From: AAAI Technical Report SS-01-05. Compilation copyright © 2001, AAAI (www.aaai.org). All rights reserved.

Learning from Uninterpreted Experience in the SSH*

Benjamin Kuipers, Patrick Beeson, Joseph Modayil and Jefferson Provost
Artificial Intelligence Lab, University of Texas at Austin, Austin, Texas 78712
{kuipers,pbeeson,modayil,jp}cs.utexas.edu

Suppose a robot is born into the world with sensors and effectors, but no knowledge of their relationship to the unknown external world. How can the robot learn an interpretation of its sensorimotor experience adequate for it to explore its environment and build a cognitive map? Pierce and Kuipers (Pierce & Kuipers 1997) solved a very restricted version of this problem, for a simulated mobile robot whose primary sensory input was a ring of sonar-like range sensors. We are working toward a more general solution, applied to physical robots with multiple sensory modalities.

We are working within the framework of the Spatial Semantic Hierarchy (SSH) (Kuipers 2000), which is a lattice of related representations for large-scale space. The SSH is conducive to research on learning from uninterpreted experience because it separates the interface representations — the representations for local control laws and local metrical models — from the symbolic representations built on them.

We use two robot platforms in this research, both with multiple, partially redundant, sensory modalities. Lassie is an RWI Magellan with laser range-finder, 16 sonar sensors, 16 IR range sensors, 16 bump sensors and odometry. Vulcan is a custom-built robotic wheelchair with two laser range-finders, binocular cameras on pan-tilt heads, 7 sonars, 12 IR proximity sensors, and odometry.

We have decomposed the problem of learning from uninterpreted sensorimotor experience into three loosely-coupled learning problems.

- **Feature learning**. Learn a hierarchy of features defined on top of the raw sensory features, and identify reliable relationships among features, and between features and motor signals.
- **Control learning**. Given a hierarchy of sensory features and primitive relationships with motor signals, learn effective control laws for achieving and maintaining desired states.
- **Place recognition**. Given sensory images from the robot's sensory system, learn to identify distinctive states, and hence places.

*This work has taken place in the Intelligent Robotics Lab at the Artificial Intelligence Laboratory, The University of Texas at Austin. Research of the Intelligent Robotics lab is supported in part by NSF grants IRI-9504138 and CDA 9617327, and by funding from Tivoli Corporation.

Feature Learning

The first subgoal of the feature learning task is to evaluate the effectiveness of Pierce and Kuipers' existing methods (Pierce & Kuipers 1997) on real sensor data collected from physical robots with multiple sensor modalities. The learning algorithm needs to group the sensors according to modality, then use similarities to map individual sensors into a spatial configuration for each group. This is a critical step in feature-generation because it changes the ontology of the sensory system by allowing spatial as well as temporal derivatives to be defined on the sensory input. This makes it possible for the agent to reason with sensory motion and sensory flow, rather than simply sensory change. It also transforms the index for an individual sensor from a simple designator into a meaningful feature value in a continuous range (the spatial configuration).

The question is whether existing methods can be extended to handle the complex systematic errors that physical sensors exhibit, and whether they can discriminate among distinct sensory modalities that provide redundant information about the same environment.

The second subgoal of feature learning is to identify the primitive relations between sensory features and motor signals. Average sensor motion vector fields can be used to identify primitive actions (Pierce & Kuipers 1997). Then the agent can identify which sensory features are most strongly related to the magnitude of each action. For example, forward motion is most closely linked to the value of the forward range sensor, and turn motion is most closely linked to the position of features in a sensor field such as a ring. (A closely related problem is identifying contextually-dependent regularities between motor signals and sensory features, but we discuss this under Control Learning.)

The third subgoal of feature learning is to identify and characterize the redundancy among sensory modalities. The simplest example is the redundancy between laser and sonar range readings. On Lassie, the laser range-finder returns 180 highly accurate distance readings over the front 180 degrees of arc, while the sonar ring returns 16 range readings over a 360 degree arc, subject to significant systematic errors due to specular reflection. A more complex redundancy is between the range sensors and odometry. Movement of the robot causes systematic changes to the odometry sensors and to the range values returned. Turning rotates the range val-

ues returned among the sensors in each modality. Travel has direct effect on forward and rear range sensors, but less or no effect on lateral range sensors. Yet another redundancy is the relation between range sensors and bump sensors. When travel reduces the forward range sensor to a certain value, a corresponding bump sensor is likely to signal.

One issue is how to represent the individual redundancies as they are discovered. A deeper issue is whether it is possible to make the abductive transition from an egocentric model of sensor input to an allocentric (world-centered) model of obstacles with locations in a fixed frame of reference, in which robot actions change the robot's position. This transition supports the local metrical map portion of the SSH Control Level. It is not absolutely necessary for exploration, since the ability to execute control laws can be achieved egocentrically without knowledge of local geometry, and can support the SSH Causal and Topological levels. However, if it can be achieved, it accounts for the origin of local fixed frames of reference, which are the foundation for local metrical maps, then the patchwork map, and finally the global metrical map.

Control Learning

Once connections have been learned between motor signals and certain sensory feature values, the goal of control learning is to construct the control laws that make up the SSH Control level, and thus support the rest of the cognitive map. We need to learn control laws for hill-climbing to locally distinctive states once in a neighborhood, and to follow trajectories from one distinctive state to the neighborhood of the next.

The first subgoal in control learning is to identify *contexts* (states where certain features have certain values) in which particular relationships between motor signal and feature values are reliably observed. For example, if a local minimum range value is directly ahead, then a travel action will reliably decrease it, and if the local minimum range value is directly in a lateral position, then a travel action will tend to leave it constant. This subgoal clearly interacts strongly with the feature-learning part of the project. A particular problem is that temporal derivatives are mostly likely to be related to motor signals, and they are notoriously noisy when computed from real data.

The second subgoal is to compose these relationships into control laws to (a) move a specified feature to a desired value, (b) select a context in which one feature value can be changed while leaving another feature value relatively unchanged, and (c) make closed-loop adjustments to keep one feature close to a desired value while changing the other.

The approach to defining control laws in (Pierce & Kuipers 1997, sect. 5) uses proportional-integral control law templates that encode a significant amount of background knowledge from control engineering. The third subgoal of this task is to approach the control learning problem using methods from reinforcement learning to achieve better results while presupposing smaller and more plausible amounts of background knowledge.

A control law typically includes several functional relationships between sensory features and the motor output. We are experimenting with ways to optimize nonlinear shapes for these functional relationships using reinforcement from overall task performance.

Place Recognition

Once the control laws are functioning as desired, the robot can move reliably from one distinctive state to others throughout the environment. At each of these distinctive states, the robot's sensory system provides a sensory *image*. There will be a significant amount of variation in the images obtained at a given distinctive state, due to a number of factors including variation in the precise position the robot is in when the hill-climbing control law terminates, differences in clutter or transient pedestrians, or even illumination differences due to different times of the day or seasons of the year.

Because of this variation, it is essential to cluster images, so the images associated with the same distinctive state fall into the same cluster, which we call a *view*. The clustering can be done with an *unsupervised* learning method, for example Self-Organizing Maps (SOMs) (Kohonen 1995).

Unfortunately, different distinctive states may well have the same view, leading to *sensory aliasing*. That is, the current view may not be sufficient to determine the current distinctive state. (This problem can never be completely solved, since two distinctive states may have identical physical surroundings.)

There are a variety of methods for using contextual information to identify the current distinctive state from its connectivity with other distinctive states in the environment (Kuipers & Byun 1991). If this is successful, we can use a *supervised* learning method such as a backpropagation network to learn the mapping directly from images to distinctive states. We hypothesize that the hidden units in the backprop network will learn features of the image that might be too subtle, or too dominated by noise, to be usable for clustering by an unsupervised learner, but which can be identified with the additional feedback provided by the correct identity of the distinctive state.

If successful, this method results in the ability to recognize the current distinctive state from the current sensory image, focusing on the features that provide useful information and ignoring larger amounts of sensory noise. We are in the midst of testing this method using laser range-finder data, and we hope to extend it to visual sensor input.

References

Kohonen, T. 1995. *Self-Organizing Maps*. Berlin; New York: Springer.

Kuipers, B. J., and Byun, Y.-T. 1991. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems* 8:47–63.

Kuipers, B. J. 2000. The spatial semantic hierarchy. *Artificial Intelligence* 119:191–233.

Pierce, D. M., and Kuipers, B. J. 1997. Map learning with uninterpreted sensors and effectors. *Artificial Intelligence* 92:169–227.