In *Proceedings of the 18th Annual Cognitive Science Conference*, San Diego, CA, July 1996

The Role of Ontology in Creative Understanding

Kenneth Moorman and **Ashwin Ram**
Georgia Institute of Technology
College of Computing
Atlanta, GA 30332-0280
{kennethm,ashwin}@cc.gatech.edu

Abstract

Successful creative understanding requires that a reasoner be able to manipulate known concepts in order to understand novel ones. A major problem arises, however, when one considers exactly how these manipulations are to be bounded. If a bound is imposed which is too loose, the reasoner is likely to create bizarre understandings rather than useful creative ones. On the other hand, if the bound is too tight, the reasoner will not have the flexibility needed to deal with a wide range of creative understanding experiences. Our approach is to make use of a principled ontology as one source of reasonable bounding. This allows our creative understanding theory to have good explanatory power about the process while allowing the computer implementation of the theory (the ISAAC system) to be flexible without being bizarre in the task domain of reading science fiction short stories.

Introduction

Over the last several years, we have been developing a functional theory of creative reading (see, e.g., Moorman & Ram, 1994a; Moorman & Ram, 1994b). The theory is being implemented in the ISAAC reading system, which reads short science-fiction stories. An important portion of the theory involves the ability to understand novel concepts, which we call *creative understanding*. The theory has now been developed to a point that we are able to carefully evaluate the precise role that ontology is playing in the process of creative understanding. The questions to answer include: *What are the theory's ontological commitments?*; *Where do they come from?*; *What power do they impart on the theory?*; and *What limitations do they create?*.

The Creative Understanding Process

In order to comprehend texts with novel concepts, we hypothesize the use of a *creative understanding process* (Moorman & Ram, 1994c), made up of four tasks: *memory retrieval*, *analogy*, *base-constructive analogy*, and *problem reformulation*. When a concept is presented to the reasoner, the memory task is used to discover if it is previously known. If it is and if the preexisting knowledge is sufficient to allow the concept to be explained or to make predictions about the concept, then understanding is successful. If there is no existing concept or if existing concepts do not allow explanation and prediction, then the remaining three tasks are called upon to produce a creative understanding of the novel concept. Analogy attempts to find a functionally-consistent mapping between concepts which were retrieved and the one being considered. Base-constructive analogy is called upon if no concepts exist with which to draw an analogy from; in this case, a new base may be dynamically created for use in an analogy by appealing to known concepts. If the reasoner realizes that the current understanding focus is not the correct one, problem reformulation is invoked; this acts to refocus the reasoner so that successful understanding is more likely to occur. The process iteratively continues until either a successful understanding is achieved or the reasoner is willing to "give up" on understanding.

While performing the creative understanding, the reasoner may need to manipulate existing concepts in order to achieve success. It is here that some constraint must be maintained to avoid the problem of bizarreness. Bizarre understanding occurs when an understanding is reached which is not useful to the reasoner. For example, suppose a reader begins a story which starts with a character, John, slapping another character, Mary. One predictive question which could be generated is why John would do such a thing. One understanding of the event is that it occurred because space aliens from the Vega system has controlled John's mind in an experiment to test human reaction to violent behavior. And, in the scope of a science-fiction story, this might even be the eventual solution. However, assuming this at this point in the story with the information provided is premature and not useful for comprehending the rest of the narrative. Thus, we classify this understanding as bizarre.

There are two primary elements which we use to guide the modification process: satisfaction and ontology. The *satisfaction* solution bounds the creative understanding process by virtue of it existing within a larger cognitive task—in our case, reading. Understanding can be said to be successful as soon as the quality of explanation or prediction provided is high enough to allow the reading process to continue. As an example, consider the world of *Star Trek*. The *warp drive* is an example of a concept which must be understood to comprehend the stories. But, if a particular story has the warp drive functioning normally, then the reader can be satisfied by simply understanding that it provides a way to get from point A to point B really, really quickly. On the other hand, if a story involves the breakdown of the antimatter in the warp core, a reader has to reach a higher level of understanding with respect to the warp drive, or the comprehension of the story will be impaired. In this fashion, the reader will never push a concept to the level at which it becomes bizarre.[1] However,

[1]Note that we are not saying that this is the only strategy that a reader may use. Readers with an interest in faster-than-light travel,

the satisfaction solution is only part of the answer. Added to it is the *ontology* solution; that is, it is necessary to create a knowledge organization which reflects the world of the reasoner, allows creative manipulations to occur, and acts to prohibit bizarre ones. This aspect is the focus of the remainder of this paper.

Motivating Examples

Consider the sentence, *John was a bear*, which could be the opening line of a story. Since there is no additional information, at this point, from the text itself, there are numerous interpretations:

- John (a human) acts like a bear.
- John (a human) has been transformed into a bear.
- John (an agent) is a were-bear.
- John (a bear) is a bear.

For comprehension to be successful, a reader must formulate an interpretation which is consistent given other information possessed by the reasoner. Although all of the above interpretations are *possible*, some may be more probable than others in a given context. How is a reasoner to select the "best" interpretation given all other information available?

A reader could simply "skip" this sentence and assume that later sentences will disambiguate the confusion. Unfortunately, this simply pushes the problem back a level—at some point, an interpretation must be created (given the problems of keeping all possible options simultaneously active in memory). Also, human readers *are* able to make a choice between the alternatives if you stop the reading process and query them. Some mechanism must be allowing the selection of a "best" choice at some point in the reading process.

One option is simply to allow an arbitrary choice. On the other end of possibilities, a designer of a reading system could encode a set of rules which would allow the system to select one choice over the other. Unfortunately, the former of these options leads to potentially bizarre results occurring (how often, after all, is a story beginning with the example sentence going to be about were-bears?); the latter option leads to a huge knowledge engineering problem and the potential to overlook possibilities which would then cause less than optimal performance. We follow a more general approach of allowing a basic ontology to constrain the problem. This allows *creative understanding* of concepts to occur, while preventing bizarre interpretations from being formed.

The Ontological Grid

All knowledge within the ISAAC system is organized by a top-level ontological grid, consisting of twenty ontological categories (Figure 1 shows this breakdown and example concepts from each category). One axis of the ontology grid represents the *type* of a concept: *action*, *agent*, *state*, and *object* (Domeshek, 1992; Schank & Abelson, 1977). The other dimension represents the *domain* of a concept: *physical*, *mental*, *social*, *emotional*, and *temporal*. A reasoner possesses knowledge concerning the cells of the grid themselves, as well as knowledge about the particular rows and

for instance, may try to extend their understanding of the warp drive even if a more in-depth understanding is not needed for a particular story.

	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 dishonor	being angry	early

Figure 1: Knowledge grid

columns (e.g., knowledge about physical types in general, or knowledge of objects in general). While performing creative understanding, a concept may need to be *transitioned* from one cell to another. If a principled method can be developed to bound these transitions based on the ontology, then a reasonable bound on the creative understanding process will then exist.

Representation of Concepts

The movements of concepts within the grid is partially dependent on the representation format we make use of in the theory. Concepts must be capable of being combined with other concepts (or parts of concepts), which means a fairly flexible knowledge representation system is required. This aspect of the work draws heavily from research in artificial intelligence concerning exactly what elements should make up a "proper" knowledge representation. Although the ideas of both semantic networks and frame representations go back decades (see, for example, Quillian, 1966 and Minsky, 1975, respectively), the structure of knowledge within many artificial intelligence systems is ad-hoc. A notable exception is that of Wilensky's work (1986) with KODIAK, which motivated much of our current approach.

While a concept is physically created and stored within the system as a frame-like entity, this is more for retrieval ease than anything else. The proper abstraction of the storage used is that a concept is represented as a node in a graph-like structure (Barsalou, 1992; Wilensky, 1986). The traditional role of slots exists in the system as data objects on the node, which contain pointers to other nodes in the memory. These design decisions were driven to allow flexible storage and retrieval of concepts via a spreading activation model of memory (Francis & Ram, 1993).

Consider the robot example in Figure 2. HEIGHT-IS is a pointer to a relationship schema containing general information about *height*. In particular, all such relationships will contain a domain and codomain, represented in the framelike representation as simply the frame-name and the filler for the indicated slot. Furthermore, the role of traditional slot-names, which are assumed to be relationships in our approach, can be described as states. Thus, HEIGHT-IS is part of the state description for an object.

ROBOT-12	
:IS-A	ROBOT
:ROLES	{INDUSTRIAL-TOOL WEAPON TRANSPORT}
:FUNCTION	INDUSTRIAL-TOOL-5
:IS-MADE-OF	TITANIUM-7
:IS-POWERED-BY	ELECTRICITY-6
:CAN-LIFT	WEIGHT-10
:COLOR-IS	GREY-8
:HEIGHT-IS	HEIGHT-11
:P-ATTRIBUTES	{:IS-MADE-OF :IS-POWERED-BY :CAN-LIFT}
:S-ATTRIBUTES	{:COLOR :HEIGHT}
:EXPLANATION	EXPLANATION-9

Figure 2: Represented concept

Additionally, each concept is tagged with the current function it is being viewed as performing, as well as a set of possible functions it is known to be capable of performing. This fact is utilized to achieve flexible memory retrieval—during one search through memory, a car and a horse might be similar; with a different function in mind, a horse and a zebra would be more closely related. The process of function tagging follows the work of Barsalou on *ad hoc categories* (1989); rather than having all categories predefined, a reasoner can create temporary categories by collecting concepts with similar functions. The *primary attributes* of a concept determine how it achieves its function, while the *secondary attributes* represent additional information. If the current function of a concept is changed (for example, a reasoner stops viewing the horse as an animal and starts viewing it as a mode of transportation), it might be necessary to repartition the primary and secondary attributes. And, by considering novel combinations of primary and secondary attributes, it is possible to hypothesize novel functions for a concept.

Manipulations of Concepts

When an existing concept needs to be manipulated during the course of a creative understanding episode, there are three basic outcomes with respect to the grid. The concept can be manipulated, yet remain in the same grid cell as when it started. For example, a reasoner may use a *horse* to understand the concept of a *zebra*. Second, a concept may transition along a row or a column. If a reasoner uses their knowledge of physical concepts in order to understand social ones (a boss blocking your promotion), this is an example of this shift. Finally, a transition may occur which moves the concept in terms of both axes. A reasoner understanding something like *His mind was a steel trap* is making use of this dual transition—a physical object is transitioned to a mental state. These three possibilities represent a simple ordering of the amount of cognitive work required to manipulate any given concept. Finally, a set of high-level heuristics is needed to bound the possible motion within the grid. These are:

- Physical types can become transitioned to other domains more easily than other domain types can be transitioned to physical. This was a recent empirical discovery resulting from experimentation performed on the system. We theorize that since humans are physical entities with a great deal of experience with other physical entities, it is "easier" to believe in the existence of a novel, non-physical entity formed from a physical analogue than it is to accept the creation of a new type of physical entity. Consider, *John saw the days fly by*. Is this a novel use of *saw* and *fly* created by altering physical concepts into the temporal domain, or is it a novel use of *days* created by considering a temporal object as a physical one?
- An object may transition to an action by creating an action which captures a function of that object and vice versa. English, in particular, tends to have many lexical examples of this. A fax is the thing you send when you fax someone. A (Star Trek) transporter is the device used to transport material from one location to another.
- An object may transition to a state by creating a state which captures a primary attribute for that object, and vice versa. Through this transition, we get many common similes and metaphors, such as *Hungry as a bear* and *As good as gold*.
- Agents and objects can easily transition between each other. This results from two observations. First, agents exist as embodied entities in the world (Johnson, 1987), explaining the agent to object transition. For example, one may treat John as a physical object. Second, it is possible to view objects as though they possess intention (Newell, 1981), enabling the object to agent transition. For instance, a thermostat may be thought of in terms of agency, i.e., it *wants* to keep the house at a constant temperature.
- Make the minimal changes necessary. This is simply a general rule, ala Occam's Razor. It results from the earlier discussed idea of satisfaction ultimately driving the creative understanding process—stop the process once you have a "good enough" understanding to allow the higher cognitive task to continue.

By combining the three basic movement types with the high-level heuristics, we get an ordering of the amount of *cognitive effort* required to manipulate concepts (from easiest to most difficult):

1. Concepts may transition within a single cell.
2. Agents may be treated as objects and objects may be treated as agents.
3. Concepts may vertically transition according to the modification heuristics.
4. Physical types may transition to other domains (horizontal motion).
5. Other domain types may transition to the physical domain (horizontal motion).
6. Combinations of 2–5 may occur.

Within this ordering, however, operations which result in the minimal changes are preferred over those which are more complex.

Implementation Details and Examples

The ISAAC system is the a computer system which instantiates our creative reading theory. ISAAC is written in Common Lisp, using the KR frame package (Giuse, 1990) for knowledge representation, the Garnet package (Myers, 1988) for graphical input and output, and the COMPERE system (Mahesh, 1993) for comprehending individual sentences. While the system is specifically designed to be a testbed for our theory of creative understanding, the complete reading theory is implemented. The theory describes reading as being made up of six *supertasks*, or large collections of functionally related tasks, including *story structure comprehension*, *scenario comprehension*, *memory management*, *sentence processing*, *explanation*, and *metacontrol*. More information on the reading aspects of the model can be found in Moorman and Ram (1994b). ISAAC is currently capable of reading three short, previously published, science fiction stories (one to three pages), as well as several paragraph synopses of *Star Trek: The Original Series* episodes (Asherman, 1989). In addition, several small examples have been studied, outside the context of reading complete stories. We now present several of these to show the extent which the ontology aids the creative understanding process.

First, we return to the original example: *John was a bear.* With all the possibilities, the least movement occurs if we simply allow *John* to be the name of a particular *bear*. Thus, this is the version that ISAAC prefers, if no additional information from the story is provided.[2]

A second example comes from the Meta-AQUA system (Ram & Cox, 1994) which reads a story involving a drug-sniffing dog. Meta-AQUA initially knows only that dogs will bark at agents which threaten them. But, in the story, a dog is barking at a suitcase. In ISAAC, the system is presented with two possibilities—its knowledge of dogs is wrong or its knowledge of suitcases is. The first involves altering an existing physical agent to create a variant of it, an intracellular movement. The second involves shifting a physical object to the physical agent cell, a vertical movement. The intracellular movement is preferred.

In the story *Men Are Different* (Bloch, 1963), a robotic archaeologist is studying the destroyed civilization of mankind. The story is presented as a first-person narrative. ISAAC is aware that narrators, archaeologists, and protagonists are all known to be human; robots are industrial tools; but the narrator, archaeologist, and protagonist of the story is known to be a robot. ISAAC can create a new type of robot which embodies agent-like aspects, or it can change the definitions of narrators, archaeologists, protagonists, and the actions in which they may participate. The new robot concept represents a more minimal change.

The final example involves the story *Zoo* (Hoch, 1978), in which the reader is presented with an intergalactic zoo which travels from planet to planet, giving the inhabitants of those planets a chance to view exotic creatures. At the end of story, however, the reader is shown that the true nature of the intergalactic ship—it is an opportunity for the "creatures" on the ship to visit exotic planets, protected from the dangerous inhabitants by the cages they are in. To understand the new zoo, the system draws an analogy between the known zoo and the novel one. The result, then, is simply a shift from one physical object to another physical one.

[2]This may seem counter-intuitive; to most people, this example would make more "sense" if interpreted as *John is bear-like*, a metaphorical usage. Remember, however, that these ontological constraints have the most effect when no other information is known. This information can be background knowledge or story knowledge, so a reasoner already familiar with the bear metaphor may retrieve that interpretation instead of this default one. It is important to note that our approach handles metaphor as a normal part of the understanding cycle. Since metaphors are pervasive in language (see Carbonell, 1982; Johnson, 1987; Lakoff & Johnson, 1980, for example), we consider this "unified" handling of non-metaphors and metaphors to be an important feature of our overall theory.

Drawbacks and Limitations

Extra Knowledge Engineering

There are three known problems with the approach which we are currently pursuing. The first is simply the need to correctly mark each concept in memory as belonging to one of the twenty top-level categories. This means some additional work is required when knowledge is being represented for ISAAC initially. However, the level of increased work has not been prohibitive to the project.

Origin of the Categories

There is some concern that the ontological categories being used may be too limited to allow an accurate model of the world. But, we have been driven by functional constraints on the creative understanding process, so we can say that our ontology is sufficient for the cognitive tasks in the theory. Added to this, we have also been motivated by prior psychological research concerning ontology. This has mainly been studied by developmental psychologists who attempt to explain the changes which take place to a child's ontology as they mature, as well as attempting to explain what, if any, ontology may exist from birth.

The first area to consider is that of what ontologies exist. Two important ontological distinctions have been studied by a number of researchers. The *physical–immaterial distinction* is an important one which arises relatively early in the development pattern of normal children, although it continues to be refined and sophisticated as they age (Carey, 1992; Carey & Spelke, 1994). Another major ontological division is the *object–event distinction*—some things in the world are objects (like rocks, people, etc.) while other things are events (such as walking) (Carey, 1992).

The second aspect of the previous research to note is the recognition that shifting basic ontological categories is difficult. Numerous researchers have noticed and theorized about this (e.g., Carey, 1992; Chi, 1993), but one of the best descriptions of the possible range of changes comes from Thagard (1992). In his framework, there are nine degrees to conceptual change, ranging from the simple addition of new instances of known concepts to the complete reorganization of the ontological hierarchy. The lower levels are far easier to perform; it is only with growing ability and sophistication that a reasoner will achieve the reorganization level.

Finally, viewing various ontological experiments over the course of the research has prompted some researchers to claim that certain types of ontological categories are fundamental to human reasoning and would, therefore, need to be accurately modeled. For example, Brewer (1993) suggests that a rich ontology is needed, consisting of (at least) natural kinds, nonexisting natural kinds, artifacts, social entities, psychological entities, and abstract entities. While our approach does not duplicate this exactly, the basic ideas are consistent.

Category Membership

Finally, there is potentially the most dangerous problem facing the theory—the problem of deciding category membership for a concept if the system has manipulated it. While it is easy in theory to say that a concept may shift across conceptual grid cell boundaries, it is sometimes difficult in practice to determine where its new location should be.[3] Consider the concept of *time travel*, for a moment. A reasoner with no prior knowledge can utilize information concerning *physical transport actions* and the *temporal column* to transform a physical transport concept into a temporal transport actions; i.e., a horizontal shift from the physical action cell to the temporal action cell. Notice that both actions have physical agents for their initiating actors and physical objects for their transported objects. Now, consider a device capable of performing time travel, namely a *time machine*. Again, a reasoner can start with knowledge of *physical transport machines* and the *temporal domain* and develop the concept of a time machine. But, in this case, a physical object has remained a physical object. Both manipulations used similar knowledge; one resulted in a horizontal shift while the other resulted in an intracellular one. For the moment, the problem is circumvented by appealing to the *minimum change* heuristic, but it is certainly an area of research to be explored.

Conclusion

The knowledge representation and ontological commitments of the ISAAC framework allows us to reasonably bound the creative understanding process without having to resort to a large list of rules and their exceptions. Additionally, there is the added knowledge engineering benefit of not having to worry about the precise placement of concepts within the implemented system—as long as the top-level ontological category is correct, the theory enables the system to function with a great deal of flexibility and robustness. As larger artificial intelligence systems are developed which are expected to be robust, real-world reasoners, it will become important for researchers to be aware of the ontological reality of the world in which their theories exist. Indeed, by taking advantage of the ontology, researchers may discover that the task of engineering large, robust systems is made somewhat easier.

Acknowledgments

This work was supported by a Fannie and John Hertz Foundation fellowship and by the Georgia Institute of Technology. The authors would like to recognize the efforts of Thomas Tiller, a visiting student who worked on aspects of the ISAAC theory, and the time and effort spent by members of the IGOR research group to read and offer useful comments on this work.

References

Asherman, A. (1989). *The Star Trek Compendium*. Pocket Books, New York.

Barsalou, L. W. (1989). Intraconcept similarity and its implications for interconcept similarity. In Vosniadou, S. and Ortony, A., editors, *Similarity and Analogical Reasoning*. Cambridge University Press, Cambridge.

Barsalou, L. W. (1992). Frames, concepts, and conceptual fields. In Lehrer, A. and Kittay, E. F., editors, *Frames, Fields, and Contrasts*. Lawrence Erlbaum Associates, Publishers, Hillsdale, NJ.

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

Brewer, W. F. (1993). What are concepts? Issues of representation and ontology. In Nakamura, G. V., Taraban, R., and Medin, D. L., editors, *Categorization by Humans and Machines*. Academic Press, Inc., San Diego.

Carbonell, J. G. (1982). Methaphor: An inescapable phenomenon in natural-language comprehension. In Lehnert, W. G. and Ringle, M. H., editors, *Strategies for Natural Language Processing*. Lawrence Erlbaum Associates, Publishers, Hillsdale, NJ.

Carey, S. (1992). The origin and evolution of everyday concepts. In Giere, R. N., editor, *Cognitive Models of Science*. University of Minnesota Press, Minneapolis.

Carey, S. and Spelke, E. (1994). Domain-specific knowledge and conceptual change. In Hirshfeld, L. A. and Gelman, S. A., editors, *Mapping the Mind: Domain specificity in cognition and culture*, pages 169–200. Cambridge University Press, Cambridge, MA.

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*, pages 312–317.

Domeshek, E. (1992). *Do the Right Thing: A Component Theory of Indexing Stories for Social Advice*. PhD thesis, Yale University, New Haven, CT.

Francis, A. G. and Ram, A. (1993). Conquering the utility problem: Designing a memory module for all seasons. Unpublished draft.

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

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

Johnson, M. (1987). *The body in the mind: Bodily basis of meaning, imagination, and reason*. University of Chicago Press, Chicago.

Lakoff, G. and Johnson, M. (1980). *Metaphors We Live By*. University of Chicago Press, Chicago, IL.

Mahesh, K. (1993). A theory of interaction and independence in sentence understanding. Technical Report GIT-CC-93/34, Georgia Institute of Technology.

Minsky, M. L. (1975). A framework for representing knowledge. In Winston, P. H., editor, *The Psychology of Computer Vision*. McGraw-Hill, New York.

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

Moorman, K. and Ram, A. (1994b). Integrating reading and creativity: A functional approach. In *Proceedings of the Sixteenth Annual Cognitive Science Conference*.

[3]A similar problem arises in research involved with the formation of novel concepts (see, e.g., Shoben, 1993; Ward, 1995; Wisniewski & Medin, 1994). If a novel concept is formed using existing concepts, what category should the new concept belong to? For example, if *pet* and *fish* are combined, where in a knowledge hierarchy should the resulting concept be placed?

Moorman, K. and Ram, A. (1994c). A model of creative understanding. In *Proceedings of the Twelfth Annual AAAI Conference*.

Myers, B. A. (1988). The GARNET user interface development environment: A proposal. Technical Report CMU-CS-88-153, Carnegie-Mellon University.

Newell, A. (1981). The knowledge level. *AI Magazine*, 2:1–20.

Quillian, M. R. (1966). Semantic memory. In Minsky, M., editor, *Semantic Information Processing*, pages 227–270. MIT Press, Cambridge, MA.

Ram, A. and Cox, M. (1994). Introspective reasoning using meta-explanations for multistrategy learning. In Michalski, R. and Tecuci, G., editors, *Machine Learning: A Multistrategy Approach*, volume IV. Morgan Kaufman Publishers, San Mateo, CA.

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

Shoben, E. J. (1993). Non-predicating conceptual combinations. In Nakamura, G. V. and Medin, D. L., editors, *Categorization by Humans and Machines. The Psychology of Learning and Motivation*, volume 29. Academic Press, San Diego, CA.

Thagard, P. (1992). *Conceptual Revolutions*. Princeton University Press, Princeton, NJ.

Ward, T. B. (1995). What's old about new ideas? In Smith, S. M., Ward, T. B., and Finke, R. A., editors, *The Creative Cognition Approach*. MIT Press, Bradford Books, Cambridge, MA.

Wilensky, R. (1986). Knowledge representation—a critique and a proposal. In Kolodner, J. L. and Riesbeck, C. K., editors, *Experience, Memory and Reasoning*, chapter 2, pages 15–28. Lawrence Erlbaum Associates, Hilldale, NJ.

Wisniewski, E. J. and Medin, D. L. (1994). The fiction and nonfiction of features. In Michalski, R. and Tecuci, G., editors, *Machine Learning: A Multistrategy Approach*, volume IV. Morgan Kaufmann Publishers, San Francisco.