Machine Learning: Proceedings of the Eighth International Conference, pp. 265–269, Chicago, IL, June 1991.

A Goal-Based Approach to Intelligent Information Retrieval

Ashwin Ram
College of Computing
Georgia Institute of Technology
Atlanta, Georgia 30332-0280
E-mail: ashwin@cc.gatech.edu

Lawrence Hunter
National Library of Medicine
Building 38A, Mail Stop 54
Bethesda, MD 20894
E-mail: hunter@nlm.nih.gov

Abstract

Intelligent information retrieval (IIR) requires inference. The number of inferences that can be drawn by even a simple reasoner is very large, and the inferential resources available to any practical computer system are limited. This problem is one long faced by AI researchers. In this paper, we present a method used by two recent machine learning programs for control of inference that is relevant to the design of IIR systems. The key feature of the approach is the use of explicit representations of desired knowledge, which we call *knowledge goals*. Our theory addresses the representation of knowledge goals, methods for generating and transforming these goals, and heuristics for selecting among potential inferences in order to feasibly satisfy such goals. In this view, IIR becomes a kind of planning: decisions about what to infer, how to infer and when to infer are based on representations of desired knowledge, as well as internal representations of the system's inferential abilities and current state. The theory is illustrated using two case studies, a natural language understanding program that learns by reading novel newspaper stories, and a differential diagnosis program that improves its accuracy with experience. We conclude by making several suggestions on how this machine learning framework can be integrated with existing information retrieval methods.

1 Inference control for IIR

An important step in the evolution of intelligent information retrieval (IIR) systems involves the loosening of the relationship between a query and information retrieved to address it. A database system finds all of the stored structures that precisely match a query. An IIR system should be able to find and process relevant stored information into an appropriate response to a query. Such a system should be guided by the meaning of its queries, not only by their syntax, and should be able to reason about information implicit in its stored structures. IIR, therefore, involves *inference*.

Systems capable of generalized inference face a problem ubiquitous in AI: a combinatorial explosion of possible inferences. Although the conclusions that can be drawn from a reasoner's knowledge and from available inputs is very large (potentially infinite), the inferential resources available to any reasoning system are limited. In general, reasoning systems simply cannot draw all justified inferences. With limited inferential capacity and very many potential inferences, reasoners must somehow control the process of inference.

Several methods of controlling inference have been proposed. Perhaps the simplest is constrained forward chaining: making as many inferences as possible within the resources constraints. For example, MARGIE [Rieger, 1975] made all the justified inferences that required the chaining of no more than 5 of 17 rules. The amount of inference scales with the number of inference rules to the power of the length of possible chains, so in practical circumstances only a small percentage of justified inferences can be drawn with this method. Empirically, many of the inferences generated this way are useless, and many useful inferences are missed because their derivations are too long. Other systems rely on backward chaining to make only the inferences that might lead to a specified outcome. Unfortunately, many valuable inferences (even quite simple ones) are overlooked by this method, since surprises are impossible; the system can only infer what it's already looking for. Still another method tries to use probability measures to draw the most likely inferences. However, some relatively unlikely inferences can be very important. Reasoners may be explicitly interested in identifying unlikely events that have significant consequences, and therefore may make inferences about, say, what they would do if they won the lottery or if the system on which they maintained their database crashed.

Our approach to this problem has been to make the *utility* of inferred assertions (i.e., beliefs) an explicit part of the inference process. Not all beliefs are equally useful to a given reasoning system. For example, accurate beliefs are usually (but not always) more valuable than inaccurate ones. Beliefs can be causally involved in success and failure of various kinds. To the degree that a belief contributes to achieving positively evaluated outcomes, it has positive utility. Although the utility of a belief is likely to vary from one believer to another, and even over time within a single individual, the fact that an assertion can be more or less valuable to a particular understander in a particular situation has important implications in the design of inferential systems.

Inference can be controlled by identifying beliefs that would be valuable if they should be inferred, and then attempting to infer them. These valuable, potentially inferrable, assertions can be reified as *explicit desires for knowledge*. The question of focus of attention is thereby transformed into two related problems: Where do these desires for knowledge come from? How can explicit desires for knowledge be used to control inference? To address these questions, we must consider the uses of knowledge and inference, and how to assess the value of knowledge in order to control inference.

This paper presents a theory of inference control for understanding and learning that is based on the notion of *knowledge goals* (or knowledge acquisition goals), a reasoner's specific desires to acquire, retrieve or otherwise infer information or knowledge. Our theory addresses the representation of knowledge goals, methods for generating and transforming these goals, and heuristics for selecting among potential inferences in order to feasibly satisfy specific desires for knowledge. Our approach to restricting potentially inferrable hypotheses is *content*-based. Explicit characterizations of *desirable knowledge* or *required knowledge* provide a principled method for restricting the realm of inputs and background knowledge considered in inference, and thereby the size of the hypothesis space that must be considered. This method of restricting the space of possibly relevant inferences is *dynamic*; the inferential bias changes depending on what is already known by the system, and on what it is trying to find out. In fact, we can carry this idea one step further: Not only can explicit goals about knowledge help appropriately bias inference, they can be used to direct *action* intended to accomplish those goals. Rather than passively waiting for useful information to show up, a system can actively pursue the knowledge it desires, using specific learning plans or instantiations of general learning strategies. In this view, intelligent information retrieval is the selection, combination and execution of (usually inferential) actions to satisfy specific goals for knowledge.

2 Computer programs with knowledge goals: Two case studies

The utility of goal orientation as a basis for theories of information processing has been proposed previously, both as a psychological model of focus of attention during text processing in people (e.g., [Zukier, 1986; Hoffman *et al.*, 1981; Srull and Wyer, 1986]) and as a computational model for artificial intelligence programs (e.g., [Ram, 1989; Hunter, 1989; Hunter, 1990a; Ram, 1990b; Ram and Hunter, 1991]). In this paper, we will discuss the general issues and then concentrate on programs performing information retrieval tasks in two broadly representative areas: natural language understanding and medical diagnosis.

In complex knowledge-based systems, it is nearly impossible to create a system that contains all the knowledge it needs in order to accomplish its goals. Instead, such systems should be able to improve their performance with experience. Both natural language texts and medical cases provide a multiplicity of possible inferences that might conceivably be useful in improving the abilities of a performance system. In both areas, the use of explicit knowledge goals helps narrow the vast space of possible inferences to a more manageable set, and to help the program make decisions about when to draw potential inferences.

Neither of the two systems presented below are information retrieval systems *per se*. However, both systems manage knowledge bases that can be used for information retrieval. These systems generate their own desires for knowledge, and manage inference in order to accomplish those goals. Theoretically, we see them as proof of the concept, rather than as prototypes for IIR systems. Pure IIR systems using this approach might take as input knowledge goals generated by users, and manage automated inferential resources, such as access to multiple databases, statistical inference or graphical displays to respond to the user's goals. User's questions that cannot be answered by the system could leave traces or *persistent goals*. When the information came available to address that goal, the user could be notified. We expect the gradual merging of IIR with knowledge-based systems, so that, for example, a scaled up version of the kind of natural language system described below could be used to actively gather desired knowledge from multiple streams of textual information.

2.1 The AQUA program

Our first example program is AQUA, a story understanding program that learns from what it reads [Ram, 1987; Ram, 1989]. In order to understand text, the performance system must integrate the text, which is often ambiguous, elliptic and vague, with its world knowledge, which is often incomplete and possibly incorrect. In order to learn from what it reads, the system must detect perceived anomalies in the text which may identify flaws or gaps in the its model of the domain, formulate explanations to resolve those anomalies, confirm or refute potential explanations, and learn new explanations or modify incorrect ones.

The process of natural language understanding generates *information goals* or *questions*, representing what the understander needs to know in order to do an understanding task, be it explanation, learning, or some other cognitive task. These questions constitute the specific knowledge goals of the understander generated during a parsing experience, and are used to focus the reasoning processes on aspects of the input that are actually relevant. These goals are also used to focus the learning process so that the understander learns what it needs to know in order to better carry out its tasks.

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. 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. 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.

2.2 The IVY program

Another, rather different, example is IVY, a program that does differential diagnoses and is intended to improve its accuracy with experience [Hunter, 1989]. The basic idea was to design a program that improves its accuracy by storing information from the cases it diagnoses correctly. The problem is that there is a huge amount of information in the correctly diagnosed cases. Very little of that information is useful for improving diagnostic performance: after all, most cases are handled correctly using the basic domain knowledge. On the other hand, there are nuggets of information in that set of cases that would be very useful for improving performance. To find the useful bits without drowning in irrelevant information, IVY generates specific goals for knowledge that is likely to improve performance, and then look for that information in its case load. IVY uses explanations of its performance failures to identify aspects of its diagnostic knowledge that are missing or incorrect. Those explanations are transformed into characterizations of information that would be useful to have to avoid the failures. That characterization is effectively a knowledge goal, and can be used to rapidly scan correctly diagnosed cases for information that would help address previously encountered problems. It is interesting to note that in order to identify the problematic knowledge and generate the goal, the learner must have at least a simple model of its own reasoning process. In order to relate a performance failure to missing or incorrect knowledge, the program must be able to do some reasoning about how it uses its knowledge.

Once analysis of failures has led to the generation of knowledge goals, the program can plan to acquire the knowledge. For IVY, the plans involved looking for specific kinds of information in one or more cases (either cases already in memory or as they arise for diagnosis), and then to transform and store the information so that it address the cause of the motivating failure. IVY's plans were capable of using case information both to supplement its general abilities with specific knowledge and to find and store exceptions to its rules.

2.3 Information seeking as a planful activity

Before exploring these two programs in more detail, let us consider the commonalities in use and generation of knowledge goals between the programs. Both programs use desires about knowledge to control potentially explosive inferential processes. For AQUA, the number of inferences that can be drawn from a story is very large. AQUA only draws those inferences that are likely to answer questions that it has, or to raise new questions that are relevant to its goals. IVY is looking for ways to improve its diagnostic performance. Any given case might be relevant to a problem that has occurred in making a diagnosis. Sometimes a relevant case may not even involve the same disease that caused the original problem. Any aspect of a new case might be relevant to any aspect of the program's diagnostic knowledge. The number of possible interactions between all aspects of all cases and all the program's prior knowledge is huge. IVY reduces this search space dramatically by characterizing the knowledge it desires as specifically as possible. Both programs instantiate our view that information seeking is a planful activity, in which the reasoner actively pursues the desired information (i.e., its knowledge goals) using available information retrieval and inference methods (i.e., knowledge acquisition plans).

It is also apparent that, for both programs, characterizing desirable knowledge requires the ability to reason about internal reasoning processes and the knowledge they use. This ability to dynamically evaluate the knowledge used by internal processing (e.g., to find gaps in that knowledge that, had they been filled, would have changed the processing) may be a general feature of many kinds of learning systems. This kind of reasoning about internal processing and knowledge is a form of introspection that we conjecture may be a necessary component of human-like learning. We argue that planful information retrieval is essentially a learning process, and that algorithms for information retrieval be integrated with corresponding methods for learning. We call the integrated method *knowledge acquisition planning.*

2.4 AQUA's knowledge goals

The questions that AQUA pursues are based on a taxonomy of types of knowledge goals (KGs). This taxonomy arises from the understanding tasks that underly AQUA's processing, as well as from a general set of "interests" that AQUA begins with.

Text goals: KGs 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: KGs 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.

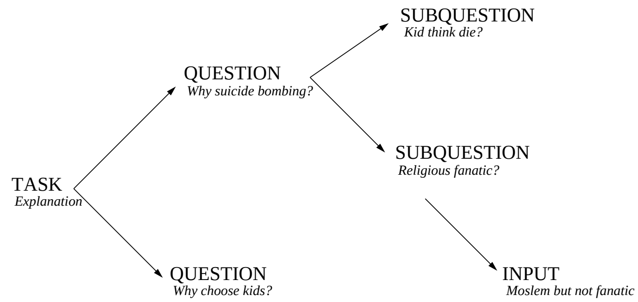

Figure 1: Questions underlying the religious fanatic explanation for suicide bombing, showing the relationships between tasks, knowledge goals, and information ultimately found in the input text.

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 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: 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.

The basic process of goal-based understanding involves the generation of knowledge goals seeking information required by various understanding tasks, the transformation of these KGs into subgoals, and the matching of pending KGs 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. Example explanation goals for a typical suicide truck bombing story are shown in figure 1, which represents the questions one might think about while reading such a story.

Each of these types of KGs is expressed as a question that focuses on a different aspect of the story. For example, explanation questions focus on different types of anomalies, and on explanations for these anomalies. It is useful to focus on KGs because they arise from a "need to learn." There are two basic ways in which a fact can turn out to be worth processing in this sense:

Top-down: A fact helps achieve a KG, or answers a pending question, is worth focussing on since it allows the reasoning system to continue the reasoning task that required the knowledge in the first place.

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

These correspond to the two diamonds in figure 2. Heuristics based on these measures is used to determine the "interestingness" or "utility" of the desired information, which in turn

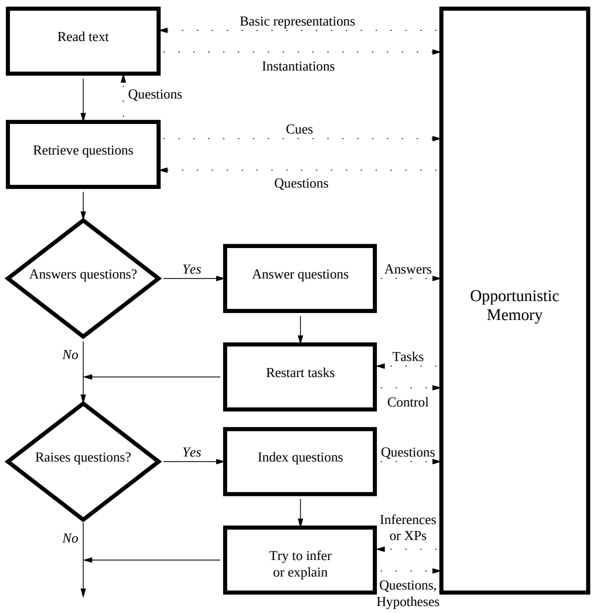

Figure 2: Control structure: The understanding cycle in AQUA. A fact is interesting if it satisfies a knowledge goal pending in memory, or if it gives rise to new knowledge goals. Uninteresting facts pass vertically down with minimal processing; interesting facts cause suspended understanding tasks to 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.

provides a principled method for controlling inferences during the overall task of extracting information from natural language texts [Ram, 1990b].

2.5 IVY's knowledge goals

KGs are generated and acted upon somewhat differently in the IVY program. IVY's task is to diagnose structured descriptions of lung tumor pathology images. These descriptions contain information at various levels of detail about populations of cells taken from lung and colored with various stains. The amount of information available in a typical input is very large (on average, IVY's case descriptions contained 116 slots), and most of it is not directly relevant to making a diagnosis. Most of the diagnoses are imprecise, and do not have definitions in terms of features that are individually necessary and collectively sufficient for their identification. There is no unambiguous mapping from characteristics to diagnoses in this domain.

IVY, like human pathologists, uses the method of differential diagnosis to arrive at a diagnosis. First, the general classes of potentially relevant diseases are identified. Then, these general classes are specified as much as possible to identify candidate diagnoses. Finally, the candidate diagnoses (together referred to as the differential) are compared, and the best is selected. A lung tumor pathology expert evaluates IVY's conclusions, specifying the correct diagnosis for each case. If the program's diagnosis was incorrect, the program explains its failure, and generates one or more goals to find knowledge that would help it avoid repeating the failure. If the diagnosis was correct, the program examines its pending KGs to see if any of them can be satisfied. The basic step in the KG generation process is explanation of a performance failure. These explanations have two components: the process that failed, and the kind of knowledge used (or not used) by that process which caused the failure. In order to determine what knowledge it needs to learn to avoid the failure, it must reason about how it used the knowledge it had. By working backwards and considering the inputs and outputs of each step, it is possible to determine in which diagnostic step the error was made, and from there, what knowledge was missing or incorrect. The decision tree for identifying the step in the diagnostic process that failed is illustrated in figure 3.

```
          <correct hypothesis> is on the <differential>?
              /                  \
           yes                    no
           /                       \
distinction   generalization of <correct hypothesis>
failure              is on the <differential>?
                                /     \
                             yes       no
                             /           \
     specification failure       recognition failure
```

Figure 3: The explanation decision tree for identifying the step in the diagnostic process that led to a failure.

IVY then turns these explanations identifying the knowledge involved in a failure into specific desires for knowledge that would ameliorate the problem. Each general explanation type (i.e., each combination of failed process and type of failed knowledge) has associated with it a KG skeleton. For example, distinction failures due to missing distinction knowledge have an associated KG skeleton for finding knowledge that can differentiate between two competing diagnoses. Newly generated KGs are then input to a knowledge acquisition planner. Each goal leads to the generation of one or more knowledge acquisition plans. A plan specifies information that would be useful in addressing a goal (the preconditions), and what to do when that information becomes available (the actions). When a plan's preconditions are met, the plan's actions are executed.

People use a wide variety of plans to achieve their KGs, ranging from simply looking up an answer in a reference book to designing and executing scientific experiments. IVY's plans identify methods for finding desired information in single cases. After each successful diagnosis, IVY compares the unsatisfied preconditions of all of its pending plans to the contents of the just diagnosed case. If the case can be used to satisfy the preconditions, the plan specifies how to store the case in memory so that the diagnostic failure that motivated the plan is addressed. IVY's planning abilities were limited to selection among and instantiation of eight different plan schemas. Knowledge goal skeletons had from one to three plan schemas associated with them. INVESTIGATOR [Hunter, 1990b; Hunter, 1990a] uses a more flexible knowledge acquisition planner. Depending on the specifics of the knowledge goal (i.e., characteristics of the variable bindings in the goal skeletons), one or more of the plan schemata are instantiated.

IVY ultimately diagnosed 118 descriptions of lung tumors, selected to be broadly representative. IVY was capable of diagnosing 95 of the 118 cases correctly without learning (about 80%). The goals generated by three of those failures could eventually be satisfied by the program, leading to four additional correct diagnoses (about 84% success). More important than any measure of percentage improvement in performance due to learning is the *quality* of the knowledge the program acquired. Two of the three "lessons" that IVY learned in response to its knowledge goals were identified by the domain expert, Dr. Yesner, as good teaching cases. One of the images IVY selected had been previously used as an example in one of Dr. Yesner's publications [Yesner and Carter, 1982]. The third case discovered by IVY was considered by Dr. Yesner "quite useful" in context. The ability of a program to independently identify cases that a domain expert considers interesting, on the basis of the program's experience and its consequent desire for information, bodes well for the potential of knowledge acquisition planning in IIR.

The potential application of the reasoning techniques demonstrated in IVY to IIR are two-fold: First, the analysis of task performance failures can be used to identify missing or incorrect knowledge. Second, automatically generated characterizations of desired knowledge can be used to screen large amounts of information that changes over time to find useful knowledge. Later work [Hunter, 1990a; Hunter, 1990b] has also shown that abstractly specified desires for knowledge can be transformed, in a process similar to subgoal decomposition planning, into a series of database queries and inferential steps that address the abstract goal. This ability to translate abstract knowledge goals into plans of (inferential) action may be useful as a component of an advanced query system.

3 Applications of knowledge goals to IIR: Some speculations

Intelligent information retrieval requires inference; unfortunately, computing the inferences entailed by even a moderately powerful system is potentially intractable. Our approach to the homologous problem in machine learning has been to trade this intractable search space for another: the space of possible plans to acquire knowledge. By identifying desired knowledge explicitly, and distinguishing among types of desires, we have the basis for building systems that are capable of effectively pursuing those desires. Heuristic methods for finding desired knowledge seem to us to work best when their decisions are based on the *content* of the desired knowledge, in addition to its form. As information retrieval systems come to embrace multiple databases and multiple inferential tools, automated decisions about which databases and which tools are appropriate for addressing a query will have to be made. In our theory, those choices are appropriately viewed as a kind of planning.

Even simple uses of explicit knowledge goals in IIR have the potential to offer innovative services. For example, the failure of a query (or perhaps a series of queries) to generate any hits can be used to leave a representation of the query and the requester at the location in the system where the desired knowledge would be stored. Should knowledge later be added to the database that matches the stored goal, the requester could be notified that there is now an answer to the query.

More sophisticated uses of knowledge goals can also be envisioned. We have demonstrated that programs can automatically generate desires for knowledge during other kinds of processing. By explaining the causes of performance failures (using simple models of internal processing), IIR systems might be able to characterize knowledge desired by their users but absent from the system. It may be possible to generalize from a large number of individual goals into characterizations of areas that the managers of an IIR system might pursue to improve the effectiveness of their systems. Finally, we see the eventual merging of IIR systems with active machine learning into systems that dynamically build knowledge bases that address user needs. Ideally, we view information retrieval as an incremental process of belief formation, involving a wide variety of interrelated inference and learning processes, driven by explicit motivations to learn. Integrated learning and retrieval systems should be able to notice interesting aspects of their experiences (both in processing queries and in acquiring information from other sources), generate knowledge goals based on those observations, and devote some of their resources to achieving those goals [Hunter, 1990b; Ram, 1990a]. The process of satisfying those goals generally involves both the selection of an appropriate learning method and the focussing of attention on potentially relevant information sources.

The theory of knowledge goals and planning presented in this paper is both cognitively motivated and functionally justified. We believe that it suggests a useful direction for the evolution of intelligent information retrieval toward systems that can represent, reason about and effectively manage their own information-seeking behavior.

Acknowledgements

Ashwin Ram's research was supported in part by the National Science Foundation under contract IRI-9009710.

References

[Hoffman *et al.*, 1981] C. Hoffman, W. Mischel, & K. Mazze. The role of purpose in the organization of information about behavior: Trait-based versus goal-based categories in person cognition. *J. of Personality and Social Psych.*, 39:211–255, 1981.

[Hunter, 1989] L.E. Hunter. *Knowledge Acquisition Planning: Gaining Expertise Through Experience.* PhD thesis, RR#678, Yale Univ., New Haven, CT, 1989.

[Hunter, 1990a] L.E. Hunter. Knowledge Acquisition Planning for Inference from Large Datasets. In *Proc. 23rd Annual Hawaii Intl. Conf. on System Sciences*, pp. 35–44, Kona, HI, 1990.

[Hunter, 1990b] L.E. Hunter. Planning to Learn. In *Proc. 12th Annual Conf. of the Cognitive Science Society*, Boston, MA, 1990.

[Ram and Hunter, 1991] A. Ram & L. Hunter. Goals For Learning and Understanding. *Applied Intelligence*, 1991. (To appear.)

[Ram, 1987] A. Ram. AQUA: Asking Questions and Understanding Answers. In *Proc. 6th Annual Natl. Conf. on Artificial Intelligence*, Seattle, WA, 1987.

[Ram, 1989] A. Ram. *Question-driven understanding: An integrated theory of story understanding, memory and learning.* PhD thesis, RR#710, Yale Univ., New Haven, CT, 1989.

[Ram, 1990a] A. Ram. Incremental Learning of Explanation Patterns and their Indices. In *Proc. 7th Intl. Conf. on Machine Learning*, Austin, TX, 1990.

[Ram, 1990b] A. Ram. Knowledge Goals: A Theory of Interestingness. In *Proc. 12th Annual Conf. of the Cognitive Science Society*, Cambridge, MA, 1990.

[Rieger, 1975] C. Rieger. Conceptual Memory and Inference. In R.C. Schank (ed.), *Conceptual Information Processing.* North-Holland, Amsterdam, 1975.

[Srull and Wyer, 1986] T. Srull & R. Wyer. The Role of Chronic and Temporary Goals in Social Information Processing. In R. Sorrentino & E. Higgins (eds.), *Handbook of Motivation and Cognition: Foundations of Social Behavior*, pp. 503–549. Guilford Press, Guilford, CT, 1986.

[Yesner and Carter, 1982] R. Yesner & D. Carter. Pathology of Carcinoma of the Lung: Changing Patterns. *Clinics in Chest Medicine*, 3(2):257–289, 1982.

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