UNIVERSITAS STUDIORUM ATHESINA

CIMeC
Center for Mind/Brain Sciences

Doctoral school in Cognitive and Brain Sciences

XXIX cycle

Ph.D. Dissertation

The impact of reward and punishment on visual attention during naturalistic visual search: valence or salience?

Ludwig Barbaro

Advisor: Prof. Clayton Hickey

Co-advisor: Prof. Marius Peelen

Table of contents

"Sen minkä ilotta oppii, sen suruttaa unohtaa.”

What one learns without joy, one forgets without grief (Finnish proverb).

Chapter 1

How do reward and punishment influence visual representation?

“The Emotions are all those feelings that so change men as to affect their judgements, and that are also attended by pain or pleasure. Such are anger, pity, fear and the like, with their opposites.” (Aristotle, Rhetoric)

Pleasure and pain have always been recognized as two fundamental determinants of learning, as we can already appreciate in these words by Aristotle. It was many years after, though, through Pavlov’s and Thorndike’s studies, that this intuition started to be investigated in a more systematic way. We owe to the former the discovery of the link between conditioned stimuli and unconditioned response, in the form of classical conditioning (Pavlov, 1927), and to the latter the law of effect, which lies at the basis of operant conditioning (Thorndike, 1911). While the connection between affect and response selection has so far been the object of a long-lasting examination, much more recent is the analysis of the parallel relationship between affect and perceptual attention. In this respect, only in the last few years it has been demonstrated that stimuli associated with reward obtain prioritized visual processing. On the contrary, less study has been devoted to the motivational counterpart of this dynamic, punishment.

In this thesis, I am going to address the issue of how both reward and punishment affect visual representation, by the use of a visual search paradigm performed in naturalistic scenes. More specifically, I am going to approach these two motivational forces by referring to their constituting dimensions, namely valence and salience. In fact, while being positioned in oppositely valenced space, these two outcomes share the burden of signaling stimuli with strong behavioral importance, coding therefore for motivational salience. After reviewing the main existing literature in this field (Chapter 1), I will describe a series of studies which examine how reward and punishment impact visual attention through the analysis of behavioral measures (Chapter 2) and fMRI activation (Chapter 3 and 4). The main idea resulting from these studies is that, in spite of what a rational approach to the problem would suggest, automatic visual attention does not process these two outcomes according to a salience, but rather through a valence pattern.

Behavioral effects of reward and punishment on visual attention

Given the limited amount of resources of our nervous system, sensory stimuli need to compete in order to undergo further processing and to reach representation. According to the biased competition model, attention is the mechanism through which this conflict is resolved in favour of one stimulus over the others (Desimone & Duncan, 1995). A series of factors modulate this process: a bottom-up attentional bias in the first place, which exploits low-level characteristics of the visual scene. In the second place, a top-down attentional control is needed in order to establish what is relevant at any given moment for a specific type of task, giving rise to the attentional template. This holds true for many different types of selection, such as selection based on spatial location, on features, or on objects. More recently, another, apparently additional, type of bias of attention has been identified, which is the one exerted by reward and in general by different types of motivational outcomes.

In general terms, the possibility of gaining a reward enhances performance in many different tasks, both in terms of improved response selection and perceptual processing. In a series of studies within the field of visual attention, reward cues have been shown to determine an increase in accuracy (Roesch & Olson, 2003) or perceptual sensitivity (Engelmann, Damaraju, Padmala, & Pessoa, 2009), or a decrease in reaction times (RTs) without a correspondent decrease of speed-accuracy tradeoff (Roesch & Olson, 2003; Taylor, Welsh, Wager, Phan, Fitzgerald, & Gehring, 2004; Small, Gitelman, Simmons, Bloise, Parrish, & Mesulam 2005). Given the proactive nature of these tasks (reward contingencies were in the above-mentioned cases known since the beginning of the trial), an important question becomes whether this behavioral facilitation arises as a result of a deliberate, strategic planning, or as an automatic tendency to prioritize rewarding stimuli. Moreover, whether volitional or otherwise, at which level of the neural processing chain does this benefit in performance arise? Is it only the product of an improvement in motor selection, configuring then itself simply as a specific type of instrumental conditioning, or does it occur because of an actual increase in the visual saliency of the stimulus? Which are the effects, and also the side effects, of this type of mechanism? Furthermore, can it be classified as an additional type of attentional control, or is it a subcomponent of one (or both) of the two already mentioned types of bias, namely the endogenous and the exogenous one?

Benefits and costs of reward in visual attention

With the aim of dissociating explicit and implicit effects of reward, Kristjánsson and colleagues (Kristjansson, Sigurjónsdóttir, & Driver, 2010) performed a study where subjects had to look for color singletons and report the orientation of a notch on the target. Two colors were paired respectively with a high or a low amount of reward in a probabilistic fashion, but, differently from the other study, this association was not communicated to participants in advance. Instead of simply considering accuracy and RTs for high and low reward trials, they focused their analysis on priming effects, i.e. on the benefit in performance when target was repeated. More specifically, priming effects in this type of parallel search, i.e. priming of pop-out (PoP), have been shown to be a product of a short-term implicit memory system, and to be resistant to top-down strategies (Maljkovic & Nakayama, 1994; 2000). First of all, subjects showed better performance for the highly rewarded color. Moreover, beyond a replication of the priming of pop-out effects after target repetition, data showed a critical interaction between reward level and target color repetition, as the product of enhanced priming after a rewarding outcome. The extent of this priming was both an effect of search history and of last trial outcome, showing a further interaction between reward expectation and actual reward achievement. Finally, all these effects tracked changes in reward contingency.

This example only shows advantageous effects due to reward, while, in a study by Krebs and colleagues, we can also appreciate detrimental effects of this implicit association in a version of the classic color-naming Stroop task (Krebs, Boehler, & Woldorff, 2010). Here, subjects viewed words representing colors, written in different inks. They had to respond to the color of the ink and at the same time ignore the semantic meaning of the word, which could have either been congruent (i.e. green written in green) or incongruent (green written in red). Among four different colors, two were linked to a potential reward, and the other two were instead neutral. The rewarded color determined a decrease in RT when present in the relevant dimension (ink color), and an RT cost when presented in the irrelevant dimension (word meaning). This is an instance of a situation where the attentional bias created by reward starts to show its negative side effects.

A perspective on automaticity is given by Bijleveld, in an experiment where subjects had to solve a mathematical problem, after receiving a cue which signaled a forthcoming high or low reward trial (Bijleveld, Custers, & Aarts, 2010). Crucially, the cue could be administered either supraliminally (cue duration 300 ms) or subliminally (17 ms). Final outcome depended both on RTs and accuracy, which were then analyzed, in both presentation conditions, for high and low reward trials. A dissociation between outcome awareness and adopted strategy was made evident: in supraliminal trials, subjects were slower but more accurate in the high reward compared to the low reward condition. On the contrary, they showed similar accuracy but faster RTs for high reward trials in the subliminal condition. The conclusion of the experimenters is that reward at stake causes an increase in effort in the task, whether outcome is consciously perceived or not. This is translated into a decrease of RTs in subliminal conditions, while, in supraliminal ones, there is an additional strategic component which modulates speed-accuracy tradeoff by increasing accuracy at the cost of RT.

Effects of inconsistent schedules of reward

Although giving hints regarding how reward works in biasing attention, all the aforementioned paradigms furnish to subjects an explicit (Krebs et al., 2010) or implicit (Kristjansson et al., 2010; Bijleveld et al., 2010) motivation to pay more attention to highly rewarded targets, leaving therefore open the possibility for an effect of exogenous attention on performance. We will now review a series of experiments where high and low rewards were administered in the absence of any consistent object- or feature-reward association.

In an experiment by Della Libera, subjects first viewed displays (primes) containing a big number composed of smaller digits, and had to focus either on the global or on the local features depending on a cue that they received at the beginning of the trial (Della Libera & Chelazzi, 2006). They received high or low reward for a correct performance, and immediately after accomplished a probe trial, where they always had to look at local features. RTs and accuracy measures were analyzed as a function of target repetition and reward level, and an interaction between reward value and priming was found for RTs. When the probe target had been the distractor in the immediately preceding prime, negative priming effects were found after high reward and positive priming effects after low reward. When the prime and probe target were the same, a benefit in RT was present after both high and low reward trials.

Coherent results were obtained by Hickey and colleagues, who developed a study based on the additional singleton paradigm of Theeuwes (1991), where subjects have to look for a shape singleton target in an array of uniformly colored distracters. A color singleton item can be also present in the array, creating an increased attentional capture with respect to other distractors. Color of targets and distracters can stay the same or swap between each other from trial to trial, causing respectively a benefit or a cost in performance known as inter-trial priming (Maljkovic & Nakayama, 1994). In this specific paradigm, participants could receive high or low-magnitude feedback following a random schedule, and subsequently the differential effects of the two types of outcome on inter-trial priming were analyzed (Hickey, Chelazzi, & Theeuwes, 2010a). High reward speeded up responses when color stayed the same, and slowed them down when color changed, while an opposite pattern was observed for low reward trials. In another experiment, this priming shown was to be correlated across subjects with high scores in a reward-seeking personality trait (Hickey, Chelazzi, & Theeuwes 2010b), as assessed by the BIS/BAS scale of Carver and White (Carver & White, 1994).

Following the same line, Hickey & van Zoest showed that reward history of a visual stimulus has a direct, low-level, and non-strategic influence on saccadic movements (Hickey & van Zoest, 2012). In every trial, participants had to orient their eyes from a central fixation point to a green or red target located at either the top center or bottom center of the screen, with correct deployment of the eyes to the target yielding a randomly assigned high or low amount of reward. Slightly to the left or right of the direct path between fixation and the target, a distractor of the opposite color was present. Here as well, color of targets and distractor could stay the same or swap from trial to trial. The influence of inter-trial priming was then analyzed, in terms of how much the distractor deviated target-directed saccades from their normal path. Trials were then binned according to two orthogonal experimental dimensions: amount of reward of preceding trial, and color stayed or swap, and were also separated into short, mid, and long latency conditions, reflecting the speed with which the saccade was initiated after stimulus onset. Short latency target-directed saccades in the high-reward/color swap condition were drawn more closely to the distractor than saccades in the high-reward/same colors condition. In contrast, long latency target-directed saccades clearly deviated further away. This shows that, at early stages of the deployment of attention, reward-associated stimuli drew attention automatically even when not task-relevant. In contrast, the opposite pattern was present in the long latency condition, probably as an effect of top-down control counteracting this automatic and, in this case, disadvantageous reward bias.

Reward-dependent learning

Another series of experiments has taken in consideration the long-lasting effects of reward delivery on visual processing, by looking at how previously learned object- and feature-reward associations are able to affect target selection during extinction. In a study by Della Libera and Chelazzi (Della Libera & Chelazzi, 2009), subjects underwent an extensive training, during which they performed a same/different judgment task on a series of shapes. Unbeknownst to participants, for some of these shapes there was a bias in the probability of receiving a high or low reward, either when acting as a target or as a distractor. Five days after training, the same shapes were used in a similar same/different judgment task and in a visual search task, but this time no feedback was delivered. Reward history affected response times during both tasks, such that, during the same/different judgment task, items which had been followed by favorable outcomes when acting as targets slowed down responses when presented as distracters, and vice versa. Conversely, items followed by high reward when presented as distracters during training were less easily selected during the same/different judgment task, and the opposite happened for shapes paired with low reward. During the visual search task, target selection was easier for shapes which had been paired with high reward and harder for the ones paired with low reward. Symmetrically, high reward associated shapes made search longer when presented as distracters, and the opposite pattern was found for low-reward associated shapes.

Along the same line, Anderson and colleagues designed a study where two colors were imbued with different amounts of reward during a training phase (Anderson, Laurent, & Yantis, 2011). This phase consisted of a visual search task for targets defined by color. Subsequently, subjects performed another visual search task, where they looked for a shape singleton among stimuli of various colors. On half of these trials, one of the two relevant colors was presented among the others. RTs were affected by presence of any of the two colors, but more dramatically by the high-value one. This effect was present even when tested several days after the initial training phase, showing a remarkable robustness.

To sum up, the attentional bias for rewarding stimuli appears to be a fast, automatic process, which acts primarily by affecting RTs (Bijleveld et al., 2010), proves to be highly efficient even when clear explicit cues regarding outcome delivery are missing (Kristjansson et al., 2010), but can at the same become possibly detrimental because of its long-lasting effects (Anderson et al., 2011) and its partial lack of specificity (Krebs et al., 2010). When administered during visual attention tasks, reward increases the saliency of objects and features paired with it. This process will, on one side, facilitate further selection of the same type of stimulus, but, on the other, it will also cause attentional capture when this stimulus will have to be ignored (Della Libera & Chelazzi, 2006, Hickey et al., 2010). The interplay between this automatic reward drive and the strategic top-down set is therefore a dynamic process, where the two components may go hand in hand and contribute synergistically to the same goal, but may also enter instead into conflict, and oppose each other. In the latter case, the reward bias has shown to manifest itself during early stages of attentional processing, while the re-establishment of a task-relevant attentional set occurs in later stages (Hickey & van Zoest, 2012). All these features suggest the idea that this bias may constitute an additional type of control with respect to bottom-up and top-down influences.

Carrot and stick

As mentioned before, in this thesis I am going to analyze how both reward and punishment are able to affect visual representation. Until here we have reviewed a series of visual attention paradigms considering how reward is able to affect behavioral performance. Much less work is present about the influence of the latter type of outcome, punishment, especially in the field of visual attention. Moreover, when taken in consideration, this motivational condition is generally intended in terms of threat. This is for example the case for a recent study by Schmidt, who investigated the effects of a threatening stimulus on the deployment of attention (Schmidt, Belopolsky, & Theeuwes, 2015a). During a fear-conditioning phase, one of two visual stimuli, a blue and an orange diamond, was paired with an electrical shock, while the other was not. Following conditioning, participants had to perform a visual search task, which is again a variant of the additional singleton task of Theeuwes (1991). In one quarter of the trials, the CS+ was present as distractor, in another quarter the CS-, and in half of the trials none of the two. The presence of any of the two CS increased RT and error rate; at the same time, CS+ slowed down responses significantly more than CS-, constituting a stronger source of capture. We also observe this automatic orienting when analyzing eye movements, similarly to what happens with reward (Hickey & van Zoest, 2012). In a follow up of the previously cited experiment, Schmidt and colleagues paired a stimulus with a mild shock, and compared how this and a neutral stimulus differently affected saccadic activity (Schmidt, Belopolsky, & Theeuwes, 2015b). With respect to the neutral object, the threat-associated stimulus elicited faster voluntary saccades, and led more frequently to erroneous responses when it had to be ignored.

Aversive stimuli do not always lead to fear though, as they can also generate a blunter feeling of discomfort or disgust, when not considered dangerous. Another important study considered how pairing this time an aversive odor with sound affected auditory perception (Resnik, Sobel, & Paz, 2011). Conditioning procedure consisted in one group in pairing an olfactory aversive unconditioned stimulus (UCS) with one of two pure tones (1 and 2 kHz). In the other group the same procedure was followed using a pleasant UCS. Before and after conditioning, auditory discrimination thresholds for both tones were tested in both groups, and only performance for the CS+ in the aversive odor group resulted deteriorated with respect to baseline, while in all the other conditions discrimination thresholds ameliorated. Results from this study are exactly the opposite with respect to the previous one, suggesting that aversive stimuli can have differential effects on perceptual discrimination depending on a series of parameters.

The whole is greater than the sum of its parts

As far as now we have seen that reward generally acts in the direction of increasing saliency, while on the contrary punishment can lead either to an increase or to a decrease of attentional deployment. I will now review a series of studies which take into consideration not only rewarding or punishing outcomes in isolation, but the two of them and also a neutral type of outcome at the same time. This simultaneous presence allows for an additional piece of analysis. In fact, reward and punishment fulfill two similar but at the same time also opposite functions. Their role in behavior is to highlight particularly important stimuli or actions, so that, once they have been clearly detected, they can either be pursued, as in the case of reward, or avoided, as in the case of punishment. It is then evident how, although sharing one component in their nature, motivational salience (or arousal), they are nevertheless representing opposite kinds of motivational valence, and antipodal types of reaction to the stimuli. The possibility of considering how both types of outcome plus also a neutral one affect attentional control, is then prone to give clearer hints about the true origin and significance of this type of bias. If both reward and punishment act toward an increase in orienting with respect to the neutral condition, then this bias is likely to reflect motivational **salience**, or arousal. If we consider appetitive and aversive stimuli as behaviorally relevant, then their prioritization with respect to neutral stimuli could also potentially be the result of strategic planning. If the effects of reward and punishment lie instead along opposite poles of the same axis, then this bias is more likely to reflect not the strategic, or generally arousing, value of the object, but rather its motivational **valence**. I would like to stress here than every time I am going to talk about salience-like models, I am going to refer to motivational salience, i.e. the model which would attribute perceptual saliency to a stimulus according to its motivational value, and not to perceptual salience by itself.

One of the pioneering studies in the field of the influence of reward in visual attention is the one by Small and colleagues (Small et al., 2005). They used a Posner-type attentional orienting task, where subjects received a spatial cue at the beginning of each trial, which could be either valid (80%) or invalid (20%). After that, they had to respond to appearance of targets (90%), or withhold responses during appearance of foils (10%). Blocks could be either win-type (money gain for correct detection), lose-type (avoiding loss for correct detection), or neutral (neither win nor lose money). Behavioral results showed an effect in response times, such that in both win and lose conditions RT were faster than in neutral blocks, displaying an effect of motivational salience on performance. Similar results were obtained by Engelmann and colleagues, who also designed a Posner-type task, where participants had to detect a degraded face stimulus present either on the left or on the right of fixation (Engelmann et al., 2009). In each trial, immediately before the target, an endogenous cue was presented, predicting target location on 70% of the trials. Motivation was manipulated in a blocked fashion by varying the magnitude and the valence of a monetary incentive linked to task performance. Here results show an increase in perceptual sensitivity for both types of incentive blocks, regardless of valence.

But again, a dual pattern arises for the attentional bias caused by aversive stimuli. Raymond and O'Brien investigated how imbuing stimuli with value would affect the recognition of those same stimuli in a subsequent rapid serial visual presentation task, with and without constraints on available attention (Raymond & O'Brien, 2009). More specifically, two faces were associated with win, two faces with loss, and two with a neutral outcome, and afterwards those same faces were presented again among others, in a task where subjects had to discriminate first the texture of an abstract object (T1) and immediately after the familiarity of a face (T2). Lag between T1 and T2 could be either long (800 ms), leaving attentional resources intact for both targets, or short (200 ms), creating a reduced-attention condition, known to determine an attentional blink for T2 (Raymond et al., 1992). They found that, in the full-attention condition, recognition was enhanced for both motivationally salient stimuli. On the other hand, when attentional resources were reduced, recognition was dramatically impaired for the loss- or no outcome-associated condition, but this was not the case for win-associated stimuli which showed no attentional blink.

Similarly, in an experiment by Laufer and Paz, subjects underwent a procedure which is a mixture between a classical and instrumental conditioning, where three tones were each separately paired with monetary gain, monetary loss or none of the two (Laufer & Paz, 2012). A tone discrimination task was performed both before and after the conditioning procedure, in order to investigate how motivational conditioning of the tones affected sensory discrimination. Data showed a decrease of performance for the loss-related tone with respect with the other two.

Finally, a recent study by Bucker and Theeuwes also supports a valence-like pattern for attentional effects of different motivational outcomes (Bucker & Theeuwes, 2016). In this spatial cuing task, two lines, one diagonal and another one which could be either vertical or horizontal, appeared in each trial on the two sides of fixation. Participants had to respond to the latter by discriminating the orientation of the line. A cue which consisted in a colored frame appeared before the target on one side of fixation (50% validity), with either a short (20 ms) or a long (810 ms) cue-target interval. Each one of the three cue colors was paired with a specific outcome at the end of the trial, such that one color was paired with monetary gain, one with loss and one was neutral. At the same time, participants knew that they had to respond accurately and fast in order to receive extra bonus, therefore the cue-outcome pairing elicited classical conditioning. Interaction between cue validity and motivational type were analyzed, at both short and long cue-target interval. At short intervals, all three types of cue facilitated responding at validly cued locations, probably due to a low-level bottom-up bias (Yantis & Jonides, 1984). At long intervals though, this facilitation remained for the gain-associated cue, disappeared for the neutral cue and reversed for the punishment-associated cue. This suggests that cues associated with an appetitive outcome can strengthen attentional capture and orienting processes, whereas cues associated with an aversive outcome reduce attentional capture and promote reorienting.

To sum up, the attentional bias for motivational stimuli appears more complex when considering rewarding and punishing stimuli at the same time. Some studies highlight a salience-like effect, both in terms of enhancement of target processing (Small et al., 2005; Engelmann et al., 2007; Engelmann et al., 2009), better stimulus recognition (full attention condition in Raymond & O'Brien, 2009), and in terms of increase of attentional capture (Schmidt et al., 2015a). This is a reasonable prediction, considering that both types of outcomes spotlight relevant events for the fitness of the individual or of the species. At the same time, other studies show an apparently counterintuitive pattern, such that stimuli paired with punishment are either not prioritized like rewarding ones (limited attention condition in Raymond & O'Brien, 2009), or they are discriminated even more poorly (Resnik et al., 2011; Laufer & Paz, 2012), and elicit even less orienting (Bucker & Theeuwes, 2016) than neutral stimuli. What causes the former or latter pattern to become the leading one in shaping the deployment of attention? And which could be the evolutionary reason for an attentional bias which only prioritizes rewarding stimuli, but does not underscore, or which even suppresses, the ones linked to potentially dangerous or noxious outcomes?

Getting some air

The paradigm that I have used throughout my project consists in a visual search task for high-level categories of objects performed in pictures of real-world environments. This is in contrast to the studies I have reviewed until now, which employed abstract and simplified displays. Visual search in real-world scenes has shown to present its own peculiarities, and therefore, in the final part of this section, I will try to briefly introduce this topic. Behavioral performance can vary dramatically depending on the type of visual search, and this variability can be addressed through search efficiency, which is measured as the relationship between RT and number of distracters in the scene. Feature-integration theory (FIT) analyses how the number of features defining the target modifies this parameter (Treisman & Gelade, 1980). A search based on a single feature (like for example a search for a color singleton), is efficient, with the addition of multiple distracters has a marginal effect in the increase of RT for target detection. In conjunction search, where multiple features have to be considered in combination, a focused deployment of attention to objects in the scene is required. This serial type of search is less efficient, and RTs increase almost linearly with the number of items in the scene.

What’s special in naturalistic search?

Visual search in real-world scenarios has to deal with a long series of issues: scenes are cluttered, with no clear pattern of organization, and the objects inside them can appear under tremendous changes in location, distance, luminance or orientation. Any high-level category presents an almost infinite number of possible exemplars, often never seen before. All these characteristics would orient to consider this type of search as a serial, inefficient one. Surprisingly, visual search in this type of context has proven to be extremely fast and efficient. In a study by Thorpe, subjects performed a go/no-go task where they had to decide whether an image presented for only 20 ms contained an animal or not (Thorpe, Fize, & Marlot, 1996). Instead of analyzing behavioral measures, which would contain not only the time needed for perceptual processing but also the interval involved in response execution, ERPs were used. They found a frontal negativity centered on frontal recording sites, which was specific to no-go trials and developed around 150 ms after stimulus presentation. This short amount of time was therefore sufficient to perform the categorization of the scene. In another experiment by Wolfe, participants searched for cued objects, in scenes characterized by different set sizes, i.e. the number of objects present in the scene (Wolfe, Alvarez, Rosenholtz, Kuzmova, & Sherman, 2011). They found that each additional object in the scene only added 5 ms to the RT needed for the search, while the correspondent increase was 40 ms per object in an artificial array.

What explains this remarkable skill? Search in naturalistic arrays typically involves scenarios with which we have developed extremely good familiarity. This is perhaps due to our continuous experience with the world, which has been forging our visual system through perceptual learning. We do not only encode the regularity of the objects that are present in the space, but also the regularity of their arrangement within the scenes. According to the phenomenon of ‘contextual cueing’, scene context creates a series of expectations about the environment which helps us during its visual exploration (Chun, 2000). These expectations regard the set of objects that we are likely to find within a specific environment, their location in the scene, and the mutual relationship of the objects between each other. The role of experience in this process is well illustrated in a study, where the visual skill of radiologists and cytologists in interpreting medical images was evaluated (Evans, Georgian-Smith, Tambouret, Birdwell, & Wolfe 2013). Surprisingly, image presentations as brief as 250 ms were sufficient to garner above-chance performance for detecting subtle abnormalities. In synthesis, this means that visual search for high-level category of objects in real-world scenes can be classified as an efficient, parallel type of search.

As specified in the beginning of the section, search for a target requires the creation of an attentional template, an internal representation of the stimulus which will have to match the sensory input in order to distinguish target from non-target. This is also the case for search in naturalistic scenes, which represent a situation where the specific features of the target are often not known in advance, given the extremely high variability in the appearance conditions of the stimulus. This gives rise to the issue of determining what kind of features could define this template. Behavioral and computational studies converge in asserting that the attentional template involved in naturalistic search most probably comprises a collection of intermediate-level category-diagnostic features (Delorme, Richard, & Fabre-Thorpe, 2010; Reeder & Peelen, 2013; Ullman, Vidal-Naquet, & Sali, 2002).

Impact of reward in real-world scenes

A study by Hickey and colleagues specifically investigates the overlap between attentional bias due to reward and visual search in naturalistic scenes (Hickey, Kaiser, & Peelen, 2015). In this experiment, participants had to detect one of three real-world categories of objects (people, cars or trees) which were cued at the beginning of every trial. A correct detection could or could not garner a rewarding feedback (+10 or +0 points, respectively), on a random schedule. The aim of the study was to look at whether a rewarding outcome would have primed the saliency of a high-level category. In order to do so, effects on priming and attentional capture were analyzed. No priming effect after target repetition was found, possibly because the 100% validity of the cue at the beginning of every trial gave the subjects the possibility to fully establish a top-down attentional set for the target, and a ceiling effect left no room for a benefit in RT or accuracy. On the contrary, the presence in the scene of a category paired with reward in the preceding trial, led both to an increase in RT and a decrease in accuracy (experiment 1) or to an increase in RT (experiment 2 and 3) with respect to the condition where the scene still contained the previous correctly detected target, which however had not garnered a rewarding outcome as a feedback. In order to understand which features of the category were effectively primed by reward, in experiment 3

Hickey dissociated two physical components of the “people” visual category. This was done by occluding the images representing people through various distractors, and leaving scenes where only head and shoulders of a person were shown, and other scenes which only contained legs and torso. They found that pairing a legs-and-torso type of scene with reward, not only primed the saliency of the corresponding midlevel feature, but also the high-level category of bodies. If the following scene contained a head-and-shoulder stimulus operating as a distractor, an increase in attentional capture was observed, in terms of an increase of RT needed for correct target detection. These results not only show that the effects of reward on attention extend to paradigms involving naturalistic scenes, but also that, during this type of task, the attentional template upon which the motivational bias is able to operate corresponds not only to the diagnostic intermediate-level features defining a high-level category on an isolated basis, but also to the category itself.

In chapter 2 of this thesis, I will present a behavioral study where participants had to search for a series of high-level categories of stimuli in briefly presented real-world scenes. These categories (people, cars, trees and houses) were imbued with different monetary payoffs: gain, loss or neutral outcome. Participants were informed at the beginning of each block about the type of feedback they would have received in case of correct performance. We subsequently looked at how these different motivational outcomes affected behavioral measures, both in terms of target selection and attentional capture. We approached our data using a valence vs salience hypothesis, bearing in mind the two opposing models that we highlighted at the beginning of the chapter, and which subtend different potential origins of motivational modulation of perception and attention.

Neural signatures of reward and punishment in perception

In the previous section, I have tried to delineate how reward and punishment affect behavioral measures of visual attention. In this section, I will review some studies which look at how this is implemented on a neural level.

Hints from the spikes

A series of electrophysiological studies have investigated how the expectation of reward modulates neuronal activity in the brain of the monkey. Roesch and Olson, for example, recorded activity from a number of frontal areas during a visual task, where the animal had to maintain in memory the location of a directional cue before making a saccade to it after a specific amount of time (Roesch & Olson, 2003). Each trial began with a central cue signaling whether correct performance would have lead to a big or small reward. As already mentioned in previous section, both RT and accuracy measures showed an improvement in the high reward condition. Moreover, many frontal regions showed a correspondent increase in firing: dorsolateral prefrontal cortex (**dlPFC**), frontal eye field (**FEF**), a transitional zone caudal to the frontal eye field (FEF/PM), premotor cortex (**PM**), and the rostral part of the supplementary motor area (**SMAr**). In a study by Platt and Glimcher, neuronal activity of lateral intra-parietal (**LIP**) area was investigated (Platt & Glimcher, 1999). Here, monkeys had to make a saccade either to an eccentric LED which elicited a movement for which the LIP neuron was maximally active (in the response field (RF)), or to another eccentric LED for which the neuron was minimally active (out of the RF). Each trial started with the fixation of a central LED, whose change in color, after a delay, instructed the monkey to make a saccade toward one of the two eccentric locations, in order to get a reward whose amount (in one case) or probability (in another case) varied from high to low across blocks. In both instances, the high (amount/probability) reward condition led to an increase in the firing rate of the LIP neuron, during the interval before the direction of the eye movement to be performed was made evident. Similar results about an increase of activation of LIP neurons during reward expectation were obtained in another relevant experiment by Sugrue (Sugrue et al., 2004).

Activity of neurons in superior colliculus (**SC**) of the monkey was recorded in a study by Ikeda & Hikosaka (2003). Here a cue, which could have been either in or out of the RF of the cell, instructed the movement to be made after a delay. In each block in an alternating fashion, only one movement direction (toward or away from RF) could garner a reward, with reward delivery orthogonal to the direction of the appropriate movement, leading therefore to four different conditions. Like before, the upcoming reward determined an increase in the firing rate of the neuron, but only when the movement had to be performed within the RF of the neuron. Interestingly, this happened either in a reactive (after the receipt of outcome) or in a proactive (before reward delivery) fashion, in gain-type and bias-type neurons respectively. The gain modulation was thought to be primarily induced by cortical inputs (FEF, SEF and LIP), and the bias modulation by the input from basal ganglia, and more specifically from the caudate nucleus (CD).

Accordingly, a couple of interesting studies address the role of the **CD** in visual attention, and more precisely of its tail region. This part of the CD is mostly unique to primates (Hjornevik et al., 2007), and receives inputs mainly from inferior temporal (IT) cortex. This region is hypothesized to be specifically relevant for object-skill, i.e. the ability to perform fast and accurate detection of a learned visual stimulus. This type of skill represents the perceptual counterpart of action-skill, i.e. the ability to fast and accurately perform a stereotyped motor response (Hikosaka, Yamamoto, Yasuda, & Kim, 2013). In a study by Yamamoto and colleagues, neurons of the caudate tail (CDt) were recorded while monkeys performed an object-directed saccade task, a passive-viewing task and a free-viewing task, always looking at the same abstract fractal figures (Yamamoto, Monosov, Yasuda, & Hikosaka, 2012). First of all, neurons of CDt encoded identity and location of the visual objects. Secondly, when stimulated, they elicited movements in the direction of the preferred location. Third, they increased their firing just before saccades were made to the same direction. This study therefore supports the hypothesis that object-skill is mediated by a neural circuit involving a specific portion of the basal ganglia. A subsequent work by the same author investigated the effects of adding reward to a subset of the stimuli (Yamamoto, Kim, & Hikosaka, 2013). Reward association procedure was either flexible, with reward contingencies changing across blocks, or stable, with fixed associations learned throughout several days. CDt neurons showed an enhanced response for stimuli previously associated with reward (even at extinction), but only in the stable association procedure. The association between stimuli and reward was learnt also in the alternative schedule, as shown by a strong gaze bias toward reward-associated stimuli in the flexible procedure but this was not reflected in an increase of neuronal firing in CDt. This suggests a specific involvement of the region in the retrieval of long-term visual associations.

Reward-related neural activation

We have briefly considered a series of studies which address how the prospect of reward is able to modulate spiking of neurons in different regions of the monkey brain. Neuronal recordings provide us with exceptional spatial and temporal resolution, and the animal experimental setting allows task manipulations which are not conceivable when working with human participants. At the same time, it is difficult to abstract these findings to the realm of visual attention, as the relevant role of oculomotor selection in all of these tasks is prone to generate possible confounds. In order to examine how reward expectation modulates activation of different areas of the human brain, we can reconsider the aforementioned study by Engelmann, where participants had to detect a degraded face stimulus presented either on the left or on the right of fixation (Engelmann et al., 2009). Target presentation occurred after appearance of a directional cue which could be valid (70%) or invalid, and motivation was manipulated block wise by varying the magnitude of cash incentive linked to performance. Beyond an increase in dprime for the incentive blocks, they also found a correspondent increase in activation in bilateral **IPS**, **FEF**, **anterior insula** and right middle frontal gyrus (**rMFG**) throughout the block. There was also an increase in activation in incentive blocks during the cue period in all the aforementioned regions and in right substantia nigra (**SN**)/midbrain, anterior cingulate cortex (**ACC**), pre-supplementary and supplementary motor area (**pre-SMA/SMA**), posterior cingulate cortex (**PCC**), **CD**, **putamen**, and thalamus.

Reward or attention?

All the previously reviewed studies represent proactive paradigms, where the response of the animals or of the participants to cues predicting the delivery of reward was analyzed. A relevant commentary by Maunsell legitimately questions whether all these experiments actually took into consideration reward or more simply attention (Maunsell, 2004). Given that subjects knew in advance the type of outcome they would have received if performing correctly, and that reward and attention were not separately manipulated, results would not allow disentanglement between these two components. A series of countermeasures have been taken in order to overcome this issue, which in part trace what we have already seen at the behavioral level, and in part differ and specifically depend on the type of technique in consideration.

Pessiglione and colleagues developed the paradigm which was subsequently adopted by Bijleveld in the abovementioned study (Bijleveld et al., 2010). In this case, instead of performing a mathematical problem, subjects needed to exert force on a hand grip, whose amount was made visible to them by the liquid line on a thermometer shown on a screen (Pessiglione et al., 2007). The higher the fluid level rose, the more of the monetary stake they would get, but the total at stake varied in each trial from one penny to one pound. Critically, the total amount they could gain was communicated at the beginning of every trial either at a supraliminal or at a subliminal level (display duration of 100 ms in the first case, and 17 or 50 ms in the second), and differences in behavior and brain activation were analyzed for the different conditions. On a behavioral level, significant effects were found at 100, 50, and even at 17 ms between the different stakes, showing that even a subliminal presentation was sufficient to energize performance. In the conscious condition, difference between the high and low reward condition garnered differential level of activation in a region which extended bilaterally in the basal forebrain, bordering several structures encompassing the ventral striatum (VS), ventral pallidum (VP), extended amygdala, and basal nucleus of Meynert. They next looked at the engagement of this circuit by subliminal incentives, and found that, over the pallidal voxels that showed significant activation in the conscious condition, the contrast between monetary stakes was significant for 50 ms but not for 17 ms. These results suggest the involvement of this basal forebrain region both in conscious and unconscious conditions.

Stimulation by reward (apparently also at a subliminal level) led therefore to an increase in activation in the above mentioned subcortical structures on a neural level, and to a decrease in RT on a behavioral level. Interestingly, basal ganglia have been highlighted as a brain area critical to the assessment of the speed-accuracy tradeoff also in other contexts. Forstmann and colleagues designed a paradigm in order to localize regions involved in facilitating decision-making under time pressure. Participants were asked to perform the same type of task (a standard moving dot task) under three different conditions: speeded, accurate or normal (Forstmann et al., 2008). They found that anterior striatum and pre-SMA selectively activated in the speeded condition, and that, across subjects, activation in these areas was inversely correlated with the individual “response caution” as estimated by a mathematical model for cognitive decision-making. It seems therefore that this relationship between increased activation in striatum and faster response is not exclusive to reward, but rather the expression of a more general mechanism of behavior regulation. Reward could then be just one of the variables triggering this system.

Another strategy to disentangle reward and attention is represented by the independent manipulation of these two variables. This was the rationale behind a study by Krebs and colleagues, where a cue at the beginning of each trial predicted on one hand the location and the difficulty level (easy vs hard) of the upcoming visual discrimination task, and on the other, orthogonally to that, the potential to win money in that trial (money vs no money) (Krebs, Boehler, Roberts, Song, & Woldorff, 2012). Actual task execution was temporally separated from this phase, so that it was possible to distinguish neural response related to anticipation of these two variables from the one related to reward receipt. Expectation of reward selectively increased activation in Nacc, PCC and calcarine sulcus (V1), while difficulty level selectively engaged dlPFC and pre-SMA. Both factors increased activity in inferior frontal gyrus (IFG), anterior insula, IPS, dorsal striatum and SC. Finally, in addition to these main effects in partly overlapping and partly selective regions, an interaction between the two effects was found in the SN/ventral tegmental (VTA) complex, the posterior thalamus/pulvinar, the right caudate body, and the right anterior midcingulate cortex (aMCC). Regarding the interaction effect, the analysis of the blood-oxygen-level dependent (BOLD) signal extracted from ROIs centered at the local activity maxima, revealed that this effect was driven in all of the regions by a selective activity increase for cues predicting both reward and high difficulty.

Anderson and colleagues performed an imaging study to look at how previous reward affects representation of visual stimuli (Anderson, Laurent, & Yantis, 2014). Here, just like in previous paradigm, they tested the effects of previously learnt reward associations during extinction, after a procedure where one of two colors was associated with a high probability, and another one with a low probability of reward. During the test phase, participants searched for a shape-defined target among stimuli of different colors, with the previously reward-associated color occasionally defining one of the distractor elements of the array. The imaging data from the training phase served as a basis to independently define three ROIs, extrastriate visual cortex, frontal-parietal regions and striatum, whose activity was then evaluated during the test phase. The analysis of the activity of these three ROIs during extinction was aimed at explaining the origin of value-driven attentional signals. If value-driven attention simply reflects increased sensitivity to reward-associated features, we would expect enhanced activation in early visual areas in the absence of any additional control signal. If this phenomenon is the consequence of perseveration in pursuing preceding goals that have been reinforced by rewarding feedback, increased activity in the frontal-parietal regions involved in top-down selection is foreseen. Finally, a third possibility conjectures value-driven attentional priority signals arising from the basal ganglia, which are implied both in reward processing and habitual responding. During presentation of a reward-associated distractor, BOLD signal displayed an increased level of activation within extrastriate regions (bilaterally and beyond whole-brain correction), in IPS (only when distractor was presented in the right hemifield) and CDt (bilaterally, but without passing whole-brain correction when distractor was presented on right hemifield). No difference was present between activation induced by high- and low-reward associated stimuli. At the same time, no increased activation in the test phase was found when a similar task was performed, if no reward was delivered during the training phase. Increased activation of IPS for reward-associated distractor is interpreted by the author as the result of competition among stimuli within the priority map (somehow in contrast with the initial motivations of the study for investigating fronto-parietal activity), while the recruitment of CDt supports a role played by this region of basal ganglia in the establishment of value-driven attention.

Overall, a selective activation by reward is mostly observed at the level of striatum (either VS (Pessiglione et al., 2007; Krebs et al., 2012) or CDt (Anderson et al., 2014), PCC and primary visual regions. Increased activation in other regions of the brain, such as dlPFC and IPS, seems to be more closely linked to the increase of attention. This dissociation on the neural level reminds of another important dichotomy: the one highlighted by Daw and colleagues with respect to behavioral control (Daw, Niv, & Dayan, 2005). They consider the opponency between the system supporting habitual or reflexive control (dopamine afferents and basal ganglia), and the one associated with more reflective and cognitive action planning (prefrontal cortex). In their work, they question in the first place why the brain should rely on different controllers, and secondly how the contrast between these two systems could be regulated in case of conflict. They propose that different classes of reinforcement learning methods are able to depict advantages and limitations of each system. Classical and instrumental learning is best modeled by temporal-difference learning. This is a model-free reinforcement learning method, which is based on ‘caching’, the association of an action or situation with a scalar summary of its long-run future value. Working with cached values is computationally simple but bears the disadvantage of rigidity: the values do not immediately change with the re-valuation of the outcome. By contrast, the function of prefrontal circuit is better represented by model-based reinforcement learning. This method involves exploring a branching set of possible future situations based on different actions, and constructing predictions of long-run outcomes on the fly. These computations can be expensive in terms of memory and time, but at the same time enable to react more flexibly to changing circumstances, as when outcomes are re-valued. Depending on the need for flexibility or speed, the brain will therefore have to determine the controller whose predictions are going to be most suitable for the situation. The authors develop a formal, computational model which suggests how this trade-off could be regulated.

In light of these ideas, the abovementioned contraposition between “reward” and “attention” could then be put into another perspective. Rather than constituting independent concepts, these two processes could instead represent situations where the same target is achieved by the means of different cognitive tools. The behavioral phenomena and the neural activation linked to the construct of “reward” could be attributed to a rudimentary, fast and automatic instrument involved in the detection and pursuit of rewarding stimuli, while “attention” would represent the product of a more elaborate, time-consuming and conscious system deployed to the same aim. We can deepen this idea by comparing the constructs of valence and salience.

Valence vs Salience

As already considered, another important strategy which has been developed in order to disentangle reward from attention is the one of using opposite types of motivational outcomes, reward and punishment, within the same paradigm.

Valence and salience across the cortex

For instance, in a follow-up of their previous electrophysiological study regarding the effects of reward, Roesch and Olson investigated the conjoint effect of reward and punishment in frontal cortex. Spiking activity of neurons of two regions, namely orbitofrontal (OFC) and premotor (PMC) cortex, were recorded in two monkeys during cue display and outcome delivery of rewards or punishments, which could have also been of either small or large entity (Roesch & Olson, 2004). A dissociation was found within the frontal lobe, such that in PMC, on the one hand, both stimuli with high salience determined an increase of the activity of neurons which encoded the movement that the monkey had to make in order to get the reward or to avoid the punishment. In OFC, on the other hand, activity of neurons was representing value in an ascending fashion from negative to neutral to reward. A salience-like type of response was also found by Kobayashi and colleagues, in a study were lateral PFC neurons were recorded during a memory guided saccade task for three different outcomes, a liquid reward, an aversive air puff and a neutral sound (Kobayashi, Nomoto, Watanabe, Hikosaka, Schultz, & Sakagami, 2006). Here, some neurons (the highest percentage) were found to be more responsive to rewarding outcomes, some to punishing ones, and some to both kinds of salient outcomes.

The same rationale led to a study where neurons in LIP were recorded in the monkey (Leathers & Olson, 2012). Here, the animal chose in each trial between cues placed in and opposite to the neuronal response field. Eight different cues signaled four possible outcomes, arising from the orthogonal combination of value and magnitude, so big and small reward, and big and small punishment. Two images of this pool were presented at the beginning of each trial, and the monkey had to make a saccade to the preferred one in order to receive the corresponding outcome. Neuronal firing during the cue period was higher for the large vs small condition, irrespective of the value of the outcome, showing a representation of salience. Importantly, monkeys consistently chose the offer with the better value, so that, for example, in the case of choice between two different levels of punishment, preferred value and magnitude of the cues differed.

Analogous results about frontal regions are present from imaging studies in humans. In a study by Litt, human participants were shown pictures of food items that ranged from being highly disliked to highly liked and were asked to make a choice whether or not they would like to eat the item after the experiment (Litt, Plassmann, Shiv, & Rangel, 2011). For each picture, participants entered their response on one of the four choices: “Strong No”, “No”, “Yes” or “Strong Yes”. These four types of responses were used to define value and salience signals. Areas showing an activation dependent on valence were medial OFC (mOFC), rostral ACC (rACC) and dorsal PCC (dPCC). Salience-based activation was instead observed in dorsal ACC (dACC), SMA and insula. Interestingly, VS showed a peculiar type of activation, which depended both on saliency and on valence. In the imaging study by Engelmann described at the beginning of the section, blocks with payoffs of increasing magnitude but also opposite valence were used, leaving therefore open the possibility to highlight valence and salience regions in a similar way (Engelmann et al., 2009). In this case, however, differently from the previous study, all the aforementioned areas (bilateral IPS, FEF, anterior insula and rMFG, ACC, pre-SMA/SMA, PCC, caudate, putamen and thalamus) showed a main effect of salience, with no area showing any effect of valence.

Valence and salience below the cortex

In the study by Litt (2011), we have seen that activation in VS showed a pattern of activation which depended both on valence and on salience. A study by Cooper and Knutson (Cooper & Knutson, 2008) was specifically aimed at investigating the relationship between valence and salience within this region. Here, each trial started with one of six cue types: two levels of certainty (“certain”/“uncertain”) crossed with three levels of reward (“gain”/“neutral”/“loss”). After a variable delay period, a visual target appeared and participants pressed a button while the target was on the screen. Valence signals were found in the NAcc when outcomes were “certain” (i.e., independent of performance) and salience signals when outcomes were “uncertain” (i.e., based on performance). A more recent study tries to disentangle representation of motivation and expected value (EV) in the striatum using a similar approach (Miller, Shankar, Knutson, & McClure, 2014). Subjects were presented with a cue at the beginning of the trial which anticipated probability of receiving reward and, orthogonal to that, the amount of reward eventually received (high or low). Task consisted in pressing a button as soon as a probe appeared on the screen, from 2 to 4.5 sec after the cue. Activation in the caudate and putamen correlated with motivation (salience), while activation in Nacc was related to differences in the subjective value of anticipated reward (valence). For comparable results see also Knutson (Knutson, Adams, Fong, & Hommer, 2001).

A similar dissociation between salience and valence has been observed in another subcortical structure whose role is fundamental for motivation, i.e. VTA/SNc. Since the groundbreaking work by Schultz (Schultz, 2000; Schultz, 2006), the activity of dopaminergic neurons has been deemed central for the comprehension of the mechanisms through which reward is capable of modulating behavior and attention. Matsumoto and Hikosaka (2009) recorded the activity of midbrain dopamine neurons during a Pavlovian conditioning procedure, and they found two different neuronal populations. Neurons of the former group, located more ventrally, in ventromedial SNc and lateral VTA, were excited by rewarding stimuli and inhibited by aversive stimuli, giving rise to a motivational value signal. Neurons of the latter group, located mostly in dorsolateral SNc, were excited by both rewarding and aversive stimuli, giving rise to a motivational salience signal. Bromberg-Martin and colleagues (Bromberg-Martin, Matsumoto, & Hikosaka, 2010) hypothesize that these two types of signal would be delivered to different parts of the brain. The motivational value signal would be retransmitted to areas implied in value learning and choice evaluation. Indeed, ventromedial SNc and VTA project to vmPFC, including OFC (Williams & Goldman-Rakic, 1998), and to VS, in particular to Nacc shell (Haber, Fudge, & McFarland, 2000). The motivational salience signal would be instead retransmitted to regions implicated in orienting, cognitive processing and general motivation. Dopamine neurons from dorsolateral midbrain do project to dorsal and lateral frontal cortex (Williams & Goldman-Rakic, 1998), and also to Nacc core (Day, Roitman, Wightman, & Carelli 2007; Anstrom, Miczek, & Budygin, 2009). This major division could explain the already mentioned dissociation of representation that we find in frontal cortex, as much as the dual representation found within VS. Finally, dorsal striatum receives inputs both from motivational value and from motivational salience neurons, so that its functions could benefit from both types of signal (Bromberg-Martin et al., 2010).

The constructs of valence and salience enable us to better characterize the contraposition of neural structures that we had already found with respect to the “reward vs attention” debate. At the cortical level, on one side, we find areas traditionally involved in the endogenous modulation of attention, such as **FEF**, **dlPFC**, **SMA** and **IPS**, and areas involved in aspecific arousal such as **dACC** and the **insula**. These areas are activated by both rewarding and punishing outcomes, and seem to be the most likely candidates for a network of regions modulated by attention. On the other side, we find cortical areas involved in the establishment of value, such as **OFC**, **rACC** and **PCC**, which are instead modulated according to a valence-like pattern. On the subcortical side, studies suggest that both types of models are represented within basal ganglia, even though the valence model seems to be preferably represented at the level of **VS**/**NAcc**, and the salience model at the level of **dorsal striatum**. Both types of models are also represented in the **SN**/**VTA** complex. This is consistent with the results of a recent meta-analysis by Bartra, which takes into consideration 206 published fMRI studies investigating neural correlates of subjective value (SV) (Bartra, McGuire, & Kable, 2013). In sum, dorsomedial PFC (dmPFC), bilateral anterior insula, bilateral thalamus and dorsal striatum show the conjunction of positive and negative effects of SV, while bilateral VS, vmPFC and PCC come out in the reward vs punishment contrast.

Competitive dynamics between valence and salience

A study by Knutson exemplifies the interplay between motivational valence and salience in an auction task (Knutson, Rick, Wimmer, Prelec, & Loewenstein, 2007). Subjects received, at the beginning of the experiment, an amount of money that they could have kept or spent on a series of articles shown during scanning. During each trial, participants were presented first with one product shown by itself (a), then with its price (b), and immediately after with the decision of whether or not to purchase the article (c). At the end of scanning, they rated products on several dimensions (i.e., desirability, percentage of retail price that they would be willing to pay for the product, and whether or not they already owned the product). Preference for the article during phase (a) and (b) was correlated with NAcc activation; price differential (i.e., the difference between what the subject was willing to pay and the displayed product price) was correlated with activation in medial prefrontal cortex (mPFC) during the price (b) period. Finally, purchasing was correlated with deactivation of the bilateral insula during the choice period (c). This study suggests an immediate affective reaction to potential gains and losses, which serve together as inputs for a subsequent, more thoughtful assessment.

Finally, another instance where regions of the salience and of the valence network work in opposition is the already mentioned study by Laufer and Paz (Laufer & Paz, 2012). Here, three tones were differentially paired either with monetary gain, with monetary loss or with none of the two. Subjects were scanned while performing an auditory discrimination task on the three tones after the conditioning procedure. As already illustrated, they showed a decrease of performance for the loss-related tone with respect with the other two. Interestingly, the higher the activity in insula, dACC and dorsomedial nucleus of thalamus during the conditioning process, the less the impairment of performance of the subject for the loss-related tone. Given that the standard response after experience with a stimulus is an improved performance (as seen with neutral and gain-related tones), an interpretation of this data is that automatic aversive responses to punishing stimuli impaired in some way their sensory processing. Higher activity in this network allowed the subject to inhibit these automatic responses and process also aversive stimuli according to a more neutral and potentially strategic approach.

Tricks of threat

Throughout this thesis, the relationship between the different motivational outcomes has always been conceived either in terms of a salience effect, with a prioritization of relevant outcomes with respect to neutral ones, or in terms of a positively valenced bias, with a stronger representation of stimuli paired with reward with respect to the ones paired with punishment. A third framework should also be taken in consideration, though, which is the one where loss presents an impact which is stronger than the one due to reward, i.e. an inverse valence effect. Evidence for a similar pattern is given by Baumeister and colleagues, who address the direct comparison between these two outcomes across a wide range of contexts (Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001). The review spans studies concerned with reaction to events, effects on relationships, appraisal and regulation of emotions, memory, development and, critically, information processing. Overall, they find a consistent prioritization of punishing stimuli with respect to rewarding ones, synthesizing this evidence in the formula according to which “bad is stronger than good”. There are clear evolutionary reasons for why this should be the case, when considering the extreme consequences of these two opposing outcomes in the natural environment. While disregard for a rewarding stimulus can indeed cause the loss of a potential gain, neglect of important dangers may nevertheless result in more serious, long-lasting effects, such as maiming or death.

Among the studies of this review, the one by Pratto and John (1991) has a similar rationale as the modified Stroop task which we have already considered in a previous section (Krebs et al., 2010). In this experiment, participants named the colors in which words indicating socially desirable and undesirable traits were printed. Across three different paradigms, color-naming latencies were longer for undesirable with respect to desirable traits. Moreover, subjects displayed a higher rate of incidental learning for undesirable traits, showing that this effect was not due to effects of cognitive effort aimed at keeping threatening material from entering consciousness. The increased incidental learning for undesirable traits suggests instead that, at the expense of reduced attention to the task-relevant dimension (color), automatic vigilance to the task-irrelevant dimension (word) was increased in the case of negatively valenced stimuli with respect to positively valenced ones.

An almond-shaped fear

Correlates of this bias have also been highlighted on a neural level. In a study by Vuilleumier and colleagues, subjects performed a matching task for pairs of stimuli presented at specific cued locations (Vuilleumier, Armony, Driver, & Dolan, 2001). In each trial, pictures of two faces and two houses were displayed, and, depending on their position relatively to the spatial cue, one of the two categories alternatively represented the task-relevant or task-irrelevant one. Moreover, faces could be fearful or neutral. Activation of left and right fusiform gyrus (FG) by faces was affected by attentional condition, as a result of top-down modulation. At the same time, right FG also showed a main effect of expression type independently of attentional condition, such that fearful faces produced increased activation irrespectively of whether the faces were attended or not. At the same time, the left amygdala showed a main effect of expression, but no main effect of attentional condition.

The FG is implicated in the interpretation of high-level visual stimuli, with a specific focus on faces in the fusiform face area (FFA) (Kanwisher, McDermott, & Chun, 1997; Elbich & Scherf, 2016). The increase in activation of FG in the presence of fearful expressions, irrespectively of the task-relevance of the type of face expression, has been interpreted as the product of an additional, automatic control, which enhances representation of threatening stimuli irrespectively of deliberate selective attention (Vuilleumier, 2005). The amygdala, which activates in response to (even not consciously perceived) aversive stimuli (Anderson, Christoff, Panitz, De Rosa, & Gabrieli, 2003) and presents a direct and strong connection with inferior temporal cortex (Amaral, Behniea,

& Kelly, 2003), has been posited as the main neural substrate for this independent attentional control. According to this view, then, the exogenous, endogenous, and emotional controls would therefore constitute three different sources of bias, operating independently from each other. Evidence for this hypothesis has been provided by Brosch and colleagues, both on a behavioral and on an electrophysiological level (Brosch, Pourtois, Sander, & Vuilleumier, 2011). In a dot-probe task, valid and invalid exogenous, endogenous and emotional cues were manipulated orthogonally. The analysis of RTs showed that attentional facilitation for the three different types of valid cues added up linearly. In a second experiment, the same paradigm was optimized to analyze the effects of the exogenous and the emotional bias at the level of ERP. The analysis showed that the two attentional effects presented non-overlapping temporal foci, which were related in one case to low-level properties of the stimulus and in the other to its emotional content.

A twisted mind

Apparently, though, the whole picture seems to be more articulated than the one presented up to now. In a dot-probe task, MacLeod and colleagues compared the attentional bias towards aversive stimuli in clinically anxious and control subjects (MacLeod, Mathews, & Tata, 1986). Two words, one of which could have been threat-related, were briefly presented on a screen along a vertical axis. Immediately after, a dot could have (or could have not) appeared in one of the two positions previously occupied by the words. Latencies for probe detection were calculated for trials where a threat-related word has appeared, and evaluated as a function of whether the aversive word had appeared in the same or in the other position with respect to the dot. Clinically anxious participants displayed an attentional bias towards threat-related words. Surprisingly though, control participants shifted attention away from the same type of stimuli, suggesting the presence of an opposite type of bias. Other studies showed similar results, such that in general this attentional bias towards aversive stimuli is stronger in high trait anxiety individuals than in low trait anxiety ones, or present only in the former group, while the latter one rather displays avoidance towards the same type of stimuli (Mogg et al., 1995; Vandenhout, Tenney, Huygens, Merckelbach, & Kindt, 1995).

This dichotomic response according to personality led to the formulation of the interaction hypothesis, according to which high trait anxious individuals show a preattentive bias *towards* aversive stimuli, while low trait anxious individuals a preattentive bias *away* from them (MacLeod & Matthews, 1988). Nevertheless, a similar framework leads to the unintuitive situation where highly aversive, and therefore also potentially highly relevant, stimuli determine a great increase of selective attention in highly anxious individuals, but at the same time a complete neglect of the same type of information in the remaining low trait anxious population.

In order to account for this discrepancy, a cognitive-motivational model was subsequently elaborated by Mogg and Bradley (1988). According to this view, the relationship between the subjective threat value and attentional bias is not linear. When the threat value of the stimuli is low, then selective attention will shift away from them; when it is high, the response will be opposite, and lead to their prioritization (Fig. 1.1). Individual differences would therefore simply lie in the threshold given by each subject, with highly anxious individuals shifting from suppression to prioritization for elements with a lower level of threat.

Interestingly, we also find evidence for an opposite influence of this personality trait on the bias for positively valenced stimuli. In another experiment by Mogg and Bradley (1999), attentional capture for different types of emotional faces (happy, neutral, threatening) was evaluated in a dot-probe task. When considering the effects of type of emotional faces (happy and threatening) and of personality trait, an interaction was found between these two factors, such that while high trait anxiety individuals showed a stronger bias for threatening faces with respect to low trait anxiety ones, an opposite trend was found for happy faces. Moreover, we have previously seen that the activation of the amygdala has been hypothesized as a potential source for this additional aversive bias. In a task where participants viewed masked and non-masked versions of fearful and neutral faces, a correlation was also found between trait anxiety level of participants, and activation of basolateral amygdala by masked expressions of fear (Etkin et al., 2004).

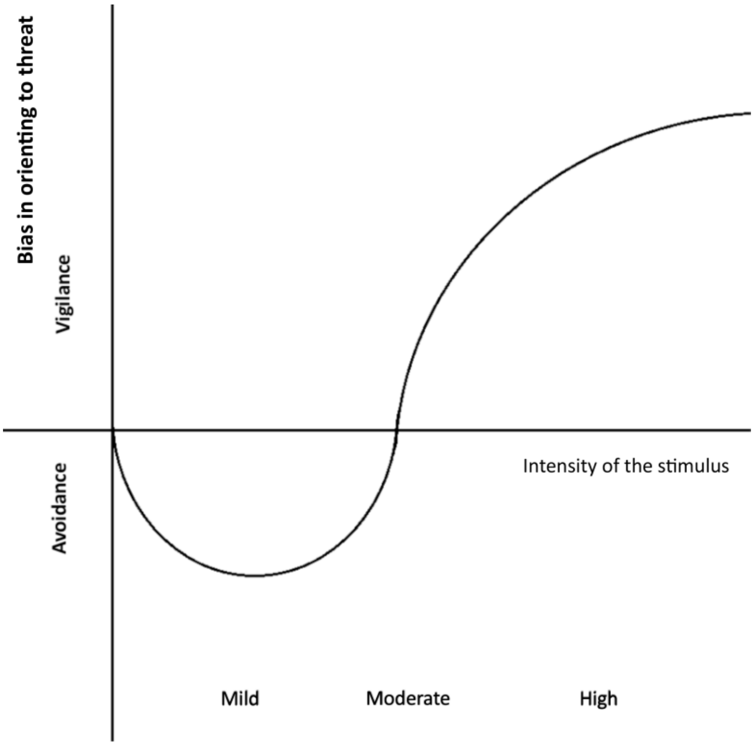

Fig. 1.1 Hypothetical relationship between the attentional bias and the subjective threat value of stimuli (adapted from Mogg and Bradley, 1988).

Do we find studies which address this question not only by relating it to personality traits of different participants, but also to the level of intensity of the stimulus? In a recent study by Simola and colleagues (Simola, Le Fevre, Torniainen, & Baccino, 2015), eye movements and eye-fixation-related potentials (EFRPs) were recorded while subjects observed emotional scenes categorized according to two orthogonal dimensions: valence (unpleasant, pleasant) and arousal (high, low). Fixations rates, gaze durations and two specific electrophysiological indices time-locked to the fixation of emotionally salient items, P2 and LPP, all displayed a similar interaction between valence and arousal. More specifically, all these parameters showed an enhanced response to pleasant stimuli with respect to unpleasant ones in the low arousal condition, and an enhanced response to unpleasant stimuli with respect to pleasant ones in the high arousal condition. Therefore, this variability in response to items of the same valence with different levels of arousal seems also to be present within subjects across stimulus intensity, and not only across subjects.

On a similar note, studies making use of different type of emotional faces (Ohman, Lundqvist, & Esteves, 2001) found that while threatening and fearful faces (negative valence, high arousal) biased attention, sad ones (negative valence, low arousal) did not. Finally, a similar interaction has been recently highlighted also on a decision-making dimension (Sussman, 2017). When people evaluated past experiences, they tended to prefer positive outcomes that had lower magnitudes of contributing gains and losses. On the contrary, when negative outcomes were considered, subjects preferred outcomes with higher magnitudes of contributing gains and losses. This finding suggests that a similar pattern may apply also to other cognitive processes beyond the realm of visual attention, and present a broader adaptational meaning.

We have already talked about the evolutionary reasons why a bad outcome could be “stronger than good”. At the same time, we also find that, depending on personality and also on the intensity of the stimulus, sometimes “good is stronger than bad”. If so, why? Taylor and Brown have addressed this question on a broad level, similarly to the wide-ranging review by Baumeister (Taylor & Brown, 1988). They claim that, contrary to the old concept of the healthy individual as someone who perceives reality accurately (Jahoda, 1958), a large amount of evidence has accrued with time which shows that the healthy individual is subject to a series of positive illusions, such as unrealistically positive vision of the self, illusions of control and unrealistic optimism. Apparently, these cognitive biases have a role in emotional regulation, inasmuch as they help coping with a potentially overwhelming environment. In fact, it has been found that individuals with low self-esteem or moderately depressed are the ones who are less prone to the abovementioned types of biases.

Patterns in the brain

Until here, my analysis has been focusing on how reward and punishment affect univariate activation in different areas of the brain. In this project, though, I am going to investigate how these two outcomes modulate the amount of category information in object-selective cortex (OSC), as assessed through Multivariate Pattern Analysis (MVPA). This type of analysis has been shown to represent a powerful tool in order to investigate how various types of stimuli and functions are represented within the brain. Seminal work by Haxby and colleagues (2001) examined how different categories of real-world objects are encoded in one of the fundamental regions dealing with visual object representation, namely ventral temporal cortex. In this study, participants performed a simple one-back task while looking at different exemplars of the same category in each block. Categories used in the study comprised natural stimuli (faces and cats), man-made objects (houses, chairs, scissors, shoes, and bottles), and nonsense, scrambled images. Ventro-temporal cortex of participants showed to reliably represent different high-level categories of objects.

MVPA of natural scenes

Again, we will now move from tasks performed in abstract, simplified arrays, to a more ecological approach based on visual search in real-world scenarios. A study by Peelen found that, when looking at complex natural scenes during a category detection task, the pattern of fMRI activity elicited in OSC retains information about the category of objects that the subject is actively looking for (Peelen, Fei-Fei, & Kastner, 2009). Data revealed that information about the target was present in OSC even from parts of the visual field which happened to be task-irrelevant and which were presented outside the focus of spatial attention. This suggests the presence of a top-down mechanism capable of biasing category representation for all areas of the visual field, whether they are actively attended or not. In another study by Peelen and Kastner, a cue at the beginning of the trial signaled the type of category to detect, but some cues were not followed by the actual presentation of the scene (Peelen & Kastner, 2011). Patterns of activation of OSC during this preparatory activity, even when considering just cue-only trials, significantly correlated with the localizer pattern of the object that had to be detected, leading to the idea that the activity of visual areas is modulated already before the scene is presented. An interpretation of these data is that the instantiation of the search template could determine the pre-activation of neuronal populations selective for the target category.

Cluttered visual scenes do not only contain the target of our search, but also and generally even more extensively other types of objects, which then act as distractors. Further studies address the issue about how different types of distractors are encoded in visual areas. One important distinction can be made between objects which have already been acting as previous targets, and categories of objects which have never been relevant. A study by Seidl compared correlation of patterns elicited by categories of objects during localizers and during blocks where they were acting as targets, as distractors (objects which had been acting as targets on previous blocks) or as irrelevant objects (categories which had never been attended by the subject during the experiment) (Seidl, Peelen, & Kastner, 2012). Beyond the enhancement of the category information of the target, an additional finding was that, with respect to a never relevant category, the distractor showed a reduced level of information in the scene. This could be interpreted as the active suppression operated by the visual system on the representation of an otherwise too salient category of stimuli. We can therefore appreciate an interaction between quality of visual representation and task-relevance, such that this is enhanced with respect to a neutral baseline if the viewed stimulus matches top-down attentional template, and it is suppressed if it does not.

Patterns of reward

But then, what if different categories of objects are paired with different levels of reward? This was the question underlying a study by Hickey and Peelen, where subjects performed a similar category detection task in pictures of outdoor natural scenes (Hickey & Peelen, 2015). Each subject had one special category of objects among three, the “rewarding” category, which allowed them to earn extra points when correctly detected. OSC showed to retain a higher amount of information for the category of objects associated with reward with respect to the neutral category, when they were both acting as targets. Symmetrically, when they were acting as distractors, there was a stronger suppression of information for the category associated with reward than for the neutral one. Reward acts then by further increasing representation of the visual category in OSC, when the stimulus matches the top-down attentional template. The interaction which we had already seen regarding how task-relevance modulates category information holds true even in this situation, as the highly represented rewarded category concurrently necessitates a higher degree of suppression when acting as distractor. The modulation of category information in OSC was correlated with univariate activity across the brain, and a series of areas implied in reward processing were found to be predictive of the increase in the quality of representation linked to reward, such as left OFC, bilateral dlPFC, ACC, inferior and superior parietal lobules, and IFG. Moreover, this correlation was also found for an area corresponding to dopaminergic midbrain. As this region constituted a special focus of investigation, this was confirmed by more specific and compelling analyses.

In chapter 3 of this thesis, I will present the imaging counterpart of the behavioral study from chapter 2. Subjects had to look for cued category of objects in naturalistic scenes, and each category was bound to a specific type of motivational feedback. A very similar paradigm is presented in chapter 4, apart from the fact that an inconsistent schedule of reinforcement is administered in this case. In both studies, we subsequently analyzed how this association affected neural representation of visual categories through the use of MVPA. Once again we approached our data with a valence vs salience hypothesis. Finally, we also performed a univariate analysis of the BOLD signal from both experiments, in order to see which regions of the brain were found to follow the former or the latter pattern.

Chapter 2

Irrational impact of reward and punishment on visual search

Abstract

Association of a stimulus with reward determines an increase in the accuracy of its perceptual processing, and renders at the same time this stimulus and its associated features more attention-drawing. Focus has been put on the origin of this effect, in order to ascertain whether it is the product of voluntary, strategic attention, or of an automatic, and also potentially counterproductive, bias. Here, we use punishment as a tool to decouple these two interacting influences. Subjects detected exemplars of cued real-world categories (people, cars, houses and trees) in naturalistic scenes. On a motivational level, rewarding and punishing blocks were highly relevant for participants, as in both cases correct detection of an exemplar of the cued category was 50 times more valuable than in neutral blocks. At the same time this constant utility was located in oppositely valenced spaces for these two circumstances, as an increase in the gain for rewarding blocks, and a decrease in loss for punishing ones. Different outcomes affected search, both in terms of target selection and of attentional capture, according to a valence, and not to a salience, pattern, such that stimuli associated with loss yielded lower performance and caused less capture than the ones associated with gain. We conclude that motivational feedback impacts search through an irrational, non-strategic bias.

Introduction

Traditional theories of visual attention have always stressed the dichotomy between the endogenous, top-down control and the exogenous, bottom-up bias. Nevertheless, a growing number of studies is suggesting that goal and physical salience may not be the sole actors involved in the process, and that an additional control may take part in the course of the deployment of attention (Awh et al., 2012). An important example of this phenomenon is given by reward, whose role has been shown to go beyond a simple modulation of top-down attentional set. In abstract arrays, reward delivery causes priming of target-associated features, even when this happens to be counterproductive (Della Libera & Chelazzi, 2006, Hickey et al., 2010a). Attentional capture has also been highlighted at the level of ocular movements, especially during early stages of the deployment of attention (Hickey & van Zoest, 2012). Moreover, similar instances of reward-associated phenomena have recently been demonstrated for visual search of real-world object categories in natural scenes. Here, categories of objects paired with reward led to a higher degree of attentional capture when acting as task-irrelevant distractors in the following trial (Hickey et al., 2015).

At present, the role of punishment in visual search has been less intensively studied. Yet, especially the knowledge about the mutual relationship between the impact of these two outcomes could represent a source of important insights into the nature of motivational modulation of perception itself. Under an evolutionary point of view, both reward and punishment act as signals for stimuli with strong biological importance. In this frame, then, their perceptual coding could follow a nonspecific affective salience, giving higher value both to rewarding and punishing stimuli with respect to neutral ones. This scheme would also represent a convenient solution to the problem of attributing value to stimuli, and therefore potentially be the result of strategic planning and voluntary attention (Fig. 2.1A). We should underline here that, when using the term salience, we are considering it under of a motivational point of view, and not under a perceptual one. On the other hand, coding of these two outcomes could follow a valence-based ranking, attributing value to stimuli in an ascending fashion from aversive, to neutral, to appetitive ones (Fig. 2.1C). This would reflect their incentive value, namely the motivational quality which makes a stimulus attractive and desired, thereby guiding selection, attention and approach. Accordingly, a punished stimulus would be less wanted than a rewarded and also than a neutral one, and consequently be ignored and avoided, although relevant. Both models have a counterpart when considering the role of distractors associated with the same types of motivational outcomes (Fig. 2.1B and 2.1D).

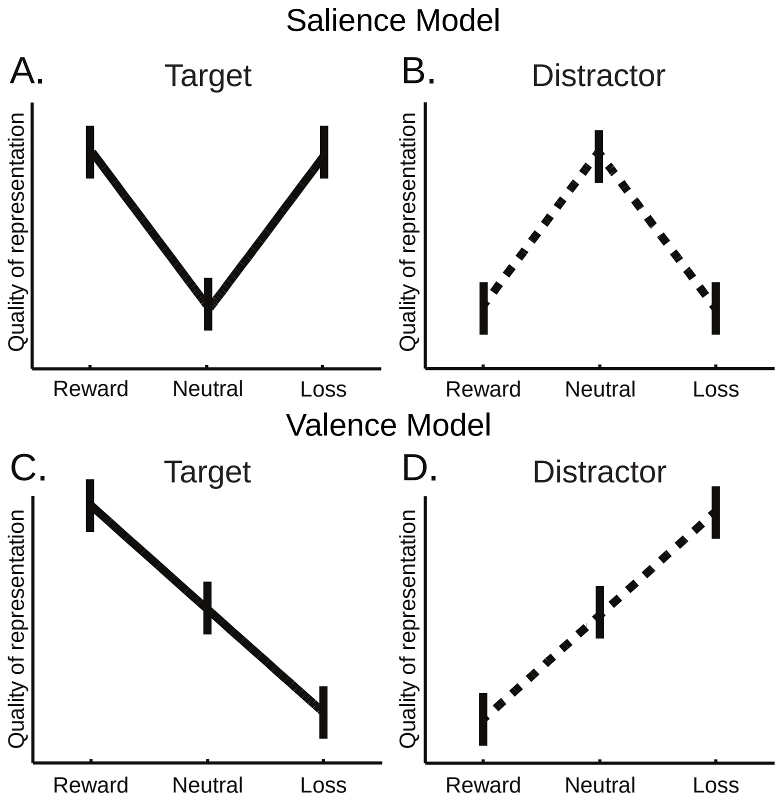

Fig. 2.1 Predictions from the two alternative models. (A) For the salience model, both reward- and loss-associated targets should be more strongly represented than neutral ones, and (B) they should draw more attention when acting as task-irrelevant distractors. (C) According to the valence model, targets associated with reward should be more strongly represented than loss-associated targets. (D) Conversely, distractors associated with reward should require stronger attentional suppression than distractors associated with punishment.

In fact, previous studies in the field of attention gave contradictory results about this issue. In some cases, punishment led, similarly to reward, to a predictable increase in the visual saliency of the stimulus (Small et al., 2005; Engelmann et al., 2009; Raymond & O'Brien, 2009; Schmidt et al., 2015). Other studies found instead a suppression in the representation of stimuli paired with this payoff (Raymond & O'Brien, 2009; Resnik et al., 2011; Laufer & Paz, 2012; Bucker & Theeuwes, 2016).

A particularly intriguing result was given by the study of Raymond and O'Brien (2009). Here, novel face stimuli were imbued with different values of expected value (EV), namely a gain or a loss. Then, recognition of these same faces was measured in a paradigm which made use of the attentional blink (AB) phenomenon (Raymond et al., 1992). When attentional resources were not constrained, then recognition of these faces followed a salience model. Crucially, when attentional resources were not fully available, recognition of these faces also followed a valence model, such that stimuli associated with reward did not show any effect of AB, while stimuli associated with loss did.

Fig. 2.2 Examples of naturalistic scenes employed for the visual search task.

In the current study, we use a category detection task to test this valence vs salience hypothesis. Participants looked for cued target categories (people, cars, houses and trees) in real-world scenes (Fig. 2.2), and received different types of feedback (rewarding, neutral or punishing) across blocks (Fig. 2.3A). In rewarding and punishing blocks, successful detection of a target category had a static relative value of +100 points. Nevertheless, this constant amount was located in a positively valenced space for rewarding blocks (+150 for correct vs +50 for incorrect) and in a negatively valenced one for punishing blocks (-50 vs -150). In neutral blocks, instead, correct target detection had a relatively negligible value of only +2 points (+1 vs -1) (2.3B). We analyzed our behavioral measures to test whether they displayed to follow the valence pattern, the salience one, or a combination of the two.

Materials and methods

Participants.

105 participants (79 female, mean age 23 ± 4.1 SD) took part in the experiment. All procedures were carried out in accordance with the Declaration of Helsinki and were approved by the ethical committee of the University of Trento. 2 participants were excluded from the analyses, as frequency of their responses was 2.5 SD lower than the mean of the population, and another one was excluded as his conditional accuracy measures differed more than 3 SD from the mean of the population. Finally, 1 participant was excluded because of inconsistent responses on a personality questionnaire.

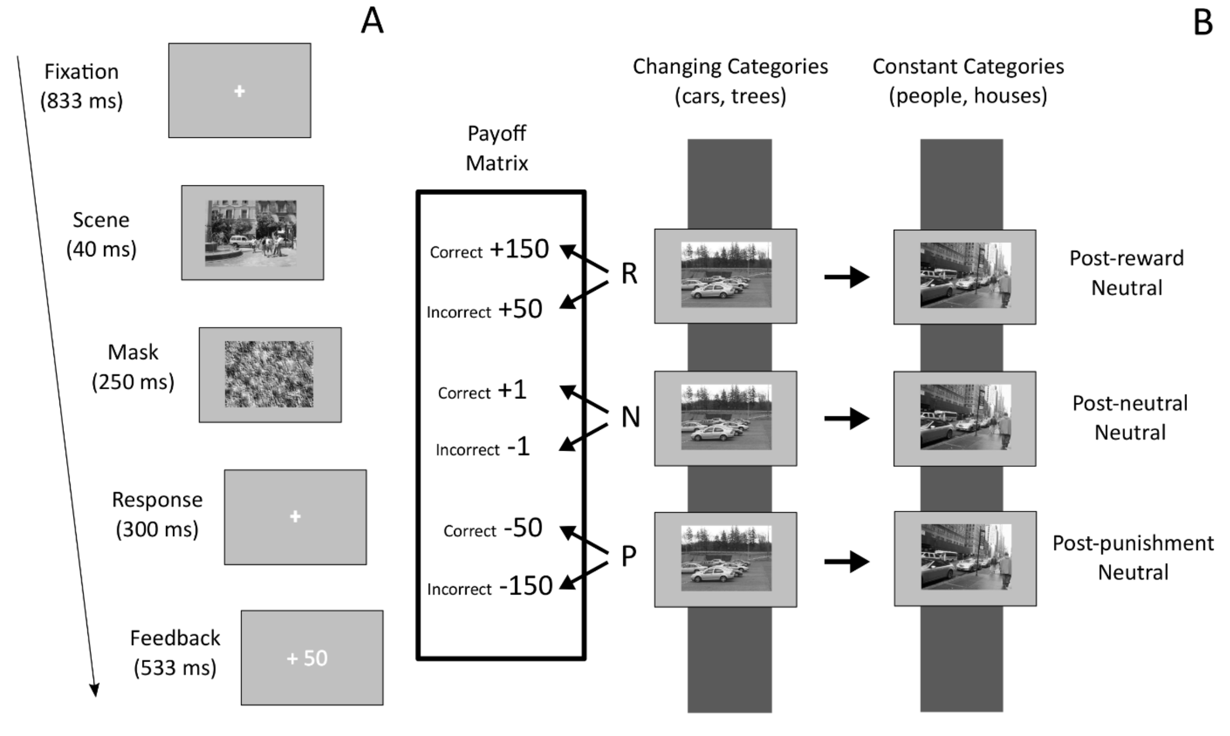

Fig. 2.3 (A) Trial structure. Subjects reported the presence of a cued target category in the scene. (B) Structure of the experiment: changing categories (cars and trees), in which payoff schedules were alternating, were always followed by constant categories (people and houses), which always maintained a neutral feedback.

Stimuli.

The stimulus set for the visual search task consisted of black and white pictures (27° × 38° visual angle) of natural scenes (n=384) selected from an online database (Russell et al., 2008). Four groups of pictures, each made up of 48 scenes, contained all categories but one. Other six groups contained all the various possible combinations of two categories. Four of these groups contained 32 pictures, while the "car-tree" and the "people-house" group contained 24 pictures. Natural scene photographs were followed by perceptual masks of equal size, generated by combining white noise at different spatial frequencies and superimposing naturalistic structure on the noise (Walther, Caddigan, Fei-Fei, Beck, 2009). Stimuli were presented using the Psychtoolbox (Brainard, 1997).

Procedure.

The experiment took place in a sound-attenuated room. Participants were presented with stimuli on a CRT monitor (51° × 65° visual angle) and asked to report the presence or absence of a category of objects via keyboard button presses (respectively, through "b" and "m" key). Each trial started with a fixation cross (833 ms) followed by the brief presentation of the scene (40 ms), a mask (250 ms), the reappearance of fixation (300 ms), and the feedback (533 ms; Fig. 2.3A). Participants reported the presence of the target category with the right index finger response and its absence with the right middle finger response. Responses given outside the 300 ms fixation interval, or absence of response, garnered an incorrect performance.

In each trial, feedback depended on block type (rewarding, punishing or neutral), trial type (target-present or target-absent) and correctness of response. We used two slightly different payoff matrices across subjects, which are both shown in Fig. 2.4. For the first 55 subjects, scores followed the matrix shown in fig. 2.4A. In rewarding blocks, subjects received, in target-present trials, 150 points for correct response, and only 50 points in case of incorrect response. In target-absent trials, they lost 0 points for correct response and 50 points for incorrect response. In punishing blocks, subjects lost, in target-present trials, 50 points for correct response, and a greater amount of 150 points in case of incorrect response. In target-absent trials, they lost 0 points for correct response and 50 points for incorrect response. For neutral targets, points were +1 for correct and -1 for incorrect response, both for target-present and for target-absent trials. Payoff matrix for the last 46 subjects is shown in fig. 2.4B. The only difference here regarded target-present trials in rewarding and punishing blocks. Correct detection of rewarded targets yielded 100 points, and failure to do so 0 points, while correct detection of punished targets determined loss of 0 points, and failure to do so the loss of 100 points. Statistical analysis demonstrated no difference in performance across these groups (for all effects involving this difference, p > 0.211), so results were collapsed and are presented together.

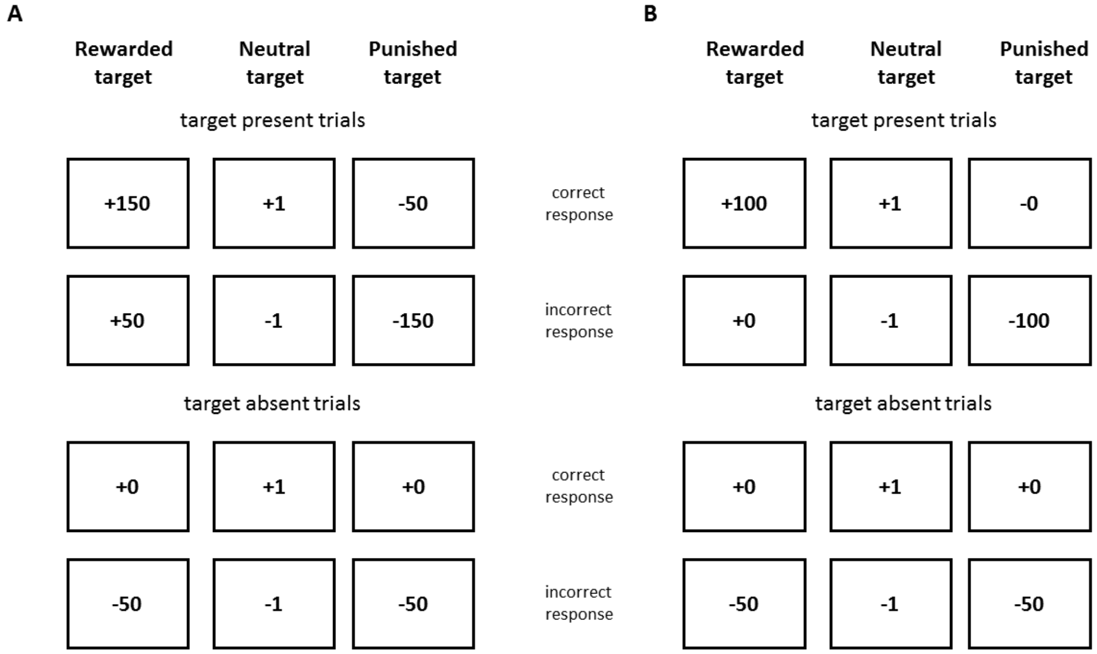

A

	Rewarded target	Neutral target	Punished target
target present trials			
correct response	+150	+1	-50
incorrect response	+50	-1	-150
target absent trials			
correct response	+0	+1	+0
incorrect response	-50	-1	-50

B

	Rewarded target	Neutral target	Punished target
target present trials			
correct response	+100	+1	-0
incorrect response	+0	-1	-100
target absent trials			
correct response	+0	+1	+0
incorrect response	-50	-1	-50

Fig. 2.4 Trial outcome depended on block type, presence of the target and correctness of response. Two payoff matrices were used in the experiment. (A) Payoff matrix used with the first 55 subjects and (B) with the last 46 subjects.

Participants performed 48 blocks of 16 trials each. Before each block started, a display was presented until subject response, indicating the target and the type of feedback for that block and the score which had been obtained up to that moment. For two categories (cars and trees, “changing categories”), the type of feedback was changing among the three motivational conditions, while for the other two (people and houses, “constant categories”) feedback was always neutral. These two types of categories were alternating, so that a changing category always preceded a constant category (Fig 3C). Each block contained 8 target present and 8 target absent trials: in both cases, 4 trials presented three-category scenes, and the remaining 4 trials presented two-category scenes. The total amount of trials was 768 per subject, so that each scene was shown twice. At the end of the task, each subject was administered a BIS/BAS scale questionnaire (Carver & White, 1994). Participants were paid based on the number of points accumulated during the experiment.

Data Analysis.

Two categories of objects – the ‘changing categories’ - (cars and trees), had shifting types of outcome, while the other two – the ‘constant categories’ - (people and houses), always had a neutral feedback. Changing and constant categories were alternating. Accuracy values for detection of changing categories were used as a measure of the direct effect of feedback on target selection (Fig. 2.5A). For constant categories, the crucial element was represented by the target of the previous block, a changing category which could have either been positively conditioned, negatively conditioned or none of the two. We then looked at how much this category disrupted search in the subsequent block, as a function of the type of feedback received. Accuracy measures for distractor-absent and distractor-present trials were calculated, separately for post-rewarding, post-neutral and post-punishment blocks, and the distractor-absent measure was subsequently subtracted from the distractor-present one. This procedure yielded three values, one per motivational condition (Fig. 2.5B). The same type of analysis was performed on RT measures (Fig. 2.6).

Finally, the six behavioral values, three for target selection and three for attentional capture, were modeled using two predictors: a former one accounting for a valence-based model of targets and distractors [+1 0 -1 -1 0 +1], and a latter one accounting for a salience-based model [+1 -2 +1 -1 +2 -1]. Values were standardized by computing the z-score for each vector, so that the first vector became [+1.12 0 -1.12 -1.12 0 +1.12], and the second one [0.6455 -1.2910 0.6455 -0.6455 1.2910 -0.6455]. Notice that, for each vector, distractor values are simply the opposite of the target values. Moreover, the vectors used for the RT analysis had opposite directionality with respect to the ones used for accuracy (therefore [+1 0 -1 -1 0 +1] for valence and so on), because of the inverse relationship expected in this case. This yielded one beta coefficient per subject per model. For each model, all the beta coefficients were then t tested against zero, to see which model could correctly fit the data. Finally, the salience model beta values were subtracted from the valence model beta values individually for each subject, and the resulting differences were again t tested against zero to make a direct comparison between the two models. Planned post-hoc comparisons were one-tailed paired sample t-tests. Effect sizes have been evaluated according to Cohen (1988).

Results

As explained above, the analysis focused on accuracy measures obtained both for changing and for constant categories. Automatic effects of reward on attention have been shown to exhibit primarily through a change in RT, while levels of accuracy would also reflect the change in the speed-accuracy tradeoff as operated by conscious effort (Bijleveld et al., 2010). We wanted to test the hypothesis that the effect on RT could be resistant to top-down strategies. Consequently, we gave participants very tight time windows to respond, so that relevant changes in RT would reflect in parallel differences in accuracy measures, affecting as a consequence monetary payoff. Therefore, if subjects were in control of this tendency they could have tried to adjust their speed-accuracy tradeoff correspondingly, in order to optimize performance.

For changing categories, mean accuracy values for the three motivational conditions were taken into account (Fig. 2.5A). For constant categories, we evaluated the degree of attentional capture exerted by previous target during search for constant categories, as a function of the type of motivational feedback received (Fig. 2.5B). We then regressed the six values coming out of this analysis with a valence and a salience model, to determine which could best fit the data. Behavioral measures followed a valence pattern.

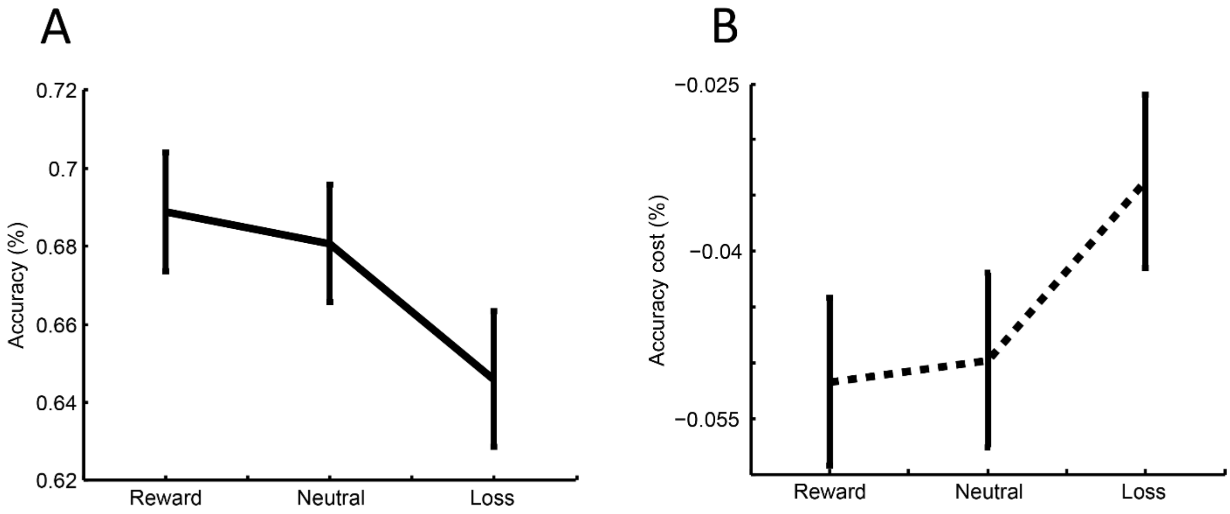

Fig. 2.5 (A) Accuracy measures for the changing categories. In line with the valence model, accuracy is better for reward-associated targets than for loss-associated targets. (B) This is paralleled at the level of attentional capture, so that categories of objects paired with reward determine a larger decrease in accuracy than categories paired with loss. Error bars reflect within-subject standard error.

Coefficients corresponding to the valence model were positive (t(100) = 4.219, $p < 10^{-5}$, Cohen's U3 = 0.322) while coefficients corresponding to the salience model were negative (t(100) = -1.693, p = 0.093, Cohen's U3 = 0.535), reflecting an inverse relationship between the model and experimental results (and thus a very bad fit). Also the direct comparison between the two models yielded a positive result (t(100) = 3.969, $p < 10^{-4}$, Cohen's U3 = 0.356). Follow-up contrasts showed that accuracy of detection of categories associated with reward was higher than accuracy measures for the ones associated with loss (t(100) = 4.820, $p < 10^{-6}$, Cohen's *d* = 0.480), and that search for a neutral constant category was disrupted more by a reward-associated distractor than by a loss-associated distractor (t(100) = 1.712, p = 0.045, Cohen's *d* = -0.170).

We subsequently analyzed RT measures in order to detect any potential effects of speed-accuracy tradeoff. For the target effect, we looked at the mean RT for the three changing categories (Fig. 2.6A). For the distractor effect, we evaluated the RT cost of previous target during search for constant categories, as a function of the type of received feedback (Fig. 2.6B). We subsequently treated the six values in the same way as we did with accuracy measures.

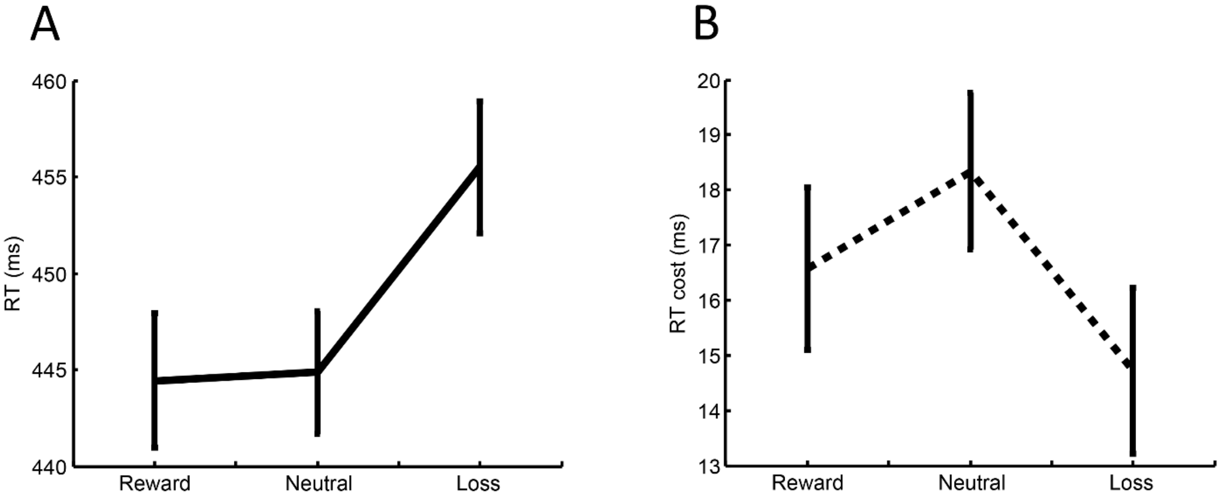

Fig. 2.6 (A) RT measures for the changing categories. In line with the valence model, RT are longer for reward-associated targets than for loss-associated targets. (B) At the distractor level, no difference is present between the two conditions. Error bars reflect within-subject standard error.

Again, we found positive coefficients for the valence model (t(100) = 4.945, $p < 10^{-6}$, Cohen's U3 = 0.277), and negative coefficients for the salience model (t(100) = -3.452, $p < 10^{-4}$, Cohen's U3 = 0.634). Finally, we found positive coefficients for the direct comparison between the two (t(100) = 5.498, $p < 10^{-7}$, Cohen's U3 = 0.247). Follow-up contrasts showed that RT for detection of categories associated with reward was shorter than RT for categories associated with loss (t(100) = 6.792, $p < 10^{-10}$, Cohen's d = -0.592), while the same comparison at the distractor level showed no difference between the two values (t(100) = 0.933 p = 0.177, Cohen's *d* = 0.093). In sum, the impact of the different motivational outcomes on RT were not suggestive of a speed-accuracy tradeoff which could explain the effects on accuracy, but rather (only at the target level) had the same directionality of those effects. Participants had a limited time for response, which means that two different types of errors were possible: explicit errors, where participants incorrectly reported the presence or absence of the target, and time-out errors, where they failed to respond in time. The presence of time-out trials renders this study unsuitable for an analysis in the framework of signal detection theory. At the same time, a more detailed examination can consider target present and target absent trials separately, also in order to understand whether the results that we find derive from a specific condition. We performed this analysis only for changing blocks, as these are the ones where a strategic approach of participants towards different types of outcomes could have been used (although our payoff matrix was specifically designed to equate the two salient conditions, reward and punishment, in terms of their strategic value). Both for target present and target absent trials, we calculated explicit accuracy, namely the amount of correct responses over the total amount or explicit responses (excluding therefore time out trials). Results are shown in Fig. 2.7. Data for explicit accuracy in target present trials followed a valence ($t(100) = 3.176$, $p = 0.002$, Cohen's U3 = 0.361), but also a salience pattern ($t(100) = 2.605$, $p = 0.011$), Cohen's U3 = 0.416), and follow-up contrasts revealed that participants responded significantly better to rewarding targets than to punishing ones ($t(100) = 3.176$, $p = 0.002$, Cohen's d = 0.623), and also better to rewarding targets than to neutral ones ($t(100) = 4.030$, $p < 10^{-5}$, Cohen's d = 0.331) (Fig. 2.7A). On the contrary, explicit accuracy in target absent trials did not follow either a valence ($t(100) = 0.698$, $p = 0.486$, Cohen's U3 = 0.564), nor a salience pattern ($t(100) = 1.367$, $p = 0.175$, Cohen's U3 = 0.386) (Fig. 2.7B).

Also the number of time-outs in target present trials was found to follow a valence ($t(100) = 5.570$, $p < 10^{-7}$, Cohen's U3 = 0.307) but not a salience model ($t(100) = -3.535$, $p < 10^{-4}$; Cohen's U3 = 0.648; again, negative values mean inverse relationship with the model), and the amount of time-outs was significantly higher in the loss condition than in the reward one ($t(100) = 5.570$, $p < 10^{-6}$, Cohen's d = -0.554) (Fig.

2.7C). Finally, the number of time-outs in target absent trials was found to follow a valence (t(100) = 2.476, p = 0.015, Cohen's U3 = 0.436) and not a salience model (t(100) = -6.067, $p < 10^{-8}$; Cohen's U3 = 0.823). Follow-up contrasts revealed that the number of time-outs for this condition was significantly lower in rewarding blocks than in punishing blocks (t(100) = 2.476, p = 0.015, Cohen's d = -0.246), but at the same time this measure was lower in neutral blocks both than in rewarding blocks (t(100) = 3.981, $p < 10^{-3}$, Cohen's d = 0.396) and in punishing ones (t(100) = 6.111, $p < 10^{-7}$, Cohen's d = -0.608) (Fig. 2.7D).

In sum, while for target-present trials the valence pattern was present both at level of explicit accuracy and of time-out trials, target-absent trials showed a different pattern. In this case, while no difference between conditions was present at the level of explicit accuracy, data for time-outs showed both a valence and an inverse salience scheme. The lack of an inverse relationship between patterns of response in target-present and target-absent trials shows that the compound accuracy results are not the product of a shift in response bias, but rather of an effect which is present for target-present, but not for target-absent trials (or, more precisely, only at the level of time-outs in this latter case).

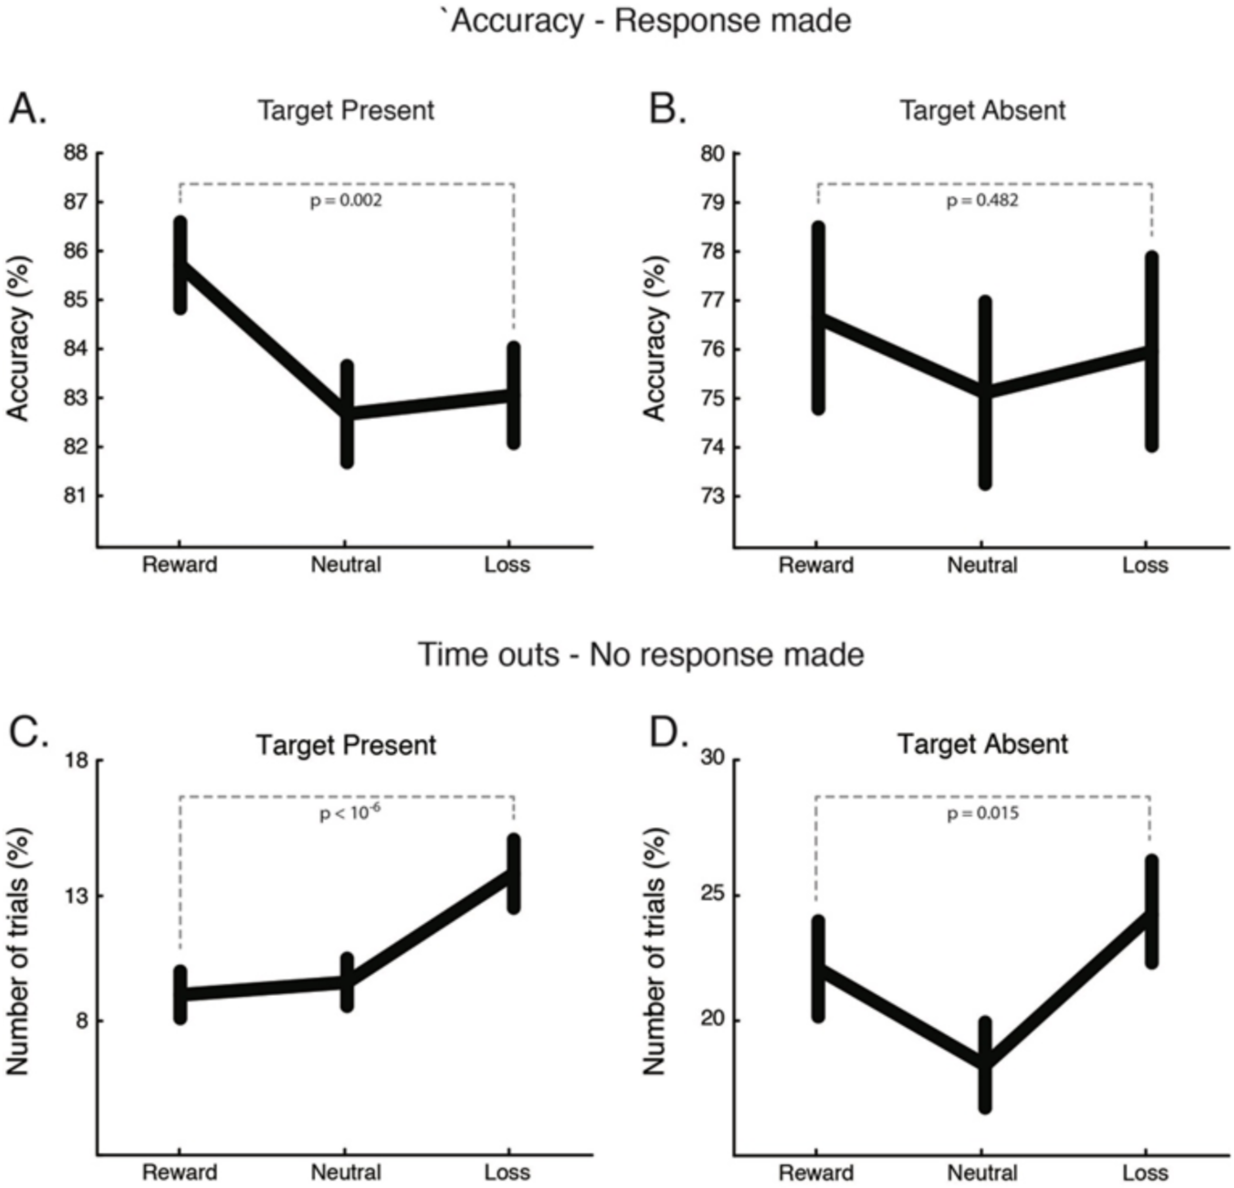

Fig. 2.7 (A) Explicit accuracy when the target was present. This analysis is based on trials where participants made an overt response within the 600 ms time limit following stimulus onset. (B) Explicit accuracy when the target was absent. (C) Number of trials where participants failed to respond within the 600 ms time limit (time-outs) when the target was present, expressed as a percentage of total trials per condition. (D) Time-outs when the target was absent.

Discussion

This study investigates the mutual relationship between the impact of reward and punishment on visual attention, and, more precisely, whether this relationship follows a valence or a salience pattern. The answer to this specific question may help shed light onto more general mechanisms underpinning motivational modulation of perception. Previous research has already aimed at disentangling the automatic component of the influence of reward on perception from the one due to the strategic allocation of attention. This was achieved by examining how stimuli previously paired with reward affected search not only when they were actively selected, but also when they acted as task-irrelevant distractors (Della Libera & Chelazzi, 2006, Hickey et al., 2010; Hickey et al., 2015a; Anderson et al., 2011). The use of punishment in this experiment provides us with another important tool in order to decouple strategic from automatic components of attention. In fact, we find that behavioral measures of target selection follow a valence-like pattern, an apparently paradoxical scheme. If participants were to be rational, a salience-like approach to the task would have resulted in a more convenient payoff, given the higher amount of points earned through correct detection of targets in rewarding and punishing blocks with respect to neutral ones. This pattern is also mirrored by measures of attentional capture, such that punished stimuli are the ones which draw less attention when acting as distractors.

These data support the concept that motivational modulation of perception does not entirely result from a deliberate allocation of resources. In this frame, then, the way in which also reward affects not only attentional capture, but even target selection, could similarly derive from a largely involuntary bias. The potential for a dissociation between endogenous and volitional control has already been the focus of a series of studies about thought suppression, stemming from the so-called white bear phenomenon (Wegner et al., 1987). More recently, Awh and colleagues have summarized situations where this type of conflict takes place at the level of visual selection. In order to solve this apparent discrepancy, they posit an alternative framework which takes into account three distinct categories of attentional controls: one due to physical salience, another one considering current top-down attentional set, and finally a third one bound to past selection history (Awh, Belopolsky, & Theeuwes, 2012). This last category would comprise simple search history bias, like for example priming of pop-out, and also involuntary biases of attention due to reward. Moreover, it would also encompass some instances of perceptual learning such as contextual cueing. All of these three phenomena are expressions of implicit memory, which could therefore represent a critical component of this new, alternative control.

Many questions remain about the matter. First of all, we have cited studies which report that reward and punishment have been shown to affect behavior in a similar or an opposing manner depending on the type of task, and sometimes depending also on the specifics of the task. An open question would therefore concern the variables which determine which is going to be the scheme which shapes behavior. One interesting perspective could regard the memory system involved in the task. A particularly enlightening example is furnished to this aim by the study of Raymond and O'Brien (Raymond & O'Brien, 2009). Here, recognition of briefly presented faces, imbued with different outcome values, was subsequently tested with and without constraints of attention. Previous studies have shown that distraction is able to bias the competition between different memory systems, namely the implicit and the explicit one, and consequently modulate the degree to which performance in a task relies on the former or on the latter system for its execution (Foerde et al., 2006). Critically, in Raymond's study, recognition followed the salience model in the full-attention condition, and the valence model in the constrained attention condition. As underlined before, many characteristics of our paradigm pushed it toward an implicit mode of execution, from the tight temporal limits available for response to the fact that category detection in naturalistic scenes relies on highly trained, experience-based skills. Could then the implicit memory system be linked to this irrational, automatic valence pattern, and the explicit system be instead driven by a strategically more convenient salience model?

Another important variable which could determine whether a valence or a salience scheme shapes attention is the type of aversive stimulus under consideration. For example, some of the studies which highlighted a salience type of pattern were using primary punishments such as electrical shocks (Schmidt et al., 2015). In our study we make instead use of monetary compensation, with a change in the nature and most likely also in the intensity of the threat. Thus, it would be important to distinguish whether the aversive stimulus evokes a feeling of actual fear or anxiety or rather a blunter one of discomfort or disgust. It could be possible that the former situation would lead to a shift from an uneven prioritization of rewarding stimuli to a more balanced bias towards all motivationally salient elements. In this case, such modulation of attention evoked by aversive stimuli could also have a separate origin, as posited by the theory of emotional attention (Brosch et al., 2011).

Another relevant issue, which is also critically connected to the previous ones, regards the neural mechanisms involved in the process. We know that dopamine release may provide the signal which allows priming of reward-associated stimuli. Theories exist about where and how this signal is implemented, but no definite answer has been drawn about the matter. In the case that the winner of the competition between memory systems defined the pattern characterizing deployment of attention, then the identification of the areas involved in each condition could clearly shed light first of all onto the specific issue of how different motivational outcomes are able to differently modulate perception. But secondly, and perhaps most importantly, this could also help understanding the way in which dopamine signal is transformed into incentive salience.

Chapter 3

Neural representation of conditioned stimuli in occipito-temporal cortex during naturalistic search follows valence, not salience

Abstract

During visual search in natural scenes, attention to a category of visual stimuli can be indexed in multi-voxel patterns of fMRI activity in occipito-temporal cortex. Encoding of a reward-associated category is enhanced with respect to a neutral one, and it is unclear whether this reflects a bias toward the positive valence of the stimulus or a broader prioritization of motivationally relevant stimuli. Here we test these two alternative hypotheses. Our task was such that correct detection of reward- and punishment-associated targets had the same relative value (100 points), but detection of a reward-associated target resulted in the receipt of 150 points (vs. 50 points for incorrect performance) whereas detection of a punishment-associated target resulted in the loss of 50 points (vs. the loss of 150 points for incorrect performance). We had two expectations: if selection is driven by the value of prior outcome, participants should preferentially encode reward-associated stimuli (valence model). However, if resources are deployed to stimuli that are motivationally relevant, reward-predictive and loss-predictive stimuli should be equally represented (salience model). We found that the amount of information for the different visual categories in occipito-temporal cortex followed the valence, and not the salience, pattern. Moreover, stimulus-evoked activity in the dopaminergic midbrain predicted the quality of these representations. At the same time, posterior parietal cortex was also found to encode information about attended categories, but this time following the latter of our expectations, namely the salience pattern.

Introduction

In the last few years, MVPA has proven to represent a powerful tool in order to investigate brain activation during visual search in naturalistic, cluttered environments. More specifically, object selective visual cortex (OSC) has shown to retain information about the high-level categories of objects that a subject is actively looking for, and to suppress information about salient, but task-irrelevant, distractors (Peelen et al., 2009; Seidl et al., 2012). Recent work has also demonstrated that this categorical visual representation is additionally modulated by motivational outcome, such that pairing a real-world category with reward acts by relatively increasing the amount of information about that category in the occipito-temporal cortical region of interest. Conversely, its representation is more strongly inhibited with respect to a neutral one, when this category turns to act as a distractor (Hickey & Peelen, 2015).

Our study stems on one side from this existing literature regarding how MVPA is able to assess visual representation during search, and on the other side from the results of an experiment that we have recently performed, which looked at how reward and punishment could impact behavioral measures during search in natural scenes (Chapter 2). We had made the two alternative hypotheses that these motivational outcomes affected search in a similar or in an opposing manner, namely according to a salience or a valence pattern, respectively. We found that accuracy measures followed, both for target selection and for attentional capture, the scheme which was less expected in terms of its strategic utility, namely the valence one, suggesting therefore a bias toward this pattern of behavior. In this fMRI study, we aim at bridging the gap between these two lines of study. On one side, we use this technique in order to better elucidate the neural mechanisms of an unexpected behavioral finding. On the other, we mean to extend previous results regarding the influence of motivational outcomes on category information in OSC. The relevance of this question lies on the idea that the increase in representation that we see for reward could come from different origins (Maunsell, 2004). On one hand, it could derive from the higher utility value that an object paired with reward retains, possibly resulting in a higher degree of voluntary attention deployed to the task. In this case, a similar motivational role would be played by an equally valuable evasion of loss. But on the other hand, it could reflect the incentive properties that reward attributes to an object, increasing its desirability and attractiveness. In this case, an object consistently paired with loss would show a relationship with reward which is exactly the opposite with respect to the one that we have just considered (Fig. 3.1).

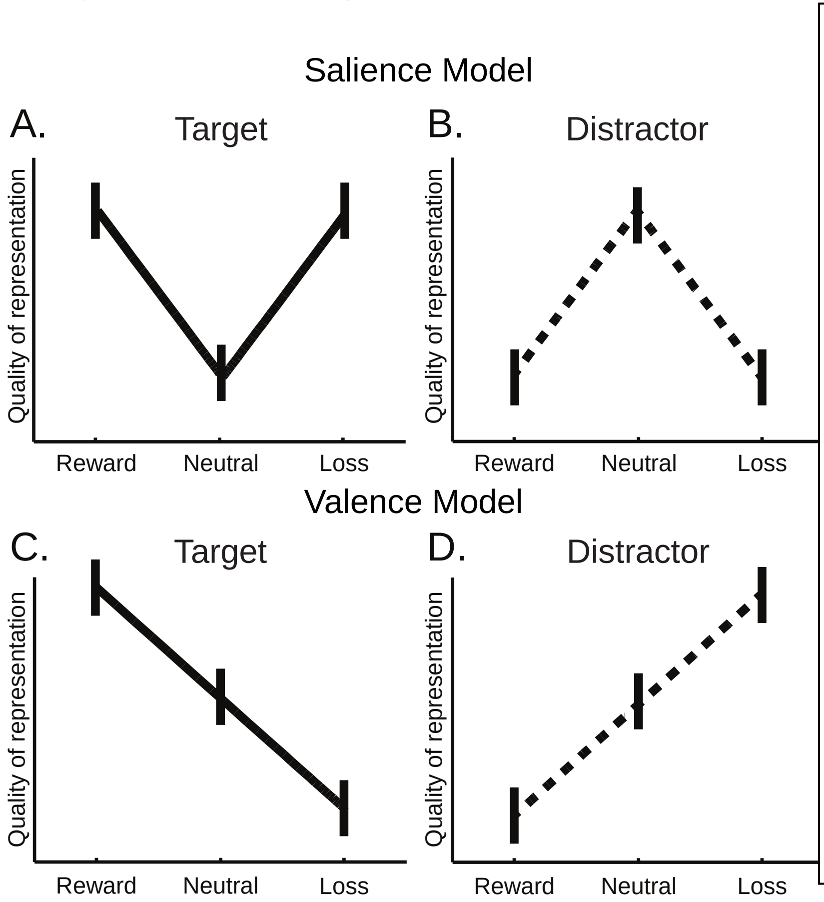

Fig. 3.1 Predictions from the salience and the valence models. (A) For the salience model, both reward- and loss-associated targets should be more strongly represented than neutral ones, and (B) they should draw more attention when acting as task-irrelevant distractors. (C) According to the valence model, targets associated with reward should be more strongly represented than neutral ones, but loss-associated targets should be less so. (D) Correspondingly, distractors associated with reward should require stronger attentional suppression than distractors associated with punishment.

To this aim, we employed a visual search paradigm where participants had to detect examples of four categories of objects (people, cars, houses and trees) in pictures of naturalistic scenes (Fig. 3.2A). The target category changed blockwise, with the other three categories acting as task-irrelevant distractors in all of the scenes of the block. For each subject, one of the four categories was associated with monetary gain, one with monetary loss, and two with a neutral outcome (Fig. 3.2B). Correct detection of the rewarding category resulted in the receipt of 150 points, while failure to do so resulted in receipt of only 50 points. For the punishing category, correct detection determined the loss of 50 points, but incorrect performance caused instead a greater loss of 150 points. Finally, neutral categories garnered only 1 point when correctly detected, and failure to do so resulted in the loss of only 1 point. Net value of correct detection of rewarding and punishing categories was therefore equally relevant (100 points), with respect to the negligible value of neutral categories (2 points).

We then looked at how different motivational outcomes affected category information on one side for categories acting as targets, and for categories acting as distractors. Once again, we approached our data considering two alternative models, the valence and the salience models, that could have possibly shaped the deployment of attention. We subsequently tested the hypothesis that dopamine released by the midbrain could be related to motivational modulation of information in OSC. Thirdly, we looked for areas of the brain which could discriminate between the categories which were attended across blocks. Finally, we performed an exploratory whole-brain univariate analysis, in order to see which areas of the brain were activated according to our contrasts of interest during the execution of this task.

Materials and methods

Participants.

Twenty-four participants (9 female, mean age 24.5 ± 4.1 SD) took part in the experiment. All procedures were carried out in accordance with the Declaration of Helsinki and were approved by the ethical committee of the University of Trento. One participant was excluded from the analyses because of low performance (d’ = 0.05).

General structure of the experiment.

The experiment consisted of three different parts: 1) the OSC localizer, 2) the category pattern localizer, and 3) the visual search task. For all three phases, stimuli were presented using Psychtoolbox (Brainard, 1997) and projected on a translucent screen at the end of the scanner bore. Participants viewed the screen through a mirror

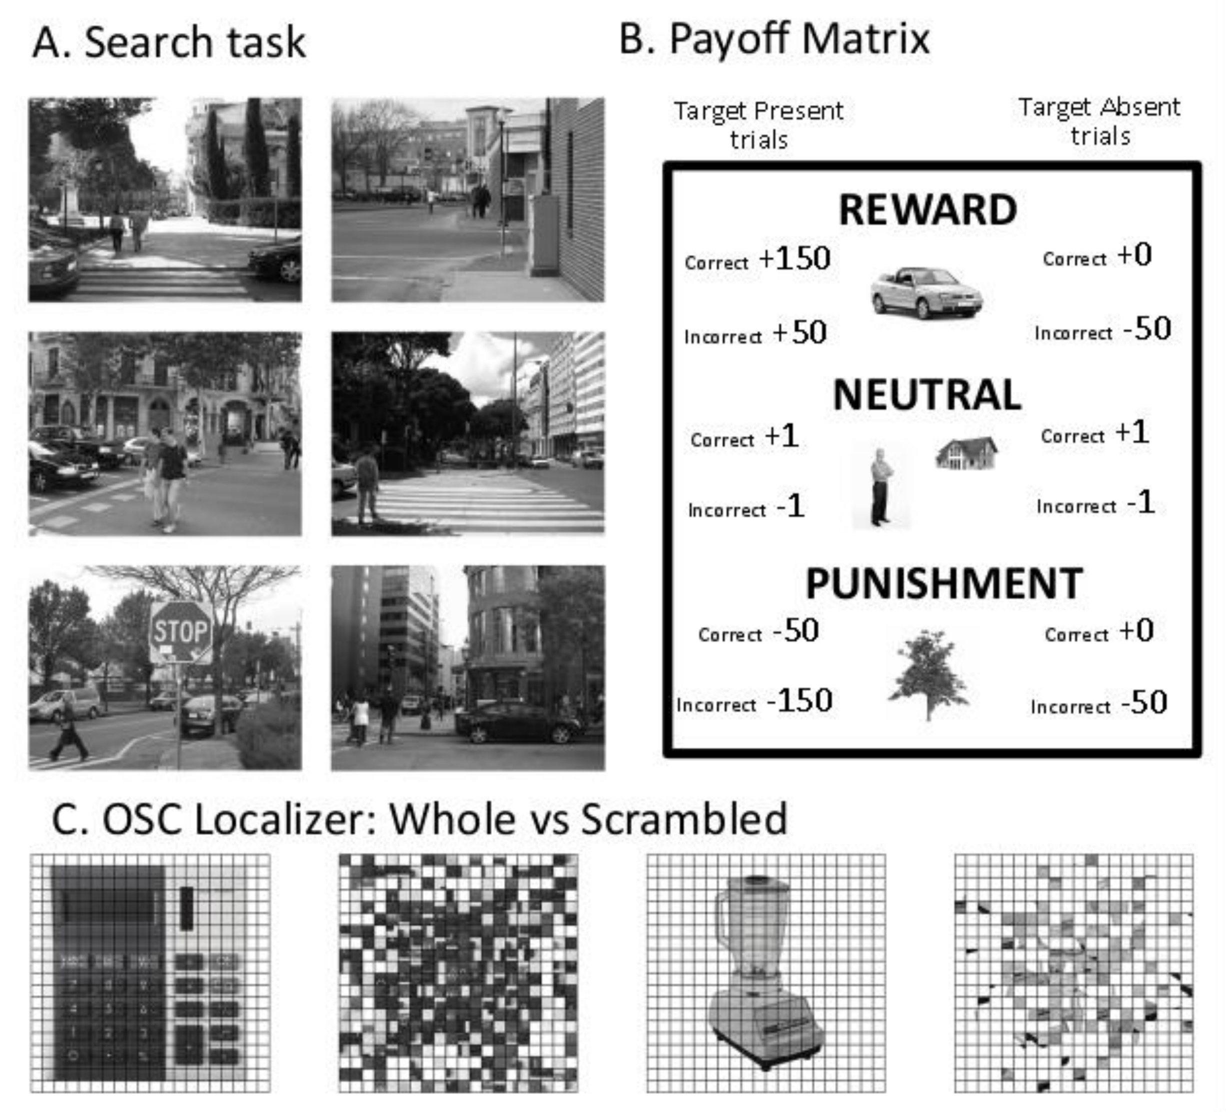

Fig. 3.2 (A) Examples of real-world scenes employed for the search task. (B) Feedback schedule. One category was paired with reward, one with punishment and the other two with a neutral feedback. Associations were counterbalanced across subjects. Feedback for each trial depended on category type, category presence and correctness of response. (C) Examples of whole and scattered object images for the OSC Localizer.

mounted on the head coil. At the end of the MR session, each subject was administered a BIS/BAS scale questionnaire (Carver & White, 1994). Participants were paid based on the number of points accumulated during the experiment.

OSC localizer.

Stimuli. The stimulus set for the OSC Localizer consisted of black and white pictures (27° × 38° visual angle) of 20 daily-life objects (e.g. telephone, cheese, alarm clock) and their scrambled version (Fig. 3.2C).

Procedure. Participants performed 2 runs of OSC localizers of 317 s duration, each containing 16 blocks of 20 trials and 3 fixation blocks. Each run began and ended with 15 sec fixation. Before the first block of the run, a 2 s display reminded the task to the subject. During each 15 s block, pictures of one type only (intact or scrambled) were presented, while subjects monitored for image repetition, which occurred once in a block. Each trial started with a fixation cross (350 ms) followed by the picture of the intact or scrambled object (400 ms). Every fifth block was a 15 s fixation block with no stimuli.

Category pattern localizer.

Stimuli. The stimulus set for the category localizer consisted of four groups of stimuli, one for each of the categories cued in the visual search task (people, cars, houses and trees). Each group consisted of 40 black and white pictures (27° × 38° visual angle) of isolated exemplars of the same category on a white background. Pictures of people were headless, as faces in the search task scenes were generally too small to be visually solved.

Procedure. Participants performed 2 runs of category localizers of 392 s duration, each containing 20 blocks of 20 trials and 4 fixation blocks. Each run began and ended with 15 sec fixation. Before the first block of the run, a 2 s display reminded the task to the subject. During each 15 s block, pictures of different exemplars of only one of the four categories were presented, while subjects monitored for image repetition, which occurred once in a block. Each trial started with a fixation cross (350 ms) followed by the picture of the intact or scrambled object (400 ms). Every fifth block was a 15 s fixation block with no stimuli.

Visual search task.

Stimuli. The stimulus set for the visual search task consisted of black and white pictures (27° × 38° visual angle) of natural scenes (n=384) selected from an online database (Russell et al., 2008). 192 pictures contained all four categories (people, cars, houses and trees), and the remaining 192 were organized into four groups (n=48), each one containing three out of four categories. Natural scene photographs were followed by perceptual masks of equal size, generated by combining white noise at different spatial frequencies and superimposing naturalistic structure on the noise (Walther et al., 2009).

Procedure. Participants performed 4 runs of 590 s duration, each containing 8 blocks of 60 s. Each run began and ended with 15 sec fixation. Before each block started, a display was presented for 10 s, indicating the target of that block, and the score which had been obtained up to that moment. Each of the four categories acted as target twice per run, and order of target types within a run was counterbalanced across runs. Each block contained 12 trials where target was present (pictures containing all four categories), and 12 trials where target was absent (pictures containing all the categories but the target) in a randomized order. The total amount of trials was 768 per subjects, so that each scene was shown twice. Each trial started with a fixation cross (833 ms) followed by the brief presentation of the scene (58 ms), a mask (325 ms), the reappearance of fixation (750 ms), and the feedback (533 ms; Fig. 3.3A). Participants reported the presence of the target category with the right index finger response and its absence with the right middle finger response. Responses given outside of the 750 ms fixation interval, or absence of response, garnered an incorrect performance. Feedback for each trial depended on type of target (rewarded, punished or neutral), type of trial (target present or target absent) and correctness of response (Fig. 3.2B). For neutral targets, points were +1 for correct and -1 for incorrect response both for target present and for target absent trials. For rewarding blocks, subjects received +150 points for correct, and +50 for incorrect response in target present trials, +0 for correct and -50 for incorrect response in target absent trials. For punishing blocks, subjects got -50 points for correct and -150 for incorrect response in target present trials, and +0 for correct and -50 for incorrect response in target absent trials. For each participant, in a counterbalanced order, one category was paired with reward, one with punishment, and the remaining two categories were neutral.

fMRI Data Acquisition.

Imaging was conducted on a Brucker BioSpin MedSpec 4T head scanner (Bruker BioSpin), equipped with an eight-channel head coil. T2*-weighted gradient-echo echo-planar images were collected as functional volumes for all the functional runs (EPI; repetition time = 2.2 s, echo time = 33 ms, flip angle = 76°, 3 × 3 × 3 mm voxel size, 0.45 mm gap, 31 slices, 192 mm field of view, 64 × 64 matrix size). A T1-weighted image (MPRAGE; 1 × 1 × 1 mm voxel size; 256 slices, 224 × 176 matrix size) was obtained as a high-resolution anatomical reference.

fMRI Preprocessing.

All neuroimaging data were analyzed using MATLAB and SPM12. The volumes were motion corrected, slice time corrected, coregistered to the structural image and spatially normalized to the Montreal Neurological Institute (MNI) template (as included in SPM12). Functional volumes were then smoothed using a 6-mm full-width half-maximum Gaussian kernel. All analyses were performed on the smoothed data.

fMRI Data Analysis.

A separate general linear model (GLM) was created for each of the three phases. The blood-oxygen-level-dependent signal of each voxel in each participant was modeled

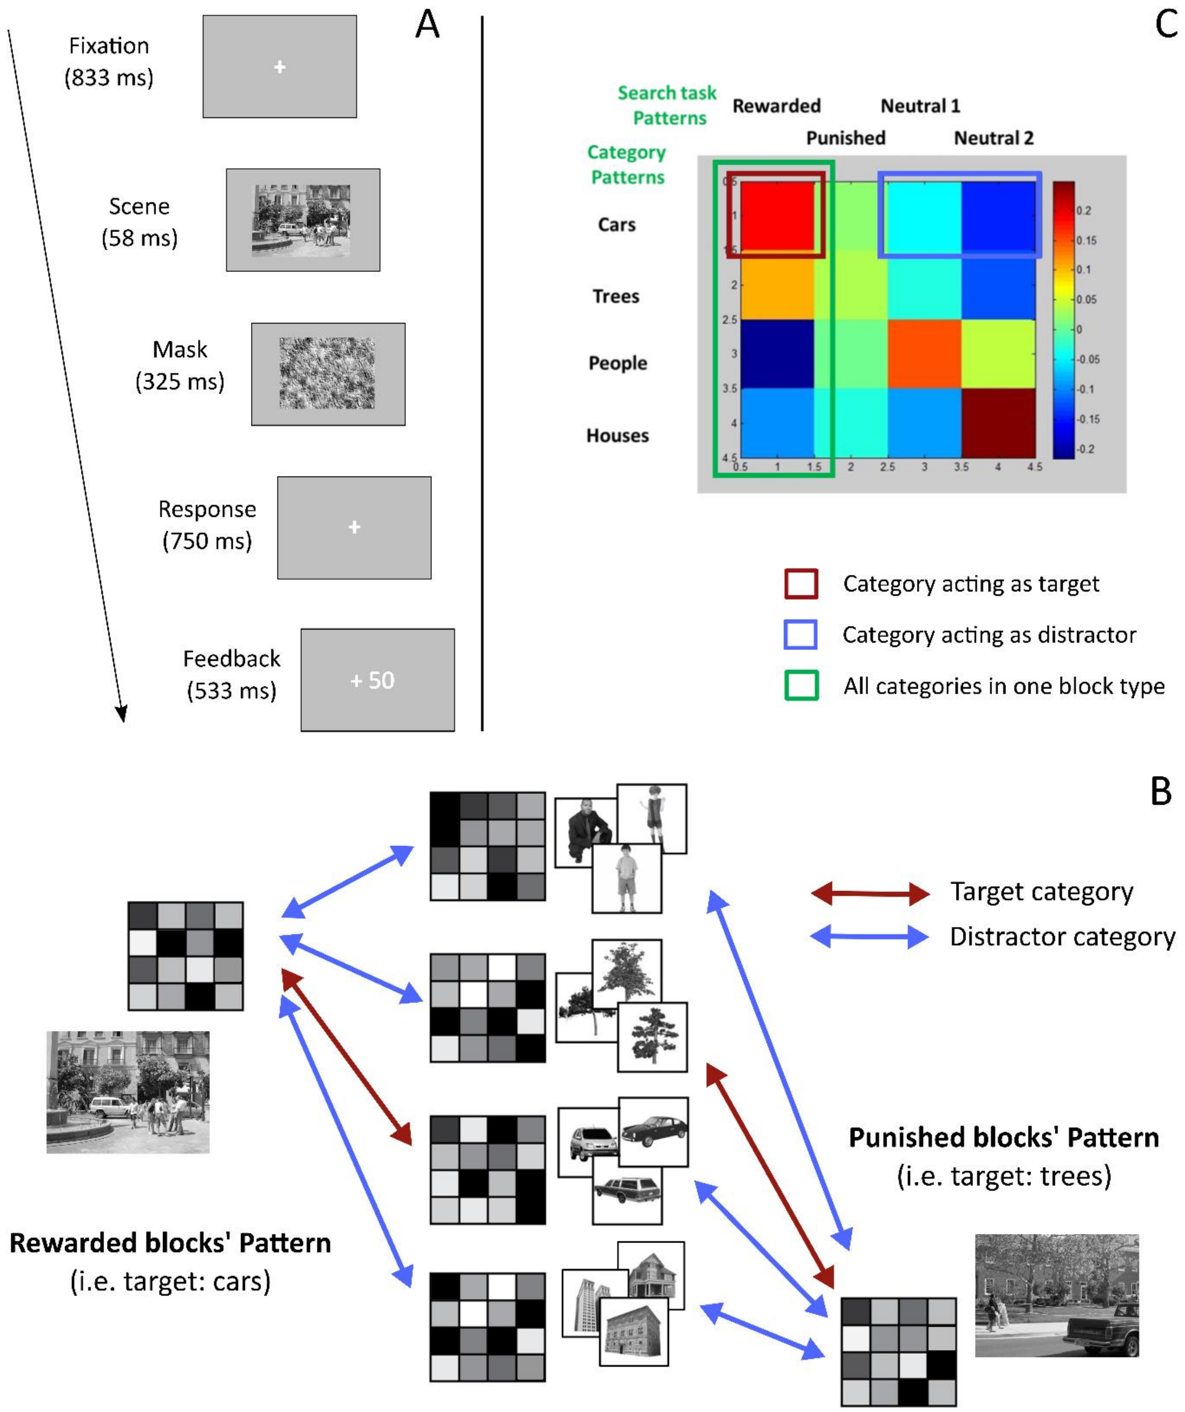

Fig. 3.3. (A) Trial structure. Subjects reported the presence of a cued target category in the scene. (B) Data Analysis. Patterns from the different search task blocks were correlated with category localizers' patterns. In each case, correlation of the scene both with target and with distractor categories were calculated. (C) This yielded a matrix were values along the diagonal index degree of correlation of the scenes with the different targets. Along a column (framed in green), we see the correlation of the scene with the target categories (framed in red), and with the three distractor categories. For each category, degree of correlation with the scene when acting as a distractor is calculated by averaging the value obtained in the two neutral blocks (framed in blue).

using two regressors of interest in the OSC localizer, one for intact and one for scrambled objects. Four regressors of interest representing the various categories (people, cars, houses and trees) were used both in the category localizer and in the visual search task, in the former case indicating the block type and in the latter the block target. In all these cases each regressor spanned a whole block. A second GLM was estimated for the search task phase, which considered target-present and target-absent trials separately for each block, leading therefore to eight regressors of interest, two per category. Finally, a third GLM considered correct and incorrect trials separately, yielding sixteen regressors of interest, four per category. In all models, six additional regressors of no interest obtained from the realignment procedure were included to account for head motion. All models also presented an intrinsic temporal high-pass filter of 1/128 HZ to correct for slow scanner drifts.

ROI definition.

OSC was functionally defined in each participant by contrasting, within occipito-temporal regions, responses evoked in the intact and in the scrambled condition ($p < 0.001$, uncorrected). Mean OSC size was 1491 voxels ± 845 SD. ROI was created using the MarsBar Toolbox (Brett, Anton, Valabregue, & Poline, 2002).

MVPA.

All pattern analyses were performed using CoSMo-MVPA (Oosterhof, Connolly, & Haxby, 2016). For each subject, four patterns (one for each object category) were obtained from the category localizer, and another four patterns (one for each target type) from the visual search task. This was done by extracting t values of each voxel for each condition (obtained from the first GLM, with regressors spanning the whole block), only for voxels comprised in the ROI of each subject. Following existing work, t values were normalized by subtracting, for each voxel, the mean value obtained for that voxel across all conditions of the single task. Normalized t values of all four conditions from the category localizer were then correlated with all four conditions of the visual search task, across all voxels of the ROI (Fig. 3.3B). This yielded a 4 × 4 correlation matrix, whose values were then Fisher-transformed and organized in terms of whether the category for that block was acting as a target or as a distractor (Fig. 3.3C). Each category was acting as distractor in three different blocks, but, for the distractor measure, we only considered values obtained in the neutral blocks. This was done to have a common baseline (neutral targets), where we would measure category information for distractors of different motivational conditions in a consistent frame. Still, there is no change in results if a mean of all three blocks where category is acting as a distractor is used for the same measure. Target and distractor values for the two neutral categories were then averaged in each subject, resulting in a 2 × 3 (task relevance × motivational condition) matrix (Fig. 3.5).

Statistical Analysis.

In order to see which model could best fit the data, we performed a multiple linear regression separately for each subject, with the two models of our main hypotheses as regressors, plus an additional regressor accounting for task relevance. The 6 MVPA values representing a) the different targets (rewarded neutral punished) and b) the different distractors (same order) were modeled using the following predictors: one accounting for a valence-based model of targets and distractors [+1 0 -1 -1 0 +1], a second one accounting for a salience-based one [+1 -2 +1 -1 +2 -1], and a third one accounting for task relevance [1 1 1 -1 -1 -1]. Values were standardized for each vector, so that the first vector would become [+1.12 0 -1.12 -1.12 0 +1.12], et cetera. Note that, for each vector, distractor values are simply the opposite of the target values. This yielded one beta coefficient per subject per model. For each model, all the beta coefficients were then t tested against zero, to see which model could correctly fit the data. Finally, the salience model beta values were subtracted from the valence model beta values individually for each subject, and the resulting differences were again t tested against zero to make a direct comparison between the two models. We used the same approach to analyze behavioral data, but using three digits vectors this time, as no measure was present for a distractor-related effect. All statistical values reported in the paper are a product of permutation analysis. For tests against the null hypothesis of zero, data-driven distributions were generated by randomly sampling from the relevant dataset 10,000 times with replacement. The likelihood of observed data given the null was calculated through comparison of observed data to these distributions. Planned comparisons were achieved by treating difference scores the same way.

Correlation between univariate activity and information.

T values from the second GLM, which considered target-present and target-absent trials separately, were used for this analysis, focusing on target-present trials only. For each subject, a single t value was calculated for a valence [+1 0 -1], and another one for a salience [+1 -2 +1] contrast of the mean univariate activity during the visual search task, for two ROIs, one corresponding to SN and the other one to Red Nucleus (Fig. 3.5A). For this regression we only considered targets, as all distractors were always present in all the scenes of each block type, and therefore no differential analysis could be conducted among them within each search task condition for univariate activity. Individual t values of both contrasts were then correlated with the individual beta coefficients of the multiple linear regression of OSC category information modulation of the corresponding model. The regression was similar to the one described for the main analysis, but again only value for targets were considered (because of the different analysis performed on univariate activity). Correlation values of these two regions were then compared through Steiger's test, which takes into account the degree of correlation between the two variables not shared in the two previous tests, i.e. univariate activity of SN and Red Nucleus. Regions were anatomically defined through WFU PickAtlas, a software which is based on the Talairach Daemon database (Maldjian, Laurienti, Kraft, & Burdette, 2003).

Searchlight analysis.

The searchlight analysis tested for regions that discriminated between the four targets based on the category localizer patterns. For each voxel in the brain, we computed voxelwise correlations in a sphere of 21-mm radius around this voxel. The correlation values from each sphere were Fisher transformed and assigned to the center voxel of this sphere. The average correlation between matching categories was contrasted with the average correlation between nonmatching categories. These net values were computed for each voxel separately, and then t tested across subjects. The threshold was set to $P < 0.001$ (uncorrected) and a minimum cluster size of 50 voxels.

Univariate analysis.

The regressors of the correct target-present trials of the four motivational categories [reward neutral1 neutral2 punishment] were contrasted according to the valence [+1 0 0 -1] and to the salience [+1 -1 -1 +1] pattern. Individual participants' contrast images entered a second level t-test (threshold set at $P < 0.001$, cluster corrected for multiple-comparisons, minimum cluster size of 20 voxels).

Results

Behavioral Analysis. Accuracy values for the three different motivational conditions are shown in fig. 3.4A, and RT values in fig. 3.4B. The three values of each type of measure were regressed with the two models, one accounting for the valence model and one for the salience model. Neither of the two models significantly fitted the data for accuracy measures (respectively, $p = 0.2053$ and $p = 0.3861$). Also for RT measures, we did not find a significant fit ($p = 0.6593$ for regression with the valence model, and $p = 0.6312$ for the one with the salience model). When looking at accuracy results for target-present and target-absent results independently, we found a pattern which looked very similar to our first behavioral study (Chapter 1). Namely, participants were better in detecting the presence of a reward-associated target than the one of a loss-associated target (0.76 vs 0.69, $p = 0.008$), but were no better in reporting the absence of reward-associated targets with respect to absence of loss-associated ones (0.69 vs 0.71, p = 0.312; interaction, p = 0.017).

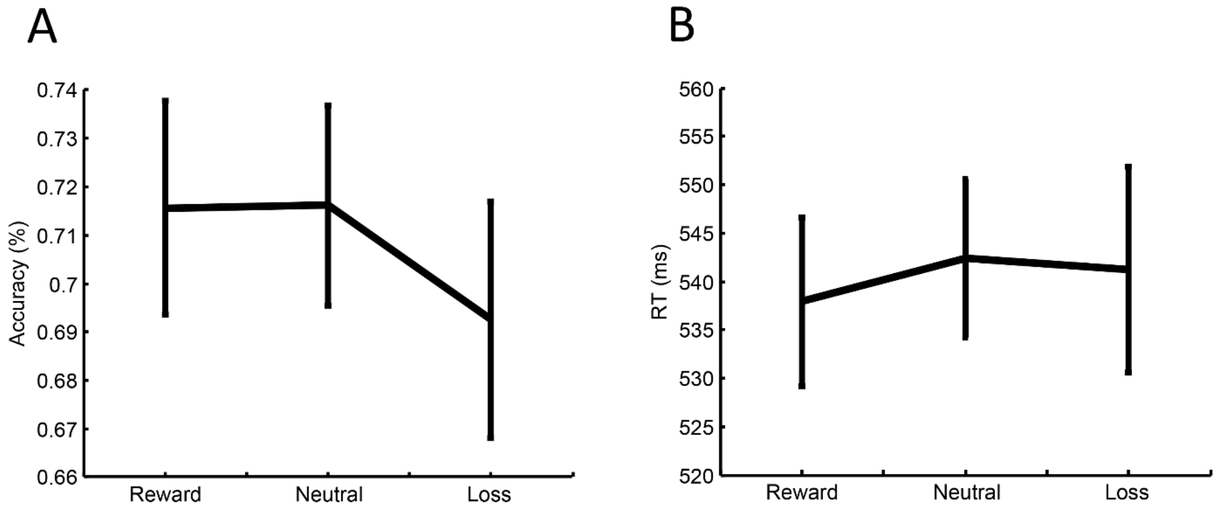

Fig. 3.4 (A) Accuracy values for detection of categories associated with different types of motivational outcomes. (B) RT values for detection of categories associated with different types of motivational outcomes.

MVPA Analysis. We looked at how different motivational outcomes affected category information on one side for categories acting as targets, and for categories acting as distractors. While target categories were present only in half of the trials of each block, the remaining three distractor categories were always present in all of the scenes of the block. Once again, we found that representation of visual categories in OSC as assessed by MVPA followed a valence, and not a salience, pattern. Category information values for the three motivational conditions when acting as targets are shown in fig. 3.5A, and for categories acting as distractors in fig. 3.5B.

These six values were regressed with three models, one accounting for task-relevance, one for valence and one for salience. Both the task-relevance and the valence-based model significantly fitted the data (respectively, p = 0.001 and p = 0.004), while the salience-based one did not (p = 0.217). Also the direct comparison between the values coming from the two models showed a significant fit in regressing the data in favor of the valence model (p = 0.006). Follow-up contrasts revealed that OSC carried more information about reward-associated targets than loss-associated targets, $p < 10^{-4}$, but less information about reward-associated distractors than loss-associated distractors, $p = 0.031$.

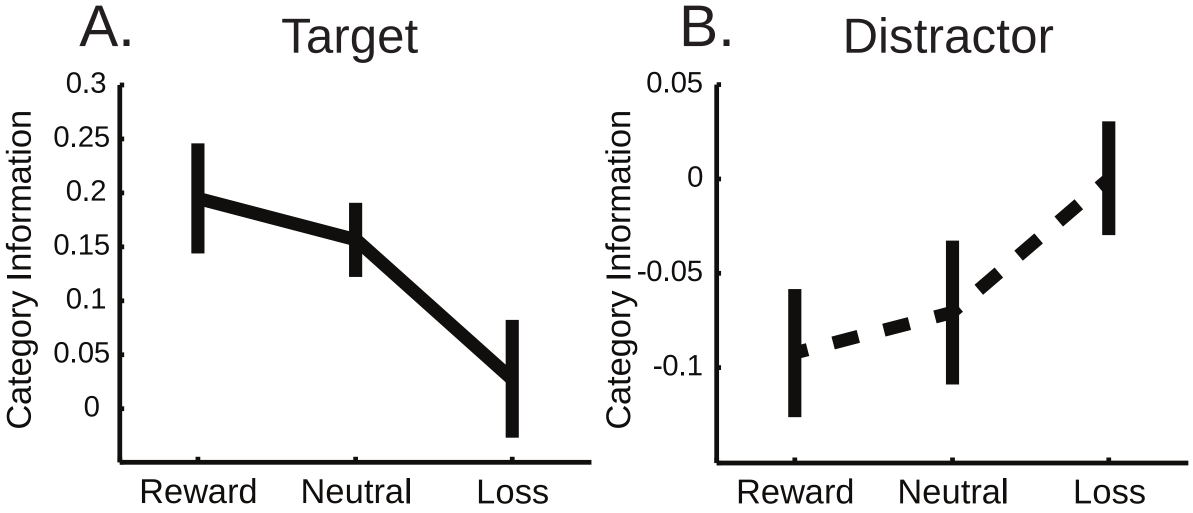

Fig. 3.5 (A) Amount of category information for categories acting as targets. OSC carries more information about reward-associated categories than about loss-associated ones. (B) Amount of category information for categories acting as distractors. Reward-associated categories are more strongly suppressed than loss-associated ones. Error bars reflect within-subject standard error.

Secondly, we investigated the hypothesis that univariate activity of midbrain dopamine during the search task could be related to modulation of category information in OSC. We used an atlas to anatomically define two regions of interest (ROIs) in the midbrain: Substantia Nigra (SN) and Red Nucleus (RN) (fig. 3.6A). The former represented the dopaminergic structure involved in our hypothesis. The second one, although located in close proximity to SN, is not involved in dopaminergic output, so it was chosen as a midbrain control region. Again, we considered two different contrasts for univariate activity, a valence [+1 0 -1] and a salience one [+1 -2 +1]. In each ROI, this yielded a t value per subject per model, representing how much, in each participant, univariate activation of that region followed a valence or a salience pattern. Similarly, the entity of modulation of category information according to a valence or a salience model in OSC obtained in the regression analysis was represented by a specific beta coefficient per subject per model. The individual beta coefficients of the two models were then correlated with the individual t values garnered through the corresponding contrasts.

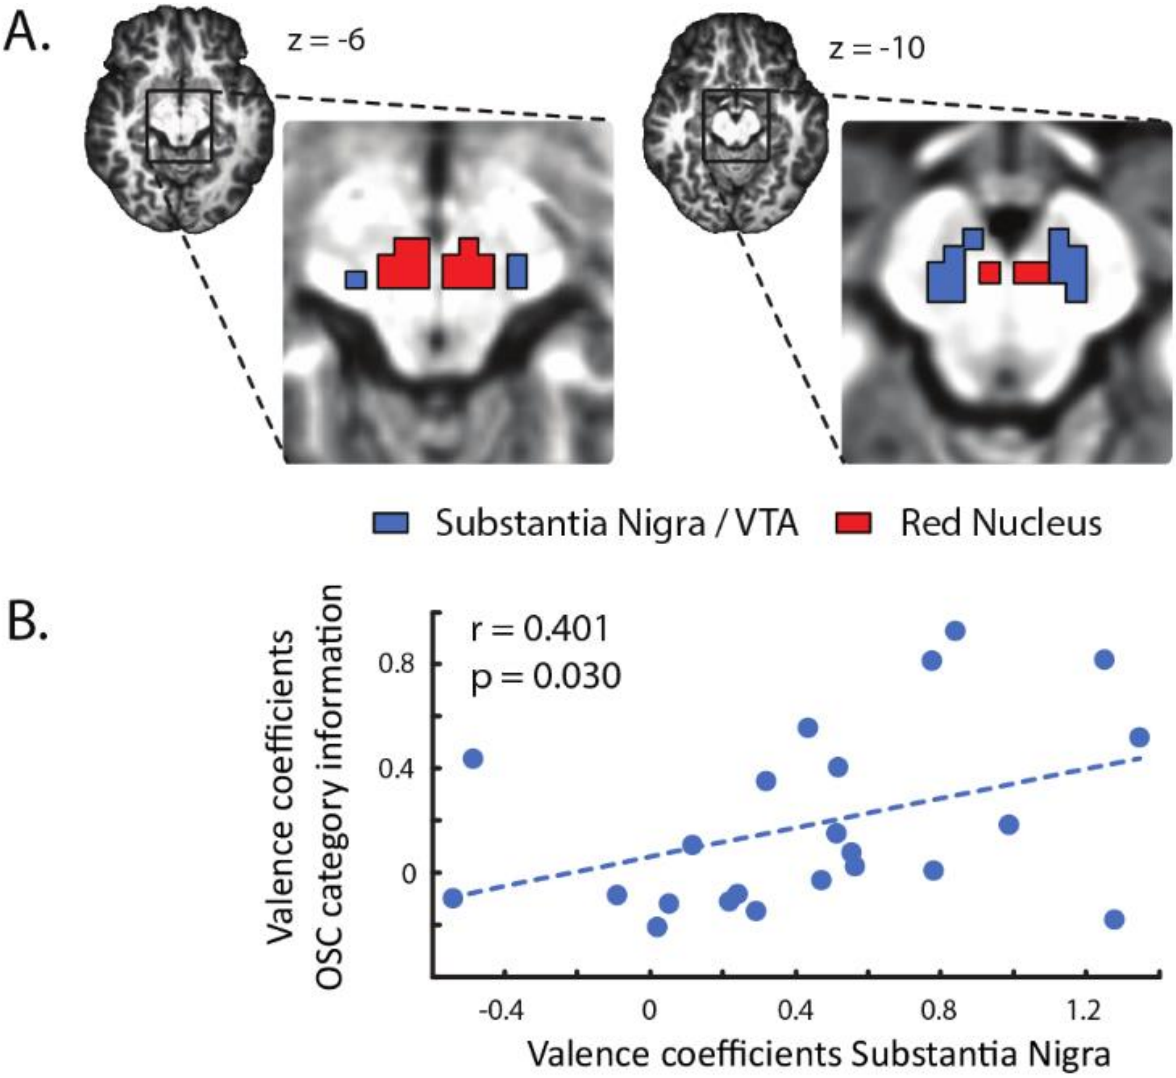

Fig 3.6 (A) Anatomically defined ROIs representing SN and RN. (B) Correlation between univariate activation of SN and categorical information in OSC.

For our SN ROI, the entity of modulation of category information in OSC according to a valence-based model correlated significantly with the univariate activity of the ROI (r = 0.401, p = 0.03 (one-tailed t-test), fig. 3.6B), but this was not the case for the salience-model correlation (r = -0.050, p = 0.545). For our RN ROI, neither of the two correlations resulted significant (valence model: r = -0.015, p = 0.39; salience model: r = -0.307, p = 0.91). Finally, we verified that the values obtained in the two regions for the valence models (r = 0.401 for SN and r = -0.015 for Red Nucleus) were significantly different. To do this, we used Steiger's test, which is suited for correlations sharing one variable in common, namely modulation of category information in OSC. The two correlations showed a trend towards a significant difference according to Steiger's test (t(20) = 1.331, p = 0.992).

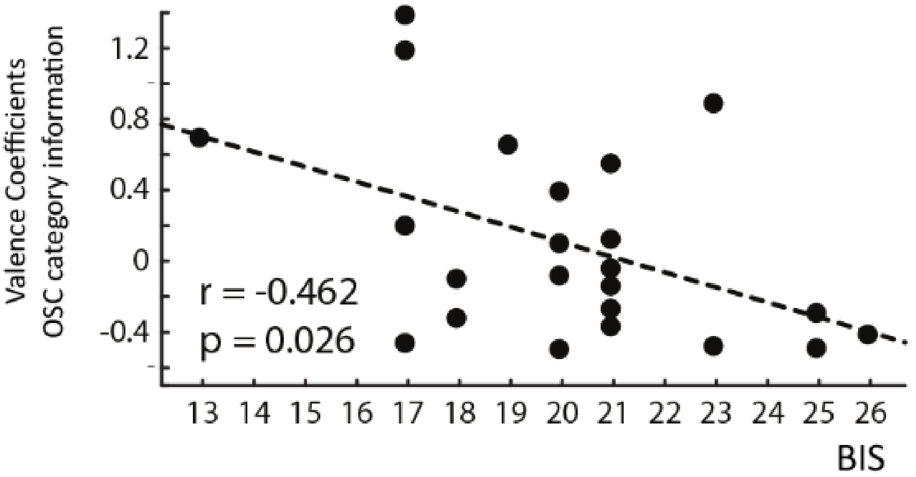

Fig 3.7 Relationship between BIS values and valence coefficients for category information in OSC.

Thirdly, we examined the relationship between the measures of OSC information content and midbrain activity described above, with scores obtained through the BIS/BAS questionnaire, a personality test administered at the end of the experiment. Individual Beta coefficients derived from regression of category information in OSC with a valence model correlated negatively with individual BIS (r = -0.462, p = 0.03, fig. 3.7). This score measures the individual's response to anxiety-relevant cues. Thus, the general overall bias for a stronger visual representation of rewarding stimuli with respect to punishing ones, was less pronounced in participants who showed higher sensitivity to aversive stimuli.

Searchlight Analysis. Additionally, we used a whole-brain searchlight in order to look for regions of the brain which retained a representation of the category of the attended target. In each voxel of the brain, we tested the degree to which multivoxel patterns in a 21-mm sphere around this voxel could discriminate the target category based on the category-specific patterns from the independent localizer. Three separate clusters came out of this analysis: two symmetrical ventral areas, and a dorsal one closer to the midline. The ventral clusters roughly corresponded to OSC, even though in our main analysis this region was evaluated on an individual basis, while in this case we created a mean region across subjects (Fig. 3.8).

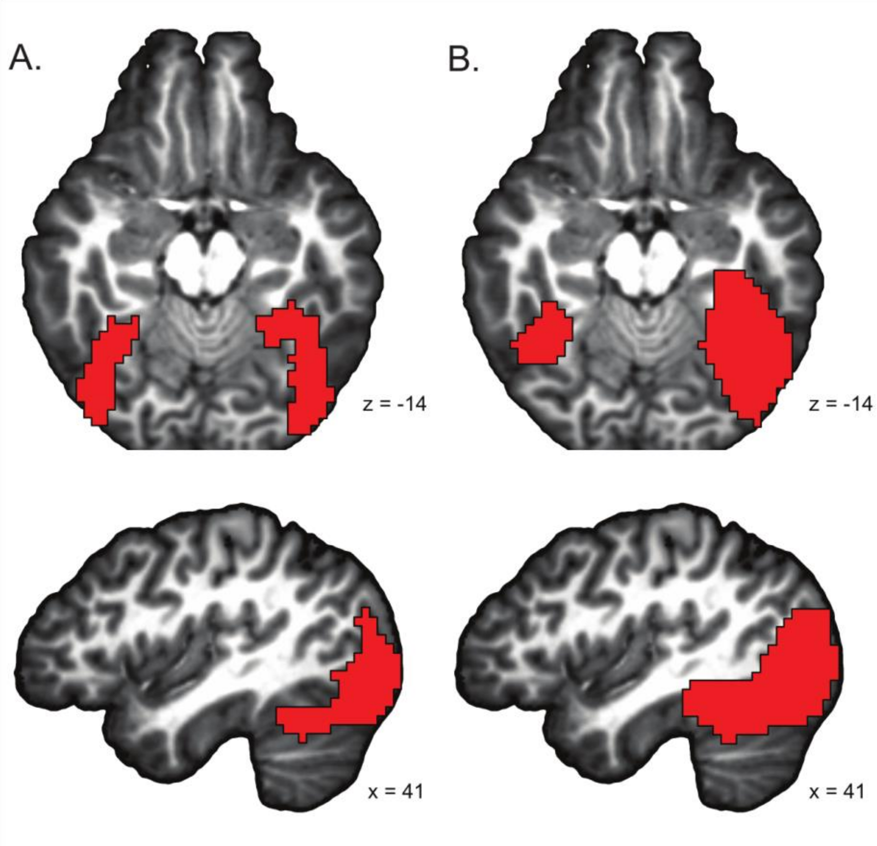

Fig. 3.8 (A) OSC as defined in the OSC localizer. Voxels identified here were present in the OSC of 16 or more of the 23 participants of the study. (B) Results from the searchlight contrast for targets vs distractors. Voxels identified here constitute the center of spheres that were selective for targets at $p < 0.001$ with a cluster threshold of 50 voxels.

The dorsal cluster was located in correspondence to right PPC (Fig. 3.9A). This is consistent with a recent study reporting the representation of abstract object identity in this region, as assessed by MVPA (Jeong & Xu, 2016). In order to test whether category information was modulated by motivational outcome also in this area, we performed, with the six values representing targets and distractors, a similar regression analysis as we did in OSC, using the task-relevance, salience and valence models as predictors (Fig. 3.9B). In this case we found, first of all, that coefficients for the valence model did not differ from zero, $p = 0.496$. At the same time, instead, the salience model was significantly represented, $p = 0.002$. Task-relevance fit was unsurprisingly positive, $p < 10^{-4}$. Follow-up contrasts revealed that there was a trend for a stronger representation of reward- than for neutral targets, $p = 0.116$, and a similar trend for a stronger representation of loss-associated than neutral categories, $p = 0.123$. More reliable results were present about distractor categories, such that reward-associated targets were more strongly suppressed than neutral ones, p = 0.020, and a similar suppression with respect to the neutral category was present for the punishment-associated target, p = 0.003.

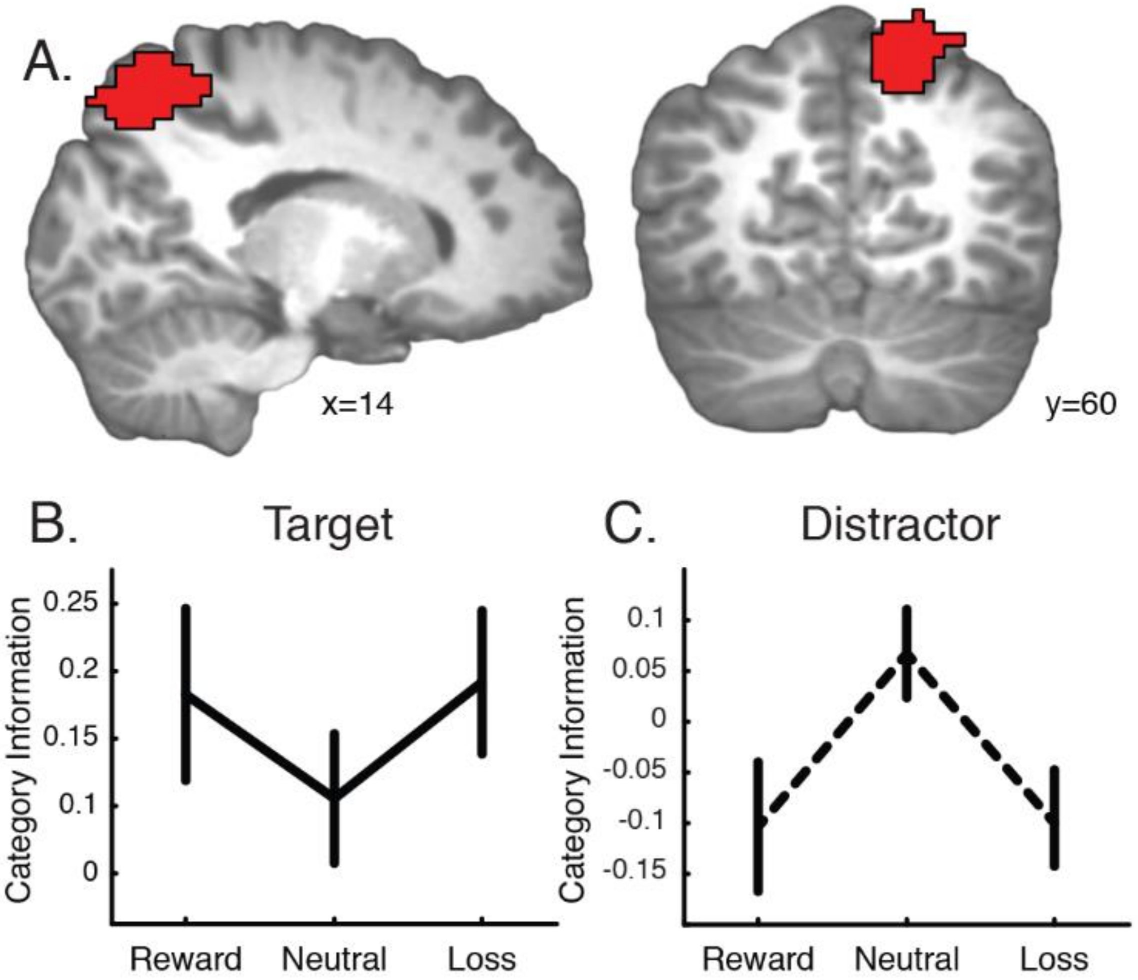

Fig. 3.9 (A) Parietal cluster identified in the searchlight analysis. The region was defined by contrasting information content for targets vs distractors. Centroid: +9, -73, +43, MNI space. (B) Representation of the different categories in the parietal cluster. This time category information for motivational categories was found to follow the salience, and not the valence, pattern.

Univariate Analysis. Finally, we also performed a whole-brain analysis which looked at the univariate activation for our two models, the valence and the salience one. The salience contrast revealed an increase in activation in bilateral middle frontal gyrus (MFG), left inferior frontal gyrus (IFG), right inferior parietal lobe (IPL), right anterior insula (AI), anterior midcingulate cortex (aMCC), posterior cingulate cortex (PCC), precuneus, bilateral fusiform gyrus and cerebellum (Table I and Fig. 3.10).

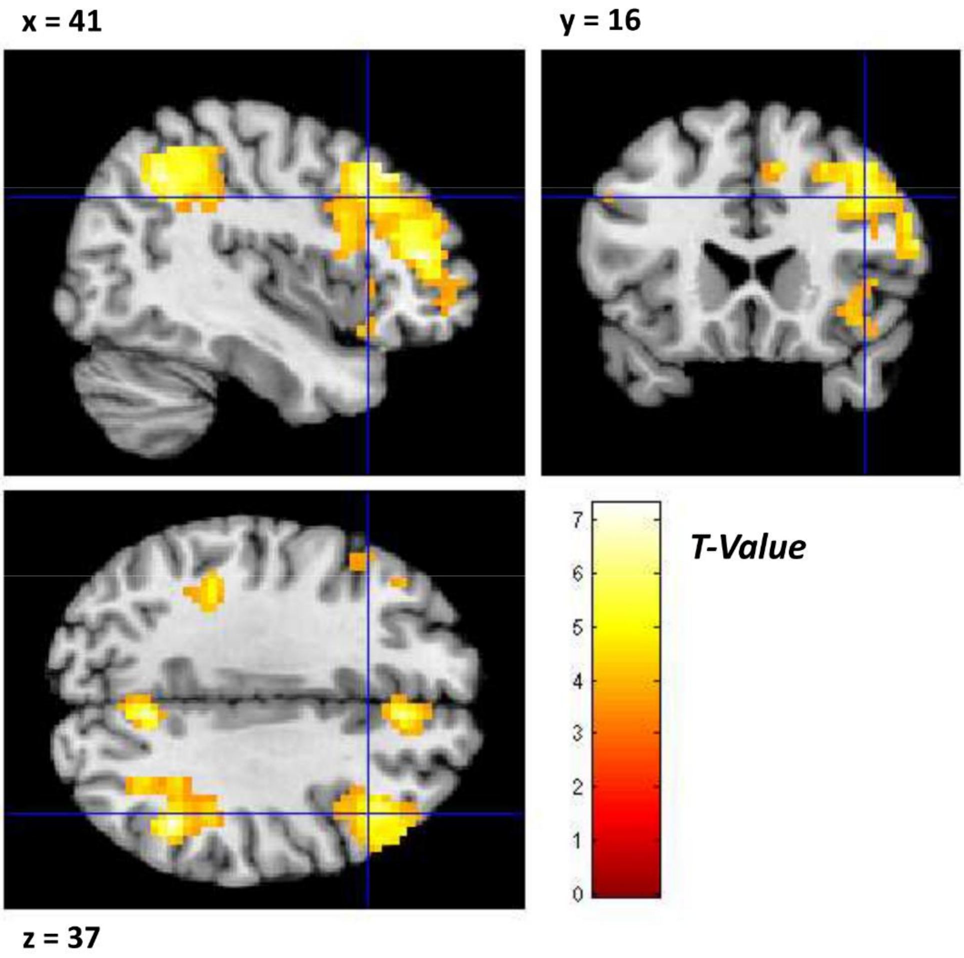

Fig. 3.10 Univariate activation for the salience contrast (display cutoff $p < 0.001$, $k > 20$ voxels).

Table I Activity clusters associated with the salience contrast

Region	L/R	k	x	y	z	T
Middle Frontal Gyrus	R	693	39	32	20	7.3
Precuneus	L/R	87	6	-64	40	7.22
Lingual	R	87	18	-97	-8	6.94
Fusiform Gyrus	L		33	-73	-18	
Inferior Parietal Lobe	R	402	45	-55	37	6.83
aMCC	L/R	140	6	29	37	6.24
Fusiform gyrus	L	86	-27	-85	-18	6.2
Superior Parietal Lobe	L	70	-24	-70	47	5.62
Middle Frontal Gyrus	L	65	-39	35	16	5.60
Anterior Insula	R	51	36	23	-8	5.08
Cerebellum	L/R	83	-6	-76	-25	5.02
PCC	L/R	43	-6	-28	30	4.87
Inferior frontal gyrus	L	45	-48	11	16	4.35

L: left hemisphere; R: right hemisphere; k: cluster size; xyz: Montreal Neurological Institute (MNI) coordinates.

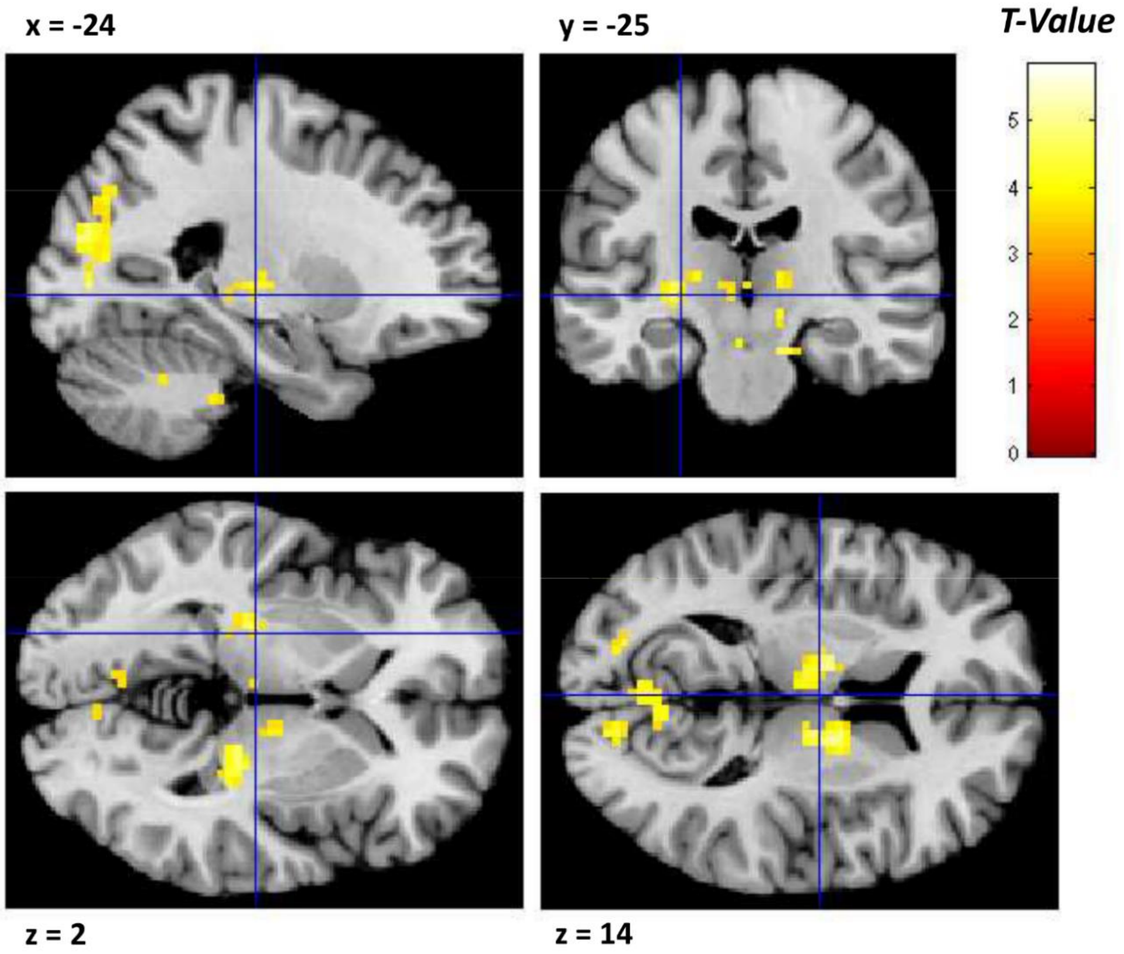

Fig. 3.11 Univariate activation for the valence contrast (display cutoff p < 0.001, k > 20 voxels).

Table II Activity clusters associated with the valence contrast

Region	L/R	k	x	y	z	T
Insula	L	64	-45	-4	44	5.83
Calcarine sulcus	L/R	362	-3	-76	-8	5.32
Cerebellum	R	40	21	-31	-25	5.19
Cerebellum	L	271	-21	-40	-32	5.17
Thalamus	R		15	-10	13	5.11
Thalamus	L		-12	-10	13	5.02
Cuneus	L	65	-21	-85	27	5.04
Precentral gyrus	R	47	45	-10	44	4.81

L: left hemisphere; R: right hemisphere; k: cluster size; xyz: Montreal Neurological Institute (MNI) coordinates.

The increase in activation for the valence contrast showed instead a completely different network, encompassing thalamus, posterior caudate nucleus, cuneus, calcarine sulcus and cerebellum (Table II and Fig. 3.11). For both patterns, no increase in activation survived threshold for the negative contrasts (punishment > reward, and neutral > salient).

Discussion

This study investigates how opposing motivational outcomes affect representation of real-world categories of objects in terms of category information in OSC. It extends previous results regarding the effects of reward in a similar task (Hickey & Peelen, 2015), and places these findings in a broader frame. A relevant issue has been raised in fact with respect to studies concerning reward by Maunsell (Maunsell, 2004). He considered the difference between effects due to actual reward and the ones due to strategic and deliberate allocation of attention. We therefore dissociated these two variables through the contemporary use of three motivational conditions. Two of them, gain and loss, constituted motivationally salient targets, although in oppositely valenced spaces. A third one, the neutral condition, presented a negligible strategic relevance, and a more blunt reward value. The salience pattern could therefore represent either the product of a general arousal effect, or a strategic approach to the task, in terms of a convenient allocation of attentional resources to different stimuli. The valence pattern tracks instead the rewarding, incentive value of the stimuli.

MVPA Analysis. We looked at how these different motivational conditions affected representation of visual objects in occipito-temporal cortex. We show that neural representation in OSC follows a valence, and not a salience, scheme. These results suggest that this portion of cortex encodes the motivational aspect of the attended categories not as a function of their strategic utility, but rather in terms of their affective valence. We also find that univariate activation of a ROI corresponding to SN, when analyzed through a valence contrast, predicts the change in the quality of this representation in occipito-temporal cortex. A similar result was obtained in another study considering the relationship between rewarding and neutral outcomes (Hickey & Peelen, 2015), but in this case also the presence of punishment is taken in consideration. This second set of data, therefore, on one hand gives further suggestion for a role played by dopamine in motivational control of visual representation, and secondly it specifies that this modulation through outcome follows a valence pattern. An important open question regards the mechanism through which dopamine release would be able to modulate cortical representation in occipito-temporal cortex. This area lacks in fact any input from the SN/VTA complex, showing the need for an intermediate actor in this dynamic.

Behavioral analysis. Behavioral results did not show any difference between conditions. At the same time, a more detailed analysis of these data revealed a pattern which looked similar to the one we found in our first behavioral study, with an increase in explicit accuracy for target present trials, and a lack of effect for target-absent trials (Chapter 2). Fewer subjects have been tested in this case, causing a decrease in the estimated power of this analysis from 0.999 of previous behavioral study (n = 101) to 0.721 for this one (n = 23). Yet, the final aim of this study was to look at the neural effects of this manipulation, rather than at the behavioral ones. The paradigm was also optimized accordingly.

Searchlight Analysis. Thirdly, we find another area, beyond occipito-temporal cortex, which discriminates the category of object that the subject is attending to, located in right PPC. This region has recently been found in another study to retain a representation of shape and object identity (Jeong & Xu, 2016). Moreover, motivational outcomes also in this case affect category representation, but, crucially, according to the latter of our a priori expectations, namely the salience pattern.

This study further corroborates, and may help explaining, the findings of a previous behavioral study, which considered how reward and punishment affected visual search on a behavioral level. We found that these two outcomes impacted accuracy measures according to a valence and not to a salience pattern, following an apparently paradoxical scheme. In this study, using an almost identical paradigm, we found that this is paralleled by an analogous pattern of representation of visual information in OSC as assessed by MVPA. At the same time, we also find an area of the brain which represents stimuli according to the second of our expectation, the salience pattern, suggesting a possible site for the alternative, rational approach to the motivational feedback. These complementary ways of processing payoffs may subsequently get integrated in order to provide a net, comprehensive account, and therefore yield a more sophisticated and articulate influence of motivation on perception.

Univariate Analysis. The whole-brain univariate analysis supports these considerations. Two distinct network of areas came out from this analysis, suggesting that the brain does actually approach different motivational outcomes according to these two alternative and opposing schemes.

The salience contrast, which subtended a strategic evaluation of the task, revealed the activation of fronto-parietal regions classically involved in the establishment of top-down attentional control (Corbetta & Shulman, 2002). This suggests that participants did not get a degraded visual representation of punished stimuli because they retained them as less important. The case was actually the opposite, as of what can be inferred through the analysis of the activation of these areas. Strategic top-down control was in fact actively deployed in the direction of giving higher priority to salient (both rewarding and punishing) stimuli with respect to neutral ones.

Another set of areas, encompassing primary visual cortex, posterior caudate nucleus and thalamus, followed instead the valence pattern, the scheme which we have also found to characterize behavioral measures and the amount of information in the occipito-temporal region. Importantly, all these areas constitute central components of the network of automatic visual attention (Kim & Hikosaka, 2015). Caudate tail (CDt) in particular has been demonstrated to present direct connections with the temporal lobe, and to be involved in visual implicit memory, such as for example category learning (Yamamoto et al., 2012; Seger, Peterson, Cincotta, Lopez-Paniagua, & Anderson, 2010). At the same time, this region has also been shown to be modulated by reward during tasks of visual attention (Yamamoto et al., 2013; Anderson et al., 2014; Anderson, 2016). Basal ganglia in general have been posited as an ideal neural substrate for reinforcement learning (Montague, Dayan, & Sejnowski, 1996). Therefore

CDt, which represents the component of striatum specifically involved in visual learning, could embody the structure encoding the implicit motivational value of visual stimuli. Due to its intermediate position between the SN/VTA complex and visual cortex, it could potentially represent the missing link in order to explain the putative influence of dopamine on the representation of conditioned visual stimuli.

A new hypothetical, additional attentional bias has recently been posited in order to account for phenomena which are not adequately explained by the traditional top-down/bottom-up dichotomy. Some of these apparently conflicting instances, most of which have been classified as examples of implicit memory, have been summarized and reunited under the common denominator of selection history (Awh et al., 2012). No structure has been suggested as a possible neural substrate for this control, but striatum could represent a reasonable candidate. This region is on one side centrally involved in the acquisition of implicit learning (Packard, Hirsh, & White, 1989; Knowlton, Mangels, & Squire, 1996), and on the other an important site for reinforcement learning. Moreover, it is the region which has been constantly found in studies which tried to disentangle the automatic component of reward from its voluntary one (Pessiglione et al., 2007; Krebs et al., 2012). An alternative source of endogenous control with respect to the classical fronto-parietal top-down bias would explain, among others, peculiar data from Rossi, where an extensive prefrontal cortex (PFC) lesion in macaques impeded subsequent endogenous selection only partially. More precisely, this deficit manifested as a function of the needed rate of task-switching, as if this was the only component of the task where the function of PFC proved to be truly essential (Rossi, Bichot, Desimone, & Ungerleider, 2007).

In this perspective, according to our data and also to a literature which reports a consistent association between these sets of regions and the two motivational schemes (Roesch and Olson, 2004; Kobayashi et al., 2006; Leathers and Olson, 2012; Engelmann et al., 2009; Litt et al., 2011), we also find that fronto-parietal regions (generally involved in the establishment of deliberate, strategic top-down selection) follow a salience scheme, while the activation of a series of subcortical structures, which we deem responsible for automatic selection history control, is shaped by a valence pattern. Why? Daw and colleagues address a similar issue at the level of action selection and decision making (Daw et al., 2005). They consider neural and behavioral data which reveal a dissociation between mechanisms which regulate on one side simple, automatic stimulus-response contingencies, and on the other side more complex, model-based patterns of behavior. The former system, relying on basal ganglia, would be generally less accurate, while the latter one, centered on prefrontal cortex, would instead have to bear with the disadvantage of a longer latency. These two systems could sometimes work synergistically, and other times be in conflict. It should therefore not surprise that the more rational, utility-based salience model is the one elaborated by the computationally more powerful neocortical system, while the faster, less sophisticated subcortical controller develops instead a less strategic account of the different motivational conditions. In this case, the striatum-based controller would not be able to disentangle the negative valence of the feedback that subjects received even in case of correct response (although less negative than in case of incorrect response) from the higher absolute value that detection of a punished category retained. Obviously, these are only very speculative hypotheses, which try to put a series of experimental data into a broader frame. Further experimental data will potentially be able to shed more light onto these questions.

Chapter 4

Neural representation of visual stimuli after inconsistent pairing with reward and punishment

Abstract

In addition to the traditionally conceived top-down and bottom-up controls, two separate sources of attention have been recently postulated, one prioritizing rewarding, appetitive stimuli and the other one threat-associated, aversive ones. Here, we investigate how the representation of visual categories in occipito-temporal object-selective cortex (OSC) as assessed by MVPA is modulated by different motivational outcomes (reward, loss and neutral outcome) during an inconsistent schedule of reinforcement. We find that two variables independently predicted the way in which positively and negatively valenced stimuli were prioritized with respect to each other across participants. On one side, the degree of responsivity to positive stimuli of Substantia Nigra, which has already been associated with the abovementioned reward-driven bias, predicted the quality of the neural representation according to a positively valenced scheme. On the other, an index of the individual sensitivity to punishment, BIS, was directly correlated with the prioritization of negatively valenced stimuli with respect to positive ones. These results support claims for the existence of these two new additional attentional controls, and also for their independent contributions to the bias of visual representation. When analyzing univariate contrasts, the salience contrast shows the activation of the central executive network (CEN) and of anterior insula. The valence contrasts shows instead the involvement of medial cortical regions and two areas of striatum, namely ventral striatum (VS) and putamen. When comparing data from this experiment and a former one where a consistent schedule of reinforcement was used, we find a dissociation between the subcortical networks activated in the valence contrast in the two conditions. While inconsistent pairing recruits ventral striatum and putamen, cue presentation after consistent conditioning triggers the differential activation of caudate tail, a region which encodes stable object values and is involved in the automatic control of visual selection.

Introduction

The release of midbrain dopamine is not static, but rather changes as a function of learning. While initially triggered by the presentation of a rewarding stimulus, this discharge shifts in time when the association with a conditioned stimulus (CS) is established, and then becomes elicited by the CS itself (Schultz, Dayan, & Montague, 1997). We have previously shown that, during a visual search task in naturalistic scenes, feedback modulates the neural representation of visual categories of objects, in terms of the amount of information as assessed by MVPA (Chapter 3). In this paradigm, each category of stimuli, irrespectively of correctness of response, cued a specific motivational outcome: a gain, a loss, or a neutral outcome. At the same time, correct response still garnered the optimal payoff for each category type, pushing therefore participants to be accurate in all conditions. Moreover, in terms of absolute value of the correct response, gain and loss were the two relevant conditions, with 100 points each per correct detection (greater reward in the appetitive condition, and evasion of a greater loss in the aversive one), while the same response only garnered a negligible amount of 2 points in the neutral condition. We had made the two alternative hypotheses that the two relevant outcomes would have affected neural representation in an opposing or in a similar manner, following respectively valence or value. In the former case this modulation would have been more closely linked to the positive, rewarding quality of the feedback. In the latter case, the increase in representation could have had two possible origins. On one side, it could have been the product of a non-specific, automatic arousal effect present for both relevant stimuli. Alternatively, it could have resulted from the strategic, rational evaluation of the utility of the various outcomes, leading to a foreseeable increase in attention and explicit motivation devoted to the two relevant categories.

Neural representation of the four visual categories was found to follow a valence pattern in the occipito-temporal object-selective cortex (OSC) of participants. Across subjects, the extent to which this representation was modulated according to a valence scheme was predicted by the degree to which univariate activation of a ROI corresponding to Substantia Nigra (SN) followed the same pattern, suggesting a role for dopaminergic structures in the establishment of this bias. In a second experiment, the paradigm was adapted in order to investigate this dynamic on a behavioral level, and a similar valence pattern was found to characterize target selection and the way in which different motivational categories captured attention when presented as irrelevant distractors. These results conflict with a rational utility model, according to which higher attentional deployment should be devoted to the two valuable conditions, gain and loss, with respect to the less relevant neutral one. Thence, these data in the first place give further support for the presence of an automatic, irrational bias toward rewarding stimuli, which has been by now well reported in the literature in a series of contexts (Anderson et al., 2011; Awh et al., 2012; Hickey & Peelen, 2015), and which is found to affect representation of these stimuli on a neural level as well. In the second place, they suggest that this irrational, reward-associated control is mainly driven by the overall positive valence of the stimuli, regardless of the absolute value of the outcome, underscoring once again its non-strategic nature.

Moreover, in spite of a vast literature reporting an attentional bias towards negative stimuli, we could not find evidence for such an effect. In this case we are not considering the loss condition as salient under a utilitarian point of view, because of its high absolute value in light of a rational evaluation. Instead, we behold it in terms of its threat-predictive characteristics, i.e. of its highly arousing nature and at the same time negative valence. Just specularly to the reward-associated control, a similarly irrational, task-unrelated gain in representation has been observed on a neural level for threat-associated stimuli, as for example in the case of emotional faces (Vuilleumier, 2001) and other types of fear-relevant stimuli (Sabatinelli, Bradley, Fitzsimmons, & Lang, 2005). A parallel bias has been highlighted also on a behavioral level, in terms of a tendency to prioritize task-irrelevant, negative stimuli with respect to neutral (Mogg and Bradley, 1998) or also to rewarding ones (Pratto & John, 1991; Baumeister et al., 2001). Importantly, though, this literature often reports a role played by personality traits in determining the extent of this bias, such that highly anxious individuals are the ones exhibiting a greater effect along this line, or in some cases the only ones showing it (Mac Leod et al., 1986). Intriguingly, we did find an inverse correlation between trait anxiety levels on a personality index score and the degree to which category information in OSC followed the inverse valence model. In other words, in spite of a general tendency toward a positive valence pattern, highly anxious individual did display a stronger representation of stimuli paired with loss with respect to the ones paired with reward.

In this study, we investigated whether the impact of feedback on stimulus representation could be appreciated also on a trial-by-trial basis, and moreover in the absence of any schedule consistency. Participants performed a similar detection task in naturalistic scenes, but in this case each trial could randomly present a rewarding, a punishing or a neutral outcome, and subject discovered the trial type only at the moment of feedback. The absence of a consistent schedule ruled out the possibility of any kind of associative learning across the different categories, and was meant to exclude the effects of specific outcomes' expectation on the neural representation of visual categories. Moreover, in this paradigm we introduced a completely task-irrelevant category, which could serve as a control for our analysis. This allowed us to rule out whether any of these correlations could be driven by an unspecific change in representation of all visual information depending on our variables of interest (SN activation and BIS score).

Materials and methods

Participants. Twenty-eight participants (12 female, mean age 23 ± 3 SD) took part in the experiment. All procedures were carried out in accordance with the Declaration of

Helsinki and were approved by the ethical committee of the University of Trento. Two participants were excluded from the analyses, as the OSC localizer did not yield any OSC region.

General structure of the experiment. The experiment consisted of three parts, in this order: 1) the visual search task, 2) the OSC localizer, 3) the category pattern localizer. Stimuli were presented using Psychtoolbox (Brainard, 1997) and projected on a translucent screen at the end of the scanner bore. Participants viewed the screen through a mirror mounted on the head coil. At the end of the MR session, each subject was administered a BIS/BAS scale questionnaire (Carver & White, 1994). Participants were paid based on the number of points accumulated during the experiment.

Visual search task.

Stimuli. The stimulus set for the visual search task consisted of black and white pictures (27° × 38° visual angle) of natural scenes (n=480) selected from an online database (Russell et al., 2008). Images were organized in six groups, each containing 80 images: three groups contained one single category in isolation, and the other three contained the possible combinations of any other possible pair of categories. For each participant, only four out of these six groups were used, depending on the specific relevant categories. Natural scene photographs were followed by perceptual masks of equal size, generated by combining white noise at different spatial frequencies and superimposing naturalistic structure on the noise (Walther et al., 2009).

Procedure. Participants performed 5 runs of 590 s duration, each containing 8 blocks of 60 s. Each run began and ended with 15 sec fixation. Before each block started, a fixation cross was presented for 6 s, and then central text for 4 s, reminding response contingencies and indicating the score which had been obtained up to that moment. From a group of 3 categories (people, cars and trees), each subject was assigned, in a counterbalanced order, one target1 (T1) and one target2 (T2) category, while the third category (T3) was never mentioned to the subject. Participants had to indicate for each trial whether T1 or T2 was present in the scene, through right index finger and right middle finger response, respectively. Each block comprised 24 trials, 12 T1-present and 12 T2-present, in a randomized order, so that one and only one of the two targets were displayed in each trial. For both targets, 6 trials contained the category by itself (T1 and T2), and 6 trials contained it alongside with T3 (T1T3 and T2T3). The total amount of trials was 960 per subjects, so that each scene was shown three times. Each trial started with a fixation cross (833 ms) followed by the brief presentation of the scene (58 ms), a mask (325 ms), the reappearance of fixation (750 ms), and the feedback (533 ms; Fig. 4.2A). Responses given outside the 750 ms fixation interval, or absence of response, garnered an incorrect performance.

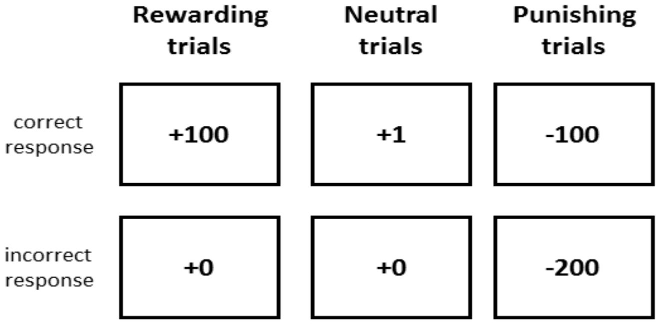

	Rewarding trials	Neutral trials	Punishing trials
correct response	+100	+1	-100
incorrect response	+0	+0	-200

Fig. 4.1 Feedback schedule during the search task. Each trial of the block was randomly assigned to one of the three different types of outcome. Participants discovered this assignment at the end of the trial, at the moment of feedback.

Each block contained 8 rewarding, 8 punishing and 8 neutral trials (4 T1-present and 4 T2-present – both targets in half of the cases by themselves and in half of the cases alongside with T3 - for each outcome type), according to a random schedule, so that participants discovered the trial type only at the moment of feedback. Correct response always garnered the highest possible payoff relatively to the trial type, which corresponded to 100 points for rewarding, 1 point for neutral, and -100 for punishing trials. Incorrect responses resulted instead in 0 points for rewarding and neutral trials, and -200 points for punishing trials (Fig. 4.1). Net value of correct response was therefore 100 points for rewarding and punishing trials, and 1 point for neutral trials.

OSC localizer.

Stimuli. The stimulus set for the OSC Localizer consisted of black and white pictures (27° × 38° visual angle) of 20 daily-life objects (e.g. telephone, cheese, alarm clock) and their scrambled version.

Procedure. Participants performed one OSC localizer run of 467 s duration, containing 24 blocks of 20 trials and 5 fixation blocks. The run began and ended with 15 sec fixation. Before the first block of the run, a 2 s display reminded the task to the subject. During each 15 s block, pictures of one type only (intact or scrambled) were presented, while subjects monitored for image repetition, which occurred once in a block. Each trial started with a fixation cross (350 ms) followed by the picture of the intact or scrambled object (400 ms). Every fifth block was a 15 s fixation block with no stimuli.

Category pattern localizer.

Stimuli. The stimulus set for the category localizer consisted of four groups of stimuli, one for each of the categories cued in the visual search task (people, cars, houses and trees). Each group consisted of 40 black and white pictures (27° × 38° visual angle) of isolated exemplars of the same category on a white background. Pictures of people were headless, as faces in the search task scenes were generally too small to be visually solved.

Procedure. Participants performed one category localizer run of 497 s duration, containing 24 blocks of 20 trials and 7 fixation blocks. The run began and ended with 15 sec fixation. Before the first block of the run, a 2 s display reminded the task to the subject. During each 15 s block, pictures of different exemplars of only one of the four categories were presented, while subjects monitored for image repetition, which occurred once in a block. Each trial started with a fixation cross (350 ms) followed by the picture of the intact or scrambled object (400 ms). Every fifth block was a 15 s fixation block with no stimuli.

fMRI Data Acquisition.

Imaging was conducted on a Brucker BioSpin MedSpec 4T head scanner (Bruker BioSpin), equipped with an eight-channel head coil. T2*-weighted gradient-echo echo-planar images were collected as functional volumes for all the functional runs (EPI; repetition time = 2 s, echo time = 28 ms, flip angle = 73°, 3 × 3 × 3 mm voxel size, 0.99 mm gap, 30 slices, 192 mm field of view, 64 × 64 matrix size). A T1-weighted image

(MPRAGE; 1 × 1 × 1 mm voxel size; 256 slices, 224 × 176 matrix size) was obtained as a high-resolution anatomical reference.

fMRI Preprocessing.

All neuroimaging data were analyzed using MATLAB and SPM12. The volumes were motion corrected, slice time corrected, coregistered to the structural image and spatially normalized to the Montreal Neurological Institute (MNI) template (as included in SPM12). Functional volumes were then smoothed using a 6-mm full-width half-maximum Gaussian kernel. All analyses were performed on the smoothed data.

fMRI Data Analysis.

A separate general linear model was created for each of the three phases. The blood-oxygen-level-dependent signal of each voxel in each participant was modeled using two regressors of interest in the OSC localizer, one for intact and one for scrambled objects. Three regressors of interest representing the various categories (people, cars and trees) were used in the category localizer. Twelve regressors of interest were used for the visual search task, corresponding to the 4 combinations of categories present in the scene (T1, T1T3, T2, T2T3), each of them in combination with the three different outcome types (rewarding, neutral or punishing). In all models, six additional regressors of no interest obtained from the realignment procedure were included to account for head motion. Finally, all models also presented an intrinsic temporal high-pass filter of 1/128 HZ to correct for slow scanner drifts.

OSC ROI definition.

OSC was functionally defined in each participant by contrasting, within temporo-occipital regions, responses evoked in the intact and in the scrambled condition ($p < 0.001$, uncorrected; voxel extent threshold $k > 20$). Mean OSC size was 686 voxels ± 682 SD. OSC was created through the MarsBar toolbox (Brett et al, 2002).

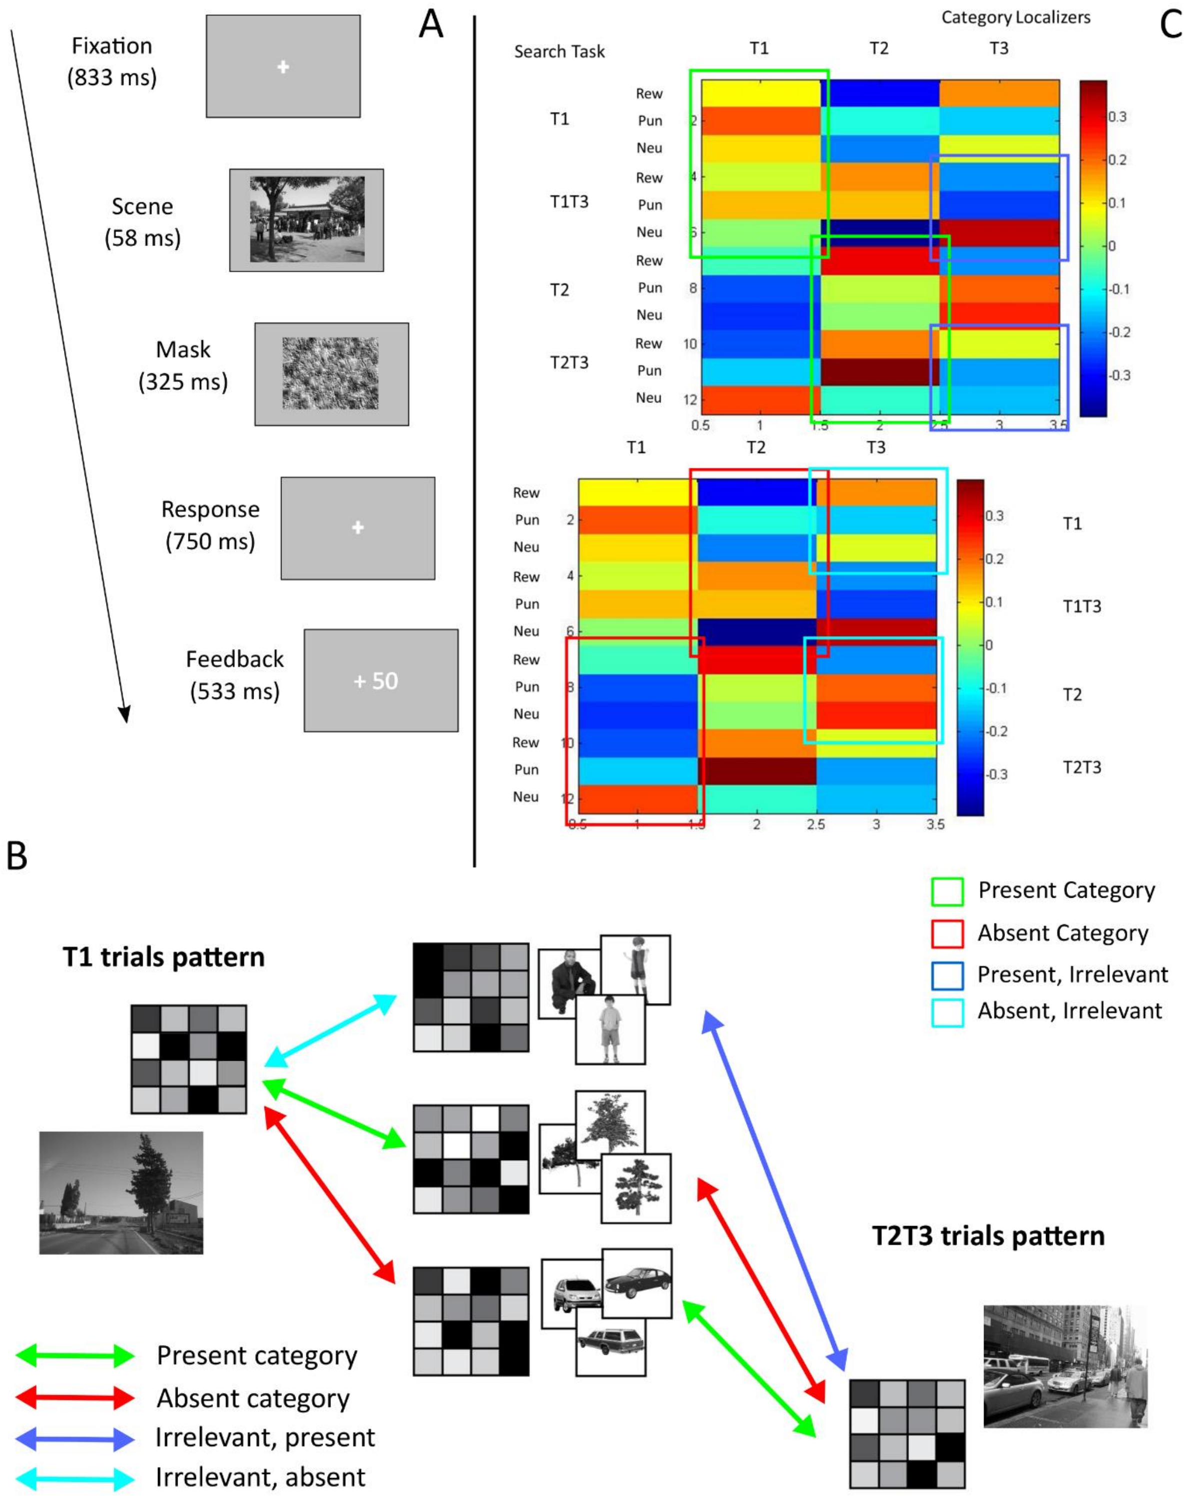

Fig. 4.2 (A) Trial structure. Participants reported which of the two cued categories was present in each scene. (B) Data Analysis. Patterns from the different trial conditions were correlated with category localizers' patterns. For each scene, one task-relevant category was present, one task-relevant category was absent, and one category unbeknownst to the participant could have been present or absent. (C) Correlation values were then organized according to outcome, category presence and task-relevance. In the example shown the participant had to look for trees and cars, while people was the irrelevant (and never mentioned) category.

MVPA.

All pattern analyses were performed using CoSMo-MVPA (Oosterhof et al., 2016). For each subject, three patterns (one for each object category) were obtained from the category localizer, and twelve patterns (one for each trial type) from the visual search task. This was done by extracting t values of each voxel for each condition, only for voxels comprised in the ROI of each subject. Following existing work, t values were normalized by subtracting, for each voxel, the mean value obtained for that voxel across all conditions of the single task. Normalized t values of all three conditions from the category localizer were then correlated with all twelve conditions of the visual search task, across all voxels of the ROI (Fig. 4.2B). This yielded a 3 × 12 correlation matrix, whose values were then Fisher-transformed and organized in terms of whether the category was present or absent in the scene (Fig. 4.2C). Values were averaged across target type for task-relevant categories (T1 and T2), and across scene type for the task irrelevant-category (T1T3 and T2T3). This yielded a 3 (feedback type) × 2 (present vs absent) × 2 (task-relevant vs task-irrelevant) matrix.

Statistical Analysis.

Differences between conditions across subjects were then tested using a three-way repeated-measures ANOVA (RANOVA), with task relevance (T1-T2 vs T3), category presence (T1 vs T2) and motivational condition (rewarding, punishing or neutral) as factors. Subsequently, a repeated-measures two-way ANOVA, with category presence (present vs absent) and motivational condition (rewarding, punishing or neutral) was performed separately on data regarding task-relevant and irrelevant categories.

Correlation analysis.

The three MVPA values representing the different outcome conditions (reward, neutral, loss) of each category type (relevant-present, relevant-absent, irrelevant-present) were modeled according to two predictors: one accounting for a valence-based model [+1 0 -1], and a second one accounting for a salience-based one [+1 -2 +1] (the values for each vector were standardized).

For each subject, a single t value was calculated for a valence [+1 0 -1], and another one for a salience [+1 -2 +1] contrast of the mean univariate activity during the visual search task, for a ROI corresponding to Substantia Nigra. The ROI was anatomically defined through WFU PickAtlas, a software which is based on the Talairach Daemon database (Maldjian et al., 2003).

Individual t values of both contrasts were then correlated (Spearman's rank-order correlation) with the individual beta coefficients of the multiple linear regression of OSC category information modulation of the corresponding model. This correlation was calculated for each of the three category types. A similar correlation was performed between participants' BIS scores and the individual beta coefficients of the multiple linear regression of OSC category information. Finally, a multiple linear regression of individual category information values was performed using both measures as predictors.

Univariate analysis.

The regressors of the target-present trials of the three motivational categories [reward neutral punishment] were contrasted according to the valence [+1 0 -1] and to the salience [+1 -2 +1] pattern. Individual participants' contrast images entered a second level t-test (threshold set at $P < 0.001$, cluster corrected for multiple-comparisons, voxel extent threshold $k > 20$).

Striatum ROI analysis.

ROIs for the three different areas of striatum coming out in the whole-brain valence contrast from the first (CDt) and the second (VS and Putamen) paradigm were created using the MarsBar toolbox (Fig. 4.9A). For each subject, T-values from all the voxels were averaged for each ROI, and the individual values of this 3 (regions) × 2 (paradigms) matrix underwent a mixed two-way ANOVA.

Connectivity analysis.

FC analyses were carried out using the CONN-fMRI functional connectivity toolbox v17 (Whitfield-Gabrieli & Nieto-Castanon, 2012). This software applies linear regression and band-pass filtering in order to remove unwanted motion, physiological, and other artifactual effects from the BOLD signal before computing connectivity measures. We considered three different sources of possible confounders: 1) BOLD signal from the white matter and CSF masks (by the use of the aCompCor strategy); 2) realignment parameters; and 3) the main condition effects (condition blocks convolved with hrf). As suggested, a frequency window of 0.01 to 0.1 Hz was used for band-pass filtering. Resting state FC (rsFC) was evaluated using functional scans from the first study (so connectivity across all scans, irrespectively of task conditions). The seed-to-voxel FC analysis was performed using each one of three regions from the ROI analysis (VS, putamen and CDt) as seeds. Seed-to-voxel FC maps were created for each participant, and then individual seed-to-voxel maps were entered into a second-level analysis. The threshold for significance was set to $p < 0.05$ whole brain cluster level FDR corrected with a cluster building threshold of $p < 0.001$ uncorrected on voxel level.

Results

Behavioral Analysis. Mean accuracy values (and their standard deviation (SD)) were 0.8662 (±0.08) for rewarding trials, 0.8717 (±0.07) for neutral trials, and 0.8673 (±0.06) for punishing trials. RT values were 588 ms (±51) for rewarding, 587 ms (±50) for neutral, and 589 ms (±52) for punishing trials. A one-way RANOVA showed the absence of any main effect of motivational condition both for accuracy ($F_{(2,25)} = 0.621$, $p = 0.541$) and for RT values ($F_{(2,25)} = 0.510$, $p = 0.603$).

MVPA Analysis. The amount of information for the various categories in the scenes was organized as a function of task-relevance, presence and motivational condition (Fig. 4.3). First we performed a three-way RANOVA, which showed a main effect of category presence, and an interaction between task-relevance and category presence (category presence, $F_{(1,25)} = 9.831$, $p = 0.004$, task-relevance × presence, $F_{(1,25)} = 4.761$, $p = 0.039$).

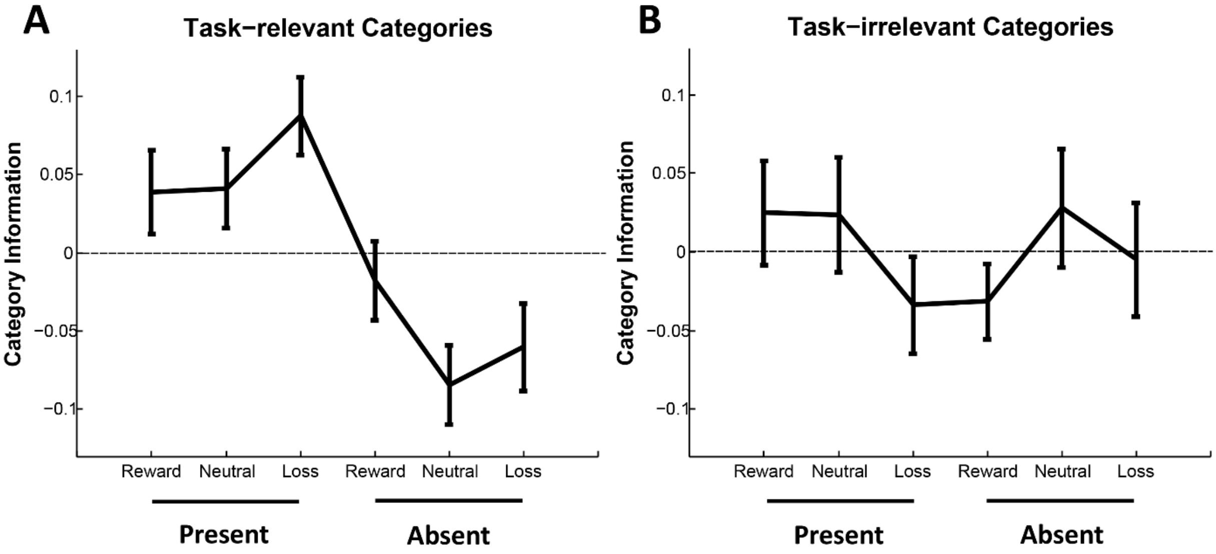

Fig. 4.3 (A) Amount of category information in OSC for task-relevant categories (T1 and T2), whether present or absent in the scene, and also as a function of the motivational outcome received after response; (B) same for task-irrelevant category (T3).

This allowed us to separate measures for task-irrelevant categories (which were never mentioned to participants) and task-relevant categories, and to look at the effect of category presence and motivational outcome for each group independently. A second two-way RANOVA conducted on task-relevant categories showed a main effect of category presence and no main effect of motivational outcome nor an interaction ($F_{(1,25)} = 12.504$, $p = 0.002$). Another two-way RANOVA conducted instead on task-irrelevant categories did not show any effect of category presence or motivational outcome (all F values < 1). This suggests that while the presence of T1 and T2 modulated category information the presence of T3 was fundamentally negligible in terms of modulation of visual representation.

Correlation analysis. Subsequently, we tested whether the change in the amount of category information across subjects could be predicted by the pattern of the univariate activation of SN, as it was the case for our previous study (Chapter 3). We calculated individual t values for the valence and for the utility contrast of the activation of this ROI and correlated them with the respective model coefficients from multivariate analysis of information in OSC. Again, we found that the t values of the valence contrast of univariate activation of the SN ROI predicted the change in category information in OSC according to the valence model (ρ = 0.399, p = 0.044; Fig. 4.4), while the same correlation did not yield any significant result for the utility model (ρ = 0.097, p = 0.635).

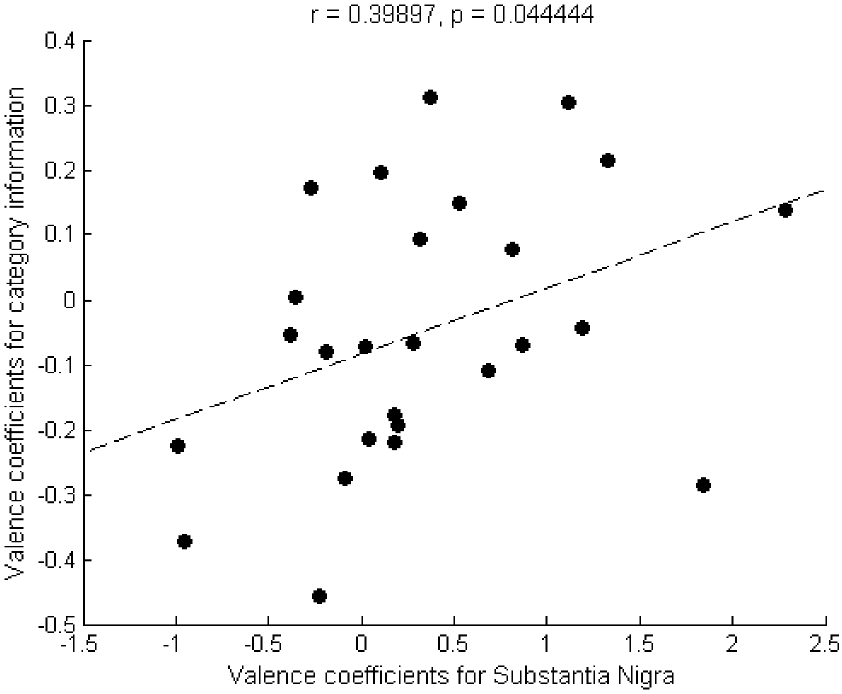

Fig. 4.4 Correlation between univariate activation of SN according to the valence contrast and valence coefficients for category information in OSC.

Moreover, in this paradigm we were able to test whether this correlation was specific to the present task-relevant category, or also extended to the modulation of the amount of category information for the other two categories of the study, the absent task-relevant one (T1 or T2), or T3 (when present). No other category beyond the task-relevant present one showed any correlation with the activation of SN, considering both possible models (correlation of OSC information regressed through the valence model with SN activation was ρ = 0.216, p = 0.287 for the task-relevant absent category, and ρ = -0.015, p = 0.941 for T3 (present); correlation of the utility coefficients was ρ = -0.128, p = 0.531 for the task-relevant absent category, and ρ = -0.105, p = 0.608 for T3 (present)).

Then, we looked at the correlation between measures of personality as assessed through the BIS/BAS questionnaire and the amount of category information in OSC. Like in previous study, we found an inverse relationship between the BIS score of participants (an index of the individual response to punishment-related cues) and the valence coefficients from the multivariate analysis of information in OSC (ρ = -0.501, p = 0.009; Fig. 4.5).

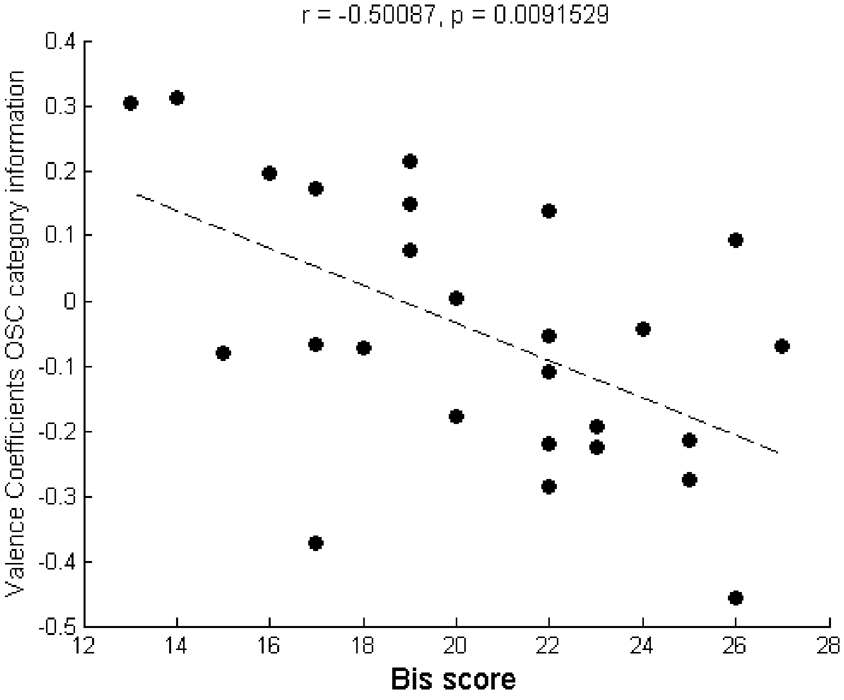

Fig. 4.5 Correlation between BIS score of participants and valence coefficients for category information in OSC.

In order to look at the independency between these two variables and at their reciprocal interaction in the modulation of category information in OSC, we subsequently performed a multiple linear regression of the individual values of this measure, using the participants' BIS scores and valence contrasts of SN activation as predictors. Both coefficient regressors showed to significantly predict category information (SN activation: $t = 2.695$; $p = 0.013$; BIS: $t = -3.690$; $p = 0.001$). Adjusted R-squared showed an increase when using both predictors together ($R^2_{adj} = 0.416$), with respect to regressions which only used one of them (SN: $R^2_{adj} = 0.110$; BIS: $R^2_{adj} = 0.264$).

For illustration purposes, we show the amount of conditional category information, after dividing the general population in four subgroups, obtained performing a median split within both abovementioned variables (Fig. 4.6). The impact of motivational outcomes on category information was clearly detectable in the two groups which presented the effects going in the same directionality for both variables (e.g. High BIS/Low SN activation, and Low BIS/High SN activation). These two groups showed a significant difference between the rewarded and the punished category (High BIS/Low SN activation: $t = -4.538$; $p = 0.006$; Low BIS/High SN activation: $t = 2.785$; $p = 0.039$), obviously in opposite directions.

Univariate Analysis. Analysis of univariate activation revealed two separate networks representing salience and valence, as in our previous study (Chapter 3). The salience contrast showed bilateral activation of anterior insula (AI), inferior occipital gyrus (IOG), and a right-lateralized fronto-parietal network (Fig. 4.7 and table I). The valence contrast showed instead an increase in activation in bilateral VS, in a big cluster spanning throughout the whole medial cortical surface, encompassing calcarine cortex, precuneus, orbitomedial prefrontal cortex (omPFC), anterior (ACC) and posterior cingulate cortex (PCC), and finally right putamen, right middle frontal and precentral gyrus (Fig. 4.8 and table II). Thus, just like in previous study, motivational modulation of striatal regions, which we have hypothesized to subtend selection history control (Chapter 3), was found to follow a valence scheme, stressing the strongly automatic and potentially irrational nature of this bias. At the same time, different areas of striatum showed increased activation in the two studies, i.e. caudate tail (CDt) in the former and VS and putamen in the latter. The different striatal region receive input from different parts of the cortex, and take part into partially independent corticostriatal loops: the motivational (which depends on VS), the executive (caudate head), the motor (putamen) and the visual loop (caudate body and CDt) (Seger, 2008). While the first loop has shown to display mainly feedback-related activity, the motor and the visual loop are thought to represent learning-related regions, which store specific motor response-outcome (motor loop) and stimulus-outcome (visual loop) associations. Consequently, we performed a ROI analysis on these regions, which confirmed their alternative preferential recruitment, depending on whether a consistent or inconsistent schedule or reinforcement was applied to visual stimuli.

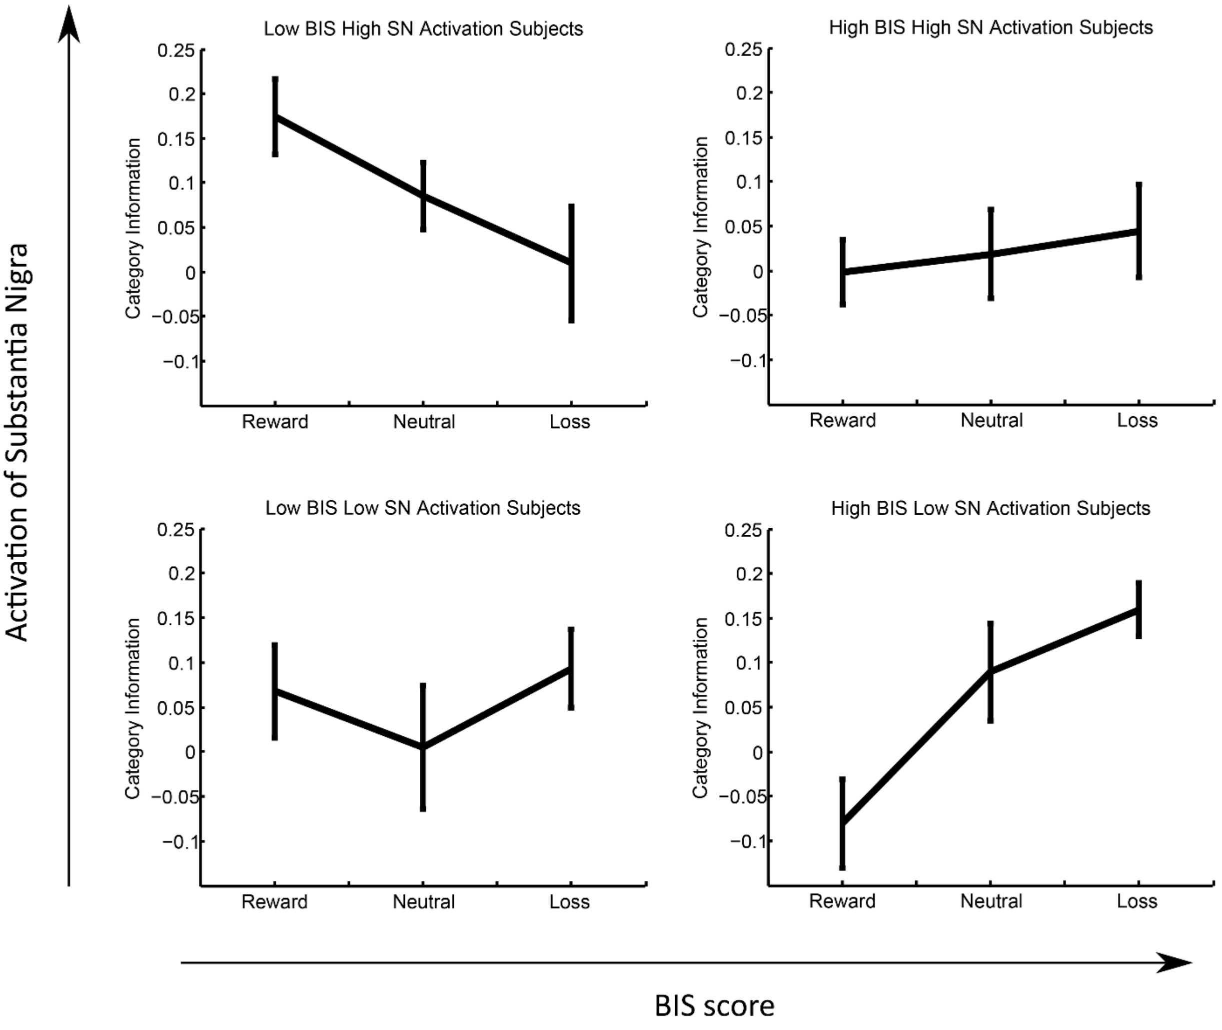

Fig. 4.6 Category information results for the four groups coming out of the median splits within both variables of interest. On the horizontal dimension, we see the effects of increasing values of BIS scores, and on the vertical dimension the effects of increasing valence coefficients for activation of SN.

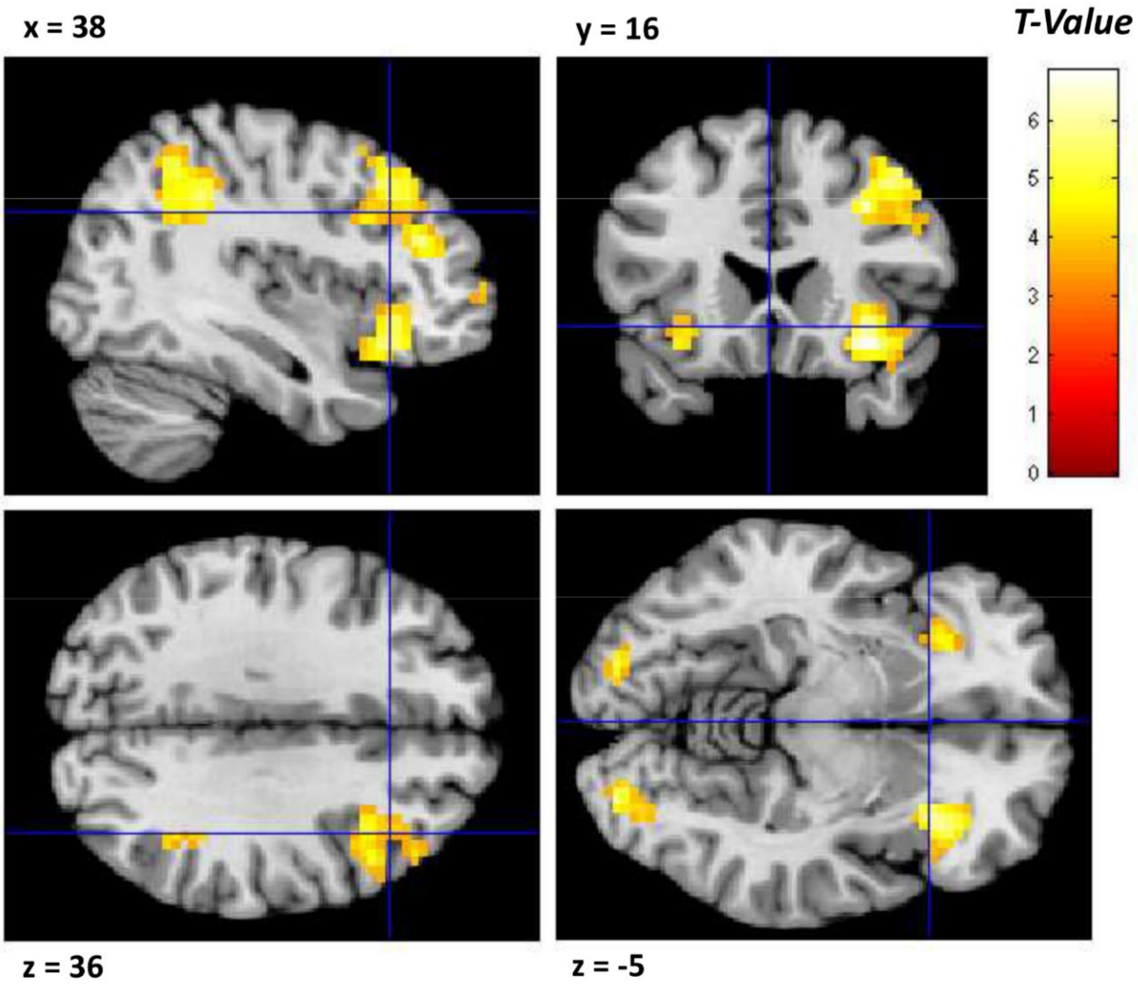

Fig. 4.7 Univariate activation for the salience contrast (display cutoff p < 0.001, k > 20 voxels).

Table I Activity clusters associated with the salience contrast

Region	L/R	k	x	y	z	T
Salient>Neutral						
Inferior Occipital Gyrus	L	34	-21	-94	-10	6.83
Anterior Insula	R	125	33	17	-10	6.69
Middle Frontal Gyrus	R	249	42	20	46	6.25
Inferior Parietal Lobe	R	188	48	-52	54	6.01
Inferior Occipital Gyrus	R	58	27	-94	-6	5.72
Anterior Insula	L	31	-30	20	-6	4.80
Neutral > salient (not shown)						
Postcentral Gyrus	R	186	66	-13	18	6.06
Inferior Temporal Gyrus	L	91	-45	-73	-2	5.35
Cuneus	R	50	15	-85	18	5.21

L: left hemisphere; R: right hemisphere; k: cluster size; xyz: Montreal Neurological Institute (MNI) coordinates.

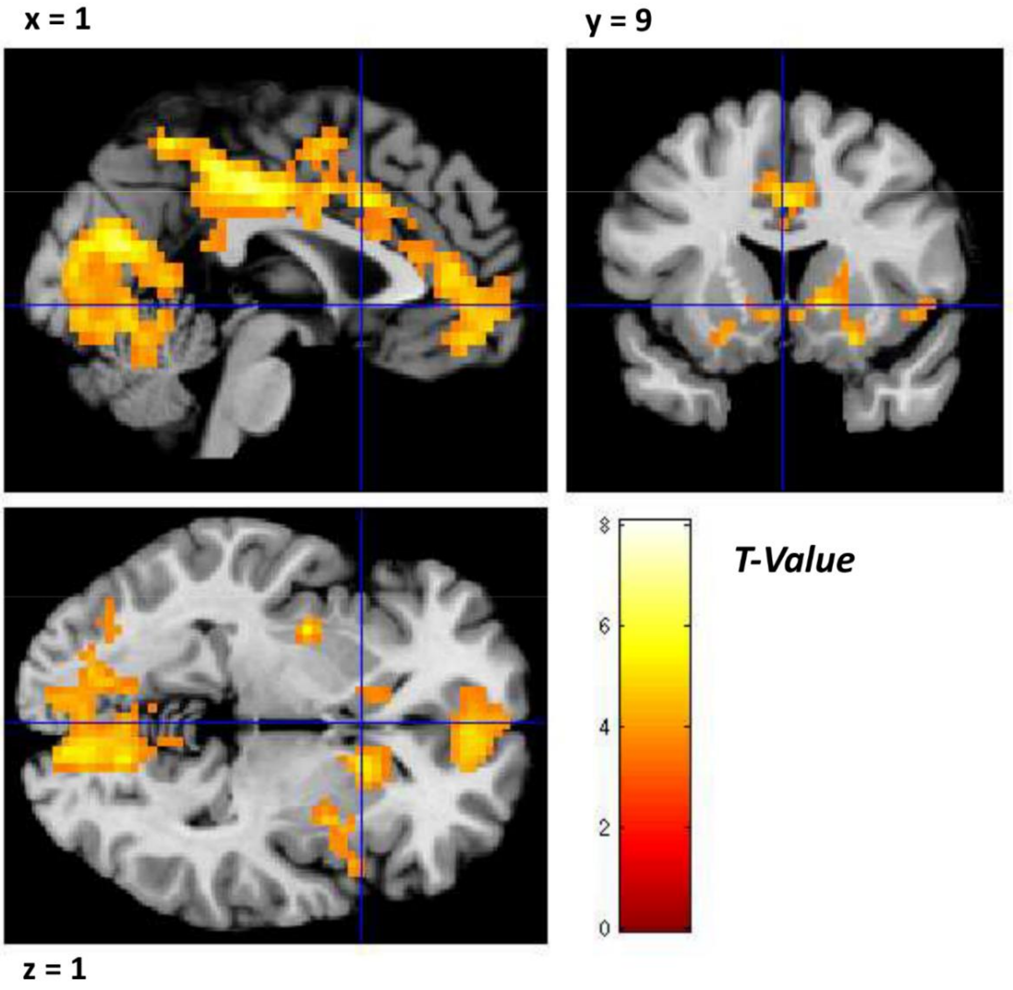

Fig. 4.8 Univariate activation for the valence contrast (display cutoff $p < 0.001$, $k > 20$ voxels).

Table II Activity clusters associated with the valence contrast

Region	L/R	k	x	y	z	T
Reward>Punishment						
Precuneus	L/R	2800	3	-43	50	8.06
Middle frontal gyrus	R	68	27	35	50	6.05
Precentral gyrus	R	50	39	-13	54	5.74
Ventral striatum	R	73	15	11	2	5.20
Superior frontal gyrus	L	78	-24	29	46	4.72
Putamen	R	43	33	-4	2	4.64
Ventral striatum	L	50	-9	-14	-2	4.29
Punishment > Reward						
None						

L: left hemisphere; R: right hemisphere; k: cluster size; xyz: Montreal Neurological Institute (MNI) coordinates.

ROI Analysis. In order to evaluate more accurately the function of striatum in the two different paradigms, we performed, for both conditions, a ROI analysis on the three regions which showed differential activation in the valence contrast of the two studies

(CDt in the former and VS and right putamen in the latter) (Fig. 4.9A). Results are shown in Fig. 4.9B for data from current study and in Fig. 4.10C for data from previous experiment. A mixed two-way ANOVA showed an interaction between striatum region and paradigm ($F = 7.661$, $p < 10^{-4}$).

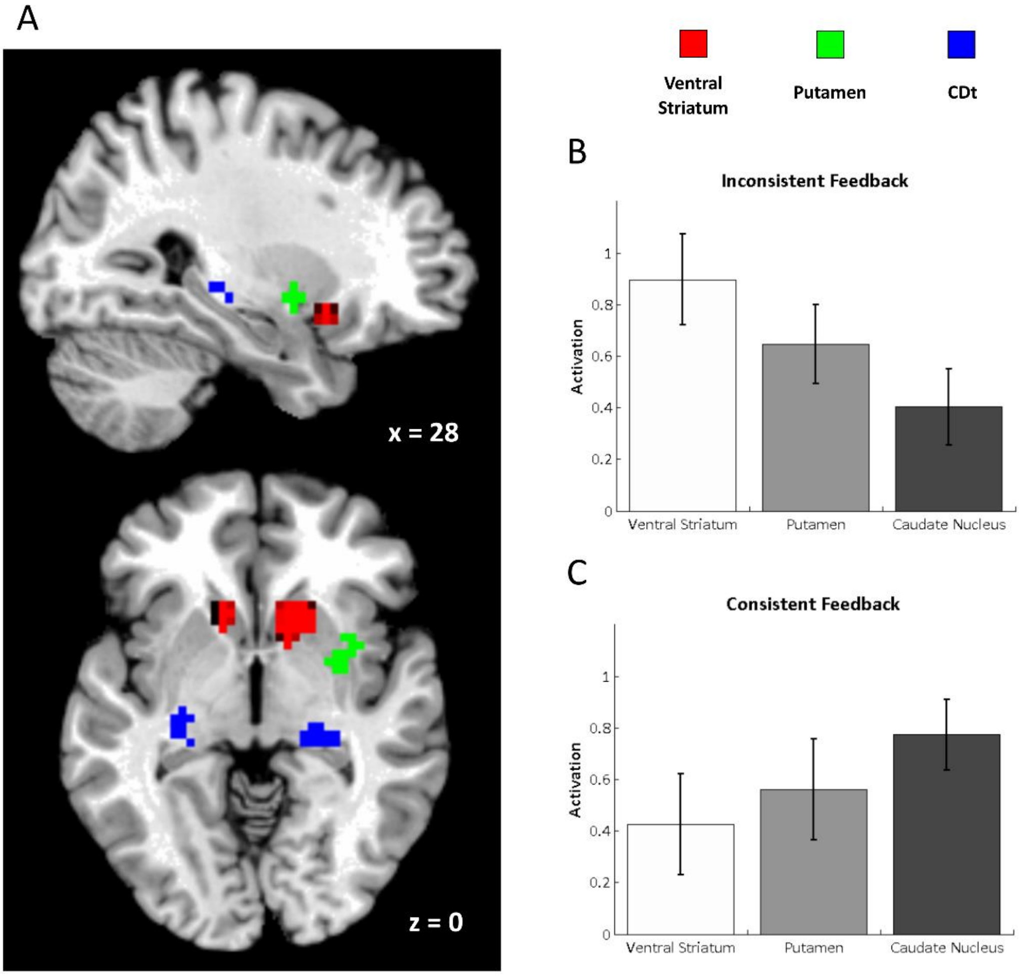

Fig 4.9 (A) ROIs representing VS, right putamen and CDt. (B) Univariate activation for the three ROIs in the paradigm following inconsistent motivational schedule. (C) Univariate activation for the three ROIs in the paradigm following consistent motivational schedule.

Functional Connectivity Analysis. We had a particular concern in ascertaining our CDt identification, given the small dimensions of this area and its relatively recent investigation in fMRI studies, which renders hard to define its localization in a rigorous way. To this aim, we evaluated resting state functional connectivity (FC) of the three striatum regions used in the ROI analysis, and checked whether they showed the expected patterns of connectivity. We used functional data from the first paradigm as our main interest lay on confirmation of the identification of the ROI which showed its activation in that study. All three ROIs showed the expected patterns of FC.

VS was preferentially connected to brainstem, bilateral frontal orbital cortex, thalamus, insular cortex (IC) and inferior frontal gyrus (Fig. 4.10 and Table III for a full list).

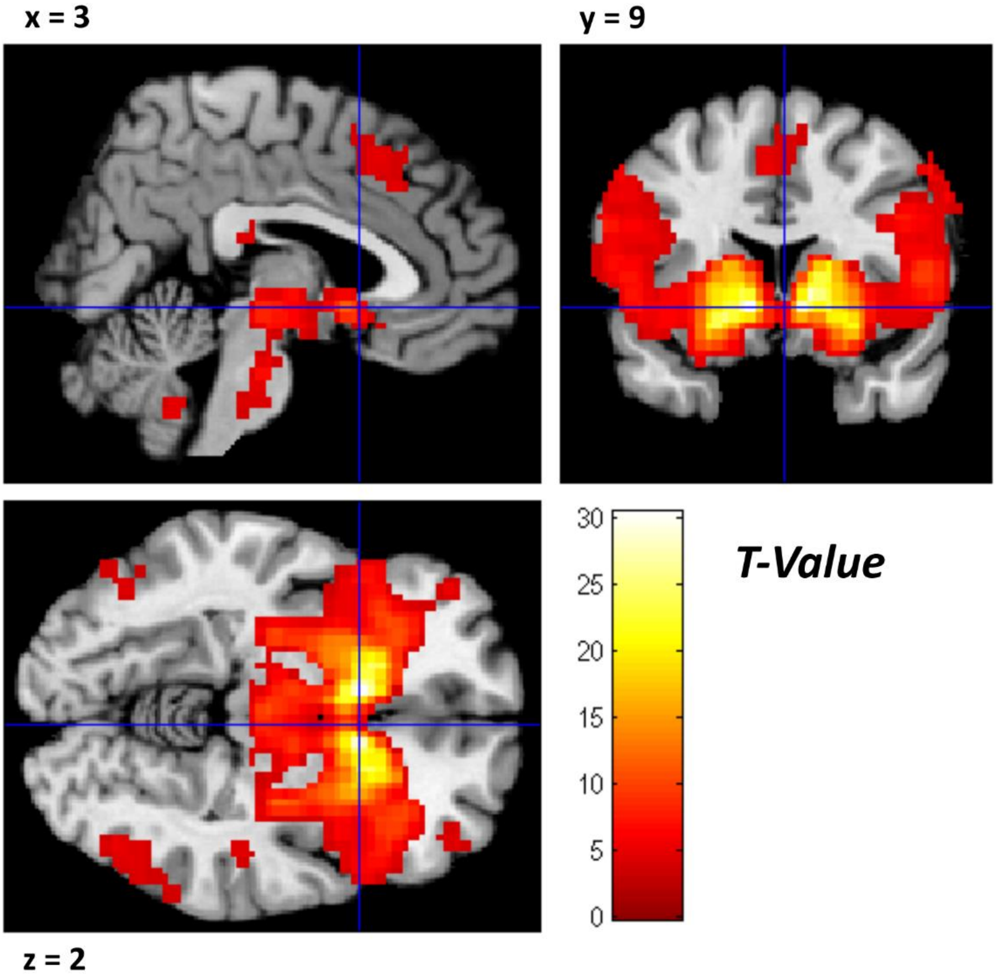

Fig. 4.10 Regions showing increased connectivity with the Ventral Striatum ROI (display cutoff $p < 0.001$, $k > 50$ voxels).

The right Putamen ROI displayed a preferential connectivity with brainstem, bilateral IC, central opercular cortex, Pre- and Postcentral gyrus, thalamus, supramarginal gyrus, planum temporale, parietal operculum, planum polare, Heschl’s gyrus and supplementary motor cortex (Fig 4.11 and table IV for a full list). Interestingly, this unilateral seed showed a perfectly bilateral connectivity. Finally, and crucially, the CDt

ROI showed a preferential connectivity with cuneus, precuneus, PCC, brainstem, and with bilateral LOC, lingual gyrus, fusiform gyrus, calcarine cortex, occipital pole, thalamus, IC, and left inferior temporal gyrus (Fig. 4.12 and Table V for a full list).

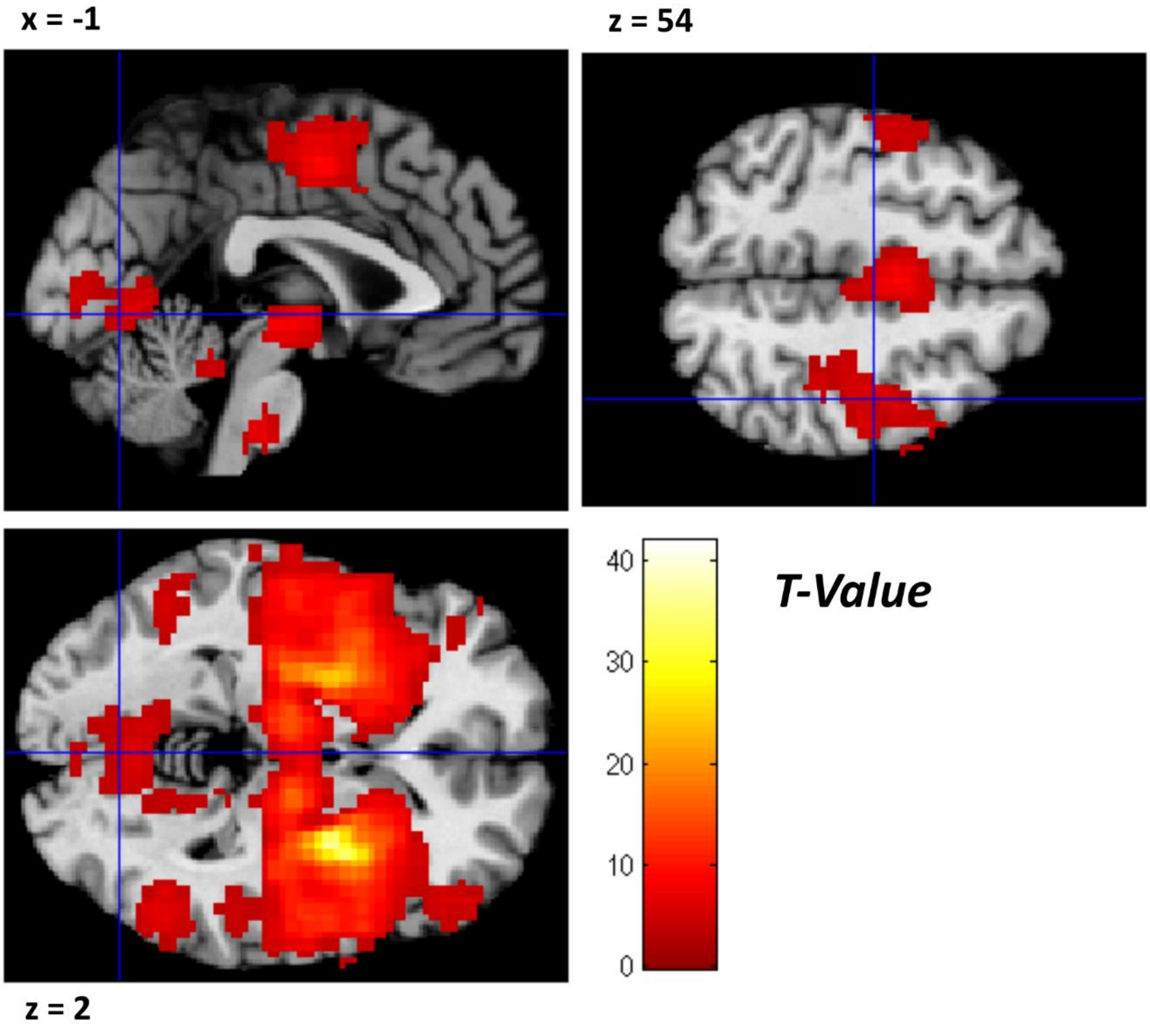

Fig. 4.11 Regions showing increased connectivity with the right Putamen ROI (display cutoff $p < 0.001$, $k > 50$ voxels).

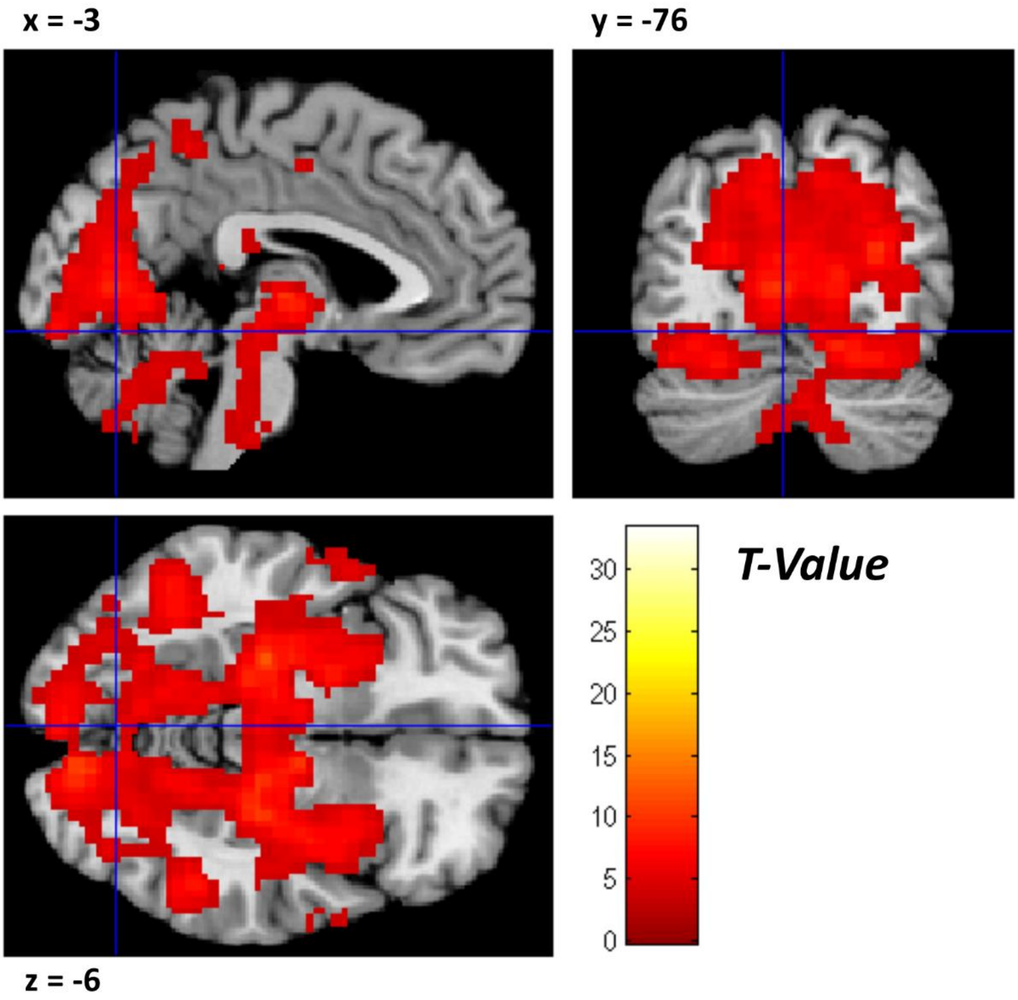

Fig. 4.12 Regions showing increased connectivity with the Caudate Tail ROI (display cutoff $p < 0.001$, $k > 50$ voxels).

Discussion

In this study we investigated whether unexpected and inconsistent motivational feedback is able to affect the trial-by-trial representation of visual stimuli, when assessed through MVPA. Such a modulation would shed light onto the mechanisms through which perceptual conditioning of stimuli is instantiated. Nevertheless, when all subjects were pooled together, we could not observe a differential effect of feedback on the amount of category information for the different motivational categories. At the same time, we did find once again that, across subjects, the modulation exerted by outcome on this visual representation was strongly influenced on one side by univariate activation of Substantia Nigra (SN), and on the other side by trait anxiety levels of participants.

On one hand, subjects whose SN followed a valence pattern more strongly, also showed a stronger representation of stimuli according to a positive valence scheme. This strengthens the idea that activation of dopaminergic areas is in some way related to the change in representation of visual information. In a previous study we formulated a hypothesis about how this could be implemented, in spite of the absence of any direct connection between visual cortex and SN (Chapter 3). Moreover, in this study we were also able to test the specificity of this correlation with respect to various stimuli present (or not) in the scene. The correlation of the activation of SN with the amount of category information was found to be restricted to the present task-relevant category. This is important when considering models of reinforcement learning which address how reward (and in this case also punishment) is able to distinctly modulate representation of stimuli which play different roles in the same perceptual context (Roelfsema, van Ooyen, & Watanabe, 2010).

On the other hand, the higher the BIS scores (an index of the individual response to punishment-related cues) of participants, the more they represented categories according to an inverse valence pattern, with an increase in representation for stimuli paired with loss and a decrease in representation for stimuli paired with reward. These two opposite influences cancelled out each other reciprocally across the population as a whole, but were clearly detectable when considering median splits of participants across these two dimensions. These two independent correlations replicated the ones we found in a previous study, where a consistent schedule of reinforcement was applied. At the same time, in the previous experiment only the gain-associated increase in representation resulted in perceptual conditioning, causing an overall positively valenced bias, while a similar type of conditioning did not occur for the loss-associated stimuli. Which could be the reason for this difference?

As briefly mentioned before, a consistent literature has grown over time with respect to an additional source of attention, the “emotional bias” (Vuilleumier, 2005; Pourtois, Schettino, & Vuilleumier, 2013). These studies have shown a prioritization of task-irrelevent threat-predictive stimuli with respect to task-irrelevant neutral ones, both on a behavioral and on a neural level (Vuilleumier, 2001; Brosch et al., 2011). At the same time, a clear influence of stimulus type and personality trait has emerged with respect to the extent of this bias across subjects. During a dot-probe task, task-irrelevant threat-related words were found to be able to capture attention, but only in clinically anxious participants. On the other hand, control subjects tended to shift attention away from the same type of stimuli (MacLeod et al., 1986; Mogg et al., 1995) Consistent results were also found in studies performed in non-clinical populations, but looking at differences in trait anxiety (Vandenhout et al., 1995; Mogg & Bradley, 1999). In order to account for this discrepancy, Mogg and Bradley (1998) developed the cognitive-motivational model. This theory posits that the relationship between the subjective threat value and attentional bias is not linear. For low values of threat, the individual will tend to ignore the stimulus; for high values of threat, selective attention will start to enhance it. Differences among individuals would therefore lie in the threshold at which selective attention will shift from suppression to prioritization. This also means that the effect of fearful and angry faces on attention may be due not simply to their aversive valence, but also to their stronger salience. As a consequence of this, the value of the stimulus will trespass the threshold which causes a shift from avoidance to capture within a bigger sector of the population under study (for an effect of different types of negative emotional expressions, see Ohman (2001)).

In view of these considerations, the lack of prioritization observed for aversive stimuli in our first MR study (chapter 3) could be due to the low level of arousal for visual stimuli which underwent conditioning with monetary rewards and punishments. With respect to this paradigm, an interpretation of the lack of a general effect for the various motivational conditions, could be then instead, that stimuli paired with loss presented in this case a higher level of arousal. The stronger effect of the aversive bias contrasted the opposite gain given by the positive bias on the reward vs punishment contrast, and the two controls cancelled each other out reciprocally across the whole population. In this case, which could be the reason for a diversity between the level of arousal of aversive stimuli from very similar paradigms (chapter 3 and 4), which only differ in terms of the consistency of the type of reinforcement? A possible explanation could lie in the arousing properties of unexpected events. Mogg and Bradley (1998) state the following about BIS:

> A primary function of the BIS is to compare actual with expected stimuli. The BIS operates in two modes. If actual stimuli are compatible with expected stimuli, it remains in `checking' mode, and control over behaviour remains with other brain systems, such as those involved in pursuing ongoing goals. However, if the actual and expected stimuli are discordant, or if the predicted stimuli are aversive, then the BIS takes direct control over behaviour, adopting a `control' model. (p.7).

This suggests that stimuli with a similar motivational value could trigger a state of higher arousal when unexpected. Another possible explanation is that, while effects from both forms of motivational controls are present at the level of the single trial, only the positive valence bias would trigger a long-lasting conditioning of the visual stimulus bound to the corresponding outcome. This is in contrast with the idea of aversive conditioning, according to which stimuli paired with punishment elicit a specific type of associative learning, but again it could be the case that only stimuli exceeding a certain threshold would undergo a similar form of plasticity. Additional studies will help answering these questions.

Results from univariate analysis in the first place replicate findings from previous studies which address the representation of valence and salience in the brain. Both electrophysiological studies in animals (Roesch & Olson, 2004; Kobayashi et al., 2006; Leathers & Olson, 2012) and imaging studies in humans (Engelmann et al., 2009; Litt et al., 2011) have highlighted a salience network in dlPFC, posterior parietal cortex (PPC), insula and dACC, and a valence one in orbitofrontal cortex (OFC), VS, rACC and PCC. This is confirmed in a meta-analysis by Bartra, which takes in consideration 206 published fMRI studies investigating neural correlates of subjective value (Bartra et al., 2013). We performed an equivalent analysis in our previous study, where a similar dissociation was found (Chapter 3). Even there, fronto-parietal regions (bilateral MFG and right IPL) showed a selective involvement in the salience contrast, while a series of subcortical and medial cortical regions were activated in the valence one. At the same time, when considering regions of striatum activated in the valence contrast, CDt was preferentially involved in the former experiment, and VS (and putamen) in the latter one.

There is an extremely long track record with regard to the connection between reward and VS, the region of striatum which is mostly linked with the hedonic aspect of reward, the liking (Berridge, 2007; Miller et al., 2014). With respect to the roles of different regions of the striatum, models of instrumental conditioning posit for VS the role of the “critic”, which learns to predict future reward, and for the dorsal striatum the role of the “actor”, which stores information about learnt stimulus-response associations (O’Doherty, Dayan, Schultz, Deichmann, Friston, & Dolan, 2004). Accordingly, VS has been found to track prediction error, and dorsal striatum to encode reward prediction (Seger et al., 2010). Other studies found feedback-related activity centered on VS and head of the caudate, and learning-related activity in the putamen and the body and tail of the caudate (Seger, 2008). All of these studies highlight a ventromedial-dorsolateral gradient which transforms reward-related information into specific stimulus-outcome associations. A similar hierarchy of information flow has been hypothesized by Haber and colleagues, who found that also at the cellular level the relationship among the different striatal regions via the midbrain dopamine cells is organized according to “an ascending spiral” (Haber et al., 2000).

This gradient has been mostly studied at the level of instrumental responses, but analogous models exist with respect to sensory stimuli. A “visual loop”, which connects temporal regions with CDt, guides visual object identification, assisting visual cortex in selection of appropriate visual representations (Seger et al., 2010). Moreover, representation of visual stimuli in this area has been shown to be modulated by reward (Yamamoto et al., 2013; Anderson et al., 2016), constituting a potential neural substrate for instances of the reward-associated attentional bias observed in human studies (DellaLibera and Chelazzi, 2009; Anderson et al., 2011). A direct comparison between the roles of the two opposite poles of caudate nucleus in processing conditioned visual stimuli has been pursued on an electrophysiological level: Kim and Hikosaka found that the head and the tail of the caudate encode flexible and stable values of visual objects, respectively (Kim & Hikosaka, 2013). The comparison between the results of our two fMRI studies represents a coherent correlate of this dynamic for human imaging. At the same time, while in the study by Kim and Hikosaka flexible stimuli retained the same value across the whole block and changed it only across blocks, in our second paradigm no consistent mapping at all was present between visual stimuli and outcome, even within the block. The two situations are thus similar, but not totally analogous. If visual stimuli retained the same value across blocks also in our case, following this line of reasoning, we could have expected increased activation in the caudate head, instead that at the level of VS. Increased activation at the level of VS then most probably reflects pure processing of the feedback, which in this case cannot be attributed to any specific visual category though.

In sum, according to this account, before a consistent association between a visual stimulus and a motivational value is established, trial-to-trial outcomes are tracked only at the level of VS. When a consistent pairing between stimulus and outcome is present, conditioning takes place, and stimulus-reward associations start to be encoded more dorsally and posteriorly. If this association is temporary and supported by working memory, it is encoded at the level of the head of the caudate (Seger, 2008; Kim and Hikosaka, 2013). When the association is prolonged in time, it starts to be represented at the level of CDt (Seger et al., 2010). Interestingly, this spatial shift in the locus of striatal activation could also possibly underpin the temporal shift of activation which we appreciate at the level of midbrain dopamine release, once a cue-outcome association is established (Schultz et al., 1997). In fact, the main input to midbrain dopamine cells is constituted by afferents from the basal ganglia. The differential recruitment that we observe would therefore underlie the process through which a previously neutral element (CS-) acquires predictive value and starts to cue the subsequent receipt of a specific outcome (CS+).

One important specificity of our studies concerns the relationship between opposite motivational outcomes. In this respect, we find that it is the valence representation of visual stimuli, not the salience one, which is progressively processed according to this implicit flow of information. A common thread in the way in which this implicit perceptual bias has been found to affect selection is constituted by its involuntary and anti-strategic nature, which could be explained by a model-free, simple controller, which cannot exploit the higher computational efficiency of neocortex. It should be pointed out that valence implies a model-free controller, but the opposite is not true: a model-free controller could potentially follow a salience-based pattern, as long as the punishment outcome starts to prioritize selection. This is indeed the case in many instances, and our model does not exclude such a situation.

Finally, another point could be made with respect to the relationship of our two models with awareness. We have hypothesized the salience network to represent deliberate, volitional top-down control and the valence one to embody the automatic and unconscious selection history bias. One obvious feature of volitional control consists in its link with awareness, and it could be interesting to notice that the areas which differentially activate for the model-based, rational scheme of behavior encompass the regions which have been more strictly connected to consciousness, namely parietal and prefrontal cortex (Rees, Kreiman, & Koch, 2002). Intriguingly, this correspondence extends to visual regions: in both paradigms, we find that primary visual regions follow the valence pattern, while bilateral fusiform gyrus exhibits a salience mode of activation. This reflects the subdivision that we also find for visual areas in terms of their contribution to conscious experience. On one side, activity in V1 has been shown to be largely independent from, and scarcely contributing to, conscious perception (He, Cavanagh, & Intriligator, 1996). On the other, activity in ventral visual cortex correlates instead rather well with conscious visual experience, and seems to be a necessary, even if not sufficient, component for awareness (Hirsch, DeLaPaz, Relkin, Victor, Kim, Li, Borden, Rubin, & Shapley, 1995). As mentioned before, the additional contribution of fronto-parietal regions constitutes an essential factor to this process, as can be seen in the case of neglect, where the sensory stimulation leads to activation of association visual areas but fails to reach consciousness (Driver & Mattingley, 1998).

Chapter 5

Summary and general discussion

Reward is an important tool to investigate the relationship between top-down control and bottom-up salience in visual attention. In fact, limits of this theoretical account have been highlighted for a series of phenomena, like priming of pop-out or value-driven attentional capture, where visual selection could not be explained either by current task goals nor by physical salience of the stimulus. An additional source of attentional control has accordingly been posited, characterized on one hand by its endogenous nature, but on the other also by its automaticity. This represents an important caveat which aims at differentiating effects of reward due to increase of attention and motivation from effects due to the appetitive nature of its positive valence. In this framework, we thought that the investigation of the joined effects of reward and punishment on visual selection could have shed more light onto the issue. These two outcomes are in fact characterized by similar motivational salience with respect to a neutral stimulus (they both signal events with strong behavioral relevance) but at the same time by opposite valence. The mutual relationship between the impact of these two outcomes on visual attention was therefore expected to follow either of these two main patterns, or a combination of them. To this aim, we employed a well-developed paradigm which investigates visual search for object categories in naturalistic scenes.

In chapter 2 of this thesis, we investigated this dynamic on a behavioral level, by looking at how different types of feedback modulated visual selection and attentional capture. Subjects had to look for instances of high-level categories of objects (people, cars, trees and houses) in briefly presented pictures of real-world scenes. These categories were imbued with different monetary payoffs: gain, loss or neutral outcome. Participants were informed at the beginning of each block about the target category, and about the type of feedback they would have received in case of correct performance. On one side, for each outcome type, target presence was always paired with the same directionality in outcome regardless of response (they always gained points with a reward category, always lost points with a punishing one, and gained a negligible amount of points with neutral categories). On the other, correct detection of the target was always the response leading to the best possible payoff. Correct detection of rewarding and punishing categories was more valuable (100 points) than detection of neutral ones (2 points), and the first two were equally valuable between each other, representing the motivationally salient targets. We found that performance was not biased in favor of the two salient categories, as could have been expected if participants showed a strategic approach to the monetary feedback, or were equally aroused by the two relevant outcomes. Instead, both measures were found to follow a valence scheme, such that categories associated with loss showed a lower level of accuracy with respect to the ones associated with gain, and correspondingly drew less attention when acting as task-irrelevant distractors.

In chapter 3, we used a very similar paradigm in an MR context, in order to see how category information about these stimuli as assessed by MVPA in occipito-temporal object-selective visual cortex (OSC) varied as a function of the same motivational outcomes. Again, participants had to detect the presence of real-world categories in natural scenes, with target type changing from block to block. Each category was tied to a specific rewarding, punishing or neutral payoff, so that subjects knew in advance the payoff they would have received in case of correct (and incorrect) performance. The amount of information in the scene was evaluated for each category both when it acted as the target and when it acted as a distractor, and as a function of motivational outcome. Consistently with previous literature within the field, information about the various categories was enhanced when they were acting as targets and suppressed when they were acting as distractors, irrespectively of motivational conditions. Moreover, we found that the representation of the various categories in OSC followed a valence-like pattern, with the punished category showing the lowest amount of increase in representation when acting as a target, and the least amount of suppression when acting as a distractor. Importantly, the degree to which this representation was following a valence pattern was found to correlate, across participants, with the degree to which univariate activation of dopaminergic midbrain regions was following a similarly valence-shaped pattern. Finally, we also found evidence for a representation of visual categories in posterior parietal cortex, which, strikingly, followed in this case the latter of our expectations, namely the salience pattern.

Don't think of a white bear

In Chapter 1, I have reviewed a series of studies which found that visual attention can be affected by opposing motivational outcomes sometimes according to a valence, and sometimes according to a salience pattern. I will now try to answer a question left unsolved from that chapter: what causes the former or the latter pattern to shape the deployment of attention? In order to answer this question, I will first try to better delineate the nature of the reward-associated attentional bias, and then to assess how punishment affects it.

A growing body of evidence has accumulated over time suggesting that the traditional dichotomy between top-down and bottom-up bias does not provide a complete account of attentional control. One example is provided by a series of studies by Maljkovic & Nakayama (1994, 2000), where they found that repetition of a salient feature of the target causes inter-trial priming in spite of task contingencies and goals. This phenomenon has been named priming of pop-out (PoP). Other instances are represented by the reward-associated phenomena that we have considered throughout this thesis. Awh and colleagues (2012) claim these and other examples to be the expression of an alternative type of control, reuniting all of these cases under a common framework. This additional bias, differently from exogenous control, depends on what has been learned about the prior trials and does not rely solely on the current form of the stimulus. As a consequence, some if its instances have been classified as top-down or endogenous (Wolfe, Butcher, Lee, & Hyle, 2003). On the other hand, this additional bias does not necessarily represent the task-relevant set, and can easily lead to misallocation of attention, suggesting its automatic and non-strategic nature. The possibility of a dissociation between endogenous and voluntary control was brought to attention long ago by means of a consideration by Dostoevskij, who noticed how hard it could result to deliberately stop thinking about a white bear, after having been told to do so. This idea led to studies about thought suppression and to the so-called white bear problem (Wegner, Schneider, Carter, & White, 1987). Awh and colleagues posit this third source of attention to be more generally based on selection history, with reward history being a specific instance of this more general category.

What are the neural structures which subtend this alternative endogenous control? DlPFC and IPS have been highlighted as the brain regions giving rise to the classically conceived, strategic top-down bias (Corbetta & Shulman, 2002). Interestingly, in a study where the behavior of highly trained animals had already been shaped by a history of rewarding feedback, it was found that PFC lesion affected subsequent endogenous selection proportionally to the rate of task-switching (Rossi et al, 2007). While the impairment was small when the target feature remained relatively fixed over trials, it increased dramatically when this had to be rapidly switched across trials. This suggests the involvement of other structures in mediating attentional bias caused by selection history, which could explain the spared performance in the more constant condition. Throughout this review, we have found striatum to be systematically involved in the occurrence of reward-associated attentional bias (Pessiglione et al., 2007; Krebs et al, 2012; Anderson et al., 2014). Although the precise functions of this structure have not yet been completely elucidated, many hypotheses exist in this regard. Redgrave proposes that basal ganglia work as a central selector among diverse cortical and subcortical circuits which cannot all be expressed in parallel (Redgrave, Gurney, & Reynolds, 2008). Moreover, striatum also represents a core region for reinforcement learning (Montague et al., 1996). One seminal theory asserts that dopaminergic input to this area provides a reward prediction error signal which allows selection (Schultz, 2000). Redgrave's theory focuses on motor selection, but at the same time we know that different parts of striatum deal with different cortical regions and functions. Correspondingly, Seger proposes for CDt, the portion of the caudate connected to temporal lobe, a role in visual selection (Seger et al, 2013). We have already considered a couple of electrophysiological studies regarding the involvement of this subregion in the deployment of visual attention (Yamamoto et al., 2012; Yamamoto et al., 2013), and other examples exist in imaging studies on humans (Seger et al, 2010).

The potential role of basal ganglia in giving rise to the selection history bias should not surprise, as this structure also enables implicit learning (Packard et al, 1989; Knowlton et al, 1996). This is a similarly automatic process, which forges simple stimulus-response associations following the delivery of some sort of feedback used as a teaching signal. Moreover, PoP itself has been explained as a form of short-term implicit memory (Maljkovic & Nakayama, 1994; 2000). Reward-associated object- and feature-priming, value-driven attentional capture, subliminal activation of motivational drive, reward-driven interference in conflict-based tasks: all these phenomena would then represent the product of reward on implicit visual memory, in the form of an involuntary deployment of attention caused by a similarly automatic, and generally unconscious, reward expectation. Reward-associated object- and feature-priming would constitute a form of short-term implicit visual memory akin to PoP, but triggered by the delivery of reward. This would follow both consistent (Kristjansson et al., 2010) and inconsistent schedules (Della Libera & Chelazzi, 2006; Hickey et al., 2010a), but with different fates. When the association is inconsistent, its effects would limit to inter-trial priming, and then vanish as a function of subsequent allocations of attention. One form of attentional capture would also be the product of a similar lingering of short-term implicit memory on a previously highlighted stimulus (Hickey & van Zoest, 2012; Hickey et al., 2015). If the association between visual stimulus and outcome is consistent and extended in time, then object-skill can develop, with an increase in the saliency, detection and recognition of the stimulus. This type of learning, just like its motor counterpart, action-skill, displays high-capacity and long-term retention (Hikosaka et al. 2013). At the same time, once established, associations formed in this way are essentially blind to recent changes in the values of individual objects, showing a lesser degree of flexibility. As a consequence, at this stage, these stimuli acquire visual saliency on their own, and can easily lead to a second form of attentional capture, when in conflict with the current attentional set (Della Libera & Chelazzi, 2009; Anderson et al., 2011).

Nevertheless, we know that reward does not only affect behavior on an implicit level, but also on an explicit one. Accordingly, the classically conceived, strategic endogenous bias is also expected to be modulated by rewarding outcomes, but in this case through an increase in attention and motivation. We have seen that many studies find, in the rewarded condition, an increase in activation in the fronto-parietal regions held responsible for top-down modulation (Roesch and Olson, 2004; Engelmann et al., 2009; Leathers & Olson, 2012; Krebs et al., 2012). This network of areas allows a conscious, model-based approach to motor planning and perceptual selection, which develops more slowly, but is less prone to erroneous generalizations and useless perseveration. A similar dissociation between a reflexive, model-based controller, mainly based on prefrontal cortex, and a reflective, model-free system, centered on basal ganglia, is also theorized by Daw and colleagues at the level of action selection and decision making (Daw et al., 2005).

Crucially, while reward has an analogous, facilitating effect on both the voluntary and the automatic type of control, the way in which punishment affects attention could differ between and also within systems. In the model-based, strategic top-down control, punishment reasonably contributes to the build-up of attention and causes an increase in the saliency of the stimulus. We have seen that fronto-parietal regions show an increase in activation not only in the case of reward (Krebs et al., 2012), but more in general for all motivationally salient outcomes (Roesch & Olson, 2004; Engelmann et al., 2009; Leathers & Olson, 2012). We also find this in both our MR studies (Chapter 3 and 4), where the salience contrast revealed increased activation in fronto-parietal regions (MFG and IPL), parts of the top-down attentional control network, indicating that salient categories were in fact recruiting regions of the brain involved in the endogenous, strategic deployment of attention.

For the automatic endogenous attentional bias, on the other hand, aversive outcomes will be instead harder to interpret, maybe because of the more basic computations that this system is capable of operating. The dilemma arising during punishment could be clarified by analyzing the different role of dopamine in appetitive and aversive contexts. Larger rewards determine an increase in dopamine release in a series of regions involved in evaluation and motivation, like for example striatum, leading in turn to increased vigor during action, and to selection of behaviors bound to reward delivery. Larger punishments decrease dopamine release, and this is particularly evident in Pavlovian conditioning paradigms, where no action is effective in preventing aversive outcome. At the same time, when punishment is avoidable, dopamine serves this behavior (Oleson, Gentry, Chioma, & Cheer, 2012). How is the contrast between this decrease of dopamine during aversive events and increases during avoidance resolved? A computational model by Dayan proposes a solution based on an adjusted EV signal (Dayan, 2012). This signal represents the level of punishment that is potentially avoidable through action, and consequently drives a dopaminergic response. A recent study by Rigoli investigates neural activation during a visual search task using this model as a framework for the analysis of imaging data (Rigoli, Chew, Dayan, & Dolan, 2016). Coherently, results showed that activity in VS and VTA/SN covaried with net EV, while activity in anterior insula, and again in VTA/SN, covaried with adjusted EV. This means that, again, the subcortical network represented outcomes according to a valence-like pattern and the cortical one according to a salience-like one.

Still, the way in which appetitive and aversive outcomes are represented in the whole striatum is not completely clear, although this structure seems to follow a ventro-dorsal gradient, with motivational valence represented ventrally and motivational salience dorsally. As a consequence, the specific attentional network, and the specific region of the basal ganglia (in case of a leading role by selection history bias) involved in the task could set the corresponding pattern of attentional deployment. Therefore, in tasks where the strategic endogenous system is prevalently active, or similarly where regions of the automatic attention network, but characterized by a salience-like pattern, are leading attention, then punishment will concur with reward in the enhancement of representation of salient stimuli (as in Raymond & O'Brien, 2009 (full attention condition) and in Engelmann et al., 2009). In instances where other regions of the automatic system, displaying in this case a valence-like pattern of activation, are instead prevalently active, punishment will fail in prioritizing aversive stimuli, even when this shows to be counterproductive for the fitness of the individual, and cause a selective suppression of representation (as in Laufer & Paz, 2012). As we have already considered in the end of previous section, visual search for high-level categories of objects in real-world scenes can be classified as a parallel, efficient type of search, probably due to the highly-automatized operations subserving this process. A leading role of the involuntary endogenous bias in this type of task would then be consistent with the automaticity of these cognitive functions on one side, and with the involvement of its putative neural substrates, basal ganglia, in action- and object-skill, on the other (Hikosaka, 2013). Indeed, Awh (2012) suggests that contextual cueing, which we have seen to be critical in improving efficiency in naturalistic search, could be one of the various instances where selection history bias manifests itself.

Unfortunately, no study addresses the specific issue of the representation of opposing motivational outcomes within the context of the CDt, the area which is more relevant to our topic. Nevertheless, this region has shown to present quite a series of differences with respect to its neighboring caudate areas, such as the CD body and head (Hikosaka et al., 2014). It is not trivial therefore to foresee what pattern of activation it could display. Interestingly, the CDt shows increased activation in our univariate analysis, and more specifically in the valence contrast, supporting a corresponding type of modulation by motivational outcomes. A similar bias for positive stimuli at the level of CDt can be inferred from a study by Kim and colleagues (Figure 3E), although the authors do not focus on this specific aspect of the data (Kim et al., 2013). If we consider that CDt presents a valence-like model of activation, and similarly beware the leading role of this structure in shaping the automatic deployment of attention, then the apparently paradoxical pattern of our (and of other similar studies) behavioral and imaging data could be explained. Another way of looking at this phenomenon is that this type of automatic, rapid visual attentional deployment is fundamentally shaped by object incentive value, rather than by its motivational, strategic or more generally arousal value.

Expect the unexpected

In chapter 4, we looked at how the same types of feedback affected cortical representation of visual stimuli when no consistent pairing was present between objects categories and motivational outcomes. Participants had to look for two categories at the same time, and for each trial discriminate which of the two targets appeared in the briefly presented real-world scene. The type of feedback in this case was not following a consistent schedule across trials, such that correct performance would have always led to the optimal feedback, but the trial could have been a rewarding, a punishing or a neutral one, and subjects discovered the trial type only at the end, at the moment of outcome delivery. In this case, we did not have any distractor category, but just a “present” and an “absent” category which alternated across trials. What we observed in this case was that the appearance of the category biased the trialwise activation of ventro-temporal cortex in favour of the multivariate representation of the present category, while the representation of the absent one was suppressed with respect to baseline, in a way which resembled what happened to salient distractors in the previous visual search paradigms. With respect to the motivational modulation of information in OSC, we observed a totally specular pattern of response across participants which depended on one side on the responsivity of SN, and on the other on a personality trait which assessed the individual sensitivity to aversive cues. More specifically, the degree to which SN displayed a valence-based activation determined the extent to which category information in OSC followed the positive valence pattern. On the contrary, the higher the BIS scores of the subjects, the higher the extent to which category information followed an inverse valence pattern. These modulations were present in both studies, and in both cases reciprocally independent. Nevertheless, while these two opposing biases were completely balanced and cancelled each other in the case of an inconsistent schedule, a net overall positive bias took over in the case of a consistent schedule. Hypotheses trying to account for this discrepancy have been illustrated in chapter 4.

We have already talked about the putative additional control due to prior association with reward; in this experiment, we also find hints for a symmetric bias towards stimuli associated with loss. Just like the reward-associated control, this bias seems to reflect an affective, irrational prioritization: the value of rewarded and punished categories is exactly the same, but high trait anxiety individuals tend to prioritize the latter over the former. It should be pointed out once again that orienting to motivational stimuli according to a salience model would have not necessarily reflected a strategic approach to the problem. It could have also resulted from an undifferentiated, automatic prioritization of all affectively relevant stimuli, rewarded and punished ones. In this case, though, the source of the attentional control would have most probably been common for both kinds of affective stimuli with respect to neutral ones. Instead, a resuming view from what has been considered throughout this thesis (both in terms of the experiments of the project, and of the other studies which have been reviewed), suggests that, although a bias towards both types of outcomes appears to be present, it does not the share the same origin. Each of these two sources of prioritization displays its own relationship with different personality traits, emotional states and task specifics (attentional load, time windows available for response, unexpectedness of stimulus type). Accordingly, each one of them also presents its own specific neural substrate: basal ganglia for the reward-associated selection history bias (Anderson, 2016), and amygdala for the emotional control (Vuilleumier, 2005). A similar segregation between a positive and a negative evaluative channel has been posited by Cacioppo and Gardner in their dissertation about different theories of emotion (Cacioppo & Gardner, 1999). These considerations picture an even more complex situation from the one we had foreshadowed before, with four different sources of attention interacting with each other independently: the positively and the negatively valenced bias, plus the traditional exogenous bias and strategic top-down control.

Previously in this chapter, we have addressed the relationship between dopamine and selection history bias, and the way in which reward and punishment could influence dopamine release. In particular, we have considered how complex it could turn out for punishment to highjack this system in order to convey signals about the strategic utility (or threatening value) of a stimulus, leading to the paradoxical situation of a selective suppression of potentially behaviorally relevant stimuli. In the perspective of an additional aversive bias, the role of other neurotransmitters, such as for example serotonin, should also be taken in consideration. Niv and colleagues developed a model which considers the tradeoff between the energetic cost for a particular behavior and the opportunity cost of time (the amount of reward lost when this behavior is not performed) (Niv, Daw, Joel, & Dayan, 2007). This model assumes that tonic levels of dopamine encode this evaluation, which would determine not only the selection of the appropriate behavior (as theories of reinforcement learning posit), but also the vigor with which this behavior has to be performed. Cools and colleagues extended this model, by considering the additional role of serotonin in this tradeoff which regulates response vigor (Cools, Nakamura, & Daw, 2011). The authors posit that this other neurotransmitter would code for punishment, and that its release would mediate punishment-induced inhibition. If we consider again the aversive bias and the neural structures which mediate it, it could be interesting to note that acute tryptophan depletion (ATD) (a treatment which is used in order to lower the concentration of serotonin of the brain) was found to enhance activation of the amygdala in response to fearful faces with respect to neutral ones (Cools et al., 2005). Moreover, getting back to our considerations with respect to the relationship between personality traits and the strength of the emotional bias, the entity of this enhancement was found to correlate with BIS scores of participants.

Finally, a direct comparison between the univariate analyses of BOLD signal of the two MR studies (Chaper 3 and Chapter 4) revealed another interesting finding. In both studies, the salience contrast displayed an increase in activation in MFG, IPL and AI, consistent with a role for fronto-parietal lobes in providing a model-based, strategic evaluation of the outcomes. The valence contrast showed instead an increased activation (among other structures) at the level of striatum. Nevertheless, the exact location of this increase varied according to the type of reinforcement schedule. Response was higher for rewarded stimuli compared to punished stimuli in ventral striatum (VS) during an inconsistent schedule, and in caudate tail (CDt) following a consistent schedule of reinforcement. This agrees with studies which posit for VS the role of the critic (feedback evaluation) and for other parts of striatum the role of the actor (storage of specific stimulus-outcome associations). Moreover, it strengthens the hypothesis which considers CDt, as part of the visual cortico-striatal loop, as the site where this type of implicit learning is stored for visual stimuli.

In conclusion, this project tried to further characterize a control system which has recently been shown to shape attentional deployment in addition to the traditional top-down and bottom-up ones, namely the reward-associated selection history bias. Our study further supports the involuntary and irrational nature of this new control, by highlighting that motivationally salient, but negatively valenced, stimuli, fail to activate it, leading to a specific neglect for punishing stimuli. We also find hints for a specular, threat-induced, “emotional” bias, which for this type of stimuli appears to be strongly dependent on the personality of the subjects. These alternative sources of attention represent instances of the many paradoxical ways in which affect can direct behavior. These examples range from trivial episodes, like our naivety to smart advertisement techniques, to serious issues like addiction, anxiety and depression. It is now becoming always clearer how these powerful, seducing, but also potentially maladaptive mechanisms, do not only affect the way we act on the world, but also the way we perceive it.

REFERENCES

Amaral DG, Behniea H, Kelly JL (2003) Topographic organization of projections from the amygdala to the visual cortex in the macaque monkey. Neuroscience. 118:1099-120.

Anderson AK, Christoff K, Panitz D, De Rosa E, Gabrieli JD (2003) Neural correlates of the automatic processing of threat facial signals. J Neurosci. 23:5627-33.

Anderson BA, Laurent PA, Yantis S (2011) Value-driven attentional capture. Proc Natl Acad Sci U S A. 108:10367-71.

Anderson BA, Laurent PA, Yantis S (2014) Value-driven attentional priority signals in human basal ganglia and visual cortex. Brain Res. 1587:88-96.

Anderson BA (2016) Reward processing in the value-driven attention network: reward signals tracking cue identity and location. Soc Cogn Affect Neurosci. Sep 27.

Anstrom KK, Miczek KA, Budygin EA (2009) Increased phasic dopamine signaling in the mesolimbic pathway during social defeat in rats. Neuroscience. 161:3-12.

Awh E, Belopolsky AV, Theeuwes J (2012) Top-down versus bottom-up attentional control: a failed theoretical dichotomy. Trends Cogn Sci. 16:437-43.

Balleine BW, Daw ND, O’Doherty JP (2008) Neuroeconomics: Decision Making and the Brain (eds Glimcher PW, Camerer CF, Fehr E, Poldrack RA) 367–389 (AcademicPress).

Bartra O, McGuire JT, Kable JW (2013) The valuation system: a coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. Neuroimage. 76:412-27.

Baumeister RF, Bratslavsky E, Finkenauer C, Vohs KD (2001) Bad is stronger than good. Review of General Psychology, 5:323-370.

Berridge KC, Robinson TE (1998) What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? Brain Res Brain Res Rev. 28:309-69.

Berridge KC (2007) The debate over dopamine's role in reward: the case for incentive salience. Psychopharmacology. 191:391-431.

Bijleveld E, Custers R, Aarts H (2010) Unconscious reward cues increase invested effort, but do not change speed-accuracy tradeoffs. Cognition. 115:330-5.

Brainard DH (1997) The psychophysics toolbox. Spat. Vis. 10:437–442.

Brett M, Anton JL, Valabregue R, Poline JB (2002) Region of interest analysis using an SPM toolbox [abstract] Presented at the 8th International Conference on Functional Mapping of the Human Brain, June 2-6, Sendai, Japan.

Bromberg-Martin ES, Matsumoto M, Hikosaka O (2010) Dopamine in motivational control: rewarding, aversive, and alerting. Neuron. 68:815-34.

Brosch T, Pourtois G, Sander D, Vuilleumier P (2011) Additive effects of emotional, endogenous, and exogenous attention: behavioral and electrophysiological evidence. Neuropsychologia. 49:1779-87.

Bucker B, Theeuwes J (2016) Appetitive and aversive outcome associations modulate exogenous cueing. Atten Percept Psychophys.

Cacioppo JT, Gardner WL (1999) Emotion Annu Rev Psychol. 50:191-214.

Carver C, White T (1994) Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: the BIS/BAS scales. J. Pers. Soc. Psychol. 67, 319–333.

Chun MM (2000) Contextual cueing of visual attention. Trends Cogn Sci. 4:170-178.

Cohen J (1988) Statistical power analysis for the behavioral sciences (2nd ed.). Hillsdale, NJ: Lawrence Earlbaum Associates.

Cools R, Calder AJ, Lawrence AD, Clark L, Bullmore E, Robbins TW (2005) Individual differences in threat sensitivity predict serotonergic modulation of amygdala response to fearful faces. Psychopharmacology (Berl). 180:670-9.

Cools R, Nakamura K, Daw ND (2001) Serotonin and dopamine: unifying affective, activational, and decision functions. Neuropsychopharmacology. 36:98-113.

Cooper JC, Knutson B (2008) Valence and salience contribute to nucleus accumbens activation. Neuroimage. 39:538-47.

Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven attention in the brain. Nat Rev Neurosci. 3:201-15.

Daw ND, Niv Y, Dayan P (2005) Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. Nat Neurosci. 8:1704-11.

Day JJ, Roitman MF, Wightman RM, Carelli RM (2007) Associative learning mediates dynamic shifts in dopamine signaling in the nucleus accumbens. Nat Neurosci. 10:1020-8.

Dayan P (2012) Instrumental vigour in punishment and reward. Eur J Neurosci. 35:1152-68.

Della Libera C, Chelazzi L (2006) Visual selective attention and the effects of monetary rewards. Psychol Sci. 17:222-7.

Della Libera C, Chelazzi L (2009) Learning to attend and to ignore is a matter of gains and losses. Psychol Sci. 20:778-84.

Delorme A, Richard G, Fabre-Thorpe M (2010) Key visual features for rapid categorization of animals in natural scenes. Front Psychol. 1:21.

Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. Annu Rev Neurosci. 18:193-222.

Dolan RJ, Dayan P (2013) Goals and habits in the brain. Neuron. 80:312-25.

Driver J, Mattingley JB (1998) Parietal neglect and visual awareness. Nat Neurosci. 1:17-22.

Elbich DB, Scherf S (2016) Beyond the FFA: Brain-behavior correspondences in face recognition abilities. Neuroimage. 147:409-422.

Engelmann JB, Damaraju E, Padmala S, Pessoa L (2009) Combined effects of attention and motivation on visual task performance: transient and sustained motivational effects. Front Hum Neurosci. 3:4.

Etkin A, Klemenhagen KC, Dudman JT, Rogan MT, Hen R, Kandel ER, Hirsch J (2004) Individual differences in trait anxiety predict the response of the basolateral amygdala to unconsciously processed fearful faces. Neuron. 44:1043-55.

Evans KK, Georgian-Smith D, Tambouret R, Birdwell RL, Wolfe JM (2013) The gist of the abnormal: above-chance medical decision making in the blink of an eye. Psychon Bull Rev. 20:1170-5.

Foerde K, Knowlton BJ, Poldrack RA (2006) Modulation of competing memory systems by distraction. Proc Natl Acad Sci U S A. 103:11778-83.

Forstmann BU, Dutilh G, Brown S, Neumann J, von Cramon DY, Ridderinkhof KR, Wagenmakers EJ (2008) Striatum and pre-SMA facilitate decision-making under time pressure. Proc Natl Acad Sci U S A. 105:17538-42.

Haber SN, Fudge JL, McFarland NR (2000) Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. J Neurosci. 20:2369-82.

Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P (2001) Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science. 293:2425-30.

He S, Cavanagh P, Intriligator J (1996) Attentional resolution and the locus of visual awareness. Nature. 383:334-7.

Hickey C, Chelazzi L, Theeuwes J (2010a) Reward changes salience in human vision via the anterior cingulate. J Neurosci. 30:11096-103.

Hickey C, Chelazzi L, Theeuwes J (2010b) Reward guides vision when it's your thing: trait reward-seeking in reward-mediated visual priming. PLoS One. 5

Hickey C, van Zoest W (2012) Reward creates oculomotor salience. Curr Biol. 22:R219-20.

Hickey C, Kaiser D, Peelen MV (2015) Reward guides attention to object categories in real-world scenes. J Exp Psychol Gen. 144:264-73.

Hickey C, Peelen MV (2015) Neural mechanisms of incentive salience in naturalistic human vision. Neuron. 85:512-8.

Hikosaka O, Yamamoto S, Yasuda M, Kim HF (2013) Why skill matters. Trends Cogn Sci. 17:434-41.

Hikosaka O, Kim HF, Yasuda M, Yamamoto S (2014) Basal ganglia circuits for reward value-guided behavior. Annu Rev Neurosci. 37:289-306.

Hjornevik T, Leergaard TB, Darine D, Moldestad O, Dale AM, Willoch F, Bjaalie JG (2007) Three-dimensional atlas system for mouse and rat brain imaging data. Front Neuroinform. 1:4.

Hirsch J, DeLaPaz RL, Relkin NR, Victor J, Kim K, Li T, Borden P, Rubin N, Shapley R (1995) Illusory contours activate specific regions in human visual cortex: evidence from functional magnetic resonance imaging. Proc Natl Acad Sci U S A. 92:6469-73.

Ikeda T, Hikosaka O (2003) Reward-dependent gain and bias of visual responses in primate superior colliculus. Neuron. 39:693-700.

Jahoda M (1958) Current concepts of positive mental health (New York: Basic Books)

Kanwisher N, McDermott J, Chun MM (1997) The fusiform face area: a module in human extrastriate cortex specialized for face perception. J. Neurosci. 1711: 4302–4311.

Kim HF, Hikosaka O (2013) Distinct basal ganglia circuits controlling behaviors guided by flexible and stable values. Neuron. 79:1001-10.

Kim HF, Hikosaka O (2015) Parallel basal ganglia circuits for voluntary and automatic behaviour to reach rewards. Brain. 138:1776-800.

Knowlton BJ, Mangels JA, Squire LR (1996) A neostriatal habit learning system in humans. Science. 273:1399-402.

Knutson B, Adams CM, Fong GW, Hommer D (2001) Anticipation of increasing monetary reward selectively recruits nucleus accumbens. J Neurosci. 21:RC159.

Knutson B, Rick S, Wimmer GE, Prelec D, Loewenstein G (2007) Neural predictors of purchases. Neuron. 53:147-56.

Kobayashi S, Nomoto K, Watanabe M, Hikosaka O, Schultz W, Sakagami M (2006) Influences of rewarding and aversive outcomes on activity in macaque lateral prefrontal cortex. Neuron. 51:861-70.

Krebs RM, Boehler CN, Woldorff MG (2010) The influence of reward associations on conflict processing in the Stroop task. Cognition. 117:341–347.

Krebs RM, Boehler CN, Roberts KC, Song AW, Woldorff MG (2012) The involvement of the dopaminergic midbrain and cortico-striatal-thalamic circuits in the integration of reward prospect and attentional task demands. Cereb Cortex. 22:607-15.

Kristjánsson A, Sigurjónsdóttir O, Driver J (2010) Fortune and reversals of fortune in visual search: Reward contingencies for pop-out targets affect search efficiency and target repetition effects. Atten Percept Psychophys. 72:1229-36.

Laufer O, Paz R (2012) Monetary loss alters perceptual thresholds and compromises future decisions via amygdala and prefrontal networks. J Neurosci. 32:6304-11.

Leathers ML, Olson CR (2012) In monkeys making value-based decisions, LIP neurons encode cue salience and not action value. Science. 338:132-5.

Litt A, Plassmann H, Shiv B, Rangel A (2011) Dissociating valuation and saliency signals during decision-making. Cereb Cortex. 21:95-102.

MacLeod C, Mathews A, Tata P (1986) Attentional bias in emotional disorders. J Abnorm Psychol. 95:15-20.

MacLeod C, Mathews A (1988) Anxiety and the allocation of attention to threat. Q J Exp Psychol A. 40:653-70.

Maldjian JA, Laurienti PJ, Kraft RA, Burdette JH (2003) An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fmri data sets. Neuroimage. 19:1233–1239.

Maljkovic V, Nakayama K (1994) Priming of pop-out: I. Role of features. Mem Cognit. 22:657-72.

Maljkovic V, Nakayama K (2000) Priming of pop-out: III. A short-term implicit memory system beneficial for rapid target selection. Vis Cognit. 7:571–595.

Matsumoto M, Hikosaka O (2009) Two types of dopamine neuron distinctly convey positive and negative motivational signals. Nature. 459:837-41.

Maunsell JH (2004) Neuronal representations of cognitive state: reward or attention? Trends Cogn Sci. 8:261-5.

Miller EM, Shankar MU, Knutson B, McClure SM (2014) Dissociating motivation from reward in human striatal activity. J Cogn Neurosci. 26:1075-84.

Mogg K, Bradley BP, Williams R (1995) Attentional bias in anxiety and depression: the role of awareness. Br J Clin Psychol. 34:17-36.

Mogg K, Bradley BP (1998) A cognitive-motivational analysis of anxiety. Behav Res Ther. 36:809-48.

Mogg K, Bradley BP (1999) Some methodological issues in assessing attentional biases for threatening faces in anxiety: a replication study using a modified version of the probe detection task. Behav Res Ther. 37:595-604.

Montague PR, Dayan P, Sejnowski TJ (1996) A framework for mesencephalic dopamine systems based on predictive Hebbian learning. J Neurosci. 16:1936-47.

Niv Y, Daw ND, Joel D, Dayan P (2007) Tonic dopamine: opportunity costs and the control of response vigor. Psychopharmacology (Berl) 191: 507–520.

O'Doherty J, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ (2004) Dissociable roles of ventral and dorsal striatum in instrumental conditioning. Science. 304:452-4.

Ohman A, Lundqvist D, Esteves F (2001) The face in the crowd revisited: a threat advantage with schematic stimuli. J Pers Soc Psychol. 80:381-96.

Oleson EB, Gentry RN, Chioma VC, Cheer JF (2012) Subsecond dopamine release in the nucleus accumbens predicts conditioned punishment and its successful avoidance. J Neurosci. 32:14804-8.

Oosterhof NN, Connolly AC, Haxby JV (2016) CoSMoMVPA: multi-modal multivariate pattern analysis of neuroimaging data in Matlab/GNU Octave. Frontiers in Neuroinformatics. 10:27.

Packard MG, Hirsh R, White NM (1989) Differential effects of fornix and caudate nucleus lesions on two radial maze tasks: evidence for multiple memory systems. J Neurosci. 9:1465-72.

Pavlov I (1927) Conditioned Reflexes (Oxford University Press).

Peelen MV, Fei-Fei L, Kastner S (2009) Neural mechanisms of rapid natural scene categorization in human visual cortex. Nature. 460:94-7.

Peelen MV, Kastner S (2011) A neural basis for real-world visual search in human occipitotemporal cortex. Proc Natl Acad Sci U S A. 108:12125-30.

Pessiglione M, Schmidt L, Draganski B, Kalisch R, Lau H, Dolan RJ, Frith CD (2007) How the brain translates money into force: a neuroimaging study of subliminal motivation. Science. 316:904-6.

Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. Nature. 400:233-8.

Pourtois G, Schettino A, Vuilleumier P (2013) Brain mechanisms for emotional influences on perception and attention: what is magic and what is not. Biol Psychol. 92:492-512.

Pratto F, John OP (1991) Automatic vigilance: the attention-grabbing power of negative social information. J Pers Soc Psychol. 61:380-91.

Raymond JE, Shapiro KL, Arnell KM (1992) Temporary suppression of visual processing in an RSVP task: an attentional blink? J Exp Psychol Hum Percept Perform. 18:849-60.

Raymond JE, O'Brien JL (2009) Selective visual attention and motivation: the consequences of value learning in an attentional blink task. Psychol Sci. 20:981-8.

Redgrave P, Gurney K, Reynolds J (2008) What is reinforced by phasic dopamine signals? Brain Res Rev. 58:322-39.

Reeder RR, Peelen MV (2013) The contents of the search template for category-level search in natural scenes. J Vis. 13:13.

Rees G, Kreiman G, Koch C (2002) Neural correlates of consciousness in humans. Nat Rev Neurosci. 3:261-70.

Resnik J, Sobel N, Paz R (2011) Auditory aversive learning increases discrimination thresholds. Nat Neurosci. 14:791-6.

Rigoli F, Chew B, Dayan P, Dolan RJ (2016) Multiple value signals in dopaminergic midbrain and their role in avoidance contexts. Neuroimage. 135:197-203.

Roelfsema PR, van Ooyen A, Watanabe T (2010) Perceptual learning rules based on reinforcers and attention. Trends Cogn Sci. 14:64-71.

Roesch MR, Olson CR (2003) Impact of expected reward on neuronal activity in prefrontal cortex, frontal and supplementary eye fields and premotor cortex. J Neurophysiol. 90:1766-89.

Roesch MR, Olson CR (2004) Neuronal activity related to reward value and motivation in primate frontal cortex. Science. 304:307-10.

Rossi AF, Bichot NP, Desimone R, Ungerleider LG (2007) Top down attentional deficits in macaques with lesions of lateral prefrontal cortex. J Neurosci. 27:11306-14.

Russel B, Torralba A, Murphy K, Freeman W (2008) LabelMe: A. database and web-based tool for image annotation. Int. J. Comput. Vis. 77:157–173.

Sabatinelli D, Bradley MM, Fitzsimmons JR, Lang PJ (2005) Parallel amygdala and inferotemporal activation reflect emotional intensity and fear relevance. Neuroimage. 24:1265-70.

Schmidt LJ, Belopolsky AV, Theeuwes J (2015a) Attentional capture by signals of threat. Cogn Emot. 29:687-94.

Schmidt LJ, Belopolsky AV, Theeuwes J (2015b) Potential threat attracts attention and interferes with voluntary saccades. Emotion. 15:329-38.

Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. Science. 275:1593-9.

Schultz W (2000) Multiple reward signals in the brain. Nat Rev Neurosci. 1:199-207.

Schultz W (2006) Behavioral theories and the neurophysiology of reward. Annu Rev Psychol. 57:87-115.

Seger CA (2008) How do the basal ganglia contribute to categorization? Their roles in generalization, response selection, and learning via feedback. Neurosci Biobehav Rev. 32:265-78.

Seger CA, Peterson EJ, Cincotta CM, Lopez-Paniagua D, Anderson CW (2010) Dissociating the contributions of independent corticostriatal systems to visual categorization learning through the use of reinforcement learning modeling and Granger causality modeling. Neuroimage. 50:644-56.

Seger CA (2013) The visual corticostriatal loop through the tail of the caudate: circuitry and function. Front Syst Neurosci. 7:104.

Seidl KN, Peelen MV, Kastner S (2012) Neural evidence for distracter suppression during visual search in real-world scenes. J Neurosci. 32:11812-9.

Simola J, Le Fevre K, Torniainen J, Baccino T (2015) Affective processing in natural scene viewing: valence and arousal interactions in eye-fixation-related potentials. Neuroimage. 106:21-33.

Small DM, Gitelman D, Simmons K, Bloise SM, Parrish T, Mesulam MM (2005) Monetary incentives enhance processing in brain regions mediating top-down control of attention. Cereb Cortex. 15:1855-65.

Sugrue LP, Corrado GS, Newsome WT (2004) Matching behavior and the representation of value in the parietal cortex. Science. 304:1782-7.

Sussman AB (2017) Valence in Context: Asymmetric Reactions to Realized Gains and Losses. J Exp Psychol Gen. [Epub ahead of print]

Taylor SE, Brown JD (1988) Illusion and well-being: a social psychological perspective on mental health. Psychol Bull. 103:193-210.

Taylor SF, Welsh RC, Wager TD, Phan KL, Fitzgerald KD, Gehring WJ (2004) A functional neuroimaging study of motivation and executive function. Neuroimage. 21:1045-54.

Thorndike EL (1911) Animal intelligence: Experimental studies (Macmillan).

Theeuwes J (1991) Cross-dimensional perceptual selectivity. Percept Psychophys. 50:184-93.

Thorpe S, Fize D, Marlot C (1996) Speed of processing in the human visual system. Nature. 381:520-2.

Treisman AM, Gelade G (1980) A feature-integration theory of attention. Cogn Psychol. 12:97-136.

Ullman S, Vidal-Naquet M, Sali E (2002) Visual features of intermediate complexity and their use in classification. Nat Neurosci. 5:682-7.

Van den Hout M, Tenney N, Huygens K, Merckelbach H, Kindt M (1995) Responding to subliminal threat cues is related to trait anxiety and emotional vulnerability: a successful replication of Macleod and Hagan (1992) Behav Res Ther. 33:451-4.

Vuilleumier P, Armony JL, Driver J, Dolan RJ (2001) Effects of attention and emotion on face processing in the human brain: an event-related fMRI study. Neuron. 30:829-41.

Vuilleumier P (2005) How brains beware: neural mechanisms of emotional attention. Trends Cogn Sci. 9:585-94.

Walther DB, Caddigan E, Fei-Fei L, Beck DM (2009) Natural scene categories revealed in distributed patterns of activity in the human brain. J Neurosci 29:10573–10581.

Wegner DM, Schneider DJ, Carter SR 3rd, White TL (1987) Paradoxical effects of thought suppression. J Pers Soc Psychol. 53:5-13.

Whitfield-Gabrieli S, Nieto-Castanon A (2012) Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks. Brain Connectivity. 2:125–141.

Williams SM, Goldman-Rakic PS (1998) Widespread origin of the primate mesofrontal dopamine system. Cereb Cortex. 8:321-45.

Wolfe JM, Butcher SJ, Lee C, Hyle M (2003) Changing your mind: on the contributions of top-down and bottom-up guidance in visual search for feature singletons. J Exp Psychol Hum Percept Perform. 29:483-502.

Wolfe JM, Alvarez GA, Rosenholtz R, Kuzmova YI, Sherman AM (2011) Visual search for arbitrary objects in real scenes. Atten Percept Psychophys. 73:1650-71.

Yamamoto S, Monosov IE, Yasuda M, Hikosaka O (2012) What and where information in the caudate tail guides saccades to visual objects. J Neurosci. 32:11005-16.

Yamamoto S, Kim HF, Hikosaka O (2013) Reward value-contingent changes of visual responses in the primate caudate tail associated with a visuomotor skill. J Neurosci. 33:11227-38.

Acknowledgments

First of all, I want to thank my advisor, Clayton, for having guided me through all this. He made me feel incredibly at ease ever since I arrived to a totally new place, doing new things, and in a different language from what I was used to. And it was so much fun having chats and insightful discussions together. I want to thank him and my co-advisor, Marius, for devoting their knowledge and their time to the development of this project. You both provided me with great teachings and skills, that I hope I could absorb to the fullest. Not to mention the lab(s') hikes, dinners and outings, which added fun to all of this!

Immediately after, I want to thank my first and irreplaceable office mate, Elisa. She has been a huge support during my first, hesitant steps in the field of research, as a personal unofficial, almost-full-time, matlab/statistics/cognitive psychology instructor. I miss the never-resting mind always ready to catch small details, the shoulder to lean on and the witty spirit with whom to share confidence and laughter.

I also want to thank people from my labs (Alisha, Christina, Daniel, Daria, Elisa, Elena, Matt, Talli, Timo, Poppy, Sushrut) and out of them (Raffaele, Scott, Nick) who all gave different and helpful contributions to this work.

I want to thank Farid for giving me joy and life throughout the whole second half of this Phd. You introduced me not only to a different and fascinating corner of the world, but also to a different side of me. You made my travelling skills grow remarkably, which turned out to be very useful throughout this year of conferences abroad! Finally, but not negligibly, the difference between valence and salience could have never been as clear in my mind without you… Mibusamet, Zanam.

Giulia, it's funny how destiny can work sometimes. And nice! I am so happy that our paths crossed again after such a long time. Thanks for having been so warm and welcoming throughout all this time, I enjoyed a lot our cozy dinners, life conversations and political discussions, also when they got to different political trends within the same party! Anyhow, fighting with you would be impossible for me, so that doesn't really matter!

Then, I want to thank Laura ♍ (third virgo in this space btw… do we have a pattern here?) for the time we have spent together, chats, drinks, and milk with mint in the middle of the night! What is Locos when you can be in Barattieri Road? At least just until 'i ne tacan! And thanks for the best ringtone ever, I am still thinking of switching back to my old mobile for that!

Laura ♌, the start of this path decreed the end of our physical closeness, but, in spite of that (or maybe because of that? ;)), you encouraged me to follow it and I appreciate it. Moreover, in spite of the distance, whenever we meet again it just takes a few minutes to feel just as if I had left yesterday. Thank you for being there.

Alessio, they say that light from the stars takes years to arrive to the ground. Sometimes even words do. So finally here are the lines which go out to my longest and farthest journey. I want to thank you for your empathy, your patience and your wisdom. I still treasure all of them just as then.

I also want to thank Irene ♉^{2} (and it could have been nothing but squared!), Irene ♌, Maria Carmela and Andrea. Each one of you has or had his own, very special place in my life. Thanks.

Last and of course not least, I want to thank my big and happy family. I mean all of it, from Archons Street, to Pelican Road (cats included? Mmm… don't know!), Lieutenant Panella Street, Very Religious the XIth Road, and Aurelia Valley Square! Jokes aside, I cannot say anything but how blessed I feel to have all of you close, although far. Thanks.