Towards Learning Classifier Systems for Continuous-Valued Online Environments

Christopher Stone and Larry Bull

Faculty of Computing, Engineering and Mathematical Sciences
University of the West of England
Bristol, BS16 1QY, United Kingdom
{christopher.stone,larry.bull@uwe.ac.uk}

Abstract. Previous work has studied the use of interval representations in XCS to allow its use in continuous-valued environments. Here we compare the speed of learning of continuous-valued versions of ZCS and XCS with a simple model of an online environment.

1 Introduction

Much current research is focussed on the asymptotic performance of Learning Classifier Systems (LCS). However, where the environment is ever changing, speed of learning and the ability of the system to recover from environmental change (i.e., the slope of the learning curve) may be a more relevant measure.

We are interested in continuous-valued environments and so use an interval representation to replace the $\{0, 1, \#\}$ classifier predicate with one representing an interval $[p_i, q_i)$ [1,5]. An interval is represented as an unordered tuple (p_i, q_i) and matches an environmental variable x_i if $p_i \leq x_i < q_i$.

Changes to the cover, subsumption and Genetic Algorithm (GA) operators must be made to ZCS [3] and XCS [4] to accommodate the interval representation. Action set subsumption is not used. We use single-point crossover between intervals for ZCS and two-point crossover between intervals for XCS. Both ZCS and XCS are run with an initially empty population [1,2].

2 Experiments

Experiments were performed on a six-bit real multiplexer [5]. Parameter settings used for ZCS were (using XCS terminology for consistency) $N = 800$, $\beta = 0.2$, $\tau = 0.1$, $\chi = 0.5$, $\mu = 0.04$, $f_I = 500$, $s_0 = 0.5$, $\rho = 0.25$, $\phi = 0.5$. Parameter settings for XCS were $N = 800$, $\beta = 0.2$, $\alpha = 0.1$, $\varepsilon_0 = 10$, $\nu = 5$, $\theta_{GA} = 12$, $\chi = 0.8$, $\mu = 0.04$, $\theta_{del} = 20$, $\delta = 0.1$, $\theta_{sub} = 20$, $p_I = 10$, $\varepsilon_I = 0$, $f_I = 0.01$, $\theta_{mna} = 2$, $m = 0.1$, $s_0 = 1$. All experiments used a fixed reward of 1000.

In an online environment there is no artificial distinction between exploration and exploitation trials, so we use a roulette wheel for action selection and permanently enable the GA and update mechanisms.

E. Cantú-Paz et al. (Eds.): GECCO 2003, LNCS 2724, pp. 1924–1925, 2003.

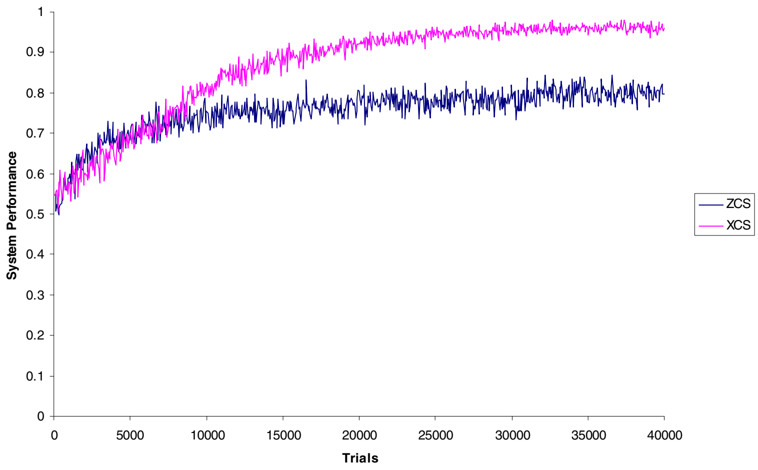

Fig. 1. Mean system performance over 10 runs of ZCS and XCS on the six-bit real multiplexer

Figure 1 shows the performance of ZCS and XCS on the real multiplexer. Similar results were obtained for a three-dimensional checkerboard problem with three divisions per dimension [2] (not shown).

We found that, although the asymptotic performance of XCS ultimately exceeds that of ZCS, ZCS performs as well as XCS during the early part of the learning curve for the problems tested. This result shows that a simple LCS architecture, such as ZCS, can compete with the more complex XCS system where speed of learning is more important than asymptotic performance.

References

1. Stone, C. & Bull, L. (2003) For real! XCS with continuous-valued inputs. To appear in *Evolutionary Computation*.
2. Stone, C. & Bull, L. (2003) Comparing Learning Classifier Systems for Continuous-Valued Online Environments. University of the West of England Learning Classifier Group Technical Report UWELCSG03-001.
`http://www.cems.uwe.ac.uk/lcsg/reports/uwelcsg03-001.ps`
3. Wilson, S.W. (1994). ZCS: A zeroth order classifier system. *Evolutionary Computation*, 2(1):1-18.
4. Wilson, S.W. (1995). Classifier fitness based on accuracy. *Evolutionary Computation*, 3(2):149–175.
5. Wilson, S W. (2000). Get real! XCS with continuous-valued inputs. In P.L. Lanzi, W. Stolzmann and S.W. Wilson (eds.), *Learning Classifier Systems. From Foundations to Applications, Lecture Notes in Artificial Intelligence (LNAI-1813)*, Berlin: Springer, pages 209–219.