

Symbol Grounding With Delay Coordinates

Michael T. Rosenstein and Paul R. Cohen

Computer Science Department, LGRC

University of Massachusetts

Box 34610

Amherst, MA 01003-4610

{mtr, cohen}@cs.umass.edu

“There is nothing more basic than categorization to our thought, perception, action, and speech” (Lakoff 1984). Moreover, categories of sensory experience provide the semantic glue between the world and the meaningless symbols often used to represent those experiences. As such, the focus of this work is an unsupervised learning mechanism for extracting categories from time series. We have in mind the situation where a sensorimotor agent, such as an infant or mobile robot, records streams of sensor readings while interacting with a complex environment. To make the leap from percepts to symbolic thought and language, the agent requires a way of transforming uninterpreted sensor information into meaningful categories. That is, the agent must solve the bottom-up version of the symbol grounding problem (Harnad 1990). The solution outlined below was inspired by the *method of delays*, a nonlinear dynamics tool for producing spatial representations of time-based data.

One general technique for discovering categories is to form clusters of points in a suitable space. This was the basis of Elman’s work on learning lexical classes from word co-occurrence statistics (Elman 1990). Elman first trained a recurrent neural network to predict successive words in a long input string. This then set the stage for hierarchical clustering of the hidden-unit activation space, where the result was groups of words that coincide with classes like NOUN-FOOD or VERB-PERCEPT. But how can we ground such syntactic classes in sensorimotor interaction with the environment?

We answer this question, not with recurrent neural networks, which require a great deal of training, but with the method of delays, which maps a snippet of time series data to a point in *delay-coordinate space*. Delay coordinates are just successive sensor readings taken at a suitable time interval (Takens 1981; Rosenstein, Collins, & De Luca 1994). For instance, imagine that our mobile robot is two meters from an obstacle and its sonar reports the value 2015 mm. From this single reading there is no way to know if the robot will crash into the obstacle. But suppose the sonar reports 2140 mm a second later. We could plot the point

(2140, 2015) in a two-dimensional space, and notice that all nearby points correspond to the situation where the robot moves away from an object. This example may seem trivial, yet it illustrates the intuition behind the method of the delays: if the state of the environment at time t is uncertain, then examination of sensor readings prior to time t will reduce the ambiguity.

Our first application of clustering with delay coordinates involved a simulated world where agents adopted one of several strategies for pursuit or avoidance of an opponent (Rosenstein *et al.* 1997; Rosenstein & Cohen 1998). Whenever two agents interacted, the simulator recorded the distance between them as well as the final outcome (contact, escape, or perpetual chase). This distance time series was the input to an incremental algorithm that clustered points based on Euclidean distance in delay-coordinate space. Nearby points were merged to create an average time series, or *prototype*. A prototype is best understood as a representative class member, and our implementation is entirely consistent with Rosch’s influential work on categories and prototypes (Rosch & Lloyd 1978).

Figure 1 shows the six prototypes derived from 100 agent interactions where the movement strategies were chosen randomly. These prototypes reflect actual differences in the simulated world, and one could describe the corresponding categories as “chase,” “contact,” contact after the agents first “overshoot” one another, and “escape” with short, medium, and long escape times. These categories not only were grounded in physical measurements, but also were formed in an entirely unsupervised manner. But where is meaning?

For this work the meaning associated with each category is the predictions it affords the agent; we call this a *predictive semantics*. So while a category may prove useful simply as an abstraction of sensory experience, the real benefit comes from its *entailments*, or consequences of category membership. In the pursuit/avoidance simulator, our agents like to predict outcomes, such as contact, escape, or chase. When the learning algorithm finds clusters of points in delay-coordinate space, it also keeps a count of the number of member interactions that finish in each of the three outcomes. The members themselves are forgotten, ex-

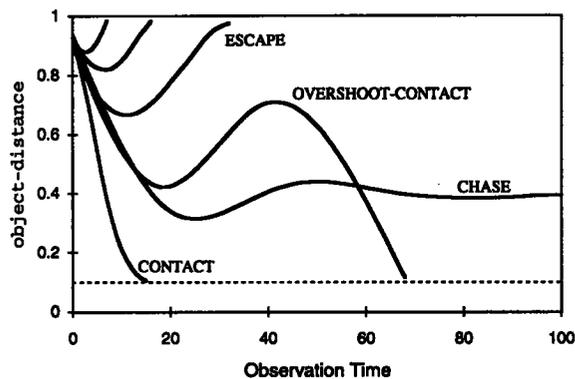


Figure 1: Cluster prototypes for 100 interactions in the pursuit/avoidance simulator.

cept for their influence on the prototype. For subsequent interactions, i.e., for new time series of distance measurements, the algorithm computes the prototype nearest the time series in delay-coordinate space and reports the majority outcome for the corresponding cluster. This straightforward prediction scheme performed quite well at recognizing new situations and predicting their outcomes.

Our present interest is the transfer of this work to Pioneer mobile robots, which come equipped with over 30 sensors including a real-time vision system and gripper. (For related work see (Schmill *et al.* 1998).) In a preliminary experiment, the robot moved around an office environment, occasionally bumping into objects. One way for the robot to sense objects is by contact with the bump switches at the tips of its gripper paddles; such events cause the corresponding sensor readings to change from low to high. In this experiment, the clustering algorithm discovered several categories, including two that involved the bump sensors. Comparison of the category prototypes showed that the robot's stall sensors triggered for one cluster but not the other. Interestingly, the former cluster was due to events where the robot hit a wall, whereas the latter was the result of all times the robot pushed a waste basket several centimeters across the floor. These prototypes suggest that the robot discovered two categories with complementary semantics: IMMOVABLE and MOVABLE.

Perhaps the most promising aspect of this work is the possibility of an agent discovering meaningful categories for itself. In particular, this work shows the value of delay coordinates for transforming time series into clusters of points. And since time series are grounded in measurements out in the world, the resulting categories provide the foundation on which to build a symbolic, grounded, conceptual structure.

Acknowledgments

This research is supported by the National Defense Science and Engineering Graduate Fellowship and by DARPA under contract No. N66001-96-C-8504. The

U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon. The views and conclusions contained herein are those of the authors and should not be interpreted as the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

References

- Elman, J. L. 1990. Finding structure in time. *Cognitive Science* 14:179-211.
- Harnad, S. 1990. The symbol grounding problem. *Physica D* 42:335-346.
- Lakoff, G. 1984. *Women, Fire, and Dangerous Things*. University of Chicago Press.
- Rosch, E., and Lloyd, B. B. 1978. *Cognition and Categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Rosenstein, M., and Cohen, P. R. 1998. Concepts from time series. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*. AAAI Press. To appear.
- Rosenstein, M.; Cohen, P. R.; Schmill, M. D.; and Atkin, M. S. 1997. Action representation, prediction and concepts. University of Massachusetts Computer Science Department Technical Report 97-31, also presented at the 1997 AAAI Workshop on Robots, Softbots, Immobots: Theories of Action, Planning and Control.
- Rosenstein, M. T.; Collins, J. J.; and De Luca, C. J. 1994. Reconstruction expansion as a geometry-based framework for choosing proper delay times. *Physica D* 73:82-98.
- Schmill, M. D.; Rosenstein, M. T.; Cohen, P. R.; and Utgoff, P. 1998. Learning what is relevant to the effects of actions for a mobile robot. To appear in *Proceedings of the Second International Conference on Autonomous Agents*.
- Takens, F. 1981. Detecting strange attractors in turbulence. *Lecture Notes in Mathematics* 898:366-381.