Modeling Invention by Analogy in ACT-R

J. William Murdock, Marin Simina, Jim Davies, Gordon Shippey
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
{murdock, marin, jimmyd, shippey}@cc.gatech.edu

Abstract

We investigate some aspects of cognition involved in invention, more precisely in the invention of the telephone by Alexander Graham Bell. We propose the use of the Structure-Behavior-Function (SBF) language for the representation of invention knowledge; we claim that because SBF has been shown to support a wide range of reasoning about physical devices, it constitutes a plausible account of how an inventor might represent knowledge of an invention. We further propose the use of the ACT-R architecture for the implementation of this model. ACT-R has been shown to very precisely model a wide range of human cognition. We draw upon the architecture for execution of productions and matching of declarative knowledge through spreading activation. Thus we present a model which combines the well-established cognitive validity of ACT-R with the powerful, specialized model-based reasoning methods facilitated by SBF.

Introduction

The cognition involved in invention may be among the rarest and most difficult to understand and model. The processes involved often take place over the course of months or even years and thus are not generally amenable to extremely fine-grained modeling, and are difficult to study experimentally. Furthermore, great inventors are rare, and the recording of their discoveries and inventions are often sketchy and incomplete. But this is not the case with Alexander Graham Bell's invention of the telephone, which is probably one of the best documented and analyzed inventions (e.g., Notebook[1]; US v. Bell, 1908; Gorman, 1997).

Let's consider the invention of the telephone's microphone. During a famous experiment on June 2, 1875, Bell realized that electromagnetic induction could be practically used for the purposes of speech transmission. By moving a magnet close to a coil, an electrical current with an amplitude proportional to the movement of the magnet could be generated in the coil circuit. Consequently, by moving the magnet in the same way as the air pressure generated by speech, speech could be transmitted electrically at distance. This is Bell's account of the invention:

> The problem that then arose in my mind was, how to move a piece of steel in the way that the air was moved by the action of the voice. *While this problem was in my mind, I was carrying on experiments with the phonautograph constructed from the human ear*... and it occurred to me that if such a thin and delicate membrane could move bones that were, relative to it, very massive indeed, why should not a larger and stouter membrane be able to move a piece of steel in the manner I desired? At once the conception of a membrane speaking telephone became complete in my mind; for I saw that a similar instrument to that used as a transmitter could also be employed as a receiver (US v. Bell, 1908).

Unfortunately, this short paragraph does not emphasize enough the *rich* description of the phonautograph in Bell's memory, nor its *level of activation* compared with other devices, which have not been relevant or reasoned about recently. But these aspects are crucial for cognitive modeling. According to Bell's account (US v. Bell, 1908), as a "teacher of the deaf", he was interested in a device able to transform the human voice in a visual form, in order to provide visual feedback to deaf students learning to speak. He started by experimenting with Leon Scott's phonautograph, which unfortunately was not sensitive enough for his purposes. While trying to improve the phonautograph design, Bell was struck by the similarity between the mechanism of the phonautograph and the mechanism of the human ear. He came to the conclusion that a phonautograph modeled after the pattern of the human ear would probably produce more accurate visual feedback, and he even built such a device. Consequently, during the quest for the microphone, Bell had a *rich* and *recent* description of the phonautograph. This afforded subtle priming effects, resulting in the retrieval of the phonautograph, while looking for a device able to move a piece of steel in the same way as the human voice. We suspect, and want to emphasize, that the retrieval of the phonautograph was not a classical long-term memory (LTM) retrieval, but a fine-grain working-memory (WM) retrieval from the "activated" part of the memory. Namely, the retrieval of the phonautograph probably was not based on its function, but on its internal behavior (Simina & Kolodner, 1995). Note that the retrieval of the phonautograph, as a source, is a prerequisite for analogy.

Let's also consider Bell's description of the phonautograph:

> The phonautograph *consisted* essentially of a stretched membrane, in front of which was arranged a conical mouth-piece, into which a person could speak. A long and light lever of wood was hinged to one edge of the membrane support, and the other end projected horizontally beyond the opposite edge, and carried at its extremity a vertical bristle ...
>
> *The operation of the instrument was as follows:* A person uttered a sound into the cone of the phonautograph, thereby causing the vibration of the membrane and its attached lever, and the bristle at the end of the lever was forced to partake of the motion (US v. Bell, 1908)...

The first paragraph of the above description emphasizes the *structure* of the phonautograph, while the second paragraph emphasizes its *behavior*. From this and other similar

[1] Alexander Graham Bell's Notebooks are available on the WWW at: *http://jefferson.village.virginia.edu/~meg3c/id/albell/homepage.html*

examples, we argue that Bell's reasoning can be modeled used the Structure-Behavior-Function (SBF) device representation (Goel et al., 1997) to reason about physical devices.

Elsewhere we have proposed a full computational architecture for modeling creative design, and in particular Bell's inventions (Simina and Kolodner, 1997; Simina et al., 1998), by relying on the SBF representation of artifacts. But to investigate, as closely as possible to human cognition, the memory issues during the invention of the microphone, we also decided to experiment with the rich memory mechanisms provided by ACT-R (Anderson, 1993) and let ACT-R's architecture guide the memory processing involved in the invention of the microphone.[2]

Figure 1 illustrates our analysis of Bell's problem solving activity in this instance. We have developed a cognitive model which instantiates this analysis. Our model gets a function specification of transforming speech (i.e., a substance[3] of type `sound`) into an electric current (i.e., a substance of type `electricity`). It remembers the device implementing electromagnetic induction, and it has to solve a simplified problem, namely transforming speech into a mechanical vibration located in a magnet. Next, the model considers a possible synergy with other information which may be active in working memory. In this analysis, we argue that design of the phonautograph is active in WM, affording priming effects, so the model can perform a fine-grained match (including internal behavioral states) between the evolving specification of the microphone and the ear phonautograph.

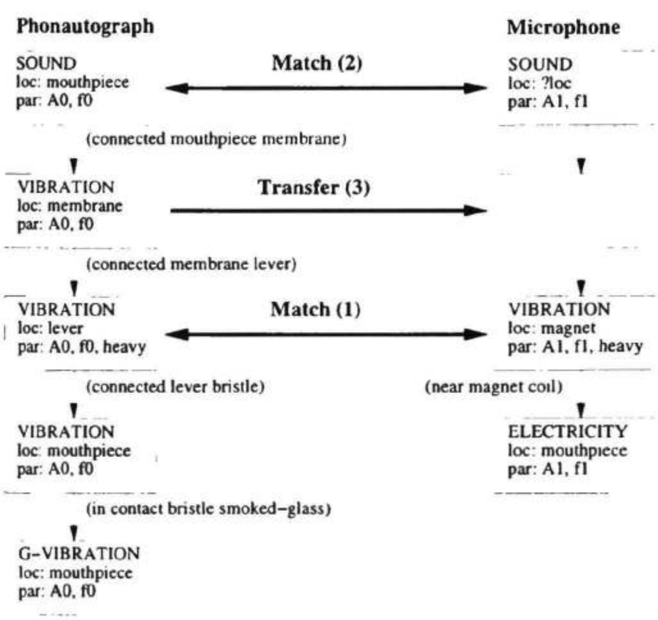

Figure 1: Analogy between phonautograph & microphone

We have developed a computational system in ACT-R which instantiates this model. The system begins operation with a nearly complete design of a microphone which is only missing a single piece, the one which transforms sound to vibration. It completes the design by performing the following operations:

- The system analyzes the relevant pieces of the existing microphone model to determine the specifications of a desired analogue.
- It retrieves a reference to a piece of the internal behavior of the phonautograph which matches this specification. Specifically, it retrieves the transition from sound to vibration.
- Finally, the system performs analogical transfer between the two models, i.e., it takes the relevant details of the transition from sound to vibration in the phonautograph and instantiates it in the microphone.

We will present the execution of this system in greater detail in a later section of this paper.

SBF Models

In order to represent the devices and components involved in this model, we chose SBF (Structure-Behavior-Function) models. One issue involved in selecting a representation within a cognitive model is the issue of whether the representation is both adequate and effective for the kinds of reasoning that humans do. It would be extremely easy to design an ad hoc modeling language which supported the reasoning capabilities displayed in this project. However, because a particular task such as this specific form of analogy imposes so few constraints on the nature of the representation, we would have very little confidence that such a language would support a plausible general model of a human designer. In contrast, the SBF language has been shown to directly support diverse and challenging design capabilities. Among these are case-based design (e.g., Kritik and Kritik2, Goel 1989; Goel et al., 1997) analogical transfer using generic abstractions (e.g., IDeAL, Bhatta et al., 1994), natural language understanding in the domain of device design (e.g., KA, Goel et al., 1996b) explanation of design process and products (e.g., Interactive Kritik, Goel et al., 1996a), creative design (e.g., ALEC, Simina & Kolodner 1997) and conceptual change in scientific theory formation (e.g., ToRQUE, Griffith et al., 1996). This is clearly a broad and interesting set of domains and problems. Thus we argue that SBF languages provide a plausible account of design reasoning, and we consider it to be a close approximation of an abstract form of human knowledge.

As the name suggests, the SBF modeling language breaks devices down into three aspects: their structure, their behaviors and their functions. The structure of a device describes the physical components of the device and how they are interconnected. Behavior consists of a causal model of the dynamic interactions between these structural elements, in the form of a sequence of behavioral states and transitions between these states. The function of the system describes the overall intensional consequences of the behavior.

Analogical matchings from one concept to another are made on the basis of spreading activation through elements of the system's behavioral properties. In the case of the phonautograph, the function is to visually record auditory information. The purpose of the microphone is to transform auditory information into electronic pulses. These functions are not sufficiently similar to support the detailed level of analogy performed in this project, and they are not used by this system. The actual analogical reminding which is done in this project is at the level of behaviors. When looking at the behavior, independent of what the function of the object is, there are more direct correspondences which define the actual

[2]The code for this cognitive model can be found at http://www.cc.gatech.edu/grads/m/Bill.Murdock/bell/

[3]The term "substance" here may be somewhat misleading. As used in the SBF language and a variety of other design modeling languages, this term refers not only to physical substances such as liquids flowing through a pipe, but also to abstract phenomena such as energy or force which can be seen as flowing through a device.

recognition of similarity done by this model (and we argue, were actually described by Bell in his journal). So Bell knew he wanted a stick to be moving up and down in accordance with the vibrations created by the voice. When considering the behavior of the phonautograph, it is apparent that a bone is moving in the same way that the stick is supposed to be moving. At this point an analogy is made between the bone and the stick, and the microphone is invented.

Theory Development/Rationale

We used the ACT-R cognitive model in order to try to replicate one of the critical steps of Bell's discovery. Specifically, we wanted to model the analogy between the phonautograph and Bell's target concept, the microphone.

Our first attempt at representing Bell's analogy led us into some stylisticly questionable constructs. We found it difficult to express our ideas within the constraints of the ACT-R theory, so some of our code relied on calls to the layer underneath ACT-R: LISP. We also found ourselves re-implementing LISP functionality in ACT-R itself. In particular we tried to search for relevant memory elements by doing a linear search across the relevant portions of our model using a mechanism similar to the member function in LISP. The matching itself in this model was done using hard coded calls to LISP within the production itself. It was relatively evident that this first cut model did not benefit from the research which went in to the development of the ACT-R cognitive model because it so thoroughly circumvented the relevant features of ACT-R.

We tried to remove as much of our "machinery" as possible, and let ACT-R do the work its way. Because ACT-R has built-in mechanisms for analogy, we considered the possibility of applying them to our problem. However we immediately realized that the ACT-R analogy system would not serve us. ACT-R draws analogies in the form of directed generalizations over inferencing mechanisms. While this is clearly one important kind of analogy, we do not believe that the protocol (described in the introductory section) suggests that this is the kind of analogy that Bell is doing in this problem. We argue that the analogy in this problem is in the form of the transfer of a small portion of a complex, *declarative* knowledge item in to another such item. In the language of ACT-R, we are examining analogy over interconnected collections of chunks while the built in analogy mechanism is interested in analogy over individual productions.

In fact, the particular task that we are addressing is very similar to the task addressed in SME (Gentner, 1983). SME also performs an analogical transfer of small portions of complex interrelated knowledge elements. It is clear that one could combine a spreading activation retrieval mechanism (such as the one which we propose later in this paper) with an SME-like analogical transfer mechanism to accomplish the results described in this paper. We do not do this. The mechanism which we use to do analogical transfer is specifically tailored to the particular representation language used in this project: the SBF language. SME is inherently a weak method; it operates completely independently of the content of the knowledge which it is performing analogy over. This makes it an extremely general and powerful method. However, this approach is an implausible account of how an *expert* inventor such as Bell reasons about devices.[4] We argue that the kinds of reasoning mechanisms (i.e., productions) that Bell used were specialized for reasoning about devices. In the system described here, specialized model-based productions are used providing a strong method of analogical transfer.

Device Representation

The language used to represent the phonautograph and the microphone combines the content of the SBF language for device representation with the syntactic mechanisms in ACT-R for specifying declarative memory elements. A device is defined as an element which contains a *structure*, a *behavior*, and a *function*. The definition of this type and an example (the phonautograph, abbreviated *pag*) follows:

```
(chunk-type device function structure
  behavior)
(add-dm
  (pag ISA device
           function pag-function
           structure pag-structure
           behavior pag-behavior))
```

The function, structure, and behavior slots are, in turn, filled by complex memory elements which in turn contain complex lower level elements. A complete description of SBF is beyond the scope of this paper.[5] For the purposes of this system, it is important to note that SBF provides a causal model of the operation of a device as a set of *states* and *state transitions* between them. Transitions describe either the movement of a substance from one location to another or the conversion of one substance into another. A transition is defined by the starting and ending state as well as numerous relevant properties such as conditions under which the transition can occur, etc. A state is defined by the adjacent states and transitions as well as a *substance state schema*, i.e., the presence and nature of a particular substance at a state. A substance state schema refers to a particular instantiation of a substance at a particular location. That instantiation further refers to the *generic substance* which it instantiates. The definitions of the memory types directly relevant to representing transitions are:

```
(chunk-type state-transition context
  previous-state next-state by-behavior
  using-function as-per-domain-principle
  parameter-relations condition)
(chunk-type state previous next enabled-by
  enabling substance-state-schema
  component-state-schema)
(chunk-type substance-state-schema location
  main-substance contained-substances)
(chunk-type substance is-a property-list)
(chunk-type generic-substance is-a)
```

An example of a transition is the first transition in the phonautograph model. The memory elements which define this transition are:

```
(add-dm
  (pag-t1 ISA state-transition
```

[4] Note that ACT-R does have built in learning capabilities for compiling applications of generally useful production rules into more specialized production rules tuned to specific situations. It seems possible that such capabilities might be able to transform a weak analogical transfer mechanism like SME into a strong, domain dependent analogical transfer mechanism like the one presented here. If this hypothesis is correct, this would provide an interesting account of a novice to expert transition in the domain of invention. Considerably more work would be needed to explore this idea.

[5] For such a description, see Goel et al. (1997).

```
        previous-state pag-s1
        next-state pag-s2
 using-function pag-cf1)

 (pag-s1 ISA state
        next pag-s2
        enabling pag-t1
        substance-state-schema pag-sss1)
 (pag-sss1 ISA substance-state-schema
          location membrane1
          main-substance sound1)
 (sound1 ISA substance is-a sound)

 (pag-s2 ISA state
        previous pag-s1
        next pag-s3
        enabled-by pag-t1
        enabling pag-t2
        substance-state-schema pag-sss2)
 (pag-sss2 ISA substance-state-schema
          location bone1
          main-substance vibration1)
 (vibration1 ISA substance is-a vibration)
```

These memory elements describe a transition, `pag-t1`, which converts sound at a membrane (`pag-s1`) to vibration in a bone (`pag-s2`) via the functionality of the membrane (`pag-cf1`). This is the transition from which the analogy to the microphone problem is drawn, as described in the next section.

Mechanism

Our model can be loosely divided into a sequence of stages which describe the overall process of recognition and application of the analogy between the microphone and the phonautograph. The input to this mechanism is a nearly complete model of a microphone which is only missing a particular component which is specified as being capable of transforming sound into vibration within a steel stick. The output of the mechanism is a complete model of the microphone in which a membrane is used to do this transformation.

Lookup

The first phase of this model is the analysis of the original microphone model in order to extract the particular features which are relevant to this problem. At this phase of the process, the system is only considering the microphone problem which it is currently working on; the current goal refers to the completion of the design of the microphone, where the particular state transition in the microphone is unknown. At this point there is no reference to the phonautograph. The result of the lookup mechanism is a retrieval cue for a related analog. This cue takes the form of a goal.

Our lookup mechanism is divided into two production rules. The first of these rules (`lookup-substance-in`) generates a cue for a state transition in an existing device whose input substance is the same as the input of the desired transition. The second (`lookup-substance-out`) does the same for the transition's output.

```
(P lookup-substance-in
 =goal> ISA goal-retrieve-component
    transition =transition
    generic-substance-in nil
  =transition> ISA state-transition
    previous-state =state
  =state> ISA state
    substance-state-schema =sss
  =sss> ISA substance-state-schema
    main-substance =substance
  =substance> ISA substance
    is-a =gen-substance
 ==>
  =goal>
    generic-substance-in =gen-substance)
```

The production (`lookup-substance-in`) finds a generic substance that corresponds to a given transition. In the telephone example, this model involved is the microphone and the particular transition being considered is the transition from sound to vibration. The production identifies the state which acts as input to the transition in the variable `=state`. In our example, this input state describes the presence and nature of the sound waves located in the preceding component. From this state is extracted the generic substance (`=gen-substance`); in the example, this generic substance is sound. Finally, a goal is generated to find a similar transition in a device which also takes this generic substance.

An important feature to observe about this production is that (unlike the retrieval production presented in the next section) it does *not* rely on spreading activation to mediate between several possible relevant memory elements. The transition type in question specifies exactly one possible binding for the variables because the transition contains a direct link to the state, which contains a direct link to a substance state schema, which specifies a single substance. Thus this production is simply analyzing a knowledge structure which is already the focus of attention rather than identifying something new to focus on.

Retrieval / Matching

Once the initial analysis of the microphone has been completed, the system needs to be able to find a related example in its memory with which it can make an analogy. One mechanism for doing retrieval and matching in a production system is to retrieve a memory that matches all of the desired features. However, in our system we have divided retrieval and matching into two separate steps. In the first of these steps, we retrieve transitions based on their similarity to *either* the input or the output of the transition being considered. In the second step, we then verify that the transition retrieved matches *both* the input and the output. This mechanism is more consistent with our understanding of complex memory retrieval in humans; many elements are retrieved based on partial relevance and then these elements are filtered through and only the fully (or almost fully) relevant ones are actually further considered. Below is one of the transitions for the retrieval step:

```
(P retrieve-substance-in
  =goal> ISA goal-retrieve-component
    generic-substance-in =gen-substance
    generic-substance-out =gen-substance-out
    retrieved-transition nil
  =transition> ISA state-transition
    previous-state =state
    - context =goal
  =goal> ISA
    goal-retrieve-component
    - transition =transition
  =state> ISA state
    substance-state-schema =sss
  =sss> ISA substance-state-schema
    main-substance =substance
  =substance> ISA substance
    is-a =gen-substance
 ==>
  =transition> context =goal
  =newgoal> ISA goal-check-transition
    generic-substance-in =gen-substance
    generic-substance-out =gen-substance-out
```

```
    matches-in =gen-substance
    transition =transition
    retrieved-transition =rt
  !push! =newgoal
  =goal>
    retrieved-transition =rt)
```

There is a corresponding `retrieve-substance-out` production as well. This production rule retrieves a transition which matches the substance input of the desired transition (there is a comparable production for retrieving transitions based on the output state). The goal for this production is that of the retrieval of an analogue for which the relevant generic substances have been identified (i.e., the lookup mechanism has already been applied). The `=transition` variable can bind to any transition in memory whose input state has the generic substance sought; in the example, it binds to the transition in the phonautograph in which the membrane converts sound into vibration.

Once such a transition is retrieved, a new goal is pushed onto the goal stack. This new goal is picked up by the production `match-substance-in` which checks to see if the transition matches the desired transition; the `matches-in` slot is filled to denote that these transitions do match with regards to their input (since that is how the transition was retrieved). Note that the presence of the `context` slot in the transition is used to prevent the rule from retrieving the same transition multiple times; once the transition has been retrieved it is marked as being identified in the context of the current goal.

In a scaled up version of this model which actually contained a very large number of models, this mechanism would become overwhelmingly cumbersome and would not *generally* support tractable retrieval. In particular, it would be expected that these models would typically not have a high enough activation level to be retrieved unless there were many strong links to the model being considered (i.e., the analogy was extremely "obvious"). We feel that this is a core theoretical commitment of our model; the recognition of analogical similarity which is being done in this particular case is *not* consistent with a complete search of all of memory. In particular, we note that in the protocol, it is stated that Bell was already thinking about and working on experiments with the phonautograph which were being driven by a detailed consideration of how the phonautograph actually worked. In the language of ACT-R, it is apparent that *the declarative representation of the behavioral elements of the phonautograph were highly activated in Bell's memory at the time of the experiment.* Thus, we argue, the generally intractable retrieval mechanism used in this model is specialized to retrieving highly activated and immediately relevant items in memory. In other words, the retrieval mechanism presented here is an *opportunistic* model of analogue retrieval, i.e., one which operates only when a particular opportunity for matching arises due to the conjunction of relevant activated memory elements.

Once elements are retrieved, they must be matched to ensure that both the inputs and the outputs are relevant to the transition being analyzed. As noted above, the production which retrieves the substance based on the input, obviously guarantees that the inputs of the transitions match. Thus that retrieval production feeds into this matching production which ensures that the outputs also match:

```
(P match-substance-out
  =goal> ISA goal-check-transition
    matches-in =gen-substance-in
    generic-substance-out =gen-substance
    transition =transition
  =transition> ISA state-transition
    next-state =state
  =sss> ISA substance-state-schema
    main-substance =substance
  =substance> ISA substance
    is-a =gen-substance
==>
  =goal>
    retrieved-transition =transition
  !pop!)
```

This transition takes as a goal a request to check a transition which has already asserted that the input substances match, as generated by the retrieval production above. The goal binds the variable `=gen-substance` to the generic substance desired as output. It checks the transition being analyzed to confirm that this generic substance is indeed the output of the transition. If this check succeeds, the transition is confirmed to have matched and the goal to check the transition is popped off the goal stack.

Placement / Specification

The actual adaptation of the matched elements is also done using strong model-based methods. A new component is generated to account for the transformation between sound and vibration. This component is placed in the particular portion of the behavior where this transformation takes place. Furthermore, the particular links between this component and other aspects of the model are specified using productions which are specific to the SBF language. At this point the new model has been fully defined. The initial design goal is satisfied and the system terminates.

The placement and specification portions of the system consist of a number of closely interrelated productions which create the new memory elements for the modified model, specify links from the existing memory elements to the new elements, and specify links from the new elements to the existing elements. One example of these productions follows:

```
(P place-component
  =goal> ISA goal-design
    retrieved-component =component
    transition =transition
    new-component nil
  =component> ISA component
    is-a =gen-component
  =transition> ISA state-transition
    previous-state =state
  =state> ISA state
    substance-state-schema =sss
  =sss> ISA substance-state-schema
    location nil
==>
  =new-component> ISA component
    is-a =gen-component
  =sss>
    location =new-component
  =goal>
    new-component =new-component)
```

This production is responsible for generating a new component. In our example, it creates a membrane analogous to the vibrating membrane in the phonautograph. The activation of this production requires that an existing component has been retrieved (i.e., that the `retrieved-component` slot in the goal is filled), that a state transition requiring a similar component has been identified (i.e., that the `transition` slot is filled), but no component has been placed in the model being designed (i.e., that the `new-component` slot has not yet been filled). The result of the production is that a new component is created and the location from which the transition takes place is identified as that component. Later rules fill in values for this component, identifying its function as analogous to the function of the existing component (the one from the phonautograph).

Discussion

Analogy is an important technique in invention (and many other tasks). However, in a given situation there may be an enormous number of potential candidates for a mapping, each of which is extremely complex. In a purely sequential retrieval system, checking all of these candidates is clearly intractable. Even when retrieval is done in parallel by spreading activation, a search of all of memory for all possible matches will generally be unsuccessful because activation is spread too thin (and thus no individual element is found) or the number of elements retrieved is enormous (and thus it is intractable to consider each of the matches).

Case-Based Reasoning (Kolodner, 1993) provides a partial answer to this question: provide each complex case with a simple, easily analyzed *index* which specifies the features which are relevant to reuse. SBF further elaborates on this idea in some domains such as design and invention, by providing accounts for what constitutes a useful index (typically, functional properties). Thus it is possible to avoid having to consider all of the details of all of memory in order to find a relevant match.

However, the historical case being analyzed here is not easily accounted for by this theory. The analogical transfer done here was not done on the basis of abstract functional characteristics of the sort that past work suggests are used to index cases in memory. Instead the transfer occurs based on very specific behavioral details. We claim that this sort of transfer is not usually possible because of the problem with too many potential matches. In this case it appears to have been done *opportunistically*; the internal (non-indexed) features of the phonautograph could be matched because they had been recently accessed, i.e., they were heavily primed within working memory. Our model not only shows how transfer from the phonautograph to the microphone could occur with the phonautograph being recently used, but also how it *would not* have occured if the phonautograph hadn't been used recently. This is the primary contribution of ACT-R to this model: an account of how knowledge which is active in memory can enable reasoning which is not possible without this activation. Because the ACT theory, as opposed to (for example) SOAR, focuses heavily on these sorts of working memory issues (Anderson, 1983; Anderson, 1993), it is particularly well suited to modeling the particular phenomenon addressed here.

The example analyzed here is a small piece relevant to the work being done in ALEC (Simina & Kolodner, 1997; Simina et al., 1998), ToRQUE (Griffith et al., 1996; Griffith & Murdock, 1998), and NICOLE (Francis and Ram, 1997). By modeling this example in ACT-R we are able to examine the issues of activation and retrieval in great depth. This analysis is potentially of great benefit to understanding the role of memory phenomena in analogy and other reasoning processes.

Acknowledgments

We thank Tony Simon, Ashwin Ram, Ashok Goel, Janet Kolodner, Mike Gorman, Todd Griffith, and Anthony Francis for their comments on this research.

References

Anderson, J.R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, 22.

Anderson, J.R (1993). *Rules of the Mind.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Bhatta, S., Goel A.K., & Prabhakar, S. (1994). Analogical design: A model-based approach. *Proceedings of the Third International Conference on Artificial Intelligence in Design*, Lausanne, Switzerland.

Francis, A.G. & Ram, A. (1997). Can your architecture do this: A proposal for impasse-driven asynchronous memory retrieval request generation. *Proceedings of the AAAI-97 Fall Symposium on ROBOTS, SOFTBOTS, IMMOBOTS: Theories of Action, Planning and Control.*

Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. *Cognitive Science*, 7.

Goel, A.K. (1989). *Integration of case-based reasoning and model-based reasoning for adaptive design problem solving.* Doctoral dissertation, Department of Computer Science, Ohio State University.

Goel, A.K., Bhatta, S., & Stroulia, E. (1997). Kritik: An early case-based design system. In Maher, M.L. & Pu, P., (Eds.), *Issues and Applications of Case-based Reasoning to Design.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Goel, A.K., Gomez, A., Grue, N., Murdock, J.W., Recker, M., & Govindaraj, T. (1996a). Explanatory interface in interactive design environments. *Proceedings of the Fourth International Conference on Artificial Intelligence in Design*, Kluwer Academic Publishers.

Goel, A.K., Mahesh, K. Peterson, J., & Eiselt, K. (1996b). Unification of language understanding, device comprehension and knowledge acquisition. *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Gorman, M. E. (1997). Mind in the World: Cognition and Practice in the Invention of the Telephone. *Social Studies of Science* 27:4.

Griffith, T. & Murdock, J.W. (1998). The role of reflection in scientific exploration. *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Griffith, T., Nersessian, N., & Goel, A.K. (1996). The role of generic models in conceptual change. *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Kolodner, J.L. (1993). *Case-Based Reasoning.* San Mateo: Morgan Kaufmann.

Simina, M., Kolodner, J.L., Ram, A., & Gorman,M. (1998). Opportunistic enterprises in invention. *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Simina, M. & Kolodner, J.L. (1997). Creative design: reasoning and understanding. In Leake, D.B. & Plaza E. (Eds.), *Case-based Reasoning Research and Development, ICCBR-97.* Springer-Verlag.

Simina, M., & Kolodner, J.L. (1995). Opportunistic Reasoning: A Design Perspective. *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (pp. 78–83). Hillsdale, NJ: Lawrence Erlbaum Associates.

US v. Bell (1908). *The Deposition of Alexander Graham Bell*, Boston: American Bell Telephone Company.