

Protosymbols that integrate recognition and response

Karl F. MacDorman^{*†}

Rawichote Chalodhorn[†]

Hiroshi Ishiguro[†]

Minoru Asada^{*†}

^{*}Frontier Research Center

[†]Department of Adaptive Machine Systems

Graduate School of Engineering

Osaka University

Suita, Osaka 565-0871 Japan

kfm@ed.ams.eng.osaka-u.ac.jp

Abstract

We explore two controversial hypotheses through robotic implementation: (1) Processes involved in recognition and response are tightly coupled both in their operation and epigenesis; and (2) processes involved in symbol emergence should respect the integrity of recognition and response while exploiting the periodicity of biological motion. To that end, this paper proposes a method of recognizing and generating motion patterns based on nonlinear principal component neural networks that are constrained to model both periodic and transitional movements. The method is evaluated by an examination of its ability to segment and generalize different kinds of soccer playing activity during a RoboCup match.

1. Introduction

Complex organisms recognize their relation to their surroundings and act accordingly. The above sentence sounds like a truism owing in part to the almost ubiquitous distinction between recognition and response in academic disciplines. Engineering has successfully developed pattern recognition and control as independent fields, and cognitive psychology and neuroscience often distinguish between sensory and motor processing with researchers specializing in one area or the other. Nevertheless, in some sense recognition and response entail one another. Recognizing an object, action, or sign is largely a matter of recognizing what it does for us and what we can do with it. Indeed, much of what we perceive can be described in terms of potential actions. *Doing* and *seeing* cannot so readily be distinguished because we acquaint ourselves with our world through what we do and our actions drive what distinctions we learn to make. None of this is meant to deny that we can experimentally isolate purely motor centers in the brain from purely sensory




Figure 1: In the proposed approach, a neural network learns each kind of periodic or transitional movement in order to recognize and to generate it. Recent sensorimotor data elicit activity in corresponding networks, which segment the data and produce appropriate anticipatory responses. Active networks constitute an organism's conceptualization of the world since they embody expectations, derived from experience, about the outcomes of acts and what leads to what. It is assumed that behavior is purposive: affective appraisals guide the system toward desired states.

ones, but rather to assert that these centers are intimately linked both in their everyday operation and in their epigenetic development. Thus, as scientists and engineers, we may have reified the distinction between recognition and response, when their main difference is merely in descriptive focus.

In this paper, we will entertain and begin to explore two controversially and, as yet, unproven hypotheses: First, there is an integrity of recognition and response. We recognize an object or event largely because it elicits expectation about what we can do with it — or at least piggybacks on those kinds of expectations. In addition, these expectations are expressed in terms of (or decontextualized from) how motor signals transform sensory data. Second, biological motion is in some sense periodic. To put it simply, patterns repeat. (If they did not, there would be little point in learning.) That is as much a function of the 'hardware' as it is the often routine nature of existence: Joints, for example, have a limited range and will

eventually return, more or less, to a given configuration. Moreover, bodies have certain preferred states: for people walking is a more efficient means of locomotion than flailing about randomly. All gaits exhibit a certain periodicity as do many gestures and vocalizations.

This paper proposes a method of generalizing, recognizing, and generating patterns of behavior based on nonlinear principal component neural networks that are constrained to model both periodic and transitional movements. Each network is abstracted from a particular kind of movement. Learning is competitive because sensorimotor patterns that one network cannot learn will be assigned to another network, and redundant networks will be eliminated and their corresponding data reassigned to the most plausible alternative. Recognition is also competitive because proprioceptive data is associated with the network that best predicts it. (The data can be purely kinematic or dynamic depending on the dimensions of the sensorimotor phase space.) Since each network can recognize, learn, and generalize a particular type of motion and generate its generalization, the integrity of recognition and response are maintained. These generalizations are grounded in sensorimotor experience. They can be varied, depending on the networks' parameterization. They may be viewed as a kind of protosymbol. While we do not claim that the networks have neural analogues, we believe the brain must be able to implement similar functions.

1.1 The emergence of signs in communication

In one vein, we are exploring the application of periodically-constrained NLPCA neural networks to vocal and gesture recognition and generation. Our aim is to develop robots whose activity is capable of supporting the emergence of shared signs during communication. Signs take on meaning in a given situation and relationship, as influenced by an individual's emotional responses and motivation (see Figure 1). They reflect mutual expectations that develop over the course of many interactions. We hypothesize that signs provide developmental scaffolding for symbol emergence. For infants, the caregiver's intentions are key to fostering the development of shared signs.

We believe that periodically-constrained NLPCA neural networks could be one of the embedded mechanisms that support the development of shared signs. We are testing this hypothesis by comparing the behavior generalized by these neural networks with motion tracking data from mother-infant interactions.¹ The results of behavioral studies are applied to the android robot, Actroid, which has 33 degrees of freedom.

¹From this we have ascertained that certain important micro-behaviors that make movement seem lifelike may have been overlooked in the approach outlined here, and we are starting to develop a micro-behavior filter.

1.2 Outline

This paper is organized as follows. Section 2 extends an NLPCNN with periodic and temporal constraints. Section 3 presents a method of assigning observations to NLPCNNs to segment proprioceptive data. Section 4 reports experimental results using NLPCNNs to characterize the behavior of a Fujitsu HOAP-1 humanoid robot that has been developed to play RoboCup soccer.

2. A periodic nonlinear principal component neural network

The human body has 244 degrees of freedom (Zatsiorsky, 2002) and a vast array of proprioceptors. Excluding the hands, a humanoid robot generally has at least 20 degrees of freedom — and far more dimensions are required to describe its dynamics precisely.




Figure 2: An NLPCA neural network with the activations of nodes p and q constrained to lie on the unit circle. Target values presented at the output layer of a nonlinear principal component neural network are identical to input values. Nonlinear units comprise the encoding and decoding layers, while either linear or nonlinear units comprise the feature and output layers.

Given a coding function $f : \mathbb{R}^N \mapsto \mathbb{R}^P$ and decoding function $g : \mathbb{R}^P \mapsto \mathbb{R}^N$ that belong to the sets of continuous nonlinear functions \mathcal{C} and \mathcal{D} , respectively, where $P < N$, nonlinear principle component networks minimize the error function E

$$\|\vec{x} - g(f(\vec{x}))\|^2, \quad \vec{x} \in \mathbb{R}^N$$

resulting in P principal components $[y_1 \cdots y_p] = f(\vec{x})$. Kramer (1991) first solved this problem by training a multilayer perceptron using the backpropagation of error.

2.1 The periodicity constraint

Kirby and Miranda (1996) constrained the activation values of a pair of nodes p and q in the feature layer of an NLPCNN to fall on the unit circle, thus acting as a single angular variable:

$$r = \sqrt{y_p^2 + y_q^2}, \quad y_p \leftarrow y_p/r, \quad y_q \leftarrow y_q/r$$

The delta values for backpropagation of the circular node-pair are calculated by the chain rule

(Kirby and Miranda, 1996), resulting in the update rule

$$\delta_p \leftarrow (\delta_p y_q - \delta_q y_p) y_q / r^3, \quad \delta_q \leftarrow (\delta_q y_p - \delta_p y_q) y_p / r^3$$

at the feature layer.

3. Automatic segmentation

We conceived of the automatic segmentation problem as the problem of uniquely assigning data points to nonlinear principal component neural networks. As the robot begins to move, the first network is assigned some minimal number of data points (e.g., joint-angle vectors), and its training begins with those points. This gets the network’s learning started quickly and provides it with enough information to determine the orientation and curvature of the trajectory. If the average prediction error of the data points assigned to a network is below some threshold, the network is assigned additional data points until that threshold has been reached. At that point, data points will be assigned to another network, and a network will be created, if it does not already exist. To avoid instabilities, only a single data point may shift its assignment from one network to another after each training cycle.

Since a network is allowed to learn more data points as long as its average prediction error per point is low enough, it may learn most data points well but exhibit slack near peripheral or recently learned data points. At the start of learning, the network should be challenged to learn data points even when its prediction error is large. As learning converges, however, the slack leads to segmentation errors. Therefore, we alter the method of segmentation once the network nears convergence so that a network may acquire neighboring points if its prediction error for those points is lower than the network currently assigned to those points.

4. Humanoid experiments

This section shows the result of automatic segmentation and neural network learning. We assess the accuracy of the result based on a manual segmentation of the data points and an analysis of how they are allocated among the networks.




Figure 3: Fujitsu HOAP-1 robots are playing RoboCup soccer.




Figure 4: Recognized motion patterns embedded in the dimensions of the first three nonlinear principal components of the raw proprioceptive data. The top and bottom plots differ only in the viewpoint used for visualization.




Figure 5: The percentage of data points allocated to each network before and after eliminating redundant networks and reassigning their data.

First, we recorded motion data while a HOAP-1 humanoid robot played soccer. Each data point is constituted by a 20-dimensional vector of joint angles. A standard (noncircular) NLPCNN reduced the dimensionality of the data from 20 to 3. We then applied our algorithm to segment, generalize, and generate humanoid motion.

Our algorithm uniquely assigned the data points among a number of circularly-constrained NLPCNNs. Each of the networks learned a periodic motion pattern by conjugate gradients. Our algorithm successfully generalized five out of six primary motion patterns: walking forward, turning right or left, and side-stepping to the right or left. It failed to generalize as a single periodic trajectory the kicking motion, which has a highly irregular, self-intersecting shape. However, human subjects were also unable to determine the kicking trajectory from the data points. Figure 4 shows that the automatic segmentation algorithm successfully employed circular NLPCNNs to separate and generalize five of the periodic motions.

We calculated statistics based on running the automatic segmentation for 20 trials. The algorithm resulted in five

decoding subnetworks for 45% of the trials, which is the most parsimonious solution. It resulted in six subnetworks for 50% of the trials, and seven for the remaining 5%.

In a separate run of the learning and segmentation algorithm, the motion sequence of recorded data during soccer playing was walking forward, turning right, turning left, walking forward, sidestepping to the right, sidestepping to the left, and kicking. We counted the number of point belonging to each network before and after removing redundant networks. Redundant networks were removed by means of linear integration. The angular value θ was varied from 0 to 2π at the bottleneck layer of one network to obtain its predicted output. This value was fed into another network to obtain its predicted value. If the integral of the sum of the squared distances of the predicted outputs was less than a threshold, one network was removed and its points reassigned to the other network (see Figure 5). This method removed all redundant networks.

5. Conclusion

Our proposed algorithm abstracted five out of six types of humanoid motion through a process that combines learning and data point assignment among multiple neural networks. The networks perform periodic, temporally-constrained nonlinear principal component analysis. The decoding subnetworks generate motion patterns that accurately correspond to the five motions without including outliers caused by nondeterministic perturbations in the data. By means of linear integration, we were able to remove redundant networks according to the proximity of their predictions.

A kind of behavior can be recognized by selecting the network that best predicts joint-angle values. It can be generated by varying the value of θ in the bottleneck layer. This shows the effectiveness of a tight coupling between recognition and response since the same networks may be used for both processes and they developed by the same mechanisms. The significance of periodicity may be more limited, however. Some motions are not periodic, and in the experiment the kicking motion, although it occurs repeatedly, was difficult to segment because of its highly irregular, self-intersecting shape.

References

- Chalodhorn, R., Aono, M., Ooga, J., Ogino, M., and Asada, M. (2003). Osaka University "Senchans 2003". In Browning, B., Polani, D., Bonarini, A., and Yoshida, K., (Eds.), *RoboCup-2003: Robot Soccer World Cup VII*. Springer Verlag.
- Fujii, T. (1987). A new approach to the LQ design from the viewpoint of the inverse regulator problem. *IEEE Transactions on Automatic Control*, 32(11):995–1004.
- Inamura, T., Toshima, I., and Nakamura, Y. (2003). Acquiring motion elements for bidirectional computation of motion recognition and generation. In Siciliano, B. and Dario, P., (Eds.), *Experimental Robotics VIII*, pages 372–381. Springer.
- Kirby, M. J. and Miranda, R. (1996). Circular nodes in neural networks. *Neural Computation*, 8(2):390–402.
- Kramer, M. A. (1991). Nonlinear principal component analysis using autoassociative neural networks. *Journal of the American Institute of Chemical Engineers*, 37(2):233–243.
- LeCun, Y., Bottou, L., Orr, G. B., and Müller, K.-R. (1998). Efficient BackProp. In Orr, G. B. and Müller, K.-R., (Eds.), *Neural Networks: Tricks of the Trade*, pages 1–44. Springer.
- MacKay, D. J. (1995). Probable networks and plausible predictions: A review of practical Bayesian methods for supervised neural networks. *Network: Computation in Neural Systems*, 6:469–505.
- Malthouse, E. C. (1998). Limitations of nonlinear PCA as performed with generic neural networks. *IEEE Transactions on Neural Networks*, 9(1):165–173.
- Okada, M., Tatani, K., and Nakamura, Y. (2002). Polynomial design of the nonlinear dynamics for the brain-like information processing of whole body motion. In *IEEE International Conference on Robotics and Automation*, pages 1410–1415.
- Ridella, S., Rovetta, S., and Zunino, R. (1995). Adaptive internal representation in circular back-propagation networks. *Neural Computing and Applications*, 3(4):222–333.
- Ridella, S., Rovetta, S., and Zunino, R. (1997). Circular backpropagation networks for classification. *IEEE Transaction on Neural Networks*, 8(1):84–97.
- Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, Massachusetts.
- Tatani, K. and Nakamura, Y. (2003). Dimensionality reduction and reproduction with hierarchical NLPCA neural networks extracting common space of multiple humanoid motion patterns. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1927–1932, Taipei, Taiwan.
- Tenenbaum, J. B., de Silva, V., and Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323.
- Zatsiorsky, V. M. (2002). *Kinematics of Human Motion*. Human Kinetics, Urbana Champaign.