

Unification of Language Understanding, Device Comprehension and Knowledge Acquisition

Ashok Goel, Kavi Mahesh, Justin Peterson, and Kurt Eiselt
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280 USA
Contact: goel@cc.gatech.edu, 404-894-4994

Abstract

Cognitive agents often acquire knowledge of how devices work by reading a book. We describe a computational theory of understanding a natural language description of a device, comprehending how the device works, and acquiring a device model. The theory posits a complex interplay between language, memory, comprehension, problem-solving and learning faculties. Long-term memory contains cases of previously encountered devices and associated structure-behavior-function (SBF) models that explain how the known device works. Language processing is both bottom-up and top-down. Bottom-up processing is done through spreading-activation networks, where the semantics of the nodes and links in the network arises from the SBF ontology. The comprehension process constructs a SBF model for the new device by adapting the known device models - we call this process adaptive modeling. This multi-faculty computational theory is instantiated in an operational computer system called KA that (i) reads and understands English language descriptions of devices from David Macaulay's popular science book *The Way Things Work*, (ii) comprehends how the described device works, and (iii) acquires a SBF model for the device.

1. Motivations and Background

Cognitive agents often acquire knowledge of complex phenomena by reading a book. For example, a naive cognitive agent may acquire knowledge of how air-conditioners work by reading a popular science book such as *The Way Things Work* by David Macaulay [1988]. In general, understanding the natural language description of a device, comprehending how the device works, and acquiring a device model, involves a complex interplay between language, comprehension, memory, problem solving and learning processes. In addition, these processes use many different kinds of knowledge including semantic knowledge of the domain, episodic knowledge from past experiences in the domain, and the information provided in the text.

But most computational models of text interpretation deal with language understanding in vacuum, in more or less complete isolation from other processes. Typically, they either propose a largely bottom-up process in which the interpretation is constructed from the text alone, or a largely top-down process in which a precompiled knowledge structure helps to generate expectations and provides a template for filling in specific details given in the text. In interpreting real texts, however, neither the text always provides sufficient information to enable the construction of a satisfactory interpretation nor does the reader always have a precompiled knowledge

structure that matches the text. Our theory of language understanding for device comprehension and knowledge acquisition not only combines bottom-up and top-down strategies for language processing, but it also integrates the language process with memory, comprehension, problem-solving and learning processes.

In contrast to multi-strategy or multi-task theories, we call our theory *multi-faculty* because it unifies multiple cognitive faculties, not just multiple tasks or strategies within a specific cognitive faculty such as language. The multi-faculty theory is embodied in an operational, but still evolving, computer program called KA.

In [Pittges *et. al.* 1993], we described an early version of the KA system that unified language, memory and comprehension processes in the service of understanding a new design problem stated in English. We also showed how past problem-solving experiences retrieved from long-term memory enable the understanding of new problems. In [Peterson *et. al.* 1994], we described a new version of the KA system that not only integrated language, memory and comprehension processes but also unified them with problem solving. We also showed how problem solving helps to evaluate the output of the language, memory and comprehension processes. The above work grew out of our earlier theory of *adaptive design* in which new design problems are solved and new designs are constructed by adapting past design cases [Goel 1991a, 1991b].

In this article, we describe new work on the KA project that differs from and adds to earlier work in two aspects. Firstly, the input to KA now is not a description of a design problem, but an English language description of a device from the book "The Way Things Work." Secondly, the new version of KA not only unifies language, memory, comprehension, and problem-solving processes but also integrates learning with them. This new work grows out of an evolving theory of *adaptive modeling* in which comprehension of the workings of a system is represented and organized in the form of a structure-behavior-function (SBF) model, and SBF model of a new device is constructed by adapting old models of familiar devices [Goel 1991b, 1996].

Since we already have described the process of language understanding in KA in earlier papers, we will not repeat it here; [Peterson, Mahesh and Goel 1994] provides a detailed account. Instead, we (i) describe our framing of the problem of device comprehension as an abduction task, (ii) present a high-level account of the knowledge and strategies KA uses for addressing this task, and (iii) discuss how KA acquires a

2. Case Study: Comprehending the Fire Extinguisher

Let us consider the task of comprehending how a fire extinguisher works from the following description that appears on page 147 of *The Way Things Work*:

An extinguisher puts out a fire by excluding oxygen so that combustion (see p.154) can no longer continue. The extinguisher must smother the whole fire as quickly as possible, and therefore produces a powerful spray of water, foam, or powder. Some extinguishers produce a jet of carbon dioxide, a heavy gas that prevents burning. A fire extinguisher works in much the same way as a spray can. The extinguishing substance, such as water, is put under high pressure inside the extinguisher, and the pressure forces the substance out of the nozzle.

This text is accompanied with a cutaway diagram of a fire extinguisher revealing its structure, and some brief descriptions of the individual components such as the gas cartridge and the release valve. Figure 1 illustrates this diagram. Note that the annotations on the diagram are more specific to the structure of the extinguisher shown in the diagram than to any part of the text itself. The text describes the behavior of the fire extinguisher, making explicit reference to the descriptions of combustion and spray cans. The reference to the concept of combustion is a forward reference; presumably the reader has not yet read it but may do so for further elaboration and specification. But the spray can is described on the previous page of the book (p. 146), just opposite to the description of the fire extinguisher.

Framing the Comprehension Problem as an Abduction Task

The input to the task of comprehending how a fire extinguisher works in KA is constituted of three elements: the above text, the annotations on the accompanying diagram, and a symbolic representation of the diagram. The symbolic representation of the diagram constitutes a structural model which specifies only the structural elements and the topology of their connections in the fire extinguisher.

But what characterizes acceptable output of the task? We view the task of comprehending how a device works from a natural language description of the device as an instance of the very general abduction task. The abduction task takes a given set of data as input and gives a "best" explanation for the data as output [Josephson and Josephson 1994]. But now the question becomes what characterizes a best explanation?

The explanation of a device must not only specify the structural elements and the functions of the device, but it must also specify how the structure results in the functions. That is, it must specify how the device structure gives rise to the causal processes that result in the device functions. Thus we characterize a device explanation as a functional and causal model of the internal workings of the device. A best explanation of a device must satisfy three properties. First, the explanation must account for as much of the input as possible - ideally, it would cover the whole input. Second, the explanation must be consistent with the input. That is,




Figure 1: The Fire Extinguisher.

no element of the explanation can be inconsistent with any element in the input. Third, the explanation must be internally consistent. That is, no two elements of the explanation can be mutually inconsistent. In sum, in KA an acceptable output of the task of comprehending how a device works from a natural language description of the device is a functional and causal model of the working of the device that accounts for as much of the description as possible, is consistent with the entire description, and also is internally consistent.

3. KA at Work

Figure 2 illustrates the general functional architecture of KA. Here we only describe the processes linked by bold-faced arrows in the figure.

The long-term memory contains episodic knowledge of previously encountered devices. Each device case has an associated case-specific structure-behavior-function (SBF) model that explains how the device works [Goel 1991a, 1991b]. The SBF model of a device explicitly represents the structural elements and their configuration, the functions, and the internal behaviors of the device. Each behavior specifies a causal process in the device; the causal processes specify how the device structure results in its functions. In particular, they specify how the device functions are composed of the functions of the structural elements of the device. The SBF model for each device case is expressed in a common ontol-




Figure 2: The KA Architecture.

ogy that arises out of earlier work on device representations [Bylander and Chandrasekaran 1985; Sembugamoorthy and Chandrasekaran 1986; Chandrasekaran, Goel and Iwasaki 1993]. The SBF ontology defines the domain concepts and the relations between them, and constitutes the conceptual knowledge of the KA system.

The language process uses lexical and conceptual knowledge to generate cues for the memory process as well as preliminary interpretations for the comprehension process. Conceptual knowledge refers to knowledge of the domain concepts and the relations between them as characterized by the SBF ontology. The language process contains a large semantic network that takes the output of the parser as input and produces conceptual interpretations. The nodes and the links in the network are based on the SBF ontology of domain concepts and the relations between them. The spreading-activation mechanism in the network uses an early-commitment processing strategy with robust error-recovery to resolve word-sense ambiguities [Eiselt, 1987].

The mechanism resolves word-sense ambiguities by considering processing choices in parallel, selecting the alternative that is consistent with the current context, and deactivating but retaining the unchosen alternatives for as long as space and time resources permit. If some later context proves the initial decision to be incorrect, retained alternatives are reactivated without reaccessing the lexicon or reprocessing the text. [Peterson, Mahesh and Goel 1994] provides details of language processing in KA.

The memory process uses cues generated by language process as probes into the long-term memory. It accesses device cases and associated case-specific SBF device models and puts them into a working memory for use by the comprehension and problem-solving processes. The memory process also stores newly learned models in the long-term memory. The device cases are indexed by the functions of the stored devices; the SBF models are indexed by the cases. This indexing scheme is borrowed from our earlier work on adaptive design [Goel 1991a, 1991b].

The comprehension process constructs a SBF model for the new device by adapting the SBF device models accessed by the memory process. It uses generic (abstract, skeletal) modification plans for the task of adapting SBF models of known devices to construct a model of the new device. The selection of relevant modification plans is based on the functional and structural differences between the SBF model of the known device and the description of the new device.

KA's method constructing the new SBF model is identical to that of adaptive modeling [Goel 1991b, 1996].

The learning process uses the SBF model of the new device to learn appropriate indices for storing the model in the long-term memory. The new indices depend both on the contents and organization of the memory and the functional and causal explanation provided by the SBF model. Again, KA's method of index learning is identical to that of adaptive modeling [Bhatta and Goel 1995]. (The problem-solving process in Figure 2 plays no direct role in this process of acquiring a SBF model of a new device from an English language description.)

The Case Study

Let us consider the comprehension process in the case study of the fire extinguisher. At this stage of processing, KA's working memory contains three elements: the interpretations of the sentences in the text generated by the language process, the SBF model of the spray can retrieved from the long-term memory, and the symbolically-represented structural model of the fire extinguisher given as part of the input to the system. The current task is to adapt the SBF model of the spray can to construct a model of the fire extinguisher.

The comprehension process notes the sentence interpretations in the working memory and the differences between these interpretations and the SBF model of the spray can. For example, it notes that the fire extinguisher and the spray can contain different substances under pressure. In addition, the comprehension process notes the structural differences between structural models of the fire extinguisher and the spray can. For example, it notes that the two contain different kinds of nozzles. It uses these differences to select generic model-modification plans that help to reduce specific differences and are indexed by the differences they can help to reduce. Examples of model-modification plans include the substance-substitution plan and the component-replacement plan. Given a specific difference between a component in the spray can and a component in the fire extinguisher, instantiating the latter plan in the context of the SBF model of spray can, for example, results in replacing each occurrence of the spray-can component in the SBF model by the corresponding component in the fire extinguisher. The application of this plan also results in the propagation of the causal effects of the new component. The structural models help to establish correspondences between the components in the spray can and the fire extinguisher. The SBF model of the spray can helps to focus the process of plan instantiation and application. The invocation and application of selected model-modification plans, one for each difference between the fire extinguisher and the spray can, results in the generation of a SBF model for the fire extinguisher. This preliminary SBF model provides a functional and causal explanation of the working of the fire extinguisher.

Next, KA evaluates the preliminary model of the fire extinguisher for both internal and external consistency. In reference to internal consistency, the comprehension process makes sure that no new element introduced into the SBF model is inconsistent with any other element. This is done by systematically tracing through the causal behaviors of the new SBF model. If an inconsistent element is detected, then the process retracts the corresponding modification from

the SBF model. In reference to external consistency, the comprehension process makes sure that no element in the SBF model is inconsistent with the output of the language process. This is done by cross-checking of the SBF model and the sentence interpretations generated by the language process. Again, if an inconsistent element is detected, then the comprehension process retracts it from the SBF model. The final SBF model is KA's best explanation of the working of the fire extinguisher.

4. Relations

Our work on KA builds on many lines of research in cognitive science including natural language understanding, device comprehension, knowledge acquisition, mental models and model-based reasoning, case-based reasoning and learning, and abductive explanations. Due to limitations of space, however, here we only outline its relationship to earlier work that lies at the intersection of language understanding, device comprehension, and acquisition of device models.

Lebowitz's [1983] RESEARCHER program read natural-language texts in the form of patent abstracts, specifically disk drive patents, and updated its long-term memory with generalizations made from these texts. Its knowledge representation scheme was oriented toward device objects and their structural relationships, which was a departure from most natural language understanding systems of that time which had typically focused on intentional actors and events. The output of the processing was a generalized representation in the form of a structural model of the disk drive which specified its components and the topological relationships among them. The system stored this structural model in its long-term memory and later used this knowledge to aid in the top-down understanding of additional patent texts. However, RESEARCHER's emphasis on components and structural relationships left it unable to build functional and causal models of the mechanisms described. In other words, the system effectively knew how a disk drive was constructed, but it did not know how it worked. In sharp contrast, KA takes a structural model of the new device as part of its input.

Dyer, Hodges, and Flowers [1987] describe EDCA, a conceptual analyzer which serves as a natural language front-end for EDISON, a naive design problem solver. EDCA uses knowledge of the function of physical devices to produce an episodic description of a device's behavior as described by an input text. This episodic description can then be used to generate a new device model to be integrated into long-term memory. The result is a much more comprehensive understanding of the device's functionality than was possible with RESEARCHER, but EDCA's analysis of the device description is not fully integrated with the processes for generating new device models and incorporating them into memory. EDCA, in other words, is but a front end to EDISON.

As Selfridge [1989] notes, separating the process of analyzing the language input from constructing and incorporating the new model is misguided --- the process of understanding a device description is the process of constructing a functional and causal model of that device. This is the approach that we have followed in our work on KA. We believe that this approach enables KA to correct the shortcomings of both

5. Discussion

KA is a computational theory of a complex cognitive phenomenon. From the viewpoint of cognitive science, one of the major advantages of building complex and elaborate, yet detailed and precise, computational theories such as KA is the identification of interesting interactions among the different processes. At the start of the KA project, we enumerated a set of ten high-level hypotheses about these interactions [Goel and Eiselt 1991]: (i) understanding natural language descriptions of physical devices enables acquisition of device models, (ii) situating language processing in problem solving identifies the meaning of the "meaning" of a device description, (iii) past cases and case-specific models, that originally provided the knowledge structures for addressing a class of design problems, also provide the knowledge structures for language processing, (iv) the SBF language for representing device models, originally developed to address design problems, provides the conceptual knowledge needed for text interpretation, (v) the model-based scheme for indexing the stored cases and case-specific models in long-term memory, again originally developed to address design problems, is appropriate for supporting language processing, (vi) the language process generates adequate cues for probing the long-term memory, (vii) the memory process retrieves relevant cases and associated models from the long-term memory into the working memory, (viii) the retrieved case-specific models act as expectation generators, (ix) the model-based expectations guide the language process, and (x) the language process generates adequate cues for guiding the comprehension process in adapting the retrieved models to construct a model for the new device.

Now at the end of this project, we can confidently assert that the KA theory helps to greatly refine these hypotheses, to make them more precise and explicit. We conclude this article with a brief discussion of how the KA theory has helped to refine the last of the ten hypotheses above because this initially surprised us. We found that language processing provides only limited guidance to the comprehension process in adapting the SBF model of a known device (e.g., the spray can) to construct a model of the new device (e.g., the fire extinguisher). The products of the language process do indicate some of the many differences between the two devices. But most of the important differences come from the structural models of the two devices. Also, the text does enable limited verification of the modified model to insure that the new model is consistent with the text. But we were initially surprised to find that language processing does not clearly indicate the precise content and form of the new device model. There are two apparent explanations for this. First, the device descriptions in Macaulay's *The Way Things Work* are coarse-grained while our SBF models, which need to support multiple reasoning processes, are fine-grained. This might be resulting in a mismatch between the text and the model so that text can provide only limited help in adapting the model. Second, the diagram that accompanies the textual description of a device is given to KA in the form of a symbolically represented structural model of the device. This might be resulting in some loss of information. Perhaps more

importantly, this may imply that the visual process, and not the language process, might be especially important for model adaptation and construction.

Acknowledgements

Kavi Mahesh now is at the New Mexico State University and Justin Peterson is with CitiBank. This work was done when all four authors were at Georgia Tech, where it was supported by Northern Telecom (research gift), the National Science Foundation (research grant IRI-92-10925), and the Office of Naval Research (research contract N00014-92-J-1234). Jeff Pittges contributed to the programming of an earlier version of the KA system described here.

References

- [Bhatta and Goel 1995] S. Bhatta and A. Goel. Model-Based Indexing and Index Learning in Analogical Design. In *Proc. 1995 Seventeenth Annual Conference of the Cognitive Science Society*, Pittsburgh, July 22-25, 1995, NJ, Hillsdale: Erlbaum.
- [Bylander and Chandrasekaran 1985] T. Bylander and B. Chandrasekaran. Understanding Behavior Using Consolidation. *Proc. Ninth International Joint Conference on Artificial Intelligence*, pp. 450-454, 1985, Los Altos, CA: Morgan Kaufman.
- [Chandrasekaran, Goel and Iwasaki 1993] B. Chandrasekaran, A. Goel, and Y. Iwasaki. Functional Representation as a Basis for Design Rationale. *IEEE Computer*, 26(1):48-56, January 1993.
- [Dyer, Hodges and Flowers 1987] M. Dyer, J. Hodges and M. Flowers. Computer Comprehension of Mechanical Device Descriptions. Technical Report UCLA-AI-87-7, University of California, Los Angeles: Artificial Intelligence Laboratory.
- [Eiselt 1987] K. Eiselt. Recovering from Erroneous Inferences. In *Proc. Sixth National Conference on Artificial Intelligence*, pages 540-544, Los Altos, CA: Morgan Kaufman.
- [Goel 1991a] A. Goel. A Model-Based Approach to Case Adaptation. In *Proc. Thirteenth Annual Conference of the Cognitive Science Society*, Chicago, August 1991, pp. 143-148, Hillsdale, NJ: Lawrence Erlbaum.
- [Goel 1991b] A. Goel. Model Revision: A Theory of Incremental Model Learning. In *Proc. Eighth International Conference on Machine Learning*, Chicago, June 1991, pp. 605-609, Los Altos, CA: Morgan Kaufmann.
- [Goel 1996] A. Goel. Adaptive Modeling. To appear in *Proc. Tenth International Qualitative Reasoning Workshop*, May 1996, Menlo Park, CA: AAAI Press.
- [Goel and Eiselt 1991] A. Goel and K. Eiselt. Mental Models, Natural Language, and Knowledge Acquisition. *ACM SIGART, Special Issue on Integrated Cognitive Architectures*, 2(4): 75:78, August 1991.
- [Josephson and Josephson 1994] J. Josephson and S. Josephson (editors). *Abductive Inference: Computation, Technology and Philosophy*. New York: Cambridge University Press.
- [Lebowitz 1983] M. Lebowitz. RESEARCHER: An Overview. In *Proc. National Conference on Artificial Intelligence*, pages 232-235, Los Altos: Morgan Kaufman.
- [Macaulay 1988] D. Macaulay, *The Way Things Work*, Boston: Houghton Mifflin Company.
- [Pittges et al., 1993] J. Pittges, K. Eiselt, A. Goel, A. Gomez, K. Mahesh, and J. Peterson. KA: Integrating Natural Language Processing and Problem Solving. In *Proc. Fifteenth Annual Conference of the Cognitive Science Society*, Boulder, Colorado, July 1993, pp. 818-823, Hillsdale, NJ: Lawrence Erlbaum.

- [Peterson *et. al.* 1994] J. Peterson, K. Mahesh, A. Goel, and K. Eisele. KA: Situating Natural Language Processing in Design Problem Solving. In *Proc. Sixteenth Annual Conference of the Cognitive Science Society*, August 1994, Atlanta, Georgia, pp. 711-716, Hillsdale, NJ: Lawrence Erlbaum.
- [Peterson, Mahesh and Goel 1994] J. Peterson, K. Mahesh and A. Goel. Situating Natural Language Understanding in Experience-Based Design. *International Journal of Human-Computer Studies*, 41: 881-913, 1994.
- [Selfridge 1989] M. Selfridge. Toward a natural language-based causal model acquisition system. In *Causal AI Models: Steps Toward Applications*, Werner Horn (editor), pages 107-128, Hemisphere Publishing Corporation.
- [Sembugamoorthy and Chandrasekaran 1986] V. Sembugamoorthy and B. Chandrasekaran. Functional Representation of Devices and Compilation of Diagnostic Problem Solving Systems. In *Experience, Memory and Reasoning*, Kolodner and Riesbeck (editors), Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1986.