Cognitive maps in some brain regions are compressed during goal-seeking decision-making

January 4 2024, by Ingrid Fadelli

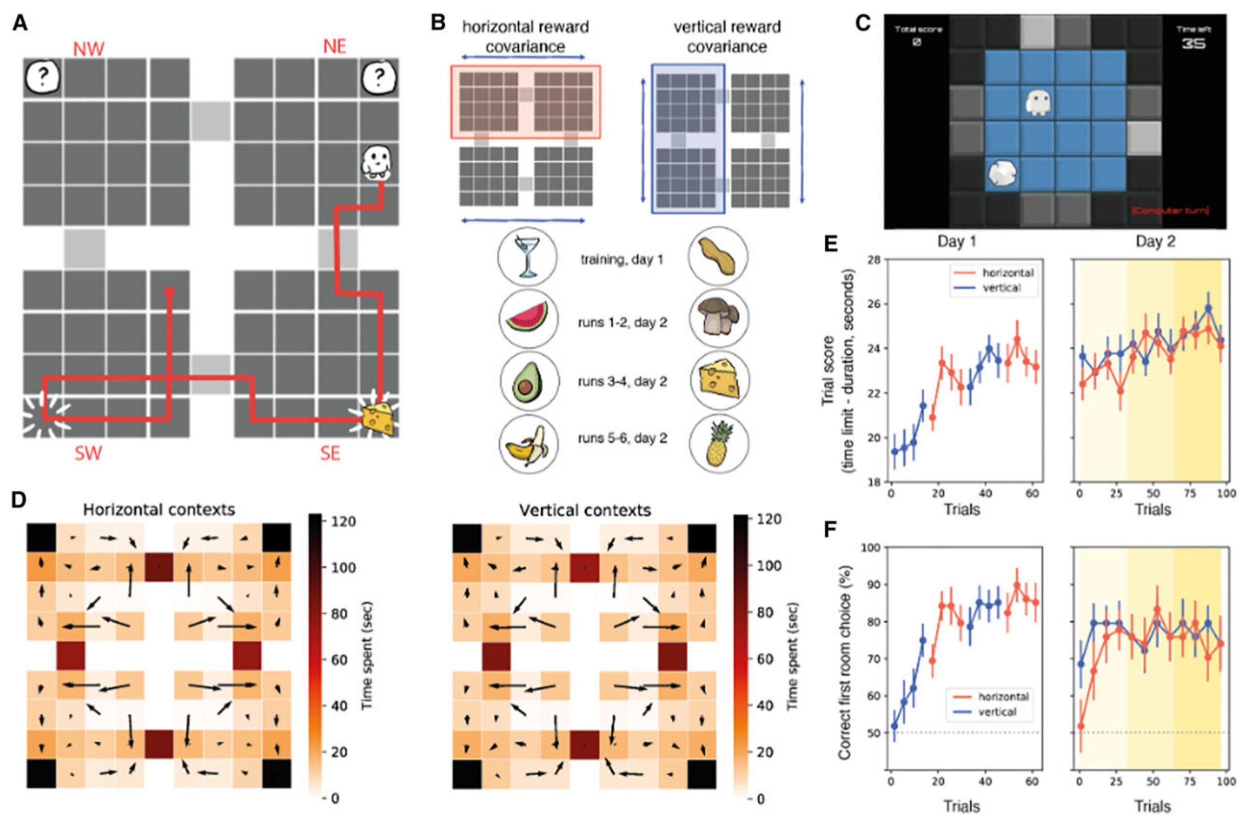

Task structure and performance. (A) Illustration of the four rooms environment and example reward locations under the vertical context. An example participant trajectory is shown overlaid in red. In this example, the agent starts in the Southwest (SW) room, explores the start boulder and does not find a cheese reward. As cheese rewards covary vertically, the two cheese rewards must therefore be in the SE and NE rooms. (B) Different contexts signaled different reward covariances: on day 1 (training) martini rewards appeared in vertically adjacent rooms, while peanut rewards appeared in horizontally adjacent rooms. On day 2, three different pairs of rewards were shown, which mapped onto the same covariance structure, and the two contexts were interleaved within a run. An example ordering of runs is shown but this was balanced across participants. (C) Participant view of exploring in one of the rooms during training. Floors of all rooms were purple in the scanner. (D) Heatmaps of the average grid square occupancy per trial in each of the two contexts. Black arrows show the average transition vector from each grid square. Data are averaged across participants. Note that the average transition vectors were also well-matched when considering only those movement periods that were controlled by the participants (see Figure S1D). (E) Participant scores on each trial on days 1 (left) and 2 (right). With training, participants get faster at finding the rewards. On day 1, contexts were blocked across trials to facilitate learning, while on day 2 they were interleaved. (F) Participants learn to preferentially search in rooms suggested by the reward structure, and this behavior generalizes to new sets of rewards associated with each context on day 2. In (D) and (E), data are shown smoothed across non-overlapping sets of 4 adjacent trials for visualization. In (E), we show only the room choices made by human participants and exclude those made by the agent. Error bars show standard error of the mean across participants. Color panels indicate different epochs with the same reward pairs. Credit: *Neuron* (2023). DOI: 10.1016/j.neuron.2023.08.021

Human decision-making has been the focus of a wide range of research studies. Collectively, these research efforts could help to understand better how people make different types of everyday choices while also shedding light on the neural processes underpinning these choices.

Findings suggest that while making instantaneous decisions, or in other words, choices that need to be made quickly based on the information available at a given moment, humans greatly rely on contextual information. This contextual information can also guide so-called sequential decisions, which entails making a choice after observing the sequential unfolding of a process.

Researchers at the University of Oxford, the National Research Council in Rome, University College London (UCL), and the Max Planck Institute for Human Development recently carried out a study exploring the impact of context on goal-directed decision-making. Their findings, published in *Neuron*, suggest that goal-seeking 'compresses' spatial maps in the hippocampus and orbitofrontal cortices in the brain.

"Humans can navigate flexibly to meet their goals," Paul S. Muhie-Karbe, Hannah Sheahan, and their colleagues wrote in their paper. "We asked how the neural representation of allocentric space is distorted by goal-directed behavior. Participants navigated an agent to two successive goal locations in a grid world environment comprising four interlinked rooms, with a contextual cue indicating the conditional dependence of one goal location on another."

To further explore what happens in the brain during goal-directed decision-making, the researchers carried out an experiment involving 27 human participants. These participants completed a task on a computer screen, which entailed navigating a virtual environment by controlling an avatar.

This avatar could move through a partially visible world represented in a grid form. This virtual world included four different rooms connected to each other, and the participants only saw the room that their avatar was occupying from above (i.e., with a bird's eye view).

During each experimental trial, the participants' avatar appeared in a room picked at random, and the participants were asked to move it using buttons on a keyboard to pick up rewards by colliding with some boulders while avoiding boulders that were empty.

At the beginning of each trial, the participants were also offered a contextual cue, which concealed partial clues hinting (but not clearly disclosing) where the rewards could be found within the virtual world. Notably, as the participants completed this task requiring goal-directed decision-making, their brain activity was recorded by an fMRI scanner.

"Examining the neural geometry by which room and context were encoded in fMRI signals, we found that map-like representations of the environment emerged in both hippocampus and neocortex," Muhie-Karbe, Sheahan, and their colleagues wrote.

"Cognitive maps in hippocampus and orbitofrontal cortices were compressed so that locations cued as goals were coded together in neural state space, and these distortions predicted successful learning. This effect was captured by a computational model in which current and prospective locations are jointly encoded in a place code, providing a theory of how goals warp the neural representation of space in macroscopic neural signals."

Essentially, Muhie-Karbe, Sheahan, and their colleagues found that the environment that was virtually accessed by participants was encoded in the form of a map in some parts of their brains, particularly the hippocampus and neocortex. Interestingly, however, these cognitive maps were somewhat compressed, coding locations that were relevant to the goal they were trying to achieve together.

These results shed new light on the neural underpinnings of goal-directed decision-making, suggesting that the brain could utilize compression mechanisms to contextually modulate sensory information during decision-making to achieve a specific goal. In the future, new studies could further investigate these compression processes, which could lead to fascinating new discoveries.

More information: Paul S. Muhle-Karbe et al, Goal-seeking compresses neural codes for space in the human hippocampus and

orbitofrontal cortex, *Neuron* (2023). DOI: 10.1016/j.neuron.2023.08.021 . www.sciencedirect.com/science/ … ii/S0896627323006323