

LISP LORE: A Guide to Programming the LISP Machine, (2nd Edition)

Hank Bromley and Richard Lamson

Kluwer Academic Publishers

\$47.50 ISBN: 0-89838-228-9

Reviewed by: Keith Price

Considering my reaction to the first edition of LISP LORE: A Guide to Programming the LISP Machine (see Sigart #100, I generally liked it, but the operating system changed), I was pleased to see a second. This new edition is based on the first with substantial additions and rewriting (by Richard Lamson of Symbolics) to account for the changes made by Symbolics between version 6.1 and 7.0 of their system. The book is somewhat expanded and rearranged (flavors are no longer the first topic that is introduced) with new example programs (as before, source is also available on tape).

LISP machines (from any of the few manufacturers) present a challenge to a beginning user since the style of effective usage is different from most other operating systems (e.g. Unix, VMS, etc.). Even though commands and features differ, the basic initial usage of the common interactive systems (and even the old card based systems) is very much the same. The LISP machine families seem to require more initial information before any usage is possible. This hurdle tends to eliminate the casual user - there is a large investment of time to know enough to do simple things so it is either used for everything or nothing. But once past this initial step, the ease of use and the overall system capabilities make such systems enjoyable to use.

As a (2-year+) user of a LISP machine, does this book teach me anything? Probably not, except for small things that become clearer, or new functions and features that are mentioned. But much the same thing happens any time I read the manuals; I find some new useful feature that I probably was not looking for and never knew I needed.

Do new users find the book helpful? Based on how the first version was used by some of the other users in our group, I think the introduction is short enough for new users to get past the initial shock of the 13 volume manual set and better organized than the introduction in the manuals. The advanced features such as windows, graphics, files and building systems are described well enough for a new user to get started and to use these features effectively. It is then easier to read the manuals to learn details after you are already using simple versions of these complex subsystems.

The book fails to answer all the obscure questions that arise through using such a system. (For example, what causes programs to get an error when adding two bignums, but not always the same bignums, or the same place, but repeatably with a given pair of bignums deep in a given program? After several days of messages, the only guess was hardware. One system was "cured," and another has had a similar problem in different programs.) It also does not address another recent (and recurring) problem where we had a complete seizure of all the systems upon attempting any network access (the solution required disconnecting taps from the physical network until the problem was found, the failure affected every system on the cable and through all the bridges), but the discussion of networks, processes, and scheduling at least helped understand why this problem occurred.

Overall I like this edition better than the first and expect that it will be used, both by current users of our systems and by new users over the next few years.

MACHINE LEARNING: A Guide to Current Research

T. M. Mitchell, J. G. Carbonell, R. S. Michalski

Kluwer Academic Publishers

\$55.00 ISBN: 0-89838-214-9

Reviewed by: Keith Price

Machine Learning: A Guide to Current Research provides a collection of current research papers mostly collected from participants at the Third International Machine Learning Workshop held in 1985. These 77 separate papers were collected and reformatted at Rutgers to produce a consistent format and a common bibliography. The range of learning topics include (taken from the preface): analogy, conceptual clustering, explanation-based generalization, incremental learning, inductive inference, learning apprentice systems, machine discovery, theoretical models, and applications. All of these topics, except incremental and apprentice systems are included as heading in the index. Other useful index headings include research institution (a few are missing, primarily those where the several authors are from more than one place) and program name. The papers are short; the average length of five pages per paper, based on the number of papers and the length of the book, overstates their length because of blank pages due to formatting considerations. In the short space for each paper it is not possible to explain in detail how anything is done. There is space only to say what is being done by the system. Thus, the bibliographic references become important for a deeper understanding of individual techniques.

As a researcher in a field other than machine learning, I can not evaluate the quality of each of the papers. The three editors are better able to make that judgment. In the preface they call this book a "snapshot" of the field and indicate that it "provides a representative sampling of the best ongoing work in the field." As such, the book serves the purpose of a research review for the entire field rather than a description of how to duplicate various research results (as should be demanded of refereed journal publications, for example). This makes the book a good initial reference, especially with the index and the single bibliography, both for researchers in machine learning and those outside the field who want to learn about it or at least keep generally aware of its major directions and concepts. Also, I see this as a research or reference book more than a text book.

The reformatting of the various submissions is a monumental task (this comes from the experience of doing similar things for the Newsletter for a few years), and the layout, etc. was generally good. The reference information and style was somewhat inconsistent, most likely reflecting what was submitted by the original authors. (In the Newsletter, I find more formatting, spelling, etc. errors in the on-line submissions or the original hard-copy than in what is re-typed here.) The book gives me a better perspective of the machine learning field and may be helpful for my long-term research since learning is ultimately the most important feature of an AI system.