

Preface

A Life of Rewards: The Collected Works of Larry Wos

Among the various aspects of computing that have proved fascinating, you may find the aspect covered in the papers in this two-volume set astounding. Specifically, the papers focus on a challenge that certain researchers—I included—accepted in the early 1960s and still accept. The challenge is that of producing a single computer program that automates logical reasoning so effectively that questions whose answers have eluded great minds for decades will frequently yield their treasure. By reading or browsing among these papers, you will see how progress occurred and what techniques were formulated to enable a program (such as William McCune's OTTER) to answer diverse open questions and solve various hard problems. You will also share a researcher's effort and excitement covering more than thirty-five years. Note that, with one exception, the order of the papers in these two volumes essentially mirrors the order in which the research occurred and the papers were written, and does not follow the accidental order in which they were published. The exception is the paper entitled "Resolution, Binary: Its Nature, History, and Impact on the Use of Computers"; in particular, the 1992 version of that paper is included, rather than the original 1987 version, to provide more up-to-date information about OTTER and its successes.

Although the papers contained in the two volumes share a common author, their style ranges from the very formal to the most friendly. You may be amused by their violation of a principle of Paul Halmos: A writer should publish his works under the same name, not vary from Lawrence Wos to Larry Wos and the like. Finally, you will find (though not obviously) evidence of how my research was influenced, guided, and inspired by the Significant Seven. You will learn in this preface who they are.

As is true of my writing style and of the series of names under which I published, the field that is the focus of the vast majority of these collected papers underwent identity changes. In the beginning (in the early 1960s) the field was known as mechanical theorem proving, then automatic theorem proving, then automated theorem proving, and finally and correctly here in 1999 as automated reasoning. (Automated deduction is a subfield, in my view, of automated reasoning, the latter name being introduced by me in 1980.) Far, far more significant than the preceding is the progress that has occurred in less than four decades. Both the amount and the rate of progress are reflected in the papers presented in these two volumes.

Indeed, in the beginning we were pleased when a reasoning program succeeded in proving the simplest of theorems, such as the classroom exercise that

asserts the provability of commutativity for groups in which the square of x (for all x) is the identity e . Such an achievement by a graduate student would provide no evidence of potential in the context of research. However, it did correctly presage what might and, eventually, did occur in the automation of logical reasoning.

For the most startling example of the fulfillment of promise, you need only glance at the marvelous success of my colleague William McCune in the area of Robbins algebra. Specifically, for more than six decades (commencing in the early 1930s), the question of whether every Robbins algebra is a Boolean algebra remained open, not even yielding its treasure to such eminent scientists as Alfred Tarski. McCune's success in answering the question is in no way isolated, which can be seen by glancing at some of my later papers and by exploring a monograph written by McCune and our colleague R. Padmanabhan.

If you are tempted to add to the successes and, therefore, wish targets, I suggest *A Fascinating Country in the World of Computing: Your Guide to Automated Reasoning* (the companion to this two-volume set), which offers many hard problems and numerous open questions, especially in Chapter 11. A large fraction of the problems and questions will appeal to you even if you have little or no interest in automation. On the other hand, if you would enjoy the companionship of a fine program (McCune's OTTER), the CD-ROM included in the cited book will prove most useful.

Now, to fulfill the promise made in the beginning, I come to the Significant Seven, in the order in which I met them. The superlatives you will encounter are more than deserved; the gratitude that is expressed is genuine; the history that is presented is accurate. What was immediately clear to me, as I met each of the Seven, was that I was in the presence of a marvelous mind.

Of the Seven, George Robinson was the first. Although he resides no longer on this planet, he is still the most logical man I have ever known. The smallest taste of his fine insights is offered in Section 1.8 of *A Fascinating Country in the World of Computing*. From him, I learned much, not the least of which concerns experimentation and the identification of hidden premisses.

Second was Dan Carson, one of the finest computer programmers that has ever existed. Were it not for a lack of money (in those early 1960s), he would have obtained a Ph.D. His theorem-proving program (of the early and mid 1960s) was clearly crucial to the field at the time and, without it, I am almost certain that the field would have died. It remained by far the most powerful—until the third member of the Seven arrived.

The person in question, Ross Overbeek, can only be described as brilliant. His reasoning programs and his ideas (that include the weighting strategy and the inference rule UR-resolution) have proved indispensable to the field. His

influence, which will forever be vastly underestimated, can be found in the best of all reasoning programs extant here in 1999, namely, OTTER.

The fourth of the Significant Seven is Robert Veroff. You will find little of his work in the literature—what a loss! He continues to have ideas that increase the power of reasoning programs. Such is not just my opinion; indeed, data strongly supports the position, including the discovery of new fine proofs and the formulation of techniques that he has freely given to colleagues. One of my failings concerns my lack of success in getting him to publish his breakthroughs.

The fifth of the Seven is Steve Winker, who toured MIT in but two and one half years, obtaining his B.A; years later, he easily obtained his Ph.D. Unfortunately, he also published rather sparsely. He played a key role in the Argonne group's attack on open questions in the late 1970s and the early 1980s. His successes in Robbins algebra were crucial to McCune's remarkable result.

Robert Boyer is the sixth of the Seven. His work (with his colleague J Moore) in program verification sets a standard that may never be matched. Boyer is the one person I might have believed about the existence of a finite axiomatization of set theory. Of course, he was right. He (along with his advisor Woody Bledsoe) caused me to write my book *The Automation of Reasoning: An Experimenter's Notebook with OTTER Tutorial*.

The seventh member of the Significant Seven is William McCune. He is deservedly a legend in his own time. His program OTTER dominates my later papers, my experimentation, and my ability to conduct research. Among other items, he has contributed strategies of note. His work (with Padmanabhan) as well as his other successes in answering open questions can be described only as beautiful.

As for future greats, I suggest you follow the efforts of my latest colleague, Branden Fitelson. In less than one year of collaboration with me, jointly and separately, he has dispatched a number of open questions. He shows more promise than any researcher I have known since I met McCune. Branden is more than a welcome member of the Argonne group, sharing our emphasis and requirement of substantial experimentation.