

ACT-R/S: Look Ma, No "cognitive-map"!

Anthony M. Harrison (anh23@pitt.edu)

Christian D. Schunn (schunn@pitt.edu)

Learning Research and Development Center, University of Pittsburgh
3939 O'Hara St., Pittsburgh, PA 15260 USA

Abstract

We present an outline of a model using ACT-R/S (spatial) in a foraging task. The model illustrates how spatially sensitive single-cell recording data can be obtained within an egocentric framework without having to rely upon an allocentric "cognitive-map"

Introduction

This paper presents the first application of a spatial reasoning extension to ACT-R (Anderson & Lebiere, 1998) to single-cell recording data in primates during navigation to demonstrate how our computational formalism provides a very different understanding of the representational assumptions required within the domain.

Since Tolman (1948) presented his cognitive map theory, representation of space has been a contentious issue. The choice of theoretical representations carries with it both explicit and implicit computational affordances and hinderances. The cognitive-map view posits the existence of externally based representations of space, either exocentric or allocentric (object centered). This requires a highly detailed and stable representation of the environment that is invariant with respect to the agent. This representational cost makes some processes such as novel short-cutting computationally trivial. By contrast, the egocentric view only requires the representation of objects relative to the agent. This representation, while easier to form and maintain, typically requires more online processing for tasks such as novel-short cutting.

The rodent place-cell and primate spatial-view cell phenomena (reviewed later) exhibit behaviors that are strongly suggestive of allocentric cognitive-maps (O'Keefe & Nadel, 1978). In this paper we argue the opposite position, detailing an egocentrically based system that not only exhibits the same neurological phenomenon, but is also representationally and computationally quite tractable, and more consistent with the finer details of the data.

The paper will present the architectural components that support navigation, followed by a description of a model utilizing the system. With these pieces in place, the spatial-view cell phenomenon will be discussed in the light of allocentric and egocentric representational schemes. The paper will then conclude with a discussion of how simulated single-cell recording in our ego-

centric framework can produce the same patterns as those seen neurologically.

Configural System

ACT-R/S (spatial) adds spatial reasoning functionality to ACT-R. Specifically, it divides space at the functional and computational boundaries of representation that are complimentary to those seen neurologically. ACT-R currently possesses a basic visual system for object identification (Byrne & Anderson, 2001). We add to this two additional systems: configural and manipulative. The manipulative system is concerned with representing the fine-grained spatial characteristics of objects to facilitate manipulation (e.g. grasping, rotation). The configural system is tuned to representing the relative, approximate configuration of objects in space.

The configural system within ACT-R/S is concerned with the configuration of objects in space relative to the self. As such it plays a key role in spatial navigation; however other configural tasks would also rely upon this system, such as facial recognition. Figure 1. presents a functional schematic of the configural system, each component of which will be detailed in subsequent sections. Conceptually, the system begins when the agent attends to an object. The configural representation is encoded and placed into the system's maintenance buffer. As the agent moves through the space, the attended representations within the buffer are updated by the path-integrator, ensuring a stable representation

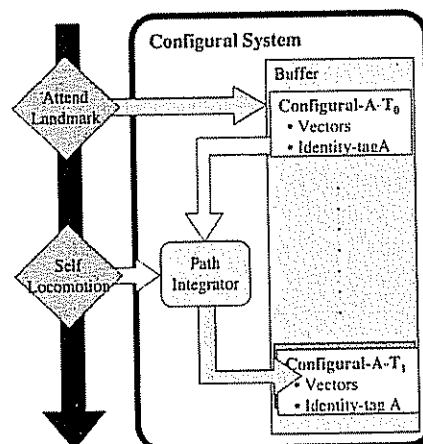


Figure 1. ACT-R/S configural system: Attending to a landmark encodes the representation into the buffer. Subsequent motion engages the path-integration system that updates/replaces active chunks in the buffer

Buffer Usage

As with ACT-R 5.0 generally, we assume the existence of a buffer associated with the configural system for temporary storage of chunks (declarative memories). Many have viewed the ACT-R buffer as a working memory store. However, recently there has been a shift to viewing the buffers as the interfaces to specific processing modules. These buffer/interfaces allow the production system to interact with the specialized modules within the architecture (e.g. the visual buffer allowing interaction with the visual system). Consistent with this framework, all manipulations of configural chunk representations take place within the configural buffer (i.e. attending, manipulating, & maintaining).

Unlike the rest of ACT-R's buffers, the configural buffer has a capacity of more than one chunk. This fairly significant divergence from traditional ACT-R is based on research showing that animals are able to use and track multiple landmarks at once (Kamil & Chang, 2001), and a computation evaluation of navigation showing the necessity of more than one landmark for the establishment of a unique location. The details of the buffer are not relevant for this discussion, so one can simply view the buffer as a "most-recently-used" cache with a capacity of two or three chunks. Once the capacity is exceeded (via attending to an additional object), the least recently used chunk will be pulled out to make room for the newest one. The old chunk will no longer be accessible through the buffer but continues to exist in declarative memory, permitting later retrieval.

Path-Integration

The next component that needs highlighting is the path-integration system. This is the component that is responsible for processing self-motion cues (in this case motor programs from the motor system) and adjusting representations based on that information. Path-integration can be viewed as the updating of one's own location in an allocentric space (see Gallistel, 1990 for review), or as the updating of multiple external egocentric representations (e.g. Klatzky, Beall, Loomis, Gollidge, & Philbeck, 1999). Since we are postulating an egocentric system of spatial representation, we implemented the latter.

At first glance this decision seems poor from a computational standpoint. Using an allocentric representation, the updating of one representation (the self in a fixed map) is much more tractable than updating all egocentric representations (i.e. all objects other than the fixed self). However, the path-integration system has no need to update all the egocentric representations. The architecture only needs to actively maintain those representations that are behaviorally significant. In other words, path-integration is only performed on those configural representations that are currently being attended

to, those that are in the buffer¹. We view the maintenance of a handful of behaviorally relevant egocentric as representations much simpler than the persistent storage and retrieval of detailed cognitive maps with multiple allocentric frames of reference.

Configural Representation

The configural system represents objects in space as egocentric vectors to the edges of an object. This gives the system just enough information to be able to move around, above, below, or to merely establish a relationship between multiple representations. The specific chunktype definition can be seen here:

Configural-rep	
Right-vector	(vector to the right side of the object)
Left-vector	(vector to the left side of the object)
Top-vector	(vector to the top of the object)
Bottom-vector	(vector to the bottom of the object)
Identity-tag	(visual-object / chunk)

One of the difficulties that an egocentric scheme presents is the existence of multiple representations of the same object, each from a different perspective. Multiple representations are linked together by the *identity-tag* slot value. This slot contains a token chunk that uniquely identifies the object being attended to in space. This chunk will often be the visual-object returned by the visual system's object identification process (which, in ACT-R parlance, would require an additional attentional request of the visual system to attend and encode the object's identity).

Foraging Model

With the architecture sketched out, we can now outline a basic model that utilizes the system. The simplest model in this situation is that of foraging, wherein the agent searches its environment for food. If food is seen, it navigates to the food and consumes it. If food is not seen, the agent either attempts to recall a previous food location or embarks on a more thorough exploration of the environment. This simple model handles two key features of navigation: basic taxon navigation (heading towards a visible location) and the retrieval of a memory of a specific location. We will sketch out each major section of the model as it usually runs and highlight the architectural activity along the way.

Food Search

The model begins by choosing between three alternatives: perform a visual search for food, try to recall a

¹ACT-R assumes that memories, once encoded, cannot be changed. In order to accommodate this constraint, the path-integration system actually replaces chunks in the configural buffer instead of updating them. This process only occurs once the difference between the integrated and the current representation exceed a certain threshold value.

successful location, or explore the environment. Given the low cost of a simple visual search, this is often chosen first. In this situation, the model will request that the configural system return to it the unattended object closest to the center of the agent's field of view. If an object is found, the representation is returned in the configural buffer. The model will then request that the visual system identify the object. The model repeats this process until a reward (food) has been identified.

Once a configural representation of the food object exists in the buffer, the agent is able to orient itself towards the destination (the location of the food). It then begins another search procedure like before, only this time it is attending to any object that will need to be navigated around (anything that intersects the volume that the agent would occupy as it moved through space to the destination). If an obstacle is found, the agent will subgoal to move to the left or right of the obstacle (all depending on its relative location to the goal) and repeat. Once there are no obstacles between the agent and the destination, it will move towards the destination.

During the movement to a location (obstacle or destination), the path-integration system will update (and possibly replace) the representations within the configural buffer. The time during active locomotion presents ample opportunity to perform two other tasks. The first is to check to see if it has arrived at its destination. This is trivial if the destination is viewable. This checking process entails an encoding which has the added benefit of correcting any representational drift (error) introduced by the path-integration system. The more difficult case of checking for arrival at an imagined location, one without an object at it, will be discussed later.

The second task is to encode new episodic traces². This allows the agent to opportunistically attend to features in the environment that might be beneficial, in this case, other landmarks that may be predictive of foods' location. The model learns what is attended to and the features that those objects have. An *experienced* model will have learned that landmarks that are large, and more distant are ideal for establishing and returning to unique positions in space. Cressant, Muller and Poucet (1997, 1999) have reported supportive findings in the place-cell literature. Rats typically navigate using distal landmarks. However, if the rats are raised in an environment where proximal landmarks are most informative, they will learn to prefer the proximal ones. This establishment of episodic encodings of landmarks with respect to the goal is critical and will be discussed further later.

² The model implements episodic encoding with symbolic chunks, but the precise mechanism is not relevant for this discussion

Retrieval and Subsequent Return

If a visual search of the environment turns up no objects that are food, the model will attempt to retrieve an episodic memory of a successful foraging experience. Having retrieved an episodic trace (which it had laid down previously during locomotive free time), it can utilize the referenced configural information to determine what landmarks it was attending to at the time. With the identity-tags of the configural chunks, it can perform another visual search of the environment to find those same landmarks. The current model assumes an episodic encoding of at least two configural chunks, this guarantees the establishment of a unique position³. We believe that this too is a learned behavior.

With access to the configural representations of the salient landmarks from the past and present, the model is able to compute a rough imaginary destination. This can be done with simple, approximate trigonometry, and is presumed to be done by the configural system itself. Using this imagined destination, the model resumes searching for obstacles and moving.

Once again, during movement time, the model checks to see if it has arrived at the destination. Path-integration operates as before, updating the imagined location. However, the arrival check will require the recomputation of the imagined destination in order to correct for path-integration drift. Since the agent will be closer than before, the approximations will be increasingly accurate. Additionally, the opportunistic episodic encoding will likely utilize the same landmarks given the previous success. This will have the effect of boosting their activation, making them even more likely to be used at a later date in the same environment.

Behavioral Patterns

There are three key features to this model that we expect would exist in all navigational models. It is these features that allow the architecture to perform best and to exhibit the behaviors that one would expect. The first of these is the opportunistic episodic attending. In navigation the majority of the time is spent actually engaging in travel. During this cognitive slack time, it makes ecological sense that the agent would attend to the environment to benefit future performance.

The episodic traces will serve to strengthen biases within the agent's behavior. If a subsequent return to a location was successful using a set of landmarks, the properties of those landmarks will be given greater weight in the future. Therefore, as the agent develops more experience navigating in various environments, it will develop a biasing towards attending to the most informative landmarks. This learned preference operates not only at the feature level, but also at the individual

³ Assuming the agent has a bias towards remaining oriented vertically away from gravity

landmark level. When presented with a similar environment, the agent will likely reattend to the same informative landmarks.

Finally, the model treats landmarks in two different, task related ways. The first class of landmarks are those that are immediately relevant to the goal: destinations and obstacles. These configural representations are in the configural buffer only during the initial attending and checking stages. Since these objects are typically directly ahead, they can usually be tracked visually (using looming heuristics), removing the constant need for the configural system. This then releases resources that can be used for the second class of landmarks: indirect but predictive landmarks. These are the landmarks that are attended to opportunistically during locomotion and are therefore in the buffer for longer periods of time, subsequently increasing their representation in episodic memory.

Spatial View Cell Phenomenon

Having laid down the foundation of the configural system and a model that utilizes it, we can now turn to the neurological phenomenon in question. The following sections will review the spatial-view and place-cell findings. The impact of allocentric and egocentric representations will be highlighted, along with links to the implemented architecture.

Some of the most productive neurological explorations of spatial ability have centered around the so called "place-cells" in the rodent hippocampus (O'Keefe & Dostrovsky, 1971). These cells show selective firing based on where the rat is within its environment regardless of orientation. This stable firing pattern was taken as evidence for the existence of allocentric (referenced to an external object) cognitive maps in the rodent (O'Keefe & Nadel, 1978). However, single-cell recording studies of primates (eg. Ono & Nishijo, 1999) revealed that the spatially sensitive cells found in the hippocampus exhibit somewhat different properties.

Spatially Selective Firing

The spatially selective cells in the primate hippocampus are what researchers have dubbed "spatial-view" cells. Unlike place-cells that fire maximally based on the location of the rat in space, spatial-view cells respond to the spatial location of the local view. In other words, they respond when the primate is looking at a particular region in space (typically occupied by an object that attracts its attention). Figure 2 shows the firing rate of one spatial-view cell plotted against the location of the monkey's center of gaze on one wall, across two different locations in space. Even if the objects along that wall are obscured, an identical pattern of activation will be seen so long as the monkey is looking at the same location in space (indicating that it is not object-dependent but spatial).

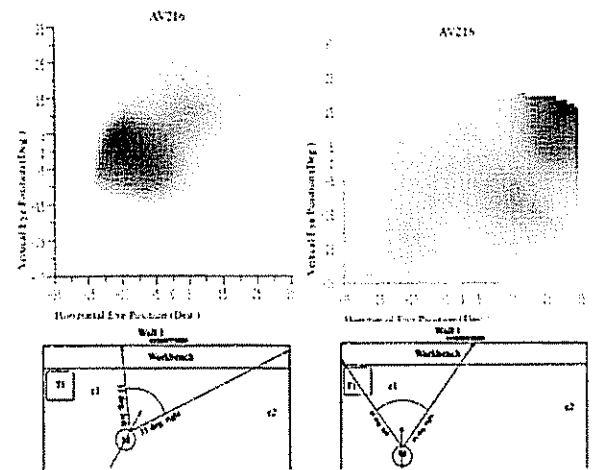


Figure 2. Spatial-view cell activation plotted against center of gaze across two different locations. Activation is maximal for the region at the center of the wall (lower portion) (adapted from Georges-Francois, et al, 1999)

The allocentric argument states that if the representations were egocentric, then there would be differing activation based on orientation (O'Keefe & Nadel, 1978). However, this definition of egocentric representation ignores the role of path-integration in spatial processing. With path-integration operating on the external representation and not the self, an external object representation would be spatially stable. Therefore, because we are unable to discern whether path-integration operates on the representation of the self or the external world, we cannot use spatial sensitivity solely to infer the representation.

On this issue, the only differences between an ego- and an allocentric representation are in the amount of processing that is performed by the path-integration system. In an allocentric scheme, the path-integration system only updates the representation of the self's location. In an egocentric representational system, path-integration updates the external representations. As mentioned previously, this seems computationally expensive, unless the representations that are to be updated are constrained in number by the architecture (i.e. buffer capacity).

Dependence

Spatial-view and place-cells respond to a spatial location even after the object at that location has been obscured (e.g. Rolls, 1999). This would suggest that the firing is not related to the object, but instead the location of it. However, an interesting phenomenon occurs in the place-cell when multiple landmarks are used. The firing fields of place-cells show a distinct pattern of being anchored to landmarks. Take an environment with two landmarks and train the rat in it. When the rat is returned to the environment after moving the landmarks

relative to each other, it is clear that the place-cells are responding to individual landmarks (see Figure 3) (e.g. Gothards, Skaggs, Morre, & McNaughton, 1996)




Figure 3. Idealized firing-field map of a place-cell's firing with respect to two translated landmarks. Note: firing fields for the cell are anchored to the landmarks

This finding presented some difficulty for the allocentric representation hypothesis. In an allocentric system, the firing fields of the place-cells should be internally stable. The modified environment should be viewed as completely new and recruit a new pattern of place-cells. Gothard, Skaggs, Morre & McNaughton (1996) hypothesized that the rat was actually maintaining and integrating multiple allocentric maps, each one focused on an individual (or cluster) of distal landmarks. Computationally, this requires more processing for storage and retrieval as well as a mechanism to merge multiple maps. However, an egocentric representation would be expected to produce such anchoring effects, as well without the complexity of representational merging.

Representational Stability

One of most fascinating aspects of place-cells that has yet to be fully tested in primate spatial-view cells is the stability of the cells across time. When returned to a familiar environment, the same set of cells will often exhibit the same firing pattern. O'Keefe & Nadel (1978) viewed this as evidence that the rodent hippocampus stored and then retrieved "cognitive maps" of the environment; in essence, a neurological topographical map of the space. Computationally, this would require the attending to, storage and retrieval of maps of arbitrary regions of space. Additionally, there would need to be a system to discern the familiarity of regions such that appropriate cognitive maps could be retrieved. These are not trivial computational concerns.

If one assumes that the stability of representations is actually just an epiphenomenon dependent upon past experience and the current task, the story changes dramatically. If attending to a certain set of landmarks has previously lead to success behaviorally, it would come as little surprise that when presented with a similar situation and a similar environment, that the attending pattern would be repeated. This hypothesis predicts that there would be tremendous variability in the stability of spatial representations as tasks change. This is in fact what is seen (e.g. Fenton & Muller, 1998) This hy-

pothesis is indifferent to the actual representation, eliminating the need for a system to store and maintain a persistent and stable cognitive map. This opens the door for a computationally simpler egocentric representational scheme.

Path-Integration

The rodent research has shown that the role of idiothetic cues (purely egocentric cues such as vestibular and proprioceptive) are also critical (e.g. Save, Nerad, Poucet, 2000). These cues allow the spatially selective cells to maintain sensitivity in situations where vision (or audition) is not sufficient to establish location. The theory is that this path-integration system uses self-motion to update its own location within the cognitive map (O'Keefe & Nadel, 1978). However, as we have already argued, path-integration may update currently active egocentric representations, yielding the same results.

Neuromimetic Phenomenology

With the neurological phenomenon, architecture, model, and behavioral consequences outlined, we can now return to the original question of how an egocentric system could produce the neurological results typically associated with allocentric representations. To do this we employ the ACT-R equivalent of single-cell recording: "single-chunk" recording. This is nothing more than recording the activation of a given declarative memory element, or in this case, the change in activation from base-level as the model progresses. For the purposes of this model, we chose to record the chunks that make up the identity-tag of the configural representations, which merely serve to link the multiple egocentric representations to a common token⁴. One can view these identity tags as specifying a unique location of such an angle and distance from landmark 1 and such an angle and distance from landmark 2 (and so on for each attended landmark).

Every time a representation is contained in the configural buffer, activation is spread to the chunks' slot values, in particular, the identity-tags. This produces an activation spike that is registered by the recorder. Since the *experienced* agent (one that has navigated across many spaces and has learned what makes good landmarks) will be biased towards attending to maximally predictive landmarks, a stable attending pattern will be seen. This attentional stability can account for the stability in the spatial-view and place-cells. Additionally, as more episodic traces are laid down, the agent has a greater probability of being able to retrieve informative landmarks when the destination is not visible (such as when a curtain is obscuring the objects that the monkey

⁴ The analogous location would be the hippocampus. The interpretation of that, however, depends on one's theoretical commitment to the hippocampus as a spatial-processing system, or an episodic encoding system

sees in spatial-view cell experiments). These then allow it to imagine the location of an object with a known identity-tag. This imagined configurational representation, when placed into the configurational buffer, spreads activation once again, reactivating previous encodings of this same (egocentric) location.

If the change in activation is then plotted with respect to the location of visual attention, a pattern such as that in the spatial-view cell literature can be seen.

Summary

In our attempt to bring ACT-R (and cognitive modeling in general) into the third-dimension, we pulled findings from numerous lines of research. Our framework posits that the mind does not contain exo- or allocentric maps. This places important constraints on spatial problem solving. Given that our assumption appears to fly in the face of a large body of research on place-cells and spatial-view cells, we have devoted our initial validation work to modeling these results. Our review of the literature leads us to believe that the allocentric assumption is not only computationally too formidable, but also not clearly supported in the data. The problem was how to formulate a tractable egocentric solution. We believe that our extension presents a solid solution to the issue of spatial-view cells. The true test will of course be its generalizability, without major modification, to place-cells and other spatial phenomenon (neurological, psychology, and behavioral).

Acknowledgements

Work on this manuscript was supported by Grant N00014-01-1-0321 from the U.S. Office of Naval Research to the second author.

References

- Anderson, J. R., & Lebiere, C. (1998) *Atomic components of thought*. Mahwah, NJ: Erlbaum.
- Byrne, M. D., & Anderson, J. R. (2001). Serial modules in parallel: The psychological refractory period and perfect time-sharing. *Psychological review*, *108*, 847-869.
- Cressant, A., Muller, R., & Poucet, B. (1997) Failure of centrally place objects to control the firing fields of hippocampal place cells. *Journal of Neuroscience*, *17*, 2531-2542.
- Cressant, A., Muller, R., & Poucet, B. (1999) Further study of the control of place cell firing by intra-apparatus objects. *Hippocampus*, *9*, 423-431.
- Eichenbaum, H. (1996). Is the rodent hippocampus just for "place"? *Current Opinions in Neurobiology*, *6*, 187-195.
- Fenton, A. A., & Muller, R. U. (1996). How two cues conjointly control hippocampal place cell firing fields. *Society Neuroscience Abstracts*, *22*, 910.
- Fenton, A. A., & Muller, R. U. (1998) Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Neurobiology*, *95*, 3182-3187.
- Gallistel, C.R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- Georges-Francois, P., Rolls, E.T., Robertson, R.G. (1999). Spatial view cells in the primate hippocampus: allocentric view not head direction or eye-position or place. *Cerebral Cortex*, *9*, 197-212.
- Gothard, K.M., Skaggs, W.E., Morre, K.M., & McNaughton, B.L. (1996). Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation tasks. *Journal Neuroscience*, *16*, 823-835.
- Kamil, A. C., & Cheng, K. (2001). Way-finding and landmarks: The multiple-bearing hypothesis. *Journal of experimental biology*, *204*, 103-113.
- McNaughton, B.L., Knierim, J.J., Wilson, M.A. (1994). Vector encoding and the vestibular foundations of spatial cognition: neurophysiological and computational mechanisms. In M. Gazzangia (Ed.), *The cognitive neurosciences*. Boston, MA: MIT.
- Mittelstaedt, M.L., & Mittelstaedt, H. (1980). Homing by path integration in a mammal. *Naturwissenschaften*, *67*, 566-567.
- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, *34*, 171-175.
- O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon.
- Ono, T., & Nishijo, H. (1999). Active septal information processing in the septo-hippocampal system. *Hippocampus*, *9*, 458-466.
- Rolls, E. (1999). Spatial view cells and the representation of place in the primate hippocampus. *Hippocampus*, *9*, 467-480.
- Save, E., Nerad, L., & Poucet, B. (2000) Contribution of multiple sensory information to place field stability in hippocampal place cells. *Hippocampus*, *10*, 64-76.
- Tolman, E.C., (1948) Cognitive maps in rats and men. *Psychological Review*, *55*, 189-208.
- Wang, R., & Spelke, E. (2000). Updating egocentric representations in human navigation. *Cognition*, *77*, 215-250.