Proceedings Paper:

Prescott, T.J. (1994) Spatial-learning and representation in animats. In: Cliff, D., Husbands, P., Meyer, J.A. and Wilson, S.W., (eds.) From animals to animats 3. Third International Conference on Simulation of Adpative Behavior, 8-12 August 1994, Brighton, UK. Complex Adaptive Systems . MIT Press , pp. 164-173. ISBN 9780262531221

In From Animals to Animats 3: The Proceedings of the Third International Conference on Simulation of Adaptive Behaviour, Brighton UK, August 1994, pp. 164-173.

Spatial Learning and Representation in Animats

Tony J. Prescott

Artificial Intelligence Vision Research Unit,
Sheffield University,
Sheffield S10 2TN,
United Kingdom.

Telephone +44 742 826547
Fax +44 742 766515

Email: T.Prescott@aivru.sheffield.ac.uk

Abstract

Animat AI has generally emphasised learning of a *dispositional*, or task-specific nature over that of a *representational* or task-independent kind. However, many animals are capable of both forms of learning, and, in particular, exploit representational learning to construct spatial knowledge that allows efficient and flexible navigation behaviour. The focus on building versatile *mobile* robots may therefore force the development of representational learning systems in animat AI. This paper considers the navigation problem and argues against the view that qualitative spatial representations, encoding principally topological relations, may necessarily be simpler to construct, store, or use than more quantitative models. It further argues against constructing a unified or global representations of space suggesting instead that knowledge should be distributed between multiple, partial, local models encoding complimentary constraints which can be combined at run-time to address a specific navigation task.

1. Learning in natural systems

Research in psychology suggests that underlying a large number of observable phenomena of learning and memory, there are two broad clusters of learning processes[1].

First, there are the *dispositional* learning processes involved in habit formation, the acquisition of motor skills, and certain forms of classical and instrumental conditioning. These processes involve incremental adaptation and do not seem to need attention or awareness. Learning is generally task-specific in that it is driven by a significant outcome in the form of a positively or negatively reinforcing event. Further, it does not seem to require or involve the acquisition of knowledge about the causal processes underlying the task that is solved.

Second, there are the *representational*[2] learning processes involved in acquiring knowledge about the relationships between events (stimuli or responses). For instance, that one event follows another (causal knowledge), or is close to another (spatial knowledge). These forms of learning appear to be have more of an all-or-none character, and may require attentional resources. They are also not directly involved in generating behaviour, and need not be acquired with respect to a specific task or desired outcome. The knowledge acquired can support both further learning or decision-making through inference.

Lesion studies with animals and patterns of learning impairment in human amnesiacs indicate that in mammals this second style of learning relies on specific medial-temporal structures in the brain, in particular, the hippocampus. In contrast the simpler associative forms of learning underlying habit and skill acquisition are not affected by damage to this brain region, but appear instead to be supported by neural systems that evolved much earlier. This view is supported by observations that all vertebrates

[1]For reviews of this extensive literature see [15, 43, 46].

[2]The terms *dispositional* and *representational* have been suggested by Thomas [47] and Morris [29] to refer to these two clusters of learning/memory processes.

and most invertebrates show dispositional learning abilities, whereas representational learning styles have evolved primarily in higher vertebrates coinciding with increased brain-size.

2. The Animat approach

The shared interest in adaptive systems, between psychologists and ethologists, on the one hand, and Artificial Intelligence researchers and roboticists on the other, has recently seen the development of a new inter-disciplinary research field. The common aim of 'Animat' (simulated animal) (Wilson [51]) research is to understand how autonomous agents—animals, simulated animals, robots, or simulated robots—can survive and adapt in their environments, and be successful in fulfilling needs and achieving goals.

Some important themes in much of this work (see, for instance [28]) are as follows: control in the agent is not centralised but is distributed between multiple task-oriented modules; there is minimal reliance on internal world models and on reasoning or planning processes; instead there is an emphasis on the role of the agent's interaction with its environment in driving the selection and performance of appropriate, generally reflexive, behaviours; perception is targeted at acquiring task-relevant information rather than delivering a general description of the current state of the perceived world .

The animat approach is thus in good accord with dispositional learning approaches (such as reinforcement learning) to the adaptation of behavioural competences. In view of the aim of building complete intelligent systems in an incremental, and bottom-up fashion this is wholly consistent with the earlier observation that learning in simpler animals is principally of a dispositional nature. However, the development of this research paradigm is already beginning to see the need for some representational learning. One reason for this is the emphasis on mobile robotics as the domain of choice for investigating animat AI.

3. Navigation as a forcing domain

The fundamental skill required by a mobile agent is the ability to move around in the immediate environment quickly and safely, this will be referred to here as *local navigation* competence. Research in animat AI has had considerable success in using pre-wired reactive competences to implement local navigation skills (e.g. [3, 10, 45]). The robustness, fluency, and responsiveness of these systems have played a significant role in promoting the animat methodology as a means for constructing effective, autonomous robots. The possibility of acquiring adaptive local navigation competences through reinforcement learning has also been investigated and has been advanced as an appropriate mechanism for learning or fine-tuning such skills [31, 40].

However, a second highly valuable form of navigation expertise is the ability to find and follow paths to desired goals outside the current visual scene. This skill will be referred to here as *way-finding*. The literature on animal spatial learning differentiates the way-finding skills of most invertebrates and lower vertebrates, from those of higher vertebrates (birds and mammals). In particular, it suggests that way-finding in most invertebrates is performed primarily by using path integration mechanisms and compass senses and secondarily by orienting to specific remembered stimulus patterns [49, 50][4-6]. This suggests that these animals do not construct models of the spatial layout of their environment and that consequently, their way-finding behaviour is relatively inflexible and restricted to homing or retracing familiar routes[3]. In contrast, higher vertebrates appear to construct and use representations of the spatial relations between locations in their environments (see, for example, [13, 32, 33, 35]). They are then able to use these models to select and follow paths to desired goals. This form of spatial learning is often regarded as *the* classic example of a representational learning process (e.g. [43]).

This evidence has clear implications for research in animat AI. First, it suggests that systems employing minimal representation and reactive competences could support way-finding behaviour similar to that of invertebrates[4]. Second, however, the acquisition of more flexible way-finding skills would appear to require representational learning abilities—this raises the interesting issue of how control and learning architectures in animat AI should be developed to meet this need.

4. How should space be represented?

In keeping with the animat approach it would seem reasonable to require the *on-line* acquisition of appropriate

[3]Gould [14] has proposed a contrary view that insects do construct models of spatial layout however, the balance of evidence (cited above) appears to be against this position.

[4]In particular it should be possible to exploit the good odometry information available to mobile robots.

spatial knowledge, and the use of representations that are simple to construct and use, cheap to store, and support a 'graceful degradation' of performance when confronted with unreliable sensory data. What forms of representational learning might satisfy these criteria?

Some recent research on robot way-finding has sought to address this challenge by advocating a substantial change in the character of the systems under investigation. Specifically, the emphasis of 'classical' AI methods on detailed path-planning using metric models of the environment (e.g. [8, 11, 16, 24, 41, 48]) has been rejected by some researchers in favour of the use of more 'qualitative' methods and models (e.g. [10, 17-19, 21-23, 25, 26, 30]). In these systems metric modelling and wayfinding is often regarded as supplementary to a core level of navigation skill based, primarily, on representations of topological spatial relations. This approach has, as part of its motivation, the perceived inadequacies of classical systems which are regarded as over-reliant on accurate sensing and detailed world models. It is suggested that such systems are both too 'brittle' in the face of degraded or missing sensory information, and too costly in terms of the computational and memory resources they require.

A second motivation for investigating topological spatial models is research on human way-finding. Much of this literature follows a theory originating with Piaget [36] that human spatial knowledge has a hierarchical structure and is acquired through a stage-like process. Specifically, Piaget, and later Siegel and White [44], have argued that a fundamental stage in the acquisition of spatial knowledge is the construction of qualitative models of the environment from more elementary sensorimotor associations. This representation is then gradually supplemented by distance and direction information to form a more detailed quantitative map. An important element of this theory is the view that a primarily topological representation can support robust way-finding behaviour in everyday environments. Computational models inspired by the human way-finding literature have been described by Leiser [22] and by Kuipers [18-21]. The latter in particular has developed a number of robot simulations of considerable sophistication and detail based on the hypothesis of a hierarchical representation of spatial knowledge. The following extract serves to illustrate this theoretical position, which has been influential in other recent work on robot way-finding (e.g. [23]):

> "There is a natural four-level semantic hierarchy of descriptions of large-scale space that supports robust map-learning and navigation:
>
> 1. *Sensorimotor:* The traveller's input-output relations with the environment.
>
> 2. *Procedural:* Learned and stored procedures defined in terms of sensori-motor primitives for accomplishing particular instances of place-finding and route-following tasks.
>
> 3. *Topological:* A description of the environment in terms of fixed entities, such as places, paths, landmarks, and regions, linked by topological relations such as connectivity, containment and order.
>
> 4. *Metric:* A description of the environment in terms of fixed entities [...] linked by metric relations such as relative distance, relative angle and absolute angle and distance with respect to a frame of reference.
>
> In general, although not without exception, assimilation of the cognitive map proceeds from the lowest level of the spatial semantic hierarchy to the highest, as resources permit. The lower levels of the cognitive map can be created accurately without depending greatly on computational resources or observational accuracy. A complete and accurate lower level map improves the interpretation of observations and the creation of higher levels of the map." ([21], p. 26)

In many respects this view is highly acceptable, for instance, the proposal that spatial knowledge is organised in distinct components encoding separate forms of constraint is a welcome contrast to the traditional approach of unitary global models. However, there are several implications of this view that are open to question. First it is important to ask in what sense the organisation of spatial knowledge should be viewed as hierarchical rather than heterarchical; second, to what degree global models, such as those described for the topological and metric levels, are required (as opposed to multiple overlapping local models); and finally, whether the emphasis on geometric content as the main distinguishing factor between models is correct.

Some of these issues can be highlighted by contrasting the literature on human way-finding with much of the research from the wider field of animal navigation.

In particular, this latter evidence suggests a discontinuity between procedural knowledge and the use of map-like metric spatial representations [13, 32, 33].

For instance, in contrast to an incremental hierarchy, O'Keefe [32, 33] has argued that there are two fundamentally independent navigation systems used by mammals including man. The first of these, which he calls the *taxon* system is supported by route-like chains of $stimulus \times action \rightarrow stimulus$ associations. Each element in such a chain is an association that involves approaching or avoiding a specific cue, or performing a body-centred action (generally a rotation) in response to a cue. Taxon strategies therefore have a similar nature to the procedural knowledge in the second level of Kuipers hierarchy. O'Keefe's second system, called the *locale* system, is, however, a 'true' mapping system in that it constructs a representational model describing the $stimulus \rightarrow stimulus$ metric spatial relations between locations in the environment. Evidence for the existence of this system and its independence from taxon strategies consists of both observational and laboratory studies of animal behaviour, and neurophysiological studies suggesting that different brain structures underlie the two systems.

Although the highest level of Kuipers' hierarchy can be identified with O'Keefe's locale system the former suggests a continuity—with assimilation of information onto 'weaker' representations to generate the metric model, whereas the latter stresses the discontinuity and apparent autonomy of the two alternative mechanisms. A further difference is that O'Keefe's theory bypasses the level of the topological map, if such a map exists it is as an abstraction from the full metric representation—this contrasts with Kuipers view in which the topological model is simpler and comes first.

Gallistel [13], who provides an extensive review of research on animal navigation, also concludes that animals make considerable use of metric data for navigation. Like O'Keefe he also proposes a modular and autonomous mapping system that stores a *metric* representation of spatial layout[5].

5. Topological and/or Metric Modelling?

The nature of the spatial relations encoded by a world model determines the type of navigation behaviour that can be supported. Procedural (or route) knowledge can only support movement along segments of known paths. Knowledge of the topological layout of the environment gains the navigator the ability to identify suitable sub-goals, and generate and follow novel routes between locations. However, because this knowledge is limited to knowing the connectivity of places navigation is constrained to using known path segments between adjacent sub-goals. A navigator with a topological map who enters unfamiliar territory can explore new paths and construct new map knowledge but cannot engage in goal-directed movement to target sites. The ability to determine short-cut or straight-line routes across un-explored terrain requires knowledge of higher-order spatial relations. Where such behaviours are observed in animals this is usually taken as strong evidence for the use of a metric model. That such skills would be very useful to an animal or robot is undeniable giving a strong incentive for constructing and using knowledge of this type.

Given the value of metric knowledge is there a justification for constructing, as the first or only form of spatial representation, models encoding weaker geometric constraints? One possible argument for such a view is the idea that a topological model could be constructed without the need to detect higher-order relations. Mathematically topological geometry is simpler and more basic than metric geometry—it requires fewer axioms. However, this mathematical simplicity perhaps belies the real difficulties of constructing topological knowledge in the absence of metric knowledge. I have argued elsewhere [39] that such a model is in general realisable only if the agent has sensory abilities that can be relied on to give accurate identification and re-identification of most locations (henceforth *place identification*), and that in practice this may require vision skills capable of object recognition or, at least, of very

[5] O'Keefe and Gallistel agree on the existence of a separate metric mapping system but largely disagree on the relative importance of dead reckoning and environmental fixes in constructing the map. This debate will be considered further below.

robust visual pattern matching[6]. This fits uneasily with the bottom-up bias of Animat AI, and, indeed, with the current perceptual abilities of animat-style robots. Many systems currently use local sonar patterns to characterise different places. However, this sort of local geometry information is not likely to be sufficiently distinctive to allow the disambiguation of similar places. For this reason sonar patterns are often supplemented by odometry information in order to make the place identification task feasible (e.g. [25]). This need to exploit metric knowledge (albeit of a rough and ready kind) demonstrates the difficulty of topological mapping with non-visual sensory data[7].

The use of some approximate metric knowledge in model-building introduces the possibility that such information might be computed and exploited in constructing representations, but might not actually be explicitly recorded or used for way-finding. Given the advantages that metric knowledge of any sort can endow the main justification for this proposal must be that the cost or complexity of building (or using) such representation would outweigh its usefulness. This supposition, which may stem from the perceived weaknesses of some existing metric-based systems, may, however, be premature. A distinctive characteristic of much research in quantitative model-building has been an emphasis on combining all available information into a unique *global model*. One aim of this paper is to suggest that it is perhaps this characteristic, rather than the use of quantitative information, that has contributed to the inflexibility and over-sensitivity to measurement error observed in some existing systems.

6. Constructing metric models

The task of building a representation of an environment that encodes distance and direction requires that places are located with respect to common coordinate frames. The coordinate frames most directly available are *egocentric*, that is, they are defined by the navigator's instantaneous position and orientation in space. However, to construct a useful model, observations from different view-points must be integrated into representations with respect to environment-centred or *allocentric* frames. That is egocentric relations must be transformed to give allocentric spatial relations.

The arguments that metric models are more complex or expensive to construct and use than more qualitative ones, generally concern the difficulties of obtaining accurate distance information (or failing that, dealing with noisy information), and the resource demands of the need for continuous transformations between egocentric and allocentric frames. However, there are representations and mechanisms that may overcome some of these objections. Specifically, distributed representations have been proposed [23, 38] in which the environment is represented using multiple models based on coordinate frames defined by distinctive local landmarks.

Two landmarks are sufficient to define a two-dimensional coordinate frame (three for 3D), however, coordinate transformations based on such frames require non-linear computations (trigonometric functions and square-roots) and further require that an arbitrary ordering of the reference points is remembered in order to specify a unique coordinate frame. However, as Zipser [52] has pointed out, if a further landmark is used to define each local frame then all the transformation calculations become linear. In [38, 39] I have described a simulation based on this proposal in which the positions of salient goals are stored with respect to the two-dimensional co-ordinate frames[8] defined by groups of three local landmarks. Multiple local frames are represented in a connectionist

[6]Some recent work on metric-free topological map-building for robot navigation has recognised this need for powerful visual processing e.g. [17] .

[7]The problem of constructing a topological map are eased considerably by introducing additional constraints to the map-building process. One possible constraint is to limit the behavioural repertoire of the robot. One way this might be achieved is to force the robot to maintain a travel-path that follows object boundaries [25, 30] . This constraint of 'wall-following' limits the connectivity of the resulting topological graph and reduces the number of true choice points (i.e. vertices in the graph of degree>2). This eases the tasks of segmenting the environment (into the regions that form graph vertices) and of place identification, and also avoids the need to represent places that lack distinctive local features. This approach has lead to some successes in building way-finding systems for indoor autonomous robots, however, it also has an obvious cost—open spaces will be poorly represented in the map and movement will be more rigidly limited to a small set of paths.

[8]These coordinates are, strictly speaking *affine* rather than metric, however, sssuming that the agent detects metric egocentric spatial relation according to some calibrated Euclidean measure, metric relations—direction and distance—will be recoverable from a stored affine model.

network in which the task of determining direction or route information to a desired goal is performed by a parallel, 'spreading activation' search. The computations required to construct these representations from (noisy) egocentric metric data require only linear mathematics (indeed a simple 'perceptron'-like learning rule will suffice) and have memory requirements roughly proportional to the number of goal-sites and landmarks stored. While following a planned route the system can also exploit run-time error-correction by incorporating egocentric fixes on sighted landmarks. This makes the route following system highly robust to noise in the representation or perceptual system.

In contrast to approaches in which the goal is to construct a permanent 'map' of environmental layout (in which position errors are minimised or explicitly modelled) this approach builds no long-term static representation of global spatial relations. Instead the store of knowledge concerning a specific goal or landmark is distributed across a number of local models in the network allowing the constraint information to be combined at run-time for any given task. Methods for combining different constraints using Kalman filtering techniques are currently being investigated.

7. A 'multiple schemata' view

This last section is an attempt to set out a perspective on the acquisition of representations of space. In contrast to Kuipers' hierarchical approach in which global topological and metric models are constructed, it proposes a heterarchy of local models in which the geometric distinction is only one among a number of characteristics identifying complimentary representational forms. Following Michael Arbib [1, 2], I call this a 'multiple schemata' view.

Arbib has proposed the use of the term *schemata* to describe active representational systems or "perceptual structures and programs for distributed motor control" ([1] p. 47). In the context of constructing and using models of space he suggests—

> "The representation of [...] space in the brain is not one absolute space, but rather a patchwork of approximate spaces (partial representations) that link sensation to action. I mean '*partial*' and '*approximate*' in two senses: a representation will be *partial* because it represents only one small sample of space (...), and it will be *approximate* in that it will be based on an incomplete sample of sensory data and may be of limited accuracy and reliability. I will suggest that our behaviour is mediated by the interaction of these partial representations: both through their integration to map ever larger portions of the territory relevant to behaviour, and through their mutual calibration to yield a shared representation more reliable than that obtained by any one alone." ([2] p. 380)

In the specific context of cognitive maps, he also suggests that:

> "There is no reason, in general, to expect the system to use Euclidean space or Cartesian coordinates for such a map. Rather, the system needs a whole array of local representations that are easily interfaced and moved between." ([1] p. 47).

The view advocated here is, I hope, in close accord with these ideas.

Spatial information can be picked-up through multiple sensory modalities in a number of different guises and forms. This information may describe spatial relations upto any level of geometric richness (topological—metric) it may also be anywhere on a scale from precise to vague. Each piece of information can be viewed as supplying a potential constraint that can assist navigation.

I propose that the critical distinction with regard to different forms of constraint information has less to do with the geometric content of the knowledge and more with *the process by which that knowledge is derived*. For instance, metric information derived from odometry is (to a large extent) independent from metric information determined by perceived distance and direction to identifiable salient landmarks. These two forms of quantitative knowledge thus provide constraints that are complimentary because they derive from *different sensory modalities*. Multiple constraints can also be obtained from within a single sensory modality by observing *different environmental cues*. For instance, the observed position of a single distant landmark (such as the sun) gives a direction constraint that is essentially independent from spatial localisation with respect to local landmarks. Indeed, different individual landmarks or landmark groups can supply separate constraints as has been demonstrated in [23, 38]. Finally, relatively independent constraints can arise within a modality by reference to the same external cues but by employing *different computational mechanisms*. It is in this sense, perhaps, that the distinction between different geometries may be most relevant. For instance, the visual characteristics of landmarks might be used to construct knowledge of topological relations that is largely independent of the mechanisms that extract distance or direction from the visual scene.

To the extent that different constraints are independent two constraints will clearly be much more powerful than one, three more than two, etc. It therefore seems reasonable to suggest that an agent should seek to detect and represent a number of independent or near independent constraints that describe the spatial relations between important places.

The emphasis of a multiple schemata approach is not on constructing unified representations such as topological or metric maps but rather on establishing multiple complimentary spatial descriptions. Each schema should exploit a different combination of *cues*, *channels*, and *mechanisms* to instantiate a set of environmental spatial relations. Thus, there will overall, be a number of relatively distinct path-ways through which knowledge is acquired. This suggests a heterarchy of models (as opposed to a hierarchy), with some, but not all, schemata sharing common sources and resources. At any time an agent should exploit the knowledge in several schemata to support its current navigation task. Although some tasks may require the temporary creation of a unified model (drawing a graphical map of the environment might constitute such a task) in general the underlying representations can remain distinct allowing the reliability of each source of information to be assessed at run-time.

Way-finding should exploit acquired schemata via *arbitration* procedures which decide on the basis of the content and accuracy of each active model the extent to which it should contribute to the decision process. This arbitration could be carried out through some fixed subsumption mechanism whereby, for instance, knowledge determined from large-scale metric relations could override taxon (route-following) strategies. Alternatively a more sophisticated system would seek to combine the constraints afforded by multiple schemata by weighting them according to their perceived accuracy or reliability. In this way, reliable identification of a highly distinctive landmark might override estimates of spatial position or orientation determined by some metric reckoning process.

These ideas are currently being investigated with respect to the distributed coding system described in [38, 39] (and, briefly, above). If a specific location is encoded by two separate schemas based on non-overlapping landmark triples then these would constitute relatively independent constraints. To the extent that landmark sets do overlap they will obviously be less independent, but will nevertheless encode partially distinct constraint information.

However, the idea of multiple schemata also generalises to encompass different coding systems. For instance, representations based principally on direction sense and odometry could be constructed which would provide a modality-independent source of information from the landmark-based coding.

An obvious argument against a multiple schemata view is that acquiring and storing spatial knowledge is not without cost. It makes demands on attention, processing and memory (there are really separate costs associated with detecting constraints, storing them, retrieving them, and combining them!). One defence against this argument is the relative independence between different schemata which will allow parallel processing to be exploited to a considerable extent. A second possibility, which is rarely explored in research with artificial animats, is that the amount of resources devoted to a given location (i.e. the number of constraints stored) may vary according to the subjective importance of being able to relocate that place or reach it quickly. We could expect, for instance, that an animal's home or nest (or a robots power source) would have the highest priority and that therefore 'homing' might be the most robustly supported way-finding behaviour.

A multiple schemata view can help in understanding the evolution of way-finding competence in animals, and may also provide support for the essentially pragmatic approach of Animat AI. In the case of the former, O'Keefe's [32] separate taxon and locale systems (which follows a very long line of research into response vs. place knowledge in animal navigation, see Olton [34, 35]) can be viewed as a distinction along these lines. However, there also seems to be a reasonable case for breaking up the 'locale' system into multiple schemata, for instance, models derived from odometry and direction senses [13, 27] and those derived principally from codings with respect to distinct local landmark groups [32].

In robotics this view suggests the abandonment of theoretical pre-conceptions about the priority, or lack of it, of different forms of geometric knowledge. It further implies that the 'brittleness' of classical approaches arises not so much from the emphasis on metric modelling but from the search for an accurate unified metric model.

Much existing work is compatible with this approach. In addition to Michael Arbib's work on schema theory [2] much work in psychology (e.g. [42]) and AI (e.g. [23]) shares similar objectives. Work in animal navigation that specifically fits this research theme has been performed by Poucet et al. (e.g. [7, 37]), Collett et al. [9] and Etienne et al. (e.g. [12]). To end this paper I would therefore like to draw upon a couple of examples from this work.

Etienne et al. [12] report that hamsters have effective dead reckoning skills which are sufficient to relocate their nest in darkness. However, in lighted conditions hamsters were found to orient primarily using visual information about local landmarks. In conflict situations, where a landmark (a single light spot) was rotated relative to the learned position, the hamsters homed using either the landmark information or their dead-reckoning sense. When the visual information and dead reckoning produced highly divergent paths dead reckoning was used, however, with smaller discrepancies visual information took priority over path integration. Etienne et al. also report that the dead-reckoning sense was more precise when used to return to the nest than when used to locate a secondary feeding site. This suggests that a dead reckoning way-finding schema maybe more available for homing than for general path-finding.

Experiments by Collett et al. [9] with gerbils suggests that these animals may encode goal positions (buried sunflower seeds) in terms of individual visible landmarks by using some form of direction sense. For instance, in one experiment gerbils were trained to locate a food cache at the centre of an array of two landmarks. When the distance between landmarks was doubled the gerbils searched at two sites each at the correct distance and orientation to one of the landmarks rather that at the centre of the two locations (as some theories of a landmark 'map' might predict). In a further experiment the gerbils were trained to go to a goal-site at the centre of a triangle of three landmarks. During testing the distance of one landmark to the centre was doubled, Collett et al. report that the animals spent most of their search time around the place specified by the two landmarks, ignoring the one that broke the pattern. They interpreted this result in the following way:

> "The gerbil is thus equipped with a useful procedure for deciding between discrepant solutions. When most of the landmarks agree in specifying the same goal, with just a few pointing to other sites, the chances are that the majority view is correct and that the additional possibilities result from mistakes in computation or from disturbances to the environment." ([9] p. 845).

Collett et al. are therefore suggesting that this multiple encoding of landmark-goal relations by hamsters occurs to provide the system with robustness. In other words, they advocate something like a multiple schemata system and give a clear example of the ability of such a hypothesis to generate interesting and testable predictions.

8. Conclusions

This paper has argued that the problem of representation for animat spatial learning may be best approached by discarding the goal of a complete global model of the environment in favour of the use of multiple, partial local models encoding complimentary constraints. This approach, I believe, has a resonance with the general ethos of animat research that opposes the need for representation for its own sake (which has often seemed to be the goal of classical AI) and is against a strong distinction between model and mechanism. This view constitutes a theoretical position that has as yet only been partly explored in simulation, it is thus proposed as a hypothesis which awaits evaluation through the construction of genuine way-finding robots.

Acknowledgements

The author is grateful to John Mayhew for advice and comments. This research was funded by the Science and Engineering Research Council.

References

1. Arbib, M.A. (1989). *The Metaphorical Brain*. New York: Wiley and Sons.

2. Arbib, M.A. (1990). Interaction of multiple representations of space in the brain, in *Brain and Space,* J. Paillard, Editor. Oxford University Press: Oxford. p. 380-403.

3. Brooks, R.A. (1989). "Robot beings". in *International Workshop on Intelligent Robots and Systems*. Tsukuba, Japan.

4. Cartwright, B.A. and T.S. Collett (1979). "How honeybees know their distance from a nearby visual landmark". *Journal of Experimental Biology*. **82**: p. 367-72.

5. Cartwright, B.A. and T.S. Collett (1983). "Landmark learning in bees: experiments and models". *Journal of Comparative Physiology A*. **151**(4): p. 521-544.

6. Cartwright, B.A. and T.S. Collett (1987). "Landmark maps for honeybees". *Biological Cybernetics*. **57**: p. 85-93.

7. Chapuis, N., C. Thinus-Blanc, and B. Poucet (1983). "Dissociation of mechanisms involved in dogs' oriented displacements". *Quarterly Journal of Experimental Psychology*. **35B**: p. 213-219.

8. Chatila, R. (1986). "Mobile robot navigation: space modelling and decisional processes". in *3rd International Symposium on Robotics Research*. . MIT Press.

9. Collett, T.S., B.A. Cartwright, and B.A. Smith (1986). "Landmark learning and visuo-spatial memory in gerbils". *Journal of Comparative Physiology A*. **158**: p. 835-51.

10. Connell, J.H. (1990). *Minimalist Mobile Robots*. Perspectives in Artificial Intelligence, Boston: Academic Press.

11. Crowley, J.L. (1985). "Navigation for an intelligent mobile robot". *IEEE Journal of Robotics and Automation*. **1**(1): p. 31-41.

12. Etienne, A.S. (1992). "Navigation of a small mammal by dead reckoning and local cues". *Current Directions in Psychological Science*. **1**(2): p. 48-52.

13. Gallistel, C.R. (1990). *The Organisation of Learning*. Cambridge, MA: MIT Press.

14. Gould, J.L. (1986). "The locale map of honeybees: do insects have cognitive maps?". *Science*. **232**: p. 861-63.

15. Hintzman, D.L. (1992). "Twenty-five years of learning and memory: was the cognitive revolution a mistake". in *Fourteenth International Symposium on Attention and Performance*. Michigan. Bradford Books.

16. Iyengar, S., *et al*. (1985). "Learned navigation paths for a robot in unexplored terrain". in *2nd Conference on Artificial Intelligence Applications*.

17. Kortenkamp, D., *et al*. (1991). "A scene-based multi-level representation for mobile robot spatial mapping and navigation". *IEEE transactions on robotics and automation*.

18. Kuipers, B. (1982). "The "map in the head" metaphor". *Environment and behaviour*. **14**: p. 202-220.

19. Kuipers, B. and Y.T. Byun (1991). "A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations". *Robotics and Autonomous Systems*. **8**, also in *Towards Learning Robots,* W.Van de Welde (ed.), MIT Press 1993.

20. Kuipers, B. and Y.T. Byun (1987). "A robust, qualitative method for robot exploration and map-learning". in *Proceedings of the Sixth National Conference on AI (AAAI–87)*. St. Pauls, Minneapolis. Morgan Kaufmann.

21. Kuipers, B. and T. Levitt (1988). "Navigation and mapping in large-scale space". *AI Magazine*. (Summer 1988): p. 25-43.

22. Leiser, D. and A. Zilbershatz (1989). "The traveller: a computational model of spatial network learning". *Environment and behaviour*. **21**(4): p. 435-463.

23. Levitt, T.S. and D.T. Lawton (1990). "Qualitative navigation for mobile robots". *Artificial Intelligence*. **44**: p. 305-360.

24. Lozano-Perez (1983). "Spatial-planning: a configuration space approach". *IEEE transactions on computers*. **C-32**(2): p. 108-121.

25. Mataric, M.J. (1990). "Navigating with a rat brain: A neurologically-inspired model for robot spatial representation". in *From Animals to Animats: Proceedings of the First International Conference Simulation of Adaptive Behaviour*. Paris. MIT Press.

26. Mataric, M.J. (1992). "Integration of representation into goal-driven behaviour-based robots". *IEEE Transactions on robotics and automation*. **8**(3): p. 304-312.

27. McNaughton, B.L., L.L. Chen, and E.J. Markus (1991). "Landmark learning and the sense of direction - a neurophysiological and computational hypothesis.". *Journal of Cognitive Neuroscience*. **3**(2): p. 192-202.

28. Meyer, J. and A. Guillot (1990). "Simulation of adaptive behaviour in animats: review and prospect". in *From Animals to Animats: Proceedings of the First International Conference Simulation of Adaptive Behaviour*. Paris.

29. Morris, R.G.M. (1990). "Toward a representational hypothesis of the role of hippocampal synaptic plasticity in spatial and other forms of learning". *Cold Spring Harbor Symposia on Quantitative Biology*. **55**: p. 161-173.

30. Nehmzow, U. and T. Smithers (1990). "Mapbuilding using self-organising networks in "really useful robots"". in *From Animals to Animats: Proceedings of the First International Conference Simulation of Adaptive Behaviour*. Paris.

31. Nehmzow, U., T. Smithers, and B. McGonigle (1992). "Increasing behavioural repetoire in a mobile robot". in *From Animals to Animats: Proceedings of the Second International Conference on Simulation of Adaptive Behaviour*. Honolulu, USA.

32. O'Keefe, J. (1990). The hippocampal cognitive map and navigational strategies, in *Brain and Space,* J. Paillard, Editor. Oxford University Press: Oxford.

33. O'Keefe, J.A. and L. Nadel (1978). *The Hippocampus as a Cognitive Map*. Oxford University Press.

34. Olton, D.S. (1979). "Mazes, maps, and memory". *American Psychologist*. **34**(7): p. 583-596.

35. Olton, D.S. (1982). Spatially organized behaviours of animals: behavioural and neurological studies, in *Spatial Abilities: Development and Physiological Foundations,* M. Potegal, Editor. Academic Press: New York.

36. Piaget, J., B. Inhelder, and A. Szeminska (1960). *The Child's Conception of Geometry*. New York: Basic Books.

37. Poucet, B., C. Thinus-Blanc, and N. Chapuis (1983). "Route planning in cats, in relation to the visibility of the goal". *Animal Behaviour*. **31**: p. 594-599.

38. Prescott, A.J. (1993). "Building long-range cognitive maps using local landmarks". in *From Animals to Animats: Proceedings of the Second International Conference on Simulation of Adaptive Behaviour*. Honolulu, USA.

39. Prescott, A.J. (1993). *Explorations in Reinforcement and Model-based Learning*. Phd thesis. University of Sheffield, UK.

40. Prescott, A.J. and J.E.W. Mayhew (1992). "Obstacle avoidance through reinforcement learning". in *Advances in Neural Information Processing Systems 4*. Denver. Morgan Kaufman.

41. Rao, N., N. Stoltzfus, and S. Iyengar (1988). "A retraction method for terrain model acquisition". in *IEEE International Conference on Robotics and Automation*. .

42. Scholl, M.J. (1992). "Landmarks, places, environments: multiple mind-brain systems for spatial orientation". *Geoforum*. **23**(2): p. 151-164.

43. Sherry, D.F. and D.L. Schacter (1987). "The evolution of multiple memory systems?". *Psychological Review*. **94**(4): p. 439-454.

44. Siegel, A.W. and S.H. White (1975). The development of spatial representations of large-scale environments, in *Advances in Child Development and Behaviour,* H.W. Reese, Editor. Academic Press.

45. Soldo, M. (1990). "Reactive and preplanned control in a mobile robot". in *IEEE Conference on Robotics and Automation*. Cincinnati, Ohio.

46. Squire, L.R. (1992). "Memory and the Hippocampus: A synthesis from findings with rats, monkeys and humans". *Psychological Review*. **99**(2): p. 195-231.

47. Thomas, G.J. and P.S. Spafford (1984). "Deficits for representational memory induced by septal and cortical lesions in rats". *Behavioural Neuroscience*. **98**: p. 394-404.

48. Turchan, M. and A. Wong (1985). "Low-level learning for a mobile robot: Environment model acquisition.". in *2nd Conference on Artificial Intelligence Applications*.

49. Wehner, R. (1983). Celestial and terrestrial navigation: human strategies - insect strategies, in *Neuroethology and Behavioural Physiology,* F. Huber and H. Markl, Editor. Springer-Verlag: Berlin.

50. Wehner, R. and R. Menzel (1990). "Do insects have cognitive maps?". *Annual Review of Neuroscience*. **13**: p. 403-14.

51. Wilson, S.W. (1990). "The animat path to AI". in *From Animals to Animats: Proceedings of the First International Conference on the Simulation of Adaptive Behaviour*. Paris.

52. Zipser, D. (1986). Biologically plausible models of place recognition and place location, in *Parallel Distributed Processing: Explorations in the Micro-structure of Cognition, Volume 2.,* J.L. McClelland and D.E. Rumelhart, Editor. Bradford: Cambridge, MA. p. 432-70.