

Learning Classifier Systems Resources

Tim Kovacs

Department of Computer Science
University of Bristol
England, BS8 1UB
kovacs@cs.bris.ac.uk

Abstract. This article lists currently available sources of information on classifier systems and classifier systems research, both on-line and in print. The need for new resources, and improvements to certain existing ones, are suggested.

Keywords: Learning Classifier Systems

1 Introduction

The term “Learning Classifier System” (LCS) covers a wide range of systems, and published work on LCS has appeared in a wide variety of conferences and journals, beginning in the 1970s. This dispersion, and the volume of the LCS literature, can make it difficult to find suitable sources on LCS, a problem which this article addresses by listing currently available sources of information on LCS and LCS research. Also significant in this respect is the LCS Bibliography [17], an attempt to compile references to all LCS publications regardless of origin.

This article is intended primarily for newcomers to the field who might otherwise be unaware of important resources. It is intended as a brief guide, from which the reader can jump to more specialised sources, which have been structured in this document as follows. Section 2 lists introductory material for those who wish to learn the basics of classifier systems. Section 3 lists both broad and specialised overviews of research in the field, suitable for those seeking more advanced information. Section 4, devoted to LCS-friendly workshops and conferences, will interest those who wish to present work on LCS, or meet LCS researchers. Section 5 lists current on-line LCS resources, while section 6 deals with the availability of LCS source code. Finally, section 7 highlights the need for certain developments and concludes with an appeal for support for a number of valuable resources.

2 Introductory material

There are currently no authored books devoted solely to LCS, although several introductory texts on evolutionary computation include material on them. For a gentle introduction to LCS the 1989 text by Goldberg [12] is recommended despite its age. For other general introductory material on LCS see ([15, 7, 25, 2,

26, 14]). Theses have always been an important source of introductory (as well as advanced) material on LCS (see [17]).

Recent introductory material on some of the currently more popular variations of LCS includes an introduction to Anticipatory Classifier Systems (ACS) [30], an introduction to Fuzzy LCS [6], and introductory material on XCS, e.g., [32, 33, 19, 8].

Given the rapid development of research in XCS perhaps a new comprehensive introductory document will soon be needed. More generally, there is certainly more than enough material to form an introductory text devoted solely to LCS, and, given the renewed interest in LCS in recent years, this may be a viable proposition.

3 Overviews of LCS Research

Broad surveys of LCS research are relatively few. The 1989 paper by Wilson and Goldberg [35] is still valuable, although now somewhat dated. Following on from this, Lanzi and Riolo survey LCS research from 1989-1999 in [20]. Midway between these two lies Fogarty, Carse, and Bull's 1994 survey [9].

Surveys of particular subjects within the LCS field include Fogarty, Ireson, and Bull's 1995 survey of industrial and commercial LCS applications [10], Wilson's 1999 survey of the state of XCS research [34], Barry's surveys of work on hierarchical LCS [3, 5], and Bonarini's survey of work on Fuzzy LCS [6].

Surveys are particularly valuable given the now considerable size of the LCS literature (the LCS bibliography [17] currently contains over 600 entries). Further surveys are to be encouraged, particularly those providing critical analysis and comparisons of alternative approaches.

4 Conferences and Workshops

LCS papers are found in the proceedings of many conferences, but often in very small numbers. LCS papers have traditionally been numerous in the International Conference on Genetic Algorithms (ICGA) series, which has now combined with the Genetic Programming (GP) series to form the Genetic and Evolutionary Computation Conference (GECCO). LCS papers are found in lesser numbers in some of the other major evolutionary computation conferences, including the Congress on Evolutionary Computation (CEC) and Parallel Problem Solving from Nature (PPSN). LCS papers also tend to be found in the proceedings of the conference on Simulation of Adaptive Behavior (SAB).

The relative scarcity of LCS papers and researchers at most conferences means that by far the best opportunity to interact with LCS researchers is at events specifically devoted to LCS. The major such event is the International Workshop on Learning Classifier Systems (IWLCS), held so far in 1992, 1999 and 2000, and widely held to be a great success. IWLCS-99 was held in conjunction with GECCO-99 and its proceedings appeared in the GECCO-99 workshop

proceedings distributed at the conference. IWLCS-2000 was a joint workshop of SAB and PPSN and 2 page abstracts of papers presented there appear in the workshop proceedings distributed at the conference. IWLCS-2001 is to be co-located with GECCO-2001. A subset of papers from each IWLCS event is available at the LCS archive [18]. Extended versions of IWLCS-99 papers form the core of [21], and extended versions of IWLCS-2000 papers will form the core of [22].

The only other event known to the author to have been devoted solely to LCS was an organised session in 2000 at the Fourth Japan-Australia Joint Workshop on Intelligent and Evolutionary Systems, entitled “Exploring New Potentials in Learning Classifier Systems” [31].

5 The Internet

Few will be surprised that the internet plays an increasing role in the dissemination of information on LCS. Of particular note are the LCS mailing list [16] and NetQ [36], which offer unprecedented access to LCS researchers, and the increasing availability of full-text papers, often from authors’ personal home pages, and from the LCS archive [18]. A collection of links to home pages can be found at the LCS Web [4].

While the internet certainly facilitates communication between LCS researchers and increases the availability of LCS publications, web pages do not have the permanence of conference and journal publications. For this reason, any collection of on-line resources, such as that presented here, is bound to quickly become dated.

Current on-line resources on LCS include the following:

- 1. The LCS Web [4].** Alwyn Barry’s LCS Web provides a wealth of information, including LCS news, a list of LCS researchers with home pages and email addresses, a collection of LCS implementations, on-line papers, and an LCS glossary. It is recommended as a first stop on the net for those interested in LCS.
- 2. The LCS Mailing List [16].** John Holmes’s mail alias, which forwards email to many LCS researchers, is an excellent way to get in touch with the LCS community to seek advice, raise issues, or advertise papers. Ideally it would be archived and made available on the net, but as yet this is not the case.
- 3. The LCS Bibliography [17].** This growing bibliography currently contains over 600 references, many with abstracts, has an on-line search interface and is freely available. Since 1999 it has formed part of Alf-Christian Achilles’s Collection of Computer Science bibliographies [1]. Adding a paper to the bibliography is an important way of advertising its existence to other LCS researchers. Contributions and corrections are always welcome and may be emailed to kovacs@cs.bris.ac.uk. Figure 1 shows the distribution of entries by year as of December 2000.
- 4. The LCS Archive [18].** This archive is an appendix to the LCS Bibliography containing electronic full-text copies of LCS papers. Download statistics

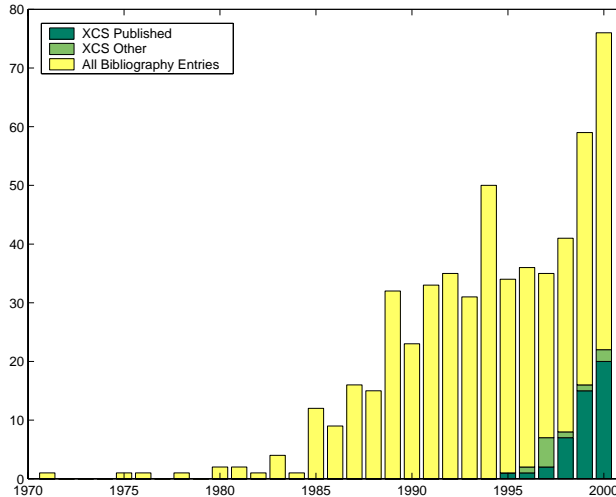


Fig. 1. Publications by year in the LCS Bibliography.

show this is an effective way of disseminating papers, and contributions are always welcome.

5. **Research Index [23, 24].** Research index offers full-text search and cross-referencing of a limited number of LCS papers. It is hoped that this project will grow, and that a greater number of LCS papers will find their way into its archive thanks to its software which autonomously seeks out electronic papers on the net. This software is aware of the LCS Archive, and any contributions to the archive should eventually find their way into the research index.
6. **Genetic Algorithms Digest List (GA-List) [27].** GA-list is a large mailing list about Genetic Algorithms which sometimes deals with LCS. Although it is a suitable forum for LCS-related material, the more specialised (though much smaller) LCS mailing list [16] is more likely to produce responses.
7. **GA-list archive [29].** The GA-list archive has back issues of the GA-list and other information on evolutionary computation.
8. **comp.ai.genetic** A USENET news group for discussion of evolutionary computation which occasionally hosts discussions of LCS. Again, using the LCS mailing list may have better results.
9. **The Hitch Hiker's Guide To Evolutionary Computation [14].** This is a valuable list of Frequently Asked Questions associated with comp.ai.genetic. It contains a section on LCS which, like most of the guide, is now somewhat dated, although it still provides much useful information, particularly for newcomers.

10. **ENCORE** [13]. ENCORE has a few LCS papers and some LCS software, but it clearly has not been updated in recent years. At present it contains a subset of the material which can be found at the LCS Web.
11. **NetQ** [36]. NetQ, recently re-established by its creator Stewart Wilson, is a scheme which allows anyone to query a paper's author on his work. An author first sets up a NetQ page for each eligible paper, and invites the public to post (optionally anonymous) questions electronically. Questions and answers are then added to the page for public display on the net. Benefiting both author and audience, NetQ has proved an excellent way of clarifying and disseminating information. We hope more authors will set up NetQ sites for their own work.

6 Software

Alwyn Barry has an impressive collection of LCS software at the LCSWeb [4], including versions of Goldberg's SCS (see [12] for documentation), Riolo's CFS-S and CFS-C systems (see [28]), Grefenstette et al.'s SAMUEL (see [11]), and several versions of Wilson's XCS (see [8] for documentation of one implementation). ENCORE [13] has a smaller collection of LCS software.

Making source code publically available is a great service to the LCS community. Source code provides a level of detail concerning implementation impossible to include in publications. It allows others to work with the exact system an author has used, making comparisons between different LCS more valid, and facilitates comparisons by eliminating the need for would-be comparators to reimplement published systems (although this can be a useful exercise). Allowing others to inspect and reuse source code makes it more likely for bugs, quirks, and discrepancies between implementations and published algorithms to be discovered, and, finally, sharing source code allows LCS implementation techniques to be shared, enriching the community's pool of expertise.

Of course the utility of source code is constrained by its quality and that of its documentation, and we should be doubly grateful to those who make the effort to provide clean, well-designed, and well-documented LCS implementations.

7 Conclusion

We hope this survey of resources will provide a useful starting point for many in their study of classifier systems, and also hope it will encourage others to provide new resources and improve existing ones. We suggested a need for more critical surveys of the literature, a comprehensive introduction to XCS and XCS research, and an introductory text on classifier systems. An upgrade of the LCS archive to provide full-text searches and crossreferences, and an archive of the LCS mailing list would also be useful.

We identified a number of valuable resources, and would like to encourage the development of personal NetQ pages, attendance at LCS-specific events, use of the LCS mailing list and contributions to the LCS archive and LCS bibliography. The more we support these resources the more useful they become!

References

1. Alf-Christian Achilles. *The Collection of Computer Science Bibliographies*. <http://iinwww.ira.uka.de/bibliography/index.html>
2. Thomas Bäck, David B. Fogel, and Z. Michalewicz, editors. *Handbook of Evolutionary Computation*. Institute of Physics Publishing and Oxford University Press, 1997. <http://www.iop.org/Books/Catalogue/>
3. Alwyn M. Barry. Hierarchy Formulation Within Classifiers System – A Review. In E. G. Goodman, V. L. Uskov, and W. F. Punch, editors, *Proceedings of the First International Conference on Evolutionary Algorithms and their Application EVCA'96*, pages 195–211, Moscow, 1996. The Presidium of the Russian Academy of Sciences.
4. Alwyn M. Barry. The Learning Classifier Systems Web. University of Bath, 2002. <http://www.cs.bath.ac.uk/~amb/LCSWEB/>
5. Alwyn M. Barry. *XCS Performance and Population Structure within Multiple-Step Environments*. PhD thesis, Queen's University Belfast, 2000.
6. Andrea Bonarini. An Introduction to Learning Fuzzy Classifier Systems. In Lanzi et al. [21], pages 83–104.
7. Lashon B. Booker, David E. Goldberg, and John H. Holland. Classifier systems and genetic algorithms. *Artificial Intelligence*, 40:235–282, 1989.
8. Martin V. Butz and Stewart W. Wilson. An Algorithmic Description of XCS. To appear in [22].
9. Terence C. Fogarty, Brian Carse, and Larry Bull. Classifier Systems – recent research. *AISB Quarterly*, 89:48–54, 1994.
10. Terence C. Fogarty, N. S. Ireson, and Larry Bull. Genetic-based Machine Learning – Applications in Industry and Commerce. In Vic Rayward-Smith, editor, *Applications of Modern Heuristic Methods*, pages 91–110. Alfred Waller Ltd, 1995.
11. John J. Grefenstette and H. G. Cobb. *User's guide for SAMUEL, Version 1.3*. NRL Memorandum Report 6820, Naval Research Laboratory, 1991.
12. David E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, MA, 1989.
13. Jörg Heitkötter. ENCORE. The Evolutionary Computation Repository Network. <http://www.cs.bham.ac.uk/Mirrors/ftp.de.uu.net/EC/clife/>
14. Jörg Heitkötter and David Beasley (editors). The Hitch-Hiker's Guide to Evolutionary Computation: A list of Frequently Asked Questions (FAQ). USENET: comp.ai.genetic <ftp://rtfm.mit.edu/pub/usetnet/news.answers/ai-faq/genetic/>
15. John H. Holland. Escaping brittleness: The possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In Mitchell, Michalski, and Carbonell, editors, *Machine learning, an artificial intelligence approach. Volume II*, chapter 20, pages 593–623. Morgan Kaufmann, 1986.
16. John Holmes. The Learning Classifier Systems Mailing List. The University of Pennsylvania, 2000. cslis@cceb.med.upenn.edu. Mail John Holmes at jholmes@cceb.med.upenn.edu to subscribe.
17. Tim Kovacs. A Learning Classifier Systems Bibliography. Department of Computer Science, University of Bristol, 2002. <http://www.cs.bris.ac.uk/~kovacs/lcs/search.html>
18. Tim Kovacs. The Electronic Learning Classifier Systems Archive. Department of Computer Science, University of Bristol, 2002. <http://www.cs.bris.ac.uk/~kovacs/lcs/lcs.archive.html>

19. Tim Kovacs. A Comparison of Strength and Accuracy-Based Fitness for Learning Classifier Systems. PhD thesis, School of Computer Science, University of Birmingham, Birmingham, U.K., 2002.
20. Pier Luca Lanzi and Rick L. Riolo. A Roadmap to the Last Decade of Learning Classifier System Research (from 1989 to 1999). In Lanzi et al. [21], pages 33–62.
21. Pier Luca Lanzi, Wolfgang Stolzmann, and Stewart W. Wilson, editors. *Learning Classifier Systems. From Foundations to Applications*, volume 1813 of *LNAI*. Springer-Verlag, Berlin, 2000.
22. Pier Luca Lanzi, Wolfgang Stolzmann, and Stewart W. Wilson, editors. *Advances in Learning Classifier Systems*, volume 1996 of *LNAI*. Springer-Verlag, Berlin, 2001.
23. Steve Lawrence, C. Lee Giles, and Kurt Bollacker. Research Index. The NECI Scientific Literature Digital Library. <http://citeseer.nj.nec.com/>
24. Steve Lawrence, C. Lee Giles, and Kurt Bollacker. Digital libraries and autonomous citation indexing. *IEEE Computer*, 32(6):67–71, 1999.
25. Zbigniew Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs*. Springer-Verlag, 1996.
26. Chilukuri K. Mohan. *Expert Systems: A Modern Overview*. Kluwer, 2000.
27. Mitchell Potter and Annie Wu (moderators). Genetic Algorithms Digest List (GA-List). To subscribe mail: ga-list-request@aic.nrl.navy.mil with “subscribe” in the body of the mail.
28. Rick L. Riolo. *CFS-C: A Package of Domain-Independent Subroutines for Implementing Classifier Systems in Arbitrary User-Defined Environments*. Technical Report, University of Michigan, 1988.
29. Alan C. Schultz. The GA-List Archive. <http://www.aic.nrl.navy.mil/galist>
30. Wolfgang Stolzmann. An Introduction to Anticipatory Classifier Systems. In Lanzi et al. [21], pages 175–194.
31. Keiki Takadama, editor, *Exploring New Potentials in Learning Classifier Systems*. A Session of the 4th Japan-Australia Joint Workshop on Intelligent and Evolutionary Systems, Ashikaga Institute of Technology, 2000.
32. Stewart W. Wilson. Classifier Fitness Based on Accuracy. *Evolutionary Computation*, 3(2):149–175, 1995. <http://prediction-dynamics.com/>
33. Stewart W. Wilson. Generalization in the XCS classifier system. In J. Koza et al., editors, *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 665–674, San Francisco, CA, 1998. Morgan Kaufmann. <http://prediction-dynamics.com/>
34. Stewart W. Wilson. State of XCS Classifier System Research. In Lanzi et al. [21], pages 63–82.
35. Stewart W. Wilson and David E. Goldberg. A Critical Review of Classifier Systems. In *Proceedings of the Third International Conference on Genetic Algorithms (ICGA-3)*, pages 244–255. Morgan Kaufmann, 1989. <http://prediction-dynamics.com/>
36. Stewart W. Wilson. *NetQ*. <http://prediction-dynamics.com/>