THE MAP-LEARNING CRITTER

Benjamin Kuipers

Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712

TR-85-33 December 1985
AI-85-17

The Map-Learning Critter

Benjamin Kuipers
Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712.

December 6, 1985

Abstract

The Critter is an artificial creature which learns, not only the structure of its (simulated) environment, but also the interpretation of the actions and senses that give it access to that environment. The Map-Learning Critter embodies a strong *a priori* hypothesis: that its environment is, at least in part, structured as a large-scale space consisting of places connected by paths. The Critter's learning strategy begins by diagnosing its actions and senses. By performing experiments and examining the periodicity of its sense-impressions, it classifies actions as "turn-like," "travel-like," and "others." After the actions are classified, it becomes possible to aggregate sets of sense-impressions to define places; then places are linked into paths. An exploration strategy assures that the entire environment will be explored and assimilated into this model. The Map-Learning Critter has been implemented and has experienced a variety of reasonable and unreasonable environments. Some implications of the results are discussed, and directions for future research are outlined.

1 The Critter Problem

The Critter Problem was proposed by Ron Rivest in Fall of 1984. Imagine a creature, born a tabula rasa into an unknown environment. The Critter is capable of performing some set of actions $a_1 \ldots a_m$. At any given moment, it receives a sense-vector $\langle s_1 \ldots s_n \rangle$ of sense-impressions. Any component of the sense-vector has only a small finite number of possible values. In the initial formulation of the problem, the environment consists of a discrete and finite set of states, and sense-impressions change only as a consequence of actions.

The Critter's predicament is that it has no interpretation for either actions or senses, and no knowledge of the structure of its environment. The problem is to write a program for the Critter allowing it to learn about its senses, actions, and environment well enough to predict the sense-vector resulting from a given sequence of actions. The problem was posed as a contest, in which Critters were to be placed into a world unknown to their programmers, and graded according to their success at learning the environment and predicting the results of actions.

One straight-forward approach is to perform actions randomly and accumulate a table of the conditional probabilities of resulting sense-vectors. The problem with this approach is that the Critter does not "understand" the structure of its environment in any useful sense. Furthermore, the combinatorics of a complex environment are such that a complete understanding can only be obtained by factoring the sets of actions and sense-impressions into manageable dimensions.

Based on my earlier work on computational modeling of the human cognitive map [Kuipers 1977, 1978, 1979, 1983], I hypothesized that the best approach for the Critter was to emulate human learning of a spatial environment. This involves acquiring from observations gathered during travel, a model of the spatial structure of the environment expressed in several different representations of progressively broader scope. However, this previous work, the TOUR model, started with a given set of actions whose general nature is known in advance, while the Critter starts with an *uninterpreted* set of action symbols.

The map-learning Critter is *model-driven*, in the sense that it has certain preconceptions about the abstract properties of possible actions, in terms of which it diagnoses its actual given set of actions. Similarly, it assumes that there will be some useful way to describe its spatial environment in terms of places and paths, and it determines the best way to assimilate its observations to that model. The learning process of the Critter decomposes into three major components which will be discussed in detail in the following sections.

1. Diagnosis of the properties of the actions and senses.

2. Assimilation of observations acquired during travel into a framework of sensory surrounds, places, and paths.

3. An exploration strategy that selects sequences of actions to allow a complete exploration of the environment.

In Kant's sense, we might see the Critter as having a *synthetic a priori* knowledge of space: *a priori* since it is built in rather than learned, and *synthetic* since it describes the real world rather than tautologies. The built-in knowledge is also of a rather abstract and mathematical character, consisting of knowledge about periodicity, prerequisites to actions, equivalence classes, and finite ordered sets.

2 Diagnosing Actions and Senses

Starting with no knowledge of its senses or environment, the Critter's first problem is to diagnose the properties of its actions and their effects on its senses. It is infeasible to look at the effects of all sequences of actions, so we want to factor the problem by identifying useful properties of individual actions. With no alternative sources of information, their effects can only be determined by looking at the sense-vectors. With no interpretation of the sense-vectors, all we can do is match for identity, and look for repetition and periodicity.

2.1 Classifying Actions

An action is diagnosed by performing it repeatedly about 40 times, then looking at the periodicity of the resulting sequence of sense-vectors. A sequence is described as having a *period*, which is the shortest cycle which yields four or more exact repetitions at the end of the sequence, and a *header*, which is the portion (if any) before the sequence becomes cyclic.

The most important category of action consists of those which produce a completely periodic sequence of sense-impressions, and can thus be thought of as *turns*. In a finite, bounded environment, an action which produces a sequence of sensory changes followed by an indefinite period of no effect (period 1), can be thought of as a *travel* action which eventually bumps into a wall or otherwise fails to satisfy the prerequisites for further motion. Actions which cannot be classified as turns or travels are distinguished according to whether they become periodic, and hence deterministic, or should be considered random. (See Table 1.)

2.2 Identifying Inverses

Once the actions have been classified, inverses are easy to look for. Only pairs of actions of the same type need be considered as possible inverses. A pair of actions are inverses if they produce a sequence of sense-impressions of period 1. Since turns are cyclic, they always have inverses, but it

Turn-Actions	$\equiv$	period > 2, header $= 0$
Travel-Actions	$\equiv$	header > 0, period $= 1$
Other-Actions	$\equiv$	period > 0.
Random-Actions	$\equiv$	no period.

Table 1: The Classification Criteria for Action Types

is also possible to identify inverse pairs such as Turn-Right and Turn-Left. If the period of a turn is even, a half-cycle can be considered a *turn-around.* Finally, this lets us identify a complex inverse for a Travel action, where

$$travel^{-1} = \langle turnaround \rightarrow travel \rightarrow turnaround \rangle.$$

2.3 Identifying Prerequisites

The Critter has only a limited capability of diagnosing the components of its sense-vector. When a travel action reaches the "end of the road", and a prerequisite is no longer satisfied, the components of the sequence of sense-vectors can be analyzed to determine whether one or more components signal the state of the prerequisite. This signal can be of two forms:

- Open/Closed: indicating whether the next action will succeed (e.g. if faced with a door).
- Bump/No-Bump: indicating whether the previous action has succeeded (e.g. indicating collision with a wall).

The remaining elements of the sense-vector are uninterpreted, and are used only in matches for identity.

Clearly, the result of this diagnostic process is dependent on the current environmental situation of the Critter. An attempt to diagnose an action whose prerequisites are currently unsatisfied will not produce useful results. Thus, the diagnosis is repeated several times, interleaved with sequences of random actions, with the possibility of revising the classification of a particular action. This provides some capability for recovery from potentially misleading diagnoses.

Once the actions have been classified, it becomes possible to build up a map of the environment.

3 Assimilating Observations into a Cognitive Map

Knowledge framed purely in terms of actions and the resulting sense-vectors is expressed at too low a level to capture the overall structure of a complex environment. Instead of this *sensorimotor knowledge*, the Critter needs knowledge expressed in terms of fixed features of the environment, *places* and their relations with each other.

But how do we define a *place*? The basic constraints on our definition are:

- a place is where you stay when you do nothing but turn;
- except at the "end of the road", a travel action necessarily takes you from one place to another;
- two places can be indistinguishable from local sensory evidence.

We define a *surround* as the cycle of sense-vectors obtained by repeating a turn action. The first time a particular surround is encountered, it is associated with a single *place*. If an existing surround is encountered from an unexpected direction, it is tested to determine whether the current place can be distinguished from previously known places associated with that surround. For example, if a travel action (not at the "end of the road") leaves the Critter within the same surround, the previous and current places are different. In such a case, a new place description is created and associated with the same surround.

A *path* is the sequence of places encountered while doing a travel action. A path has the property that the same sequence should be encountered in reverse order when it is traveled in the reverse direction. The two directions of travel on a path are therefore distinguishable by the order in which places are encountered. In summary:

- A *surround* is the cycle of sense-vectors obtained by repeating a turn action.
- A *place* describes a state of being within a given surround distinguishable from other places in the same surround.

- A *path* is a sequence of places reached by successive executions of a particular travel action.

The length of the cycle of sense-vectors represented by a surround defines the number of orientations radiating from a place. Each of those orientations at a place *may* permit travel along a path in a given direction. A complete place-and-path network describes all the paths radiating from each place and all the places on each path, linked by turn and travel actions. Actions other than turns and travels may also create links between places, and it is possible for the place-and-path network to have disconnected components connected only by these other actions.

The assimilation of observations into surrounds, places, and paths uses a data-driven, opportunistic algorithm, implemented as a set of rules that examine the previous and current state of the Critter after each action. The state of the Critter (adapted from the "You Are Here" Pointer in the TOUR model [Kuipers 1977, 1978]) is described in terms of its current sense-vector, surround, place, path, orientation at place, and direction of travel on path. The orientation rules allow the state of the Critter to be filled out from information already in the cognitive map, and the assimilation rules extend the cognitive map from the information in the state description. These orientation and assimilation rules build place and path descriptions incrementally under the control of the exploration strategy component.

4 Exploration Strategy

Several levels of exploration strategy are interleaved with the assimilation rules discussed in the previous section.

4.1 Determining the Current Place

After taking an action, the Critter's first priority is to determine where it is, and to compare that with where it expected to be. It executes a series of turns to collect a cycle of sense-vectors and establish the current surround without otherwise changing its state (since it knows that turns are periodic and do not change the current place). If the map is already sufficiently complete to predict the place the Critter should be at, this can be verified from the surround. If the surround is new, a new place description is created.

If the current surround has been previously encountered, it must point to at least one place, and the Critter's goal is to determine whether that place is the current place. The Critter is able to construct a reversible route between that place and other known *and distinctively recognizable* places, staying within known territory. It then tries to follow such a route and return to the questionable spot. If the route performed as predicted, then the current place matches the previously known one, and the Critter knows where it is. If not, then the current place must be a new place with the same surround as the old one. This strategy is slightly complex, but quite reliable[1], and allows the Critter, for example, to recognize that two ways around a block bring it to the same place on the other side, not just places that appear similar.

At a place, the orientation is simply the point in the cycle of sense-vectors that the Critter is currently facing.

[1]provable?

4.2 Exploring the Current Path

Once the current place and orientation are known, the Critter wants to know what path it is on in that direction. Its strategy for exploring a path is to follow it to the end (recognized as a prerequisite violation), turn around, follow it in the reverse direction to the other end, turn around again, and return to the starting point. During this tour, the assimilation rules automatically construct and verify the path description, linking it appropriately to the sequence of places. Since the exploration of an individual path, like exploring a place, returns the Critter to its original position, it is easy to explore all the paths that intersect at the current place.

4.3 Exploring the Network

The overall exploration of the network is a random sequence of actions, alternating with exploration of all paths intersecting at the current place. A random sequence of actions is guaranteed to reach all portions of a finite network in reasonable expected time [ref?], and the exploration of all paths at a place quickly constructs the complete place-and-path network. When the place-and-path network is complete, the exploration strategy can concentrate on recording the consequences of the non-turn, non-travel actions until the entire environment is known. Other deterministic actions are easy to explore, and random actions are repeated until the number of times each has been executed in each context passes a threshold.

5 Assessment of the Map-Learning Critter

5.1 Results of the Contest

The actual environment on which the Critters were tested (in February 1985) had the following properties, many of which conflicted with the map-learning Critter's assumptions about a spatial environment.

- A cyclic path of five rooms. This meant that the Critter diagnosed Go-Forward as a turn action, and that a state of the five rooms was considered a surround, and corresponded to a single place description.
- No turn actions. The action Turn-Around has period 2, and was diagnosed as an Other-Action, interpreted by the Critter as connecting two adjacent places.
- A movable landmark: a stone which could be thrown. Since the configuration of observable properties of the five rooms constitutes the Critter's concept of a place, throwing the stone was interpreted as motion to a different place.
- Fixed landmarks with changeable state: lightbulbs which could be on/off and fixed/broken by various actions. These similarly are interpreted as additional places in the map.
- Modifiable internal state: a "happiness" parameter whose state was determined by changes to the current sensory input. This also resulted in new places being added to the map.

Not surprisingly, the map-learning Critter scored less well than the conditional-probability Critter in the actual contest. Neither program did particularly well.

- The conditional-probability Critter accumulated a table of conditional probabilities, and predicted possible outcomes of actions based on the table. Since the environment included spatial and non-spatial state variables not captured by the table, its performance was mediocre.

- The map-learning Critter attempted to build a map of the environment. Since non-spatial state variables in the environment were mapped as additional sets of places, the map quickly grew too large for the Critter to explore. When a proposed sequence of actions fell within the map, the Critter made a single correct prediction. More often, the sequence of actions fell outside the map, and the Critter made no prediction at all. Its performance was poor.

5.2 A Successful World for the Map-Learner

An environment where the map-learning Critter's strategy is successful is a four-by-four grid where the four outer walls have different colors, and the actions are Turn-Right, Turn-Left, and Step. Its senses describe the color directly ahead, and a bump on stepping into a wall. Many situations, for example the four central squares, are indistinguishable both from the current sense-vector and from the current surround. The Critter successfully explores this space, creating descriptions for the nine surrounds, sixteen places, and eight paths in the complete spatial map.

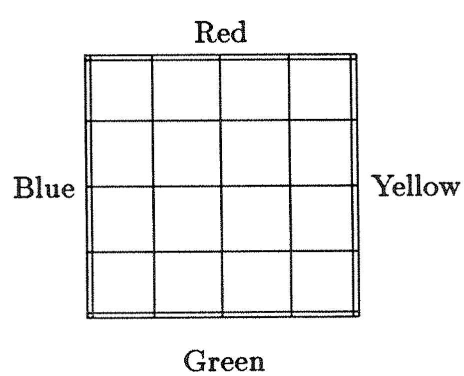

It is worth emphasizing here that the Critter does not make any metrical assumptions about the two-dimensionality of the environment, or about the magnitudes of steps or turns, so it does not impose a 2-D Cartesian model on the space. The Critter's approach, starting with a sensorimotor description and building a topological map, is more robust in the face of complex environments, uncertain information, and possibly limited computational resources. The TOUR Model [Kuipers, 1977, 1978] on which the map-learning Critter is based, builds a metrical description on top of the topological one, but that feature was not included here.

5.3 The Critter's Assumptions

It is difficult to make sense of the "blooming, buzzing confusion" of the world. If its world is sufficiently complex, it is not surprising that the Critter is unable to create a useful model of its regularities, starting where it does. The map-learning Critter is an expert on actions and their effects, and on topological spatial relations. Its model-driven approach is based on the need to factor the actions and sensory inputs into categories reflecting the underlying structure of the environment. In a case such as the above, where the Critter lacked the model elements needed to cope with certain aspects of the environment, it attempted to treat them spatially, as additional sets of places, and was overwhelmed.

When the Critter does succeed in mapping its environment, it is also not surprising that its map can be significantly different than that of the programmer. In retrospect, the Critter's interpretation of Go-Forward as a turn in a cyclic world, and its treatment of the configuration of the five rooms as a single place, is entirely reasonable. Similarly, in a test example involving an elevator which moves only when its doors are closed, the Critter's map consists of a linear "Main Street" with dead-end cross streets one block long corresponding to the Door-Open/Door-Closed actions.

It is instructive to consider different types of worlds for which the Critter creates a non-standard interpretation.

- Suppose there are several distinct turn actions (not just two inverses). Consider the four-by-four grid described above, but with edges identified so that the space is toroidal. All three actions would then be diagnosed as Turns, and there are three distinct notions of surround and place, corresponding to turning in place at one spot, or following a "great circle" in either of two directions around the torus. From the Critter's limited perspective, there is no reason to pick any of the given actions as the "real" turn.

- In a world with no turn actions, such as the elevator, the definition of a place, and orientation at the place, is more difficult but not impossible.
- Suppose the sensory information is misleading about the true period of turn actions. Consider the four-by-four grid with the walls painted in only two alternating colors. The actions Turn-Right and Turn-Left would have period two, and therefore not be diagnosed as turns, so the map of the environment would be quite different. The problem here is that diagnosis of the actions precedes exploration of the environment, and information about the structure of paths is needed to determine that those actions *really* have period four.
- An environment of tangled paths leads to odd behavior, since the Critter expects that paths are linear, finite, and invertible.

6 Future Directions

6.1 Smarter Critter

Various minor improvements would significantly improve the performance of the Critter on problems such as this contest. There are more substantial research questions buried under some of these issues, however:

- Action diagnosis should include a more general theory of prerequisites, context-sensitivity, and inverses. The Other-Actions and Random-Actions can doubtless be more satisfactorily described.
- A more sophisticated approach to the Critter problem would attempt to separate out the spatial aspects of the environment and understand them independently of non-spatial aspects. The spatial model could then serve as a skeleton on which to explore the other aspects of the environment. However, since the distinction between spatial and non-spatial aspects is not explicit at the sense/action level of description, it is not clear how to do this factoring.

6.2 Theorems Relating the Map and the World

Given an appropriate axiomatization of the environment, it should be possible to prove by induction that the map created by the Critter actually reflects the structure of the environment. That is, at each step, given that the current map accurately reflects the environment, a new place or path must also accurately reflect the environment.

However, environments can have different levels of accessibility to the map-learning Critter, depending on the richness of the sensory images it provides.

1. every sensory image is distinguishable;
2. every place is distinguishable (i.e. has a unique surround);

3. there is some place which is uniquely distinguishable;

4. there is no uniquely distinguishable places (i.e. every surround has at least two places).

A class of map-learning Critters could also have several levels of representation, enabling it to deal with environments of different levels of complexity.

1. A place is identified with an equivalence class of sensory images under turn actions;

2. There are distinct concepts of place and surround, with explicit correspondences as in the Critter described here;

3. A more elaborate diagnostic process is required to determine place from surround, requiring commitment to the current place to be withheld for significant periods of time.

It should be possible to set up appropriate levels of both map-learner and environment, and prove necessary and sufficient theorems about the level of representation needed to construct a valid map of each level of environment.

On the other hand, what are the properties of the invalid map produced by a map-learner which is weaker than its environment requires?

6.3 Final Reflections

It may be argued that there should be some simple "magic bullet" algorithm that will solve the Critter Problem, but this appears to me to be a reemergence of the naive optimism that characterized AI in its first decade. My conviction is that the Critter Problem is very hard, and will be solved only with a complex, model-driven approach that builds in certain assumptions about the types of regularities the world can contain. This conviction stems from several observations about spatial knowledge in humans.

- Spatial knowledge provides a fundamental context within which most other kinds of knowledge are acquired and interpreted. It is also frequently applied metaphorically to non–spatial relationships.
- People take many years to develop an adequate understanding of their spatial environments, passing through many distinct states of partial knowledge on the way.
- The human representation of spatial knowledge consists of successive layers, first abstracting away from a sensorimotor description to a description of fixed environmental features, then building metrical on top of topological spatial relations.
- While there is a great deal of individual variation in spatial representations, this variation fits within a simple framework, leading one to suspect that there may be no other cognitively feasible ways to describe a large-scale spatial environment.

7 References

B. J. Kuipers. 1977. Representing Knowledge of Large-Scale Space. Cambridge, MA: MIT Artificial Intelligence Laboratory TR-418.

Kuipers, B. J. 1978. Modeling spatial knowledge. *Cognitive Science*, **2**: 129-153.

Kuipers, B. J. 1979. Commonsense knowledge of space: Learning from experience. In *Proceedings of the Sixth International Joint Conference on Artificial Intelligence.* Stanford, California.

Kuipers, B. J. 1983. Modeling human knowledge of routes: Partial knowledge and individual variation. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-83).* Washington, D.C.

Acknowledgements

I am very much indebted to Tom Russ, the winner of the First Annual Critter Contest, the implementor of the conditional-probability Critter, and the creator of numerous worlds to test the map-learning Critter. Without those worlds to learn from, and the surprises they contained, this Critter would be even more simple-minded than it turned out to be.