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

Luianta Verra (verra@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 (nicolas.schuck@uni-hamburg.de)
Universität Hamburg
Von-Melle-Park 5, 20146 Hamburg, Germany

Abstract

This study examines the effect of reward on non-spatial cognitive maps in humans, building on previous research showing reward influences spatial maps in animals. Seventy-two participants (38 undergoing MRI), completed a perceptual discrimination task pre- and post-reward learning. Post-reward exposure, performance improved in previously rewarded areas, with effects generalizing to non-rewarded map areas. Behavioral findings suggest reward learning alters psychological distances between stimuli, corresponding to simulations of place field shifting towards rewarded locations, akin to gravitational pulling. Preliminary fMRI data supports this interpretation, with similar representational shifts in hippocampal representations, but mixed results in the medial-orbitofrontal and visual cortex. This suggests reward affects non-spatial cognitive maps and neural representations.

Keywords: Cognitive maps; Representational change; fMRI;

Introduction

Humans and animals maintain an internal representation of the environment, known as a cognitive map, through mediotemporal and medio-prefrontal representations, including hippocampal place cells and entorhinal grid cells (O'Keefe & Dostrovsky, 1971; Moser, Kropff, & Moser, 2008; Doeller, Barry, & Burgess, 2010). These cell types collectively form a coordinate system for both spatial and non-spatial navigation (Behrens et al., 2018; Kaplan, Schuck, & Doeller, 2017). Cognitive maps dynamically adapt to environmental changes, with place and grid cells adjusting to new shapes and salient locations (O'Keefe & Burgess, 1996; Krupic, Bauza, Burton, & O'Keefe, 2018), reflecting a subjective, experience-driven representation (Stachenfeld, Botvinick, & Gershman, 2017; Ginosar, Aljadeff, Las, Derdikman, & Ulanovsky, 2023). Reward prediction is crucial for goal-directed behavior and plays an important role in cognitive maps (Moneta, Grossman, & Schuck, 2024), potentially integrated into this system by increasing firing rates and shifting representational fields towards the reward (Zeithamova, Gelman, Frank, & Preston, 2018; Sosa & Giocomo, 2021; Boccara, Nardin, Stella, O'Neill, & Csicsvari,

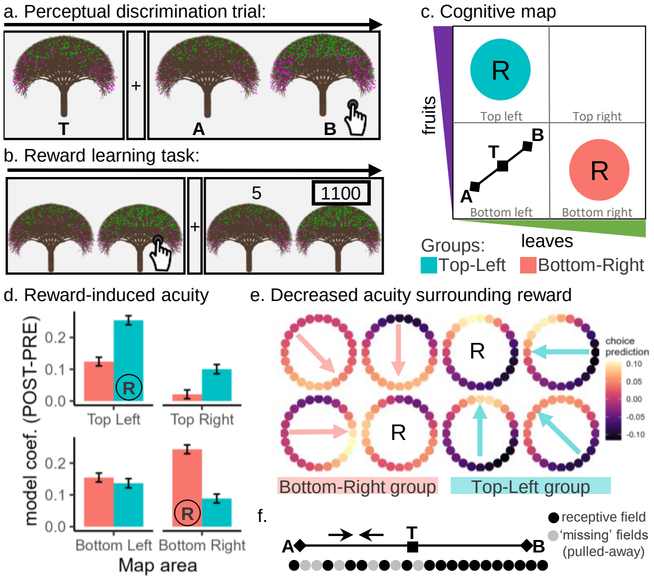

Figure 1: **a.** In the perceptual discrimination task (1st and 4th sessions), participants chose which of two reference trees (A, B) was more similar to a previously presented target tree (T). **b.** In the 2nd and 3rd sessions only, participants learned to associate specific trees with reward through trial and error in a two-choice task. They were presented with two trees, selected one and received feedback as a reward. **c.** Each tree is mapped by the number of leaves (x-axis) and fruits (y-axis). We sampled the space equally in four areas (gray lines), ensuring equal exposure across the map. Participants were split into two groups with rewarding trees (only in 2nd & 3rd sessions) in either the Top-Left area (teal) or Bottom-Right area (red). **d.** Contrasted model coefficients (y-axis, POST-PRE), split by map areas (panels) for each group (x-axis), showed a stronger increase in perceived distances in each group's rewarded area (marked by (R)). **e.** Contrasting angle-dependent model prediction (POST-PRE) revealed an increase in choice bias in non-rewarded areas pointing towards the reward. Arrows within each circle indicate the angle pointing towards the reward where participants' choices *increased*. **f.** We reasoned that a decrease in field density would decrease discriminability, making trees appear more similar, e.g., perceiving A as closer to T although objectively A and B are equidistant to T.

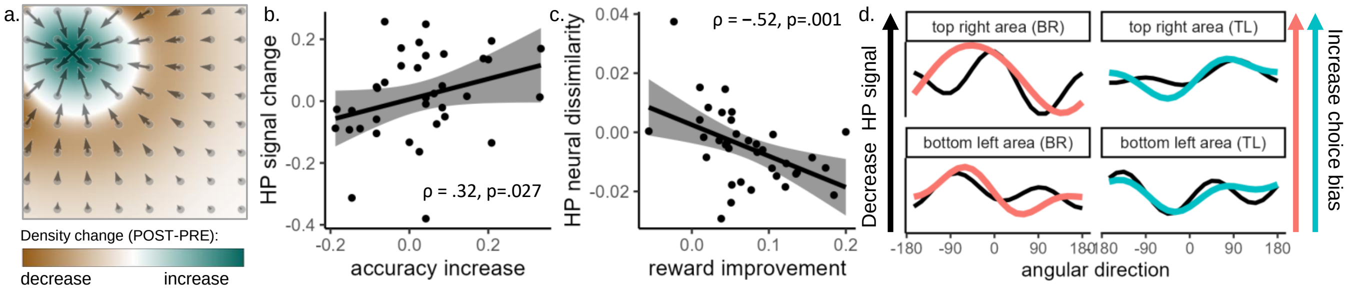

Figure 2: **a.** Simulation of hippocampus place cells showed that after reward-driven gravitational pull (illustrated by arrows), the increase in field density around the reward (teal) is accompanied by a decrease in adjacent areas (brown) compared to further areas (gray). **b.** Participants with a stronger increase in hippocampus univariate signal in the rewarded area (averaged betas: POST-PRE, rewarded area - opposite, y-axis) also showed increased behavioral accuracy for the same trial comparison (x-axis). Stats indicate a one-sided Spearman's rank correlation test (non-parametric, robust). **c.** Participants with a stronger decrease in neural dissimilarities around the reward location (1-Pearson, same comparison as b, y-axis) also showed a stronger improvement in reward learning across days (day2 - day1 accuracy, x-axis). Stats as b. **d.** In reward-adjacent areas (upper row: top-right area, lower row: bottom left area), hippocampus showed a decrease in univariate signal (y-axis, black line) for trees pointing at the same angular direction (x-axis) where behavior showed a decrease in perceived distances (y-axis, choice bias by group, for colors see Fig.1c). Mean-Squared-Error between lines is significantly below permutation tests' 5th percentile (not presented).

2019). Our study aims to investigate how reward learning morphs abstract cognitive maps by measuring perceived similarities of non-spatial stimuli.

Results

Over two days, 74 participants (38 with MRI) completed a perceptual task (PRE), a reward learning task, and another perceptual task (POST). We examined behavioral and neural changes in the main perceptual task, PRE, and POST reward exposure, where participants judged tree similarity in a 2D map of leaf/fruit space (Fig. 1a-c.). Participants exceeded chance in both sessions (PRE: $\mu = .765$, $\sigma = .046$, POST: $\mu = .787$, $\sigma = .044$) and improved from PRE to POST (paired t-test: $t(73) = 3.9$, $p < .001$). They also learned to identify the most rewarding tree more accurately from day one to two (paired t-test: $t(73) = 11.05$, $p < .001$).

Reward exposure enhanced local acuity at the expense of its immediate surroundings. As predicted, exposure to reward increased local perceptual acuity in the POST task, with accuracy improving significantly only in the rewarded area for each group (PRE-POST change: bottom right group: $\mu = 0.038$, $\sigma = 0.016$, p=.017, top left group: $\mu = 0.035$, $\sigma = 0.016$, p=.03, Bonferroni corrected, $p > .05$ for all other areas). A logistic regression model fitted to participants' choices included several factors for each area of the map. First, we computed the distance between tree B and the target minus the distance of tree A to the target. This regressor indicates changes in sensitivity in each area, where a higher fitted coefficient indicates participants were more sensitive to even more minor perceptual differences (similar to a temperature parameter). Confirming the accuracy results, comparing coefficients POST minus PRE revealed an increase in sensitivity specific in the rewarded area of each group, akin to perceiving trees there as more distant (post-hoc marginal means z-test $\text{est} = 2.348$, $SE = 0.286$, $p < .001$, Fig. 1d.).

Second, directional regressors for each area captured biases to choose trees in specific directions, *irrespective of correctness*. Model predictions indicated a choice bias for trees in directions pointing towards the reward in adjacent areas, occurring despite the absence of rewards in these areas in previous sessions and specifically in the POST session, dependent on group-specific reward location (Fig. 1e, arrows point towards the reward direction in a group & area-specific manner). This angular choice bias suggests perceived distances in specific directions are smaller, potentially corresponding to a decrease in field density in the same areas (Bellmund et al., 2020, Fig. 1f).

Simulation of reward as gravitational pull best explains behavior. We simulated potential representational changes elicited to density of place cells by reward. We predicted that reward at a stable location would elicit a gravitational pull, attracting receptive fields of cells (inspired by Ginosar et al., 2023). In addition to the increase of field density in the rewarded area, the pulling hypothesis predicts an angle-dependent *decrease* in field density pointing towards the reward in all non-rewarded areas (brown areas in Fig. 2a), which, as predicted, corresponds to the behaviorally observed decrease in perceived distances at the same directions.

Preliminary fMRI results show hippocampus signal best explained by the pulling hypothesis. We analyzed fMRI data recorded during target tree presentation and found that hippocampus representational changes support predictions made by simulating the pulling hypothesis. First, although we did not find a significant increase in the group-level ROI, across subjects, an increase in signal in the rewarded area related to a rise in accuracy among trials in that area (Fig. 2b). Second, comparing neural dissimilarities change in different areas of the map, we found a significant decrease in dissimilarity among trees close to the reward area ($t_{(35)} = -2.38$,

p=.012). This decrease was even stronger for participants who showed higher across-day improvement in reward learning (Fig. 2c). Lastly, focusing on reward-adjacent areas only, we found a group-dependent *decrease* in hippocampus signal which significantly related to the same angular-dependent behavioral choice bias (Fig. 2d.), supporting a decrease in field density in areas adjacent to reward, pointing towards it.

Conclusion

In summary, both behavioral and fMRI results in the hippocampus indicate that reward enhances perceptual discrimination by distance-scaled pull of representational fields toward it. This leads to systematic changes in an abstract cognitive map, which generalized both in time, to subsequent tasks where the reward wasn't present, and in space, to areas of the map that were never rewarded. Future work will focus on modeling uni- and multivariate signals and exploring grid-like coding and reward influences during decision-making.

Acknowledgments

N.W.S. is funded by a Starting Grant from the European Union (ERC-2019-17 StG REPLAY-852669) 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 partially funded by the Einstein Center for Neuroscience (ECN) Berlin. LV is funded by the International Max Planck Research School on the Life Course. 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-CoGGEOCOG 724836), the Kavli Foundation, the Jebsen Foundation, Helse Midt Norge and The Research Council of Norway (223262/F50; 197467/F50

References

Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018, October). What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron*, *100*(2), 490–509. doi: 10.1016/j.neuron.2018.10.002

Bellmund, J. L., De Cothi, W., Ruiter, T. A., Nau, M., Barry, C., & Doeller, C. F. (2020). Deforming the metric of cognitive maps distorts memory. *Nature human behaviour*, *4*(2), 177–188.

Boccara, 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.

Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature*, *463*(7281), 657–661.

Ginosar, G., Aljadeff, J., Las, L., Derdikman, D., & Ulanovsky, N. (2023). Are grid cells used for navigation? on local metrics, subjective spaces, and black holes. *Neuron*, *111*(12), 1858–1875.

Kaplan, R., Schuck, N. W., & Doeller, C. F. (2017). The role of mental maps in decision-making. *Trends in Neurosciences*, *40*(5), 256–259.

Krupic, J., Bauza, M., Burton, S., & O'Keefe, J. (2018). Local transformations of the hippocampal cognitive map. *Science*, *359*(6380), 1143–1146.

Moneta, N., Grossman, S., & Schuck, N. W. (2024). Representational spaces in orbitofrontal and ventromedial prefrontal cortex: task states, values, and beyond. *Trends in Neurosciences*.

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., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, *381*(6581), 425–428.

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain research*.

Sosa, M., & Giocomo, L. M. (2021). Navigating for reward. *Nature Reviews Neuroscience*, *22*(8), 472–487.

Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature neuroscience*, *20*(11), 1643–1653.

Zeithamova, D., Gelman, B. D., Frank, L., & Preston, A. R. (2018). Abstract representation of prospective reward in the hippocampus. *Journal of Neuroscience*, *38*(47), 10093–10101.