

A GENERAL ARCHITECTURE FOR CLASSICAL CONDITIONING OF PERCEPTUAL-MOTOR SEQUENCES

Michael R. Blackburn

Undersea AI and Robotics
Naval Ocean Systems Center
San Diego, California 92152-5000

Abstract

A general neural network architecture, loosely modeled on the cerebral cortex, for the classical conditioning of perceptual-motor sequences is described. The utility of such an algorithm in robotics applications lies in its potential to adaptively order gross and fine motor actions under sensory control.

Introduction

Most useful behavior requires the sequencing of several motor actions under sensory control. Robotic systems to be governed autonomously by artificial neural networks are cases in point.

Classical (Pavlovian) conditioning is a learning paradigm that associates environmental events with reflex responses. We have used classical conditioning to generate perceptual-motor sequences in an artificial visual system [1]. We describe here a general architecture for classical conditioning and its application. The basic elements of classical conditioning are 1) the unconditioned stimulus (UCS) that consistently and reflexively elicits 2) an unconditioned response (UCR); 3) an initially nonprovocative but perceptible event that becomes, after repeated presentations with the UCS, a conditioned stimulus (CS) with the ability to produce 4) a conditioned response (CR) that has the characteristics of the UCR.

A simple neural circuit for classical conditioning is diagramed in Figure 1. The UCS has a fixed weight connection to the UCR element (A). Activity in the UCS is always able to produce a threshold response in A. The hidden element (B) also has a fixed weight connection to element A but the activity in B is initially too low to overcome the threshold of A. The CS has a modifiable connection to B that is initially small. The CS connection weight to B grows when both the CS and the US are active. In time, after many repeated pairings of CS and UCS, the transfer of CS activity through B to A is sufficient to cause a threshold response.

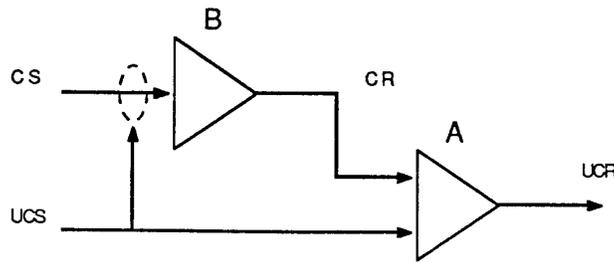


Figure 1

Elements of classical conditioning. The UCS facilitates the increase in conduction from the CS to the CR through processing element B, which has a fixed-weight, topographical relationship with the UCR element A.

The biological significance of the CS is to predict the UCS and elicit the UCR in advance of the UCS. If the UCS subsequently fails to occur, as it must if it is successfully anticipated and avoided, then prediction becomes less critical and the connection weight at B can decrease. Such an extinction of the conditioned response is observed in conditioned animals when the CS is presented without the UCS.

Prediction is fundamental to successful adaptive behavior. Prediction is also at the foundation of communication. This text would be useless to a reader who was not able to anticipate the next words (and ideas) in each series. The series history provides a context which directs the trajectory of the activity from sensors to effectors. If the context is sports, then the cue "base" reliably evokes "ball". If the context is architecture, the the same cue will evoke something else, such as "ment" or "board". When the predicted stimulus is encountered, we can smoothly move on to the next, confident that our knowledge is up to the task. If we are surprised by an unanticipated event, we may have to backtrack in time or space to make sense out of it. Knowledge, whether symbolic or practical, is demonstrated by predicting correctly, therefore, if we want to develop an autonomously intelligent machine, we must provide it with the ability to predict consequences of its actions. Classical conditioning is a mechanism to accomplish such predictions.

The problem for biological species, and by analogy for robotic species, is not so much to classify stimulus patterns as is the case with many conventional and neural network pattern recognition systems, but to generate a sequence of actions that produces something useful, such as food, or safety from being a source of food. Conveniently, perception and movement are functionally linked [2]. The perceiving system is successful when its motor activity results in an input pattern that closely matches the expected. The degree of agreement between expected and observed induces the next action in the sequence. A series of correct perceptions will follow from and contribute to a series of appropriate motor patterns. Useful products are a consequence of this process because the fundamental reflexes promote wellbeing.

Architecture

A general architecture can meet the requirements of an adaptive and predictive perceptual-motor system. The architecture includes 1) an input field of the observed pattern, 2) short-term storage of the input history, 3) an output field of the predicted pattern, 4) a comparator of the observed and predicted patterns, 5) an error field to store the differences, and 6) association matrices of the input history with a) the current input, b) the current error, and c) the next motor command.

The architecture is schematically represented in Figure 2. Only two processing elements with all of their connections are indicated in each field for clarity. Classical conditioning occurs in the network when the weights of the association matrices are modified. Algorithmic rules similar to those described by Pavlov [3] are applied. When the CS is active, the connection weights are increased in the presence of the UCS, and decreased in its absence.

The UCS for matrix A is a collateral from the reflex element that is currently in execution. The UCS for matrix B is the current sensory input, and the UCS for matrix C is the current error. In each case, the UCS is the event to be predicted. A predicted motor response will execute the pattern in advance of the reflex. The predicted pattern will interact constructively with the observed input and facilitate feature extraction, decreasing processing time between search movements. The predicted error will lead to error minimization as the accuracy of the search path increases - similar to the ease with which fluid prose begets fluid prose, or a virtuoso pianist completes her performance.

Learning in the present system is by self-organization induced by the order or correlations in the external environment. The environment is thus the teacher, but the system will learn whether or not the lesson is correct. The obvious opportunity for those who wish to control the learning in this self-organizing system is to manipulate the order of the environment.

The environment initially orders the system by acting on basic reflexes. Through classical conditioning, stimulus patterns become associated with the reflexes and are later able to generate the behavior in advance (and consequently in the absence of) the evoking stimulus. For example, in vision, the CS can be the features of the pattern present on the fovea when the UCS, which can be a peripherally detected event, causes a reflex adjustment of fixation and changes the pattern. Subsequently, a particular pattern falling on the fovea, such as the image of someone's nose, may cause a change in fixation to a point where an eye or ear would be expected. The subsequent observation of the predicted feature can release other associated motor sequences indicative of recognition, such as the vocalization of a greeting or a name.

Predictions about the environment are based on an interaction of the recent history with

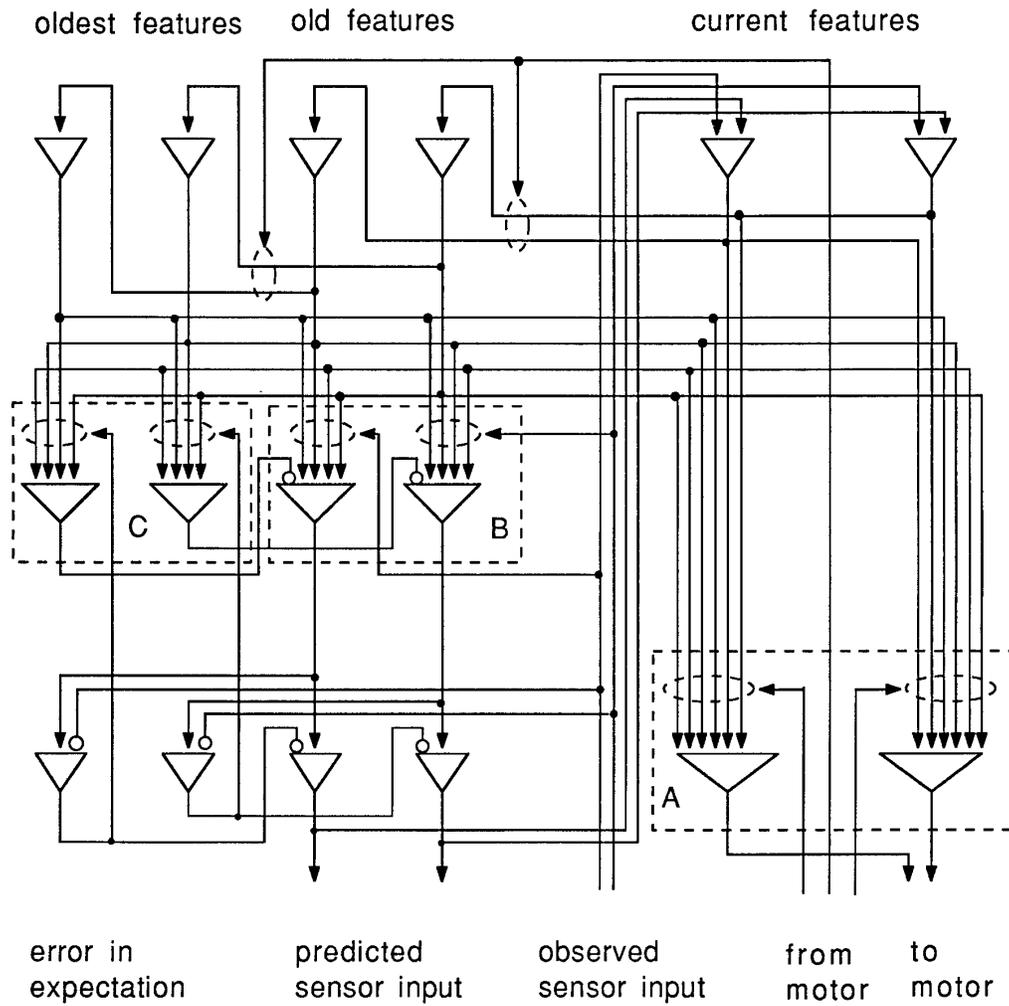


Figure 2

General model for classical conditioning of perceptual-motor sequences.

Connections:

—○ inhibitory —→ excitatory —(---)→ modifiable with UCS

the stored experiences of many perceptual-motor sequences. These predictions are continuously being tested against reality and are subject to modification if differences are found. The modifications are incorporated into long-term memory after feature extraction. The output patterns, justified with reality, represent the success or failure of prediction, or the validity of the stored knowledge. Feature extraction of the output from accurate predictions will reinforce the features extracted from the current input and accelerate the search for the next pattern in the sequence. Again, it is the search itself that is of potential benefit for the system. The reception of predicted patterns simply facilitates a behavioral sequence.

The shift registers that maintain a history of the feature space in the present model were inspired by the short range arcuate fibers that communicate bidirectionally between slabs in the cerebral cortex. Propagation delays of activity communicated out and back over these fibers can provide for the effects of near-term history. The possibility that motor commands influence the shift of activity across these registers was suggested by the distribution of pulvinar projections to cortical layers I and II [4].

Applications

The present model has been applied to machine vision, providing a method for the generation of scan paths that uniquely sample dominant features of simple images. The saccadic eye (or camera) movement in a scan path is a motor action that changes the input pattern predictably. The new input pattern is compared with a learned expectation of the pattern to be encountered, and a good match generates further scanning eye movements. If changes in visual fixation would be supplemented with dexterous manipulation of objects in the visual field, then the perceptual-motor mechanisms that we have described here could yield useful work.

The auditory system has many of the same performance and processing requirements as vision. The primary similarity is time dependency. A sequence of sounds defines a song, a rhythm, a statement or a query. To implement the present model of classical conditioning in an artificial auditory system, one needs to define the reflexes that initially order learning of features from the input space. A candidate reflex is vocal mimicry. Examples are the choruses of crickets, frogs, dogs, and men. The vocal apparatus is released by sound, that of itself and of its conspecifics. As rhythm is the product of sound motion, rhythm and rhythmic tracking or entrainment can enable the conditioned reflex. Some features in the auditory domain are pitch and harmony. The shaping and sequencing of that feature space by mechanisms described here could result in the perception and reproduction of language and music. We are currently exploring these hypotheses.

The maintenance of a near-term history of the input space that is used in the present algorithm has some computational advantages over a system that responds only to the current input. First it greatly increases the number of discriminable patterns that are possible with a given size of input vector. For example, if there are n binary input

elements then there will be $(2^{**n})-1$ unique patterns that can influence an output. If however, there are m temporal samples of the same input space, each of which has an influence on the current output, then there will be the potential to discriminate $((2^{**n}) - 1) ** m$ unique patterns. Biological neural systems appear to take advantage of this temporal leveraging of the input. In vision, the number of input elements is large while the short-term (iconic) memory is short, whereas in audition, where the input vector is relatively small, the short-term (echoic) memory is relatively long. For machine vision with a sensor field of $1024*1024$ and a temporal memory of 3 frames, the number of unique representations is approximately $1.5 * 10^{**18}$. A machine auditory system with only 32 frequency bands would need a temporal memory of only 12 frames to achieve a similar discriminability of the input space. Second, the increase in the numbers of discriminable patterns does not entail an exponential increase in the number of connections because the elements are not completely interconnected. Rather, influences converge from several delimited fields upon output elements and the associations between temporal fields are topographical and quite sparse.

Conclusion

Perceptual-motor sequences are the venue of adaptive behavior. These sequences can be organized and maintained by mechanisms of classical conditioning. The mechanisms and their supporting architectures are general in the biological nervous system, participating in all sensor modalities and effector subsystems. We propose that they should provide similar utility to robotics systems when applied.

References

- [1] Blackburn, M.R. and Nguyen, H.G. "Modeling the biological mechanisms of vision: Scan paths." In X. Avula (ed.) Mathematical and Computer Modeling in Science and Technology. New York: Pergamon Press, in press.
- [2] Varela, F.J. Principles of Biological Autonomy. New York: North Holland, 1979.
- [3] Pavlov, I. Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex. London: Oxford University Press, 1927.
- [4] Mansfield, R.J.W. "The role of the striate cortex in pattern perception in primates." In D.J. Ingle, M.A. Goodale and R.J.W. Mansfield (eds.) Analysis of Visual Behavior. Cambridge, MA: MIT Press, 1978.

Supported by the Independent Research (IR) program at the Naval Ocean Systems Center. Opinions expressed in this paper are not to be considered as official or reflecting the views of the Department of the Navy.