A theory of questions and question asking*

Ashwin Ram

College of Computing
Georgia Institute of Technology
Atlanta, Georgia 30332-0280
(404) 853-9372
E-mail: ashwin@cc.gatech.edu

Abstract

This article focusses on knowledge goals, that is, the goals of a reasoner to acquire or reorganize knowledge. Knowledge goals, often expressed as questions, arise when the reasoner's model of the domain is inadequate in some reasoning situation. This leads the reasoner to focus on the knowledge it needs, to formulate questions to acquire this knowledge, and to learn by pursuing its questions. I develop a theory of questions and of question asking, motivated both by cognitive and computational considerations, and I discuss the theory in the context of the task of story understanding. I present a computer model of an active reader that learns about novel domains by reading newspaper stories.

* *The Journal of the Learning Sciences*, 1(3&4):273–318, 1991.

1 Introduction

In this article, I discuss a theory of questions, viewed as a basis for understanding and learning. Most teachers have had the experience of thinking that their students understood some material because they were asking the "right questions." Children ask questions constantly in an attempt to understand and learn about the world around them. Even as adults, we express our curiosity in the form of questions, often to ourselves, as we wonder about novel situations, explore new hypotheses, and become interested in various issues. The ability to ask questions, it seems, is central to the processes of reasoning, understanding and learning. In this article, I formalize the basis for these processes by developing a theory of questions and question generation. I explore different kinds of questions, how a reasoner might come to ask them, and the effects of having asked these questions. Also presented is a computer program that asks intelligent questions in an attempt to reason and to learn about its domain.

Story understanding programs are often designed to answer questions to demonstrate that they have adequately understood a story (e.g., [Lehnert, 1978]). In contrast, I claim that asking questions is central to understanding. The underlying theme of this research is the goals of the reasoner and the interaction of these goals with reasoning processes. In particular, I focus on *knowledge goals*, that is, the goals of a reasoner in learning by acquiring and organizing new knowledge and by reorganizing existing knowledge in memory. I argue that knowledge goals, often expressed as questions, arise when the reasoner's model of the domain is inadequate in some reasoning situation. This leads the reasoner to focus on what he or she needs to know, to formulate questions in acquiring this knowledge, and to learn by pursuing these questions.

Rather than presenting details of learning algorithms (see [Ram, 1992]), this article is concerned more with the relationship between questions and learning and with the nature of the questions themselves. I discuss the sources of questions, the types of questions arising from different reasoning tasks, the process of generating questions, and the process of learning by answering questions. I also discuss a computer model of question-driven understanding and learning, and the model's implementation in a natural language understanding domain. The main point of the research is to create a model of a dynamic understander that is driven by its questions or goals to acquire knowledge. Rather than being "canned," the understander is always changing as its questions change. Such an understander reads similar stories differently and forms different interpretations as its questions and interests evolve. The intent is not to design a system that can acquire the "right" understanding of a topic, but one that is able to wonder and to ask questions about the unusual aspects of its input. As it learns more about the domain, the system asks better and more detailed questions. This kind of questioning forms the origins of creativity; rather than being satisfied with available explanations, a creative person asks questions and explores the explanations in novel ways.

In my model, question-driven information-seeking forms the basis for active, goal-based learning processes. Question generation is the process of identifying what the reasoner needs to learn. Learning occurs incrementally as the reasoner's questions are answered through experience. My system's "experience" corresponds to reading newspaper stories. Although the computer model is being used to explore cognitive issues such as the ones previously mentioned, there are also practical benefits of a system that can represent and reason explicitly about its own goals. Such a system can focus its limited resources on relevant aspects of its environment while paying less attention to irrelevant ones. This allows it to spend more time drawing inferences that are relevant and useful to its goals. This is important in reasoning situations in which the reasoner might draw a combinatorially large set of inferences and also in learning situations in which it is impractical to focus attention on every aspect of a situation and remember every novel aspect. The reasoner needs a principled way to determine which inferences are worth drawing or what is worth learning. In order to ensure that the system does not spend its limited resources trying to infer everything it can, its knowledge goals are used to focus the inferencing and learning process on information that is useful to the goals of the system. What the system learns from the story depends on what it needs to learn, that is, on its knowledge goals. Our computer system reasons about its own reasons for learning and, hence, does not need to generalize everything in the hope that it might eventually be useful.

There is an additional benefit to be derived from the study of questions, especially those questions that are asked in novel reasoning situations. A theory of question generation will help us to develop improved theories of teaching. Questions arise from an interaction between the interests and goals of the understander and the information provided by the environment. In an educational situation, the "environment" may be a teaching program. To design educational environments that facilitate learning, we need to pay attention to the question-asking process that students go through when interacting with the teaching program. For example, in Scardamalia and Bereiter's [1991] Teacher C model, the teacher is concerned with "helping students formulate their own goals, do their own activating of prior knowledge, ask their own questions, direct their own inquiry, and do their own monitoring of comprehension" [p. 39]. We also need to consider the motivational aspects of teaching. If we can put students into situations in which they *want* to find something out, they will be better motivated and better able to focus their attention on the relevant information. The solution to this problem lies in understanding the nature of the "learning goals" of the student. These goals are often manifested as *questions*, and may be thought of as information goals or knowledge goals. This ties into my formulation of questions as expressions of goals to collect knowledge, which arise from underlying needs to know that piece of knowledge.

Thus understanding the nature of questions, and their role in learning, is an important and fundamental problem in cognitive science and has implications for theories of learning and education.

1.1 The role of questions in understanding and learning

The basic assumption of our theory, which is called *question-driven understanding*, is that asking questions is central to understanding. To illustrate what this means, consider the following story (New York Times, April 14, 1985):

> **S-1: Boy Says Lebanese Recruited Him as Car Bomber.**
> JERUSALEM, April 13 — A 16-year-old Lebanese was captured by Israeli troops hours before he was supposed to get into an explosive-laden car and go on a suicide bombing mission to blow up the Israeli Army headquarters in Lebanon. ...
> What seems most striking about [Mohammed] Burro's account is that although he is a Shiite Moslem, he comes from a secular family background. He spent his free time not in prayer, he said, but riding his motorcycle and playing pinball. According to his account, he was not a fanatic who wanted to kill himself in the cause of Islam or anti-Zionism, but was recruited for the suicide mission through another means: blackmail. [p. A1]

If one wants to learn more about the motivations of the terrorists in the Middle East, this story is interesting because it is anomalous. The usual stereotype of the Shiite religious fanatic does not hold here. Instead, this story raises many new questions. Some of the questions that were voiced by a class of graduate students when this story was read to them:

1. Why would someone commit suicide if he was not depressed?
2. Did the kid think he was going to die?
3. Are car bombers motivated like the Kamikaze?
4. Does pinball lead to terrorism?
5. Who blackmailed him?
6. What fate worse than death did they threaten him with?
7. Why are kids chosen for these missions?
8. Why do we hear about Lebanese car bombers and not about Israeli car bombers?

9. Why are they all named Mohammed?
10. How did the Israelis know where to make the raids?
11. How do Lebanese teenagers compare with American teenagers?

Some of these questions seem reasonable, (e.g., "Did the kid think he was going to die?"), but some are rather silly in retrospect (e.g., "Does pinball lead to terrorism?"). Some, though perfectly reasonable, are not central to the story but relate to other issues that a given student was reminded of, was wondering about, or was interested in (e.g., "Why do we hear about Lebanese car bombers and not about Israeli car bombers?").

The claim is that an understander has questions already extant in memory before it begins to read a story. These questions are left over from the understander's previous experiences. As the understander reads the story, it remembers these questions and thinks about them again in a new light. This raises further questions for the understander to think about. Many of these questions seek *explanations*, which are knowledge structures that allow the understander to answer its questions based on a causal understanding of the situation (e.g., "Kids are chosen because they are more gullible"). Explanations, in turn, can give rise to further questions (e.g., "Are Lebanese teenagers more gullible than American teenagers?").

Ultimately, the understander is left with several new questions that it may or may not have asked before. Certainly, after reading the blackmail story, one expects to have several questions representing issues one was wondering about that were not resolved by the story. For example, in this story, it turns out that the boy was blackmailed into going on the bombing mission by a terrorist group that was threatening his parents. This makes one think about the question "What are family relations like in Lebanon?" This question remains in memory after reading the story. To the extent that one is interested in this question, one will read stories about the social life in Lebanon, and one will relate other stories to this one. To cite another example, one of the students in the class repeatedly related the story to his readings on the IRA because he was interested in similar issues about Ireland.

Understanding is a process of relating what one reads to the questions that one already has. These questions represent the knowledge goals of the understander, i.e., the things that the understander wants to learn [Dehn, 1989; Hunter, 1989; Ram, 1987; Ram, 1989; Ram, 1990c; Schank, 1986; Schank and Ram, 1988]. The purpose of reading is to find answers to these questions and, thus, to arrive at a more complete understanding of the issues one is interested in. However, while doing this, many new questions are often raised. These questions are stored in memory and, in turn, guide the understanding of future stories and affect the interpretations that are drawn. This process is shown in figure 1.

Although this type of reasoning may not be conscious, learning is motivated by a reasoner's goals and interests. When the reasoner encounters difficulties during understanding, planning, or any other task, it remembers the nature of these difficulties and learns in order to perform its tasks better in the future. The knowledge goals of the reasoner, which arise from these very difficulties, are used to focus the learning process. Our model is very different from other approaches that rely on properties of the domain to determine what needs to be learned because it relies on the goals of the reasoner. For example, one might propose a rule, similar to that discussed by DeJong [1983], that the understander generalize a new schema whenever it reads a story in which a preservation goal (P-GOAL) is violated in a novel manner. But this should be so only if noticing violations of this P-GOAL is actually useful to the program. Any such rule must make a statement about the goals of the program, not just about the content of the domain. A similar argument can be made for the use of knowledge goals, or questions, to focus inference generation for understanding, explanation or diagnosis [Ram, 1990c; Ram and Hunter, 1992; Ram and Leake, 1991].

In the following sections, I argue that a goal-based model of learning is a plausible account of human behavior, and has computational advantages for the design of learning programs as well. The model raises two sets of issues:

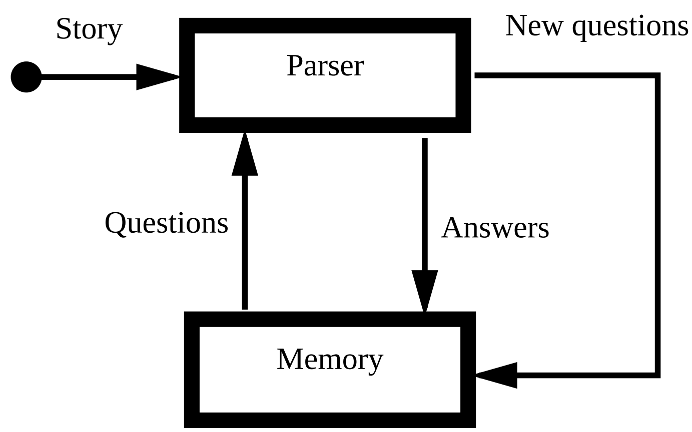

Figure 1: Question-driven understanding. In contrast to the traditional view of understanding as a "story in, representations out" process, we view understanding as a "questions + story in, answered questions + new questions out" process. This process is elaborated in figure 2.

- **Content:** What kinds of questions are there? How does the reasoner know which questions to ask?
- **Process:** What difference do questions make? What effect do they have on the understanding process? How do they affect what one learns? How are questions managed in memory?

These issues are discussed from a cognitive perspective as well as a computational one.

1.2 Cognitive motivations: Questions as a basis for learning

When comparing the way people read newspaper stories with the way computer programs typically read them, notice the following differences:

Subjectivity: People are biased. They interpret stories in a manner that suits them. They jump to conclusions. Computer programs, on the other hand, are usually designed to read stories in an objective manner, and to extract the "correct" or "true" interpretation of a story to the extent that they can.

Variable-depth parsing: People do not read everything in great detail. They concentrate on details that they find relevant or interesting and skim over the rest. In contrast, computer programs are designed to attend to every aspect of a story that is within the scope of their knowledge structures. Consequently, they either process the entire story in great depth, or skim everything in the story. They cannot decide which aspects to process in detail and which ones to ignore.

Learning and change: People change as they read. They never read the same story twice in the same way. They notice different things during the second reading, or they simply get bored. After reading a story, they interpret other similar stories differently. Typical computer programs, in contrast, are not adaptive; they always read a given story the same way.

What makes people different from computer programs? What is the missing element that our theories do not yet account for? The answer is simple: People read newspaper stories for a reason: to learn more about what they are interested in. Computers, on the other hand, do not. In fact, computers do not even have interests; there is nothing in particular that they are trying to learn when they read. If a computer program is to be a model of story understanding, it should also read for a "purpose."

Of course, people have several goals that cannot be attributed to computers. One might read a restaurant guide in order to eventually satisfy hunger or entertainment goals or to find a good place for a business lunch. Computers do not get hungry, and computers do not have business lunches.

However, these physiological and social goals give rise to several intellectual or cognitive goals. A goal to satisfy hunger gives rise to goals to find information: the name of a restaurant that serves the desired type of food, how expensive the restaurant is, the location of the restaurant, and so on. These are goals to acquire information or knowledge and are knowledge goals. These goals can be held by computers too; a computer might "want" to find out the location of a restaurant and might read a guide in order to do so in the same way as a person would. Although such a goal would not arise out of hunger in the case of the computer, it might well arise out of the "goal" to learn more about restaurants.

In other words, knowledge goals also arise from the desire to learn, to pursue one's intellectual interests, to improve one's model of the world. These goals can be viewed as questions about the domain of interest. To be interested in terrorism, for example, is to have many questions about the various aspects of terrorism and to think about these questions in the context of input data, such as newspaper stories about terrorist incidents. The point of reading these stories is to answer one's questions as well as to reveal flaws or gaps in one's model of terrorism in order to improve this model. These gaps give rise to new questions that, in turn, stimulate further interest in terrorism. In this sense, both computers and people can be "interested" in terrorism.

In contrast with people, a computer has only one underlying goal: to learn and improve its world model.[1] However, this (and, in the case of people, other physical and social goals) gives rise to knowledge goals that then drive the understanding process. Understanding consists of using questions to focus attention on relevant or interesting aspects of the input, answering these questions using information provided by the input, and asking new questions based on unusual or unexpected aspects of the input. This process is illustrated in figure 2.

My theory of questions and question-asking is based on a functional account of the role of questions in understanding and learning. However, similar arguments have been made based on empirical results in psychology and education. For example, Ng and Bereiter [1991] identified three types of learning goals in students undergoing a course in BASIC programming: task-completion goals, instructional goals, and knowledge-building goals. Students oriented towards knowledge-building goals (similar to what I am calling *questions*) were distinctive in many ways; they "actively constructed learning agendas for themselves, used prior knowledge to make sense of what they were learning, and used their new learning in turn to reconsider their prior knowledge." Scardamalia and Bereiter's [1991] educational environment, CSILE, is based on an analysis of the types of questions asked by students, which is similar to the analysis of questions in this paper. These empirical results are consistent with my theory of questions and question-driven understanding, and provide further support for my claim that better question asking leads to better learning.

Although, in this article, I use story understanding as the reasoning task, a similar argument can be made for other reasoning tasks as well. Understanding in general is also subjective and adaptive, as are other cognitive processes such as problem solving and design. A similar functional analysis could be performed for these tasks as well, which yields taxonomies of questions similar to that presented in this article. These question taxonomies overlap with the one in this article to the extent that the same subtasks (e.g., explanation) are used in different reasoning situations (e.g., diagnosis). Similar questions or knowledge goals can then arise in service of different reasoning tasks, and similar mechanisms can then be used for question

[1] Because computers will eventually be expected to interact with the physical world (e.g., robots) and the social world (e.g., employees), they will also be expected to have some of the physical or social goals that we currently attribute only to people.

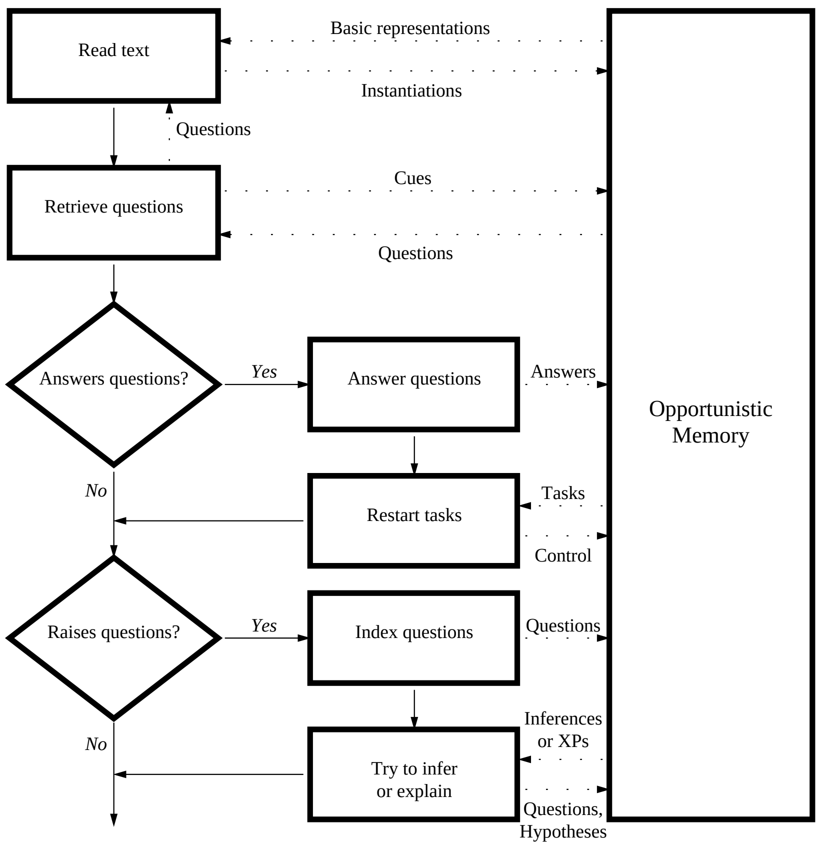

Figure 2: Control structure: The understanding cycle. A fact is interesting if it satisfies a knowledge goal currently pending in memory, or if it gives rise to new knowledge goals. Uninteresting facts pass vertically down with minimal processing; interesting facts cause suspended understanding tasks to be restarted or new tasks to be created. New tasks can give rise to new knowledge goals, which are suspended along with the tasks if answers are not yet known and cannot be inferred. This process implements the question-driven understanding model of figure 1.

management and learning. For example, Birnbaum [1986] arrived at similar conclusions about the indexing and opportunistic retrieval of pending goals during real-time planning and plan execution, based on an analysis of Zeigarnik's [1927] experiments demonstrating enhanced memorability of unsatisfied goals during the performance of various tasks.

1.3 Computational motivations: What are the knowledge goals of an understanding program?

Learning, then, can be viewed as the pursuit of one's interests or questions. However, it defeats our theoretical purpose to build a "question-asking" or "interest-pursuing" program per se. Instead, questions and interests should arise naturally as cognitive goals of the program during various stages of the reasoning process. This means that the program should ask a question only when it needs to acquire that piece of knowledge. For example, in the case of a story understanding system, a knowledge goal should be formulated only when the system requires the answer for the purposes of understanding the story. In other words, knowledge goals should be functionally useful to the overall goals of the system.

Similarly, if a student is to ask a question or formulate a knowledge goal, the teacher must provide a context in which the student requires the answer for a particular purpose. Again, this requires an analysis of the functional role of questions in the context of a learning task involving understanding, problem solving or design. Although problem solving or design contexts may be more appropriate to some teaching goals (e.g., scientific discovery), the task of understanding novel and unusual situations can also provide interesting contexts (e.g., teaching geography in the context of a mystery story). Understanding and problem solving tasks may also be combined (e.g., teaching economics in the context of a market situation [Shute *et al.*, 1988]). In this paper, I focus on an analysis of understanding tasks from a point of view of question generation. A similar approach could be used for the analysis for problem solving or design tasks as well.

My theory of questions is based on a theory of *understanding tasks*, the basic tasks of an understander. In addition to parser-level tasks such as noun group connection, pronoun reference, and so on, these tasks include higher-level tasks such as the integration of facts with what the understander already knows, the detection of anomalies in the text that identify flaws or gaps in the understander's model of the domain, the formulation of explanations to resolve those anomalies, the confirmation and refutation of potential explanations, and the learning of new explanations for use in understanding future situations. These are the basic tasks that an understander needs to be able to perform.

In order to carry out these tasks, the understander needs to integrate the text, which is often ambiguous, elliptic and vague, with its world knowledge, which is often incomplete. In formulating an explanation, for example, the understander may need to know more about the situation than is explicitly stated. However, it is impossible to anticipate when a particular piece of knowledge will be available to the understander because the real world (in the case of a story understanding program, the story) will not always provide exactly that piece of knowledge at exactly the time that the understander requires it. Thus, the understander must be able to suspend questions in memory and to reactivate them just when the information needed becomes available. In other words, the understander must be able to remember what knowledge is needed and why.

Furthermore, the system's understanding of any real world domain can never be quite complete. Conventional script-, frame- or schema-based theories assume that understanding means finding an appropriate script, frame or schema in memory and fitting it to the story. Schemas in memory are assumed to be "correct" in the sense that they are completely understood and constitute a correct model of the domain. If an applicable schema is found, an instance of the schema is created and applied to the story. The story is then assumed to be "understood." However, this model is inadequate because an understander's memory is always incomplete. Knowledge structures often have gaps in them, especially in poorly understood domains. These gaps correspond to what the understander has not yet understood about the domain. Even if a schema appears to be correct, novel experiences or stories may reveal flaws in the schema or a mismatch with the real world. Furthermore, the schema may not be indexed correctly in memory.

Understanding tasks, therefore, generate information subgoals or *questions*, which represent what the understander needs to know to perform the current task, be it explanation, learning, or any other cognitive task. These questions constitute the specific knowledge goals of the understander. Learning is a process of seeking answers to these questions in the input, which in turn raises new questions while answering old ones.

For example, in order to understand the blackmail story S-1, the system must understand the motivations of the would-be suicide bomber (an explanation task). In other words, it must formulate the question "Why would the boy have done the suicide bombing?" The desired explanation for the suicide bomber's actions constitutes an answer to this question. The explanation task gives rise to further questions and subquestions. Ultimately, the system finds an answer to a question in the input story, which enables it to complete its explanation task. Thus understanding can be viewed as a process of question transformation, as shown in figure 3.

If novel answers to questions are found, or if standard hypotheses are refuted, the system can learn a new fact or can modify an old belief. When a new piece of information comes in, the understander determines

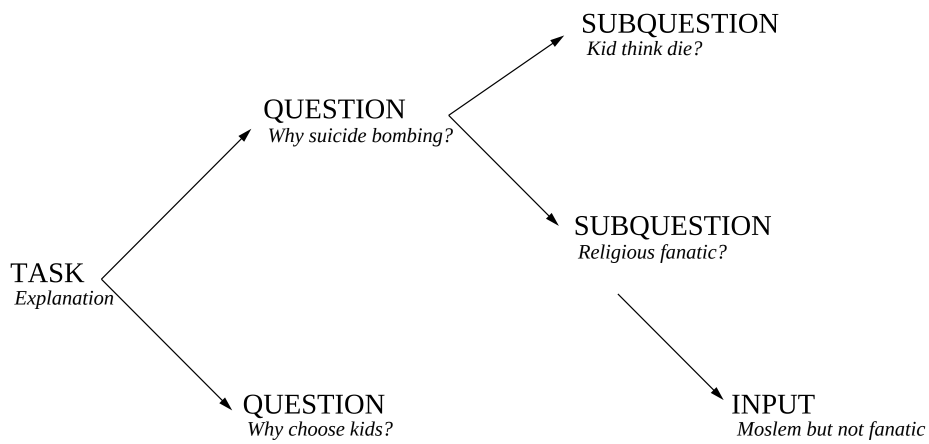

Figure 3: Inferences supporting the religious fanatic hypothesis, illustrating the idea of understanding as a process of question transformation.

whether a question pending in memory can be answered. If a question is answered, the source of the question determines the further processing that needs to be done on the new piece of information. In other words, the task that gave rise to the question now receives the new information, which is then processed by that task.

For example, consider the following sentence:

S-2: The boy was 16 years old.

What inferences should be drawn from this sentence? Suppose that this is used to answer the question, "Was the boy a teenager?" This now tells the understander that it was the fact that he was a teenager, as opposed to being male or 16 years old, that is important. Furthermore, the understander also knows why this is important because it knows why the question was asked in the first place. For example, if the question arose from a memory-level similarity-based generalization task, the understander might now try to construct a new generalization in memory, say, that suicide bombers in Lebanon are often teenagers. Alternatively, if the understander was trying to confirm the hypothesis that the boy was recruited because he was a gullible teenager, the understander would then try to confirm that explanation. This requires the understander to determine whether the boy was gullible. Because it is computationally intractable to draw all possible inferences from every fact, it is essential for the understander to be able to focus attention on inferences that are likely to be useful in learning.

This example illustrates how an understander can use its questions to focus the reasoning process. The computational advantage of this approach is it provides the system with a principled way of determining what it should learn and which inferences it should draw so that its attention can be focussed on the knowledge needed for a given task.

1.4 Question-driven learning

I argued that modelling goal-driven behavior is a reasonable task from the point of view of cognitive modelling and also that the ability to suspend and reactivate questions is essential to deal flexibly with real-world input.

This view has strong implications for theories of learning. I view learning as an incremental process of theory formation, involving both case-based reasoning and explanation-based learning processes. In general, any learning system would have incomplete knowledge of its domain because, by definition, it is still learning about that domain. These gaps give rise to difficulties during processing and drive the system to learn. The system learns in an incremental manner, by noticing interesting aspects of novel stories and generating questions, and by filling in the gaps in its memory which correspond to questions answered from previous experiences.

Learning can be broadly classified into two types: the acquisition of new knowledge and the reorganization of existing knowledge. Corresponding to these types of learning, we can distinguish between *knowledge acquisition goals* and *knowledge organization goals*, respectively. In the former case, the system learns by acquiring a new piece of knowledge (e.g., by acquiring a new knowledge structure or by acquiring new knowledge that fills a gap in an existing knowledge structure). In the latter case, the system learns by organizing or reorganizing what it already knows (e.g., by learning a new way to index an existing knowledge structure). Each type of knowledge goal corresponds to a class of gaps in the system's knowledge.

What is done with the newly learned information depends on the kind of "knowledge gap" the system is trying to fill. In an explanation-based understander, the new piece of knowledge may result in a new explanation in memory; it could be used to fill in a gap in an existing explanation; it could be used to elaborate an existing explanation if that explanation was not detailed enough to deal with the new situation; or it could be used to reorganize or reindex knowledge in memory to allow the reasoner to use what it already knows in novel situations to which that piece of knowledge had not been applied before [Ram, 1992]. Each type of learning leaves the system a little closer to a complete understanding of its domain. Each type of learning may also result in a new set of questions as the system realizes what else it needs to learn, which, in turn, drives the system toward further learning.

2 A computer model of question-driven understanding and learning

I have developed a computational model of this process in a natural language understanding domain. AQUA (Asking Questions and Understanding Answers) is a question-driven story understanding program that learns about terrorism by reading newspaper stories covering unusual terrorist incidents in the Middle East [Ram, 1987; Ram, 1989; Schank and Ram, 1988]. AQUA builds causal and motivational explanations for the events in the story in order to understand why the characters acted as they did or why certain events did or did not occur.

AQUA's basic goal in reading is to answer its questions and to improve its understanding of the domain (terrorism). AQUA's output consists of answers to old questions about the domain plus, of course, new questions. A specific example illustrating questions generated and answered during the reading of two stories is shown in figure 4. AQUA is driven by its questions or knowledge goals. It is a dynamic program, and it reads similar stories differently and forms different interpretations as its questions and interests evolve. AQUA would reread a story differently from the way it first read the story because the questions and explanations generated during the first reading affect the questions raises on the second reading.

My model of question-driven understanding and learning is summarized in table 1. The chief aspects of this process are the algorithms for the tasks of text interpretation, inference control, explanation, and learning. These algorithms are integrated using the generation and answering of questions as a unifying framework. Each task is sensitive to the current set of questions of the system and can raise new questions during its execution. Rather than being pre-programmed, questions are generated dynamically based on the requirements of the current task. Questions are indexed in an opportunistic memory, with suspended understanding tasks that are awaiting the answers. This means that the system knows not only what it

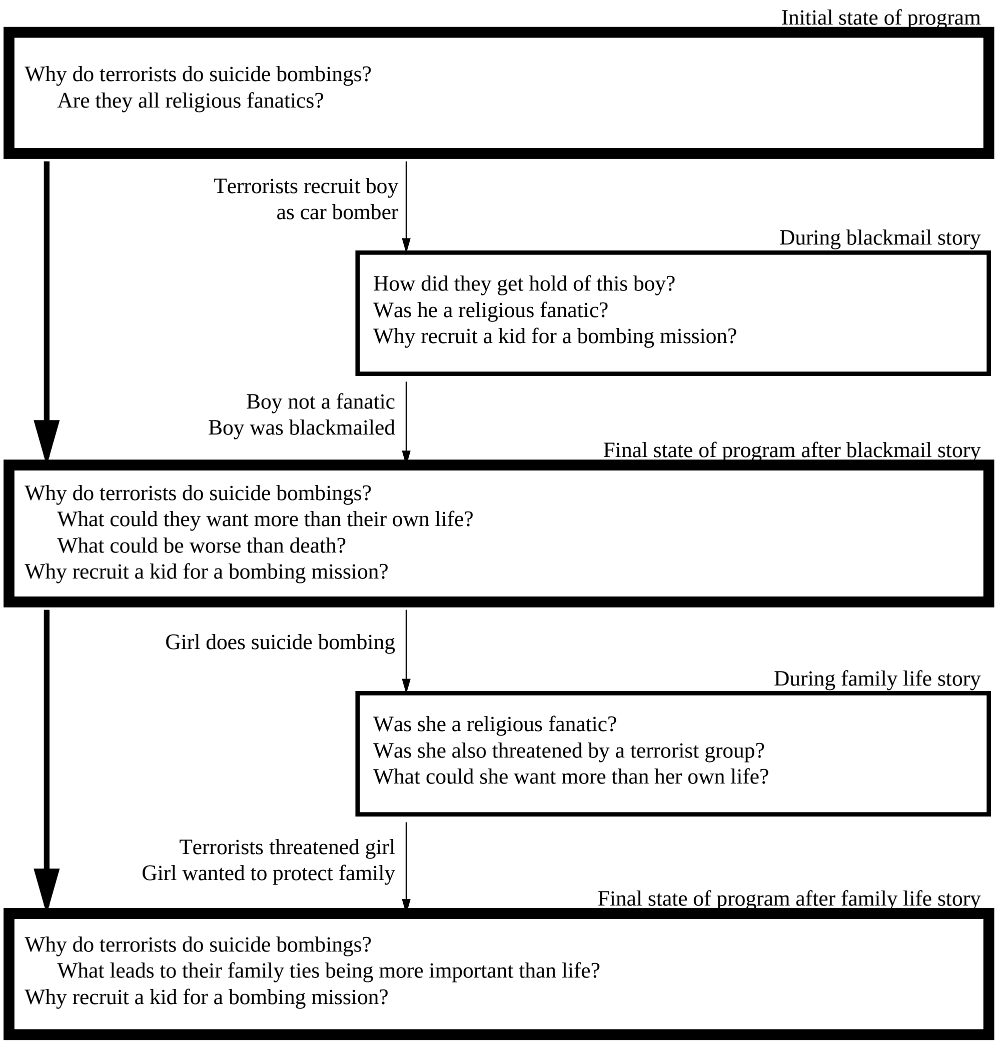

Figure 4: Questions in, questions out. Question-driven understanding is a process of asking questions and trying to answer them by reading a story. The understander starts out with a set of questions. At it reads, some of these questions are answered, and new questions are raised. After reading the story, the understander is left with a new set of questions that are the starting point for reading future stories. Here, the understander has read two stories, one about a boy being blackmailed into going on a suicide bombing mission in which no further details are given and another about a girl being “persuaded” to commit a suicidal terrorist attack by a terrorist group who threatened her family.

Read some text, focussing attention on interesting input as determined below. Build minimal representations in memory.

Determine interestingness of new input based on previous questions (the **retrieve questions** step below) and on new questions (the **ask questions** and **build explanations** steps).

Retrieve questions indexed in memory that might be relevant. Use these questions as an interestingness measure to focus the **read** step above.

Answer questions retrieved in the previous step, and restart the suspended process that was waiting for this piece of information. Learn from novel answers to questions.

Ask questions based on new input. Index new questions in memory, and use these questions as an interestingness measure to focus the **read** step above.

Build explanations to explain anomalous input identified by anomaly detection questions. Index new questions raised by this process in memory, and use these questions as an interestingness measure to focus the **read** step above.

Table 1: The processing cycle in AQUA. Learning occurs in the **answer questions** step, in which AQUA learns from novel answers to its questions, and in the **ask questions** step, in which AQUA "learns" new questions to ask in different situations.

wants to find out, but why it needs that piece of information. The architecture of the system is shown in figure 2.

Memory plays an important role in the construction of explanations for the purposes of understanding. As is discussed later in section 4, the understander relies on its memory of past experiences to provide potential explanations for new stories. However, the system's memory does not always contain "correct" cases or "correct" explanations, but rather one or more hypotheses about what the correct explanation might have been.[2] These hypotheses often have questions attached to them, representing what is still not understood or verified about those hypotheses. As the system reads new stories, it is reminded of past cases, and of old explanations. In attempting to apply these explanations to the new situation, it also remembers previously unanswered questions. The system's understanding of its cases gradually gets refined as these questions get answered [Ram, 1989; Ram, 1990b; Ram, 1992].

Much of real-world learning is an incremental process of this type. A reasoner learns by modifying what it already knows using pieces of new information that it acquires during its experiences. AQUA, the program described in this paper, implements the model of question generation and learning. The basic process in AQUA is one of question transformation (figure 4). The questions that AQUA asks, and the hypotheses that it formulates, change as it reads. As AQUA learns more about the domain, it asks better and more detailed questions, both "basic information questions" and "wonderment questions" [Scardamalia and Bereiter, 1991]. AQUA also gets "bored" if a story does not raise any new questions that it finds interesting by virtue of its interestingness criteria. Curiosity and interest are fundamental components of human learning; AQUA's question transformation process is an attempt to model this kind of learning.

Although this model was developed in the context of a story understanding task, I expect my theory of active question-driven learning to be applicable to a wide range of cognitive tasks. For example, Hunter presents a theory of learning based on "knowledge acquisition goals" for the diagnosis of lung tumors [Hunter,

[2] Actually, a single story or episode can provide more than one "case," each case being a particular interpretation or dealing with a particular aspect of the story. In AQUA, each anomaly in a story, along with the corresponding set of explanatory hypotheses, can be used as a case.

1989] and for scientific discovery in molecular biology [Hunter, 1990]. In addition to modelling reasoning tasks, my theory can also help in the design of educational environments. For example, Scardamalia and Bereiter [1991] examine concrete issues that have arisen from their attempts to implement a particular innovative educational environment. The environment they have created, CSILE, encourages students to collaborate with each other as they are learning, and in particular, encourages the asking of questions. Scardamalia and Bereiter examine whether students are capable of creating their own curricula by first asking questions, and then determining which set of questions will be productive to follow up on. The design of such environments depends on a theory of questions. In the AQUA model, many interesting questions arise from the detection of anomalies and the formulation of explanations to resolve these anomalies. This suggests, for example, that it may be useful to put students into situations in which they encounter anomalies or unexpected events, and to allow them to pose questions in the pursuit of explanations, and to explore different hypotheses as they try to make sense of the situation.

3 The nature of questions

I argued that questions play a central role in learning and presented both cognitive and functional arguments for a theory of question-based understanding and learning. In this section, I consider the nature of questions in more detail. Further details of the understanding and explanation aspects of AQUA are presented in the next section.

Much of the discussion here focusses on processing issues in question asking and learning. In addition, I argue that it is important for the understander to ask the "right" questions to make sure that it focuses on relevant issues and does not miss the point of the story. The depth of understanding that the understander achieves depends on the questions that it asks. Thus, it is also necessary to develop a content theory of questions and a taxonomy of the types of questions that an understander might ask. Of particular interest to the problem of understanding goal-based stories is the underlying model of motivational explanations, which will be discussed along with the questions associated with this model.

3.1 Sources of knowledge goals

Questions arise from gaps in a reasoner's domain knowledge. This could happen in three ways:

Novel situation: In a truly novel situation, an applicable case or schema may not be available. The reasoner simply does not have a prior experience that provides it with a case that is relevant to the current situation.

Mis-indexed knowledge: The reasoner may have a case or schema that is applicable to the current situation, but it may be unable to retrieve it since the case is not indexed under the cues that the situation provided.

Incorrect or incompletely understood knowledge: Previous experiences, especially in novel and complex domains, may not have been completely understood, and so cases or schemas corresponding to them may be incomplete or incorrect.

In all these circumstances, the understander has an opportunity to learn by further refining or altering the case or schema, by reindexing it in memory, or by answering the questions attached to it. I consider these situations in turn.

3.2 Knowledge structures are not perfect

Conventional script, frame or schema-based theories assume that understanding means finding an appropriate script, frame or schema[3] in memory and fitting it to the story (e.g., [Cullingford, 1978; DeJong, 1979]). Schemas in memory are assumed to be "correct" in the sense that they are completely understood and constitute a correct model of the domain. If an applicable schema is found, it is instantiated and the story is assumed to be "understood."

However, this model is inadequate since in practice an understander's memory will always be incomplete. Knowledge structures often have gaps in them, especially in poorly understood domains. These gaps correspond to what the understander has not yet understood about the domain. Even if a schema appears to be correct, novel experiences or stories may reveal flaws in the schema or a mismatch with the real world. Furthermore, the schema may not be indexed correctly in memory.

3.3 Misindexed knowledge structures

If a ready-made schema cannot be found that fits the story perfectly because the schema is not indexed correctly in memory, a conventional schema-based understander would simply fail. However, in a realistically sized memory, knowledge structures would not always be indexed with every conceivable and relevant scenario; also, the understander could not be expected to try all possible structures in every situation. For example, one would not expect the blackmail schema to be associated with suicide bombing a priori, nor would one want to draw on the blackmail hypothesis each time a suicide bombing is encountered.

3.4 Gaps in knowledge structures

Even if a relevant schema is found, the understander's knowledge of the scenario is often incomplete. Clearly, it is impossible to input all knowledge about every situation into our computers. Instead, we that expect many schemas would be incomplete, that many would have questions attached to them representing what the understander did not understand about those situations in the past, or that many would be simply very sketchy.

In the blackmail story S-1, for example, even if one assumes that blackmail was known to be a possible explanation for suicide bombing, it could hardly be a fully fleshed out causal chain. Instead, there would be many gaps in that chain, many questions attached to that explanation: How does one blackmail a child? What could anyone want more at the expense of his or her own life? In such a situation, there is more to understanding a story than fitting it into a schema; the understander has the opportunity to answer some of these pending questions if the story happens to provide these answers and, thus, improve its understanding of the scenario represented by the schema.

3.5 Incremental refinement of knowledge structures

Even if the schema appears to be perfectly self-contained and complete, new stories that differ in small but interesting ways can raise new questions about the scenario, questions that the understander could not have anticipated in advance. The religious fanatic explanation for suicide bombings in the Middle East, for example, is one that most people raise and apply automatically, but there are many stories about fanatics that are just a little different from the standard scenario. For example, consider applying the religious fanatic explanation to the following story (Los Angeles Times, April 10, 1985):

[3] We will use the neutral term *schema* for knowledge structures that encapsulate "canned" or stereotypical information about some situation, when for the purposes of the discussion at hand the distinctions between various kinds of schemas (scripts, frames, memory organization packets (MOPs), explanation patterns (XPs), etc.) or cases are irrelevant.

S-3: 4 Die in Suicide Attack Against Israelis.

JERUSALEM — A young guerrilla driving a car filled with explosives blew it up in a suicide attack Tuesday against a group of Israeli guards in Lebanon, killing three and wounding two others, an Israeli military source confirmed.

Besides the driver, identified by the guerrilla group as a 16-year-old girl, the blast killed two Israeli soldiers and a Druze member of the Israeli-backed South Lebanon Army. The wounded were Israeli soldiers, the source said. ...

A spokesman for the group identified the driver as Sana Mohaydaleh. In an interview taped before the incident and shown Tuesday after the attack on Beirut television, a girl identified as Mohaydaleh said: "I hope that my soul will unite with the souls of all the martyrs before me. I decided on martyrdom to free our land because I saw the misery of my countrymen under the occupation. I hope I will be successful and able to kill the highest number possible of our enemies."

Assuming that the understander has previous knowledge about religious fanatics, this story fits pretty well into the mold of a stereotypical scenario. Conventional understanders would be able to apply the schema to the story, but they would ignore the novel aspects of the story that are often the most interesting aspects. Instead, an understander should be able to focus on these aspects of the story and learn from them by modifying its schemas to accommodate these variations. This kind of incremental learning is an essential part of the understanding process.

While reading the story S-3, an understander might think about why these suicide bombers are all teenagers or about the strange fact that a television interview with this girl was videotaped before the mission. Is there something general to be learned by reading this story? Certainly, these issues are more interesting than other more stereotypical facts in the story, such as what kind of car was used or how many people were killed. These are the issues that the understander should focus on in this story.

3.6 Interestingness and the focus of attention problem

Combinatorial explosion of inferences has always been one of the classic problems in AI. Resources are limited, and inferences potentially infinite; a reasoner needs to be able to determine which inferences are useful to draw from a given piece of text. However, unless one considers the goals of the reasoner, it is very difficult to give a principled definition of what it means for an inference to be "useful."

The pragmatic need to focus attention provides a functional context in which to develop a theory of questions. Knowledge goals, the goals of a reasoner to acquire some piece of knowledge required for a reasoning task, can be used as the focussing criteria for inference control. Knowledge goals correspond to the interests of the reasoner. This approach, then, provides the basis for a theory of interestingness that is functionally motivated by consideration of the needs of the reasoner.

A program that uses knowledge goals to guide understanding is an improvement over one that processes everything in equal detail, that is, one that is completely text driven. An understander that is completely text driven processes everything in detail in the hope that it might turn out to be relevant. To avoid this, the understander should draw only those inferences that help it determine what it needs to know. In other words, the understander should use its knowledge goals to focus its attention on the interesting aspects of the story, where "interesting" can be defined as "relating to something the understander wants to find out about."[4] This idea is similar to that of "goal-guided inference" in social cognition [Zukier, 1986], to the

[4] Clearly, the author of a story also has control over what the understander finds interesting at any point. There is a tension between the conveyance of the story and the subjective bias of the reader in reading the story (Alterman, personal communication, March 15, 1991). The interaction between the two is complex. For example, there is some evidence that the subjective bias of the reader may cause the reader to perceive biases on the part of the author (e.g., the "hostile media phenomenon" [Vallone *et al.*, 1985]). This issue is beyond the scope of this article.

"goal satisfaction principle" of Hayes-Roth and Lesser [1976], which states that more processing should be given to knowledge sources whose responses are most likely to satisfy processing goals, and to the "relevance principle" of Sperber and Wilson [1986], which states that humans pay attention only to information that seems relevant to them. These principles make sense because cognitive processes are geared to achieving a large cognitive effect for a small effort. To achieve this, the understander must focus its attention on what seems to it to be the most relevant information available [Sperber and Wilson, 1986].

Why would an understander need to find something out in the first place? Ultimately, the point of reading is to learn more about the world. Questions arise when reading a story reveals gaps or inconsistencies in the world model. It is useful to focus attention on such questions because they arise from a "need to learn." For example, questions arising from anomalous facts are more useful than those arising from routine stereotypical facts because in the former case the understander may learn something new about the world.

There are two basic ways in which a fact can turn out to be worth processing, corresponding to the two diamonds in figure 2:

Top-down: A fact that answers a question is worth focussing on because it helps to achieve a knowledge goal of the understander, which in turn allows the understander to continue the reasoning task that was awaiting the answer.

Bottom-up: A fact that raises new questions is worth focussing on if the questions arise from a gap or inconsistency in the understander's knowledge base because the understander may be able to improve its knowledge base by learning something new about the world.

These can be viewed as heuristics for focussing attention. The decision to focus attention corresponds closely with the notion of "interestingness." When an understander focuses on a particular fact and processes it in greater detail, it can be said to be "interested" in that fact.[5] For this reason, focus of attention heuristics can also be thought of as interestingness heuristics. These heuristics provide a functional definition of "interestingness" as a criterion for focussing attention. Interestingness is a guess at what one thinks one might learn from paying attention to a fact or a question. The guess must be made without processing the fact or question in detail because otherwise the purpose of focussing attention to control inferences would be defeated. Thus, the interestingness heuristics used in AQUA are indeed heuristics rather than precise measures of the value of thinking about a fact or a question. Readers interested in this aspect of question-driven understanding are referred to Ram [1990c] and Ram and Hunter [1992] for more details.

3.7 Opportunism and the memory model

It would, of course, be impractical to maintain a list of "pending questions" and to check every question in that list every time a new fact was encountered. Thus, questions, with suspended understanding tasks that gave rise to them, must be indexed in memory. AQUA's memory model is based on the theory of dynamic memory [Kolodner, 1984; Lebowitz, 1983; Schank, 1982]. Rather than build an independent representation, which then has to be integrated into memory and related to previous episodes as a separate stage in the understanding process, an understander parses text directly into memory and relates it to its questions. In turn, this memory and, in particular the questions currently in memory are actively used to guide the parser as previously described.

Because questions are indexed in memory, it is quite likely that an understander will find the answers to questions other than the ones it is thinking about while reading. In other words, knowledge goals can be

[5] Because interestingness depends on one's goals, the heuristics implemented in AQUA do not cover interests arising from goals that lie outside the scope of the basic understanding and learning tasks that AQUA performs. For example, a parent would be interested in the report card of his or her child. Since AQUA's goals do not include caring for children, AQUA does not have any reason to be interested in a report card, unless the report card was anomalous with respect to AQUA's beliefs.

satisfied *opportunistically* during the course of understanding. Birnbaum [1986] made a similar argument for the opportunistic pursuit of goals during planning. Questions that are formed during understanding may be ones raised earlier in the same story or even during some other previously read story.

For example, one of the questions that is raised during many suicide bombing stories is: Why are children chosen for these missions? One doesn't actively and constantly think about this question, of course. However, answers to questions like these often are forthcoming even when the question is only "at the back of one's mind." Consider the following story (New Haven Register, August 6, 1988):

> **S-4: 200,000 children in world's armed forces, study finds.**
> GENEVA — An estimated 200,000 children are under arms worldwide, most of them forcibly recruited but some urged to enlist by their parents. ...
> According to the survey, examples of the problem reportedly included:
>
> - Illegal street roundups in Afghanistan to recruit youths under 15.
> - Abduction of boys under the legal draft age of 18 by army recruiters in El Salvador.
> - Lowering of Iran's conscription age to 13, with voluntary enlistment by parental consent for younger children.
> - Introduction by South Africa of compulsory military training at age 16. ...
> - Use of volunteers under age 15 in Honduras and Morocco. [p. 5]

This story provides an answer, in fact, several possible answers, to questions of why children are recruited for suicide bombing missions in Lebanon. If the understander is interested in and is reminded of this question when reading this story, this opportunity can be used to answer the question or to propose new hypotheses. Thus questions can be answered opportunistically, either within the same story in which they were raised or during future stories.

4 A question-based model of interpretation and explanation

I now turn to the implications of the model for the tasks of text interpretation and explanation in the context of story understanding and learning, and I discuss the implementation of the model in the AQUA program. Methodologically, I used AQUA as a testbed for exploring issues of interpretation, learning, explanation, and interestingness in an integrated framework. The domain of AQUA is that of newspaper stories about terrorism. In addition to the "factual" stories that many understanding programs deal with, AQUA reads "human interest" stories in the terrorism domain. It can achieve a better level of understanding of these stories than conventional script- and MOP-based understanders, since its questions need not only be generated from slots in stereotypical scripts or MOPs.[6] In fact, many interesting questions about a story arise from the detection of anomalies in the story and from the construction of explanations for those anomalies. These questions are used to guide understanding in the same way as questions from other sources.

Because questions represent the knowledge goals of the understander, they also provide the focus for learning. In addition to asking questions, therefore, AQUA can learn from answers to these questions. As implemented, AQUA improves its explanatory knowledge of its domain by incremental refinement of this knowledge using answers to questions that arise from the explanation process [Ram, 1990b; Ram, 1992]. In this section, we describe the underlying model of question-driven text interpretation and explanation that form the basis for the AQUA program.

[6] A MOP is a memory structure which contains information about how other memory structures are linked together in frequently occurring combinations. MOPs are composed of *scenes* which describe how and where a particular set of actions take place, which in turn can point to *scripts* that embody specific aspects of these scenes [Schank, 1982].

4.1 Question-driven interpretation of natural language text

In my theory, reading is viewed as a "plan" to learn more about the world. The understander's world model grows more complete as it reads a story and relates it to gaps and questions in memory. Ideally, only those inferences should be drawn that lead to conclusions required by the program. In practice, however, this is not always possible. Given that the basic task of a question-based understanding program is to answer questions in its memory by reading stories, there is an obvious choice to be made in the design of the program as characterized by the following extremes:

Text-driven program: A program that is totally text driven will read the text, build representations for it, and then process it to determine whether it addressed any questions in memory.

Question-driven program: A program that is totally question driven will select the most interesting or urgent question in memory and answer it via inference, reading text, or any other available method.

Each of these approaches has its disadvantages. The text- or data-driven method relies completely on bottom-up processes. It processes everything in detail in the hope that everything will be relevant. Instead, we want the program to concentrate on those aspects that are of immediate interest to it. In other words, we want the process to be driven by the interests or goals of the understander.

The goal- or question-driven method, in the purest sense, relies too heavily on top-down processing. It sees only what it is looking for already. Such a program is not be able to learn from unexpected input. Furthermore, another disadvantage of the purely question-driven method is that it expends resources in trying to answer a question immediately; therefore, it might be advantageous to design a program that indexes the question in memory, to be answered later when the opportunity arises. Last, so as not to overlook obvious information and to be sufficiently sensitive to the exigencies of the input, the process should incorporate some features of the data-driven method. The interaction between these requirements is non-trivial.

This dilemma has long been recognized in theories of text understanding, and various compromises between the two approaches have been proposed. For example, Kintsch's [1988] construction-integration approach involves bottom-up processing to construct a "text base" that is then integrated into a coherent whole. This approach is in contrast to top-down approaches in which language analysis proceeds in a top-down predictive manner, and bottom-up processing is invoked only when expectations fail (e.g., [Schank, 1978]). However, these theories do not take into account the learning goals of the understander. There is a static notion of what it means to "understand" a story, usually defined in terms of coherence relationships at some arbitrary level of detail. Typical story understanding systems are usually designed either to read stories in depth (e.g., BORIS [Dyer, 1982; Lehnert *et al.*, 1983]) or to skim stories (e.g., FRUMP [DeJong, 1979]). However, these systems cannot decide the depth to which the story should be processed or which inferences should be drawn during the understanding process because they do not maintain an explicit model of their learning goals. In other words, such systems are not trying to learn about anything in particular; they are merely reading.

At some level, of course, the task of the understanding process is to integrate the components of the text with the relevant components of the understander's memory into a coherent whole, given some suitable definition of coherence, for example, "event concept coherence" [Alterman, 1985]. However, in an active reader, this integration is directed by the goals, or questions, of the system. Questions, then, should be used to focus attention on those aspects of the story that would allow the understander to learn something of interest to its needs. This is the basis of AQUA's question-driven understanding algorithm.

AQUA is designed as a compromise between the bottom-up and top-down approaches. The basic understanding cycle that it uses is as follows. The parser reads the story word by word, trying to build a basic conceptual structure to represent the input. As quickly as possible, this structure is related to the questions in memory. If the new structure answers a question in memory, the inferences that were awaiting that answer are restarted. Thus, the program only draws those inferences that are required to match the new structure to the question and those that are demanded by the task that generated the question in the first place.

This method is called *variable-depth parsing*. The process is data driven to the extent that pieces of the input for which there are no explicit expectations but which are likely to be relevant are processed to the extent necessary to determine their relevance. In practice, this means that there is a set of bottom-up questions that the program always asks of incoming text. Further processing of the input is goal driven and is done only if these questions raise other questions that need answers or if the input turns out to answer some question already in memory. These correspond to the two diamonds in figure 2. These diamonds attempt to determine which facts the understander should focus on.[7]

4.2 Theory of explanation

The major emphasis in the development of AQUA was on the questions that arise during, and in support of, the explanation process. AQUA's theory of explanation is based on the claim that new explanations are built not by chaining inference rules together, but by reusing explanations that have been encountered in previous situations and are already known to the system. Previously encountered explanations are represented as stereotypical patterns of causality, known as *explanation patterns* (XPs; [Schank, 1986]). AQUA builds on Schank's theory of explanation patterns in three ways. First, a content theory of volitional explanations for motivational analysis is proposed. Second, a graph-based representation of the structure of explanation patterns is introduced. Third, the process of case-based explanation, while similar to that used by the SWALE program [Kass *et al.*, 1986], is formulated in a question-based framework. My emphasis is on the questions that underly the creation, verification, and learning of explanations, and not on the creative adaptation process described in Kass, Leake and Owens [1986]. The main points of my theory are summarized below; further details may be found in Ram [1989; 1990a].

4.2.1 Content theory of volitional explanations

The content of explanations is specific to the domain of interest. AQUA's task of understanding human interest stories depends on a theory of motivational analysis, the construction of volitional explanations to describe the planning behavior of agents. The representation of explanatory cases in this domain is based on the theory of decision models, which describe the planning process that an agent goes through when considering whether to perform an action. A decision model relates the actions in which the characters in the story are involved to the outcomes that those actions had for them; to the goals, beliefs, emotional states and social states of the characters as well as to the priorities or orderings among the goals; and to the decision process that the characters go through in considering their goals, goal-orderings and likely outcomes of the actions before deciding whether to perform those actions. A detailed volitional explanation involving the planning decisions of a character is called a *decision model* and is illustrated in figure 5. An example of a decision model representing the stereotypical religious fanatic explanation is shown in figure 6.

4.2.2 Structure of explanation patterns

AQUA has several XPs indexed in memory, representing its causal knowledge of the terrorism domain. These XPs are represented as graphs in figures 5 and 6. XPs have four main components:

[7] To improve on this even further, AQUA tries to determine which of its questions are interesting and worth pursuing. This decision is made using a set of interestingness heuristics. Details of these heuristics are beyond the scope of this article (see [Ram, 1990c]). For the purposes of this article, we may assume that AQUA pursues all its questions. Even this simplified heuristic is more efficient than pursuing all possible inferences.

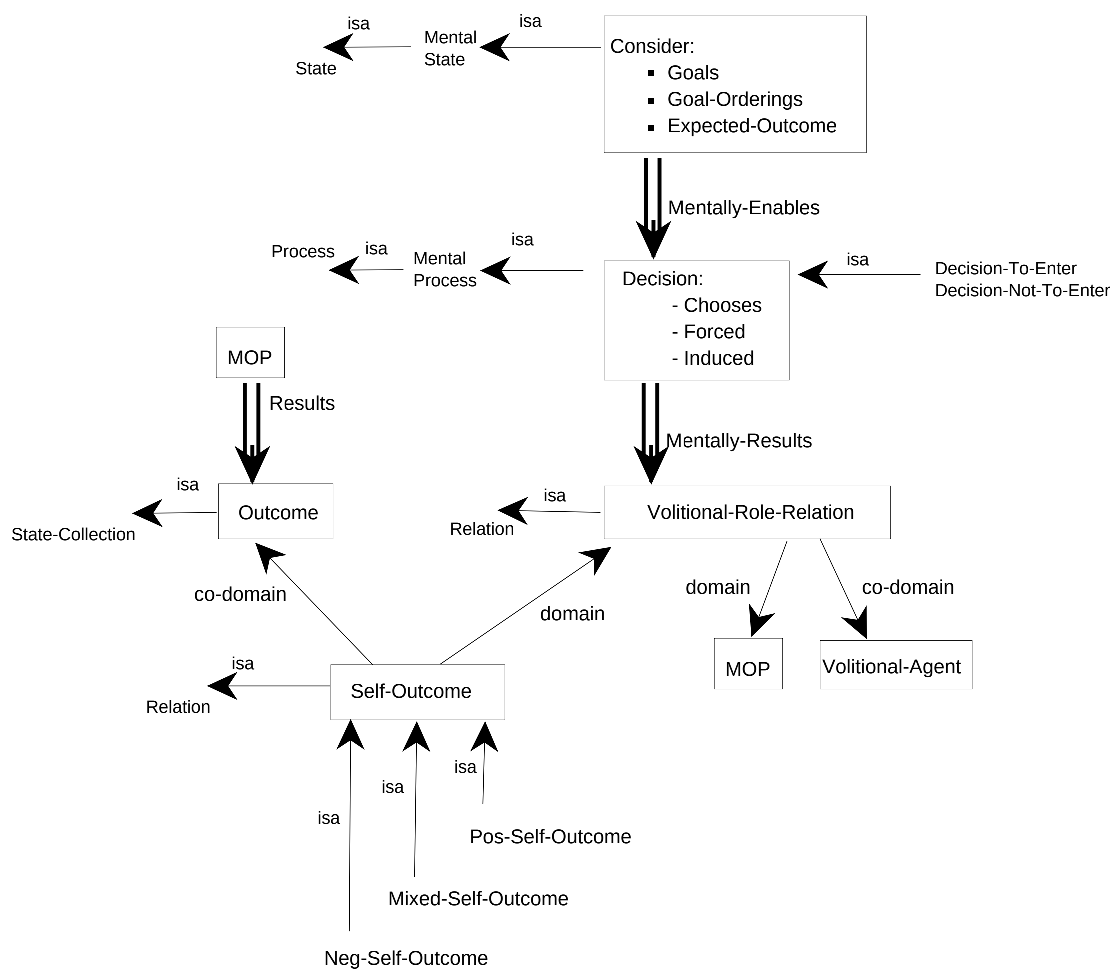

Figure 5: The structure of volitional explanations. A `volitional-agent` participates in some `volitional-role` in a `mop`, which then `results` in an `outcome` (a `collection` of `states`). Before this, the `volitional-agent` undergoes a `decision` process in which he `considers` his `goals`, `goal-orderings` and `expected-outcome`, which then `mentally-results` in the `volitional-role-relation` being considered becoming true (`in`) or false (`out`) depending on the outcome of the `decision`.

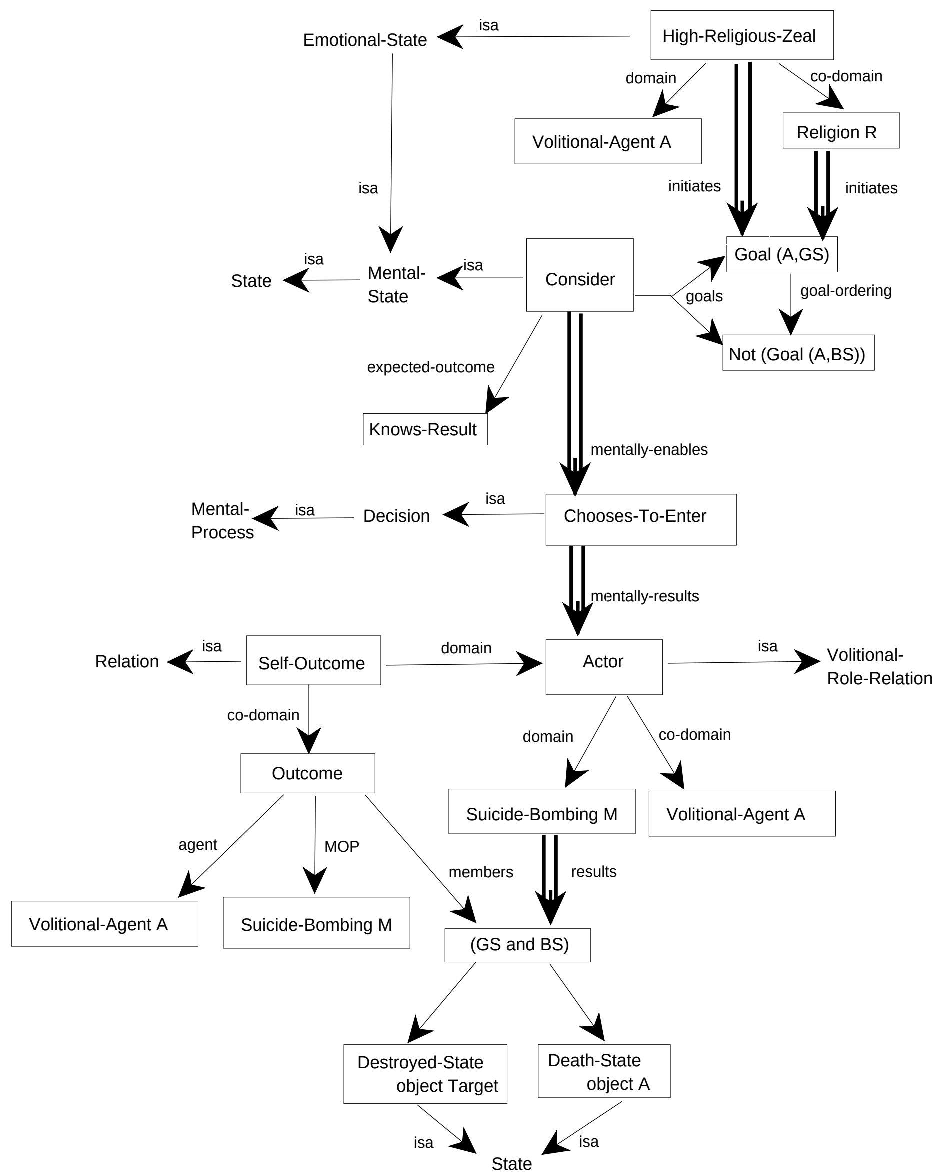

Figure 6: The religious fanatic explanation pattern. A is the agent, R his religion, M the action he chooses to do, and GS and BS the good and bad outcomes for A as a result of doing that action. A considers his goals to achieve GS and to prevent BS, the relative priorities of the two goals, and volitionally chooses to perform M knowing both expected outcomes of M, the `death-state` of A and the `destroyed-state` of the `target`.

- **PRE-XP-NODES:** Nodes that represent what is known before the XP is applied. One of these nodes, the EXPLAINS node, represents the particular action being explained.
- **XP-ASSERTED-NODES:** Nodes asserted by the XP as the explanation for the EXPLAINS node. These make up the premises of the explanation.
- **INTERNAL-XP-NODES:** Internal nodes asserted by the XP to link the XP-ASSERTED-NODES to the EXPLAINS node.
- **LINKS:** Causal links asserted by the XP. These, together with the INTERNAL-XP-NODES, are also called the internals of the XP.

An explanation pattern states that the XP-ASSERTED-NODES lead to the EXPLAINS node (which is part of a particular configuration of PRE-XP-NODES) via a set of INTERNAL-XP-NODES, the nodes being causally linked together via the LINKS (which, in turn, can invoke further XPs). In other words, an XP represents a causal chain composed of a set of nodes connected together using a set of LINKS (causal rules or XPs). The "antecedent" (or premise) of this causal chain is the set of XP-ASSERTED-NODES; the "internal nodes" of the causal chain are the INTERNAL-XP-NODES of the XP; and the "consequent" is the EXPLAINS node. The difference between XP-ASSERTED-NODES and INTERNAL-XP-NODES is that the former are merely asserted by the XP without further explanation, whereas the latter have causal antecedents within the XP itself.

4.2.3 Process model for explanation

An explanation-based understander must be able to detect anomalies in the input, and resolve them by building motivational and causal explanations for the events in the story in order to understand why the characters acted as they did or why certain events did or did not occur. This process characterizes both "story understanders" that try to achieve a deep understanding of the stories that they read and programs that need to understand their domains in service of other problem-solving tasks. Explanations are constructed by retrieving XPs from memory, applying them to the situation at hand, and verifying or evaluating the resulting hypotheses. This process consists of the following five steps:

Anomaly detection: Anomaly detection refers to the process of identifying an unusual fact that needs explanation. The anomalous fact may be unusual in the sense that it violates or contradicts some piece of information in memory. Alternatively, the fact may be unusual because, although there is no explicit contradiction, the reasoner fails to integrate the fact satisfactorily in its memory.

Explanation pattern retrieval: When faced with an anomalous situation, the reasoner tries to retrieve one or more explanation patterns that would explain the situation. Ideally, an XP should be indexed in memory so that it is retrieved only in those situations in which it is applicable. In practice, the indices to an XP represent the system's best guess at characterizing the situations to which the XP is likely to be applicable. In AQUA, indices to volitional XPs consist of stereotypical situations or contexts in which the XP might be encountered and of stereotypical categories of actors to whom the XPs might be applicable. These are called *situation indices* and *character stereotype indices*, respectively. A third type of index, the *anomaly category index*, represents the category of the XP required to explain a given type of anomaly.

Explanation pattern application: Once a set of potentially applicable XPs is retrieved, the reasoner tries to use them to resolve the anomaly. This involves instantiating the XPs, filling in the details through elaboration and specification, and checking the validity of the final explanations. The problem arises when the XP being applied has pending questions attached to it. If this occurs, these questions are instantiated and used to focus the understanding process. If the instantiated questions are answered by reading the story, answers to the questions are generalized and used to modify the original XP by answering the general questions attached to the XP.

Hypothesis verification and evaluation: The final step in the explanation process is the confirmation or refutation of possible explanations, or, if there is more than one hypothesis, discrimination between the alternatives. A hypothesis is a causal graph that connects the premises of the explanation to the conclusions via a set of intermediate assertions. At the end of this step, the reasoner is left with one or more alternative hypotheses and a set of new questions raised by these hypotheses.

4.2.4 Questions and explanation:

From the point of view of questions, the process model for the task of explanation can be formulated as follows:

Anomaly detection

- Ask anomaly detection questions based on the goals, goal-orderings, plans, beliefs, and decisions represented in decision models.

XP retrieval

- Ask XP retrieval questions based on the indices used by AQUA and attempt to match the current situation to the PRE-XP-NODES of an available XP.
- Retrieve specific XPs based on XP retrieval questions.
- Apply specific XPs or abstract XPs if no specific XPs are found.

XP application

- Ask XP applicability questions based on the INTERNAL-XP-NODES and LINKS of the XP. Suspend XP application if necessary.
- Instantiate nodes of the XP.
- Instantiate links of the XP.

Hypothesis confirmation

- Ask hypothesis verification questions (HVQs) based on the XP-ASSERTED-NODES of the XP.
- Suspend hypothesis verification if necessary.
- Confirm/refute hypothesis later when HVQs are answered.

At the end of this process, AQUA is left with one or more alternative hypotheses, each with its own set of HVQs. Partially confirmed hypotheses are maintained in a data dependency network called a *hypothesis tree*, along with questions (HVQs) representing what is required to verify these hypotheses (see figure 7).

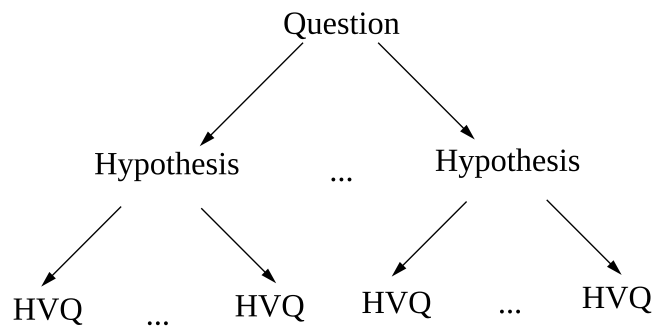

Figure 7: The structure of a hypothesis tree.

5 Types of questions

A functional theory of questions, as argued earlier, must be based on a taxonomy of types of knowledge goals that arise from the underlying understanding tasks that the questions serve. To develop a taxonomy of these knowledge goals, I asked several subjects to voice the questions that occurred to them as stories were read out to them. I then analyzed those questions and grouped them according to the understanding task (e.g., hypothesis verification) that they were relevant to. The groupings were revised based on a functional analysis of the knowledge required for the subtasks in my theory of story understanding and explanation, the subtasks, in turn, being mutually refined based on my analysis of the question data.

It is interesting to note that although my main taxonomic criteria were functional, the taxonomy fits my data well. Thus, I hypothesize that my theory, although intended as a computational model of an active reader, is also a plausible cognitive model. This is supported by the fact that my model is consistent with psychological data on question asking reported by Scardamalia and Bereiter [Scardamalia and Bereiter, 1991]. My goal-based approach is also consistent with psychological data on goal orientation in learning (e.g., [Ng and Bereiter, 1991]) and in focus of attention and inferencing (cf. review by Zukier [1986]).

I propose the following taxonomy of knowledge goals for story understanding:

Text goals: Knowledge goals of a text analysis program, arising from text-level tasks. These are the questions that arise from basic syntactic and semantic analysis that needs to be done on the input text, such as noun group attachment or pronoun reference. An example text goal is to find the referent of a pronoun.

Memory goals: Knowledge goals of a dynamic memory program, arising from memory-level tasks. A dynamic memory must be able to notice similarities, match incoming concepts to stereotypes in memory, form generalizations, and so on. An example memory goal might be to look for an event predicted by stored knowledge of a stereotyped action, such as asking what the ransom will be when one hears about a kidnapping.

Explanation goals: Goals of an explainer that arise from explanation-level tasks, including the detection and resolution of anomalies, and the building of motivational and causal explanations for the events in the story in order to understand why the characters acted as they did or why certain events did or did not occur. An example explanation goal might be to figure out the motivation of a suicide truck bomber mentioned in a story.

Text questions:

Was the car the target or the instrument of the bombing?
(Task: Interpret the noun phrase "car bombing".)

Memory questions:

Why are they all named Mohammed?
(Task: Notice similarities.)

Explanation questions:

Did the kid think he was going to die?
Why are kids chosen for these missions?
(Task: Explain anomalies.)

Relevance questions:

How does Mohammed Burro compare with a typical American teenager?
(Task: Determine personal relevance.)

Table 2: Examples of questions.

Relevance goals: Goals of any intelligent system in the real world, concerning the identification of aspects of the current situation that are "interesting" or relevant to its general goals. An example is looking for the name of an airline in a highjacking story if the understander were contemplating travelling by air soon.

A detailed taxonomy is presented in the Appendix. Table 2 presents examples of questions corresponding to these types of knowledge goals. Each question focuses on a different aspect of a story. For example, explanation questions focus on different types of anomalies, and on explanations for these anomalies. Asking an anomaly detection question is essential to detecting the corresponding anomaly. For example, asking the question "Does the actor want the outcome of this action?" is essential to the detection of a goal violation anomaly in the sense that the program will not notice the anomaly if it does not focus on the goals of the agent, that is, if it does not think of asking the question.

To put this another way, the questions asked by the understander influence its final understanding. Thus it is important for the understander to ask the "right" questions in order to achieve a detailed understanding of the situation. For the purpose of understanding stories involving motivations of people, I developed a taxonomy of motivational questions that focus on those motivational aspects of stories that are needed to build volitional explanations based on decision models. Figure 8 illustrates a series of anomaly detection questions based on questioning the goals, plans, beliefs, and decisions of an agent. Each question, if its answer seems anomalous, raises further questions. A more detailed taxonomy of explanation questions is presented in the Appendix.

In addition to their theoretical role in my model of inference control and interestingness, knowledge goals have also played an implementational role in my research by providing a uniform mechanism for the integration of various cognitive processes. For example, knowledge goals arising from, say, memory tasks are indexed in memory and used in the same way as knowledge goals arising from explanation tasks. A knowledge goal generated from one task may be suspended and satisfied opportunistically during the pursuit of some other task at a later stage or even during the processing of a different story. Implementational details may be found in Ram [1989].

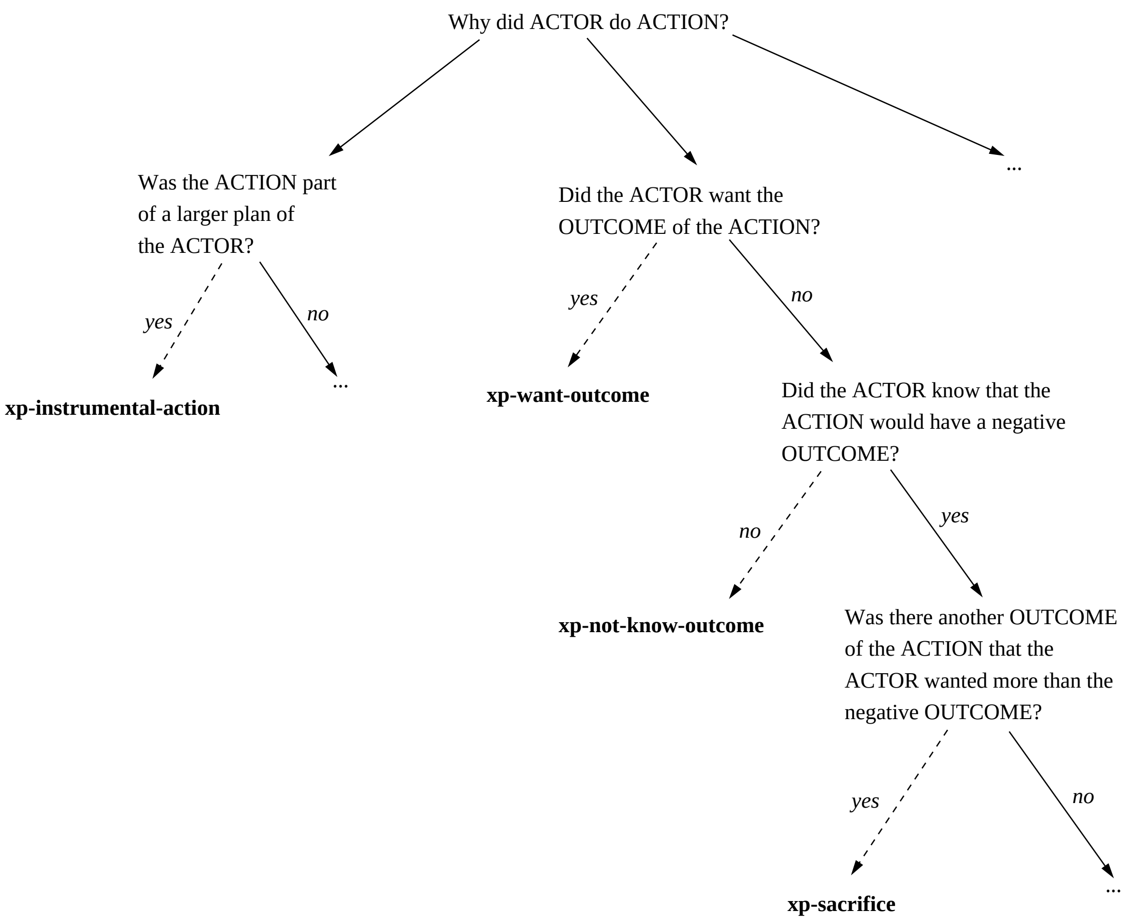

Figure 8: Anomaly detection questions arise from questioning different parts of the planning/decision model used to represent volitional explanations. These questions can be viewed as comprising a discrimination net of volitional questions, as shown in simplified form. This net corresponds to the hierarchy of abstract XPs represented in the program; the questions shown here are generated using the PRE-XP-NODES of abstract XPs in this hierarchy. If an anomaly is detected, the discrimination net also determines the category index of the XP required.

6 Representation of questions

When an understander is trying to reason about something (e.g., it is trying to explain something that seems anomalous, and it needs some piece of information that is not present in memory), it formulates a question that is indexed in memory where it expected to find the information. The question consists of two parts:

- **Concept specification:** The object of the question, that is, the desired information. This is represented using a memory structure that specifies what is minimally acceptable as an answer to the question. A new piece of knowledge is an answer to a question if it matches the specification completely. The answer can specify more than the question required, of course.
- **Task specification:** What to do with the information once it comes in, which depends on why the question was generated. This may be represented either as a procedure to be run or as a declarative specification of the suspended task. When the question is answered, either because the program actively pursued it or opportunistically answered it while processing something else, the suspended process that depends on that information is restarted.

In a sense, a question is similar to an open "slot" in a memory structure. AQUA's initial processing could be viewed as being similar to that of typical "script-based" understanders: Words in the input text are used to instantiate memory structures, and open slots in these memory structures are used as predictions for the rest of the story. However, there are three main differences (expressed here in a "slot-filling" terminology for comparison):

1. Typically, all open slots in newly instantiated structures are used as "requests" or "predictions" and cause the understander to look for fillers for those slots. AQUA, however, uses its interestingness heuristics to mark interesting slots to be used as predictions or questions. In addition, slots can be marked as being interesting by understanding tasks when their values are needed but not yet known.
2. Open slots arise not only from scriptlike knowledge structures, but also from causal or explanatory structures.
3. Typically, the ultimate task of the understander is to fill in as many of these open slots as possible. However, the action of filling in a slot does not do anything more than provide a value for that slot. In AQUA, however, slots are not filled for their own sake, but rather for the sake of performing some kind of reasoning with that value (e.g., confirming a hypothesis).

Thus AQUA subscribes to the basic slot-filling idea but extends this idea by selecting which slots are worth filling, by using different kinds of knowledge structures to provide slots, and by remembering why particular slots need to be filled so that it can use the filled values when they become known. The uniform representation of questions generated by different processes allows us to design an integrated system in an easy and natural manner.

AQUA's memory is built on top of an opportunistic memory architecture [Birnbaum, 1986; Birnbaum and Collins, 1984; Dehn, 1989; Hammond, 1988; Hayes-Roth and Hayes-Roth, 1979; Ram, 1989], which provides a uniform way to integrate question-based processing for different types of understanding tasks. The underlying memory model is based on the theory of dynamic memory [Schank, 1982], such as that used by IPP [Lebowitz, 1983] or CYRUS [Kolodner, 1984]. In addition to MOP-based episodic structures similar to these programs, AQUA's memory contains XPs indexed by situation, character stereotype, and anomaly category indices. The memory also contains questions indexed with the concepts in memory, each with its own reason for being asked. When these questions are answered, perhaps opportunistically, AQUA can restart the suspended computation. This requires the following supporting mechanisms:

- **Knowledge goal retrieval:** Finding suspended knowledge goals, even if not currently in "focus," that a new piece of knowledge might satisfy.
- **Knowledge goal indexing:** Storing knowledge goals in memory so that they are found almost only when they are relevant without having to look through long lists of pending questions.
- **Process scheduling:** Suspending tasks when there is insufficient information to execute them and restarting suspended tasks that depend on knowledge goals when the required information becomes available.
- **Hypothesis management:** Deleting alternative knowledge goals and hypotheses when a knowledge goal is satisfied because their likelihood of being useful decreases because an alternative has been found.

Implementation details of these mechanisms are discussed in Ram [1989]. These mechanisms allow AQUA to reason about what it knows, pose questions about what it needs to learn during the performance of an understanding task, and manage alternative hypotheses and explanations. These processes are fundamental to what Ng and Bereiter [1991] call the "knowledge-building goal orientation" in learning and are similar to the processes for self-management of learning goals used in their study.

7 Conclusions

Question generation and question answering are central to the processes of reasoning, understanding and learning. The point of reading is to answer one's own questions or fill in gaps in one's memory structures, that is, to learn. Based on this premise, I developed a model of questions and question asking in which questions are viewed as goals to learn. Questions are not just open "slots" in memory structures, but rather gaps that a reasoner needs to fill, or information that it needs to gather, for the purposes of carrying out various reasoning tasks.

My theory of questions has three components: a content theory describing the nature of questions and the actual questions that a reasoner might ask in a given situation; a computational process theory describing the processes of question generation, management of questions in memory, and interaction of questions with learning and understanding; and an implementational model in which I developed a program that embodies the content and process theories. The model is based on empirical results in psychology, social cognition, and education as well as research in artificial intelligence on case-based reasoning, explanation, machine learning, natural language understanding, and opportunistic planning.

The taxonomy of explanation questions presented in this article is appropriate for domains involving reasoning about goals, plans, and motivations of people. Similar taxonomies may be developed for other domains if explanatory theories are available for the domain. The taxonomy of text level questions is appropriate for natural language understanding tasks and is independent of the domain. The main claims of my process theory are summarized below:

- Questions are subgoals of understanding tasks and arise as knowledge goals when the reasoner needs a piece of knowledge to carry out some task that it does not already have.
- Questions are indexed in memory.
- Questions can be used to focus attention on interesting facts and limit the inferences drawn from them. In story understanding, questions determine what to look for in the story.
- Questions may be answered opportunistically. The explicit representation of questions facilitates the discovery of unexpected opportunities to answer the questions.

- Questions (and memory) are dynamic.
- Learning occurs incrementally through the answering of questions and the generation of new questions.

The ultimate focus of my research is on the relationship between questions and learning. I believe that learning is an active process involving the generation of questions and the pursuit of inferences in the context of these questions. Interestingness is viewed as a heuristic that allows the system to focus its attention on what it needs to learn. Learning, in turn, is viewed an incremental process of both question generation and question answering. I have demonstrated that this model is viable by implementing a computer program, AQUA, whose source of power (to borrow a term from Lenat [Lenat and Brown, 1984]) resides in its question-asking capabilities. The learning process in AQUA is incremental and dynamic, involving both goal-driven (top-down or active) and data-driven (bottom-up or opportunistic) processes. AQUA gradually evolves its understanding of a domain through experience with different situations. Its ability to ask questions is central to this process.

The issues in this article have implications not only for theories of learning, but for theories of education. While much research has focussed on question answering (e.g., what information to provide a student, how people learn, algorithms for machine learning), question asking is also important and involves issues such as when students ask questions, what kinds of questions they ask, what information to provide students relative to their pending questions, and so on. Research in learning and education needs to focus on the relationship between the learning goals of the student and the information provided by the environment. My theory models active-learning behavior in a dynamic environment and is a step towards understanding this relationship.

8 Acknowledgements

Research for this article was supported in part by the National Science Foundation under contract IRI-9009710. Part of the research described was conducted while the author was at Yale University, and supported by the Defense Advanced Research Projects Agency and the Office of Naval Research under contract N00014-85-K-0108 and by the Air Force Office of Scientific Research under contracts F49620-88-C-0058 and AFOSR-85-0343.

I thank Roger Schank for his guidance during this research, and Rick Alterman, Jeremiah Faries, and an anonymous reviewer for their detailed comments on an earlier draft of this article.

References

[Alterman, 1985] R. Alterman. A Dictionary Based on Concept Coherence. *Artificial Intelligence*, 25:153–186, 1985.

[Birnbaum and Collins, 1984] L. Birnbaum and G. Collins. Opportunistic Planning and Freudian Slips. In *Proceedings of the Sixth Annual Conference of the Cognitive Science Society*, pages 124–127, Boulder, CO, 1984. Institute of Cognitive Science and University of Colorado, Boulder.

[Birnbaum, 1986] L. Birnbaum. *Integrated Processing in Planning and Understanding.* Ph.D. thesis, Yale University, Department of Computer Science, New Haven, CT, December 1986. Research Report #489.

[Cullingford, 1978] R. Cullingford. *Script Application: Computer Understanding of Newspaper Stories.* Ph.D. thesis, Yale University, Department of Computer Science, New Haven, CT, 1978. Research Report #116.

[Dehn, 1989] N. Dehn. *Computer Story Writing: The Role of Reconstructive and Dynamic Memory.* Ph.D. thesis, Yale University, Department of Computer Science, New Haven, CT, 1989. Research Report #792.

[DeJong, 1979] G. F. DeJong. Prediction and Substantiation: A New Approach to Natural Language Understanding. *Cognitive Science*, 3:251–273, 1979.

[DeJong, 1983] G. F. DeJong. An Approach to Learning from Observation. In R. S. Michalski, editor, *Proceedings of the 1983 International Machine Learning Workshop*, pages 171–176, Monticello, IL, June 1983. Department of Computer Science, University of Illinois, Urbana-Champaign.

[Dyer, 1982] M. G. Dyer. *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension.* Ph.D. thesis, Yale University, Department of Computer Science, New Haven, CT, May 1982. Research Report #116.

[Hammond, 1988] K. J. Hammond. Opportunistic Memory: Storing and Recalling Suspended Goals. In J. L. Kolodner, editor, *Proceedings of a Workshop on Case-Based Reasoning*, pages 154–168, Clearwater Beach, FL, May 1988. Morgan Kaufmann, Inc., San Mateo, CA.

[Hayes-Roth and Hayes-Roth, 1979] B. Hayes-Roth and F. Hayes-Roth. A Cognitive Model of Planning. *Cognitive Science*, 2:275–310, 1979.

[Hayes-Roth and Lesser, 1976] F. Hayes-Roth and V. Lesser. Focus of attention in a distributed logic speech understanding system. In *Proceedings of the IEEE International Conference on Accoustics, Speech and Signal Processing*, pages 416–420, Philadephia, PA, April 1976. IEEE, New York, NY.

[Hunter, 1989] L. E. Hunter. *Knowledge Acquisition Planning: Gaining Expertise Through Experience.* Ph.D. thesis, Yale University, Department of Computer Science, New Haven, CT, January 1989. Research Report #678.

[Hunter, 1990] L. E. Hunter. Knowledge Acquisition Planning for Inference from Large Datasets. In B. D. Shriver, editor, *Proceedings of the Twenty Third Annual Hawaii International Conference on System Sciences*, pages 35–45, Kona, HI, 1990. IEEE Computer Society Press, Los Alamitos, CA.

[Kass *et al.*, 1986] A. Kass, D. Leake, and C. Owens. *SWALE: A Program That Explains*, pages 232–254. Lawrence Erlbaum Associates, Hillsdale, NJ, 1986.

[Kintsch, 1988] W. Kintsch. The Role of Knowledge in Discourse Comprehension: A Construction-Integration Model. *Psychological Review*, 95(2):163–182, 1988.

[Kolodner, 1984] J. L. Kolodner. *Retrieval and Organizational Strategies in Conceptual Memory: A Computer Model.* Lawrence Erlbaum Associates, Hillsdale, NJ, 1984.

[Lebowitz, 1983] M. Lebowitz. Generalization from Natural Language Text. *Cognitive Science*, 7(1):1–40, 1983.

[Lehnert *et al.*, 1983] W. G. Lehnert, M. G. Dyer, P. N. Johnson, C. Yang, and S. Harley. BORIS — An Experiment in In-Depth Understanding of Narratives. *Artificial Intelligence*, 20:15–62, 1983.

[Lehnert, 1978] W. G. Lehnert. *The Process of Question Answering.* Lawrence Erlbaum Associates, Hillsdale, NJ, 1978.

[Lenat and Brown, 1984] D. B. Lenat and J. S. Brown. Why AM and EURISKO appear to work. *Artificial Intelligence*, 23:269–294, 1984.

[Ng and Bereiter, 1991] E. Ng and C. Bereiter. Three Levels of Goal Orientation in Learning. *The Journal of the Learning Sciences*, 1(3), 1991. To appear.

[Ram and Hunter, 1992] A. Ram and L. Hunter. The Use of Explicit Goals for Knowledge to Guide Inference and Learning. *Applied Intelligence*, 1992. To appear. Also available as Technical Report GIT-CC-92/04, College of Computing, Georgia Institute of Technology, Atlanta, GA.

[Ram and Leake, 1991] A. Ram and D. Leake. Evaluation of Explanatory Hypotheses. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, Chicago, IL, August 1991.

[Ram, 1987] A. Ram. AQUA: Asking Questions and Understanding Answers. In *Proceedings of the Sixth Annual National Conference on Artificial Intelligence*, pages 312–316, Seattle, WA, July 1987. Morgan Kaufman Publishers, Inc., Los Altos, CA.

[Ram, 1989] A. Ram. *Question-driven understanding: An integrated theory of story understanding, memory and learning.* Ph.D. thesis, Yale University, Department of Computer Science, New Haven, CT, May 1989. Research Report #710.

[Ram, 1990a] A. Ram. Decision Models: A Theory of Volitional Explanation. In *Proceedings of the Twelvth Annual Conference of the Cognitive Science Society*, pages 198–205, Cambridge, MA, July 1990. Lawrence Erlbaum Associates.

[Ram, 1990b] A. Ram. Incremental Learning of Explanation Patterns and their Indices. In B. W. Porter and R. J. Mooney, editors, *Proceedings of the Seventh International Conference on Machine Learning*, pages 313–320, Austin, TX, June 1990. Morgan Kaufman Publishers, Inc.

[Ram, 1990c] A. Ram. Knowledge Goals: A Theory of Interestingness. In *Proceedings of the Twelvth Annual Conference of the Cognitive Science Society*, pages 206–214, Cambridge, MA, July 1990. Lawrence Erlbaum Associates.

[Ram, 1992] A. Ram. Indexing, Elaboration and Refinement: Incremental Learning of Explanatory Cases. *Machine Learning*, 1992. To appear. Also available as Technical Report GIT-CC-92/03, College of Computing, Georgia Institute of Technology, Atlanta, GA.

[Scardamalia and Bereiter, 1991] M. Scardamalia and C. Bereiter. Higher Levels of Agency for Children in Knowledge Building: A Challenge for the Design of New Knowledge Media. *The Journal of the Learning Sciences*, 1(1):37–68, 1991.

[Schank and Ram, 1988] R. C. Schank and A. Ram. Question-driven Parsing: A New Approach to Natural Language Understanding. *Journal of Japanese Society for Artificial Intelligence*, 3(3):260–270, May 1988.

[Schank, 1978] R. C. Schank. Predictive Understanding. In R. Campbell and P. Smith, editors, *Recent Advances in the Psychology of Language — Formal and Experimental Approaches*, pages 91–101. Plenum Press, New York, 1978.

[Schank, 1982] R. C. Schank. *Dynamic Memory: A Theory of Learning in Computers and People.* Cambridge University Press, New York, NY, 1982.

[Schank, 1986] R. C. Schank. *Explanation Patterns: Understanding Mechanically and Creatively.* Lawrence Erlbaum Associates, Hillsdale, NJ, 1986.

[Shute *et al.*, 1988] V. Shute, R. Glaser, and K. Raghavan. Inference and Discovery in an Exploratory Laboratory. Technical Report 10, Learning Research and Development Center, University of Pittsburgh, Pittsburgh, PA, February 1988.

[Sperber and Wilson, 1986] D. Sperber and D. Wilson. *Relevance: Communication and Cognition.* Language and Thought Series. Harvard University Press, Cambridge, MA, 1986.

[Vallone *et al.*, 1985] R. P. Vallone, L. Ross, and M. R. Lepper. The Hostile Media Phenomenon: Biased Perception And Perceptions Of Media Bias In Coverage Of The Beirut Massacre. *Journal Of Personality And Social Psychology*, 49(3):577–585, 1985.

[Zeigernik, 1927] B. Zeigernik. Das Behalten Erledigter und Unerledigter Handlungen. *Psychologische Forschungen*, 9:1–85, 1927. Reported in Birnbaum [1986].

[Zukier, 1986] H. Zukier. The Paradigmatic and Narrative Modes in Goal-Guided Inference. In R. Sorrentino and E. Higgins, editors, *Handbook of Motivation and Cognition: Foundations of Social Behavior*, pages 465–502. Guilford Press, Guilford, CT, 1986.

A Appendix: A taxonomy of questions

A.1 Anomaly detection questions

Anomaly detection questions can be categorized as follows:

Decision questions: These questions focus on the decision that the actor took when he decided to do the action. Therefore, this is also a taxonomy of the planning decisions one would consider when deciding to do an action.

- Personal goals:

- Does the actor want the outcome of this action?
- Does the actor want to avoid a negative outcome of not doing this action?
- Does the actor want a positive outcome of this action more than he wants to avoid a negative outcome of doing the same action?
- Does the actor enjoy doing that action?
- Does the actor habitually do this action?

- Instrumentality:

- Is this action instrumental to another action that the actor wants to carry out?
- Is this action part of a larger plan that the actor is carrying out?

- Interpersonal goals:

- Does the actor want a positive outcome of this action for someone he likes?
- Does the actor want to avoid a negative outcome of this action for someone he likes?
- Does the actor want a negative outcome of this action for someone he dislikes?
- Does the actor want a positive outcome of this action for a group that he belongs to?
- Does the actor want to avoid a negative outcome of this action for a group that he belongs to?
- Does the actor feel gratification in doing good for others?

- Social control:

- Did someone with social control over the actor ask him to perform the action?
- Did someone with social control over the actor force him to perform that action?

- Knowledge and beliefs:

- Did the actor know the probable outcomes of the action?
- Did the actor believe that the action would have a positive outcome for him?
- Did the actor know about the possible negative outcome of the action?

Interference questions: These questions focus on possible interference from external sources.

- Did someone want to block the actor's goal?
- Did someone want to prevent this state of the world? Would this state of the world violate this person's goals?
- Did someone want the actor to be involved in this action?
- Did the actor accidentally get involved in this action?

Planner questions: Given an action that was planned and executed by different people or groups:

- Did the action result in a positive outcome for both the planner and the actor?
- Did the planner select the actor knowing that the action would result in a negative outcome for the actor?

Physical anomaly questions: These questions focus on the physical causality underlying the observed events. For example, if this is the first suicide bombing story one has read:

- How can a car be used as a bomb?

A.2 XP retrieval questions

The taxonomy at the level of XPs representing abstract decision models mirrors the taxonomy of anomaly detection questions. Since the particular XP retrieval questions at the level of stereotypical XPs depend on the stereotypes currently in memory, this category is illustrated using examples rather than a taxonomy.

Abstract XP retrieval questions:

Decision anomalies:

- Anomaly: **goal-violation**
  Situation: Actor does action that results in a negative outcome.
  Questions:
  - Did the actor actually want this outcome (i.e., did we misperceive his goals?)
  - Did the action result in another outcome that the actor wanted even at the expense of the negative outcome?
  - Was the actor forced into doing this action?
  - Did the actor know that the action would have this negative outcome?
    - Did the actor have enough information about the environment?
    - Did the actor project the effects of the action correctly?
- Anomaly: **goal-violation** or **unusual-goal-ordering**
  Situation: Actor does action that results in a positive outcome and a negative outcome.
  Questions:
  - Does the actor prefer to achieve the positive outcome even at the expense of the negative outcome?
  - Does the actor actually want to avoid the negative outcome?
  - Can the goal violated by the negative outcome be pursued later?

  XPs:
  - Goal priority elevation in particular contexts.
    - Goal violated by negative outcome can be pursued later.
    - Goal of positive outcome is temporarily urgent.
    - Short term goals preferred to longer term goals.
    - Personal goals preferred to group goals.
    - Difficult goals postponed.
    - New goals from wanting what others have.
  - Actor's goal priorities were misperceived.
    - Personal differences (individual, parental).
    - Group differences.

    * Cultural differences.

- Anomaly: **bad-plan-choice**
  Situation: Actor does action to achieve a goal even though another action looks better.
  Questions:
  - Did the second action have a negative side effect for the actor?
  - Did the actor know about the second action?
  - Was the actor capable of performing the second action?
  - Is the first action better in the long run?
    * Is the cumulative effect of the first action better?
    * Does the first action keep more options open?
  - Does the actor enjoy doing the first action?
- Anomaly: **failed-opportunity**
  Situation: Actor doesn't do action that would have resulted in a positive outcome for actor.
  Questions:
  - Did the actor actually want this outcome (i.e., did we misperceive his goals)?
  - Did the actor know that the action would result in the positive outcome?
  - Did the action also result in a negative outcome for the actor?
  - Was the actor capable of performing that action?
- Anomaly: **unmotivated-action**
  Situation: Actor does an action that doesn't satisfy any of his goals.
  Questions:
  - Did the actor think that the action would satisfy one of his goals?
  - Does the action actually satisfy a goal (i.e., did we misperceive the situation)?

Planner anomalies:

- Anomaly: **malicious-intent**
  Situation: Planner knowingly recruits actor for action that results in negative outcome for actor.
  Questions:
  - Was the planner's real intention to achieve a negative outcome for the actor?
  - Did the planner want some other outcome of the action, but also wanted to avoid the negative outcome from happening to himself?
  - Was the planner willing to sacrifice the actor's goal to achieve a goal of his own?
  - Did the planner want both the outcomes, i.e., was he killing two birds with one stone?
  - Was the planner in turn forced to make this decision?

Stereotypical XP retrieval questions:

- XP: Religious fanatic does suicide bombing.
  Questions:
  - Was the actor religious?
  - (If in the Middle East) Was the actor a Shiite Moslem?
- XP: Depressed teenager commits suicide.
  Questions:
  - Was the actor a stereotypical teenager?
  - Was the actor depressed?

A.3 XP application and hypothesis verification questions

XP application questions arise from the internal nodes and links of XPs, and hypothesis verification questions arise from XP-ASSERTED-NODES (abductive assumptions). These questions depend on the particular XPs in memory.

A.4 Text-level questions

Actor-action identification:

- Who was the actor of the action?

Attachment questions:

- Noun group connection:

- Is this a noun phrase?
- How do I connect the two nouns together?

- Adjectives, prepositional attachment: These result in syntactic predictions being made, which are represented as questions. For example, on reading "recruit boy as ...":

- For what purpose did the terrorists recruit the boy?

Reference questions: A referential description could be a pronoun (e.g., "he") or a definite reference (e.g., "the teenage bomber").

- Who or what does this description refer to?

Disambiguation and specialization questions: These questions try to disambiguate or refine vague descriptions. For example, on reading "get into":

- What is the destination of "get into"? (If the destination turns out to be a car, which is a type of container, "get into" can be specialized to "enter.")

A.5 Memory-level questions

Journalism questions: These questions seek the who, what, when and where of the story.

Reference questions: At the text level, reference questions are triggered by pronouns and definite noun phrases, but referential cues may be non-textual too. Some examples of reference questions are:

- Is this bombing the same bombing that was predicted?
- Is the 16-year-old teenager the same as the boy that was recruited?
- Is the explosive-laden car the same car that was implicitly part of the suicide bombing mentioned in the headline?

Attachment questions: Attachment questions try to determine the connection between two concepts.

Stereotype activation questions: These questions arise from the activation of stereotypes (stereotypical features of concepts, stereotypical activities and goals of people, or stereotypical motivational explanations for actions.)

- Does this instance fit this stereotype?

- Is this a typical role filler for this stereotype?
- Does this person have these stereotypical goals?
- Can I view this instance as one that matches this stereotype?
- Does this group of instances follow a pattern?

MOP and scene inference questions: These questions are based on the following tasks:

- **Predicting the scenes of a MOP.** For example, `recruit` predicts `locate` and `persuade`.
- **Finding the correct context for a scene.** For example, `blow-up-bomb` is part of `bombing`.
- **Activating definitional inferences.** For example, `suicide` by definition results in the `death-state` of the `actor`.
- **Specializing or refining vague descriptions.** For example, a `person` with `age` = `16` is a `teenager`; a `ptrans` with `instrument` = `car` is a `drive`.

For example, on reading about the recruiting of a boy for a suicide bombing mission, AQUA asks:

- How did they locate the boy?
- How did they persuade the boy to do the suicide bombing?

Reminding and memory search questions: A dynamic memory must be able to recognize unique combinations, notice similarities (both superficial and structural), and retrieve memory structures given a specification. Some examples of these questions are:

- Why are they all named Mohammed?
- Why do we hear about Lebanese car bombers and not about Israeli car bombers?
- Are there any other instances of teenage suicide bombers in memory?

Analogy and generalization questions: Remindings give rise to questions that seek generalizations. For example:

- Are car bombers motivated like the Kamikaze?

A.6 Relevance questions

All intelligent systems in the real world are, in some sense, constantly asking questions about personal relevance:

Why does this matter to me?

How does this relate to my goals?

Is this story about anyone I know or have heard of?

Is this story about a place that I know of or have been to?

Is this story about goals that I share?