

A Point of Leverage for Artificial Intelligence

Jon Doyle

Laboratory for Computer Science
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139, USA
doyle@lcs.mit.edu

<http://medg.lcs.mit.edu/people/doyle/doyle.html>

Abstract

AI should swiftly educate workers in other fields of science, engineering, art, etc. about the techniques it possesses for structuring and organizing knowledge, as many fields are now working in ignorance of AI techniques to develop their own representation “standards” embodying many of the mistakes discovered and corrected long ago by AI.

The current efforts to make all information available in a global information network offers a large opportunity to artificial intelligence, as means for representing bodies of general and personal information will be absolutely crucial in making this information useful. Servers used by automated agents and by humans will need formalized, easily communicable and easily explicable versions of the knowledge and information involved in all the processes and activities of society and everyday life. More importantly, the automated agents employed by people and by organizations will need accurate representations or embodiments of the evolving plans, preferences, and beliefs of the people or organizations they serve. For example, the *Guardian Angel* project at MIT¹ seeks to improve the delivery of health care to individuals by augmenting the fragmented, incomplete, and often inaccessible medical records currently maintained (or supposedly maintained) by each provider with lifelong personal medical records carried with each person (as “credit cards” or “dogtags” or “bracelets”) that immediately provide routine and emergency care providers with the information they need, including both information about the person’s medical history and about the person’s preferences about treatment and outcomes. These personal representations of medical information naturally call for external representations of general medical knowledge (for use in consultation and explanation) and organizational records (for each hospital, physician, etc). While the textual records currently offered by medical handbooks and reference books are useful for these purposes, convenient and automated exploitation of such knowledge calls for nontextual representations that more clearly reflect the structure and use of the knowledge. Simply positing hypertext versions of extant books serves the real needs only poorly, since useful organizations need not pure pointers to information but descriptive or prescriptive directions that indicate the practical and theoretical relation of the current information to the indicated information.

Recognizing this need, many fields are now working to formalize taxonomies of their concepts, and standard definitions of all these concepts, in order that the work of formalizing their domain of expertise may be divided. Unfortunately, many such efforts have been undertaken with no awareness of even the existence of the AI field of knowledge representation. These efforts have only the most rudimentary notions of taxonomy to work with, and to the extent they try to improve on these, seem to be reinventing many things AI research on taxonomic representation systems found

¹P. Szolovits, J. Doyle, W. J. Long, I. Kohane, and S. G. Pauker. *Guardian Angel: Patient-centered health information systems*. Technical Report MIT/LCS/TR-604, Massachusetts Institute of Technology, Laboratory for Computer Science, 545 Technology Square, Cambridge, MA, 02139, May 1994.

to be flawed years ago. Reinvention of mistakes seems especially likely as these efforts move beyond formalizing mere taxonomies of terms to formalizing theories and data.

There is no good purpose served by permitting development of dozens (or even hundreds) of incompatible and flawed systems for representing knowledge in different fields when AI already has a long history of exploring these issues and, in some cases, has developed accepted and powerful techniques for expressing and organizing the knowledge of interest. This is not to say AI has solved all representational problems, or that efforts by “amateurs” in other fields cannot make contributions. But since the mistakes will be with us a long time, it seems best to avoid unnecessary ignorance the better to further the efforts in other fields at uncovering problems still needing solution rather than preserving mistakes.

There are several areas in which AI could immediately provide help to efforts in other fields.

- To help with the initial work of formalizing taxonomies and theories in other fields, AI should vigorously publicize and educate others about the expressive, robust, and meaningful representation languages and systems now used and in development in AI. The fair convergence of ideas in description logics offers useful tools ready to communicate to others, and the ongoing implementations (such as Classic, Loom, Krep) offer tools ready to use to others.
- Connected with this, many fields organized their everyday knowledge using simple default rules to express standard expectations and exceptions. While these are not yet fully part of the concept languages just mentioned, AI provides clear and useful theories of these constructs, along with many examples of how to use them in fruitful ways. These too then offer value to other fields in formalizing protocols and other process descriptions along with the exceptions that routinely arise.
- AI work on representing probabilistic information in Bayes nets is already used heavily in many practical systems and has immediate relevance to representing and reasoning with many domains of knowledge.
- AI also has several substantial proposals for standard languages for communicating information among different representations. While there is not yet complete agreement about these communication standards, something like them may offer a better alternative than each field of human activity inventing its own flawed system with no means of clear communication. AI has a lot to offer here, with many good techniques for open communication of general and specific information, and clear understandings on many of the common and subtle pitfalls to be avoided.

In summary, the substantial progress made by AI on knowledge representation offers almost every field substantial shortcuts and useful tools for formalizing, standardizing, and representing the concepts and information of the field. AI should take the generally accepted core techniques it now possesses for structuring and organizing knowledge and advocate these techniques to other fields and the general scientific and informatics community. Failure to do this will mean an information infrastructure more fragile and less easily usable than might be otherwise. The advocacy of basic knowledge representation techniques offers an immediate way for AI to change the way other fields think about their own problems and results.