Latent cause inference as an efficient and flexible learning rule for cognitive graphs

Jungsun Yoo (jungsun.yoo@uci.edu)
Department of Cognitive Sciences
University of California, Irvine
Irvine, CA 92697

Dale Zhou (dale.zhou@uci.edu)
Department of Cognitive Sciences
Department of Neurobiology and Behavior
University of California, Irvine
Irvine, CA 92697

Aaron M. Bornstein (aaron.bornstein@uci.edu)
Department of Cognitive Sciences
Center for the Neurobiology of Learning and Memory
University of California, Irvine
Irvine, CA 92697

Abstract

In complex and continuous environments, agents acquire knowledge from experience by extracting structural – either spatial or temporal – regularities from sequential observations. Cognitive graphs based on Hidden-Markov Models (HMM) offer an efficient framework for elucidating the mechanisms of how agents learn latent environment structure. However, a leading algorithm implementing this approach ("Clone-Structured Cognitive Graphs"; CSCG) assumes a fixed allocation of neural resources to this problem, which may undermine biological plausibility and prove inefficient/inflexible for learning environments of unknown complexity. Here, we replace the fixed allocation of neural resources with a rational procedure that adapts the complexity of the internal representation according to nonparametric inference of latent structure. We demonstrate that, on the same benchmarks used to validate the original algorithm, our modification enhances efficiency without sacrificing performance. Our result suggests that this adaptive construction of cognitive graphs could potentially benefit learning in environments with unknown state-space complexity, and may thus provide a better explanation of behavior in resource-constrained biological organisms.

Keywords: clone-structured cognitive graphs; latent cause inference; Hidden Markov Models;

Introduction

In complex and continuous environments, agents acquire knowledge from experience by extracting structural – e.g. spatial or temporal – regularities from sequential observations. Hidden-Markov Models (HMM) offer an efficient framework for capturing how agents learn such latent structure (Mark, Moran, Parr, Kennerley, & Behrens, 2020; George et al., 2021). Recently, an extension of the HMM framework, called Clone-Structured Cognitive Graphs (CSCG), captures the ability to infer latent structure in the presence of aliased observations, using *clones* to distinguish latent states that have overlapping perceptual attributes; this model demonstrates the capability to transfer inferred structure across related tasks (George et al., 2021).

However, while this model captures behavior and function with remarkable efficacy, it assumes a fixed allocation of neural resources ("clones") to the problem (George et al., 2021), which raises concerns in at least two aspects. First, the optimal number of clones in each environment is unlikely to be known *a priori*. Second, we show here that learning efficacy depends greatly on this number of pre-allocated clones per observation. A potential solution to these issues is for individuals to flexibly allocate neural resources in response to experience.

We introduce a more flexible variant of the HMM-based cognitive graphs which creates new latent states as a function of their statistical likelihood, and allocates representational resources correspondingly. Specifically, we add an allocation rule based an infinite mixture model. This procedure has been shown to be a resource-optimal approach to clustering observations with unknown category structure (Dasgupta & Griffiths, 2022). Within the context of CSCG, whenever an observation is encountered, the mixture model guides the allocation of new resources according to a Chinese Restaurant Process (or Dirichlet prior) that dictates the likelihood of the observation being assigned to a new latent cause (hence, allocating a new clone).

We examine the impact of our model extension on efficiency and accuracy, in a benchmark simulation environment described in the original study (George et al., 2021). We show that our extended model shows comparable performance, operationally defined as the similarity of the inferred cognitive graphs to the ground-truth transitional structure of the environment, while being less demanding on memory resources (number of allocated clones). As a result, we believe that our approach provides a new avenue for biologically plausible, adaptive, representation learning which could be potentially applied to explaining behavior in complex and continuous environments.

Methods

To test the assumption that cognitive graphs created from scratch can still recover the ground-truth environmental structure, we devised a series of experiments where we compared our proposed vs. benchmark model under a spatial environment which had been used for validating the benchmark model (George et al., 2021).

Models

We compare two models – the benchmark HMM cognitive graphs (CSCG) and our extension incorporating latent cause inference.

Benchmark model Clone-structured cognitive graphs (CSCG) as described in George et al. (2021) was used as a benchmark model (refer to the Methods section of that paper for details of this model). We assigned 70 clones per each unique observation (280 in total) for the benchmark model, as reported in the original paper.

Our proposed model At the beginning, each model started with a randomly initialized transition matrix of dimension ($A \times T \times T$; where A = actions, T = latent causes of the observations), with one latent cause (or clone) assigned for each observation. Upon perceiving a sequence of observations during training, each new observation was stochastically allocated to a new latent cause (or "clones") according to a Chinese Restaurant Process prior (Gershman & Niv, 2012), expressed as:

$$P(c_t = k|c_{1:t-1}) = \begin{cases} \frac{N_k}{t-1+a} & \text{if } k \text{ is an old cause} \\ \frac{a}{t-1+a} & \text{if } k \text{ is a new cause,} \end{cases} \quad (1)$$

, where c_t=the latent cause for data at time point t, and

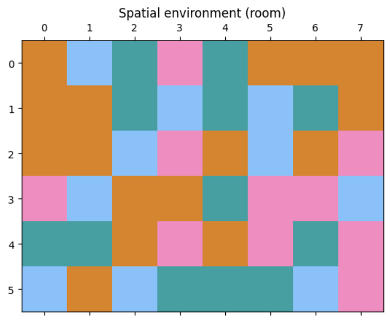

Figure 1: Spatial environment

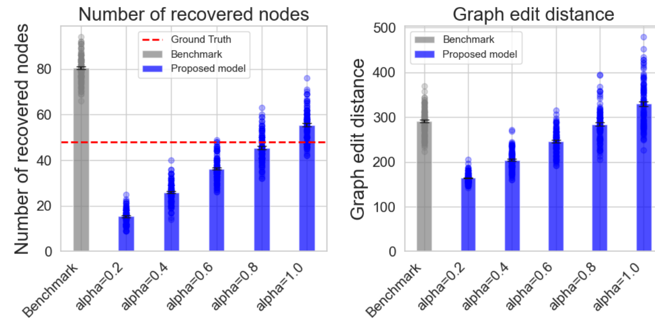

(a) Number of latent states (locations) recovered

(b) Graph edit distance (lower is better)

Figure 2: Performance of proposed vs. benchmark model. Individual data points represent the results from each experiment, and the error bars indicate standard error.

N_k=total number of allocations in cause k. For the simulations described here, we examined behavior of the model under five different values of the α parameter (0.2, 0.4, 0.6, 0.8, 1.0), which governs the propensity to create a new cluster (i.e., a higher α value is more likely to assign an observation to a new latent cause).

Experiments

Environments Each model was simulated in a spatial room that had 6×8 unique locations (cells; Fig. 1). In each cell, the agent could perceive one of four unique observations associated with that cell (orange, light blue, teal, and pink), thus making it a critical challenge for the agents to correctly identify the latent state (i.e., location) given aliased (perceptually overlapping) observations and actions. There were 4 actions that the agents could take in this environment (up, down, left, right), which were modeled as edges that link the latent states in cognitive graphs.

Training procedure The training data consisted of a sequence of 50,000 observations (i.e., pink, teal, sky blue, ...) and actions (i.e., up, left, down, ...) produced by random walks through the spatial environment. The models were trained using an expectation-maximization algorithm, where the log-likelihood of the inferred latent states given observations was maximized (George et al., 2021). Specifically, each model (our proposed and benchmark) was trained for 100 iterations. Each experiment was performed with 100 different random seeds, and the mean of the 100 simulations is reported.

Results

Performance

The performance of each simulation was assessed by (a) **accuracy**: the similarity of the constructed latent cognitive graph of the models with the ground-truth environment that produced the observations, and (b) **efficiency**: the judicious use of neural resources in capturing the underlying structure. The first was measured with two criteria: the similarity in the number of latent states recovered in the models vs. the ground-truth number of states, and the graph edit distance (structural similarity) between the inferred and ground-truth graph (Abu-Aisheh, Raveaux, Ramel, & Martineau, 2015). Efficiency was measured as the total number of allocated clones (latent states).

We find that, for most values of the critical parameter (α), our model recovered the underlying structure with comparable accuracy to the original model, while consistently using fewer neural resources.

Accuracy For nearly all tested values of α, the extended model also resulted in a more accurate, or equivalently accurate, reconstruction. Most of our model resulted in a smaller absolute difference in the number of nodes, relative to the standard CSCG (Fig. 2a; $\alpha_{0.4} : t_{(198)} = 13.47, p < .001; \alpha_{0.6} : t_{(198)} = 25.7, p < .001; \alpha_{0.8} : t_{(198)} = 37.2, p < .001; \alpha_{1.0} : t_{(198)} = 27.83, p < .001$). This was accompanied by generally lower graph edit distances for our model, relative to the standard CSCG (Fig. 2b; $\alpha_{0.2} : t_{(198)} = 38.42, p < .001; \alpha_{0.4} : t_{(198)} = 22.58, p < .001; \alpha_{0.6} : t_{(198)} = 10.84, p < .001; \alpha_{0.8} : t_{(198)} = 1.47, p = .07$; for $\alpha_{1.0}$ graph edit distance was greater than the benchmark: $t_{(198)} = 6.41, p < .001$).

Efficiency Across all tested values of α, the extended model used fewer nodes than the benchmark model (280 nodes; $\alpha_{0.2} : t_{(198)} = 797.78, p < .001; \alpha_{0.4} : t_{(198)} = 528.05, p < .001; \alpha_{0.6} : t_{(198)} = 446.52, p < .001; \alpha_{0.8} : t_{(198)} = 353.32, p < .001; \alpha_{1.0} : t_{(198)} = 304.08, p < .001$).

Discussion

Augmenting CSCG with an adaptive procedure for adjusting to complex structure improved its resource efficiency, while not sacrificing accuracy of the inferred structure. This work raises the possibility of a more biologically-plausible implementation of the algorithm, which may prove a superior explanation of human and animal learning. Future work will examine empirical fit to human behavior in exploring large, graph-structured spaces, as well as further augmenting the algorithm to improve efficient coding of environmental regularities.

References

Abu-Aisheh, Z., Raveaux, R., Ramel, J.-Y., & Martineau, P. (2015). An exact graph edit distance algorithm for solving pattern recognition problems. In *4th international conference on pattern recognition applications and methods 2015.*

Dasgupta, I., & Griffiths, T. L. (2022). Clustering and the efficient use of cognitive resources. *Journal of Mathematical Psychology*, *109*, 102675.

George, D., Rikhye, R. V., Gothoskar, N., Guntupalli, J. S., Dedieu, A., & Lázaro-Gredilla, M. (2021). Clone-structured graph representations enable flexible learning and vicarious evaluation of cognitive maps. *Nature communications*, *12*(1), 2392.

Gershman, S. J., & Niv, Y. (2012). Exploring a latent cause theory of classical conditioning. *Learning & behavior*, *40*, 255–268.

Mark, S., Moran, R., Parr, T., Kennerley, S. W., & Behrens, T. E. (2020). Transferring structural knowledge across cognitive maps in humans and models. *Nature communications*, *11*(1), 4783.