Bellcore Workshop on High-Performance Information Filtering, Morristown, NJ, November 1991.

Interest-based information filtering and extraction in natural language understanding systems

Ashwin Ram

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

Abstract

Given the vast amount of information available to the average person, there is a growing need for mechanisms that can select relevant or useful information based on some specification of the interests of a user. Furthermore, experience with natural language understanding and reasoning programs in artificial intelligence has demonstrated that the combinatorial explosion of possible conclusions that can be drawn from any input is a serious computational bottleneck in the design of computer programs that process information automatically. This paper presents a theory of interestingness that serves as the basis for two story understanding programs, one that can filter and extract information likely to be relevant or interesting to a user, and another that can formulate and pursue its own interests based on an analysis of the information necessary to carry out the tasks it is pursuing. We discuss the basis for our theory of interestingness, heuristics for interest-based processing of information, and the process used to filter and extract relevant information from the input.

1 Motivation

With the advent of newspapers, television, computers, multimedia, and world-wide communication services, we are fast entering what might be called the "Information Age." There has been a radical change in the availability of information to the average human being. A large portion of our time is spent processing the large amount of information impending on us, and much time is wasted on processing information that is not interesting or useful. Although artificial intelligence programs may help automate some of this process, these programs too will spend a lot of time processing vast amounts of irrelevant information if they are not focussed in some manner.

Clearly, we need mechanisms to sort through the plethora of information in a fast, efficient manner and pull out just the facts, or draw just the conclusions, that are of interest to a user. In this paper, we present two natural language understanding programs that read newspaper stories. Unlike conventional story understanding programs, these systems have been designed to focus on the interestingness of the input in the context of the goals and interests of the user. One of the systems uses a model of the user to filter and extract relevant information from the input. In the other system, the "user" is the system itself; it determines its own interests based on the overall tasks it is pursuing.

Our approach is based on subjective, interest-based analysis of the input text, and goes beyond simple keyword search which involves no actual understanding of the stories. This allows the system to determine interestingness based on an analysis of the relationship of the content of the input to the user's interests, rather than on the occurrence of a given set of keywords as in conventional database query systems.

Our approach also has advantages over conventional story understanding systems in artificial intelligence. The first advantage is computational: Since it is computationally intractable to pay attention to all possible aspects and consequences of all input facts, an interest-based story understanding system is more efficient since it only processes those facts that are likely to be relevant or useful, and only to the extent that they are likely to be relevant or useful. The second advantage is cognitive: Since people do not process all the information in a story in equal depth or detail, a system that processes information based on its interests and goals is closer to being a model of human cognition.

2 Theoretical basis

The primary element in our theory is an analysis of the *knowledge goals* that underly the understanding process. People read newspaper stories for a reason: to learn more about what they are interested in. Typical computer programs, on the other hand, don't. In fact, computers don't even have interests; there is nothing in particular that they are trying to find out when they read. If a computer program is to be a model of story understanding, it should also read for a "purpose."

Such a purpose might be the goal to learn more about a given domain, to improve its model of the world, or to collect some information for some other purpose. In Ram [1990], we presented a model of understanding as the pursuit of goals to find and infer information from the input (in the case of a story understanding program, a story). The goal of a reasoner to acquire some piece of knowledge is called a *knowledge goal*. Such goals arise from the basic tasks that the reasoner is trying to perform. In formulating an explanation, for example, the reasoner may need to know more about the situation than is explicitly stated before it can decide which is the best explanation. However, the knowledge base of the reasoner is never quite complete: knowledge structures may have "gaps" in them, or they may not be indexed correctly in memory. Furthermore, it is impossible to anticipate when a particular piece of knowledge will be available to the understander, since the real world 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 reactivate them at the right time when the information it needs becomes available.

In other words, the understander must be able to remember what it needs to know, and why. This "need to know" forms the basis of our theory of knowledge goals [Ram, 1989; Ram, 1990]. In addition to representing its goals for acquiring knowledge, the understander must be able to use these goals to guide and focus the process of information gathering. For example, in a story understanding situation, the understander must be able to use its knowledge goals and interests to extract that information that is likely to be relevant and useful. We call this process goal-based or interest-based understanding.

The same process can be used in an information filtering system which extracts relevant information based on the interests and needs of an outside user. In order to go beyond keyword search, even if an initial set of questions or interests is input by the user, the system would still need to generate a new set of knowledge goals to pursue based on an analysis of the user's questions and the information provided by the input. Our model treats knowledge goals from all sources uniformly, and thus provides a nice framework for integrating different methods of generating and pursuing knowledge goals.

3 Two case studies in interest-based understanding

To summarize briefly, a knowledge goal can be viewed as a question about the world, or a desire to find information. These questions determine the intellectual interests of the understander, since they represent what the understander is interested in finding out. In a database system, such a question may be directly input by the user as a query to the system. In a reasoning system, a question may arise from the need to know some fact for the purposes of a reasoning task. In either case, knowledge goals can be used as a theoretical basis for an interest-based story understanding system.

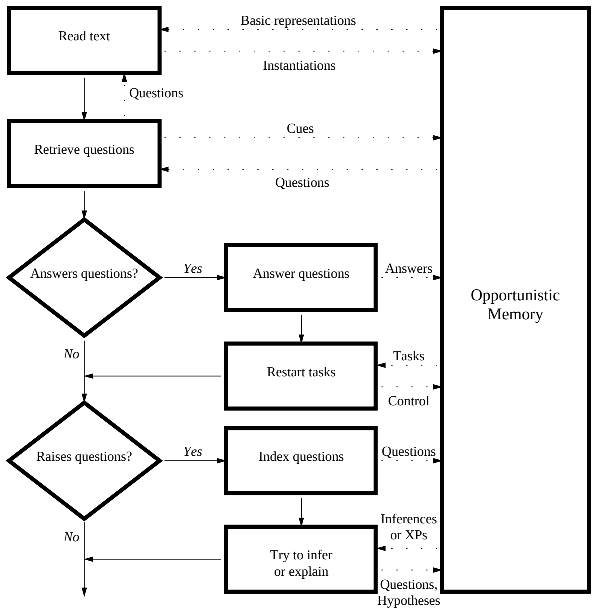

Figure 1: Control structure: The understanding cycle in AQUA. A fact is interesting if it satisfies a knowledge goal pending in memory, or if it gives rise to new knowledge goals. Uninteresting facts pass vertically down with minimal processing; interesting facts cause suspended understanding tasks to 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.

An example of such a system is the AQUA (Asking Questions and Understanding Answers) program [Ram, 1987; Ram, 1989]. AQUA is a dynamic story understanding program that is driven by its questions or goals to acquire knowledge. Rather than being "canned," the program is always changing as its questions change; it reads similar stories differently and forms different interpretations as its questions and interests evolve. AQUA learns about suicide bombing by reading newspaper stories about unusual terrorist incidents in the Middle East, and gradually refines both its model of terrorism as well as its questions about different aspects of terrorism. Thus AQUA can formulate its own set of interests and use them to extract the desired information from newspaper stories.

The control structure used in AQUA is shown in figure 1. Although AQUA extracts information from newspaper stories, this approach can be extended to a more traditional database query task. The system would understand textual information using a set of questions stored in memory, rather than by using syntactic parsing rules or keyword search. It would read text from the point of view of this set of questions, retrieve information that were relevant to these questions, and be able to learn more about the domain by finding answers to these questions in the understood text as well as by generating possible questions to pursue further. The initial set of questions may be input by human users, may be persistent goals left over from previous queries, or may even be new questions generated by the system itself. Our goal-based approach to information retrieval is discussed in Ram & Hunter [1991].

Our second program, written by Philip Gordon, is a "personalized newspaper" program called PIES (Personal Interest Engine for Stories), which can read and summarize stories from the point of view of the interests of the user. PIES is designed to scan several stories from different sources in an efficient manner, and produce a condensed, personalized newspaper for the user. This system is still under development. As currently implemented, PIES can understand simple stories about sports events that have been hand translated into conceptual representations, and select those stories that are likely to be of interest to the user. As in the AQUA system, PIES does not rely on keyword search to tag "interesting" concepts, but rather tries to understand the input subjectively as the user might, based on interest and relevance of various aspects of the input. The control structure used in the PIES system is shown in figure 2. The process of using interests (or questions) to focus and guide the story understanding process is similar to that used in AQUA (figure 1). However, in PIES we have emphasized the information filtering aspects of the theory rather than

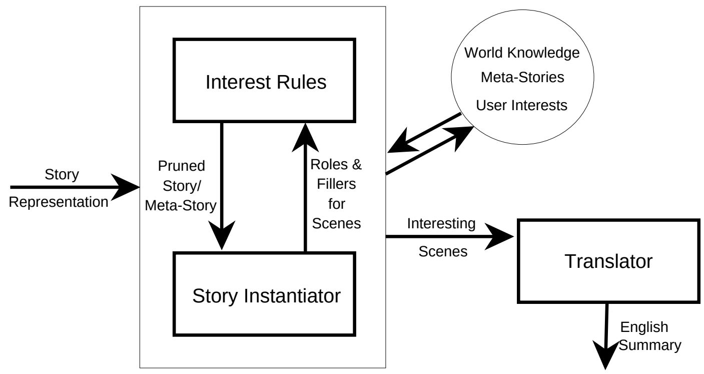

Figure 2: Control structure: The understanding cycle in PIES. The user's interests are represented as a set of interest rules, and used to prune the knowledge structures that will be applied to the input story. The story instantiator understands the story using these knowledge structures, and the representation of the understood story is further pruned so as to extract only the interesting scenes. These are then summarized in simple English.

---

the question asking and answering aspects which were the theoretical focus in AQUA.

4 Theory of interestingness

4.1 Knowledge goals

Our theory of interestingness is based on the functional role of questions, or knowledge goals, in the process of understanding [Ram, 1990; Ram, 1991]. We characterize knowledge goals according to the type of understanding task that they arise from:

Text goals: Knowledge goals of a text analysis program, arising from basic syntactic and semantic analysis that needs to be done on the input text. 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 such as noticing similarities, matching incoming concepts to stereotypes in memory, and forming generalizations. An example memory goal might be to look for an event predicted by stored knowledge of a stereotyped action, such as wondering about what the ransom will be when one reads about a kidnapping.

Explanation goals: Knowledge 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 occurred or did not occur. An example explanation goal might be to figure out the motivation of a suicide truck bomber mentioned in a story.

Relevance goals: Knowledge 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 here might involve looking for the name of an airline in a hijacking story if the understander were contemplating travelling by air soon. In the sports domain, an example might be a user's interest in Bret Saberhagen, a baseball player for the Atlanta Braves, because he was a friend of the user.

4.2 Interestingness heuristics

In AQUA, knowledge goals are used to determine which input facts the understander should focus on. To improve on this even further, the understander needs a way of determining which knowledge goals worth pursuing and which ones are not. Not all questions are equally important, nor are all answers equally valuable. The understander needs to be able to determine the priorities of its knowledge goals, depending on how likely the understander is to learn something by thinking about these knowledge goals. For example, an unexpected situation is more interesting than a stereotypical one, because it raises new questions for the understander to think about.

Similarly, in PIES, although the user model represents the current set of interests of the user, the system needs to determine which aspects of the input are likely to address these interests. For example, if the user lives in Atlanta and is interested in sports, the system needs to be able to determine whether the user is more likely to be interested in the Atlanta Braves or the Chicago Cubs.

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. For this reason, heuristics for focussing 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 interest-based understanding is indeed a *heuristic* process rather than a way of deriving precise measures of the value of thinking about a fact or a question.

We use a set of interestingness heuristics, described in Ram [1990], to focus attention. An example of an interestingness heuristic is the goal relevance heuristic:

> A fact that could be instrumental to or could hinder a goal of the understander is more interesting than one that has no relevance to the understander's goals. A fact that directly matches or conflicts with a goal of the understander is very interesting.

Some heuristics are specific to particular domains. For example, here is one of the interest rules used in PIES:

> If x is a sports team, then x is more interesting if it is a world champion team.

Both AQUA and PIES combine the recommendations of applicable heuristics to determine the overall interestingness of the input.

4.3 Knowledge goals as the basis for interest-based understanding

The basic process of interest-based understanding involves the generation of knowledge goals seeking information required by various understanding tasks, the transformation of these knowledge goals into subgoals, and the matching of pending knowledge goals to information in the story. One might think of this as a process of question transformation, in which the reasoner generates questions which then trigger, or get transformed into, other questions. Thus one major aspect of understanding is the ability to "plan" to look for, retrieve, or infer information that might be relevant [Ram and Hunter, 1991].

Some knowledge goals arise from an intellectual need to find information, such as the need to determine the religion of a suicide bomber for the purposes of determining the applicability of the religious fanatic explanation in a given story about terrorism. Other knowledge goals arise from personal relevance, such as the goal to determine Bret Saberhagen's score in a recent baseball game because one knows Saberhagen personally. We argue that knowledge goals, whether they arise from intellectual needs or personal relevance, form the basis for defining "interest" for interest-based understanding.

Goals can also arise from a combination of two or more heuristics. For example, a user may be interested in baseball in general and read stories about baseball to keep up with recent scores of various teams. A user who lives in Atlanta may be particularly interested in the performance of the Atlanta Braves (personal relevance to known location), and even more so in Saberhagen's scores (personal relevance to known person). Alternatively, interest in the scores for a Braves game may arise because one needed to know this score for a different purpose, say, for the purpose of convincing someone that Atlanta has a good baseball team, even if one did not personally care about the score. Interest-based understanding is the process of reasoning about, and building representations for, input information based on the system's interest in different aspects of the input that are likely to be relevant.

4.4 Interest-based information filtering and selection

The above notion of goal-based understanding is necessarily a subjective one. There is no correct understanding of a story, and no correct depth of processing for the story. Rather, understanding is viewed as a process of interest-based filtering and selection of input information likely to be relevant or useful. In addition, the determination of interestingness must be done without detailed analysis of all possible aspects of the input, which would quickly get computationally intractable as all possible consequences and inferences were considered. These aspects of the problem were explored in the AQUA program.

There is another aspect of the problem of interest-based information filtering and selection which has largely been ignored thus far. In addition to the identification of interesting aspects of the input, this process also requires the identification of interesting or relevant aspects of the system's own knowledge structures that ought to be brought to bear on the input. For example, in a script-based understanding system, the system may be able to prune its own scripts so as not to consider uninteresting scenes. This would allow it to skip over aspects of the story without even trying to fit it into those scenes of the script, thus saving time and memory resources. Thus even though a script for a terrorist attack has a scene for the method of attack (e.g., suicide car bombing), there is no point in trying to match and instantiate this scene if all one is trying to determine is the name of the terrorist group which organized the attack.

We use the term interest-based information filtering and selection to apply not only to interest-based understanding of the input, but also to interest-based pruning of the system's own knowledge structures. This idea has been explored in the PIES program. The basic process of understanding used in PIES is similar to conventional script-based or MOP-based understanding programs (e.g., [Cullingford, 1978; DeJong, 1979; Schank, 1982]). However, PIES understands its input from the point of view of the user. PIES uses its representation of the user's interests, not only to identify relevant aspects of the input, but also to prune its own knowledge structures before they are applied to the input. This pruning is done dynamically, based on the particular set of interests being pursued. The pruned knowledge structures are then applied to the input, and the resulting representation is further pruned using interestingness heuristics such as those discussed earlier.

An open issue in this research is the amount of pruning that should be done based on the specified interests. Clearly, the depth of processing exhibited by human readers, and the amount of irrelevant information they are willing to read, depends not only on their interests, but also on other factors such as how much time they have, how much detail they want to see, and so on. In our current implementation of PIES, we have lumped these factors together into a user-supplied "interest threshold", which is used to control the depth of pruning.

5 Program examples

5.1 Example from the AQUA program

Consider the following story (New York Times, April 14, 1985) from the domain of AQUA:

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 [by the Islamic Jihad] through another means: blackmail.

AQUA uses its interestingness heuristics to judge the interestingness of this story. As discussed in Ram [1990], AQUA determines that this story, although not personally relevant, is interesting from the point of view of human interest. The story discusses novel explanations for the motivations behind the actions involving the violation of a shared thematic goal, `preserve-life`. Suppose AQUA has read several religious fanaticism stories, but has not encountered any coercion stories so far. After reading the above story, AQUA will be left with several questions, including:

- What did Mohammed value more than his own life?
- Why did the Islamic Jihad plan this mission?
- Why did the Islamic Jihad choose a teenager for the mission?

These questions are represented as knowledge goals in AQUA's memory. On the basis of these knowledge goals, AQUA will now be interested in stories involving the people or institutions it has questions about, such as:

- Another story about Mohammed Burro (relevance to known person)
- Another story about the Islamic Jihad (relevance to known institution)

AQUA will also be interested in stories involving the newly learned blackmail explanation, such as:

- Another story about someone being blackmailed into a suicide bombing mission (relevance to novel explanation)

AQUA will also be interested in a story involving goals or goal priorities that it has questions about, such as:

- Another story about someone valuing something over their own life (relevance to goal)

5.2 Example from the PIES program

The above example illustrates how a program could determine its own interests in different aspects of a situation, and use these interests to guide processing of the current and future stories about similar situations. In the PIES program, interest values and personal relevance values of different concepts are determined directly by the user. The system starts with a model of the user, and uses that model to guide the interest-based information filtering and extraction process. For example, consider the following story:

S-2: The Chicago Cubs played the Atlanta Braves last night at Wrigley Field. Bret Saberhagen pitched for the Cubs, while Tom Glavine started for the Braves. Saberhagen had 3 strikeouts and allowed 2 earned runs, one coming off a solo homerun by Ron Gant in the 3rd inning.

Suppose that PIES is reading this story from the point of view of a reader who is a fan of the Atlanta Braves, is interested in baseball plays and scores, and knows Bret Saberhagen personally. PIES first pre-processes the MOP that is used to represent baseball games, pruning away scenes that are unlikely to be interesting. For example, in this story the location of the game is unlikely to be interesting, since it does not relate to the user's interests in any way. Thus there is no point in looking for, parsing, and representing the location information in this story.

Next, PIES uses the pruned representation of the baseball game MOP to process the story. This process extracts information from the story that matches the representational elements it is looking for. At this point, the story representation may still contain uninteresting facts. For example, if the Cubs pitcher was someone other than Saberhagen, the name of the pitcher would no longer be interesting. This determination can only be made after the story has been processed since if the system were to prune the pitcher information from the baseball MOP ahead of time, it would not be able to represent a situation in which Saberhagen (who the user is interested in) was the pitcher. Thus the final step is one in which PIES post-processes the understood story, pruning away elements that are unlikely to be interesting to the user.

The pruning process is based the program's knowledge of the user, as well as on a user-supplied interest threshold that is used to represent the current inclinations of the user. For example, if the interest threshold is high, the user wants to see only those aspects of the story that are likely to be very interesting. In such a situation, the program prints out the following synopsis of the above story:

- The Chicago Cubs played the Atlanta Braves.
- Bret Saberhagen had 3 strikeouts and allowed 2 earned runs.
- Ron Gant had a home run in the 3rd inning against Bret Saberhagen.

A lower interest threshold may be used if the user has more time to browse through the details of the story. In this situation, more details of the above story would be included in the synopsis.

6 Discussion

Our theory of interestingness is based on a functional analysis of the problem of controlling inferences for cognitive tasks such as story understanding. The theory relies on an analysis of the goals and interests of the reasoner (human or machine), and provides a method for dynamically determining the interestingness of various aspects of the input based on the knowledge goals and interests that the reasoner currently has. Interest in some input arises both through intellectual relevance of that input to the questions that the reasoner is thinking about, or the information that the reasoner needs to collect for some purpose, and also through the personal relevance of the input to the reasoner, the goals of the reasoner, the people and places known to the reasoner, and so on.

The functional utility of our models arises from efficient, goal-directed inferencing, which provides a computational justification for our methods. It is too expensive to use all aspects of all knowledge structures to process every input, and to draw all possible conclusions from every input fact. Based on the current goals and interests of the reasoner, knowledge structures are selected and pruned so that irrelevant aspects of the input will not be processed. For example, if one does not care about where a baseball game is played, there is no point in trying to look for a match for the location scene of the baseball game MOP in the story, since even if such a match is found the information derived will not be useful. However, it is not possible to simply not represent that particular scene in the MOP; this determination must be done dynamically because in a different situation the location may well be important. After the knowledge structures are pruned, they are applied in a manner similar to most script-based or MOP-based understanding systems to yield a representation of the input.

The pre-processing step in PIES is essentially an efficiency heuristic. However, since it is only a heuristic, the system must check the actual representation of the understood story to select the interesting or relevant aspects of the story. Thus while the pre-processing step is an interest-based filter, the post-processing step is an interest-based selector. If efficiency was of no concern, one would simply process the story in great depth, and relegate all interestingness processing to the post-processing step since that is the most accurate way of determining interestingness. However, part of our claim is that interestingness provides a way to extract information likely to be useful without having to process everything in great detail.

Similarly, if more processing were moved into the pre-processing step, PIES would be concerned more about efficiency of processing rather than about the accuracy or functional plausibility of its interestingness determinations. This corresponds to the situation where the program becomes more top-down, guided more by what it wants to find out and less by the information provided by the particular input at hand. More research needs to be done in the area of balancing these two aspects of our system; at present, the balance is determined empirically by the system designer.

More research also needs to be done into dynamic updating of the interestingness values in PIES. For example, a user's interest in baseball games, or in different aspects of baseball games, may change over time. We would like the system to learn what the user does or does not find interesting, and update its user model automatically. AQUA's notion of interestingness is dynamic since it is based on the current set of questions that are active in the system's memory. As these questions change, AQUA's interests change correspondingly. However, although the set of concepts AQUA is interested in changes as it reads, its interest in any given concept is either present or absent; it cannot be incrementally a little more or less interested in that concept in the way that PIES can. We are currently exploring a combination of methods from the two systems to address this problem.

7 Conclusions

The AQUA and PIES projects explore cognitive issues in learning and understanding, and are based on a sophisticated model of interestingness and interest-based control of processing. Interestingness is neither inherent in the information nor in the system, but rather arises from the interaction between the two. It arises from the interaction between the stimulus and the goals of the system. A system with no goals would have no reason to find any input more interesting than any other, nor would any particular piece of information be universally interesting for all systems unless they shared the same goals.

This is a functional approach to the problem of interestingness [Hidi and Baird, 1986; Schank, 1979]. The theory is related to the "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]. Our ideas are consistent with psychological research on goal direction in focus of attention, particularly from social psychology (e.g., see Zukier's [1986] review), and with empirical results on active and goal-oriented learning in teaching situations (e.g., knowledge-building goals [Ng and Bereiter, 1991]).

Although our model is cognitively inspired, the functional justification for the theory is based on computational arguments. There are practical benefits of a program that can represent and reason about its knowledge goals and interests. Such a program can able to focus its limited resources on the relevant aspects of its environment while paying less attention to irrelevant ones. This allows it to spend more time drawing those inferences that are relevant and useful to its goals. Such a program can also form the basis for interest-based information processing applications that can help us deal with the information overload we face today.

References

[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.

[DeJong, 1979] G. F. DeJong. *Skimming Stories in Real Time: An Experiment in Integrated Understanding*. Ph.D. thesis, Yale University, Department of Computer Science, New Haven, CT, May 1979. Research Report #158.

[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.

[Hidi and Baird, 1986] S. Hidi and W. Baird. Interestingness — A Neglected Variable in Discourse Processing. *Cognitive Science*, 10:179–194, 1986.

[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, 1991] A. Ram and L. Hunter. A Goal-Based Approach to Intelligent Information Retrieval. In *Machine Learning: Proceedings of the Eighth International Workshop*, Chicago, IL, June 1991. Morgan Kaufman Publishers, Inc.

[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, 1990] 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, 1991] A. Ram. A Theory of Questions and Question Asking. *The Journal of the Learning Sciences*, 1(3&4), 1991. In press.

[Schank, 1979] R. C. Schank. Interestingness: Controlling Inferences. *Artificial Intelligence*, 12:273–297, 1979.

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

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