

Development of structure inference contributes to age-related differences in exploration

Nora C. Harhen (nharhen@uci.edu)

Department of Cognitive Sciences, University of California, Irvine
Irvine, CA, 92697, USA

Catherine A. Hartley* (cate@nyu.edu)

Department of Psychology, New York University
New York, NY, 10003, USA

Aaron M. Bornstein* (aaron.bornstein.edu)

Department of Cognitive Sciences, University of California, Irvine
Center for the Neurobiology of Learning and Memory, University of California, Irvine
Irvine, CA, 92697, USA

*denotes equal senior authorship

Abstract

How we represent an environment’s structure influences how we explore it. Here, we asked whether developmental differences in structure learning contribute to known age-related differences in exploration. Participants between ages 8 and 25 completed a patch foraging task previously shown in adults to elicit individual differences in structure inference. We compared their behavior to the predictions of a Bayesian structure learning model. We found that, consistent with normative models, all participants were able to adapt their exploration with respect to their uncertainty about the environment’s structure. However, we found that older participants overharvested more when it was more profitable, consistent with their representing the patch structure in a more complex, but ultimately veridical, fashion, compared to younger participants. These results suggest a role for changes in representation in underlying the developmental shift away from exploration towards exploitation.

Keywords: Foraging; exploration; structure learning; development

Introduction

Childhood and adolescence are developmental stages characterized by heightened exploration. While previous research has primarily focused on how changes in reinforcement learning algorithms might account for developmental differences in exploration (Giron et al., 2023; Schulz, Wu, Ruggeri, & Meder, 2019; Somerville et al., 2017; Meder, Wu, Schulz, & Ruggeri, 2021), learning to represent the structure of the environment being explored is another potential source of difference. Structure learning and mental model formation exhibit developmental changes into young adulthood (Schlichting, Guarino, Schapiro, Turk-Browne, & Preston, 2017; Pudhiyidath, Roome, Coughlin, Nguyen, & Preston, 2020). However, conventional exploration tasks lack the environmental complexity necessary to measure interactions between structure inference and exploration. Patch foraging tasks may present a more complex and “naturalistic” choice context than standard decision making tasks as they allow for environments with richer structure (Mobbs, Trimmer, Blumstein, & Dayan, 2018). Here, using a combination of behavioral methods and computational modeling, we ask – do developmental differences in structure learning and use underlie differences in exploration?

Methods

Participants and task 252 participants between the ages of 8 and 25 completed a patch foraging task used in a prior adult study (Harhen & Bornstein, 2023). Participants traveled to different planets to mine for space

gems (Fig 1a-c). On each planet, participants decided between staying to dig from the current planet’s depleting mine or incurring a time cost to travel to a new planet with a replenished mine. Mimicking natural environments, planets varied in their quality, as defined by their average depletion rate per dig. Planet quality was correlated in time – new planets were most likely to belong to the same type as the last planet (non-switch), but occasionally, there were rare “switches” in planet type. Critically, this structure was not explicitly signaled to participants, requiring them to infer it solely from the sequence of rewards they received.

Models In our model (Harhen & Bornstein, 2023), we relax MVT’s assumption that the forager has perfect knowledge. We use a Chinese Restaurant Process (CRP; Aldous, 1985) to model a forager rationally inferring the environment’s underlying structure. The CRP is distinguished by its prior which always maintains the possibility of a new planet type being encountered, with probability proportional to the parameter α . When $\alpha > 0$, the forager’s representational complexity is allowed to grow as experience warrants it. This supports more precise predictions of how the planet will deplete over time, informing the value they estimate for staying. When $\alpha = 0$, the forager assumes a simplistic structure, representing a single planet type. The forager can never be fully certain that their representation is truly accurate. In our model, this uncertainty is counteracted with an adaptive planning horizon (Jiang, Kulesza, Singh, & Lewis, 2015). Previously, we have shown that the dynamics of participants’ choices in this task can be explained by a reduction in planning horizon at points of uncertainty. While this computation is included in our model, due to space constraints, here we only report findings with respect to the structure learning computation.

We compared how well two variants of the model described participants’ choices using 10-fold cross-validation: one variant in which $\alpha = 0$ and another in which $\alpha = 1$. We also compared these models against two other models considered in Constantino & Daw (2015). One model couples a MVT-based decision rule with error-driven learning of the environment’s overall distribution of rewards and the other uses a temporal-difference algorithm to learn state-action values.

Results

Our Bayesian structure-learning model predicts that a forager’s response to planet richness depends on how they represent the environment (Fig 1d-e). Notably, rich planets should produce the most patent response differences – foragers representing multiple planet types should overharvest while those representing a single type should underharvest. Examining effects of age and




Figure 1: **A. Patch foraging task** On each trial, participants decided between staying to dig from a depleting gem mine or incurring a time cost to travel to a new planet. **B. Environment structure** Planets belonged to one three planet types each with their own characteristic distribution over depletion rates. **C. Environment dynamics.** A new planet had an 80% probability of being the same type as the prior planet (“no switch”). However, there was a 20% probability of transitioning or “switching” to a planet of a different type. **D. Structure learning** With our structure learning model, we simulated agents with different settings of α completing the task. Bar heights indicate the agent’s planet residence time (PRT) and the dotted lines indicate the MVT-optimal PRT. Our model predicts that foragers’ inferring multiple planet types (structure learners with $\alpha > 0$) should overharvest on rich planets while foragers inferring a single type (non-structure learners with $\alpha = 0$) should underharvest. **E. Uncertainty adaptive planning** Our model predicts that a forager who adapts their planning horizon with respect to their uncertainty will increase their overharvesting in response to switches in planet type. Foragers who do not adapt will show no modulation of their overharvesting.

its interactions, we found that on rich planets older participants explored less than younger participants, overharvesting more ($\beta=0.36, p = .0078$). Based on our model predictions, this is consistent with age-related improvements in structure inference.

If participants engage in uncertainty-adaptive planning, then their overharvesting should increase at points when the planet type switches (Fig 2b). Indeed, participants’ overharvesting did increase at switch points ($\beta=0.31, p < .001$). We did not observe any age-related differences in sensitivity to switch points ($\beta=-0.0094, p = .81$).

Model comparison revealed that the $\alpha = 1$ model provided the best fit for the greatest proportion of participants in all age groups (children: 51%; adolescents: 59%; adults: 66%). Notably, however, a smaller proportion of children were best fit by this model. To explore this further, we took the difference in cross-validation scores between the $\alpha = 0$ and the $\alpha = 1$ models for each participant. A positive value would indicate the the participants’ choices were better described by the $\alpha = 1$ model. We found that the difference in scores grew increasingly positive with age (Kendall’s $\tau = .18, p < .001$), further substantiating an improvement in structure inference with age.

Conclusion

In all age groups, participants’ choices revealed a sensitivity to uncertainty, being less likely to explore a new option when uncertain. However, older participants choices’ were more aligned with a complex, multi-planet representation of the environment. Taken together, our results suggest that with age people represent novel en-




Figure 2: **A.** Relative to MVT-optimality, participants overharvested across all planet types but to varying extents. On rich planets, older participants overharvested to a greater extent than younger participants. **B.** Participants in all age groups increased the extent of their overharvesting following switches in planet type.

vironments with greater complexity, consequently altering their exploratory decisions.

Acknowledgements

This work was supported by the NIMH – R01 MH126183 to CAH, P50 MH096889 to AMB (PI: TZ Baram), F31 MH134620 to NCH – the Templeton World Charity Foundation (to CAH), and the Department of Defense (NDSEG fellowship to NCH).

References

- Aldous, D. J. (1985). Exchangeability and related topics. In D. J. Aldous, I. A. Ibragimov, & J. Jacod (Eds.), *École d'été de probabilités de Saint-Flour XIII — 1983* (pp. 1–198). Springer Berlin Heidelberg.
- Giron, A. P., Ciranka, S., Schulz, E., van den Bos, W., Ruggeri, A., Meder, B., & Wu, C. M. (2023, August). Developmental changes in exploration resemble stochastic optimization. *Nat Hum Behav*.
- Harhen, N. C., & Bornstein, A. M. (2023, March). Overharvesting in human patch foraging reflects rational structure learning and adaptive planning. *Proc. Natl. Acad. Sci. U. S. A.*, *120*(13), e2216524120.
- Jiang, N., Kulesza, A., Singh, S., & Lewis, R. (2015). *The dependence of effective planning horizon on model accuracy*.
- Meder, B., Wu, C. M., Schulz, E., & Ruggeri, A. (2021, July). Development of directed and random exploration in children. *Dev. Sci.*, *24*(4), e13095.
- Mobbs, D., Trimmer, P. C., Blumstein, D. T., & Dayan, P. (2018, July). Foraging for foundations in decision neuroscience: insights from ethology. *Nat. Rev. Neurosci.*, *19*(7), 419–427.
- Pudhiyidath, A., Roome, H. E., Coughlin, C., Nguyen, K. V., & Preston, A. R. (2020). Developmental differences in temporal schema acquisition impact reasoning decisions. *Cogn. Neuropsychol.*, *37*(1-2), 25–45.
- Schlichting, M. L., Guarino, K. F., Schapiro, A. C., Turk-Browne, N. B., & Preston, A. R. (2017, January). Hippocampal structure predicts statistical learning and associative inference abilities during development. *J. Cogn. Neurosci.*, *29*(1), 37–51.
- Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019, November). Searching for rewards like a child means less generalization and more directed exploration. *Psychol. Sci.*, *30*(11), 1561–1572.
- Somerville, L. H., Sasse, S. F., Garrad, M. C., Drysdale, A. T., Abi Akar, N., Insel, C., & Wilson, R. C. (2017, February). Charting the expansion of strategic exploratory behavior during adolescence. *J. Exp. Psychol. Gen.*, *146*(2), 155–164.