Neuropsychologia 48 (2010) 2912–2921

Contents lists available at ScienceDirect

Neuropsychologia

journal homepage: www.elsevier.com/locate/neuropsychologia

A virtual reality-based FMRI study of reward-based spatial learning

Rachel Marsh[a,*], Xuejun Hao[a], Dongrong Xu[a], Zhishun Wang[a], Yunsuo Duan[a], Jun Liu[a], Alayar Kangarlu[a], Diana Martinez[c], Felix Garcia[a], Gregory Z. Tau[a], Shan Yu[a], Mark G. Packard[b], Bradley S. Peterson[a]

[a] The MRI Unit and Division of Child & Adolescent Psychiatry in the Department of Psychiatry, the New York State Psychiatric Institute and the College of Physicians and Surgeons, Columbia University, New York, NY, United States
[b] The Department of Psychology, Texas A&M University, College Station, TX, United States
[c] The Department of Psychiatry, the New York State Psychiatric Institute and the College of Physicians and Surgeons, Columbia University, New York, NY, United States

ARTICLE INFO

Article history:
Received 7 December 2009
Received in revised form 26 May 2010
Accepted 28 May 2010
Available online 4 June 2010

Keywords:
Memory systems
fMRI
Spatial learning
Reward

ABSTRACT

Although temporo-parietal cortices mediate spatial navigation in animals and humans, the neural correlates of reward-based spatial learning are less well known. Twenty-five healthy adults performed a virtual reality fMRI task that required learning to use extra-maze cues to navigate an 8-arm radial maze and find hidden rewards. Searching the maze in the spatial learning condition compared to the control conditions was associated with activation of temporo-parietal regions, albeit not including the hippocampus. The receipt of rewards was associated with activation of the hippocampus in a control condition when using the extra-maze cues for navigation was rendered impossible by randomizing the spatial location of cues. Our novel experimental design allowed us to assess the differential contributions of the hippocampus and other temporo-parietal areas to searching and reward processing during reward-based spatial learning. This translational research will permit parallel studies in animals and humans to establish the functional similarity of learning systems across species; cellular and molecular studies in animals may then inform the effects of manipulations on these systems in humans, and fMRI studies in humans may inform the interpretation and relevance of findings in animals.

Published by Elsevier Ltd.

1. Introduction

Previous fMRI studies of spatial navigation within virtual reality (VR) environments have consistently identified a network of regions – including the hippocampus, parahippocampus, retrosplenial cortex, and posterior parietal cortex – that are engaged during virtual navigation relative to non-navigational control tasks (Aguirre, Detre, Alsop, & D'Esposito, 1996; Hartley, Maguire, Spiers, & Burgess, 2003; Hassabis et al., 2009; Iaria, Fox, Chen, Petrides, & Barton, 2008; Spiers & Maguire, 2006b; Wolbers & Buchel, 2005). However, many of these control tasks differed from the navigation tasks on salient features, such as the form of visual stimuli, motor demands, task instructions, or cognitive effort required for navigation, thereby precluding full isolation of the neural correlates of spatial learning in standard functional imaging subtraction paradigms. In addition, none of the prior fMRI studies have directly assessed the role of reward in spatial learning.

Our VR task for spatial learning is directly analogous to the standard "win-shift" radial-arm maze paradigm that has been used to study spatial learning and memory in animals. This paradigm involves navigation of a maze with 8 identically appearing arms extending outward from a central platform that has extra-maze objects visible from within the maze. Rodents obtain food rewards by visiting each arm of the maze once, with re-entries into maze arms previously visited scored as errors. Because performance on this task requires rodents to remember those arms that have been previously visited, it is regarded as a prototypical test of spatial memory or cognitive mapping. Extensive evidence from animal studies employing lesions (Olton & Samuelson, 1976; Packard, Hirsh, & White, 1989) and pharmacological stimulation (Packard & White, 1991) have demonstrated that the hippocampus contributes to performance on the win-shift radial maze task.

Analogous fMRI studies in humans can complement these radial-arm maze studies of spatial learning in animals. For example, fMRI studies can more readily and simultaneously assess the differential involvement of brain regions other than the hippocampus in spatial learning. In addition, fMRI studies can disentangle the roles of reward and learning in reward-based spatial learning, and they more readily can study of de novo learning, without requiring prior training as in animal studies.

* Corresponding author at: Columbia University and the New York State Psychiatric Institute, 1051 Riverside Drive, Unit 74, New York, NY 10032, United States. Tel.: +1 212 543 5384; fax: +1 212 543 0522.
E-mail address: MarshR@childpsych.columbia.edu (R. Marsh).

0028-3932/$ – see front matter. Published by Elsevier Ltd.
doi:10.1016/j.neuropsychologia.2010.05.033

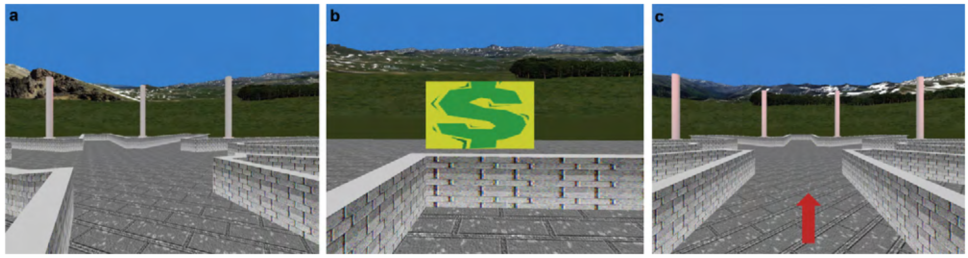

Fig. 1. (a) A view of the virtual reality environment, (b) the U.S. dollar signs denoting rewards that were hidden at the ends of the maze arms, and (c) a view of the 'trail-following' control condition C.

We report an fMRI study of reward-based spatial learning in healthy adults performing a novel VR task. Participants used an MRI-compatible joystick to navigate through an 8-arm radial maze that was surrounded by a naturalistic landscape consisting of mountains, trees, and flowers that had to be used for spatial navigation if participants were to avoid re-entries into previously visited arms. Hidden at the end of each arm of the maze was a reward that participants could earn and keep only if they did not re-enter an arm already visited. To isolate the neural correlates of spatial learning, we included a rigorously defined control condition that shared all salient features with the active learning task except for the use of the extra-maze cues for navigation. Our *a priori* hypothesis was that participants would engage medial temporal and parietal areas during spatial learning in the active compared with the control condition. Based on prior findings of reward-based learning in animals and humans (Adcock, Thangavel, Whitfield-Gabrieli, Knutson, & Gabrieli, 2006; Ito, Robbins, Pennartz, & Everitt, 2008; Vanni-Mercier, Mauguiere, Isnard, & Dreher, 2009), we also hypothesized that activation of mesolimbic areas would be associated with the receipt of rewards in the active compared to the control condition. In addition, we explored whether mesolimbic activation was associated with the anticipation of rewards, and whether it depended on the actual receipt of rewards or on the involvement of learning processes. Finally, to ensure that any confusion or additional cognitive effort in the control condition was not producing activation in the active condition, we compared brain activity during spatial learning and the receipt of rewards in the active condition with brain activity in a 'trail-following' condition that eliminated the attempted use of the extra-maze cues for navigation.

2. Methods

2.1. Participants

Participants were recruited through flyers posted in the local community to serve as healthy control participants for a neuroimaging study of the neural basis of substance dependence. The Structured Clinical Interview for DSM-IV-TR Axis I disorders (First, Spitzer, Gibbon, & Williams, 2002) was administered to all participants. Any who met DSM-IV criteria for a current Axis I disorder, or who had a lifetime history of substance abuse disorder, neurological illness, past seizures, head trauma with loss of consciousness, mental retardation, or pervasive developmental disorder were excluded from participating. The Institutional Review Board of NYSPI approved this study, and all participants gave informed consent prior to participating.

2.2. Spatial learning paradigm

Creation of the virtual environments was accomplished through use of software that was built upon C++ and OpenGL, an API (Application Programming Interface) of graphics. The virtual environments were composed of an 8-arm radial maze with a central starting location and a low outer perimeter wall. The maze was surrounded by a naturalistic landscape (e.g., mountains, trees and flowers) that constituted the extra-maze cues that could be used for spatial navigation (Fig. 1a). Prior to scanning, participants underwent a 10-min training session on a desktop computer to practice using a joystick to move and navigate freely around a virtual maze that was similar in appearance to the experimental environment used for scanning, but with different extra-maze cues.

Stimuli during scanning were presented through non-magnetic goggles (Resonance Technology, Inc.). Participants used an MRI-compatible joystick (Current Designs Inc.) to navigate the maze. Before entering the scanner, participants were informed that they would find themselves in the center of a virtual maze with 8 identical runways extending outwards, and that hidden rewards (U.S. $, Fig. 1b) would be available at the very end of the runways. They were instructed to try to find the rewards and were told that they could earn and keep as many as they found, but that they would loose money if they revisited arms they had already visited. They were told that they would complete several sessions of the task. They were not told that the conditions of the sessions would differ from one another, and therefore they believed that they would be performing the same task multiple times. Participants were not given actual monetary rewards for their performance on the task, but rather paid a set amount for their participation in the study.

The first session (condition A) required spatial learning. To find and retain all 8 rewards, participants had to learn to use the extra-maze cues to avoid revisiting previously visited arms. Each trial began at the center platform. After reaching the end of an arm, participants reappeared automatically in the middle of the center platform to initiate a new trial, with the initial viewing perspective randomly oriented across the entire 360° of the circular maze, which compelled participants to use extra-maze cues to orient themselves for subsequent navigation. The randomized orientation at the start of each trial also prevented use of S-R strategies, such as "chaining" rules (e.g., exploring arms positioned successively to the left or right of the last arm entered, regardless of extra-maze cues), when performing the task.

After obtaining all 8 rewards in condition A, participants saw a black screen with white written text indicating that a new session was about to begin. In this first control condition (B), the same extra-maze cues as used in condition A were randomized among locations after each trial: the mountains that surrounded the maze were cut into several pieces that were then shuffled and re-stitched into whole mountains, while the individual trees and flowers were themselves shuffled among locations. This randomization of the cues together with the randomized directional orientation of participants with respect to the various arms of the maze at the start of each trial destroyed any possibility of using the spatial layout of the cues (spatial learning) to perform the task. As in condition A, the randomized orientation at the start of each trial also prevented the use of S-R or other procedural learning strategies. To control for the frequency of reward and punishment with the spatial learning task, however, participants were rewarded at the same frequency as in the previous spatial learning condition, though now without regard to any characteristics of their actual performance. Participants were not told that the cues would be randomized in this session or that anything else would differ from the previous session. This control condition thus shared all salient features with the spatial learning task (A), including all the lower-order stimulus features, such as color, texture, hue, luminance, the form and movement of visual stimuli, as well as many higher-order task features, such as the motor demands, task instructions, cognitive effort in searching for rewards, and reward and punishment experiences. Condition B was pre-programmed to terminate according to the number of trials that a given participant needed to obtain all 8 rewards in condition A. For instance, if a participant required 10 trials to find all 8 rewards in condition A (i.e., 8 correct and 2 error trials), they would be given 2 unbaited trails randomly in condition B. After completion of condition B, another black screen appeared that instructed participants to "follow the arrows" during the next session (condition C).

Condition C was a second control condition to contrast with condition A. Condition C was added to ensure that the use of the cues or the additional confusion or cognitive effort in condition B was not producing any observed activation in condition A (when compared to B). Condition C was a "trail-following" condition in which the extra-maze cues were randomized among locations in the same manner as in control condition B, but a red arrow at eye level indicated the path to follow (the arm to be traversed, Fig. 1c). Because participants were told simply to follow the arrow on each trial, they were presumably not *trying* to use the extra-maze cues for navigation, and in fact they were *unable* to use those cues because of the randomization of the extra-maze cues upon the initiation of each trial. Because the allocation of rewards was randomized, however, the arrow did not always point to the correct arm (where a reward was hidden). This control condition shared many salient features with both the spatial learning task (A) and the control condition B, including the form and movement of visual stimuli, motor demands, and reward and punishment experiences. It did not require as much cognitive effort in searching for rewards as did condition B, the task instruction ("follow the arrow") differed