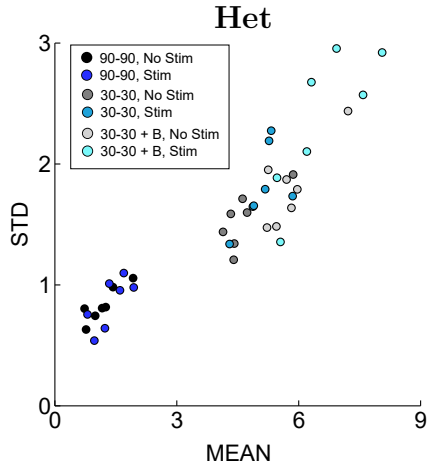

Figure 5: Linear correlation between the Mean and STD values of COALR in the ACC inactivation experiment. Linear correlation indicates scale invariance preservation of the collected dataset (n = 7 Het animals). Conservation of this property for the 90-90 protocol is not so evident, what is probably related to the binning of the data for this protocol [52].

3.4. Orbitfrontal Cortex inactivation disrupts strategy in value-based decision-making behaviour in mouse rodent model

3.4.1 OFC inactivation alters behavioural strategy in mouse value-based decision-making behaviour

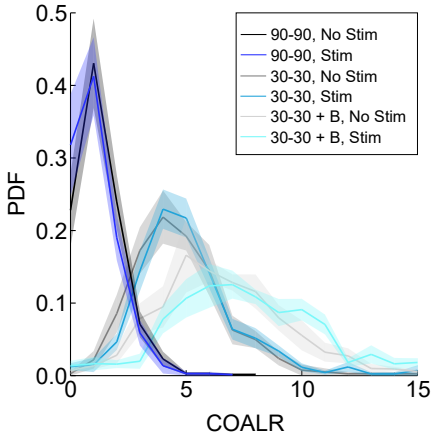

(a) Het.

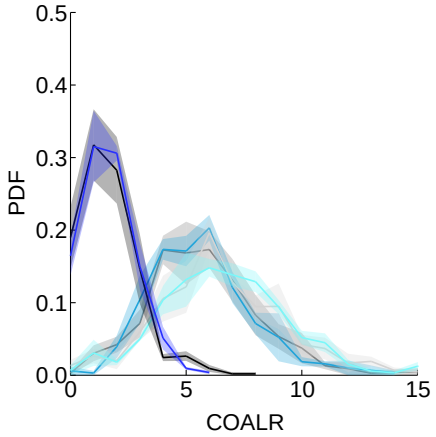

(b) Control.

Figure 6: OFC inactivation alters mouse behavioural strategy in the flipping task.PDF of the consecutive omissions after the last reward for the Het (a) and Control (b) mice Stimulated trials are represented in blue shades and unstimulated trials in black shades. The presented error shades correspond to s.e.m across mice (n = 6 Het + 4 Control).

When comparing the stimulated and the unstimulated trials for the Het animals (see Figure 6), the effect interacts differently with each protocol. The strongest and clearer effect is the increase in the number of COALR in the 30-30 + Barrier protocol (30-30 + Barrier: $p - value = 1e{-}4$). Contrarily, the distribution of COALR suggests a slight decrease in the same parameter for the 90-90 protocol (90-90: $p - value = 1e{-}2$). There was no significant effect of the stimulation in the baseline protocol (30-30: $p{-}value = n.s.$). Control animals show no effect with the stimulus [see Figure 6(b)], as expected.

3.4.2 OFC inactivation disrupts mouse hidden-state notion in value-based decision-making behaviour

Focusing attention on Figure 6, the preservation of the hidden-state notion is not apparent for the Het animals [see Figure 7(a)], as the slope that relates the COALR with the position of the last reward is not horizontal. This happens for the 30-30 + Barrier protocol, as well as for the 90-90 protocol, where the slope also seems to have

changed when compared to the control group.

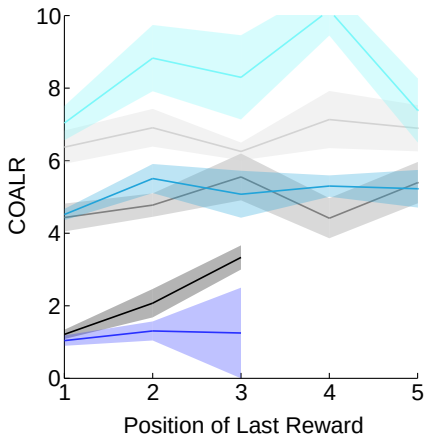

(a) Het.

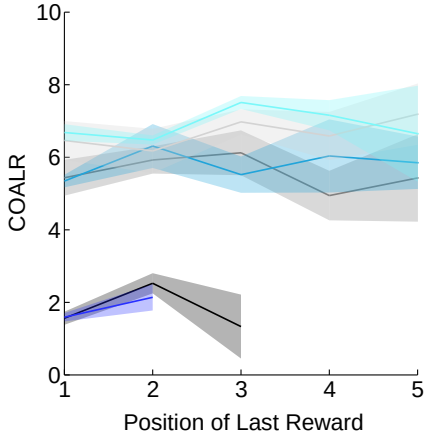

(b) Control.

Figure 7: OFC inactivation disrupts mouse hidden-state notion in the flipping task. Representation of the consecutive omissions after the last reward as a function of the position of the last reward for Het (a) and Control (b) animals performing the flipping task during OFC inactivation sessions. Stimulated trials are represented in blue shades and unstimulated trials in black shades. The presented error bars correspond to s.e.m across mice (n = 7 Het + 4 Control).

The normative solution involves a complete reset in the integration of COALR after a reward delivery. Therefore, the presence of a non-flat slope in Het animals [see Figure 7(a)] indicates that the mice strategy to change sides has changes and is now probably also depending on the reward history.

Indeed, for the 30-30 + Barrier protocol, it is visible an increase in the number of COALR whenever the mice received a bigger total number of rewards per trial (Figure 8). Also this effect seems to interact non-linearly with the difficulty of the protocol. The effect of inactivation of the OFC in the 90-90 protocol appears, once again, reversed – rewards seem to be more important for the mouse inference of the hidden-states of the task when the OFC is intact.

Moreover, the OFC inactivation does not seem to have any effect in the 30-30 protocol, suggesting the interesting possibility that a specific quantity computed by OFC plays a more or less central role depending on the environment statistics, in a way that does not scale linearly with the difficulty.

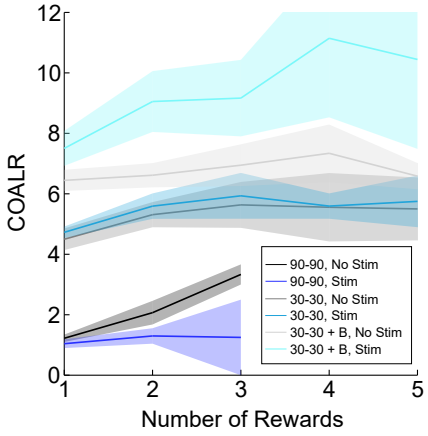

Figure 8: OFC inactivation alters mouse notion of the hidden-state. The number of COALR varies with the variation of the number of consecutive rewards the mouse receives in a trial (n = 6 Het). This effect can be observed for the 30-30 + Barrier in the positive direction and for the 90-90 protocol in the opposite direction.

In face of these results, we can assume that the strategy of the mice when performing the flipping task strongly involves the OFC region to be computed. The simple inactivation of this area enables to quickly investigate and potentially exclude the involvement of different brain regions in a given task. Nevertheless, to better define and precisely clarify what role this region plays in decision process, it requires the use of different approaches capable of monitoring the natural activity of the neurones during the task.

3.5. Histology Results

In this work, histological evaluation of the fibre implant location was performed through MRI scanning of the mice brains. For each mouse, coronal- (see Figure 9)and transverse-section image slices of the brain were fitted to the atlas [46] to determine the exact coordinates of the cannulae implantation.

ACC cannula coordinates were identified to be ranging from AP +1.68 to +1.93, ML ±0.20 to ±0.80 and DV +1.75 to +2.00, while OFC cannula coordinates were identified to be ranging from AP +2.50 to +3.00, ML ±1.05 to ±1.27 and DV +1.70 to +1.90, which correspond to the targeted areas.

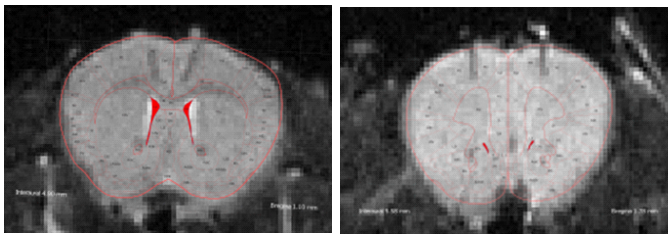

(a) Coronal section. (b) Coronal section.

Figure 9: MRI histology example images for ACC and OFC inactivation experiment mice. Histology was performed in living mice after all the experimental tests using a T2 Rare customized program. Transverse and coronal images for examples mice of the ACC and OFC inactivation experiments were used to confirm cannulae positioning in the brain.

4. Conclusions

Work from the World Health Organization has shown that neurological disorders are the leading source of years of life lost to death or disability, affecting up to one billion people [1, 41].

Decision-making is recently a major focus in cognitive neuroscience, in part because its impairment is strongly associated to several mental disorders [2]. Tremendous progress has been made in comprehending the mechanisms underlying decision-making – the elaboration of successive paradigms that allow the transposition of the essential questions to the laboratory [17, 53, 54], the insurgence of brain manipulation techniques that enable the testing of those questions [40, 55, 56], the restriction of search to a small number of brain regions [22, 25, 57] and construction of computational models that help in the understanding of the steps realized in the brain [27, 58]. Nonetheless, connections between these elements, to enable possible therapeutic applications, remain unclear.

Therefore, this work aimed to test a novel foraging task's capacity to reliably distinguish different brain region's roles in decision-making behaviour. For that, two PFC regions pointed to be relevant for the performance of the behaviour of interest were chosen - the anterior cingulate cortex (ACC) and the orbitofrontal cortex (OFC) [59, 60]. Optogenetic inactivation (see section 1.4.1) of both areas was conducted in separate batches of a transgenic VGat-ChR2 mouse line (see section 2.1).

For the ACC inhibition experiment, the animals increased significantly the number of omissions performed after the last reward (COALR) for the three presented protocols [see Figure 3(a)]. Additionally, it was possible to verify that the animals did not lose their notion of the existence of hidden-states in the task, as they continued to reset omission-pokes counting for every received reward [see Figure 4(a)]. These results revealed ACC to be a modulator of the accumulation of evidence process, that favours the default action / behaviour when inactivated [17, 60].

In the OFC inactivation experiment, animals significantly increased the number of COALR for the harder protocol and decreased it for the easier protocol presented [see Figure 6(a)]. There were no significant changes in behaviour for the baseline protocol (30-30). Also, the hidden-state notion of the OFC-inactivated animals were corrupted [see Figure 7(a)], as the reset of evidence accumulation was no longer verified for every collected reward. The value animals attributed to omissions and rewards changed – on the harder protocol, animals associated an increased value to the rewards when stimulated, whereas the opposite occurred for the easier protocol (see Figure 8). Hence the OFC appears to compute some central quantity whose absence has strong impact in the computation of the decision process in ways that change in respect of the environment difficulty.

In face of these results, we can conclude that the flipping task has shown to be sensitive enough to discriminate the contribution of distinct brain regions of a rodent performing the same value-based foraging task. Furthermore, it allows the observation of the different strategies adopted in the different situations (probability of reward and travel cost variations) and signals up to what extent the studied brain regions interfere in the process of decision-making.

Eventually, this task would need some adjustments to enable the study of this behaviour in a real-time scale. This could be done by introducing animal's positioning and time tracking during the performance, so to calculate interpoke-intervals and travel speed between ports, imediately providing more and valuable information about the decision-making behaviour. Analysis of such factors could be useful to explore animal's confidence level associated to their choice [35, 37].

Nevertheless, even with these modifications, the flipping task could not reveal ACC and OFC true relevance for the decision-making process. To better define and precisely clarify what roles these regions play in this behaviour, it requires the use of different approaches and tasks, capable of monitoring the natural activity of the neurones.

Future goals in the ACC research could contemplate a task design capable of disambiguate foraging and difficulty values representation, two popular yet contradictory theories proposed for its function [17, 26]. For the OFC, one can consider that the reason why its inactivation disrupted mouse behaviour lies with the learning and relearning paradigm [31, 32]. Then, it would be interesting to test for OFC response in a reversal learning environment. Electrophysiology recordings in both areas would be highly beneficial for understanding these regions.

To conclude, this work is not just a contribution for the differentiation of two PFC regions in a value-based foraging task. It is also an input in characterizing the elements that compose decision-making behaviour and one of the billions of steps in the long path to unravel the real mechanisms and functions of the brain.

Acknowledgements

The author would like to acknowledge the Principal Investigator (PI) of the Systems Neuroscience Lab, Zachary Mainen, for the amazing opportunity, and the supervisors Eran Lottem, for the frequent guidance and trust, and Patrícia Figueiredo for the support. To all members of Mainen Lab, in particular Dario Sarra, for the introduction to the experimental techniques and assistance during lab experiments, and Pietro Vertechi, for the introduction to computational modelling and data analysis. The author also wants to thank the PI of the Neuroplasticity and the Neural Activity Lab, Noam Shemesh, and his team, for the MRI teachings and collaboration.

References

[1] WHO Geneva. WHO — Neurological disorders affect millions globally: WHO report. *WHO*, 2007.

[2] P. S. Sachdev and G. S. Malhi. Obsessive – compulsive behaviour : a disorder of decision-making. *Australian and New Zealand Journal of Psychiatry*, 39: 757–763, 2005.

[3] D. Kramer. Foraging Behavior. *Evolutionary Ecology Concepts and Case Studies*, 00:232–238, 1961.

[4] D. W. Stephens and J. R. Krebs. *Foraging Theory*, volume 121. Princeton University Press, 1986.

[5] P. W. Glimcher. Indeterminacy in Brain and Behavior. *Annual Review of Psychology*, 56(1):25–56, feb 2005.

[6] K. Doya. Modulators of decision making. *Nature Neuroscience*, 11(4):410–416, apr 2008.

[7] J. V. Neumann and O. Morgenstern. *Theory Of Games And Economic Behavior*. Princeton University Press, 1944.

[8] M. L. Platt and S. A. Huettel. Risky business: the neuroeconomics of decision making under uncertainty. *Nature Neuroscience*, 11(4):398–403, apr 2008.

[9] D. J. Levy and P. W. Glimcher. The root of all value: a neural common currency for choice. *Current Opinion in Neurobiology*, 22(6):1027–1038, dec 2012.

[10] J. I. Gold and M. N. Shadlen. Representation of a perceptual decision in developing oculomotor commands. *Nature*, 404(6776):390–394, 2000.

[11] M. C. Dorris and P. W. Glimcher. Activity in Posterior Parietal Cortex Is Correlated with the Relative Subjective Desirability of Action. *Neuron*, 44(2):365–378, oct 2004.

[12] L. P. Sugrue, G. S. Corrado, and W. T. Newsome. Matching behavior and the representation of value in the parietal cortex. *Science (New York, N.Y.)*, 304 (5678):1782–1787, jun 2004.

[13] J. M. Wolfe. When is it time to move to the next raspberry bush? Foraging rules in human visual search. *Journal of Vision*, 13(3):10–10, jan 2013.

[14] W. Schultz, P. Dayan, and P. R. Montague. A neural substrate of prediction and reward. *Science (New York, N.Y.)*, 275(5306):1593–9, mar 1997.

[15] L. Tremblay and W. Schultz. Reward-related neuronal activity during go-nogo task performance in primate orbitofrontal cortex. *Journal of Neurophysiology*, 83(4):1864–76, 2000.

[16] K. Watanabe, J. Lauwereyns, and O. Hikosaka. Neural correlates of rewarded and unrewarded eye movements in the primate caudate nucleus. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 23(31):10052–7, nov 2003.

[17] N. Kolling, T. E. J. Behrens, R. B. Mars, and M. F. S. Rushworth. Neural Mechanisms of Foraging. *Science*, 336(6077):95–98, apr 2012.

[18] S. R. Heilbronner and B. Y. Hayden. Dorsal Anterior Cingulate Cortex: A Bottom-Up View. *Annual Review of Neuroscience*, 39(1):149–170, 2016.

[19] D. Thaler, Y.-C. Chen, P. D. Nixon, C. E. Stern, and R. E. Passingham. The functions of the medial premotor cortex. *Experimental Brain Research*, 102 (3):445–460, 1995.

[20] K.-i. Amemori and A. M. Graybiel. Localized microstimulation of primate pregenual cingulate cortex induces negative decision-making. *Nature Neuroscience*, 15(5):776–785, apr 2012.

[21] S. W. Kennerley, M. E. Walton, T. E. J. Behrens, M. J. Buckley, and M. F. S. Rushworth. Optimal decision making and the anterior cingulate cortex. *Nature neuroscience*, 9(7):940–7, jul 2006.

[22] D. G. Tervo, M. Proskurin, M. Manakov, M. Kabra, A. Vollmer, K. Branson, and A. Y. Karpova. Behavioral Variability through Stochastic Choice and Its Gating by Anterior Cingulate Cortex. *Cell*, 159(1): 21–32, sep 2014.

[23] T. Kawai, H. Yamada, N. Sato, M. Takada, and M. Matsumoto. Roles of the Lateral Habenula and Anterior Cingulate Cortex in Negative Outcome Monitoring and Behavioral Adjustment in Nonhuman Primates. *Neuron*, 88(4):792–804, nov 2015.

[24] B. Y. Hayden, J. M. Pearson, and M. L. Platt. Neuronal basis of sequential foraging decisions in a patchy environment. *Nature Neuroscience*, 14(7):933–939, jun 2011.

[25] M. E. Walton, J. T. Devlin, and M. F. S. Rushworth. Interactions between decision making and performance monitoring within prefrontal cortex. *Nature neuroscience*, 7(11):1259–65, nov 2004.

[26] A. Shenhav, M. A. Straccia, J. D. Cohen, and M. M. Botvinick. Anterior cingulate engagement in a foraging context reflects choice difficulty, not foraging value. *Nature Neuroscience*, 17(9):1249–1254, jul 2014.

[27] M. M. Botvinick and J. D. Cohen. The Computational and Neural Basis of Cognitive Control: Charted Territory and New Frontiers. *Cognitive Science*, 38(6):1249–1285, aug 2014.

[28] R. D. Baler and N. D. Volkow. Drug addiction: the neurobiology of disrupted self-control. *Trends in molecular medicine*, 12(12):559–66, dec 2006.

[29] G. Schoenbaum, a. a. Chiba, and M. Gallagher. Neural encoding in orbitofrontal cortex and basolateral amygdala during olfactory discrimination learning. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 19(5):1876–84, mar 1999.

[30] A. N. Hampton, P. Bossaerts, and J. P. O'Doherty. The Role of the Ventromedial Prefrontal Cortex in Abstract State-Based Inference during Decision Making in Humans. *Journal of Neuroscience*, 26(32): 8360–8367, aug 2006.

[31] G. Schoenbaum, S. L. Nugent, M. P. Saddoris, and B. Setlow. Orbitofrontal lesions in rats impair reversal but not acquisition of go, no-go odor discriminations. *Neuroreport*, 13(6):885–890, may 2002.

[32] Y. Chudasama and T. W. Robbins. Dissociable contributions of the orbitofrontal and infralimbic cortex to pavlovian autoshaping and discrimination reversal learning: further evidence for the functional heterogeneity of the rodent frontal cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 23(25):8771–80, sep 2003.

[33] C. Padoa-Schioppa and J. A. Assad. Neurons in the orbitofrontal cortex encode economic value. *Nature*, 441(7090):223–226, may 2006.

[34] P. H. Rudebeck, T. E. Behrens, S. W. Kennerley, M. G. Baxter, M. J. Buckley, M. E. Walton, and M. F. S. Rushworth. Frontal cortex subregions play distinct roles in choices between actions and stimuli. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 28(51):13775–85, dec 2008.

[35] A. Lak, G. M. Costa, E. Romberg, A. A. Koulakov, Z. F. Mainen, and A. Kepecs. Orbitofrontal Cortex Is Required for Optimal Waiting Based on Decision Confidence. *Neuron*, 84(1):190–201, oct 2014.

[36] R. C. Wilson, Y. K. Takahashi, G. Schoenbaum, and Y. Niv. Orbitofrontal Cortex as a Cognitive Map of Task Space. *Neuron*, 81(2):267–279, jan 2014.

[37] A. Kepecs, N. Uchida, H. A. Zariwala, and Z. F. Mainen. Neural correlates, computation and behavioural impact of decision confidence. *Nature*, 455 (7210):227–231, sep 2008.

[38] A. Schnider. Orbitofrontal Reality Filtering. *Frontiers in Behavioral Neuroscience*, 7, 2013.

[39] F. H. C. Crick. Thinking about the Brain. *Scientific American*, 241(3):219–232, sep 1979.

[40] E. S. Boyden, F. Zhang, E. Bamberg, G. Nagel, and K. Deisseroth. Millisecond-timescale, genetically targeted optical control of neural activity. *Nature Neuroscience*, 8(9):1263–1268, sep 2005.

[41] K. Deisseroth. Controlling the brain with light. *Scientific American*, 303(5):48–55, nov 2010.

[42] K. Deisseroth. Optogenetics: 10 years of microbial opsins in neuroscience. *Nature Neuroscience*, 18(9): 1213–1225, aug 2015.

[43] S. Katzner, L. Busse, and M. Carandini. GABAA inhibition controls response gain in visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 31(16):5931–41, apr 2011.

[44] B. V. Atallah, W. Bruns, M. Carandini, and M. Scanziani. Parvalbumin-Expressing Interneurons Linearly Transform Cortical Responses to Visual Stimuli. *Neuron*, 73(1):159–170, jan 2012.

[45] The Jackson Laboratory. Mouse Strain Datasheet - 014548 - B6.Cg-Tg(Slc32a1-COP4*H134R/EYFP)8Gfng/J, 2017. URL `https://www.jax.org/strain/014548`.

[46] K. Franklin and G. Paxinos. *The Mouse Brain in Stereotaxic Coordinates, Compact: The Coronal Plates and Diagrams*. Academic Press, 3rd edition, 2007.

[47] E. Lottem, D. Banerjee, P. Vertechi, D. Sarra, M. oude Lohuis, and Z. F. Mainen. Activation of serotonin neurons promotes active exploitation in a probabilistic foraging task. *doi.org*, page 170472, jul 2017. doi: 10.1101/170472.

[48] G. Lopes, N. Bonacchi, J. FrazÃ£o, J. P. Neto, B. V. Atallah, S. Soares, L. Moreira, S. Matias, P. M. Itskov, P. A. Correia, R. E. Medina, L. Calcaterra, E. Dreosti, J. J. Paton, and A. R. Kampff. Bonsai: an event-based framework for processing and controlling data streams. *Frontiers in Neuroinformatics*, 9, apr 2015.

[49] The Julia Language. MixedModels.jl - Linear Mixed-Effects Models, 2017. URL `http://dmbates.github.io/MixedModels.jl/latest/optimization.html{#}Linear-Mixed-Effects-Models-1`.

[50] C.-H. Luk and J. D. Wallis. Choice Coding in Frontal Cortex during Stimulus-Guided or Action-Guided Decision-Making. *Journal of Neuroscience*, 33(5):1864–1871, jan 2013.

[51] M. Walton, S. Kennerley, D. Bannerman, P. Phillips, and M. Rushworth. Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. *Neural Networks*, 19(8):1302–1314, oct 2006.

[52] M. J. Berry, D. K. Warland, and M. Meister. The structure and precision of retinal spike trains. *Proceedings of the National Academy of Sciences*, 94(10): 5411–5416, may 1997.

[53] D. Stephens, J. Brown, and R. Ydenberg. *Foraging: Behaviour and Ecology.*, volume 1. Foraging Behavior and Ecology Edited by David W. Stephens, Joel S. Brown, and Ronald C. Ydenberg The University of Chicago Press, 2007.

[54] J. I. Gold and M. N. Shadlen. The Neural Basis of Decision Making. *Annual Review of Neuroscience*, 30 (1):535–574, jul 2007.

[55] I. C. Farber and A. Grinvald. Identification of presynaptic neurons by laser photostimulation. *Science (New York, N.Y.)*, 222(4627):1025–7, dec 1983.

[56] E. M. Callaway and L. C. Katz. Photostimulation using caged glutamate reveals functional circuitry in living brain slices. *Proceedings of the National Academy of Sciences*, 90(16):7661–7665, aug 1993.

[57] T. a. Stalnaker, N. K. Cooch, and G. Schoenbaum. What the orbitofrontal cortex does not do. *Nature Neuroscience*, 18(5):620–627, apr 2015.

[58] A. Rangel and T. Hare. Neural computations associated with goal-directed choice. *Current Opinion in Neurobiology*, 20(2):262–270, apr 2010.

[59] L. K. Fellows. The Cognitive Neuroscience of Human Decision Making: A Review and Conceptual Framework. *Behavioral and Cognitive Neuroscience Reviews*, 3(3):159–172, 2004.

[60] M. F. S. Rushworth, T. E. J. Behrens, P. H. Rudebeck, and M. E. Walton. Contrasting roles for cingulate and orbitofrontal cortex in decisions and social behaviour. *Trends in Cognitive Sciences*, 11(4):168–176, apr 2007.