

Reward morphs non-spatial cognitive maps in humans

Nir Moneta (moneta@mpib-berlin.mpg.de)

Max Planck Institute for Human Development
Lentzeallee 94, 14195 Berlin, Germany

Christian F. Doeller (doeller@cbs.mpg.de)

Max Planck Institute for Human Cognitive and Brain Sciences
Stephanstraße 1a, 04103 Leipzig, Germany

Charley M. Wu (charley.wu@uni-tuebingen.de)

University of Tübingen
Maria-von-Linden-Str. 6, 72076 Tübingen, Germany

Nicolas W. Schuck (schuck@mpib-berlin.mpg.de)

Universität Hamburg
Von-Melle-Park 5, 20146 Hamburg, Germany

Abstract

Animal work has shown that neural representations of a spatial cognitive map are affected by reward. Here we ask how non-spatial cognitive maps in humans are affected by reward. Sixteen participants performed two sessions of a perceptual discrimination task, before and after extensive reward learning. To identify changes in perceptual cognitive maps post reward learning we used a computational approach of participants behavior inspired by Maximum Likelihood Difference Scaling models. This showed that reward learning increased the spacing between stimuli in the rewarding area, akin to increased perceptual discrimination. These results indicate that reward affects non-spatial cognitive maps and suggest accompanying neural changes.

Keywords: cognitive maps; reward; grid-like encoding;

Introduction

Animals and humans maintain a cognitive map of the environment through a set of medio-temporal and medio-prefrontal representations, including hippocampal place cells and entorhinal grid cells (O'Keefe & Dostrovsky, 1971; Moser, Kropff, & Moser, 2008; Doeller, Barry, & Burgess, 2010). Collectively, these cell types are seen as a coordinate system for spatial as well as non-spatial navigation (Constantinescu, O'Reilly, & Behrens, 2016; Kaplan, Schuck, & Doeller, 2017; Sharpe et al., 2019; Theves, Fernandez, & Doeller, 2019).

Animal studies have found that activity of place and grid cells might adapt to changes in the environment and be influenced by reward locations. Place cells can conditionally respond only in a certain context, a larger number of place cells represent areas around reward versus non-reward locations, and grid cells change in firing rate and location too (Moita, Rosis, Zhou, LeDoux, & Blair, 2003; Dupret, O'neill, Pleydell-Bouverie, & Csicsvari, 2010; Boccara, Nardin, Stella, O'Neill, & Csicsvari, 2019; Sosa & Giocomo, 2021). Here we ask whether such changes occur in humans performing a non-spatial task, and how they affect behavior.

The goal of our study was therefore to investigate how reward learning morphs abstract cognitive maps by measuring perceived similarities of non-spatial stimuli.

Methods

Participants ($N = 16$) performed four sessions across two days. The 1st session consisted of a perceptual discrimination task (PRE-session, Fig. 1a). On each trial, first a target

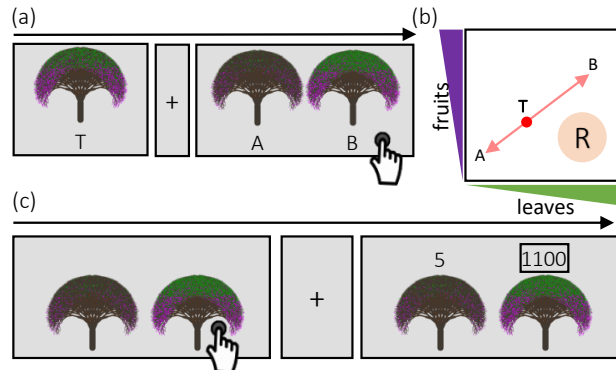


Figure 1: **Task design.** a) Perceptual discrimination task (1st and 4th sessions). On each trial, participants selected which of two reference trees (A, B) was more similar to a target tree (T). b) We hypothesize an underlying cognitive map, where each tree can be mapped by the number of leaves (x-axis) and number of fruits (y-axis). c) In the 2nd and 3rd sessions participants learned to associate a specific region (marked as R in b) in the perceptual space with reward through trial and error in a two forced choice task. They were asked to select one of two trees, and received feedback in the form of reward.

tree was presented, followed by two reference trees, and participants judged which of the two reference trees was more similar to the previously presented target. Trees were characterized by their number of leaves and fruits, such that each tree can be conceptualized as a particular point in a leaf/fruit space (Fig. 1b). Afterwards, participants performed two sessions where they learned through trial and error in a two forced choice task to associate a specific location in the perceptual space with reward, i.e. some fruit/leaf proportions were rewarded while others were not (akin to introducing a reward to a location in the cognitive map, Fig. 1c). Participants were exposed equally to the entire map during reward learning, preventing over-exposure to the rewarding area. In the final POST session, participants repeated the initial discrimination task. Our main question is thus how the reward learning changed behavior in the POST compared to the PRE session. We acquired fMRI during all sessions.

Results

We fitted perceptual models to participants' choices to capture the subjective perceptual distance between trees, separately for the PRE and POST sessions. Specifically, we fitted a logistic regression in a two-step manner. Inspired by Maximum

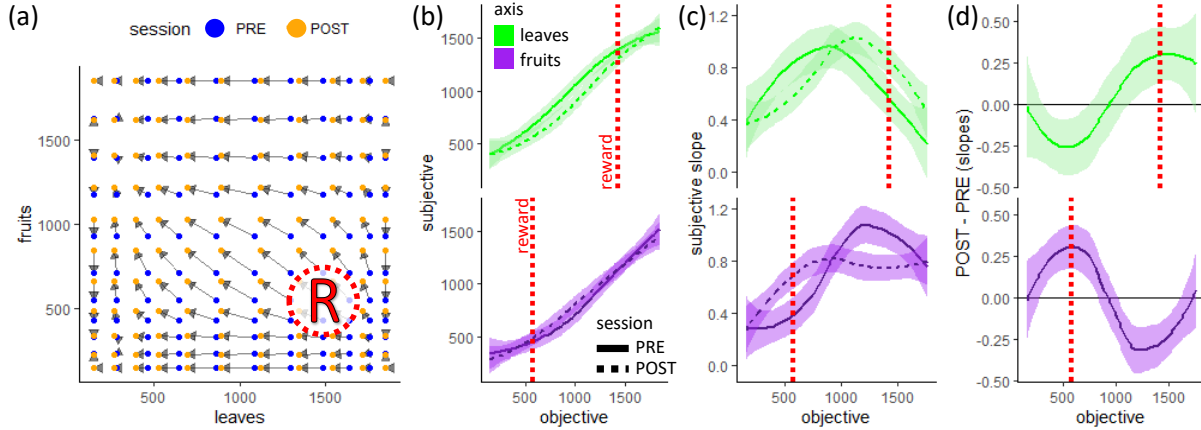


Figure 2: **Behavioral results, N=16** **a)** Construction of the cognitive map based on participants' choices in the PRE (blue) and POST (orange) sessions. Gray arrows indicate the change between sessions and rewarded area is marked with R. **b)** Averaged translation function across participants, separate for leaves (green, x axis in the 2D cognitive map) and fruits (purple, y axis in the 2D cognitive map) for PRE-session (solid line) compared to POST (dashed line). The x-axis represents the objective space (the true number of leaves/fruits) and the y-axis the subjective space (the perceived distance). Red vertical lines indicate the location of the most rewarding tree on each axis in the 2nd and 3rd sessions. **c)** The slope of the translation function in panel b. Areas with higher slope indicate increased subjective distance since a change in objective distance results in a bigger change in subjective distance. **d)** Comparing the slope of the subjective distance between PRE and POST sessions taken from panel c reveals an increase in slope for each axis, specifically in the most rewarding area.

Likelihood Difference Scaling (MLDS; Knoblauch & Maloney, 2008), we utilized **translation functions** using four free parameters to capture the subjective perceptual distances separately for each axis:

$$f_x = \frac{1}{1 + e^{-p_1(x-p_2)}}, \quad f_y = \frac{1}{1 + e^{-p_3(y-p_4)}}, \quad (1)$$

where p_1, p_2 govern the subject-specific transformation of perceptual distances for leaves (x-axis), while p_3, p_4 is for fruits (y-axis).

We then fitted a logistic regression to model participants' choices as a function of the transformed distance D:

$$\log\left(\frac{P_A}{P_B}\right) = \beta_0 + \beta_1 D, \quad (2)$$

where D represents the euclidean distance of tree A to the target minus the euclidean distance of tree B to the target, such that higher D indicates that the distance from B to T is greater than A to T. In other words, D captures evidence for selecting tree A, because it is more similar to the target. Note that the locations used to compute D are the fitted locations from Eq. 1.

Our main question was whether PRE vs POST changes in participants' cognitive maps indicate a morph around the location of the reward. Since the reward was located in the bottom right corner, we expected that the fitted translation functions indicated a greater slope in the same area. In line with this idea, we found that the model derived translation functions revealed that trees in the POST-session increased their distances in the area of the previously experienced reward (Fig. 2a). Further investigation of the translation functions confirmed that the increase in perceived distances appeared on

both axes, as expected (Fig. 2b-d). Note that half the participants had the rewarding region in the top left region, and results were flipped accordingly for better visualization of group effects.

Specifically, based on the fitted parameters, we extracted the translation function of each subject and each session for each axes which indicates the transformation from objective difference in leaves and fruits to the subjective, perceived, difference. We averaged the translation functions across all subjects, separated for session and axis (Fig. 2b) and examined the slope of the function (Fig. 2c). Because of the nature of the translation function, higher slope indicates that a change in the objective distance results in a higher change in perceived distances. Inspection of the change in slope from PRE to POST session revealed an increase in slope (i.e. increase in perceived distances) for each axis, specifically in the area that was previously rewarded (Fig. 2c, most rewarded area marked by dashed red line). We formally tested this change in slope and found that the slope difference of POST minus PRE, averaged for both axes, was significantly greater than 0 ($t_{15} = 2.04, p = .029$).

Conclusions

Our result show that reward learning increased the spacing between stimuli in the rewarding area. This suggests that the underlying neural representation of the cognitive map might change as well, akin to the findings from animal literature on the influence of reward on representations during spatial navigation.

Acknowledgments

NWS was funded by an Independent Max Planck Research Group grant awarded by the Max Planck Society (M.TN.A.BILD0004) and a Starting Grant from the European Union (ERC-2019-StG REPLAY-852669) and the Federal Ministry of Education and Research (BMBF) and the Free and Hanseatic City of Hamburg under the Excellence Strategy of the Federal Government and the Länder. NM was funded by and is grateful for Einstein Center for Neuroscience (ECN) Berlin throughout this study. CMW is supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A and funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy–EXC2064/1–390727645. CFD's research is supported by the Max Planck Society, the European Research Council (ERC-CoG GEOCOG 724836), the Kavli Foundation, the Jebsen Foundation, Helse Midt Norge and The Research Council of Norway (223262/F50; 197467/F50). We thank Gregor Caregnato for help with participant recruitment, Lydia Brundisch, Neele Elbersgerd, Sonali Beckmann and Nadine Taube for help with data acquisition and the Neurocode lab for support and contributions throughout this project.

References

- Boccaro, C. N., Nardin, M., Stella, F., O'Neill, J., & Csicsvari, J. (2019). The entorhinal cognitive map is attracted to goals. *Science*, *363*(6434), 1443–1447.
- Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science*, *352*(6292), 1464–1468.
- Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature*, *463*(7281), 657–661.
- Dupret, D., O'Neill, J., Pleydell-Bouverie, B., & Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nature neuroscience*, *13*(8), 995–1002.
- Kaplan, R., Schuck, N. W., & Doeller, C. F. (2017). The role of mental maps in decision-making. *Trends in Neurosciences*, *40*(5), 256–259.
- Knoblauch, K., & Maloney, L. T. (2008). Mlds: Maximum likelihood difference scaling in *r*. *Journal of Statistical Software*, *25*, 1–26.
- Moita, M. A., Rosis, S., Zhou, Y., LeDoux, J. E., & Blair, H. T. (2003). Hippocampal place cells acquire location-specific responses to the conditioned stimulus during auditory fear conditioning. *Neuron*, *37*(3), 485–497.
- Moser, E. I., Kropff, E., & Moser, M.-B. (2008). Place cells, grid cells, and the brain's spatial representation system. *Annu. Rev. Neurosci.*, *31*, 69–89.
- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain research*.
- Sharpe, M. J., Stalnaker, T., Schuck, N. W., Killcross, S., Schoenbaum, G., & Niv, Y. (2019). An integrated model of action selection: distinct modes of cortical control of striatal decision making. *Annual review of psychology*, *70*, 53–76.
- Sosa, M., & Giacomo, L. M. (2021). Navigating for reward. *Nature Reviews Neuroscience*, *22*(8), 472–487.
- Theves, S., Fernandez, G., & Doeller, C. F. (2019). The hippocampus encodes distances in multidimensional feature space. *Current Biology*, *29*(7), 1226–1231.