

Introduction to the special issue on learning classifier systems

Larry Bull · Pier Luca Lanzi

Published online: 7 February 2009
© Springer Science+Business Media B.V. 2009

It is now 30 years since John Holland presented the first implementation of his learning classifier system (LCS) framework (Holland and Reitman 1978). This “Cognitive System Level 1” used a genetic algorithm (Holland 1975) to learn appropriate rules of behaviour in one-dimensional, dual-objective maze navigation tasks with a form of reinforcement learning assigning utility to the rules. Holland later revised the algorithm to define what would become the standard system (Holland 1980, 1986). However, Holland’s full system was somewhat complex and practical experience found it difficult to realize the envisaged behaviour/performance (e.g., Wilson and Goldberg 1989) and interest waned. Some years later, Wilson presented the “zeroth-level” classifier system, ZCS (Wilson 1994) which “keeps much of Holland’s original framework but simplifies it to increase understandability and performance” (ibid.). But ZCS did not reach optimality in the most common reinforcement learning sense. Accordingly, Wilson introduced a form of LCS which altered the way in which rule fitness is calculated—XCS (Wilson 1995). XCS also makes the connection between LCS and temporal difference learning (Watkins 1989) explicit with, in its standard form, its ability to represent the state-action value map in a rule form thereby enabling compaction through generalization.

Shortly after Holland had formulated the general framework, Stephen Smith (1980) presented a modification wherein a traditional genetic algorithm was used to design a complete set of rules. That is, Smith’s poker playing “Learning System 1” avoided the need to assign utility to individual rules.

The subsequent years have seen a resurgence of LCSs as XCS in particular has been found able to reach optimality in a number of difficult benchmark problems. Perhaps more importantly, XCS has also begun to be applied to a number of hard real-world problems such as data mining, simulation modeling, robotics, and adaptive control (see Bull 2004 for an overview)—where excellent performance has often been achieved. A theoretical basis

L. Bull (✉)
University of the West of England, Bristol, UK
e-mail: larry.bull@uwe.ac.uk

P. L. Lanzi
Politecnico di Milano, Milano, Italy

for LCS is also beginning to emerge so as to provide the foundations to a principled approach to classifier system design (see Bull and Kovas 2005 for an overview).

This special issue brings together work using both versions of Holland's general framework which clearly demonstrate this new-found potential of the approach. The interested reader is referred to the proceedings of the annual International Workshop on LCSs (e.g., Kovacs et al. 2007), in particular.

The guest editors wish to thank the Editor of Natural Computing for the opportunity of publishing this special issue on LCSs, and the referees who helped in the reviewing process.

References

- Bull L (ed) (2004) Applications of learning classifier systems. Springer, Berlin
- Bull L, Kovas T (eds) (2005) Foundations of learning classifier systems. Springer, Berlin
- Holland JH (1975) Adaptation in natural and artificial systems. University of Michigan Press, Ann Arbor, MI
- Holland JH (1980) Adaptive algorithms for discovering and using general patterns in growing knowledge bases. *Int J Policy Anal Inf Syst* 4(3):245–268
- Holland JH (1986) Escaping brittleness: the possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In: Michalski RS, Carbonell JG, Mitchell TM (eds) Machine learning, an artificial intelligence approach. Morgan Kaufmann, San Francisco
- Holland JH, Reitman JH (1978) Cognitive systems based in adaptive algorithms. In: Waterman DA, Hayes-Roth F (eds) Pattern-directed inference systems. Academic Press, NY
- Kovacs T, Llorà X, Takadama K, Lanzi P-L, Stolzmann W, Wilson SW (eds) (2007) Learning classifier systems: international workshops, IWLCS2003–2005, revised selected papers. Springer, Berlin
- Smith SF (1980) A learning system based on genetic adaptive algorithms. PhD Dissertation, Pittsburgh University
- Watkins CJ (1989) Learning from delayed rewards. PhD Dissertation, Cambridge University
- Wilson SW (1994) ZCS: a zeroth-level classifier system. *Evol Comput* 2(1):1–18
- Wilson SW (1995) Classifier fitness based on accuracy. *Evol Comput* 3(2):149–176
- Wilson SW, Goldberg DE (1989) A critical review of classifier systems. In: Proceedings of the 3rd international conference on genetic algorithms. Morgan Kauffman, San Francisco, pp 244–255