Tech Report GIT-COGSCI-97/04

Invention as an Opportunistic Enterprise

Marin Simina, **Janet Kolodner**, **Ashwin Ram**
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
{marin, jlk, ashwin}@cc.gatech.edu

Michael Gorman
TCC, SEAS
University of Virginia
Charlottesville, VA 22901
meg3c@virginia.edu

Abstract

This paper identifies goal handling processes that begin to account for the kind of processes involved in invention. We identify new kinds of goals with special properties and mechanisms for processing such goals, as well as means of integrating opportunism, deliberation, and social interaction into goal/plan processes. We focus on *invention goals*, which address significant enterprises associated with an inventor. Invention goals represent "seed" goals of an expert, around which the whole knowledge of an expert gets reorganized and grows more or less opportunistically. Invention goals reflect the idiosyncrasy of thematic goals among experts. They constantly increase the sensitivity of individuals for particular events that might contribute to their satisfaction. Our exploration is based on a well-documented example: the invention of the telephone by Alexander Graham Bell. We propose mechanisms to explain: (1) how Bell's early thematic goals gave rise to the new goals to invent the multiple telegraph and the telephone, and (2) how the new goals interacted opportunistically. Finally, we describe our computational model, ALEC, that accounts for the role of goals in invention.

Introduction

This paper investigates the enterprises of invention. We focus on invention goals, which address those enterprises of an inventor resulting in the creation of novel and useful devices. Invention goals provide one way to explain how creative and innovative ideas are generated, evaluated and further pursued by expert reasoners, such as inventors.

In our view, a professional inventor pursues his invention ideas to satisfy his/her curiosity (i.e., learn more) and/or to get social recognition (e.g., to get famous and wealthy). Whenever the inventor comes across a new idea (e.g., through experimentation or social interaction) that is interesting for him and/or for society, the preliminary preconditions for generating an invention goal are met. But the inventor must also estimate his chances of success, before allocating significant cognitive resources for pursuing a new idea. The estimation is based on the inventor's knowledge and also on his/her available cognitive resources. Bell's quest for the multiple telegraph and for the telephone[1] (US v. Bell, 1908) count among invention goals. These goals help to expand and reorganize the whole expertise of the inventor, by providing a method to try new approaches and learn more when other methods don't work. In particular, invention goals may generate expertise goals (i.e., learning goals to get domain expertise), while the knowledge gained through expertise goals may be used to generate new invention goals.

Due to their high risk, and to the possibility of cross-fertilization among overlapping goals, an inventor tends to pursue several invention goals in parallel. Gruber (1981) used the term "network of enterprises" to describe the way scientists such as Darwin pursued a set of related enterprise goals. However, since the successful pursuit of an invention goal may often demand most of the recognition and processing capabilities available to the inventor, the number of invention goals pursued in parallel tends to remain low.

Our exploration is based on a well-documented example: the invention of the telephone by Alexander Graham Bell. We started by analyzing Bell's invention process in terms of themes[2], goals and plans (Schank and Abelson, 1977), but we realized that some of the thematic goals[3] got refined and achieved a preferential status in an unexpected way, that further guided the whole invention process. As an example, Bell became obsessed by the idea of the telephone, even if it was not in his thematic expertise domain (i.e., acoustics):

> "If I could make a current of electricity vary in intensity, precisely as the air varies in density during the production of a sound, I should be able to transmit speech telegraphically" (Watson, page 8, 1913).

Social interaction, a phenomenon neglected in traditional accounts of goal/plan behavior, played an important role for the invention of the multiple telegraph and telephone. By social interaction, Bell learned new plans to achieve his invention goals. Reading electricity books and performing electrical experiments were among those plans. But it was the opportunistic recognition of information interesting to his invention goals that guided Bell's learning processes. Unlike traditional goals, invention goals may remain active after finding several (sketchy) alternative model/design solutions for them.

Whenever Bell came across a better model/design in the service of an invention goal (suspended or active in the background), he learned/assimilated the new alternative. As an example, Bell's initial conceptual model for the telephone (i.e., the harp apparatus[4]) was based on the idea that speech must

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

[2]A *theme*, according to Schank, is a generator/predictor of thematic goals.

[3]*Thematic goals* are those goals that are "around" all the time, and can be predicted for an individual.

[4]The *harp apparatus* was Bell's mental model for a device that could transmit either musical tones or speech.

be decomposed explicitly in its harmonic tone constituents for electrical transmission at distance. Later, during an experiment with the multiple telegraph, Bell came across a new design able (in principle) to transmit voice at distance without decomposing it explicitly in its constituent tones. At once, the goal of inventing the telephone was remembered and the new design was considered as a more promising solution for it.

Whenever a recently considered/processed artifact was primed during the elaboration of the current invention goal, that artifact was considered as a potential alternative to satisfy the invention goal. For example, while working on the telephone microphone, Bell was also performing experiments with the ear phonautograph, a mechanical device for visualizing speech. When Bell tried to remember a device able to move a piece of steel in the way that the air was moved by the action of the voice, he immediately remembered that the ear phonautograph provided the required behavior[5].

This idiosyncratic sensitivity for invention goals behaved like a kind of "knowledge lens", which helped Bell to deal with large amounts of information, by focusing him only on the relevant parts. The relevant information was constantly reorganized and learned as new cases (i.e., chunks of knowledge). Consequently, Bell achieved a goal-directed expertise in electricity, very efficient for his goals, but which sometimes ignored the traditional view of the domain. Moses Farmer offered a now famous characterization of Bell's idiosyncratic expertise:

> "If Bell had known anything about electricity he would never have invented the telephone"(Wattson, 1913).

This paper identifies a process of goal/plan handling that begins to account for the kinds of goal processing that inventors and expert researchers do. We also identify a new kind of goal with special properties (i.e., invention goals) and mechanisms for processing such goals as well as means of integrating opportunism, deliberation, and social interaction into goal/plan processes.

Based on Bell's case study, we are developing a computational architecture, ALEC, which accounts for the role of invention goals and social interaction in invention and creative design. Our computer model extends the memory architecture presented in Simina & Kolodner (1995).

Cognitive Issues in Invention

How do interesting ideas give rise further to new invention goals (i.e., goals to invent new artifacts)? What knowledge and processing is relevant for pursuing an invention goal? What is the role of social interaction in invention? These are some of the issues that a cognitive model of invention must address. Our analysis of Bell's quest for the multiple telegraph and telephone has identified ways in which existing goal processing methods (Schank & Abelson, 1977; Wilensky, 1983; Hammond et al., 1992) must be modified or augmented to handle invention.

[5]Note that the required behavior was not reflected in the phonautograph function, namely to transform the speech in a graphical representation. Consequently, Bell could not remember the phonautograph based solely on its function. This issue is addressed in Simina & Kolodner (1995).

A Critical review

We started by analyzing Bell's reasoning in terms of goals, plans and themes, according to Schank & Abelson's (1977) computational model of goal generation and refinement. Using Schank and Abelson's model, we could easily identify Bell's main *life theme* as a "teacher of the deaf", which generated the goal to invent machines that would make it easier for the deaf to hear and to learn to speak (e.g., to "visualize speech" by providing visual feedback). But we could not identify a straightforward theme-goal-plan chain to account for Bell's goal to invent the multiple telegraph or the telephone. After all, Bell was an expert in acoustics and his thematic goals, as taught by his father, had nothing to do with electricity. Nor is it clear why a teacher of the deaf would want to invent devices such as a multiple telegraph or telephone. According to our analysis, the generation of invention goals is a *deliberative* process, in which the rule-based mechanisms for goal generation proposed by Schank & Abelson (1977) is only part of the story. Even if an inventor finds some design ideas interesting for him and/or society, he will not decide instantly to pursue these ideas. Before that, an inventor must also estimate if he has the prerequisite knowledge and cognitive resources to pursue the idea successfully. In this respect, invention goals are similar to knowledge goals (Ram & Hunter, 1992).

Once an invention goal is generated, what processing mechanisms are relevant for pursuing it further? Once Bell decided to pursue an invention, he allocated significant cognitive resources to it. But a recently processed invention goal may still remain active in background for a while (Simina & Kolodner, 1995). By strategically choosing the current invention goal, an expert reasoner (e.g., inventor) may maintain a relevant network of active goals in the background. Similar reasoning processes were identified by Gruber (1974) in Darwin's work:

> "The fact that he was all these things at once meant that a unique and productive intersection of many enterprises could occur in his thinking. At the same time, the existence of this ensemble was not an accident but the deliberately cultivated fruit of Darwin's work."

Research on predictive encoding (Hammond & al., 1992) grew up from the difficulties associated with active goals (e.g., it is unlikely that all the goals are active due to computational demands). Basically, goals that cannot be satisfied immediately are associated at the time of encoding in memory with features of the environment in which goal achievement would likely to be possible. But in invention the structure of goals is more complex and it is difficult to enumerate, in advance, all the features of the environment in which the satisfaction of the goal might be possible.

Fortunately, we can maintain a small number of goals active at any time (e.g., a network of enterprises, according to Gruber), without excessive computational demands. The issue is to identify the most likely constituents of this set. Simina & Kolodner (1995) presents a memory model which accounts for the existence of active goals and postulates their limited number, but it does not indicate which types of goals are more likely to be active and why.

While the explicit purpose of invention is designing a novel or useful artifact, learning plays an important part in invention. An inventor often must learn new concepts (sometimes by consulting others; sometimes by experimentation and exploration) in order to move forward with designing. Ram & Hunter (1992) describe the role of knowledge goals to guide inference and learning, in the context of story understanding and problem solving. While invention and understanding have common elements, experimentation is specific to invention (Gorman, in press). The whole invention process may be driven sometimes only by experimentation (i.e., "tinkering"). This requires specific inference mechanisms which distinguish invention goals from knowledge goals.

The role of social interaction was in general minimized in cognitive models of invention. Both IMPROVISER (Wills and Kolodner, 1993) and IDeAL (Bhatta & Goel, 1993) propose an "oracle" for modeling external interaction. In IDeAL, an oracle supplies the needed information when the system itself fails to solve a problem and the system assimilates the new information so as not to fail for the same reason in future. Thus IDeAL's interaction with other agents is quite limited. Our analysis of Bell's notebooks suggests that external interaction in invention takes the elaborated form of social interaction. An expert may request assistance from his peers, but an answer is not guaranteed. New processes are responsible for handling social interaction. We can present now our approach to invention.

Our Approach: Invention Goals

To deal with the above invention issues, we identified a distinguished class of goals, namely *invention goals*. Most of the enterprises associated with an inventor are invention goals (i.e., goals to identify and generate novel artifact ideas, to be pursued as design problems). We associate a *level of activation* with invention goals, to account for their importance to the inventor. While a low importance goal may be satisfied as soon as we find a first solution for it (even imperfect), a high importance goal may remain unsatisfied even after accumulating several alternative solutions for it. Let's explore how invention goals may help us to deal with the above invention issues.

How do interesting ideas give rise further to new invention goals? In our view, an idea is interesting for an inventor when: (1) it is instrumental to satisfying some of his higher thematic goals, and (2) he/she has some idea about how to go about investigating it. Some of the higher thematic goals are personal (e.g., scientific curiosity), while others have to do with what society will value. If an invention idea is judged to be interesting for the inventor and/or the society, the inventor must also deliberate if he has the knowledge and the cognitive resources for pursuing it further. In the favorable case, a new invention goal is generated and a level of activation is associated with it, depending on its importance for satisfying higher thematic goals. The essential aspect of invention goal generation is that it is a *deliberative* process in invention.

Let's exemplify the above ideas with the episode when Bell decided to approach the invention of the multiple telegraph (Bruce, 1973). In October 1872, Bell read a newspaper article[6] describing the impact of a new telegraph system able to transmit simultaneously two telegraphic messages over the same wire (Stearn's duplex system). The article also suggested that fame and fortune awaited the inventor of a telegraphic system able to transmit more than two messages simultaneously. The multiple telegraph idea was definitely interesting for society, and this could be an incentive for Bell to check the relevance of the multiple telegraph idea for his own research. Bell assessed the problem as follows: (1) multiplex multiple messages and send them over telegraphic wire, and (2) demultiplex the messages at destination. Bell remembered that he had some experience with telegraphic equipment, while he was trying to understand Helmholtz's Apparatus[7]. According to Bell's understanding, Helmholtz's Apparatus was able to unscramble (demultiplex) multiple tones sent over a single wire, by using tuned receivers. Since both devices performed demultiplexing, Bell could use his understanding of Helmholtz's Apparatus for the multiple telegraph idea. Bell estimated that he also had the prerequisite knowledge for inventing the multiple telegraph. Helmholtz's Apparatus provided an easy solution for half of the problem (demultiplexing), while the multiplexing part looked even easier: just add several tone generators in series with the one existing already in Helmholtz's Apparatus. In this moment Bell decided to pursue the invention of the multiple telegraph (i.e., an invention goal), for which he already had a theoretical model and a partial implementation.

What knowledge and processing are relevant for pursuing invention goals? Invention goals are goals that identify novel and useful artifact ideas (functions or behaviors), to be pursued further as design goals. Consequently, invention goals may contain a behavioral device specification, possibly incomplete and inconsistent. The synthesis of a structural solution may be facilitated by evolving in parallel several alternative solutions to the design specification (Wills & Kolodner, 1993; Grubber, 1974). As a side effect of pursuing invention goals, the expertise of the inventor increases. Namely, if invention goals cannot be pursued due to the lack of expertise, they may spawn expertise goals (broader learning goals) to get that expertise.

In our view, two complementary and interacting processes are responsible for incrementally evolving an invention (Kolodner & Wills, 1993): (1) Artifact Generator, which proposes new design solutions to the current design specification, and (2) Evaluator, which critiques the current design solution, by updating the design specification. Both processes rely heavely on the previous experience of the inventor. Artifact Generator relies on a library of known artifacts and techniques for adapting them to fit the current design specification, while Evaluator relies on simulation, experimentation and knowledge to interpret the experimental results. The critique provided by Evaluator may suggest a divide and conquer strategy for pursuing an invention goal by decomposing it in subgoals. Gorman (1992) identifies that Bell used such a strategy to elaborate independently critical subparts (identified by their behavior) of the telephone. In our view, the subgoals of an invention goal inherit part of the invention goal status (i.e., a slightly smaller level of activation).

But what processing may be responsible for generating such subgoals? Our exploration of the invention of the multiple telegraph and of the telephone suggested to us a possible

[6] In *Boston Transcript*, the newspaper in which Bell advertised his speech lessons.

[7] Helmholtz's Apparatus was an electrical device for producing artificial vowels, relevant to Bell's acoustical research.

approach. The inventor chooses a main design alternative responsible for satisfying the main, or most difficult to satisfy, constraints. Then he generates subproblems to satisfy the other secondary constraints, in the framework of the main design alternative. For the invention of the telephone, the main constraint was to transmit an "undulatory current" (i.e., both pitch and amplitude) over a telegraphic wire. Once Bell identified such a design, he generated subproblems to improve its secondary characteristics. Those subproblems resulted in the design of the telephone's transmitter and receiver.

Let's address also some of the processing differences between invention goals and thematic goals. While themes may characterize the goals common to most of inventors, they don't say anything about how inventors *idiosyncratically* refine these themes and why different individuals working in the same domain may chose different paths to achieve their (similar) thematic goals. Also, the traditional view of thematic goals does not explain why and how inventors allocate more computational resources to some thematic goals (keeping them more active), but not to others. Invention goals are intended to explain the idiosyncrasies of individual experience among experts, providing a "knowledge lens" for efficient interaction with the events noticed in the world. Our analysis of Bell's diaries and related literature suggest the following mechanism. Once an expert evaluates the potential of a goal to fulfill at least one of his major thematic goals, he activates and focuses mostly on that (invention) goal. The expert affords to ignore for a while his other thematic goals, since the achievement of an invention goal will satisfy implicitly his major thematic goal.

What is the role of social and environmental interaction in invention? In our view, a cognitive model for invention must account for social and environmental interaction. Otherwise, the model fails to give a plausible account of the huge computational resources involved in acquiring the "right" knowledge for making the invention possible. A real expert takes advantage of the knowledge and processing available elsewhere through strategic social and environmental interaction.

An expert is part of a society of experts, who may communicate among themselves. Each expert has similar cognitive abilities, but different knowledge. Consequently, particular invention goals may be addressed by some experts, but not by others. Each expert can "ask" (i.e., make queries) to: (1) retrieve knowledge available elsewhere, and (2) request external evaluation of a proposed design. Once an expert learns to design a new artifact, all of the other experts in the society have access to that design through the publications and social interactions of that inventor.

By reading a newspaper article (i.e., environmental interaction), Bell learned about the importance of inventing the multiple telegraph. The Goal Generator process found the problem interesting and feasible, so it generated the goal to invent the multiple telegraph. What about the role of social interaction? If the Artifact Generator process cannot propose a design solution to the current specification, it may ask other experts for a solution. In particular, this was the reason why Bell hired Wattson for developing electrical artifacts. Also, if a design solution proposed by the Artifact Generator process cannot be evaluated by the Evaluator process (due to his lack of previous experience), it may ask other experts to perform the evaluation. For example, when Bell had doubts about the physical principles involved in telephony, he requested help from distinguished scientists, such as Joseph Henry.

Computational Model

Architecture

Our computational architecture for invention is implemented in a computer program called ALEC (see Figure 1). We propose a distributed model of invention, where each network node represents an expert. Each expert has the same architecture, but different knowledge. Experts communicate by sending messages among them.

Since our view of the invention process relies heavily on event detection and processing, we opted for an event-driven architecture. A WORKING MEMORY (WM) keeps track of all the state information which may generate internal or external events. Since the processes (represented as gray rectangles in Figure 1) operating on the WM rely heavily on previous experience, we need also a LONG-TERM MEMORY (LTM) to account for the role of experience in invention. Our model builds on the blackboard model of WM (Hayes-Roth & Hayes-Roth, 1979) and on the dynamic memory model of LTM (Schank, 1982). Basically, changes in the WM generate events used to activate or generate new goals in the AGENDA. A Strategic Control process selects the next CURRENT GOAL from the AGENDA, based on a CONTROL PLAN.

Let's elaborate the LTM of our architecture. In an event-driven architecture, the monitored events may remember SUSPENDED GOALS and THEMATIC PATTERNS (i.e., the generators of thematic goals). The SUSPENDED GOALS are indexed in LTM in terms of their *preconditions*, represented as monitored events. To facilitate the decomposition of goals in subgoals, via plans, the LTM contains a collection of PLANS, used by the Plan Scheduler to retrieve an appropriate plan, given a goal. To facilitate the generation of design ideas, given their intended function, the LTM contains a collection of CASES, ARTIFACTS, MODELS, indexed by their function. To facilitate internal evaluation of the CURRENT SOLUTION proposed by the Artifact Generator, the LTM contains SIMULATOR KNOWLEDGE used by the Evaluator. To facilitate a further analysis of the invention process, LTM contains a NOTEBOOK, which provides a derivational record of the invention.

Once an expert learns about a new expert (and its expertise domain), its External Interaction process must be able to use this knowledge, if needed, to send request messages. Consequently, LTM should be also a repository of knowledge about OTHER EXPERTS, indexed by their expertise domains.

The WM represents the "activated" part of LTM. To simulate *priming effects*, WM keeps track of the recently processed design goals, BACKGROUND GOALS, and of the recently processed artifacts, BACKGROUND ARTIFACTS. A decay process (not represented in the Figure 1) limits the number of the activated items in WM. Based on the events noticed in WM, the Goal Generator may update the AGENDA by: (1) generating thematic goals, (2) activating suspended goals, or (3) generating invention goals, instrumental for satisfying thematic goals. The Strategic Control selects a CURRENT GOAL from the AGENDA. The CURRENT GOAL may contain a CURRENT SPEC of a design (possible incomplete and inconsistent), a CURRENT SOLUTION, and a set of ALTERNATIVES (used by the

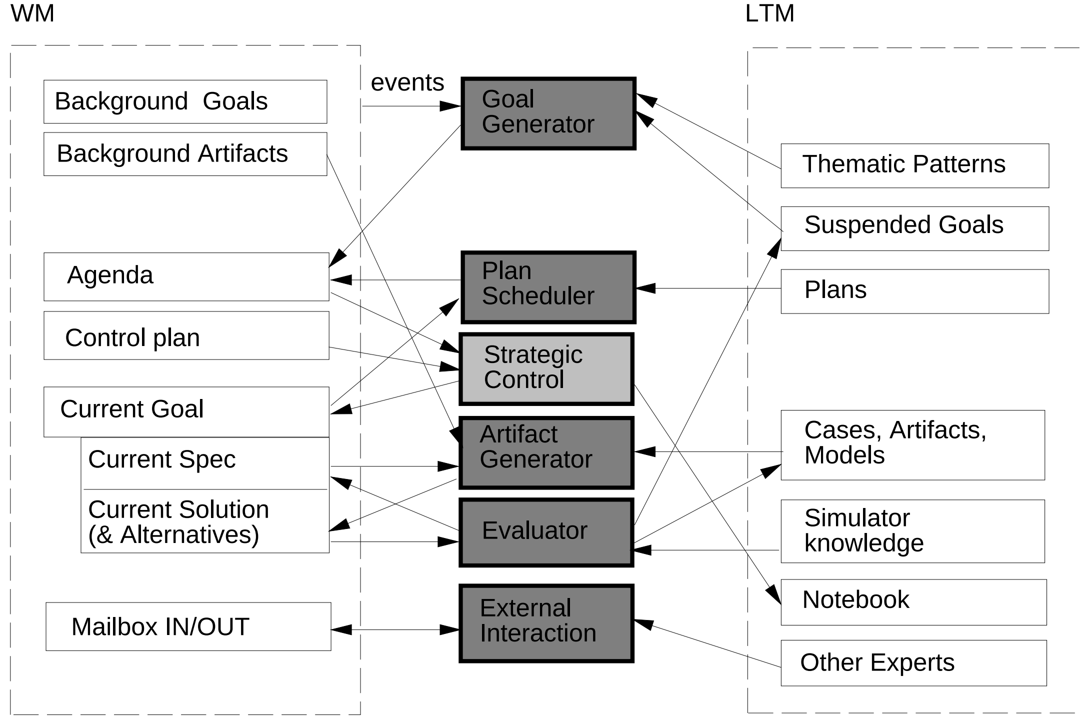

Figure 1: ALEC: a Framework for Invention

Artifact Generator to propose new designs). The MAILBOX is used for communication with other experts, via External Interaction. If a design goal cannot be achieved using the knowledge available locally, a request is sent to relevant experts. Since an answer is not guaranteed and it may come asynchronously, the system must be able to encode predictively the goal in memory and to approach other goals from the AGENDA in the meantime.

Algorithms

Now we can describe the main algorithms for creating and manipulating invention goals, essential to our theory of invention.

Invention goal generation is part of the Goal Generator process. Whenever ALEC learns about a new design idea (by social interaction or experimentation), it performs the following steps: (1) identify if the idea is *interesting* (does the implementation of the idea result in postconditions that match those of important unsatisfied thematic goals?), (2) estimate its own expertise to implement the idea as an artifact (given the idea's rough design spec, can the Artifactor Generator generate an artifact, judged/simulated as promising by the Evaluator?) (3) generate an invention goal (instrumental for satisfying specific thematic goals), if the results of steps (1) and (2) are positive.

Given an invention goal specification, the Artifact Generator uses the following retrieval algorithm: (1) if the design specification matches any of the BACKGROUND ARTIFACTS, return it, (2) otherwise, if LTM retrieval is successful, return the remembered devices, (3) otherwise, *ask* (send message to) other experts for such an artifact and suspend the invention goal in memory, indexed by its design specification.

If the Artifact Generator is unable to retrieve an artifact that satisfies all the design spec constraints, it attempts to decompose the initial invention goal in subgoals to facilitate the further synthesis of the desired artifact. The algorithm for decomposing an invention goal in subgoals, given the design specification of the artifact is: (1) find a design alternative that satisfies the *main* constraint, (2) identify which components of this design alternative are responsible for satisfying the secondary constraints, (3) generate subproblems for designing components, to satisfy (better) the secondary constraints.

But Evaluator is essential to validate the design solutions proposed by the Artifact Generator. Here are the steps: (1) perform simulation of the artifact and critique it, if ALEC has enough domain knowledge, (2) otherwise, implement the artifact, perform experiments and interpret the results, if this is possible, (3) otherwise, send message to relevant experts to remotely perform the evaluation.

Step (3) of the Artifact Generator retrieval and Evaluator algorithms may result in the deliberative generation of an expertise goal (i.e., a learning goal to provide a broader expertise). The details are beyond the scope of this paper.

An Example: the Multiple Telegraph Idea

After reading a newspaper article that fame and fortune awaited the inventor of a multiple telegraph, Bell realized that he could design a multiple telegraph based on his knowledge of Helmholtz's Apparatus and acoustics. Moreover, the goal to invent the multiple telegraph became Bell's invention goal for the next years. We present a possible analysis of this episode, in terms of our computational model (ALEC) presented in Figure 1.

Reading the newspaper article about the multiple telegraph is equivalent in our system to receiving a message (in the MAILBOX), containing a design goal (i.e., Multiple Telegraph), characterized by a sketchy DESIGN SPECIFICATION and a *post-condition* describing that the inventor will become *famous & wealthy*. This event is analyzed by the Goal Generator, which identifies that *famous & wealthy* is a postcondition associated with some of Bell's thematic goals (i.e. designing the multiple telegraph might provide an *novel* alternative plan to satisfy thematic goals) and is also a precondition to other thematic goals (e.g., *get married*). Consequently, the Multiple Telegraph is *interesting*, and if ALEC could relate it to its own expertise, the Multiple Telegraph would become a good candidate for an invention goal. Based on the specification of the *Multiple Telegraph* (i.e., multiplex telegraphic messages over a single wire and demultiplex them at destination), the Artifact Generator searches the memory of CASES, ARTIFACTS, MODELS and retrieves Helmholtz's Apparatus, which was able to demultiplex a collection of harmonic tones. After a quick case-based *adaptation*, involving using several generators instead of only one, the Artifact Generator has a design solution. The Evaluator simulates that indeed the new adapted design satisfies the CURRENT SPEC and it reports the result to the Goal Generator, which promotes the Multiple Telegraph as an invention goal, by associating with it a computed *level of activation* (based on the importance of the related thematic goals) and an initial (theoretical) solution. The new invention goal is put in the Agenda and will be further selected for processing.

Discussion

ALEC, our model for invention, addresses some of the control problems associated with the invention processes. In our opinion, a realistic computational model for invention should emphasize both the deliberative and opportunistic components of the control in invention. We intend for invention goals to provide appropriate mediation between the deliberative and opportunistic components of the reasoning process in invention. Namely, generating an invention goal is a deliberative process, intended to further guide the opportunistic evolution of the invention.

Previous event-driven architectures (Hayes-Roth & Hayes-Roth, 1979; Wilensky, 1983) emphasized the opportunistic aspect of control and goal generation, but did not provide enough specific tools for balancing the deliberative aspect.

Other cognitive models for invention, like IDeAL (Bhatta & Goel, 1993), ignore completely the opportunistic aspect of control, but rely on the existence of an "oracle", which provide "the right information at the right time". Our model recognizes the limited computational power available to an inventor, but provides something more realistic than an oracle, namely social and environmental interaction. As our Bell's case study shows, an inventor which acts in a society is necessarily opportunistic.

ALEC's architecture grew up from the difficulties encountered with IMPROVISER's opportunistic control (Wills & Kolodner, 1994). Simina & Kolodner (1995) postulated the existence of "active goals", as a mechanism to explain priming effects, but it did not characterize which goals are more likely to be active. ALEC goes further, identifying invention goals as the most likely active goals in a cognitive model for invention.

Of course, ALEC is still in its infant stages (as a Macintosh Common Lisp prototype), and we can't say yet that it provides a complete or correct model of goal generation in invention. We must still test it on a wide variety of examples (the invention of the telephone offers plenty, beside those mentioned briefly here), and we must investigate more deeply: (1) the synergy among active goals, (2) the relation between invention/enterprise goals and expertise goals, (3) ALEC's ability to simulate inventors with different backgrounds.

Acknowledgements

This research was funded in part by ONR Grant No. N00014-92-J-1234 and in part by NSF Grant No. DIR-8722002 and DIR-90-12311. We thank Ashok Goel, Hari Narayanan and Nancy Nersessian for their comments on this research.

References

Bhatta S. & Goel A. (1993) Learning Generic Mechanisms from Experiences for Analogical Reasoning. *Proceedings of the Fifteenth Conference of the Cognitive Science Society.* LEA (pp. 237–242).

Bruce, R.(1973). *Bell: Alexander Graham Bell and the Conquest of Solitude*, Little Brown & Co.

Gruber, H. (1981). *Darwin on Man: A Psychological Study of Scientific Creativity* (2nd ed.). University of Chicago Press.

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

Gorman, M. (1992). *Simulating Science: Heuristics, Mental Models and Technoscientific Thinking.* Bloomington: Indiana University Press.

Hammond, K., Converse, T., Marks, M. & Seifert, C. (1993). Opportunism and Learning. *Machine Learning*, 10 (pp. 279-309).

Hayes-Roth, B. & Hayes-Roth, F. (1979). A cognitive model of planning. *Cognitive Science* 3 (pp. 275–310).

Kolodner, J. & Wills, L. (1993). Case-based Creative Design. *AAAI Spring Symposium on AI and Creativity.* Stanford, CA. March 1993.

Ram, A. & Hunter, L. (1992). The Use of Explicit Goals for Knowledge to Guide Inference and Learning. *Journal of Applied Intelligence*, 2(1) (pp.47–73).

Schank, R. (1982). *Dynamic memory*. Cambridge University Press.

Simina, M., Kolodner, J. (1995). Opportunistic Reasoning: A Design Perspective. *Proceedings of the Seventeenth Conference of the Cognitive Science Society.* LEA (pp. 78–83).

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

Watson, T. (1913). *The Birth and Babyhood of the Telephone.*

Wilensky, R. (1983). *Planning and Understanding.* Addison-Wesley.