Knowledge Systems Laboratory
Technical Report KSL 89-87

November 1989

The Use of Formally-Represented Engineering Knowledge to Support Human Communication and Memory

by

Thomas R. Gruber

Presented at the 1990 AAAI Spring Symposium on Knowledge-based Human-Computer Communication, Stanford University, March 27-29, 1990.

KNOWLEDGE SYSTEMS LABORATORY
Computer Science Department
Stanford University
Stanford, California 94305

The Use of Formally-Represented Engineering Knowledge to Support Human Communication and Memory [1]

Thomas R. Gruber
Knowledge Systems Laboratory
Stanford University

What can Artificial Intelligence do for engineering? The standard answer is given without hesitation: “automate it!” This ambition underlies the paradigm of automated reasoning: that what AI can do is automate processes currently performed by thinking human beings, such as design. I question the dominance of this view, and propose to emphasize another goal for AI. AI's *unique* contribution to science is knowledge representation: techniques for stating human knowledge in a formal, explicit, and operational manner. The contribution is not domain knowledge itself or alternate representations of physical mechanisms, but representations of knowledge that underlies reasoning. Engineers don't need AI people to come up with new ideas to *do* engineering, and they don't need new representations to *automate* well-understood engineering procedures. They already know about mechanics and thermodynamics and analysis of beams. Engineers can program, too.

So what can AI representation technology bring to engineering? To address this question we ask another: what is knowledge representation good for? Again, the standard answer is narrow: “representations are used by intelligent agents to reason about things.” How about them intelligent agents! How about us? People use external representations all the time, as cognitive aids to memory and inference, and as media for communication. Engineers are no exception. In fact, engineers have plenty of useful ways of modeling the world of machines: representations of physical devices and how they work, such as differential equations, geometric drawings, and natural language. Of engineering representations, some are machine understandable (equations). For engineering tasks that can be modeled with these representations, such as kinematic analysis, tools have been built to automate the task. What engineering representations do not capture (and are not designed to represent) is the human knowledge with which devices are designed, manufactured, tested, operated, diagnosed, and repaired. Modeling this knowledge is the province of AI.

AI representations are appropriate for modeling the knowledge underlying human engineering practice. Importantly, these representations can be used for purposes other than automated engineering tasks. AI representations can be used as aides for memory and communication, where written and spoken natural language is used today. Knowledge bases can serve as personal and institutional memory, if made accessible to

---

[1] Position statement for the AAAI Symposium on Knowledge-based Human-Computer Communication, March 27-29, 1990, Stanford University. Available from the Stanford Knowledge System Laboratory as KSL Memo 89–87, November, 1989.

human users. The requirement that the knowledge be represented formally and explicitly can be valuable constraints for building shared models for human consumption. AI representations are also useful for modeling integrated human-machine systems, which are completely outside the expressive power of artifact-centered engineering descriptions. Furthermore, because the domain of AI representations is knowledge, they are useful for modeling the intellectual activity of model formation itself, such as the process by which engineers come up with equations to analyze a circuit.

Based on a vision of comprehensive, shared knowledge bases of engineering knowledge, we have begun a program of research aimed at developing representations for building human-scale models of physical devices and the ways the devices interact with the (human) environments of use. The project is called How Things Work, and it is a multi-year, collaborative effort involving AI researchers, engineers, and physicists from within the KSL and from other centers at Stanford. The goal is to build knowledge-based models about how engineered devices work, and how people and machines interact in a context of use. A key objective is to make the models that are both machine-understandable and human-usable.

We have identified classes of knowledge that would be useful to represent and that push on existing representation technology. For example, among the ways that people and machines interact is design. While low-level CAD systems allow engineers to work in a visual medium, the representation fails to capture the knowledge that goes into a design. Knowledge about a design is useful for communicating with other designers, and with people in downstream operations such as manufacturing, and for accumulating institutional memory on design expertise. One important kind of knowledge about design is *design rationale*: an explanation of why a particular piece of a device is designed as it is, in terms of its function in the context of use (i.e., its purpose) and design objectives and constraints that describe utility (e.g., cost, weight, manufacturability, etc.). To model design rationale requires representations for device structure (to describe the design to be justified), multiple levels of abstraction (to be able to identify the relevant level of description of a complex device model), teleology (to talk about the purpose of a structure, which is more than pointing to the resulting behavior), and utility (qualitative descriptions of the relative or absolute value of a design choice over alternatives).

Another type of knowledge to model is the process of *formulating* equation models for simulation, which is a task requiring engineering expertise. While much has been done to automate the use of equations for simulation, little is said about where the equations come from. Often they come from well-understood principles of physics applied to a partial design in not-so-well-understood ways. To capture model formulation knowledge we need to represent “first principles” and other engineering knowledge (e.g., laws of mechanics, energy conservation, etc.), the purposes of a simulation (what kind of question to answer by doing the simulation), properties of available components (information found on engineering spec sheets), and operational knowledge of algebra and calculus.

If the role of AI theory in engineering is to invent the representations for modeling human engineering knowledge, what is the role of automation? In keeping with the view outlined above, we are focusing on the tasks of model acquisition and model explanation.

Model acquisition can be supported by tools that 1) present the modeling technology to a user in a comprehensible and efficient interface, 2) provide feedback on the semantics of what has been said in the representation (e.g., static consistency and completeness analysis, and dynamic qualitative simulation of systems described by the models), and 3) generate more "compiled" models such as sets of equations from more complete and declarative models such as descriptions of device structure and the background knowledge mentioned above.

Model explanation is the description of the structure, behavior, or function (purpose) of some modeled system. Explanations of behavior given structure are presentations of the output of qualitative simulation. Machine-generated explanations reveal what has been represented in the machine, with human-readable presentations, such as hypertext and interactive diagrams showing causal pathways and change of state over time.

Many possible applications are enabled by developing AI knowledge representation technology for human consumption. Knowledge bases about how devices are built and used can serve as common models for computer-mediated dialogs among people distributed throughout an institution. The combination of qualitative simulation and explanation facilities would enable the production of focused simulation scenarios for training, interactive documentation, and group presentations. The ability to represent human-scale models at multiple levels of abstraction, with automated model formulation assistance, would enable design and test engineers to rapidly explore spaces of designs. It is intended that a long term effect of this technology would be to improve the overall productivity of American industry.