# Introduction to the Special Issue

This special issue is dedicated to *The First International Workshop on Learning Classifier Systems*, which was held at Johnson Space Center in Houston, Texas, October 6–9, 1992. The workshop was a self-organized meeting of researchers with long-term interests in the complex issues involved in learning classifier systems (LCSs) (Goldberg, 1989; Holland, Holyoak, Nisbett, & Thagard, 1986). Given the relative success of GAs in optimization, LCS issues have often been treated with less interest. The workshop was conceived as a meeting where interested parties could focus exclusively on LCS-related issues. The workshop featured invited talks and extended brainstorming sessions. At the workshop's conclusion, several of the attendees extended the work they presented, incorporating information from the workshop's discussions. This special issue contains some of the results from those efforts. Other papers from the workshop will appear in a later issue of *Evolutionary Computation*.

The workshop brainstorming sessions were fast-paced and stimulating, and they revealed a great deal about the LCS zeitgeist. Some of the most interesting conversations focused on defining the LCS itself. To paraphrase an assertion made by Lashon Booker at the workshop, the discussions indicated that the LCS was more of an *approach* than a *method* (Booker, 1992). That is, the LCS is a set of conceptual details, rather than a set of algorithmic details. Clarification of these conceptual details became a central focus of the workshop.

In workshop discussions, the following key LCS issues emerged:

**Cooperation:** Discussions indicated that the cooperative aspect of an LCS population was a defining characteristic of the LCS approach. When GAs are used in optimization, population members are functionally independent. These population members only interact in a competitive fashion, through the GA selective process. In an LCS, population members are interdependent. They must cooperate to improve overall system performance, while also competing in the selective process. Therefore, the demands placed on the GA in an LCS are quite different from those placed on a GA used in optimization. In an LCS, a GA must maintain a functional, diverse population of individuals that are *co-adapted* to one another. In addition, the GA must discover new individuals that will augment the existing population's performance.

**Representation:** Workshop participants indicated that there was a great deal of latitude in possible representations for the LCS. Surprisingly, the typical 1, 0, # representation was not felt to be essential to the LCS paradigm. The consensus was that representation could be determined for the problem at hand. However, choice of representation in the LCS remains a key issue because it affects cooperation, the GA discovery process, and overall system performance.

**Credit Assignment:** Typically, LCSs are associated with *reinforcement learning problems* (Barto, 1990). However, workshop discussions did not indicate that this was a defining feature of the LCS approach. The *bucket brigade algorithm* is commonly associated with temporal credit assignment and reinforcement learning in LCSs. Workshop discussions often referred to the growing body of literature on reinforcement learning. Workshop participants agreed that this body of work has a significant bearing on

LCSs. The bucket brigade was simply seen as one of a growing class of temporal credit assignment schemes. Although credit assignment emerged as a central issue for LCSs, no particular credit assignment scheme was seen as defining to the LCS approach.

**Internal Messages:** LCSs examined in most research to date have operated in *stimulus-response* mode. However, workshop discussions indicated that a defining (albeit little examined) aspect of LCSs is the ability to perform complex, memory-exploiting behavior through internal message passing.

The papers in this issue consider some of these key LCS concerns. In the first paper, Wilson examines a "zeroth-level" classifier system. This simplified system allows for a clear introduction to LCS and careful examination of several aspects of its behavior. Wilson considers the relationship between LCS credit assignment and *Q-learning*, a popular technique from the reinforcement learning literature. Through proposed extensions of this simplified LCS, Wilson also makes interesting suggestions about LCS representation, cooperation in the GA discovery process, and internal message processing.

In the second paper, R. Smith and Cribbs consider whether an LCS can be accurately described as a type of neural network (NN). The LCS/NN analogy presented leads to implementation of an LCS-like network. This system can also be viewed as an alternate LCS representation. Experimental results indicate that this type of system has potential as a new use of co-adaptive GAs for neural network training. The paper also suggests other ways in which neural networks and LCSs can influence and benefit one another.

In this issue's third paper, Horn, Goldberg, and Deb carefully consider the cooperative action of the GA under LCS credit assignment schemes. This study's analyses show how division of reward in the LCS is an implicit form of *fitness sharing*. This helps to formally explain how the LCS can maintain the diverse, cooperative population. By establishing this connection, this paper allows us to understand both LCS reward division and sharing within the same theoretical framework.

In the fourth and final paper in this special issue, Greene and S. Smith consider the COGIN (COverage-based Genetic INduction) system in light of the cooperation necessary in LCSs. COGIN is a practical system for constructing rule-based classification models. COGIN's utilization of cooperative population diversity gives the system performance advantages over other induction methods. The paper introduces two modifications to COGIN that improve the system's population diversity and its performance. COGIN also expands the definition of the LCS, through its application to a supervised learning problem.

Other LCS workshop-related papers are in progress for a later special issue of EC. These papers will further consider cooperation in LCSs, alternate LCS representations, credit assignment, and internal message processing. It is hoped that these papers will relay some of the excitement shared by the workshop attendees, and spur interest in the fascinating, complex, and promising evolutionary computation systems called LCSs.

Robert E. Smith
Department of Engineering Science
and Mechanics
The University of Alabama
Box 870278
Tuscaloosa, Alabama 35487
Phone: (205) 348-1618
Email: rob@comec4.mh.ua.edu

Manuel Valenzuela-Rendón
ITESM, Campus Monterrey
Centro de Inteligencia Artificial
Sucursal de Correos J, C.P. 64849
Monterrey, N.L., Mexico
Phone: (52 83) 58-1400 EXT. 5135
Email: mvalenzu@mtecv2.mty.itesm.mx

## References

Barto, A. G. (1990). *Some learning tasks from a control perspective* (COINS Technical Report No. 90-122). Amherst: University of Massachusetts.

Booker, L. B. (1992, October). *Viewing classifier systems as an integrated architecture.* Paper presented at The First International Conference on Learning Classifier Systems, Houston, Texas.

Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning.* Reading, MA: Addison-Wesley.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986). *Induction: Processes of inference, learning, and discovery.* Cambridge, MA: MIT Press.