frontiers
IN HUMAN NEUROSCIENCE

Updating Mental Models: Insights from Playing Rock, Paper, Scissors with Stroke Patients

Britt Anderson and James Danckert

Journal Name:	Frontiers in Human Neuroscience
ISSN:	1662-5161
Article type:	Hypothesis & Theory Article
Received on:	09 Sep 2010
Frontiers website link:	www.frontiersin.org

Updating Mental Models: Insights from Playing Rock, Paper, Scissors with Stroke Patients

Britt Anderson[a,b,1], James Danckert[a,1]

[a] *Dept. of Psychology, University of Waterloo*

[b] *Centre for Theoretical Neuroscience, University of Waterloo*

Abstract

Individuals with focal brain injury often display disorders of updating. For example, in spatial neglect, patients have a bias to begin visual searches on their right, and moreover, even after engaging in a prolonged, fruitless search, neglect patients persist in their repetitive, unproductive strategy. This is one example of an updating impairment and updating is at the core of our ability to navigate a complex and changing world; we need to recognize static circumstances so that we can employ efficient, automatic responses; we need to recognize change so that we can adapt; and we need to be able to do all this without an explicit reward signal. In this paper, we develop the concept of updating as a core part of the process of building and using mental models. We elaborate on the cerebral structures (inferior parietal lobe, insula, and striatum) that seem to be critical for normal updating performance, and we use these ideas to suggest what sort of tasks would be good for assessing updating ability. Lastly, we elaborate on how our use of the word updating fits with other related concepts.

Keywords: Games, Probability, Brain, Parietal Lobe, Neglect, Updating, Representation, Model

---

1. Introduction

The children's game Rock, Paper, Scissors (RPS) is simple, enduring, and popular. Each player makes a gesture of a fist (rock), two fingers (scissors), or their palm (paper) and the winner is rock > scissors > paper > rock. Versions of RPS exist in many cultures, websites present strategies (Walker, 2004), and tournaments for cash prizes take place (Maki, 2010). Despite its simplicity, RPS produces complicated dynamics (Sato et al., 2002) and relies on a number of cognitive processes. The game requires performing a sequenced motor act under temporal constraints, requires vision to see an opponent's gesture, and requires the application of a rule to determine the result of a round of play. For these reasons, RPS has been used to investigate motor systems (Shimada and Abe, 2010), set switching (Matsubara et al., 2004), and transverse patterning (Leirer et al., 2010).

However, there are other aspects of playing RPS that are strategic and require a model of the opponent's play. A strategy of selecting uniformly and randomly from each of the three options guarantees a player 1/3 each of wins, draws, and losses. If a player is engaged in repetitive play against an opponent, and wishes to improve his proportion

---

[1] Supported by a grants from the Heart and Stroke Fund (Canada), the Canadian Institute for Health Research, the University of Waterloo, and NSERC CRC program (JD)

of wins beyond 1/3, he must takes advantage of any regularities or biases in his opponent's play. The player's only evidence of these biases and patterns is the history of play itself. However, when a player deviates from a uniform, random strategy, he exposes himself to risk and must be alert to detect any change in the opponent's play that would signal a change in the opponent's strategy. All this estimation needs to take place on a noisy background where any individual outcome and opponent action could happen by chance one third of the time.

These aspects of the RPS game, which are distinct from the processes of gesture production and result evaluation, capture a concept we call updating. These aspects of RPS make the game a good empirical tool for probing the cognitive processes and neural structures that support updating, and we return to the practical application of RPS as a research tool below. First, we wish to develop at an abstract level the concept of updating and how it fits into our competing needs for behavioral stability and behavioral flexibility in a complex world. After outlining our model of updating, we give an example of an updating impairment. We focus on spatial neglect since it is the condition that led us to consider this concept in the first place. Third, we discuss, based on lesion data from the neglect syndrome patients, what might be the critical brain structures for updating. Fourth, we present examples of empirical tasks and data that demonstrate how one might test the updating function. Lastly, we place our ideas in context by considering other, related concepts.

2. The Concept of Updating and the Need for Mental Models

The metaphor of the brain as a builder of mental models is an old, persistent, and fecund one in psychology and cognitive neuroscience (Craik, 1952; Churchland and Churchland, 2002; Friston, 2009). Mental models represent an effort to instantiate the rules that govern the external world, in analogy to an orrery and the solar system. Orreys are not made randomly, but are built to careful specifications. The specifications are computed based on careful observations in combination with astronomical models. As a result, an orrery can be used to make predictions about the future positions of the planets. Those predictions will reflect both the data gathered and the model used. Ptolemaic and Copernican conceptions will give rise to very different machines, and, potentially, very different predictions. Just the same, our predictions for future states of the world will reflect our mental model. But our world is a variable place and our capacities are limited. Circumstances and goals change. How will we know when our current mental model is inappropriate or incorrect? Updating as we define it refers to the processes of recognition and revision. First, recognizing that our mental model no longer fits our observational data and second constructing a new mental model that better fits our new observations. Gathering the sensory evidence that supports such a recognition, and engaging in the motor behavior that reflects model revision are other components of the global system, but they do not reflect updating per se.

The sensory data upon which mental models are built have a very definite character: they are probabilistic. Our estimates of the world reflect repeated experience. This is consistent with the abundance of recent research that has demonstrated our sensitivity to the statistical structure of our environments. For examples, animals and humans reflect reward probability in their choice probability (Vulkan, 2000). We estimate the prior likelihood of events and combine that with sensory evidence. Such considerations can account for data from the fields of sensori-motor learning (Körding and Wolpert, 2004), to the metrics of inhibition of return (Farrell et al., 2010), and even to the statistics of neural firing (Beck et al., 2008). Why isn't a completely data driven, model-free method sufficient for capturing the world's regularities? Why should we need or benefit from mental models?

The principal reason is that the computational requirements of model-free mechanisms make them a poor match to our most pressing needs. Model-free probability processes are iterative and accumulative. They are slow and computationally demanding. Where, but in a Skinner box or a casino, would one face the consistent repeatable experiences needed to fully shape behavior? We are frequently required to make decisions and to act in situations that are unique. These situations may resemble circumstances we have faced before, but neither we nor anyone else may have faced exactly such a situation before, and how many examples are we likely to get? Once we make our decision to act, the action itself frequently makes it impossible that we will ever face just such a situation again. This is not mysticism. If I play a certain choice in RPS, you will see me, and the result will be known to us both. I cannot repeat an identical play a few tens of times to estimate the probability for each of the actions you are able to play. My repetition would alter exactly the probabilities I want to estimate.

The reason we need to build mental models is because our brains are puny and our lives are short. Similar considerations drove the early statisticians to adopt model-based statistical tests. R. A. Fisher advocated for parametric statistical tests not because he could not appreciate the asymptotian ideal of non-parametric statistics, but because he recognized that their computational demands made them impractical. Similarly, refined by time, our brains may have resolved to the same approach: assume a model for the data and act under that assumption; update your model only when you reject the null hypothesis that random variation of the variety specified by your model explains the variation in your sensory data. In this way we can still use and rely on experience to guide us, but we make simplifications. This may, at times, lead us to inaccuracies or to adopt sub-optimal strategies, but we can still use the iterative, Bayesian idea in those circumstances where similar events are frequently repeated; unfortunately these circumstances are the exception. Gradualist characterizations of learning are insufficient to explain discontinuities in problem solving behavior (Köhler, 1925/1999). Our own "Aha" experiences are evidence of the sudden switch in perspective that are more easily reconciled with the updating of mental models (Tolman and Honzik, 1930; Epstein et al., 1984) than the gradual modifications brought about by reinforcement learning.

Figure 1 is a schematic outline of our updating model. The world evolves according to unknown rules and with a history dependence potentially extending all the way back to $t = 0$. Our knowledge of states of the world is indirect and incomplete. Our observations are filtered through a tunable sensory apparatus. As an obvious example, what we observe depends on where we look and where we look depends on where we think we are likely to see something important. Our observations are constrained though by what we sample (e.g. we cannot see into the infrared) and how much we sample. The tuning component of our model allows us to prioritize categories of input.

Our schema does include a role for accumulated observations in building a probabilistic representation of our environment, but we use a bell-curve (Figure 1) to symbolize that we are estimating parameters for a model. We are not building a non-parametric distribution. The uses of such a mental model are three fold. First, it determines sensory tuning ("look here", "listen for that", "pay attention to the red 'X's' "); second, it guides action selection by weighting the likely consequences of possible actions; and third, and most critically to this paper, it provides the metric for checking if the model is good enough or if it needs updating? Do our observations match our expectations? Are we surprised by what we see? If not, then the current model is good enough. If so, then a model revision is required.

In Figure 1 we depict that surprising data can be built up in different ways: first, we might see an extremely unlikely event (*a* in the Figure) that is individually so improbable as to provoke us to update our model, or second, we might see repeated improbable events (*b* in the Figure) that in the aggregate, but not individually, are sufficiently, collectively, unlikely that we are surprised and report a mismatch between our model and our observations.

What happens when we are surprised by data or detect a mismatch between expectations and observations? We suggest that there are two functionally distinct responses. First, we can make subtle refinements in our expectations. This is analogous to moving the mean of our bell-shaped curve a bit to the left or right; we may have, for this example, assumed that a normal distribution is the right form for our model, but that we need to update our estimates of the mean and variance parameters. A second response to surprising data may be the generation of a completely new model; rather than fine tune parameters we abandon the normal distribution in favor of an exponential (depicted in Figure 1 as point 2 of the mental model).

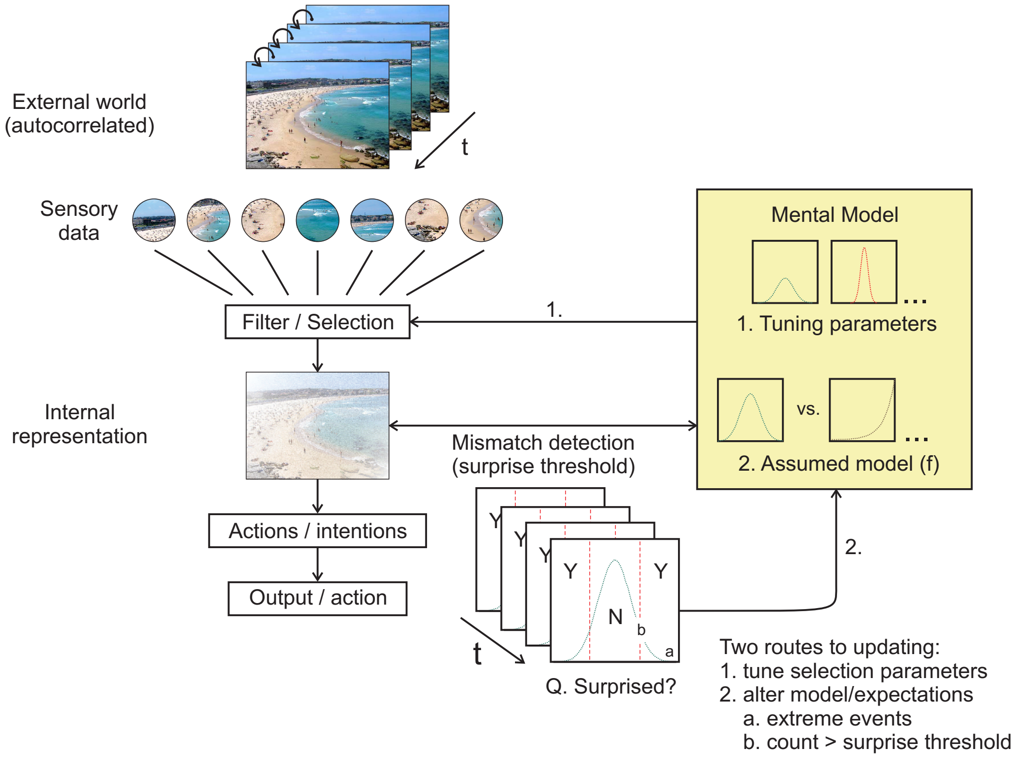

Figure 1: Our model of updating is schematic. It includes a representation of the world that is recursive and obeys unknown laws (the beach scene upper left). Our view of the world is built up from incomplete, filtered snapshots. We use this mental model to plan our actions, to direct our sensory sampling, and as the substrate for self-diagnostics. How do we know if our mental model is adequate? We see if the sensory data coming in is explained well by our current model. We presume that our model has a probabilistic representation (*symbolized* by the bell-shaped curve) that allows us to test how surprising incoming sensory data is. As explained in the text, a mismatch between observations and models can be triggered by individual extreme observations (letter *a* lower right), or an accumulation of improbable events (letter *b* lower right). If a mismatch is detected, different degrees of model revision can take place: either a recalibration or substitution. Damage to different model components will result in different deficits, and though not depicted here, these different elements are likely to be associated with different structural brain regions.

Our model is highly schematic and overlaps with other cognitive processes. For example, we suggest that one signal for mental model revision is an accumulation of data that is, in the aggregate, surprising. This process will depend on working and recent memory. Also, we state that mental model updating can be either a revision or a replacement. The former might show considerable overlap with systems supporting reinforcement learning. Below, we discuss further the areas of overlap and uniqueness, but first we would like to make our notion as concrete as possible. To do so, we will retrace the steps that led us to the idea of updating, that is, as an explanation for impairments seen in people with brain injury. This example also makes it easier to see how updating is similar to and different from other related cognitive processes.

3. Neglect — A Disorder of Updating?

It was our puzzling over the behavior of patients with the neglect syndrome that led us to our concept of an updating impairment. In this section, we briefly introduce the classic neglect findings and more recent experimental results of non-spatial impairments. We then consider the structures associated with neglect and highlight two that seem particularly relevant to the function of mental model updating.

The clinical syndrome of neglect has been recognized for more than 100 years (Langer, 2009) and two features are particularly striking: the lack of concern and awareness many patients display and the pronounced spatial modulation of their difficulties. The syndrome has been extensively studied and its basic features are well established (for two recent reviews, see Danckert and Ferber, 2006; Adair and Barrett, 2008). The brain lesions producing neglect are usually right sided with the behavioral impairment maximal on the left. Requested to bisect a line down the middle, neglect patients err to the right. Shown a collection of short lines scattered across a page, and asked to cross-out each one, the patient will concentrate on the right sided items to the exclusion of others, often revisiting and re-marking perseveratively previously cancelled items (Na et al., 1999). Even at a representational level, neglect patients are impaired, often describing only the left side of imagined locations (Marshall and Halligan, 2002) and functionally they may fail to shave the left side of their face or eat food on the left side of their plate (Danckert and Ferber, 2006). These findings naturally lead to the traditional emphasis on the spatial aspects of the syndrome. More recently we, and others, have been struck by the non-spatial impairments in neglect (e.g. Husain and Rorden, 2003).

The increased duration of the attentional blink seen with neglect is one robust example of a non-spatial deficit (Husain et al., 1997). In a typical attentional blink task, the participant is shown a series of centrally presented targets and distractors with instructions to report the targets. The attentional blink is a normal phenomenon that denotes the decreased probability of accurately reporting a second target when it appears shortly after a first, successful, target detection (Shapiro et al., 1997). In neglect, the attentional blink is substantially prolonged even for centrally presented stimuli that, individually, are detected normally. The loci of injury associated with a prolonged attentional blink have been the inferior parietal, superior temporal lobe, and periopercular structures, and with damage to these structures a prolonged attentional blink may occur in the absence of clinical neglect (Rizzo et al., 2001; Shapiro et al., 2002).

A prolonged attentional blink implies slow or inefficient sampling in a region of space where basic detection is normal. Other disorders of temporal processing that are distinctly non-spatial have also been found in neglect. Danckert et al. (2007) studied eight participants with clinical neglect, six right brain damaged controls without neglect and eight controls. All participants estimated intervals of five, fifteen, thirty and sixty seconds. While participants in all groups tended to underestimate all intervals, the performance for the participants with neglect was particularly striking, with a mean estimate of the sixty second interval of 8.89 seconds (SD = 3.18) while it was 49.26 (SD = 7.99) for the neurologically healthy controls. This impairment of temporal estimation is not vision specific as auditory temporal estimation is also impaired neglect (Merrifield et al., 2010).

If updating involves accumulating sensory data and comparing estimates of its likelihood to an estimate of an internal model, working memory would also seem to be an important functional component, and again, a working memory impairment is seen in participants with the neglect syndrome. Ferber and Danckert (2006) assessed spatial working memory in four neglect participants by presenting three squares on the participants' right (good) side and then, after a three second delay, probing with a circle. Participants were required to report if the circle appeared in a location previously occupied by a square. A control verbal working memory task used sequences of single digit numerals and a single test digit. The neglect participants were dramatically impaired compared to right brain damaged controls without neglect and normal older healthy controls. A spatial working memory impairment also accounts for the frequent repeated fixations of right sided target locations that patients with neglect make during visual searches (Husain et al., 2001).

These spatial working memory impairments do not preclude the learning of environmental contingencies (e.g., we know that neglect subjects show both color and location repetition priming in right *and* left space (Kristjánsson et al., 2005), however these non-spatial deficits render neglect patients less sensitive to the statistical structure of their environments. Shaqiri and Anderson (2010) tested three subjects with current or recent spatial neglect on a simple visual discrimination task that required pressing one of two buttons to indicate the color of a small circle displayed on a computer screen. Two areas of the screen, one left and one right, were more likely to be target locations than the rest of the screen and the left sided “hotspot” was three times as likely to be the target location when compared to the similarly sized right sided region. While control subjects improved their speed to classify targets from the left sided “hotspot,” the neglect subjects did not, although the neglect subjects did show some improvement for the ipsilateral, right sided, higher probability region.

At this point we can return to consider a fundamental question for neglect: what is it that makes a person who has spatial neglect so much more impaired than a person with only a left hemianopia? Both may begin their searches on the right, but the neglect patient never revises his estimate of where the object may be; he does not update his search to consider other regions, other locations. He is is persistently in error. Certainly it is disadvantageous to have a right sided bias for orienting and responding, but if, failing to find what was sought, the patient with neglect updated his search accordingly there would not be the same severe functional consequences. Spatial impairments may drive early orienting responses, but when those reactions are wrong it is the inability to benefit from this failure that makes neglect so disabling. It is the inability to take a proper account of when things happened (temporal impairment, prolonged attentional blink) and what has happened recently in what locations (spatial memory and environmental statistics). As a result, even when one can change motor behavior one may not change performance. A picturesque example is Ferber et al. (2003). Ferber et al. (2003) had patients with neglect view two vertically oriented chimeric pictures. In each image, one half of the picture was a smiling face and the other half was emotionally neutral. One image in each pair had the smiling half on the right and the other on the left. Their location, top or bottom, was random. Subjects reported which picture seemed happier. Neglect subjects routinely only fixated the right half of the pictures and always selected the image with the right half smiling as the happier picture. Next, neglect subjects underwent prism adaptation after which they were shown to fixate both left and right sides of the images. Thus, their ocular exploratory behavior was changed, but there was no change in the behavioral report. They still were almost uniform in designating as happier that image with the right half smiling.

In trying to reconcile these findings, we have come to regard the additional source of failure seen in many neglect subjects, that which separates them from an hemianopic, as an updating impairment. Either they fail to take account of new information indicating the need to revise an internal mental model or they cannot use their recognition to replace the current model. If these deductions are accurate, the locations of brain lesions that cause neglect may help us identify the brain structures important for such updating functions.

The right inferior parietal lobe has been the classic site of injury associated with clinical neglect (Vallar and Perani, 1987). But, as in many others areas of clinical neuroanatomy, the lesions that yield neglect have become better understood with MRI scanning. Subcortical, temporal lobe, and parietal lobe structures are all strongly linked to clinical neglect.

Newer studies confirm a prominent role for parietal lobe damage in neglect (Mort et al., 2003; Verdon et al., 2010), and thereby indicate a potential role for the parietal lobe in updating. Supportive evidence for a role of the parietal lobe in updating is its connectivity. The parietal lobe receives both sensory and visual afferents (Andersen et al., 1990; Lewis and Van Essen, 2000), Functionally there is parietal lobe activity during many memory tasks (Cabeza et al., 2008), and evidence that the parietal lobe signals an expectancy for upcoming sensory data (O’Connor et al., 2010)

In addition to the parietal lobe, other structures have been implicated in causing neglect. The superior temporal lobe has been emphasized by Karnath et al. (2001) though this finding has produced some controversy. Mort et al. (2003) emphasizes that Karnath et al. (2001) excluded subjects with hemianopia, in an effort to define a purer clinical population, and thereby may have induced biases. When Mort et al. (2003) looked at their clinical cohort, without this selection criterion, they found that in about half of their cases of middle cerebral artery territory strokes producing neglect there was no lesion of the superior temporal gyrus.

In addition, most studies on lesion locations in neglect find damage to subcortical structures. Karnath et al. (2004) reviewed the lesions in 140 consecutive patients with right hemisphere strokes and used a voxel-wise statistical analysis to look for regions of the right hemisphere statistically associated with the neglect syndrome; he confirmed the involvement of the superior temporal gyrus, and highlighted subcortical areas including the insula and the striatum.

While the insula has not traditionally been emphasized as a causative lesion in the neglect syndrome, it is frequently damaged by the strokes that cause neglect (Karnath et al., 2002, 2001; Mort et al., 2003; Ferber and Danckert, 2006; Danckert et al., 2007). Anterior insular involvement is especially prevalent in the minority of neglect cases where the functional impairment is chronic (Doricchi and Tomaiuolo, 2003; Bartolomeo et al., 2007). These findings prompt consideration of the insula as a component of an updating network. Structurally, the insula denotes the neocortex hidden (insulated) by the frontal, temporal and parietal opercula. The human insula possesses components that have no clear correlate in the monkey and the anterior insula also possesses a variety of neuron, the von Economo neuron, that has been found almost exclusively in higher primates (Allman et al., 2010). For many years the main functions of the insula were held to be primarily interoceptive (Penfield and Faulk Jr, 1955), but more recently a role for the insula in human awareness has been suggested (Critchley et al., 2004; Craig, 2009). The role of the insula in salient target detection (Corbetta and Shulman, 2002) fits within our updating schema, and other research has shown that the anterior insula is part of the system activated by the implementation of task sets (Dosenbach et al., 2006).

Connecting the roles played by the right inferior parietal lobe and the anterior insula in updating is a combined functional and DTI study (Umarova et al., 2010). Umarova et al. (2010) used a Posner type cuing task to develop seeds for a structural white matter connectivity study. Umarova et al. (2010) reported that the anterior insula and parietal lobes were co-activated by the attentional tasks and linked by subcortical white matter tracts. Damage to the same subcortical white matter pathways alone may cause neglect (Karnath et al., 2009), and neglect itself may arise as a disconnection syndrome (He et al., 2007; Bartolomeo et al., 2007).

The anterior insula is near ventral striatal structures and these structures, too, are often injured in the strokes that cause neglect. They too deserve consideration as having a role in an updating impairment, especially as the striatum (e.g. Balleine et al., 2007; Pennartz et al., 2009) has been repeatedly linked to the broad domain of reward computation and decision making (Sutton and Barto, 1998).

The prevalent version of this model has a prediction error being signaled by dopaminergic neurons that project from the midbrain to the striatum, with the striatal neurons representing an association of actions and states. The conventional role for the striatum is therefore one of gradual, incremental change (Johnson et al., 2007). While the role of striatal structures in procedural learning seems well established, the role of the striatum in planning processes is less certain. Recently, van der Meer and Redish (2010) reported that ventral striatal neurons recorded from rats running a T-maze increased their firing at the final choice point in the maze, consistent with a striatal role in planning the action that will lead to reward. Using ensembles of neurons the authors were able to show a representation of reward at the final choice point. Thus, there is good evidence for striatal structures in decision making, associating actions, states, and reward, and therefore striatal injury could also be important for the updating impairment seen in subjects with spatial neglect.

4. Assessments of Updating

The model outlined in Figure 1 has several sub-components that cooperate to implement a global updater. The updating process needs to signal both when environmental statistics do not match mental model predictions (updating is required), and when environmental statistics match the current model (no model updating is required). A good updating task should reveal these events, be simple enough for brain damaged participants, and flexible enough to permit combining with other investigatory methods, such as fMRI.

There is a large literature demonstrating that people and animals can learn probability distributions (Estes, 1972) as inferred by changes to patterns of reward. Modern versions of tasks frequently used to study decision making, such as the multiarmed bandit problem (Zhang and Lee, in press), are basically human Skinner boxes. Experimental subjects make repeated choices from a finite set of selections where feedback is provided on each trial and the relationship between choice and reward is typically probabilistic. Although solving such a task requires, in some sense, a model of the world, (i.e. is the relations between states and actions), it seems impoverished in a few important ways. First, is that it presumes that experience will provide a large number of stereotyped trials, however, the world is frequently not so kind. Two, history is only important in the aggregate; it doesn't matter exactly how I came to my present state (Markov assumption), only that I am here and the number of times I have been in State S_i, taken Action A_j, and received Reward R_k. This does not match common experience. Third, and most importantly, there is always a clear, explicit reward, but frequently we learn about our world in the absence of reward (or via implicit judgments of value).

It is worth reminding ourselves that learning does not require reinforcement (Tolman, 1948). Reinforcement is the method we use to demonstrate that learning has occurred. Rats allowed to explore and become familiar with an empty maze seem no different from other rats, but when reward (food) becomes available they reveal their knowledge. Our concept of a mental model is like Tolman's map; it is an abstraction that productively summarizes observable behavioral results. Much of what we learn about the world in our daily life seems more like the period of a rat's purposeless meandering in a maze than it does the timed driven race to a food reward or massed sequential pulls on one of two slot machines. Thus, while reinforcement learning might be a subset of what we have described as updating, it does not subsume it.

A good task for assessing updating mechanisms needs to demonstrate the acquisition of new knowledge in response to changing environmental contingencies. It also needs to be simple enough to be used in subjects with updating impairments.

Demonstrating this approach, Geng and Behrmann (Geng and Behrmann, 2002, 2006, 2005) employed probability cuing, by altering slightly the structure of a conventional cuing task. In a conventional cuing task, a cue appears at the center of a computer screen and indicates some aspect of the subsequent target, usually where it is likely to appear. Subjects are quicker to report the target when it appears in the cued location (valid trials) then when it appears in another location. What Geng and Behrmann (2002, 2006, 2005) reported was that probability could serve as the cue. Subjects simply reported the presence of a target that appeared on a computer screen. There was no ambiguity about action selection nor a variation in success. The task itself is trivial for normal participants (the authors also tested subjects with brain damage; for this population success is not assured). When certain locations were selected to be more likely as target locations, subjects reported targets at those locations more quickly. Others have also used this approach, Walthew and Gilchrist (2006) have found similar results, though Walthew and Gilchrist (2006) have argued that the learning of the statistical relationships of positions with targets is the result of short time frame repeats (e.g. the last 4 trials). We have also looked at the ability for environmental statistics to influence behavioral performance in a similarly simple visual discrimination task (Druker and Anderson, 2010, and Figure 2). We too find that subjects adapt to the statistical structure used to select target locations, and thus respond more quickly and no less accurately.

We also undertook to use this type of task in brain damaged patients (Shaqiri and Anderson, 2010), the types of patients we asserted above have an updating deficit. We confirmed that older control participants could learn to

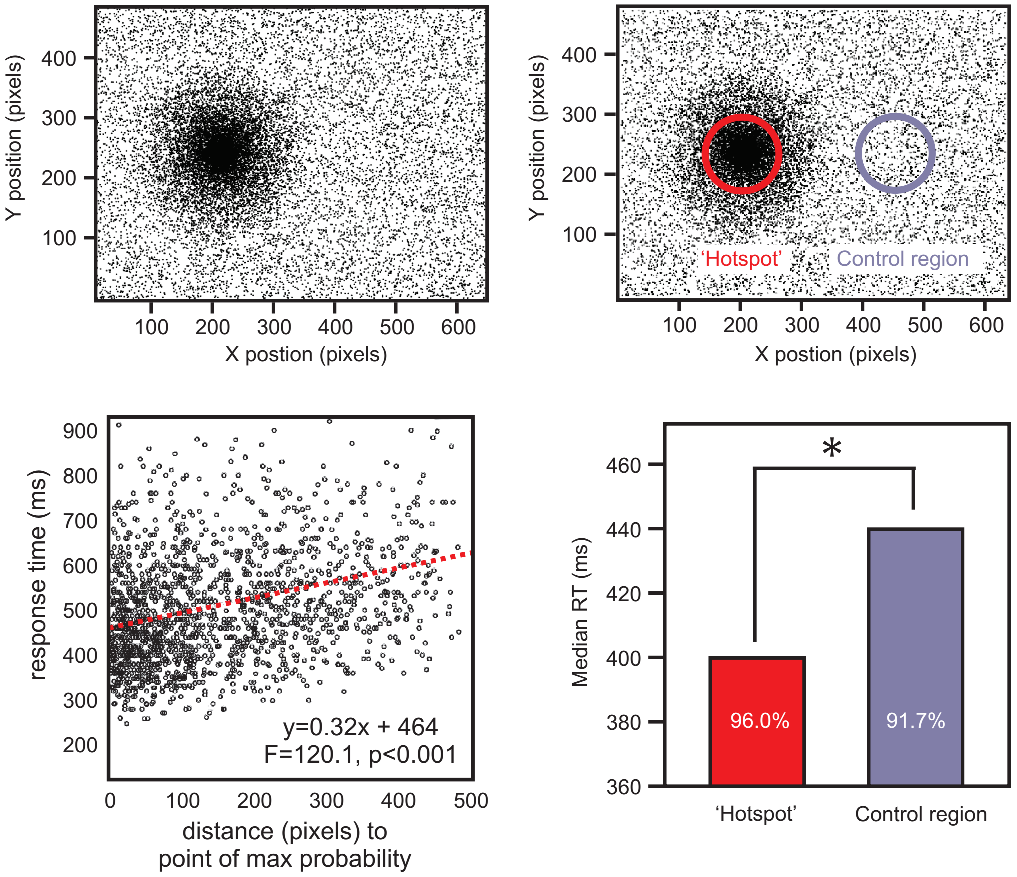

Figure 2: Data from the "hotspot" experiment of Druker and Anderson (2010) is presented. In this task, participants made a simple color classification: one button for a red spot and another button for a green spot. Unknown to the participants was the fact that the probability for where a target could appear was not uniform across the screen. There was a point of maximal probability, the "hotspot" which declined steeply, but continuously to a baseline value (see Druker and Anderson, 2010, for full methodological details). The consequences of using this distribution for selecting target positions is shown in the upper left where the target positions for all participants across all trials is collapsed. There is a central point of highest density and then target locations thin out. For analyses it is possible to look at the distance to this high probability "hotspot" or to define screen regions equidistant from fixation for the selection of subsets of trials in high and low probability regions (upper right panel). Showing that the effect of the probability manipulation is graded with probability, data from one participant is plotted in the lower left panel. The RT for each correct trial for this subject is plotted on the y axis and the distance to the point of maximal probability is plotted on the x axis. There is a clear relation between the distance and response time measures, highlighted here by plotting the linear regression line. In addition, across participants, there is a benefit in terms of both response time and accuracy of classification for trials when targets appear in the high probability region of the screen are compared with trials in the low probability region (lower right panel). These results are consistent with participants having a representation of the probability of where targets will appear and weighting the perceptual evidence accordingly.

use the probability of target locations to aid classification and we also showed that this ability was impaired in participants with a history of right hemisphere injury and neglect and that this probability cuing impairment was spatially modulated (Figure 3 and (Shaqiri and Anderson, 2010).

Many other research groups have also developed tasks that manipulate attention via probability effects. (e.g. Lambert, 1987; Hoffmann and Kunde, 1999; Ciaramitaro et al., 2001; Carreiro et al., 2003; Liston and Stone, 2008). As examples Farrell et al. (2010) reported that manipulating the statistics of where participants will be directed to look, influences the saccadic dynamics of an eye movement analogue of inhibition of return (Klein, 2000) and Eckstein et al. (2006) used a visual search task to demonstrate that eye movements were consistent with a bias towards where

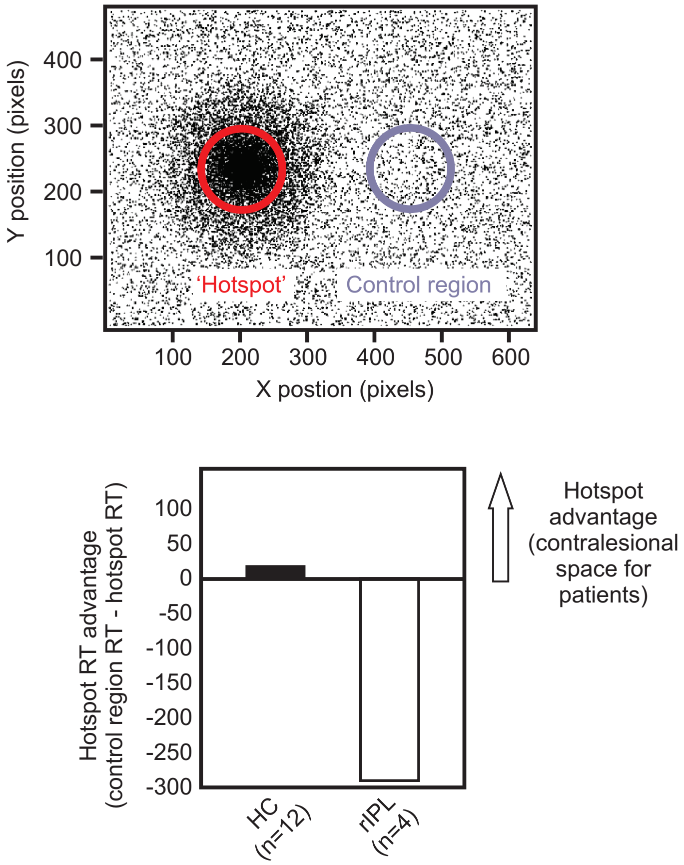

Figure 3: Sixteen subjects where given a modified version of the "hotspot" task illustrated in Figure 2. Four of the participants had a history of parietal lobe injury. When comparing trials where the target appeared in the high probability region of space with a similarly sized region equidistant from fixation on the low probability side (upper panel), we demonstrated a significant reaction time cost for targets in the high probability location for the patients compared to a slight benefit for the twelve healthy control participants (these data were collected in a paradigm similar to (Shaqiri and Anderson, 2010), but have not been previously presented).

we expect things to be based on experience (e.g. chimneys belong on houses).These data fit into our updating schema by showing that environmental statistics are reflected in performance, perhaps through the "tuning" of the sensory apparatus.

Another tactic for investigating updating is to use a simple game. Games have a long history in psychological research. Examples of games include bandit problems (Zhang and Lee, in press) and choice paradigms (Brown and Steyvers, 2009). One particular game popular in cognitive neuroscience research has been RPS. RPS has simple rules, but can achieve complex, chaotic dynamics (Sato et al., 2002; Salvetti et al., 2007). RPS has been used in a variety of functional imaging studies, but typically with the idea of studying gestures, set switching, and transverse patterning (Dinstein et al., 2008; Matsubara et al., 2004; Kadota et al., 2009; Leirer et al., 2010).

Paulus and colleagues have reported two fMRI studies where subjects played RPS and which are directly relevant for the issues being discussed (Paulus et al., 2004, 2005). Both studies used the same basic task: subjects were allowed to choose their response by pressing one of three buttons. Subsequently, they received an auditory report of the outcome and their choice and the computer's choice were displayed visually. A running counter of wins and losses was also presented incrementing, decrementing, or remaining the same depending on the outcome. Subjects were not informed that there was a preferred, neutral, and anti-preferred action for blocks of trials (16 or 20 trials long). If the subject chose the preferred action, he would win 90% of the time or lose 90% of the time if choosing the anti-preferred. Success was 50% with the neutral choice. The two studies differed slightly in the length of blocks, the timing of feedback, and the total number of trials (96 or 120). Paulus et al. (2004) was able to observe an effect on subject behavior. Subject choices of preferred stimuli increased in the second halves of blocks to 44% and the choices of anti-preferred actions decreased. Correlating with the acquisition of the preferred choice was an increase in activity in prefrontal cortex. Paulus et al. (2005) was not able to show any behavioral acquisition of the preferred/anti-preferred choices. The authors attributed this to fewer trials overall and a slower pacing resulting in more delayed feedback. However, when looking at their data based on whether the phase of the block required assessment/action selection or outcome evaluation, the researchers found activation in the anterior insula and superior temporal gyrus. For outcome evaluation there was greater activation in the inferior parietal and superior frontal regions.

In addition, to these human based functional studies, neurophysiological responses have been recorded in monkeys playing similar games, e.g matching pennies (Lee et al., 2004, 2005; Lee and Seo, 2007; Seo et al., 2009).

To make a version of RPS accessible to patients with brain damage we simplified the task demands. Instead of making gestures, subjects pushed a button, or stated their choice verbally (and a technician pushed the button). Blocks were long (200 trials) and self-paced, with choices displayed as pictures. The task did not require subjects to adopt a non-traditional strategy (e.g. trying to lose) or to switch sets. No explicit feedback was given. Without informing participants, we varied the computer's strategy from selecting a choice uniformly (trials 1 - 200) to one favoring one of the options 50% of the time (trials 201 - 400), and then to a more obvious favoritism where the computer selected one particular option 80% of the time (trials 401 - 600; in each case of favoritism the other two choices were equally likely).

The result was a heterogeneity of performance with a rather clear partition between the normal control subjects and the subjects with focal right hemisphere injury. Figure 4 provides an example of the diversity of performance that we found. Normal controls subjects chose randomly, with some detecting the mild shift in computer strategy. They typically detected the extreme shift in computer strategy within a few trials. Brain damage subjects revealed very different performances. Some brain damaged subjects perseverated on particular choices. Sometimes this led them to maximize their wins, their performance by this criterion exceeding the normal controls, and at other times it resulted in substantially fewer wins than following a uniform random strategy. Many subjects seemed to persist with a uniform random strategy without variation (Fig 4D).

The variety of responses provoked by the shifts in the computer's strategy reveal the richness of this simple approach. At present we cannot say if these examples represent points on a continuum of impairment or are examples of subtypes of impairment. In our experience, it is rare to see control subjects adopt the maximizing strategy (Vulkan, 2000) that we have seen in participants with frontal injury.

The change in the control subjects' choice probabilities when the computer uses a biased strategy, coupled with a failure of the control subjects to adopt a maximizing strategy, suggests that the control subjects have a model for their computer opponent that assumes more frequent variations in strategy than was in fact the case. This suggests that our control participants are trying to adapt their play to what is, in fact, noise.

To examine this idea, and to gain another perspective on the flexibility of participant choice, we plot the empirical entropy of the computer's selections over short blocks of trials (Figure 5). Even though there is a fixed entropy associated with each computer strategy, selecting small blocks of trials, (e.g., 20) will yield different empirical estimates of the entropy. We also did the same computation for the sequences of choices made by the brain damaged participants. Entropy is an information theoretic concept which we use here as a proxy measure for randomness. The greater the entropy the more random a sequence; lower entropy implies more structure (e.g. a series of 20 rocks in a

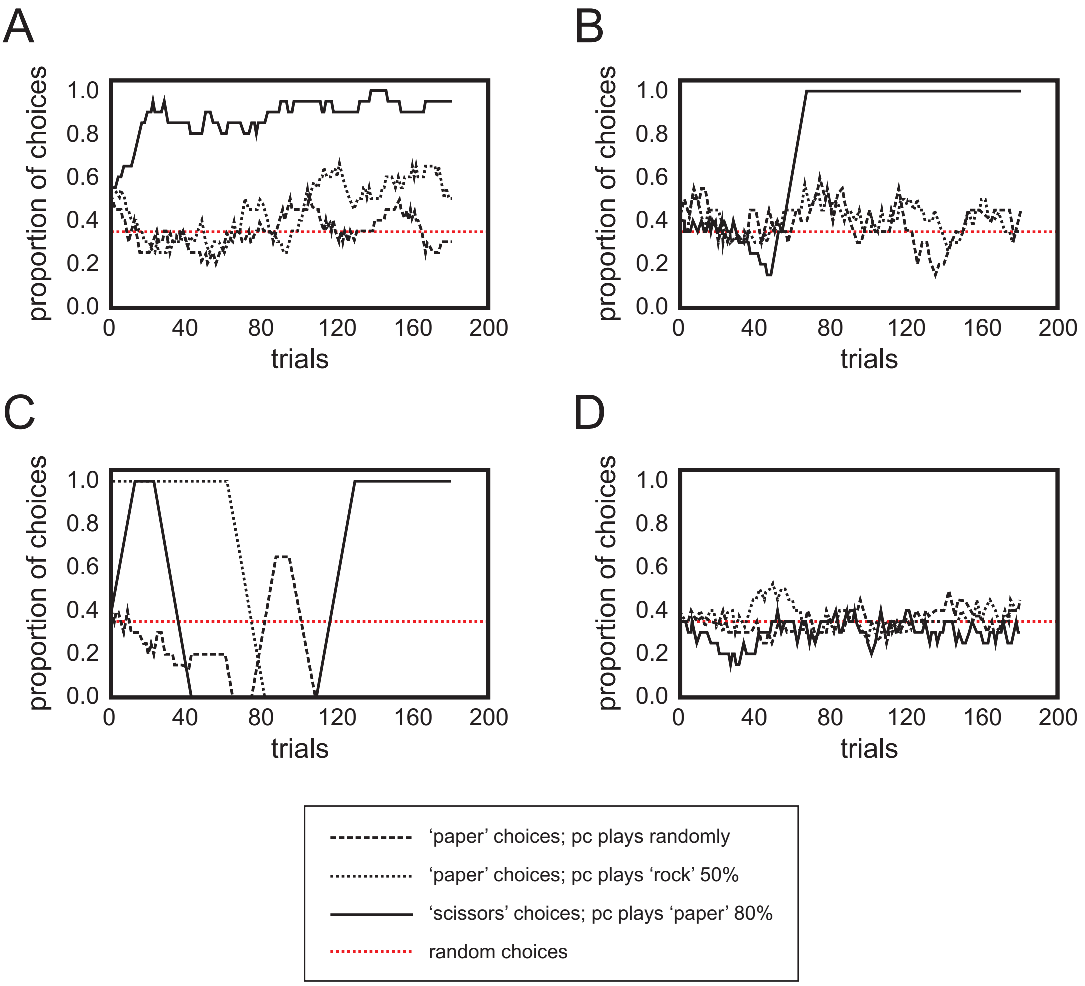

Figure 4: RPS performance for one control and three patients to demonstrate the heterogeneity of performance that can be seen with a fairly simple paradigm. Proportion of responses is computed from a sliding window of twenty trials. The control participant (upper left; A) shows three principle results. First, when the computer selects uniformly so does the control. Second, when the computer begins to pick one item 50% of the time, the control participant takes over a 100 trials to begin responding to this shift. Many control participants never alter their play when the computer chooses one option 50% of the time. Third, when the computer's choices are more extreme, the control participant rapidly exploits this strategy, within 10 - 20 trials. It is also typical that the control participants demonstrate a matching behavior rather than a maximizing strategy (Vulkan, 2000). Three different patient participants are shown to highlight the variety of abnormal performance that can be observed. First, one patient (upper right; B) uses a maximizing strategy. This leads to this participant winning more times than the controls over the last 100 trials of the 80% condition. However, as is typical, this participant takes longer to discover this strategy (40 - 50 trials). Another participant (lower left; C) shows a tendency, similar to the participant in panel B, to select one item 100% of the time, but this participant's selections are largely unrelated to the computer's strategy and do not result in an improvement in the number of wins. Lastly, the most typical pattern seen in the impaired participants, is that they select their choices randomly, and never seem to change their selections as a result of the computer's choices, even when the computer may choose the same item for many consecutive trials.

row is not "random" at all when compared to one that is a mixture of 7 rocks, 7 scissors and 6 papers). What we can see in Figure 5 is that the brain damaged participants do not show any real correlation between the entropy of their choice sequences and those of the computer (panels B, C, and D). However, the entropy of the control participant in panel A, tracks rather closely the entropy of the computer's sequence, implying that the control participant may be trying to adapt his choices on a shorter time scale than we, as the task designers, know is relevant. These data might provide a method for probing the time scale over which updating is taking place.

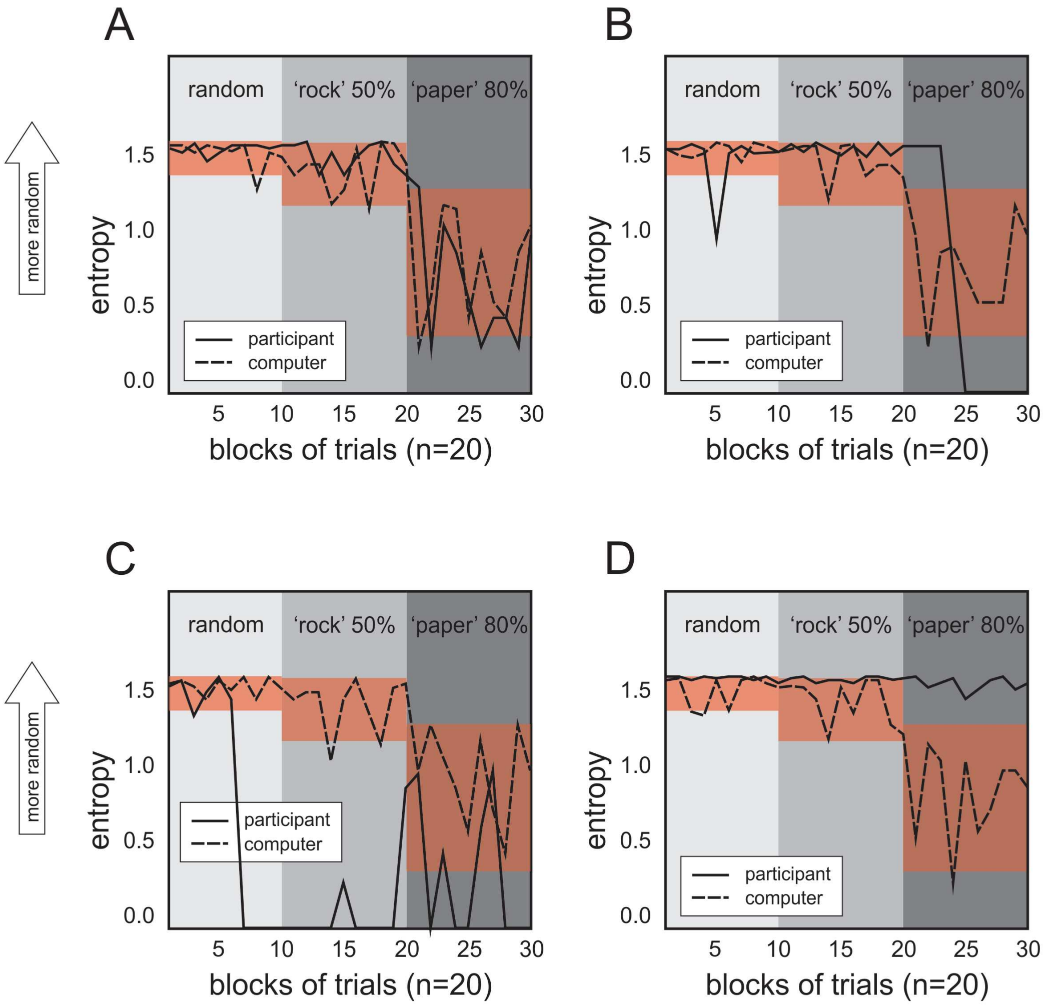

Figure 5: Short sequence entropy. Entropy is a metric of the randomness of a sequence and we compute it here with a simple histogram method whereby the proportion of each response in the sequence of 20 trials is used as the probability for calculating entropy. This figure presents each participant's responses as one continuous sequence. The empirical entropy for each block of twenty trials (no overlap) is presented against the block number. The three different computer choice conditions are shown by different shaded regions. The red shaded region highlights the empirical entropy values that would be commonly observed (upper bound 95%; lower bound 5%) due to sampling 20 trials from the fixed distribution. The computer's empiric entropy changes only modestly for the switch from uniform to 50%, and then decreases further for the 80% condition. Due to the random sampling there is substantial variability in how "random" given sequences of trials appear. For the control subject, (upper left; A) the estimated entropy of the control participant tracks the sawtooth course of the computer entropy reasonably well. The patients (panels B, C, and D) are the same participants as in Figure 4.

The experiments described in this section are ongoing. Preliminary results (Danckert and Anderson, in preparation), from participants with different locations of brain injury show that dividing participants into groups based on lobe of injury does a poor job of partitioning performance (e.g. the number of wins) on this RPS task. If participants are divided into those with and without neglect there is a substantially better separation. If we operationally define subjects based on the performance of the controls, we see an overlap of lesions involving the insula, deep subcortical white matter, and striatum (always on the right since subjects were selected for right hemisphere lesions).

Currently, we are exploring versions of the RPS game where the rate of change of the computer's strategy can be manipulated from gradual to abrupt to explore the claim that there are two discrete and different systems for detecting surprising departures from the current mental model. In addition, because the selection of the computer's choice is hidden from participants, we can also systematically vary the number of trials that participants win, lose, and draw, to examine how the consequences of a given sequence of play interacts with environmental regularities to provoke mental model updating.

5. Updating Updating

We settled on the word "updating" because its everyday meaning agreed with our sense of what we were attempting to define and study. Other people working in other areas of human neuroscience have had the same experience. As a result the term updating is used in multiple settings and for multiple functions. In this section, we highlight some of the different ways the term is used.

Updating has been used to label an aspect of working memory. Shimamura (2000) describes updating as the modulation of information in short term memory to alter activation. For example, in an n-back memory task, the nth item in short term memory is enhanced for making the comparison to the presently displayed item. The emphasis here is on short term modulation of individual items for a particular purpose. Our notion of updating is of a longer time frame process where we are more interested in whether the likelihood of certain items have changed or the rules for transitioning between them have changed. While it seems likely that working memory would be necessary for this, it does not sound like the same thing and one can imagine someone being able to report short term recall, such as the neglect patients described by Ferber and Danckert (2006), but not be adequately able to update the large scale structure relevant for a task, like the RPS scenario detailed above. Thus, while working memory might be a necessary component in the system for detecting these departures from a predictable pattern, the updating of short term memory denotes a distinct concept.

Another use of the term updating is in the area of reinforcement learning (Sutton and Barto, 1998; Daw et al., 2006). Reinforcement learning ideas have been used for decades to model classical conditioning experiments (Rescorla and Wagner, 1972; Miller et al., 1995) and are also useful for modeling many aspects of human decision making (Schonberg et al., 2007). Adjusting a set of weights or a strategy for selecting actions in response to environmental feedback is certainly a reasonable use of the term updating. However, the situations to which reinforcement learning is applicable differ from those where representational updating is useful. Reinforcement learning is applied for procedures that require stereotyped, repetitive choices with frequent reward signals, (e.g., drawing cards from decks with different probabilities of winning). The invocation of representational updating occurs in an unstructured setting, with potentially no feedback, and where the data may be observational. Where reinforcement learning uses the term updating in a dynamic programming sense of updating weights and parameters of a specific model, we use it to describe a more global process of recognizing the need to update the model.

Updating has also been used to describe executive functions. This is perhaps the oldest sense of the term in human neuroscience (Eling et al., 2008). When one needs to switch from one task to another or to switch modes of responding within a task, one needs to do something which is well summarized as updating, and the term is used to describe these scenarios (e.g. Barcelo et al., 2006). The prototypical example of this type of task is the Wisconsin Card Sorting Task (WCST). In WCST-like tasks the signal for the need to switch is clear and unambiguous and the feedback signal is overt and deterministic. But what about when it is not clear that the task or set require switching? What about when the data is noisy? What about when you are not the actor, but the observer? These, too, are part of human experience, arguably the greater part. There is also a need for a term to describe the revising of one's models of the rules and transformations that occur generally and always. The executive updater is the organizer for executing behavioral tactics, but there must also be a strategic updating that occurs as well, and it is not clear that those two things must be the same at either the conceptual level or the neural level.

This discussion with its references to latent learning and the detection of surprising events defined by calculation from mental models, prompts us to highlight some other recent studies. Strange et al. (2005) demonstrated that there are behavioral and functional imaging correlates of the statistical concepts of surprise and entropy. These are exactly the sorts of calculations that a concept of updating like the one we describe would require. Strange et al. (2005) emphasized the hippocampus as the locus for the activations. While the hippocampus was not mentioned in our review of locations above, we do not mean to imply that the only locations relevant to our variety of updating must be found to be damaged in spatial neglect. Another result showing that brains compute a surprise signal is that of Mars et al. (2008). These researchers did not focus on location, but rather establishing if there was a quantitative relation between different degrees of surprise and the neural activity giving rise to the P300. A point of emphasis for us is that updating mental models will need to occur in situations without an explicit task or clear reward. This type of latent learning is similar to the observational learning studied by Burke et al. (2010). In their recent functional imaging study they conducted a “bandit” type task, like those used for reinforcement learning modeling, but in addition, incorporated the notion of an observation where the subject in the experiment had the chance to gain information from observing another’s choice. This protocol presents an intermediate step between the classic reinforcement learning paradigm and a completely unstructured observational setting.

No one of these is uniquely entitled to the term updating. Each use is consistent with common usage, but highlights different aspects of an experimental technique or theoretical perspective. Until we know more about all of them, it will not be possible to say exactly where they may be redundant. At the present, there seems to be space for the additional sense of updating that we develop here.

6. Summary

We need to detect regularities in our environment and exploit them, but we also need to detect changes in our environment and update our mental models to reflect our new circumstances. As a first step towards a concrete and computational model of updating, we have offered a schematic of the updating process (Figure 1). This schematic provides specific components whose functions are sufficiently specific to be the object of empirical testing. We suggest some methods for this testing, and show how even simple tasks, like observing the sequence of choices made while playing the RPS game, can afford a window on the dynamics of the updating process. In addition, we compare our use of the term updating to other applications in human neuroscience. Combining these ideas with clinical observations from brain damaged patients leads us to suggest that the right inferior parietal lobe and anterior insula are two candidate structures likely to be central to an updating network.

7. Conflict of Interest

Neither of the authors has any conflicts to report.

8. References

Adair, J., Barrett, A., 2008. Spatial Neglect: Clinical and Neuroscience Review. Annals of the New York Academy of Sciences 1142, 21–43.

Allman, J., Tetreault, N., Hakeem, A., Manaye, K., Semendeferi, K., Erwin, J., Park, S., Goubert, V., Hof, P., 2010. The von Economo neurons in frontoinsular and anterior cingulate cortex in great apes and humans. Brain Structure and Function 214, 1–23.

Andersen, R.A., Asanuma, C., Essick, G., Siegel, R.M., 1990. Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. Journal of Comparative Neurology 296, 65–113.

Balleine, B., Delgado, M., Hikosaka, O., 2007. The role of the dorsal striatum in reward and decision-making. Journal of Neuroscience 27, 8161 – 8165.

Barcelo, F., Escera, C., Corral, M., Periáñez, J., 2006. Task switching and novelty processing activate a common neural network for cognitive control. Journal of Cognitive Neuroscience 18, 1734–1748.

Bartolomeo, P., Thiebaut de Schotten, M., Doricchi, F., 2007. Left unilateral neglect as a disconnection syndrome. Cerebral Cortex 17, 2479 – 2490.

Beck, J., Ma, W., Kiani, R., Hanks, T., Churchland, A., Roitman, J., Shadlen, M., Latham, P., Pouget, A., 2008. Probabilistic population codes for bayesian decision making. Neuron 60, 1142–1152.

Brown, S., Steyvers, M., 2009. Detecting and predicting changes. Cognitive Psychology 58, 49–67.

Burke, C., Tobler, P., Baddeley, M., Schultz, W., 2010. Neural mechanisms of observational learning. Proceedings of the National Academy of Sciences 107, 14431 – 14436.

Cabeza, R., Ciaramelli, E., Olson, I., Moscovitch, M., 2008. The parietal cortex and episodic memory: an attentional account. Nature Reviews Neuroscience 9, 613–625.

Carreiro, L.R.R., Haddad, H., Baldo, M.V.C., 2003. The modulation of simple reaction time by the spatial probability of a visual stimulus. Brazilian Journal of Medical and Biological Research 36, 907–11.

Churchland, P., Churchland, P., 2002. Neural worlds and real worlds. Nature Reviews Neuroscience 3, 903–907.

Ciaramitaro, V., Cameron, E., Glimcher, P., 2001. Stimulus probability directs spatial attention: an enhancement of sensitivity in humans and monkeys. Vision Research 41, 57–75.

Corbetta, M., Shulman, G., 2002. Control of goal-directed and stimulus-driven attention in the brain. Nature Reviews Neuroscience 3, 201–215.

Craig, A.D., 2009. How do you feel: now? The anterior insula and human awareness. Nature Reviews. Neuroscience 10, 59–70.

Craik, K., 1952. The Nature of Explanation. Univ. Press Cambridge, Eng.

Critchley, H., Wiens, S., Rotshtein, P., Öhman, A., Dolan, R., 2004. Neural systems supporting interoceptive awareness. Nature neuroscience 7, 189–195.

Danckert, J., Ferber, S., 2006. Revisiting unilateral neglect. Neuropsychologia 44, 987–1006.

Danckert, J., Ferber, S., Pun, C., Broderick, C., Striemer, C., Rock, S., Stewart, D., 2007. Neglected time: impaired temporal perception of multisecond intervals in unilateral neglect. Journal of Cognitive Neuroscience 19, 1706–1720.

Daw, N., O'Doherty, J., Dayan, P., Seymour, B., Dolan, R., 2006. Cortical substrates for exploratory decisions in humans. Nature 441, 876–879.

Dinstein, I., Gardner, J., Jazayeri, M., Heeger, D., 2008. Executed and observed movements have different distributed representations in human aIPS. Journal of Neuroscience 28, 11231 – 11239.

Doricchi, F., Tomaiuolo, F., 2003. The anatomy of neglect without hemianopia: a key role for parietal-frontal disconnection? Neuroreport 14, 2239 – 2243.

Dosenbach, N., Visscher, K., Palmer, E., Miezin, F., Wenger, K., Kang, H., Burgund, E., Grimes, A., Schlaggar, B., Petersen, S., 2006. A core system for the implementation of task sets. Neuron 50, 799–812.

Druker, M., Anderson, B., 2010. Spatial probability aids visual stimulus discrimination. Frontiers in Human Neuroscience 4, 1 – 10.

Eckstein, M.P., Drescher, B.A., Shimozaki, S.S., 2006. Attentional cues in real scenes, saccadic targeting, and Bayesian priors. Psychological Science 17, 973–980.

Eling, P., Derckx, K., Maes, R., 2008. On the historical and conceptual background of the Wisconsin Card Sorting Test. Brain and Cognition 67, 247–253.

Epstein, R., Kirshnit, C., Lanza, R., Rubin, L., 1984. 'Insight'in the pigeon: antecedents and determinants of an intelligent performance. Nature 308, 61 – 62.

Estes, W., 1972. Research and theory on the learning of probabilities. Journal of the American Statistical Association 67, 81–102.

Farrell, S., Ludwig, C., Ellis, L., Gilchrist, I., 2010. Influence of environmental statistics on inhibition of saccadic return. Proceedings of the National Academy of Sciences 107, 929 – 934.

Ferber, S., Danckert, J., 2006. Lost in space–The fate of memory representations for non-neglected stimuli. Neuropsychologia 44, 320–325.

Ferber, S., Danckert, J., Joanisse, M., Goltz, H., Goodale, M., 2003. Eye movements tell only half the story. Neurology 60, 1826 – 1829.

Friston, K., 2009. The free-energy principle: a rough guide to the brain? Trends in Cognitive Sciences 13, 293–301.

Geng, J.J., Behrmann, M., 2002. Probability cuing of target location facilitates visual search implicitly in normal participants and patients with hemispatial neglect. Psychological Science 13, 520–525.

Geng, J.J., Behrmann, M., 2005. Spatial probability as an attentional cue in visual search. Perception & Psychophysics 67, 1252–1268.

Geng, J.J., Behrmann, M., 2006. Competition between simultaneous stimuli modulated by location probability in hemispatial neglect. Neuropsychologia 44, 1050–1060.

He, B., Snyder, A., Vincent, J., Epstein, A., Shulman, G., Corbetta, M., 2007. Breakdown of functional connectivity in frontoparietal networks underlies behavioral deficits in spatial neglect. Neuron 53, 905–918.

Hoffmann, J., Kunde, W., 1999. Location-specific target expectancies in visual search. Journal of Experimental Psychology. Human Perception and Performance 25, 1127–1141.

Husain, M., Mannan, S., Hodgson, T., Wojciulik, E., Driver, J., Kennard, C., 2001. Impaired spatial working memory across saccades contributes to abnormal search in parietal neglect. Brain 124, 941 – 952.

Husain, M., Rorden, C., 2003. Non-spatially lateralized mechanisms in hemispatial neglect. Nature Reviews Neuroscience 4, 26–36.

Husain, M., Shapiro, K., Martin, J., Kennard, C., 1997. Abnormal temporal dynamics of visual attention in spatial neglect patients. Nature 385, 154–156.

Johnson, A., van der Meer, M., Redish, A., 2007. Integrating hippocampus and striatum in decision-making. Current opinion in neurobiology 17, 692–697.

Kadota, H., Nakajima, Y., Miyazaki, M., Sekiguchi, H., Kohno, Y., Kansaku, K., 2009. Anterior prefrontal cortex activities during the inhibition of stereotyped responses in a neuropsychological rock-paper-scissors task. Neuroscience letters 453, 1–5.

Karnath, H., Ferber, S., Himmelbach, M., 2001. Spatial awareness is a function of the temporal not the posterior parietal lobe. Nature 411, 950–953.

Karnath, H., Fruhmann Berger, M., Kuker, W., Rorden, C., 2004. The anatomy of spatial neglect based on voxelwise statistical analysis: a study of 140 patients. Cerebral Cortex 14, 1164 – 1172.

Karnath, H., Himmelbach, M., Rorden, C., 2002. The subcortical anatomy of human spatial neglect: putamen, caudate nucleus and pulvinar. Brain 125, 350 – 360.

Karnath, H., Rorden, C., Ticini, L., 2009. Damage to white matter fiber tracts in acute spatial neglect. Cerebral Cortex 19, 2331 – 2337.

Klein, R., 2000. Inhibition of return. Trends in Cognitive Sciences 4, 138–147.

Köhler, W., 1925/1999. The Mentality of Apes. Routledge.

Körding, K.P., Wolpert, D.M., 2004. Bayesian integration in sensorimotor learning. Nature 427, 244–247.

Kristjánsson, Á., Vuilleumier, P., Malhotra, P., Husain, M., Driver, J., 2005. Priming of color and position during visual search in unilateral spatial neglect. Journal of Cognitive Neuroscience 17, 859–873.

Lambert, A., 1987. Expecting different categories at different locations and spatial selective attention. The Quarterly Journal of Experimental Psychology Section A 39, 61–76.

Langer, K., 2009. Babinski's anosognosia for hemiplegia in early twentieth-century French neurology. Journal of the History of the Neurosciences 18, 387–405.

Lee, D., Conroy, M., McGreevy, B., Barraclough, D., 2004. Reinforcement learning and decision making in monkeys during a competitive game. Cognitive Brain Research 22, 45–58.

Lee, D., McGreevy, B., Barraclough, D., 2005. Learning and decision making in monkeys during a rock-paper-scissors game. Cognitive Brain Research 25, 416–430.

Lee, D., Seo, H., 2007. Mechanisms of reinforcement learning and decision making in the primate dorsolateral prefrontal cortex. Annals of the New York Academy of Sciences 1104, 108–122.

Leirer, V., Wienbruch, C., Paul-Jordanov, I., Kolassa, S., Elbert, T., Kolassa, I., 2010. Hippocampal activity during the transverse patterning task declines with cognitive competence but not with age. BMC neuroscience 11, 113 – 122.

Lewis, J., Van Essen, D., 2000. Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. The Journal of Comparative Neurology 428, 112–137.

Liston, D.B., Stone, L.S., 2008. Effects of prior information and reward on oculomotor and perceptual choices. Journal of Neuroscience 28, 13866–13875.

Maki, A., 2010. ESPN gives credibility to RPS. http://www.theglobeandmail.com/sports/article746716.ece.

Mars, R., Debener, S., Gladwin, T., Harrison, L., Haggard, P., Rothwell, J., Bestmann, S., 2008. Trial-by-trial fluctuations in the event-related electroencephalogram reflect dynamic changes in the degree of surprise. Journal of Neuroscience 28, 12539.

Marshall, J., Halligan, P., 2002. Whoever would have imagined it? Bisiach and Luzzatti.(1978). On representational neglect in patients, IG, & NV. Classic Cases in Neuropsychology 2, 272–274.

Matsubara, M., Yamaguchi, S., Xu, J., Kobayashi, S., 2004. Neural correlates for the suppression of habitual behavior: a functional MRI study. Journal of cognitive neuroscience 16, 944–954.

van der Meer, M., Redish, A., 2010. Covert expectation-of-reward in rat ventral striatum at decision points. Frontiers in Integrative Neuroscience 4, 1 – 12.

Merrifield, C., Hurwitz, M., Danckert, J., 2010. Multimodal temporal perception deficits in a patient with left spatial neglect. Cognitive Neuroscience In press.

Miller, R., Barnet, R., Grahame, N., 1995. Assessment of the Rescorla-Wagner model. Psychological Bulletin 117, 363–386.

Mort, D., Malhotra, P., Mannan, S., Rorden, C., Pambakian, A., Kennard, C., Husain, M., 2003. The anatomy of visual neglect. Brain 126, 1986 – 1997.

Na, D., Adair, J., Kang, Y., Chung, C., Lee, K., Heilman, K., 1999. Motor perseverative behavior on a line cancellation task. Neurology 52, 1569 – 1576.

O'Connor, A., Han, S., Dobbins, I., 2010. The inferior parietal lobule and recognition memory: expectancy violation or successful retrieval? Journal of Neuroscience 30, 2924 – 2934.

Paulus, M., Feinstein, J., Leland, D., Simmons, A., 2005. Superior temporal gyrus and insula provide response and outcome-dependent information during assessment and action selection in a decision-making situation. Neuroimage 25, 607–615.

Paulus, M., Feinstein, J., Tapert, S., Liu, T., 2004. Trend detection via temporal difference model predicts inferior prefrontal cortex activation during acquisition of advantageous action selection. Neuroimage 21, 733 – 743.

Penfield, W., Faulk Jr, M., 1955. The insula: further observations on its function. Brain 78, 445 – 470.

Pennartz, C., Berke, J., Graybiel, A., Ito, R., Lansink, C., van der Meer, M., Redish, A., Smith, K., Voorn, P., 2009. Corticostriatal interactions during learning, memory processing, and decision making. Journal of Neuroscience 29, 12831 – 12838.

Rescorla, R., Wagner, A., 1972. Variations in the Effectiveness of Reinforcement and Nonreinforcement.

Rizzo, M., Akutsu, H., Dawson, J., 2001. Increased attentional blink after focal cerebral lesions. Neurology 57, 795 – 800.

Salvetti, F., Patelli, P., Nicolo, S., 2007. Chaotic time series prediction for the game, Rock-Paper-Scissors. Applied Soft Computing 7, 1188–1196.

Sato, Y., Akiyama, E., Farmer, J., 2002. Chaos in learning a simple two-person game. Proceedings of the National Academy of Sciences 99, 4748 – 4751.

Schonberg, T., Daw, N., Joel, D., O'Doherty, J., 2007. Reinforcement learning signals in the human striatum distinguish learners from nonlearners during reward-based decision making. Journal of Neuroscience 27, 12860.

Seo, H., Barraclough, D., Lee, D., 2009. Lateral intraparietal cortex and reinforcement learning during a mixed-strategy game. Journal of Neuroscience 29, 7278 – 7289.

Shapiro, K., Hillstrom, A., Husain, M., 2002. Control of visuotemporal attention by inferior parietal and superior temporal cortex. Current Biology 12, 1320–1325.

Shapiro, K., Raymond, J., Arnell, K., 1997. The attentional blink. Trends in Cognitive Sciences 1, 291–296.

Shaqiri, A., Anderson, B., 2010. Spatial probability in hemispatial neglect. Poster at: 33rd European Conference on Visual Perception .

Shimada, S., Abe, R., 2010. Outcome and view of the player modulate motor area activity during observation of a competitive game. Neuropsychologia 48, 1930 – 1934.

Shimamura, A., 2000. The role of the prefrontal cortex in dynamic filtering. Psychobiology 28, 207–218.

Strange, B., Duggins, A., Penny, W., Dolan, R., Friston, K., 2005. Information theory, novelty and hippocampal responses: unpredicted or unpredictable? Neural Networks 18, 225–230.

Sutton, R., Barto, A., 1998. Reinforcement Learning: An Introduction. The MIT press.

Tolman, E.C., 1948. Cognitive maps in rats and men. Psychological Review 55, 189 – 208.

Tolman, E.C., Honzik, C.H., 1930. "Insight" in rats. University of California, Publications in Psychology 4, 331 –351.

Umarova, R., Saur, D., Schnell, S., Kaller, C., Vry, M., Glauche, V., Rijntjes, M., Hennig, J., Kiselev, V., Weiller, C., 2010. Structural Connectivity for Visuospatial Attention: Significance of Ventral Pathways. Cerebral Cortex 20, 121 – 129.

Vallar, G., Perani, D., 1987. The anatomy of spatial neglect in humans. Advances in Psychology 45, 235–258.

Verdon, V., Schwartz, S., Lovblad, K., Hauert, C., Vuilleumier, P., 2010. Neuroanatomy of hemispatial neglect and its functional components: a study using voxel-based lesion-symptom mapping. Brain 133, 880 – 894.

Vulkan, N., 2000. An economist's perspective on probability matching. Journal of Economic Surveys 14, 101–118.

Walker, D., 2004. The Official Rock Paper Scissors Strategy Guide. Fireside Books.

Walthew, C., Gilchrist, I.D., 2006. Target location probability effects in visual search: an effect of sequential dependencies. Journal of Experimental Psychology. Human Perception and Performance 32, 1294–1301.

Zhang, S., Lee, M., in press. Optimal experimental design for a class of bandit problems. Journal of Mathematical Psychology Corrected Proof, Available online 24 September 2010, ISSN 0022-2496, DOI: 10.1016/j.jmp.2010.08.002.

Figure 1.JPEG

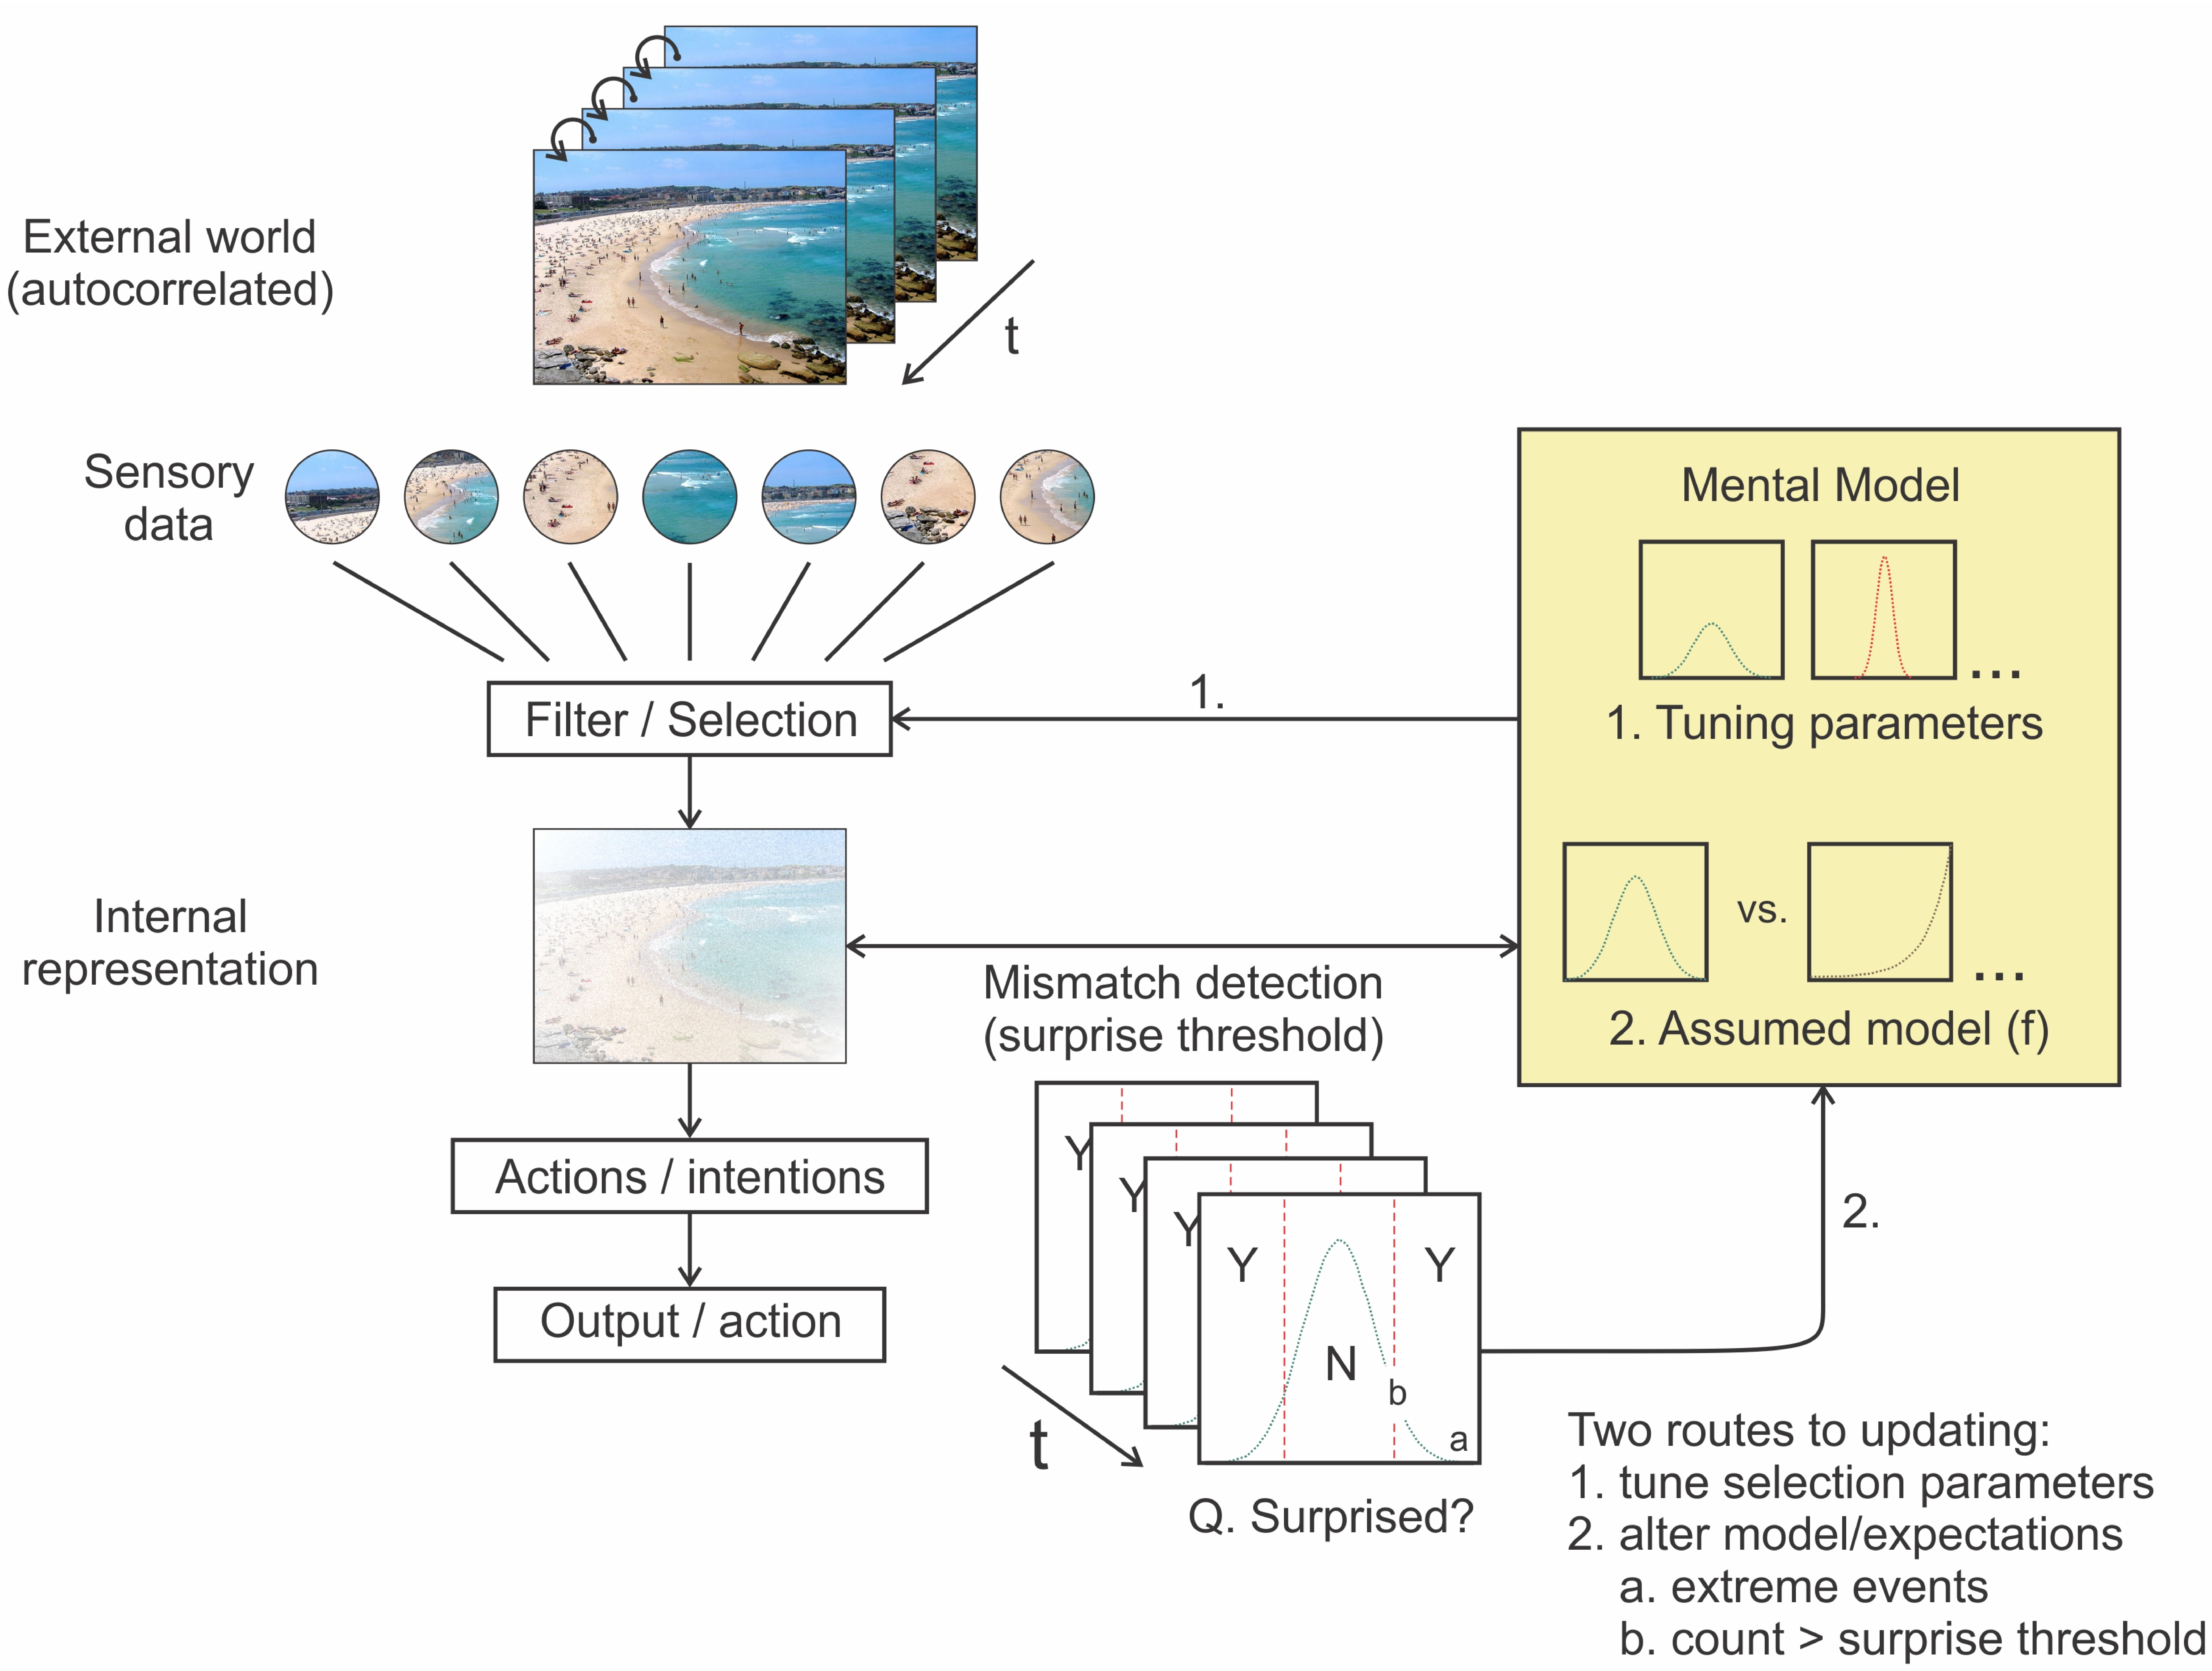

Figure 2.JPEG

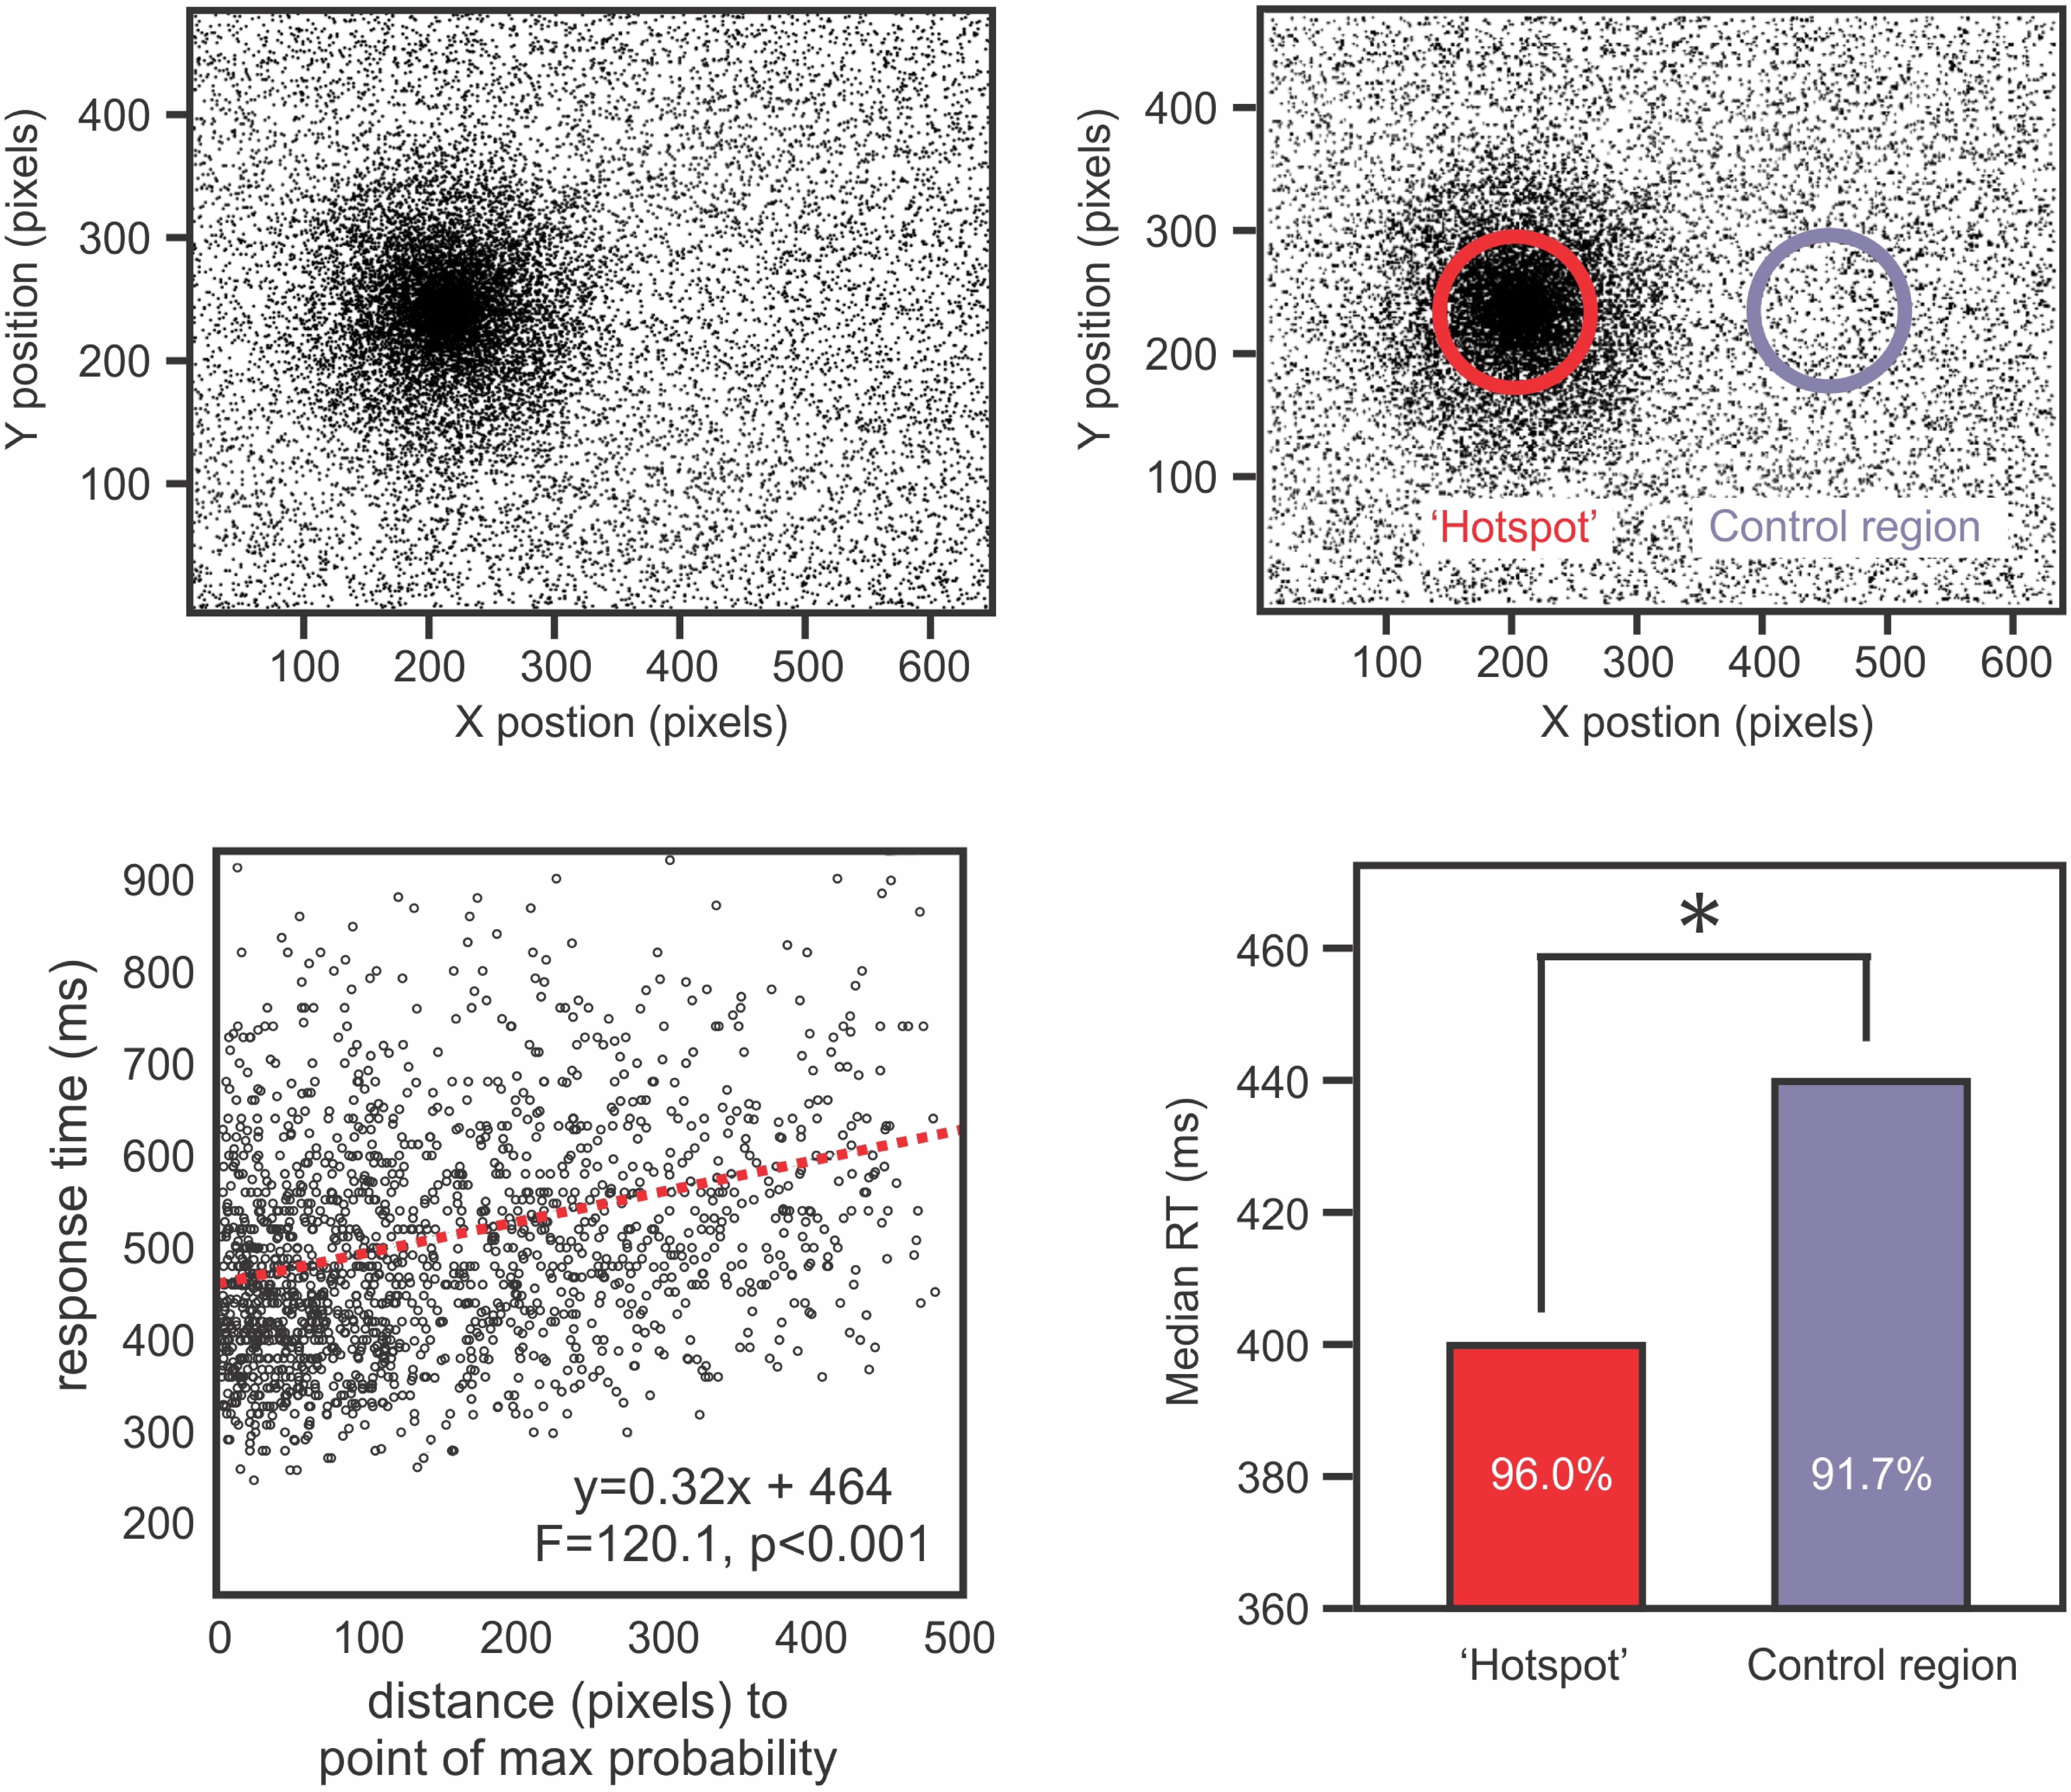

Figure 3.JPEG

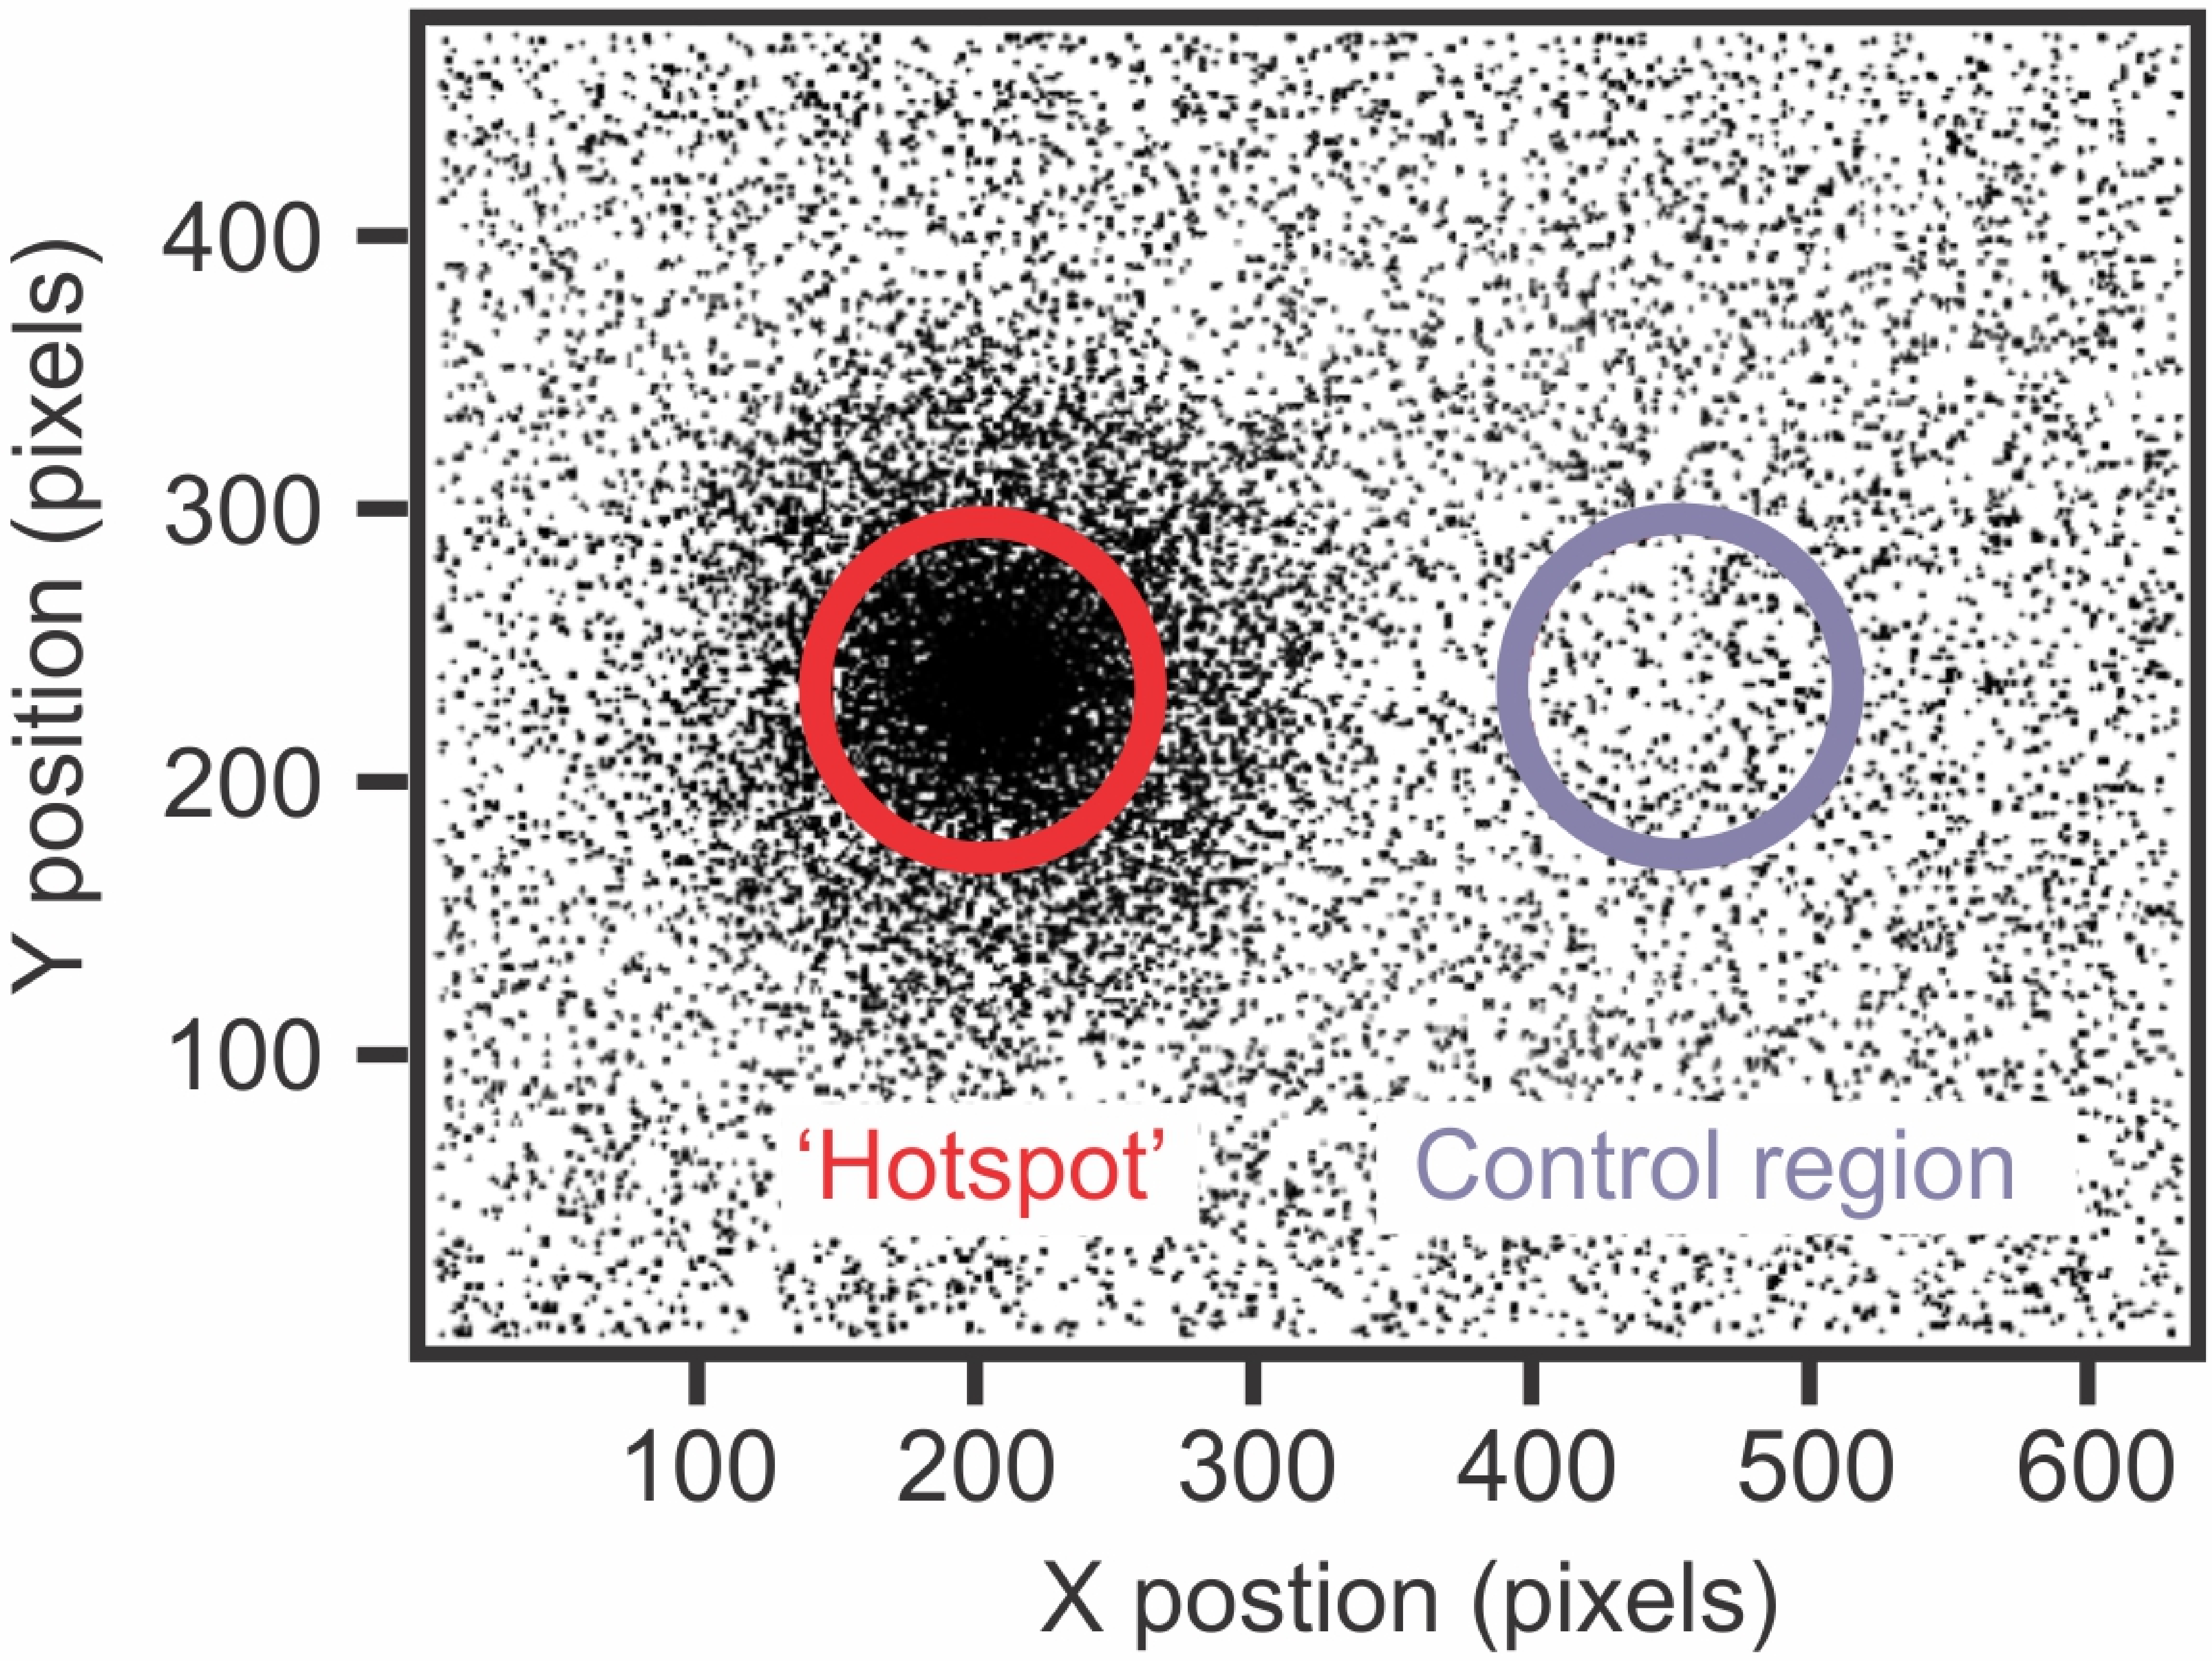

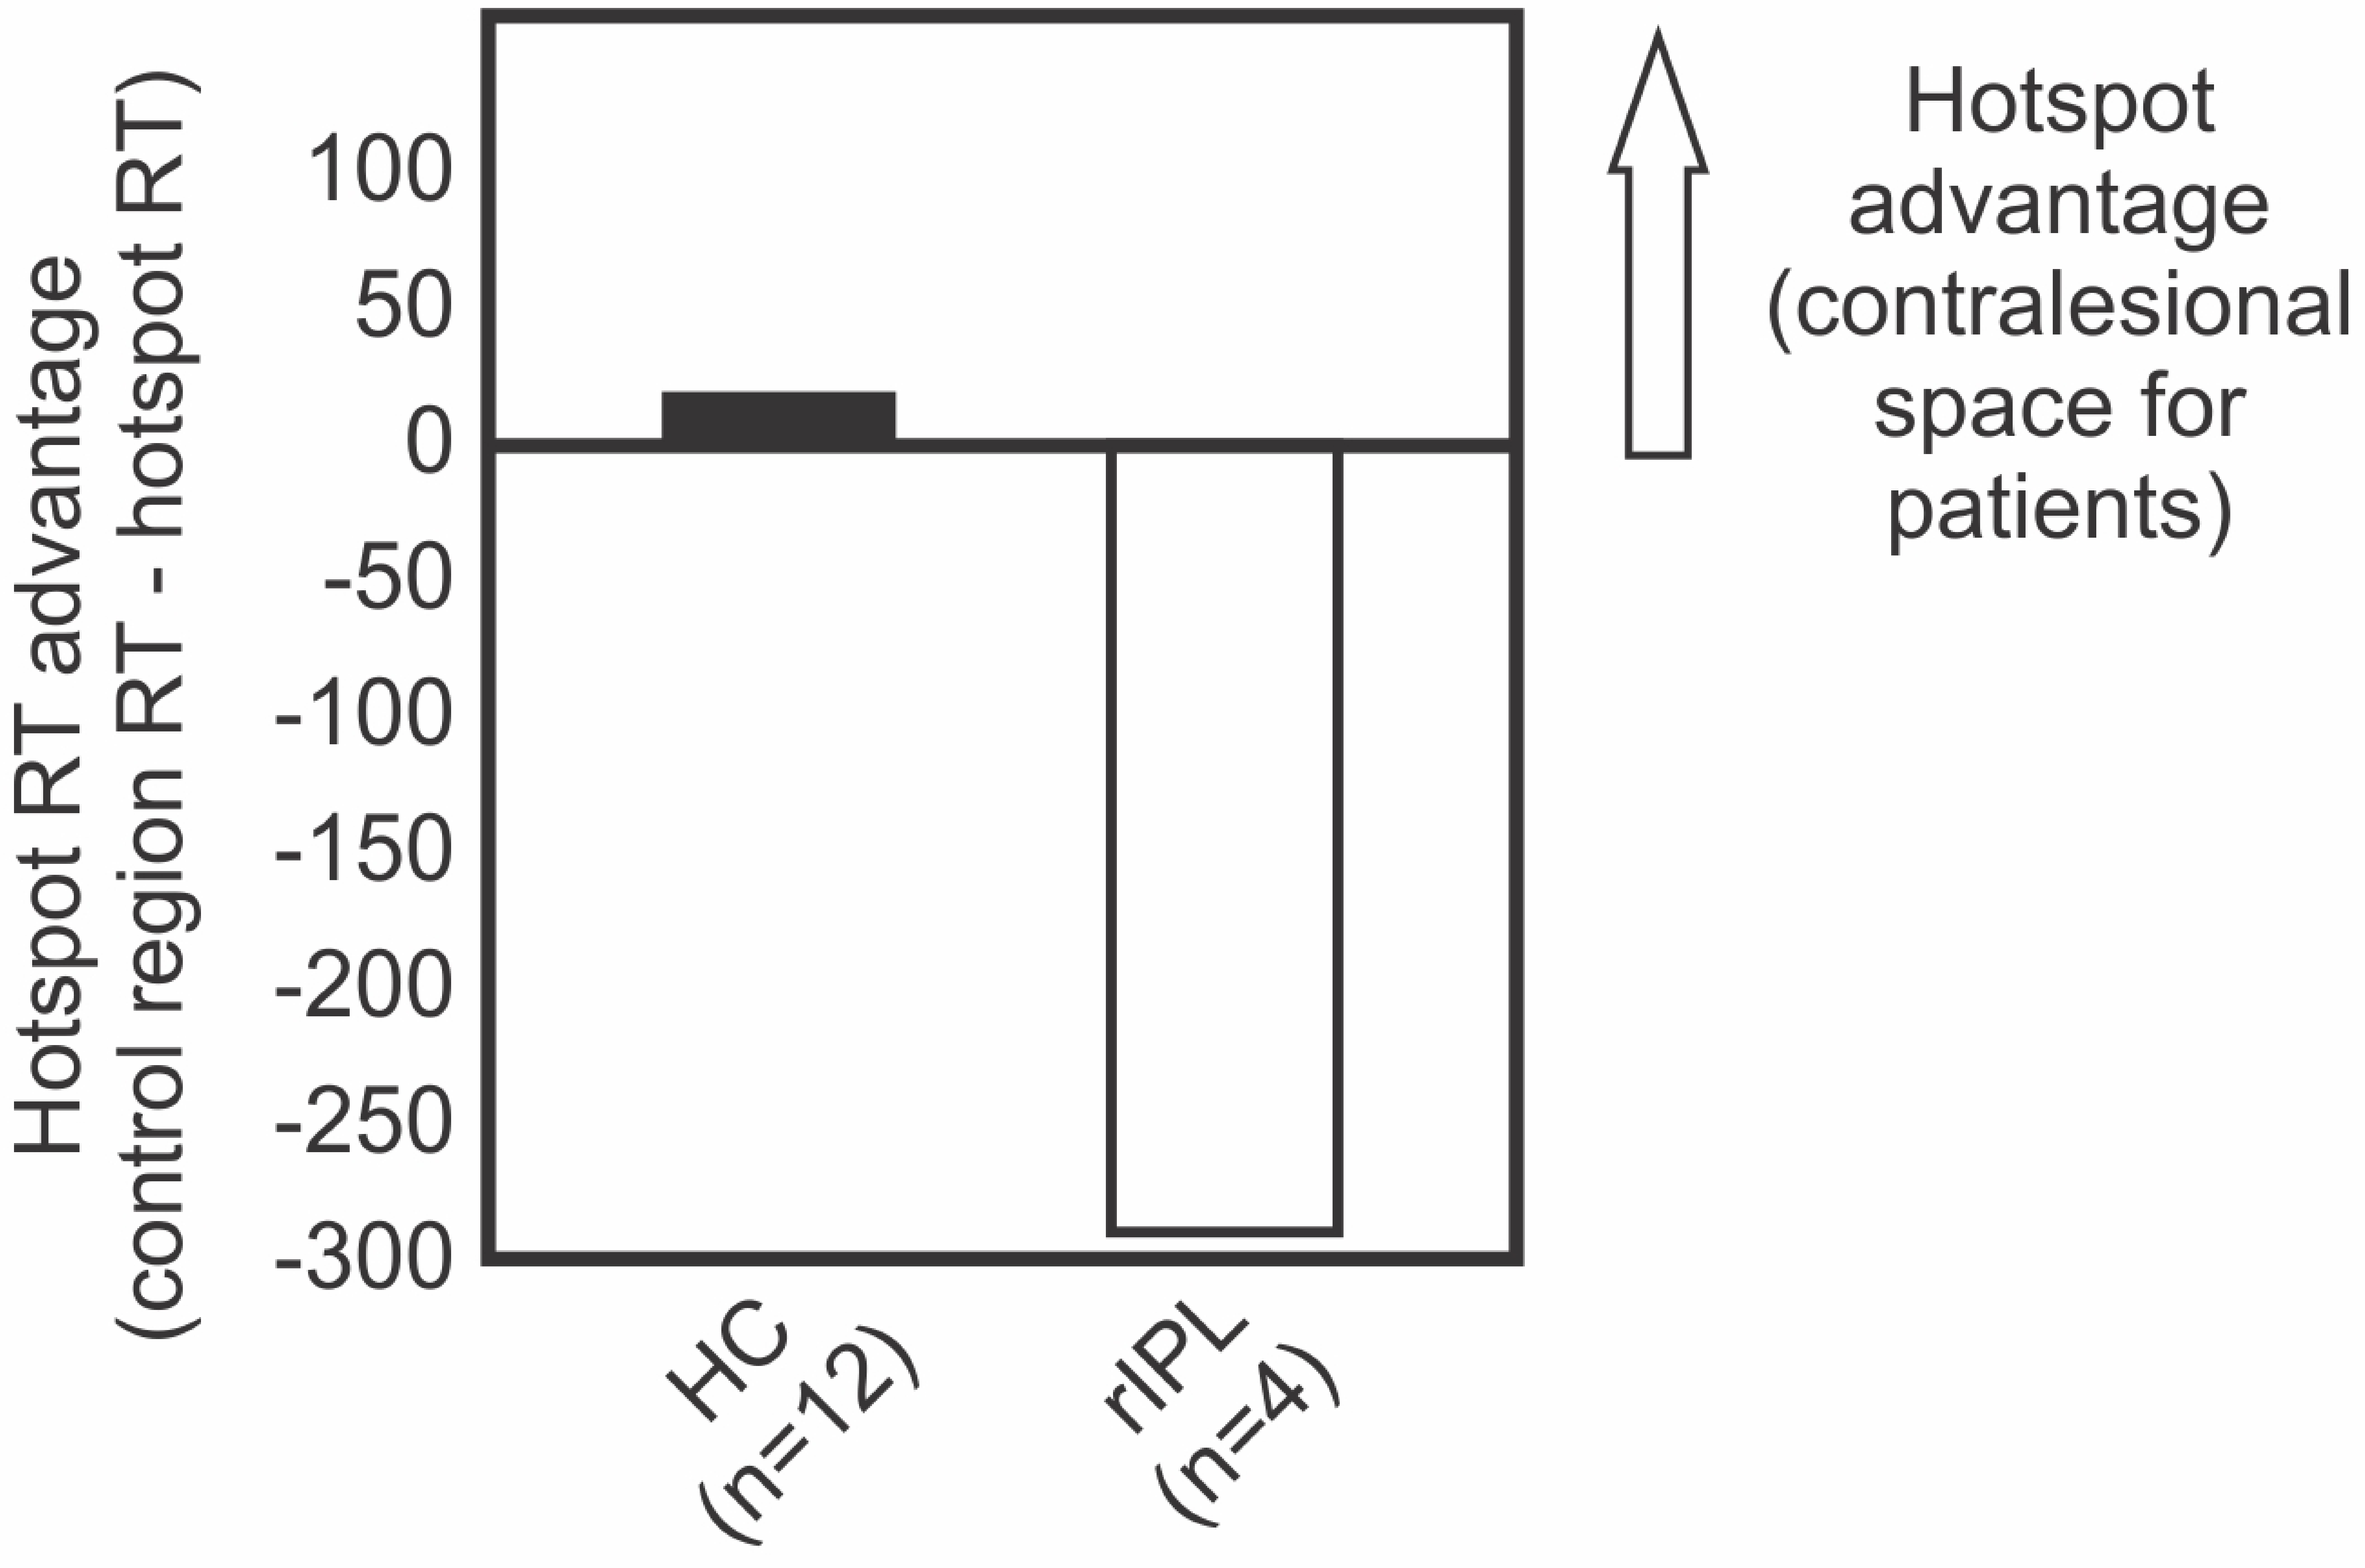

Figure 4.JPEG

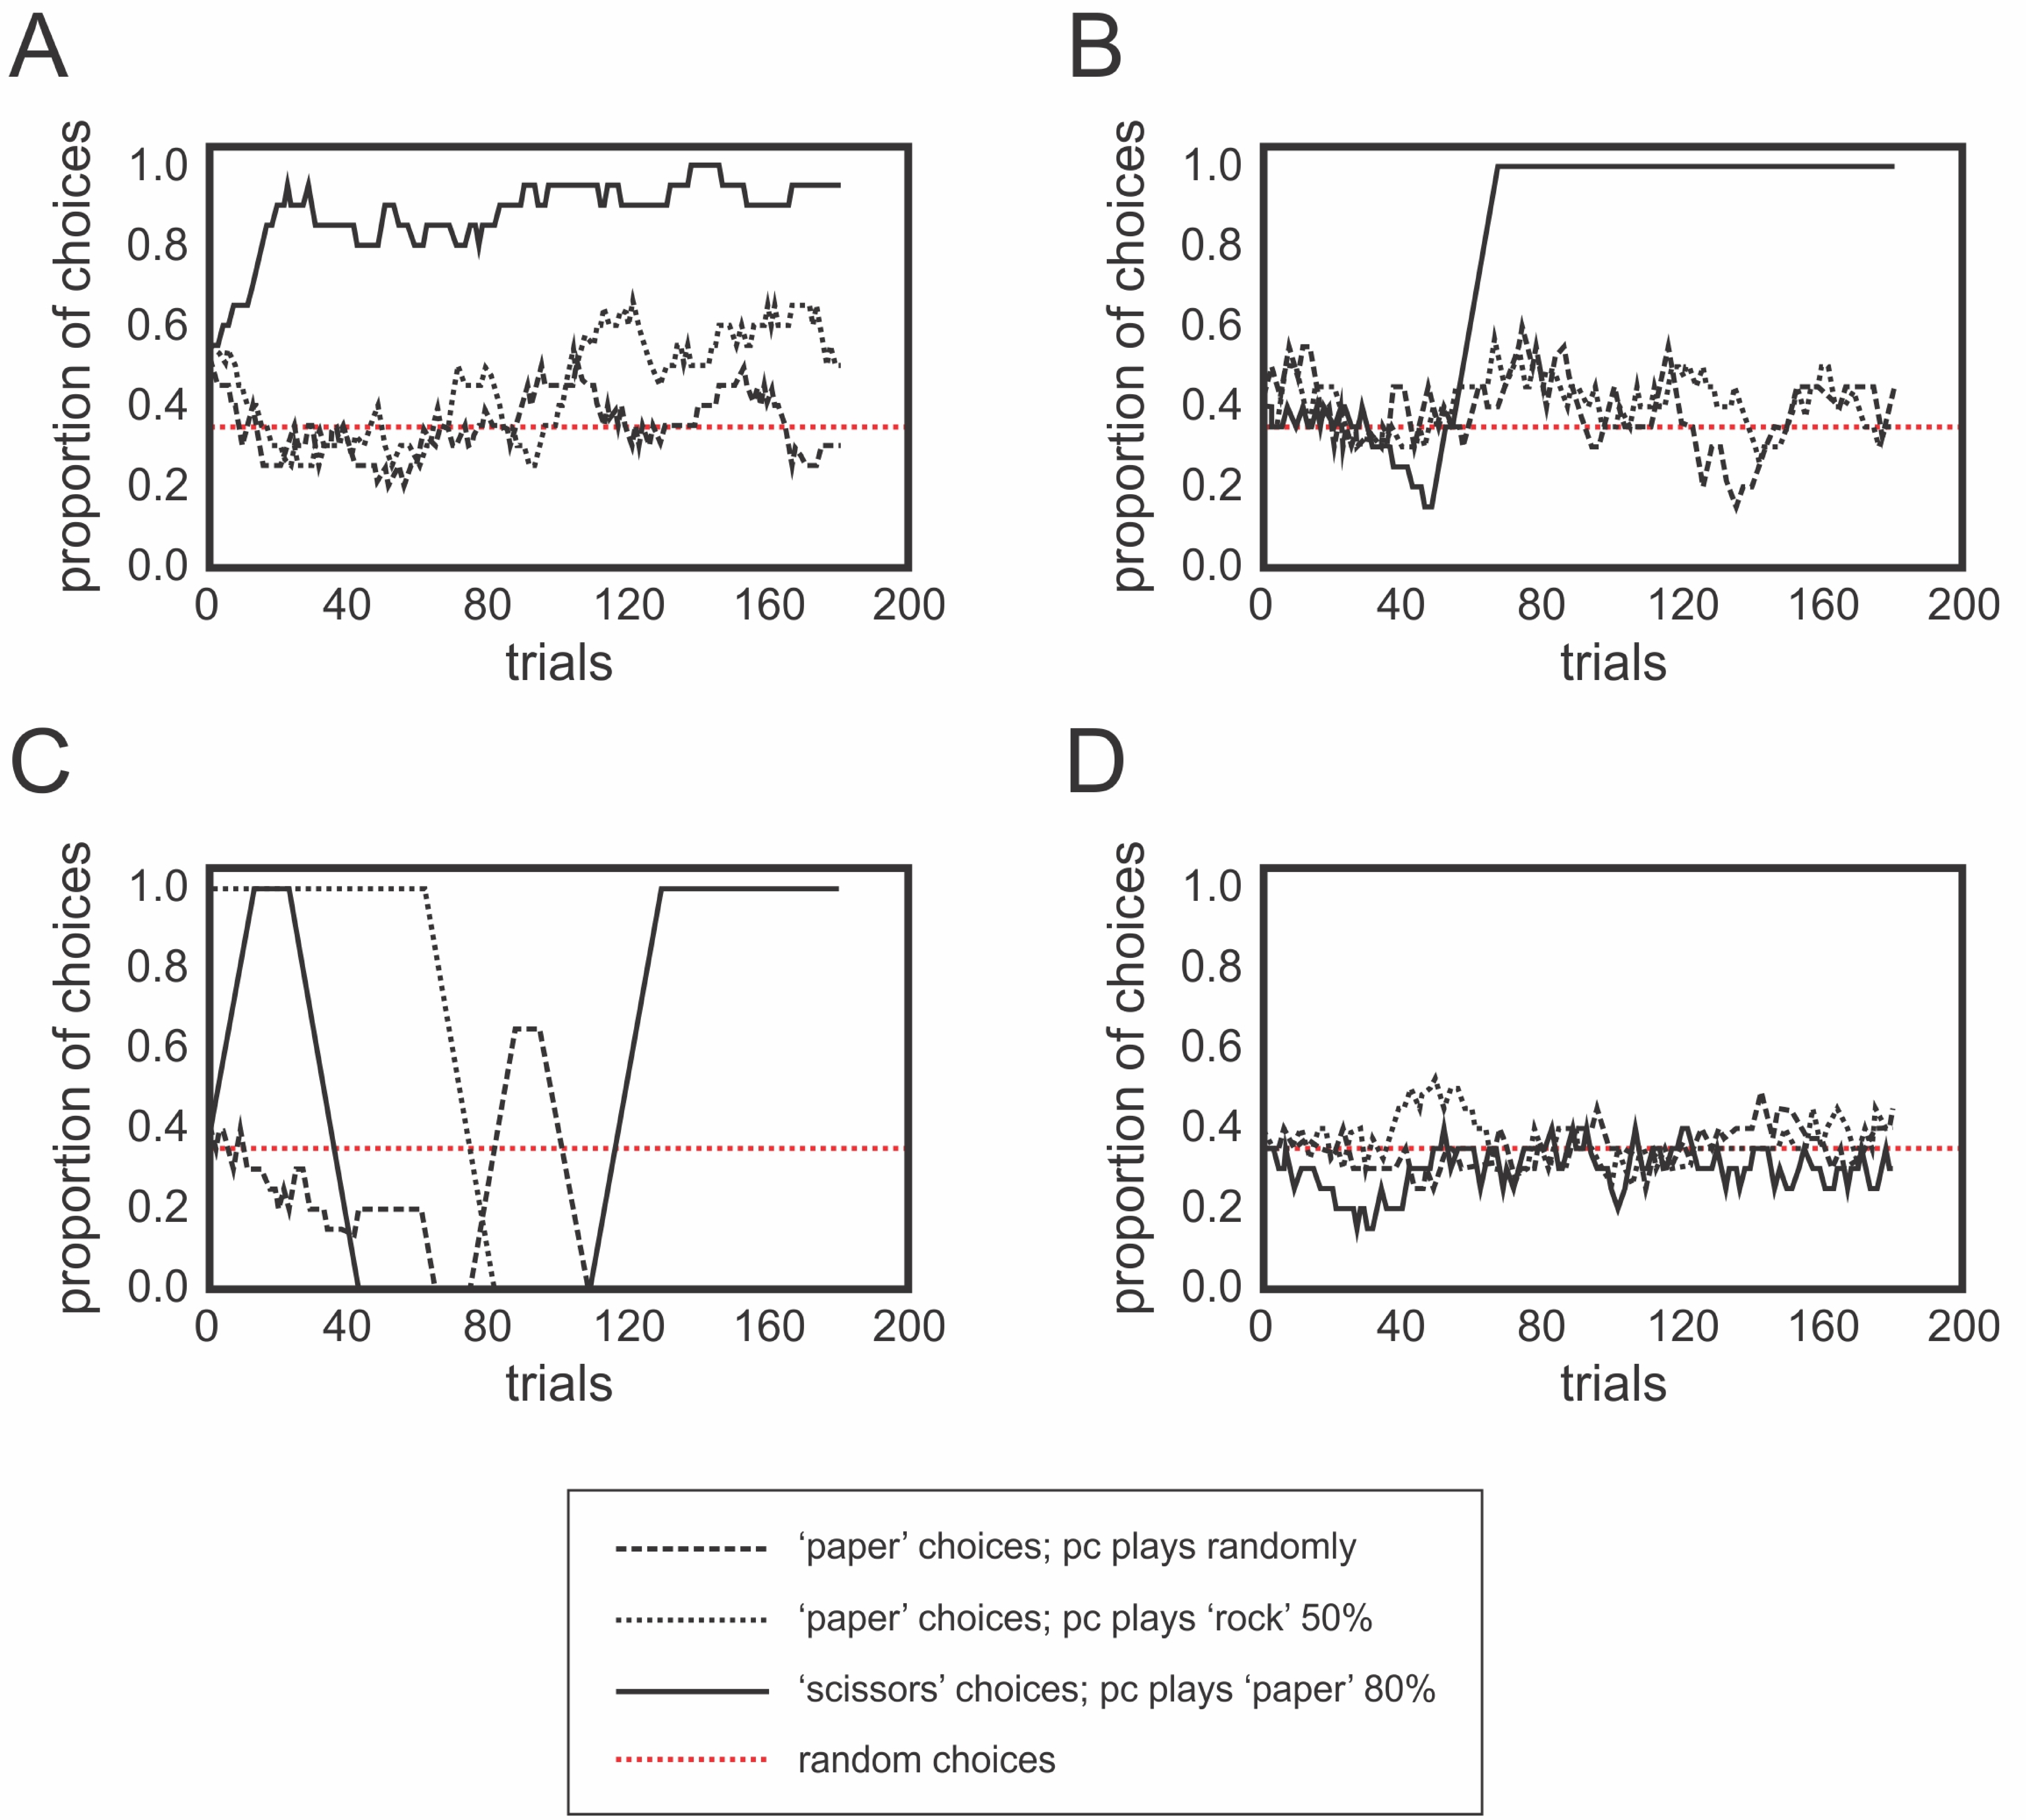

Figure 5.JPEG