

University of East London Institutional Repository: <http://roar.uel.ac.uk>

This paper is made available online in accordance with publisher policies. Please scroll down to view the document itself. Please refer to the repository record for this item and our policy information available from the repository home page for further information.

Author(s): Palmer-Brown, Dominic; Lee, Sin Wee.

Article title: Continuous Reinforced Snap-Drift Learning in a Neural Architecture for Proxylet Selection in Active Computer Networks

Year of publication: 2005

Citation: Palmer-Brown, D.; Lee, S.W. (2005). "Continuous Reinforced Snap-Drift Learning in a Neural Architecture for Proxylet Selection in Active Computer Networks." International Journal on Simulation: Systems, Science and Technology, 6 (9) 11-21.

Link to published version:

<http://ducati.doc.ntu.ac.uk/uksim/journal/Vol-6/No.9/Paper2.pdf>

CONTINUOUS REINFORCED SNAP-DRIFT LEARNING IN A NEURAL ARCHITECTURE FOR PROXYLET SELECTION IN ACTIVE COMPUTER NETWORKS

DOMINIC PALMER-BROWN and SIN WEE LEE

School of Computing and Technology,
University of East London,
Longbridge Road,
RM8 2AS, UK.
<http://www.uel.ac.uk/scot>
d.palmer-brown@uel.ac.uk

Abstract: A new continuous learning method is used to optimise the selection of services in response to user requests in an active computer network simulation environment. The learning is an enhanced version of the ‘snap-drift’ algorithm, which employs the complementary concepts of fast, minimalist (snap) learning and slower drift (towards the input patterns) learning, in a non-stationary environment where new patterns arrive continually. Snap is based on Adaptive Resonance Theory, and drift on Learning Vector Quantisation. The new algorithm swaps its learning style between these two self-organisational modes when declining performance is detected, but maintains the same learning mode during episodes of improved performance. Performance updates occur at the end of each epoch. Reinforcement is implemented by enabling learning on any given pattern with a probability that increases linearly with declining performance. This method, which is capable of rapid re-learning, is used in the design of a modular neural network system: Performance-guided Adaptive Resonance Theory (P-ART). Simulations involving a requirement to continuously adapt to make appropriate decisions within a BT active computer network environment, demonstrate the learning is stable, and able to discover alternative solutions in rapid response to new performance requirements or significant changes in the stream of input patterns.

Keyword: Computational Intelligence, Artificial Neural Networks, Category Learning, Reinforcement Learning.

INTRODUCTION

The target application is representative of many situations in which machine learning is required to continuously and rapidly seek to improve performance in accordance with evolving circumstances. We describe this type of learning situation as having three related characteristics: provisional learning, fast learning, and performance feedback.

Provisional Learning

The adaptive systems of interest are not required to solve an optimisation problem in the traditional sense; they are involved in a continuous search for good solutions (solutions that are fit for purpose according to the chosen criteria of the target application) in a hyperspace that may contain many plausible solutions. There may an objective function of some kind, which the system tries to ‘optimise’. The classic example is error minimisation, in which in general the data is imperfect, e.g. limited, sparse, missing, error-prone, and subject to change (non-stationary). Therefore, the error minimum is really just a local minimum: local to a subset of data and an episode of time. It is only a global minimum in practice if the entire data-set is knowable,

and this is never possible because the data is changing. Whilst this does not preclude the discovery of solutions that work for all data-time, it does mean that such generalisation involves extrapolations and assumptions that cannot be justified on the sole basis of the available information. In such circumstances, it is reasonable, when a new candidate solution is found, for it to be held – as a provisional hypothesis – until or unless it is rejected, or until it can be replaced by a stronger hypothesis. This provisional learning approach is suited to contexts such as the BT active computer network scenario in which we require continuous learning that is sufficiently responsive to the speed of changes in a non stationary environment.

Fast Learning

Slow, iterative and intensive sampling based methods (eg. Gradient descent methods, and Bayesian methods [Barber and Bishop 1998] involving Monte Carlo and related methods) are *inherently* non-real-time, in the sense that they require multiple presentations of sets of patterns or samples, and therefore they cannot respond to the changing environment *as it is* changing. This contrasts sharply with the human case. Humans learn

'as they go along', to a significant extent, without the need for multiple presentations of each exemplar or pattern of information. In the BT application we also

require the learning to adapt to patterns it may only encounter once.

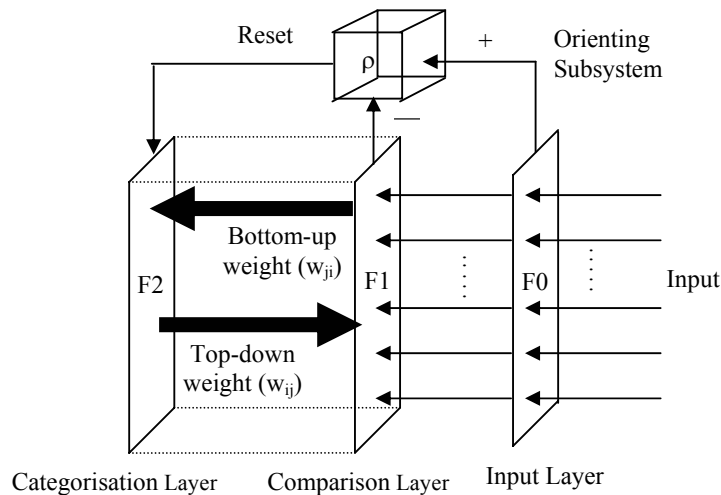


Figure 1: Architecture of an ART1 network.

Performance-guided Learning

An important concern in artificial intelligence is how to combine top-down and bottom-up information. This applies to learning systems. For example, reinforcement learning is an effective top-down approach to rewarding successful strategies, or moves, during learning; supervised learning is a powerful means of modifying an ANN when it makes mistakes; and genetic algorithms are effective at selecting for improvement across generations of solutions. These are important and effective approaches, not to be dismissed simply because they are not fast, or because they are computationally intensive. Fascinating results and innovations are still occurring with these approaches [Vieira et al 2003, Andrews 2003, Lee et al 2003]; and although unsupervised learning, which does not harness top-down information, is an extremely useful tool, for example as an alternative or complement to clustering, in its purest form it does not (by definition) make use of information on the current performance of learning in order to guide adaptation in appropriate directions. Ideally, learning should be rapid, and yet capable of taking external indicators of performance into account; and it should be capable of reconciling the data (bottom-up) with feedback concerning how the ANN is organising the data (top-down). We seek to enable fast bottom-up learning to freely discover solutions that are then filtered by a top-down quality assurance process.

The Adaptive Resonance Theory (ART) Network

Because Adaptive Resonance and Learning Vector Quantization have complementary strengths as

unsupervised learning methods, in our approach they are combined. Developments [Carpenter and Grossberg 1987a] of the original ART [Grossberg, 1976a; Grossberg 1976b] networks include ART1 that self-organises recognition categories for arbitrary sequences of binary input sequences; and ART2 which does the same for either binary or analogue inputs [Carpenter and Grossberg 1987b]. Subsequently, ART3 [Carpenter and Grossberg 1990] has been used to implement parallel searches of compressed or distributed recognition codes (output categories) in a neural network hierarchy. Following the successful implementation of the theory in real-time applications, further development has seen the creation of ART2-A [Carpenter et al. 1991a], which is 2 or 3 orders of magnitude faster than ART2. Fuzzy ART [Carpenter et al. 1991b] the fuzzy extension of ART, incorporated computations from fuzzy set theory. Extensions to ART networks to allow supervised learning were also introduced [Palmer-Brown 1992]; and ARTMAP [Carpenter et al. 1991c] and Fuzzy ARTMAP [Carpenter et al. 1992] autonomously learn to classify based on predictive success. Furthermore, there are several other versions of ART network [Tan 1997; Carpenter et al. 1998; Bartfai and White 2000], including supervised multi-layer, self-growing systems [Palmer-Brown 1992].

The ART1 Architecture and Learning Principles

ART1 networks are capable of fast and stable learning by categorising arbitrary binary input patterns using the basic principles of self-organization. The ART1

network (similar to dP-ART architecture in Figure 1) consists of 3 layers: the input layer (F0), the comparison layer (F1) and the categorisation layer (F2) with N , N , and M number of nodes respectively. Each node of the input layer is connected via non-modifiable links to its corresponding node in the comparison layer (there is a one to one mapping between F0 and F1 nodes). The F1 and F2 layers are interconnected using bottom-up and top-down modifiable weighted links that are adapted during the learning stage.

The learning process of the network can be described as follows: Upon the presentation of a binary input pattern I ($I_j \in \{1,0\}$, $j = 1, 2, 3, \dots, N$), the network attempts to categorize it by comparing it against the stored knowledge of the existing categories of each F2 node. This is achieved by calculating the bottom-up activation, which can be expressed as

$$T_i = \frac{|w_i \cap I|}{\beta + |w_i|} \quad (1)$$

where β is the constant that enables the larger magnitude prototype (weights encoding a minimal set of features that describe all input or ‘member’ patterns allocated to that category node) vector to be selected when there exist multiple prototype vectors that are subset of the binary input pattern.

The F2 node with the highest bottom-up activation, i.e. $T_i = \max \{T_i \mid I = 1, 2, \dots, M\}$ is then selected. If a category is found with the required matching level, known as the *vigilance level* and represented by the vigilance parameter ρ where $0 < \rho < 1$ and expressed by (2), then F2 node J will enter into a resonant state whereby it learns by modifying its prototype (weights encoding a minimal set of features that describe all input or ‘member’ patterns allocated to that category node), to retain only the critical features for the selected output category. This adjustment process is expressed by (3):

$$\frac{|w_i \cap I|}{|I|} \geq \rho \quad (2)$$

$$w_{ij}^{(new)} = \eta (w_{ij}^{(old)} \cap I) + (1 - \eta) w_{ij}^{(old)} \quad (3)$$

where η is the learning rate ($0 < \eta < 1$). All other weights in the network remain unchanged.

If no existing matching prototype is found, i.e. when the stored w_j do not match the input sufficiently, then the winning F2 node is reset and another F2 node with the highest activations is selected based on the similarity between its prototype and the current input, and so on. When no corresponding output category (F2 node) can be found, the network considers the input pattern as novel, and generates a new output category that learns the current input. Essentially, it is a fast adaptive form of competition-based learning [Carpenter & Grossberg 1988].

Limitations

There are limitations of ART networks in non-stationary environments where self-organisation needs to take account of periodic or occasional performance feedback:

- The ART network tends to organize itself into a stable state during fast learning whereby the weights stop changing, even in the presence of any new inputs.
- There is no external feedback to improve the performance of the network, even when it stabilises with poor performance.

PERFORMANCE-GUIDED ART (P-ART)

P-ART Architecture

The P-ART network proposed is a modular, multi-layered architecture as shown in Figure 2. It is composed of 3 modules, a Distributed P-ART (dP-ART) network, a Selection P-ART (sP-ART) network and a Kohonen Self-Organising Map. The $F1_1 \leftrightarrow F2_1$ connections of the dP-ART network and $F1_2 \leftrightarrow F2_2$ of the sP-ART are interconnected through weighted bottom-up and top-down connections that can be modified during the learning stage. For clarity, only the connections from the F1 layer to the active (winning) F2 node in each P-ART module are shown. The $F0_1 \rightarrow F1_1$ and two P-ART modules connected through $F2_1 \rightarrow F1_2$ are unidirectional, one to one and non-modifiable. Each of the $F2_2$ nodes is hard-wired onto a specific pre-trained region of the Kohonen Feature map where similar available proxylets (the target outputs) are spatially organised on the 2-D map according to their featural similarity.

Overview of the Operation of the System

On presentation of an input pattern at the input layer $F0_1$, the dP-ART will learn to group the input patterns according to their general features using the novel

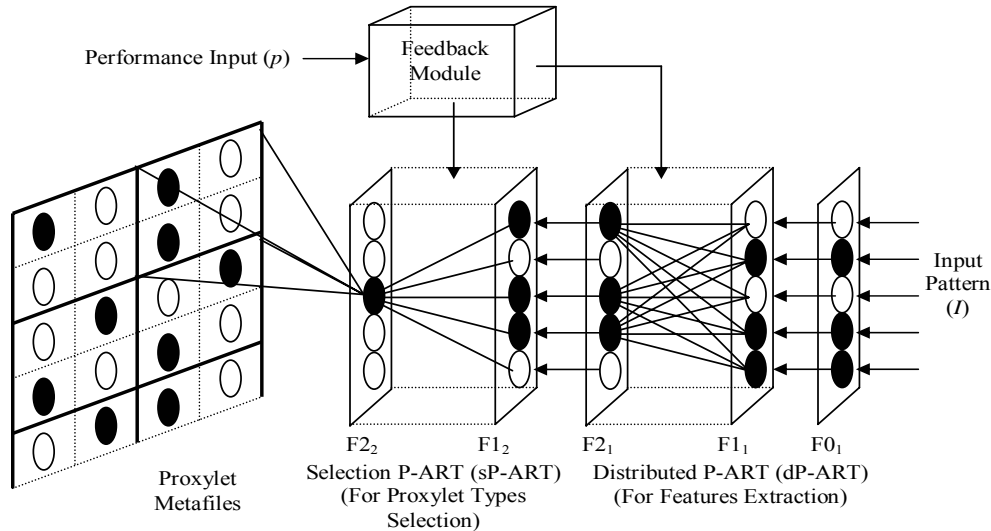


Figure 2: Architecture of the P-ART Network

learning principles developed in this work from the snap-drift' algorithm recently developed [Lee et al. 2002; Lee et al. 2003; Lee et al. 2004]. The latest version has several improvements over the previous one in terms of the normalization process, and the synchronization of learning between the s and d P-ARTs; but the two key differences are performance guided toggling of learning between snap and drift, and the introduction of a probabilistic aspect to enhance reinforcement and stability. The standard matching and reset mechanism of ART [Carpenter and Grossberg 1987a] is retained: If no existing matching prototype is found, i.e. when the stored pattern prototypes are not a good match for the input, the winning $F2_1$ node is reset and another $F2_1$ node is selected. When no corresponding output category can be found, the network considers the input as novel, and generates a new output category node that learns the current input pattern.

The three winning $F2_1$ nodes, whose prototypes are best match to the current input pattern, are used as the input data to the P-ART module for selecting an appropriate output type (called a proxylet in the target application). For the purpose of selecting the required proxylet, the proxylet type information indicated by the P-ART references pre-trained locations on the Kohonen Self-Organising Map (SOM) [Kohonen 1982; Kohonen 1990a], which represent specific proxylets. If the proxylet is unavailable, one of its neighbours is selected (the most similar alternative available).

A non-specific performance measure is used because, as in many applications, there are no specific performance measures (or external feedback) in response to each individual output decision. This measure is used to

encourage or discourage reselection of outputs (proxylet types) to occur in order to improve the performance of the neural system. The continuous learning method is the snap-drift algorithm. It involves toggling between snap and drift modes depending on performance changes. Snap and drift are alternative forms of adaptation, and they are described in the next section, **THE LEARNING**. Table 1 shows a summary of the steps that occurs in P-ART.

Table 1: Pseudo code of the Snap-Drift Algorithm for d-PART.

```

Step 1: Initialise parameters: ( $\alpha = 1, \sigma = 0$ )
Step 2: For each epoch, t
    Measure or calculate performance in the range
    {0,1} over the last epoch,  $P(t)$ .
    Performance improvement,  $PI = P(t) - P(t-1)$ 
    Set probability of learning,  $PL = 1 - P(t)$ 
Step 3: For each new input pattern
    Find the D winning nodes with the largest
    input (or create new nodes for mismatches)
    Set learn (adapt) true with probability PL.
    If learn is true test learning strategy condition:
    IF ( $PI \leq 0$ ) THEN
        Weights of d-PART adapted according to the
        alternate learning procedure: ( $\alpha, \sigma$ )
        becomes Inverse ( $\alpha$  and  $\sigma$ ) in equations (4
        and 10) below
    ELSE
        Weights of d-PART adapted according to the
        same procedure as in the last epoch: ( $\alpha, \sigma$ )
        unchanged.
Step 4: Process the output pattern of  $F2_1$  as input pattern
of  $F1_2$ 
    Find winning node (just one) in  $F2_2$ 
    
```

Weights of s-PART are adapted according to the same learning probability and strategy conditions as above, except that the in the first half of the learning epoch, both dP-ART and sP-ART learn, whereas in the second half of the epoch, only sP-ART learns. This allows relearning of the mapping from features to selections without the moving target problem of those features changing simultaneously.

THE LEARNING

Snap-Drift

In an environment where new patterns are introduced over time, the learning utilises a novel snap-drift algorithm based on fast, convergent, minimalist learning (snap) and cautious learning (drift) when the performance is good. Snap is based on a modified form of ART; and drift is based on Learning Vector Quantization (LVQ) [Kohonen 1990b]. The two forms are combined within a semi-supervised learning system that shifts its learning style whenever it receives a drop in the performance feedback. So, in general terms, the snap-drift algorithm can be stated as:

$$w = \alpha(\text{Fast_Learning_ART}) + \sigma(\text{LVQ}) \quad (4)$$

where α and σ are determined by performance feedback. In previous simulations, α and σ were real values [Lee et al. 2002; Lee et al. 2003; Lee et al. 2004]. In this paper, α and σ are set to (0, 1) or (1, 0) depending on changes in performance, and the learning is then enabled probabilistically.

Probability of Learning

Human and animal learning research provides inspiration for this aspect of the learning algorithm. Under certain human and animal experimental circumstances, learned decisions have about the same probabilities of occurrence as the chances of successful outcomes resulting from those decisions. This effect is called probability matching [Anderson 2000; Sutherland & Mackintosh 1971] where the probability of learning is equal to the probability that recent learning has resulted in good decisions. The interpretation of probability matching deployed here, is that learning via adaptation on a winning neuron is enabled with a probability inversely proportional to the current overall performance of the neural network. This probability function can be stated as:

$$P(\text{learning}) = P(1 - p) \quad (5)$$

Therefore, during the learning phase, if the performance of the system is good, e.g. > 0.8 , then the probability of the system learning (adapting in response to) the current pattern is < 0.2 . Conversely, poor performance guarantees that learning will proceed, thus encouraging

changes to be made to the weights.

Input Encoding

A form of coarse coding [Eurich et al. 1997] is used to represent proportional differences between numeric data encoded within the input patterns, e.g. the representation of the value 15 must be closer in input space to the representation of value 20 than that of, say 30. The input pattern is arranged in a 25 bit vector. Each property, such as bandwidth, time, file size, loss and completion guarantee, occupies 5 bits of the overall pattern. Table 2 shows the realistic range for each of the request properties. The coding of the user request is performed as illustrated in Table 3, across a different range for the 5 bits in the case of each property. The input patterns are generated by maintaining the coding of each field in turn and randomly generating the codes for rest of the fields for every 20 patterns, giving 1000 patterns in all.

Table 2: Value Ranges of User Request Properties

Properties	Ranges
Bandwidth	10Kb/s \rightarrow 2000Kb/s
Time	1ms \rightarrow 1000ms
Loss	20% \rightarrow 60%
Cost	0.1p \rightarrow 100p
Completion Guarantee	40% \rightarrow 100%

Table 3: Example Coding of Bandwidth in User Requests

Ranges (Kb/s)	User Request
200 \rightarrow 400	10000
800 \rightarrow 1000	01100
1800 \rightarrow 2000	00001

Weights Initialisation

The weights are calculated as floating point and are initialised at the beginning of the simulations. Top-down weights are set randomly to either 0 or 1.

$$w_{ji}(0) = [0,1] \quad (6)$$

Thus, a simple distributed affect will be generated at the output layer of the network, with different patterns tending to give rise to different activations across F_2 from the start. The bottom-up weights w_{ij} are assigned initial values corresponding to the initial values of the top-down weights w_{ji} . This is accomplished by equation (7):

$$w_{ij}(0) = \frac{w_{ji}(0)}{|w_{ji}(0)|} \quad (7)$$

The Distributed P-ART (dP-ART) Learning

On presentation of input pattern, the bottom-up activation is calculated using (8). Then the D number of F2₁ nodes with the highest bottom-up activation, using (9), are selected.

$$T_j = \sum |w_{ij} \cap I| \tag{8}$$

$$T_j = \max\{ T_j \mid J = 1,2,\dots, M \} \tag{9}$$

D is set to 3 in this application. If the distributed output categories are found with the required matching level, the three F2₁ nodes will enter into resonant state and learn using (10):

$$w_{ji}^{(new)} = \alpha(I \cap w_{ji}^{(old)}) + \sigma(w_{ji}^{(old)} + \beta(I - w_{ji}^{(old)})) \tag{10}$$

where w_{ji} = top-down weights vectors; I = binary input vectors, and β = the drift speed constant = 0.5. When $\alpha = 1$, (10) can be simplified to:

$$w_{ji}^{(new)} = (I \cap w_{ji}^{(old)}) \tag{11}$$

This invokes fast minimalist learning, causing the top-down weights to reach their new asymptote on each input presentation:

$$w_j \rightarrow I \cap w_j^{(old)} \tag{12}$$

In contrast, when $\sigma = 1$, (10) simplifies to

$$w_{ji}^{(new)} = (w_{ji}^{(old)} + \beta(I - w_{ji}^{(old)})) \tag{13}$$

This causes a simple form of clustering or LVQ at a speed determined by β . As describe in the pseudo code show in Table 1, learning is a combination of the two forms of adaptation, because the mode is toggled between snap and drift whenever performance has deteriorated during the previous epoch. In addition, whether adaptation occurs or not on a given pattern is a probabilistic decision, whereby the probability of the snap or drift occurring is proportional to declining performance. The novel bottom-up learning of the P-ART is a normalised version of the top-down learning:

$$w_{ij}^{(new)} = \frac{w_{ji}^{(new)}}{|w_{ji}^{(new)}|} \tag{14}$$

where $w_{ji}^{(new)}$ = top-down weights of the network after learning. Poor performance can occur when the final selection of proxylet type is wrong, *even if* the general features extracted by dP-ART are valid. To cope with this, there should be a dissociation between d and s-PART learning. Hence, dP-ART learning is toggled on-

off every half-epoch so that sP-ART can readjust its learning of selections without modification of the general features in dP-ART, thus resolving a moving target problem.

The Selection P-ART (sP-ART) Learning

The outputs produced by the dP-ART act as input to the sP-ART. The behaviour of sP-ART is the same as that described in section **P-ART Architecture**, with one exception; only the F2 node with the highest activation is adapted. Each output node of the sP-ART points to a set of available application-specific groupings (in this case proxylet types). The proxylet type data, containing attributes of the types, is used as off-line training data for the SOM so that it forms a map with similar proxylets placed on adjacent nodes. This allows each output node of the sP-ART to be ‘hardwired’ onto regions of the SOM. The task of the sP-ART is therefore to learn to associate the correct group of input patterns with an output node that is hardwired to a region of the SOM. The effect of learning and relearning within the sP-ART module is that specific output nodes will relate different groups of input patterns to different regions of the SOM until the performance feedback indicates that it is indexing the SOM regions that select the most appropriate proxylets. In that event, the learning probability is low, so that even if the snap-drift has not yet converged, further adjustment is slow.

The Performance Feedback

The external performance feedback into the PART reflects the performance requirement in different circumstances. Various performance feedbacks profiles in the range {0, 1} are fed into the network to evaluate the dynamic stability and effectiveness of the learning. Initially, some very basic tests with performances of 1 or 0 were evaluated in a simplified system [Lee et al. 2002; Lee et al. 2003; Lee et al. 2004]. Here, the simulations involve computing the performance based on a parameter associated with the winning output neuron. The winning output neurons represent proxlet selections which are either good or poor selections for the current input request, and hence the accumulation of these over time, averaged, gives a performance indicator in the range {0,1}.

BRITISH TELECOM (BT) APPLICATION

Application Layer Active Network (ALAN)

British Telecom (BT) is the main data network provider in the UK. At present, most applications are run on edge devices (which send and receive data, but do not route third party data), such as servers, PCs and WAP enabled devices. There are strong arguments for moving as many of these applications as possible into the network

[Tennenhouse and Wetherall 1996], thereby ensuring optimal placement of applications with respect to performance, version synchronicity (so that more users have the same version), and increased security. 'Active Networking' [Tennenhouse and Wetherall 1996] aims to achieve this application migration into the network by running code within the network on specialised routers. It gives users the ability to load software components onto network devices dynamically without explicit references to any third party. There are three types of active networks:

1. *Capsule* – This approach is used to enable the active service codes to run on the network devices, such as servers, that the packets encounter.
2. *Programmable* – This approach allows the clients to download their own active service codes onto network devices before using their application [Campbell et al. 1999].
3. *Active Bridging* – This approach allows the network operators¹ the freedom to choose appropriate active service codes of their own [Alexander 1996]. Despite the difficulties of the approach (Marshall 1999a), it has highlighted the important requirements for a feasible active network and encouraged other researchers to develop better alternatives to resolve them.

The ALAN architecture [Fry and Ghosh 1999] enabled the user to supply JAVA based active-service codes known as proxylets that run on a network device. Each networked server runs the 'Execution Environments for Proxylets' (EEPs) that contains the user supplied software. The purpose of the architecture is to locate the software at optimal points of the end-to-end path between the server and the clients.

Automated Active Network Management using Distributed Genetic Algorithm (GA)

The original ALAN proposal, the management system supports conventional management agent interfaces [Marshall 1999b; Marshall et al. 2000] that respond to instructions from the system operators. Each application is individually placed in the network. However, since ALAN with the potential for an enormous range of services, it is necessary to combine the active services with an automated and adaptive management solution. Recently, a novel adaptive approach, a Distributed Genetic Algorithm (GA) solution was introduced by BT Research Laboratories [Marshall and Roadknight 2001]. It performs proxylet placement. Here, P-ART provides a means of finding a set of conditions that produce optimum proxylet selection in an EEP containing the frequently requested proxylets that have been placed. Continuous performance guided adaptation of the mapping of input patters, which contain the main

attribute values of user proxylet requests, performs intelligent proxylet type selection.

THE P-ART SIMULATION

P-ART is used for learning and mapping user requests onto appropriate proxylets. The test patterns consist of 1000 input vectors. Each test pattern characterizes the properties of a network request, such as bandwidth, time, file size, loss and completion guarantee. These test patterns are presented in random order with 10 patterns per epoch for 100 epochs where the performance, p , is calculated according to the average bandwidth of selections. This on-line continuous random presentation of test patterns simulates the real world scenario in which pattern presentation order is random, so that a given network request might be repeatedly encountered while others are not used at all.

ANALYSIS OF RESULTS

Results are presented in Figures 3, 4 and Table 4 and 5. They are representative of many simulations that have been run. Performance feedback is updated at the end of each epoch of 10 patterns. Much longer epochs are less effective. The best results are for the shortest epochs for which the performance estimate remains a reasonable estimate of overall performance, which of course it would not be for a very small number of patterns. In this application, although there are 1000 patterns, there are only 100 general types, and hence 10 is approximately the smallest reasonable sample for which updates in performance may be trusted to increase or decrease with true overall performance. This will clearly be different for each application.

Figure 3 represents a typical run in which performance converges towards 100% and settles with some remaining jitter. In the other results tables below, learning is actually over by about epoch 75, after which no new selections of proxylets occur until the criteria change to low bandwidth. After 75 (750 pattern presentations), all the performance variation between epochs (the jitter in the performance curve) is due to the epochs being short (in other words, samples of 10 give approximately 70% accurate true performance estimates), and hence the performance over 1000 patterns is actually constant at about the average of the values of the table values from 70-100, which is just under 80%.

In Figure 4 and Table 5, the performance criterion is swapped from high to low bandwidth after 100 epochs, and we see relearning and re-stabilisation occur, with similar performance to before the swap being restored once convergence has re-occurred.

¹ Network operators are those companies (such as BT, NTL) who control the public network and allow Internet Service Providers (ISPs) (such as Freeserve, BT, AOL) to provide additional Internet services.

CONCLUSIONS & FURTHER WORK

In conclusion, the snap-drift algorithm learning stabilizes reliably and is able to map the inputs onto appropriate proxylets. The simulation results show that the snap-drift algorithm is able to provide continuous probabilistic real-time learning in order to improve the performance, based on the external performance feedback. On-going research is advancing on three fronts, with unsupervised and reinforcement versions of snap-drift:

- Phonetic Feature Discovery in Speech waveforms:** An fully unsupervised snap-drift algorithm is being used for discovering the internal features in speech utterances and to characterise the acoustic properties of the normal and stammering speaker groups. The learnt features will be studied, in particular the features occurring in utterances of both normal and stammering speakers, and the distinct features which only exist one of the two groups, as a means of correlative data analysis.
- Cluster Analysis using Iris Data:** The unsupervised snap-drift algorithm is also being used for clustering the well known Iris dataset [Fisher 1936] and has been widely used in cluster analysis. This provides the opportunity to test the snap-drift algorithm on a well-defined, well understood problem and thereby enable a comparative evaluation with other types of clustering algorithms. It will then allow an investigation of the affect on snap-drift clustering of the reinforcement method used in this paper.
- Phrase Recognition for a Connectionist Language Parser:** The emphasis of this research is to use the snap-drift algorithm for recognition of phrases extracted from the Lancaster Parsed Corpus (LCS) in order to improve the overall performance a corpus-based parser [Palmer-Brown et al. 2002; Tepper et al. 2002], which involved phrase segmentation and recognition stages.

In terms of the snap-drift algorithm itself, we are exploring the role of individual reinforcement feedback onto each neuron to control its learning mode, as compared to the current method, which uses the same performance measure for every neuron. This will enable a snap-drift network to function as a classifier.

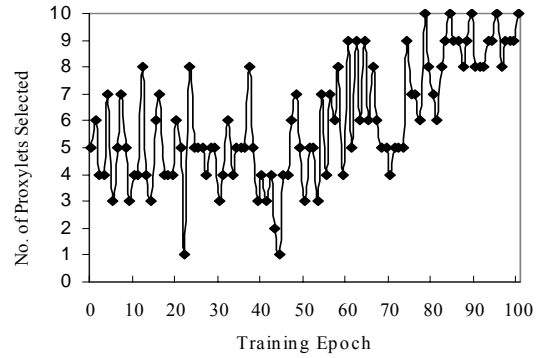


Figure 3: The Selection Frequency of the Proxylet Type. E.g. Bandwidth Bands: Low Bandwidth Proxylet: 0 → 1000 Kb/S and High Bandwidth Proxylet Type: 1001 → 2000 Kb/S

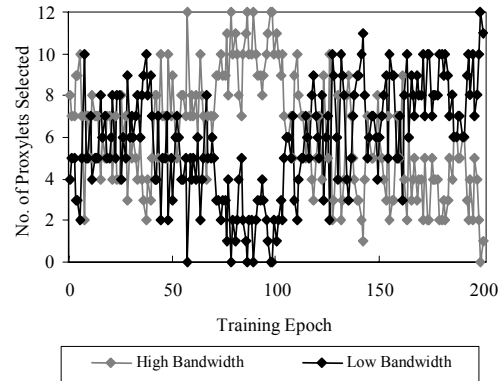


Figure 4: The Selection Frequency of Proxylet Type.

Table 4: Performance of P-ART

Epoch	Average No. of High Bandwidth Proxylet Selected (/10)	Performance (%)
1 – 10	4.9	49
11 – 20	4.8	48
21 – 30	4.9	49
31 – 40	4.8	48
41 – 50	4.0	40
51 – 60	5.2	52
61 – 70	6.8	68
71 – 80	6.6	66
81 – 90	8.5	85
91 – 100	8.7	87

Table 5: Performance of P-ART with Switching of Performance Criteria @ epoch 100.

Epoch	Average No. of High B/W Proxylet Selected (/12)	Average No. of Low B/W Proxylet Selected (/12)	High B/W Proxylet Selection (%)	Low B/W Proxylet Selection (%)
1 – 10	6.08	3.92	50.69	32.64
11 – 20	5.25	4.75	43.75	39.58
31 – 40	3.50	6.50	29.17	54.17
41 – 50	6.00	4.00	50.00	33.33
51 – 60	6.33	3.67	52.78	30.56
61 – 70	5.83	4.17	48.61	34.72
71 – 80	7.83	2.17	65.28	18.06
81 – 90	8.42	1.58	70.14	13.19
91 – 100	8.33	1.67	69.44	13.89
101 – 100	6.67	3.33	55.56	27.78
111 – 120	5.17	4.83	43.06	40.28
131 – 140	4.42	5.58	36.81	46.53
151 – 160	3.75	6.25	31.25	52.08
161 – 170	2.83	7.17	23.61	59.72
171 – 180	3.42	6.58	28.47	54.86
181 – 190	2.83	7.17	23.61	59.72
191 – 200	0.08	9.92	0.69	82.64

REFERENCES

- Alexander D. S. 1996. "Active Bridging" *Computer Communication Review*, Vol. 26, No. 2, 101 – 111.
- Anderson, J.R. 2000. "Learning & Memory". 2nd, NY, Wiley.
- Andrews, S. G. 2003. "Novel neural network methods for describing attributes contained within lesions images". In *Proceeding of the European Simulation Multi-conference*, 142 – 147, Nottingham, UK.
- Bartfai, G. and R. White. 2000. "Incremental Learning and Optimization of Hierarchical Clustering with ART-based Modular Networks." In *Innovations in ART Neural Networks*, L. C. Jain, B. Lazzarini and U. Halici (Eds.). Physica-Verlag, 87 – 132.
- Barber, D and Bishop C. 1998. "Ensemble Learning in Bayesian Neural Networks." In *Neural Networks and Machine Learning*, Bishop C (Ed.). Springer-Verlag, 215-237.
- Campbell, A. T. et al. 1999. "Open Signalling for ATM, Internet and Mobile network", *Computer Communication Review*, Vol. 29, No. 1.
- Carpenter, G. A. and S. Grossberg. 1987a. "A Massively Parallel Architecture for a Self-Organising Neural Pattern Recognition Machine," *Computer Vision, Graphics and Image Processing*, Vol. 37, 54-115.
- Carpenter, G. A. and S. Grossberg. 1987b. "ART2: Self-Organization of Stable Category Recognition Codes for Analogue Pattern." *Applied Optics*, Vol. 26, 4919 - 4930.
- Carpenter, G. A. and S. Grossberg. 1988. "The ART of Adaptive Pattern Recognition by a Self-Organising Neural Network." *IEEE Computerm* Vol. 21, No. 3, 77 - 88.
- Carpenter, G. A. and S. Grossberg. 1990. "ART 3: Hierarchical Search Using Chemical Transmitter in Self-Organizing Pattern Recognition Architectures." *Neural Network*, Vol. 3, No. 4, 129 – 152.
- Carpenter, G. A.; S. Grossberg and D.B. Rosen. 1991a. "ART 2-A: An Adaptive Resonance Algorithm for Rapid Category Learning and Recognition." *Neural Networks*, Vol. 4, 493 - 504.
- Carpenter, G. A.; S. Grossberg and D.B. Rosen. 1991b. "Fuzzy ART: Fast Stable Learning and Categorization of Analogue Pattern by an Adaptive Resonance System." *Neural Networks*, Vol. 4, 759 - 771.
- Carpenter, G. A.; S. Grossberg and J. H. Reynold. 1991c. "ARTMAP: Supervised Real-Time Learning and Classification of Nonstationary Data by a Self-Organizing Neural Networks." *Neural Networks*, Vol. 4, 565 – 588.
- Carpenter, G. A.; S. Grossberg; A. Markuzon; J. H. Reynold and D. B. Rosen. 1992. "Fuzzy ARTMAP: A Neural Network Architecture for Incremental Supervised Learning of Analogue Multidimensional Maps." *IEEE Transactions in Neural Networks*, Vol. 3, No. 5, 698 – 713.
- Carpenter, G. A.; B. Milenova and B. Noeske. 1998. "dARTMAP: A Neural Network for Fast Distributed Supervised Learning." *Neural Networks*, Vol. 11, 793 – 813.
- Eurich, C. W.; H. Schwegler and R. Woesler. 1997. "Coarse Coding: Applications to the Visual System of Salamanders." *Biological Cybernetics*, Vol. 77, 41 – 47.
- Fisher, R. A. 1936. "The use of multiple measurements in taxonomic problems." *Annual Eugenics*, Vol. 7, No. II, 179 – 188.
- Fry, M. and A. Ghosh. 1999. "Application Layer Active Network." *Computer Network*, Vol. 31, No. 7, 655 – 667.
- Grossberg, S. 1976a. "Adaptive Pattern Classification and Universal Recoding. I. Parallel Development and Coding of Neural Feature Detectors." *Biological Cybernetics*, Vol. 23, 121 – 134.
- Grossberg, S. 1976b. "Adaptive Pattern Classification and Universal Recoding. II. Feedback, Expectation,

Olfaction, and Illusions.” *Biological Cybernetics*, Vol. 23, 187 – 202.

Kohonen, T. 1982. “Self-Organized Formation of Topologically Correct Feature Maps.” *Biological Cybernetics*, Vol. 43, 53 – 69.

Kohonen, T. 1990a. “The Self-Organizing Maps.” In *Proceeding of IEEE*, Vol. 78, No. 9, 1464 – 1480.

Kohonen, T. 1990b. “Improved Versions of Learning Vector Quantization.” In *Proceeding of the International Joint Conference on Neural Networks*, Vol. 1, 545 – 550.

Lee, S. W.; D. Palmer-Brown; J. Tepper and C. M. Roadknight. 2002. “Performance-guided Neural Network for Rapidly Self-Organising Active Network Management.” In *Soft Computing Systems: Design, Management and Applications*, A. Abraham J. Ruizdel-Solar and M. Köppen (Eds.). IOS Press, Amsterdam, 21 – 31.

Lee, S. W.; D. Palmer-Brown; J. Tepper and C. M. Roadknight. 2003. “Snap-Drift: Real-Time Performance-guided Learning.” In *Proceeding of International Joint Conference on Neural Networks*, Vol. 2, 1412 – 1416. Portland, Oregon, U.S.A

Lee, S. W.; D. Palmer-Brown and C. M. Roadknight. 2004. “Performance-guided Neural Network for Rapidly Self-Organising Active Network Management.” *Neurocomputing*, Vol. 61C, 5 – 20.

Marshall, I. W. 1999a. “Active Network - Making the Next Generation Internet Flexible”, *British Telecom Engineering*, Vol. 18, 2-8.

Marshall, I. W. 1999b. “Application Layer Programmable Internetwork Environment.” *British Telecom Technology Journal*, Vol. 17. No. 2, 82 – 94.

Marshall, I. W.; J. Hardwicke; H. Gharid; M. Fisher and P. Mckee. 2000. “Active Management of Multi-Service Networks.” In *Proceeding of the IEEE Network Operations and Management Symposium* (Hawaii). 981 – 983.

Marshall, I. W. and C. M. Roadknight. 2001. “Provision of Quality of Service for Active Services.” *Computer Networks*, Vol. 36, No. 1, 75 – 85.

Palmer-Brown, D. 1992. “High Speed Learning in a Supervised, Self Growing Net.” In *Proceeding of the International Conference of Artificial Neural Network*, Vol.2, 1159 – 1162.

Palmer-Brown, D., Tepper, J. and Powell, H. 2002. “Connectionist Natural Language Parsing.” *Trends in Cognitive Sciences*, Vol. 6, No. 10, 437 - 442.

Sutherland, N.S. and Mackintosh, N.J. 1971. “Mechanisms of Animal Discrimination Learning”. NY, Academic Press.

Tan, A-H. 1997. “Cascade ARTMAP: Integrating Neural Computation and Symbolic Knowledge Processing.” *IEEE Transactions in Neural Networks*, Vol. 8, No. 2, 237 – 250.

Tennenhouse, D. and D. Wetherall. 1996. “Towards An Active Network Architecture.” *Computer Communication Reviews*, Vol. 26, No. 2, 5 – 18.

Tepper, J., H. Powell and D. Palmer-Brown. 2002. “A Corpus-based Connectionist Architecture for Large-scale Natural Language Parsing.” *Connection Science*, Vol. 14, No. 2, 93 – 114.

Vieira, C. , P Mather, P Alpin. 2003. “Improving Artificial Neural Network Performance by Using Temporal-Spectral Features for Agricultural Crop Classification”. In *Proceeding of the European Simulation Multi-conference*, 124 – 130. Nottingham, UK.

AUTHOR BIOGRAPHIES



DOMINIC PALMER-BROWN is professor of neurocomputing and Associate Head, School of Computing and Technology, University of East London, UK. In recent years he has maintained active research links with several organisations, including British Telecom Research Labs, The Centre for Ecology and Hydrology, and with several universities. A key focus of his research is neurocomputing and related methods of adaptation and learning in cognitive science, intelligent data analysis, and pattern recognition. Dominic has published about 60 international conference and journal papers and supervised 12 PhDs, since completing his own PhD on an adaptive resonance classifier in 1991. His interests have principally concerned supervised and performance-guided ART, enhanced MLPs for intelligent data analysis, and architectures incorporating MLPs and SRNs for thematic knowledge extraction and natural language processing. He was editor of the review journal *Trends in Cognitive Sciences* during 2000-2 before rejoining Leeds Metropolitan University. In 2005 he was appointed to his present post.



SIN WEE LEE was born in Melaka, Malaysia, in 1976. He graduated with first class honours in electronics and computing engineering from the Nottingham Trent University, United Kingdom, in 1999. His PhD's thesis focuses on the development of performance-guided neural network for active network management. From 2000 to 2001, he was a systems engineer at Malaysia Multimedia University in Malaysia. In December 2001, he joined the School of Computing, Leeds Metropolitan University, Leeds, United Kingdom, with a research scholarship from EPSRC/BT Research Laboratories in Neural Networks. As a research assistant, he works with Dominic Palmer-Brown, on the improvement and development of phrase recognition for a connectionist language parser, feature discovery in phonetics data and cluster analysis using iris data using snap-drift algorithm.