In *Proceedings of the 12th National AAAI Conference*, Seattle, WA, August 1994

A Model of Creative Understanding*

Kenneth Moorman and Ashwin Ram

Georgia Institute of Technology
College of Computing
Atlanta, GA 30332-0280
{kennethm,ashwin}@cc.gatech.edu

Abstract

Although creativity has largely been studied in problem solving contexts, creativity consists of both a generative component and a comprehension component. In particular, creativity is an essential part of reading and understanding of natural language stories. We have formalized the understanding process and have developed an algorithm capable of producing creative understanding behavior. We have also created a novel knowledge organization scheme to assist the process. Our model of creativity is implemented as a portion of the ISAAC (Integrated Story Analysis And Creativity) reading system, a system which models the creative reading of science fiction stories.

Introduction

Creativity remains a largely unexplained facet of human intelligence; neither psychologists nor artificial intelligence researchers have produced complete theories of it. While most creativity researchers have investigated the behavior in a problem solving context, we are more interested in how creativity is manifested during comprehension. We thus distinguish between two forms of creativity: a generative type, *creative invention* or *creative problem solving* (e.g., Hofstadter & McGraw 1993; Kolodner & Wills 1993); and an explanatory type, *creative understanding* (e.g., Kass, Leake, & Owens 1986; Ram 1993). Past research has shown the value of exploiting the relationship between problem solving and understanding (e.g., Wilensky 1983; Birnbaum 1986); likewise, our study of creative understanding should aid general creativity research.

Our model of creative understanding is functional in nature. A cognitive *process* can be explained by describing the *function* of each of its *tasks*, the *relationships* between them, the *mechanisms* which accomplish them, and the *knowledge* required. The resulting *functional theory* can then be used to guide implementation of a process model. We have identified four tasks which are sufficient for the modeling of creative understanding. For three of these, we extended well-known mechanisms from traditional problem solving and comprehension domains. The final one, problem reformulation, relies on a new mechanism, function-driven morphological synthesis. While developed in a reading context, this is a general method which can apply to other creativity modeling. The creative understanding processes, mechanisms, and a novel knowledge organization scheme which supports the processes are implemented in the ISAAC (Integrated Story Analysis And Creativity) system which creatively reads science fiction stories.

Creative Understanding

The reading of any story requires some level of creative understanding; it is when normal processing fails that creative understanding is most necessary. Consider the science fiction story, *Men Are Different* (Bloch 1963), seen in Figure 1. If a reader is unfamiliar with the concept of sentient robots, the story is impossible to appreciate. While the next section explains exactly what happens when such a need for creative understanding arises, the remainder of this section presents a discussion of creativity in order to provide a common ground to an elusive and often controversial concept.

Creativity

Creativity is a directed, internal process of a cognitive agent which results in an artifact which is both novel and useful. This is intended as a working definition; the following elaborations should clarify our intended meaning.

Process: For a process to be creative, it must be both *internal* and *directed*. The internal restriction ensures that the reasoner is not simply repeating a piece of knowledge just received from another source. The directedness restriction ensures that the reasoner is not simply a random generator of solutions, where one may eventually be "creative" through sheer chance.

Artifact: An *artifact* results from a mental process by a reasoner and may be either physical or mental. Any artifact can be described by a set of *attributes* which define its characteristics. One of these, *function*, represents the

*This work was supported by a Fannie and John Hertz Foundation fellowship and by the Georgia Institute of Technology.

> I'm an archaeologist, and Men are my business. Just the same, I wonder if we'll ever find out about Men—I mean *really* find out what made Man different from us Robots—by digging around on the dead planets. You see, I lived with a Man once, and I know it isn't as simple as they told us back in school.
>
> We have a few records, of course, and Robots like me are filling in some of the gaps, but I think now that we aren't really getting anywhere. We know, or at least the historians say we know, that Men came from a planet called Earth. We know, too, that they rode out bravely from star to star; and wherever they stopped, they left colonies—Men, Robots, and sometimes both—against their return. But they never came back.
>
> Those were the shining days of the world. But are we so old now? Men had a bright flame—the old word is "divine" I think—that flung them far across the night skies, and we have lost the strands of the web they wove.
>
> Our scientists tell us that Men were very much like us—and the skeleton of a Man is, to be sure, almost the same as the skeleton of a Robot, except that it's made of some calcium compound instead of titanium. Just the same, there are other differences.
>
> It was on my last field trip, to one of the inner planets, that I met the Man. He must have been the last Man in this system, and he'd forgotten how to talk—he'd been alone so long. I planned to to bring him back with me. Something happened to him, though.
>
> One day, for no reason at all, he complained of the heat. I checked his temperature and decided that his thermostat circuits were shot. I had a kit of field spares with me, and he was obviously out of order, so I went to work. I pushed the needle into his neck to operate the cut-off switch, and he stopped moving, just like a Robot. But when I opened him up he wasn't the same inside. And when I put him back together I couldn't get him running again. Then he sort of weathered away—and by the time I was ready to come home, about a year later, there was nothing left of him but bones. Yes, Men are indeed different.

Figure 1: *Men Are Different*

best-known uses of the artifact. The remaining attributes are divided into *primary attributes*, which contribute to an explanation for why the artifact can perform its function; and *secondary attributes*, which are the rest.

Novel: There are many arguments concerning whether a creative act must be *novel* and from whose perspective this novelty is judged (e.g., Boden 1991; Stewart 1950; Thurstone 1952). In reference to a given perspective, there are two important ways in which an artifact (M) may be novel with respect to a goal (G); to see these, consider a longsword used for combat.

- Evolutionary Novelty (E-Novel): M is defined to be E-Novel iff M is unknown and M accomplishes G in a better way than other examples of artifacts which accomplish G. This is generally the result of altering one of the primary attributes of M. A shortsword, a bastard sword, or a two-handed sword would all be E-Novel.
- Revolutionary Novelty (R-Novel): M is defined to be R-Novel iff M is unknown and M accomplishes G in a different way than other artifacts which accomplish G. For this, secondary attributes may need to be altered to cause them to participate in the function of M. The light saber from Star Wars is an R-Novel variation of the original longsword.

Useful: *Useful* is somewhat easier to define for an artifact. A reasoner attempting to perform creatively will have a certain task to accomplish. In order for a creative process to exist, it must produce an artifact which accomplishes this task in some manner.

Creative understanding specification

The presented definition of creativity is simply descriptive and is insufficient for implementing a model of creative behavior. For this, we need to develop an algorithmic view of creativity as it exists within creative understanding.

Formalization of problem solving Since understanding can be described as the complementary operation to problem solving, a formal view of problem solving can be helpful in developing an understanding formalization. *Problem solving* begins with the reasoner in an *initial state*. A reasoner knows of operations that it can perform which will move it through a *search space*. This stops when a *goal state* is achieved. The output from the process is a *solution path* which takes the reasoner from the initial state to the goal state (Newell & Simon 1972). The idea of *constraints* on this process is also important (Sacerdoti 1974; Sussman 1973). These are conditions which cannot be violated in the final solution. Finally, there are times when the reasoner may already possess a solution. If so, problem solving can discover a better solution if the reasoner possesses a critique of why the current solution is not a viable one (Hammond 1989). The complete formulation is shown in Figure 2.

Formalization of understanding Based on Figure 2, a formal specification of understanding can be developed (Figure 3). Using the example of *Men...*: if a reader sees a robotic character "turn off" a man and then open him, they may understand the episode by reasoning that the robot had the goal of repairing the man (*abduction*);

INPUT:
Initial state (I) Goal state desired (G) Set of constraints (C) (optional) Current solution (S) (optional) Critique why S is not good enough solution (K) (optional)

OUTPUT:
Solution path (S′) which achieves G given I and does not violate C

Figure 2: FUNCTION **Problem Solver**

a reader who learns that Mankind has become extinct and that the remaining robots are curious as to the fate of their creators may understand this by reasoning about upcoming story actions (*prediction*); finally, the reader may attempt to understand why the robot felt that field repairs was a good solution to the man's discomfort (*explanation*).

INPUT	OUTPUT	Behavior
Solution (S)	Goal (G)	Abduction
Goal (G)	Solution (S)	Prediction
Solution (S) and Goal (G)	Critique (K) of why S is a good solution	Explanation

Figure 3: FUNCTION **Understander**

The creative understanding process

If only known concepts are given to the understanding process, no creative behavior is necessary. If, however, a reasoner uses the understanding *process* to comprehend *novel artifacts* in a way which is *useful*, creative understanding occurs. Notice that the context of reading provides an exact meaning for the final part of this requirement, usefulness. Since the reasoner has the task of comprehending some piece of text, an understanding of an artifact from it is useful if this understanding allows comprehension to occur and reading to continue. This is in marked contrast with "traditional" approaches to creativity which must make the context of usefulness more explicit.

The tasks involved in creative understanding are carried out by a core set of cognitive mechanisms. This *cycle of creative understanding* (CUP) is shown in Figure 4 (depicted for an abduction task, the same approach is used for explanation and prediction). Mundane understanding exists if the reasoner only considers steps 1 and 2; that is, if the reasoner performs only memory retrieval and incorporation. Each cycle increases the potential for successful creative understanding. At some point the reasoner will be so far removed from the original concept that further iterations will be useless; still, there is no theoretical limit to the "amount" of creativity generated

INPUT:
Solution (S)

OUTPUT:
Goal state desired (G) Constraint set (K)

PROCESS:
REPEAT 1. Perform *memory retrieval* 2. If (1) fails ⇒ attempt *incorporation* 3. If (2) fails ⇒ attempt *baseless analogy* 4. If (3) fails ⇒ *reformulate the problem* UNTIL (successful OR concept is too bizarre)

Figure 4: FUNCTION **Creative Understander**

The four steps of CUP

The first step of the CUP algorithm involves a *memory retrieval*. If concepts are retrieved which cause understanding of the novel concept, the cycle ends successfully. If nothing is returned which is immediately useful, processing continues. This may occur if nothing is available in memory or if the proper items are simply not returned (e.g., due to an indexing problem). In *Men...*, the reasoner will be confronted with the idea of a sentient robot. If only an industrial robot exists in memory, understanding will fail. The concept of *industrial robot* is insufficient to explain the robot's actions.

If normal memory retrieval fails to produce adequate understanding, the algorithm will attempt to perform *incorporation*. Incorporation involves the discovery of relations between concepts which were retrieved from memory and the concept being explained. If a relationship can be discovered that explains the new concept, understanding is successful. This stage of the cycle may result in understanding which appears either mundane or creative, depending on exactly what was retrieved and what sorts of relations were discovered. In the example of *Men...*, incorporation fails because sentience is in conflict between the concepts of *industrial robot* and *story robot*.

If incorporation fails, the CUP algorithm attempts a technique known as *baseless analogy*. *Analogy* (e.g., Falkenhainer 1987; Gentner 1989) attempts to explain a concept (the *target*) by appealing to known information about an analogous concept (the *base*). However, if no existing base exists, it may be possible to dynamically build the base within a given domain if the reasoner possesses a great deal of information about the target's domain and the intended base domain (e.g., Clement 1989). For example, atomic structure as in the Bohr model can be understood within the framework of gang warfare (WKRP Episode 60 1980). If the reasoner has a great deal of conceptual background knowledge concerning sentience and robots, it is possible that an understanding of the story robot may result.

Finally, if all of the above steps have failed to produce a satisfactory understanding, the reasoner must resort to *problem reformulation*. There are some cases in which the initial statement of a problem is not the one which will lead to an optimal solution. By recasting the problem in a new way, a reasoner may gain insights into a possible solution. For *Men...*, the reader can attempt to take the retrieved concept of *industrial robot* and manipulate it with the goal of explaining the *story robot*. The method which accomplishes problem reformulations is function-driven morphological synthesis.

1. Consider an artifact in the world, designated as M.
2. Let f be defined as the function which returns the function of an artifact.
3. Let C be the class of functions which alter an object, either by changing some attribute of that object or by adding a new attribute to the object.
4. C_1 through C_n are a set of n such functions.
5. Thus, a set of objects S_{all} can be created by $\cup_{i=1\ldots n} C_i$.
6. Consider the subset, S_f defined as $\{s | s \epsilon S_{all}$ and $f(s) = f(M)\}$
7. Finally, consider the subset S_c defined as $\{s | s \epsilon S_f$ where s is unknown $\}$
8. The items in S_c are useful (they fulfill the same role as the original object M) and they are novel to the reasoner. Therefore, they are creative.

Figure 5: Function-driven Morphological Synthesis

Function-driven morphological synthesis

In order to model problem reformulation, we developed a new mechanism—*function-directed morphological synthesis* (FMS), depicted in Figure 5. It is assumed that the reasoner has an artifact that needs to be understood. The reasoner applies a set of manipulator functions to the artifact, altering its attributes and producing new artifacts. The artifacts which possess the original functionality and are novel to the reasoner are considered to be creative ones. The FMS technique was inspired by Allen's *morphological synthesis* (cited in Finke, Ward, & Smith 1992), in which a reasoner manipulates combinations of primary attributes to produce potentially creative results. Since only primary attributes were modified, Allen's technique could not result in an R-Novel item. By removing this restriction, FMS *is* able to produce such novelty. Finally, FMS can exist in both a strong form and a weak form. Strong-FMS performs the manipulations by examining other objects with the same functionality to see how they accomplish their tasks. Weak-FMS does away with this constraint and guides the manipulation through the reasoner's knowledge of given attributes and possible values.

In *Men...*, FMS must attempt to understand the *story robot*. It knows that the robot is a willful agent, which violates the concept of *industrial robot*. Memory retrieval produces a concept which is a willful agent, but is not a robot—man. Thus, FMS begins with *man* as an input (whose "purpose" is willful agency). FMS is aided by the fact that it has the goal of understanding *story robot*. Thus, when it manipulates the attributes of *man*, it does so by examining the concepts of *man* and *industrial robot* (an example of strong-FMS). It adds attributes to *man* from *industrial robot* and deletes attributes which are no longer needed. The result is a man-like robot which can be used to explain the actions of the *story robot*.

When to say when

A central issue in concept manipulation systems is how does a reasoner know which manipulations are good ones and which ones are potentially dangerous. While a *willing suspension of disbelief* (Corrigan 1979) is required for creativity, too much suspension will lead to ridiculous outcomes. Various approaches have been taken to minimize this problem in other creativity models. One possibility is to have the system do little self-monitoring; this approach can be seen in Lenat's AM system (1990). While AM did create a number of creative concepts, it also created a much larger number of worthless concepts which were filtered out by the human researcher. Alternatively, other systems exist which possess a large number of programmer-supplied heuristics for deciding what to manipulate and how far to carry the alterations; Turner's MINSTREL system (1992) and SWALE (Kass, Leake, & Owens 1986) are examples. The problem with this approach is that no techniques have been proposed which would allow the straightforward creation of these heuristics; instead, developing them is a "black-art." Both approaches have been argued against from a theoretical standpoint (Birnbaum 1986); we prefer a technique which would allow flexible modifications, decided by the system itself. As part of this, a new method of knowledge organization is needed.

Our knowledge organization scheme resembles a standard semantic network, but knowledge is tagged through the use of a multidimensional grid, as shown in Figure 6. One axis of the grid represents a Schankian breakdown: *action*, *agent*, *state*, and *object* (Schank & Abelson 1977). The other dimension represents a natural breakdown: *physical*, *mental*, *social*, *emotional*, and *temporal*. For example, a TRANSfer is a generic action. In the physical column is PTRANS, the mental column contains MTRANS, and the social column contains ATRANS. Our extended representation also includes emotional TRANSfers (the giving of one's love); and temporal TRANSfers (March getting closer to us).

Another difference between a standard semantic network and our organization scheme is the *function tagging* of each concept. Each concept within the knowledge system is tagged with a set of its common functions. If the reasoner is searching for similar concepts, possibly for use in an FMS attempt, this search can be aided by these tags. For example, in one memory retrieval, a car and a horse might be similar; with a different goal, a horse and a zebra would be a closer pair. This tagging allows a more flexible organization of knowledge than previous methods which were forced to carefully place items into the network to ensure that proper similarities were

	Physical	Mental	Social	Emotional	Temporal
Agents	person	consciousness	boss	Ares	entropy
Actions	walking	thinking	selling	loving	getting closer to March
Objects	rock	idea	teacher-student relationship	hatred	second
States	young	lack of knowledge	public dishoner	being angry	early

Figure 6: Knowledge representation grid

revealed when needed.

The knowledge organization scheme allows concept manipulation which is bounded in a reasonable fashion. Each change may leave a concept in the same conceptual grid cell (an *intracellular shift*) or it may cause the concept to cross a cell boundary (an *extracellular shift*). The system is biased against boundary crossings. As a result, conceptual movement within the same grid cell is the cheapest type to perform. Movement along *either* a row or a column is more difficult, and movement which must go along both is the most difficult. These costs act as a heuristic which guides a reasoner performing creative understanding—the greater the cost, the more conceptual movement has occurred. High amounts of conceptual movement indicates that it is likely that the result will be more bizarre than creative. This restricts how many iterations are permissible in the CUP algorithm. Each successive cycle creates concepts which are more and more distant from the original one. The first few iterations will result in concepts which fall within the same grid cell as the original concept. More cycles will create concepts which are shifted in the grid with respect to the beginning concept. By tracking this movement, the reasoner can decide when creative understanding has become too expensive to continue, based on the goals of the reasoner.

Implementation

The ideas discussed above are embodied in the ISAAC reading system, currently implemented at a level of functionality capable of reading the science fiction story *Men Are Different* (Figure 1). It is built in Common Lisp and runs on RS/6000 machines. ISAAC uses the KR frame package (Giuse 1990) for knowledge representation. More details of the ISAAC system, and its general reading capabilities can be found in (Moorman & Ram 1994).

Upon beginning to read *Men Are Different*, ISAAC realizes that the concept of robot it knows (an industrial tool) is insufficient to produce an understanding of the story robot. This realization arises when ISAAC is unable to reconcile its current definition of robot with the actions of the robot in the story, i.e., the predictions made by the existing conceptual definition are failing. Creative understanding is given the task of explaining the novel robot, with the existing concept and the novel one passed to the routine. Incorporation fails to create an understanding since there is little similarity between the two concepts of robots. Next, baseless analogy is attempted. ISAAC attempts to transfer the concept of *industrial robot* to the domain of volitional agents (since this is how the story robot appears to be acting). Unfortunately, ISAAC does not have enough background knowledge to succeed in this case. ISAAC then attempts problem reformulation, using FMS to produce a merged concept containing elements of the current robot concept and the best volitional agent it can retrieve from memory—a man. The result is a man-like, intelligent, volitional robot. The new concept maintains some characteristics of the original robot (it is made of metal, resistant to damage, uses sensors and feedback as a control mechanism, and so on.), but is sufficient to understand the robot in the story. ISAAC stores the new concept in memory as a *story-robot*.

ISAAC needs to perform similar creative understanding on the man in the story. In this case, ISAAC is attempting to explain what goals the robot might have possessed to cause it to act in the fashion that it did. As a result of this understanding attempt, ISAAC understands that the robot is seeing the man as more similar to itself than is warranted. The irony in the story can be seen as a dual shift within our knowledge grid. First, ISAAC is presented with a robot character acting as an *agent* rather than as a *physical object*; the ending is ironic because the narrator then treats the Man, a *physical agent*, as a *physical object* and disassembles him, thereby killing him.

Conclusions

While the CUP algorithm has been successful in assisting ISAAC in the understanding of a single story, more work must be done. Additional stories need to be added to the system in order to evaluate the impact of possessing unneeded information during creative understanding. It is already decided that the next story will be *Zoo* (Hoch 1978), a story similar to *Men...* in several ways but different enough to test expansion possibilities. Objective evaluation of ISAAC's creative performance is also needed. The best such test currently available, the *Torrance Tests of Creative Thinking* (Torrance 1988), is considered flawed in significant ways. While it is an objective test, opponents suggest that what it measures may be too abstract to apply to general creative behavior. We also intend to explore the relationship of *conceptual change* to creative understanding (e.g., Ram 1993; Chi 1993). Finally, we plan to demonstrate our claim that creative understanding issues are important to creative design by implementing a design algorithm based on the CUP and FMS algorithms.

While many researchers have tried to "demystify" creativity, most models of the process which have been put forth have been too vague to allow implementation. Seeing creativity as having two aspects, a generative side and an understanding side, has permitted us to explore new issues in creativity and develop interesting results. The four steps of the CUP model are sufficient for producing behavior which is judged creative. Additionally, the described FMS algorithm is a mechanism which has proven capable of producing novel understanding of concepts. In addition, our reading area has forced us to confront real-world applications of creativity. Humans are certainly aided in reading comprehension by possessing creativity (Popov 1993); by incorporating creative understanding into artificial systems, they gain the ability to learn from experiences with novel concepts and thus grow in scope.

References

Birnbaum, L. 1986. *Integrated Processing in Planning and Understanding*. Ph.D. Dissertation, Yale University. Research Report 489.

Bloch, A. 1963. Men Are Different. In Asimov, I., and Conklin, G., eds., *50 Short Science Fiction Tales*. New York: MacMillan Publishing Co.

Boden, M. A. 1991. *The Creative Mind: Myths and Mechanisms*. New York: BasicBooks.

Chi, M. T. H. 1993. Barriers to conceptual change in learning science concepts: A theoretical conjecture. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, 312–317.

Clement, J. 1989. Learning via model construction and criticism: Protocol evidence on sources of creativity in science. In Glover, J. A.; Ronning, R. R.; and Reynolds, C. R., eds., *Handbook of Creativity*. New York: Plenum Press.

Corrigan, R. W. 1979. *The World of the Theatre*. Glenview, IL: Scott, Foresman and Co.

Falkenhainer, B. 1987. Scientific theory formation through analogical reasoning. In *Proceedings of the Fourth International Workshop on Machine Learning*.

Finke, R. A.; Ward, T. B.; and Smith, S. M. 1992. *Creative Cognition: Theory, Research, and Applications*. Cambridge: MIT Press.

Gentner, D. 1989. Mechanisms of analogical learning. In Vosniadou, S., and Ortony, A., eds., *Similarity and Analogical Reasoning*. London: Cambridge University Press.

Giuse, D. 1990. Efficient knowledge representation systems. *The Knowledge Engineering Review* 5(1):35–50.

Hammond, K. J. 1989. CHEF. In Reisbeck, C., and Schank, R., eds., *Inside Case-Based Reasoning*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Hoch, E. D. 1978. Zoo. In Asimov, I.; Greenberg, M. H.; and Olander, J. D., eds., *100 Great Science Fiction Short Short Stories*. Garden City, NY: Doubleday.

Hofstadter, D., and McGraw, G. 1993. Letter Spirit: An emergent model of the perception and creation of alphabetic style. Technical Report 68, Indiana University.

Kass, A. M.; Leake, D. B.; and Owens, C. C. 1986. SWALE: A program that explains. In Schank, R. C., ed., *Explanation Patterns: Understanding Mechanically and Creatively*. Hillsdale, NJ: Lawrence Erlbaum Associates. 232–254. Appendix.

Kolodner, J. L., and Wills, L. M. 1993. Paying attention to the right thing: Issues of focus in case-based creative design. In *Proceedings AAAI-93 CBR Workshop*.

Lenat, D. 1990. The ubiquity of discovery. In Shavlik, J. W., and Dietterich, T. G., eds., *Readings in Machine Learning*. San Mateo, CA: Morgan Kaufmann Publishers, Inc.

Moorman, K., and Ram, A. 1994. A functional theory of creative reading. Technical Report GIT-CC-94/01, Georgia Institute of Technology.

Newell, A., and Simon, H. A. 1972. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

Popov, A. 1993. Creativity and reading comprehension. *Journal of Creative Behavior* 26(3):206–212.

Ram, A. 1993. Creative conceptual change. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, 17–26.

Sacerdoti, E. D. 1974. Planning in a hierarchy of abstraction spaces. *Artificial Intelligence* 5(2):115–135.

Schank, R., and Abelson, R. 1977. *Scripts, Plans, Goals, and Understanding*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Stewart, G. W. 1950. Can productive thinking be taught? *Journal of Higher Education* 21:411–414.

Sussman, G. J. 1973. A computational model of skill acquisition. Technical Report AI TR-297, MIT.

Thurstone, L. L. 1952. *Applications of Psychology*. New York: Harper & Row. chapter Creative Talent.

Torrance, E. P. 1988. The nature of creativity as manifest in its testing. In Sternberg, R. J., ed., *The Nature of Creativity*. Cambridge: Cambridge University Press.

Turner, S. 1992. *MINSTREL: A Computer Model of Creativity and Storytelling*. Ph.D. Dissertation, University of California, Los Angeles.

Wilensky, R. 1983. *Planning and Understanding*. Reading, MA: Addison-Wesley.

WKRP Episode 60. 1980. Venus and The Man. Written by Hugh Wilson, Directed by Rod Daniel.