A Universal Automated Information System for Science and Technology

Peter B. Andrews

Carnegie Mellon University, Pittsburgh, PA, U.S.A.
andrews@cmu.edu
http://gtps.math.cmu.edu/andrews.html

Abstract. It may be centuries before there is a well developed Universal Automated Information System for Science and Technology, but such a system is starting to grow even now. We consider various aspects of such a system.

Let us try to look ahead a few hundred years and imagine what we may find. Surely there will be computer-based information sources which can be used as aids in answering questions from all realms of mathematics, science, applied science, and technology. Just as those searching for information on the web today tend to regard "the web" as one monolithic information source, those who use these information sources of the future are likely to think of them as comprising a Universal Automated Information System for Science and Technology. What will this system be like? One possibility is that it will be a huge kluge of systems designed by specialists from various fields using whatever techniques and programming methodologies they found most immediately convenient. Another possibility is that it will be a well organized union of information systems in which formal logic and techniques of automated deduction play significant roles.

When one uses an information system, one may be seeking information which is not explicitly in the system, but which can be derived from information which is there. Thus, it is desirable that information systems be organized so that one can apply logical inferences to them. This is particularly important if it is a computer, rather than a person, that is looking for the information. People often make inferences almost subconsciously, but reasoning by computers must be explicit. The need to derive consequences of information suggests that the information should be represented in a form which can at least be readily translated into a formal language for which there are well known rules of logical inference, i.e., a formal logical language.

The Automated Information System of the future will undoubtedly have many components which have been contributed by many different people and projects at many different times. It seems desirable that the content of information in the system be separated from mechanisms for retrieving and deriving information, so that improvements in the retrieval and inference mechanisms

[0] This work was supported by NSF grant CCR-0097179.

can be made in a modular way which will benefit users of all components of the system.

Sometimes one asks a question which can be answered only by using knowledge from several disciplines. Thus, it is highly desirable that the Automated Information System be well integrated to facilitate dealing with such questions.

Let us consider as a very grand challenge the design and development of a Universal Automated Information System for Science and Technology having at least the following features:

- There are formal logical systems (formal languages with rules of reasoning) in which one can represent, in useful and natural ways, all the statements which might be made in the fields of mathematics, computer science, physics, chemistry, biology, engineering, medicine, and other physical sciences and applications of physical science. In particular, all information which is used in these fields is representable in these logical systems.
- There are good interfaces for these logical systems. There are algorithms for translating statements between different formal languages, and between formal languages and natural languages.
- Virtually all of the knowledge which constitutes the fields mentioned above is expressed in these formal systems. This includes not only factual data, but also definitions and fundamental principles which one may regard as axioms, hypotheses, or theorems. This knowledge, as so represented, is stored in computerized information libraries.
- The system grows continually as new knowledge is made available to it.
- There are Automated Reasoning Systems which can access this knowledge and apply logical inferences to derive answers to a great variety of questions. Calculations, algorithms, decision procedures, and special techniques for special problems are applied appropriately. These Automated Reasoning Systems can be applied to work automatically, semi-automatically, or interactively.

Of course, many people have been thinking about such a system, in at least a casual way, for many years. Indeed, the QED project [7] is concerned with building such a system for mathematics, and many of the same considerations apply to both. A system having all the features described above may take centuries to develop, but the Automated Information System of the future is growing right now, and it is important to encourage growth in the right directions. Good plans about how such a system should grow could have enormous future benefits. Let's think about the design of such a system.

It is not easy to make good plans about complex projects, and people will disagree about many details. One should try to reach as much consensus as possible on basic ideas and principles underlying the project, and explicitly plan for ways of handling disagreements. For example, different choices of formal languages may be accommodated by suitable interfaces and translation mechanisms. In the long run, the system will evolve in ways that no one can foresee or control.

Most information is written in a natural language (such as English or French), but to obtain the benefits of automated deduction one needs (some of) this information to be expressed in a formal logical language. Enormous progress could be made in developing automated information systems if suitable portions of natural language could be translated automatically into symbolic logic. For example, if one could translate all the theorems and definitions in Bourbaki's *Elements of mathematics* [4] into machine-readable formulas of symbolic logic, one would have a very impressive mathematical library for automated reasoning systems to use.

Formalization of the knowledge in any intellectual discipline is an enormous task. However, handling the details involved in such a task and testing various aspects of the formalization are greatly facilitated by using a general-purpose theorem-proving system (such as TPS [10]) which can be used in a mixture of automatic and interactive modes and has suitable library facilities. Automated deduction makes formalization much more tractable than it has ever been, and formalization will continue to become more tractable as theorem-proving systems improve.

Mathematics, which is often referred to as the language of science, plays a role in all technical disciplines, so a language which is used to formalize a technical discipline should at least be adequate for formalizing mathematics. Type Theory [1] and axiomatic set theory [9] have been studied extensively as general purpose languages for expressing mathematics. If one has an information system which is based on first-order logic, one can simply regard it as being based on type theory, which includes first-order logic. If one uses some extension of axiomatic set theory, it must be an extension of a formulation such as [9] which accommodates entities which are not sets.

As far as formalization is concerned, mathematics has received much more attention than other disciplines. Still richer languages than those needed for mathematics may be needed to formalize certain disciplines in suitably natural ways. For example, it may be desirable to incorporate into the formal language of science some representation of the diagrams of molecular structure which chemists use. Actually, diagrams are used in mathematical reasoning too[1], and the relevant theory for Venn diagrams is well developed [6][8]. However, at present, facilities for using Venn diagrams are not integrated into most general-purpose automated theorem proving systems.

Answering a question or solving a problem involves much more than deriving a statement which expresses the answer from relevant information. One must *find* the answer as well as prove that it is correct. This may involve search, logical deduction, and a variety of special techniques. We mention [5] and [11] as examples of research in this area. One way of testing problem-solving techniques is to formalize the information in relevant textbooks, and see if the problem-solving techniques are adequate for automatically doing the exercises in those textbooks.

[1] See [3] for persuasive arguments about using diagrams in mathematical proofs in ways that involve no compromise with complete rigor.

Finding relevant information becomes an increasingly serious problem as the amount of available information increases. There will be an ongoing need for research on this problem.

One well known method of finding relevant information is to classify items of information, and retrieve items with relevant classifications. Existing classification systems need to be extended so that (in mathematics, for example) one can classify not only books and articles, but also theorems, definitions, and examples. Current classification systems leave much to be desired, at least in mathematics. Mathematicians sometimes have trouble finding out whether a theorem they have proven has been proven before. The most practical method of answering this question is usually to rely on the the memories of experts in the field.

One would hope to at least have effective ways of determining whether a specified theorem is already in the Automated Information System. This can be a nontrivial problem, since there are many ways of expressing mathematical theorems. For example, the statements "Every function which has an inverse is bijective" and "A function which is not bijective cannot have an inverse" are clearly two ways of expressing the same theorem. One can think of a variety of ways to make trivial changes in the ways theorems are expressed without changing the essential mathematical content of the statements. Can we find some natural equivalence relation on the set of statements of a formal language for expressing mathematical theorems such that equivalent statements should be regarded as expressing the same theorem? Clearly logical equivalence is much too broad for this purpose, since all provable statements are logically equivalent to each other.

A classification scheme for mathematics should be robust. Different ways of expressing a theorem should not lead to different classifications for it.

One would like to have a classification scheme for mathematics which is based on some fundamental understanding of the structure of the field of mathematics, and which is objective. We don't want to have to convene a committee of mathematicians to decide how a new theorem should be classified. Unfortunately, we don't yet have any such fundamental understanding of the structure of mathematics.

Of course, classification is another area where there can be disagreements. Different classification schemes may be most appropriate for different purposes, or simply preferred by different people. One solution to such disagreements is to permit many classification schemes to be available simultaneously, and permit users to decide which classification to use at any particular time. Such a multiple classification system has been implemented for the TPS library by Chad Brown; it is described briefly in [2]. It remains to be seen whether complications arise when one tries to establish multiple classification schemes for enormous amounts of information.

Since the information in the Automated Information System will be contributed by many people and processes at many different times, some of it is likely to be unreliable or obsolete. Systems which provide some degree of authentication for certain components of the entire information system will be needed, and automated reasoning systems will need to keep track of the sources of information they use.

The development of an Automated Information System such as we have been discussing is indeed a very grand challenge. Developing such a system would not only provide an intellectual tool which would elevate our ability to reason reliably about complex technical questions to a new level, it would also lead to a deeper understanding of the nature and structure of our knowledge. Developing such a system involves working on a number of tasks which are themselves grand challenges. We close by summarizing some of these:

- Develop general standards which would permit the many components of the Automated Information System to be used together in a harmonious and efficient way.
- Promote widespread discussion and adoption of these standards.
- Study what extensions of existing formal languages may be needed for the formalization of various scientific and technical disciplines. Investigate whether various formal languages satisfying these requirements may be regarded as specializations of a single more general formal language.
- Formalize the knowledge in all scientific and technical disciplines.
- Develop automated systems for translating between various formal and natural languages.
- Develop improved methods of retrieving relevant knowledge.
- Develop good classification schemes for knowledge in all technical disciplines.
- Improve theorem-proving systems.
- Improve and extend question-answering and problem-solving systems which are based on automated deduction.
- Develop mechanisms for safeguarding, verifying, monitoring, and assessing the reliability of information obtained from the Automated Information System.
- Find ways to influence the education of students in scientific disciplines so that eventually workers in these fields will be familiar with formal logic and automated reasoning tools, and will be able to effectively use and contribute to those components of the Automated Information System which are relevant to their special interests.

References

1. Peter B. Andrews. *An Introduction to Mathematical Logic and Type Theory: To Truth Through Proof.* Kluwer Academic Publishers, second edition, 2002.
2. Peter B. Andrews, Matthew Bishop, Chad E. Brown, Sunil Issar, Dan Nesmith, Frank Pfenning, and Hongwei Xi. TPS *User's Manual*, 2003. 103+iv pp. Available from http://gtps.math.cmu.edu/tps.html.
3. Jon Barwise and John Etchemendy. Visual Information and Valid Reasoning. In Walter Zimmermann and Steve Cunningham, editors, *Visualization in Teaching and Learning Mathematics*, pages 9–24. Mathematical Association of America, 1991.

4. Nicolas Bourbaki. *Elements of mathematics.* Hermann, 1966. 10 volumes. Translation of *Elements de mathematique.*
5. Cordell Green. Application of Theorem Proving to Problem Solving. In Donald E. Walker and Lewis M. Norton, editors, *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 219–239, Washington, D.C., 1969.
6. Eric M. Hammer. *Logic and Visual Information.* CSLI Publications & FoLLI, Stanford, California, 1995.
7. The QED Manifesto. In Alan Bundy, editor, *Proceedings of the 12th International Conference on Automated Deduction*, volume 814 of *Lecture Notes in Artificial Intelligence*, pages 238–251, Nancy, France, 1994. Springer-Verlag.
8. Sun-Joo Shin. *The Logical Status of Diagrams.* Cambridge University Press, 1994.
9. Patrick Suppes. *Axiomatic Set Theory.* D. Van Nostrand Company, Inc., Princeton, N.J., 1960.
10. TPS and ETPS Homepage. http://gtps.math.cmu.edu/tps.html.
11. R. Waldinger, P. Jarvis, and J. Dungan. Pointing to Places in a Deductive Geospatial Theory. In A Kornai and B Sundheim, editors, *Workshop on Analysis of Geographical References; Human Language Technology Conference*, Edmonton, Canada, May 2003. NAACL.