

IN SUPPORT OF RESEARCH IN DEDUCTIVE SYSTEMS

It is evident that deductive inferences are an important element of automatic problem solving procedures. Their need arises in many AI areas, e.g., natural language understanding, scene analysis, robot planning, automatic program synthesis, etc. Two of the approaches to inference making which have received wide attention are formal theorem proving procedures and the deductive processes embedded in the AI languages. In fact, the motivation behind the development of AI languages was, in part, a desire to overcome certain deficiencies of formal theorem proving methods in their state of development circa 1970. Three major deficiencies of these methods were (1) the inability to deal with a dynamically changing world, or with a multiplicity of worlds which could possibly contradict each other; (2) the need to represent the world axiomatically was inefficient, if not totally inadequate for certain kinds of information which might otherwise be more efficiently represented procedurally; and (3) the inability to integrate and utilize semantic information in the deductive process made theorem proving procedures especially susceptible to the combinatorial explosion associated with the deductive process.

With regard to deficiency (1), this is a serious problem which cannot be solved from within the theorem proving procedure. That is, one must go outside the procedure and manage the progression of worlds from one state to another. This problem is elegantly solved by the context mechanisms of CONNIVER and QA4. The conclusion that one must draw from this is not that theorem proving methods are no good, but only that they are not enough in and of themselves to deal with a dynamic world. In fact, it is important to draw

a distinction between static and dynamic worlds. The solution to problems in a dynamic, or changing, world may be obtained by linking together the solutions to problems in a sequence of static worlds. Thus, we need methods for specifying and controlling the generation of static worlds, and for conducting a search through a space of worlds. The AI languages are well-suited to this task. On the other hand, we also need effective methods for controlling and limiting the (deductive) search for solutions to problems in a given static world. This is an important problem which neither the AI language researchers nor the theorem proving researchers have solved. It is a problem addressed by the proposed research in deductive systems.

With regard to deficiency (2), the AI language approach guarantees that any information which might be best expressed procedurally will be done so. In fact, all axiomatic information is expressed this way. We would suggest that perhaps this is just as extreme a restriction as requiring all axiomatic information to be syntactically described. In fact, the problem of deduction is far from solved by this approach. The combinatorial explosion in the number of syntactic entities has been replaced by a combinatorial explosion in the number of procedure calls. Thus, even with the AI languages, all but the simplest of deductive problems are still intractable because of all the thrashing around induced by a deductive search.

With regard to deficiency (3), the possibility of a combinatorial explosion is what makes deduction a hard problem. The deductive mechanisms described in the NSF proposal do integrate well with semantics as well as procedural knowledge. We take a very pragmatic view of the deductive process which does not fit into any stereotype of earlier theorem proving research. Thus, through the use of semantics, procedural knowledge, well conceived search

procedures, redundant inference elimination and learning, we hope to be able to deal with rather complex searches in quite large data bases (or static worlds). And all of this we plan to embed into a programming language, perhaps an AI language, so that these deductive processes may be applied to problems in a succession of static worlds.

To conclude, we believe that what is needed is a better understanding of the deductive process and how it may be effectively controlled and limited. The particular formalism which should be employed cannot be determined a priori but rather it must be shaped by the knowledge of what is required.

The proposed research in deductive systems seeks to understand, and thereby learn to control, the deductive process. The results of the research will have something to say to researchers in theorem proving, but it will also have something to say to developers and, perhaps users, of AI languages.

In fact, we have already learned much about the deduction process. As a result of our theorem-proving experiments with a large data base containing as many as 500 axioms, it has been discovered that numerous refinements of resolution developed over the years, some of which make inferences very selectively on certain problems, will yield poor results when applied to large data base problems. What is more, it is clear why this is the case. It is also clear what properties are required of inference methods which do not suffer from their faults. Such methods have been proposed for development in the NSF proposal.

What we are also learning about the deduction process are characteristics that a search strategy must possess to effectively control the search, and where, and in some cases, how semantics may be used to limit and guide the deduction process. These observations do not apply only to theorem prov-

ing, but to deductive processes in general. In particular, they have much to say about the control structures in AI languages and how to improve their dynamic behavior.