Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, 1993, pp. 24–29

Focusing Construction and Selection of Abductive Hypotheses*

David B. Leake
Department of Computer Science
Indiana University
Bloomington, IN 47405, U.S.A.

Abstract

Many abductive understanding systems explain novel situations by a chaining process that is neutral to explainer needs beyond generating *some* plausible explanation for the event being explained. This paper examines the relationship of standard models of abductive understanding to the *case-based explanation* model. In case-based explanation, construction and selection of abductive hypotheses are focused by specific explanations of prior episodes and by goal-based criteria reflecting current information needs. The case-based method is inspired by observations of human explanation of anomalous events during everyday understanding, and this paper focuses on the method's contributions to the problems of building good explanations in everyday domains. We identify five central issues, compare how those issues are addressed in traditional and case-based explanation models, and discuss motivations for using the case-based approach to facilitate generation of plausible and useful explanations in domains that are complex and imperfectly understood.

1 Introduction

The success of abductive understanding systems depends on their methods for performing the two fundamental steps of abductive explanation: generating candidate explanations and choosing between competing alternatives. Many models of abductive understanding construct explanations by chaining from scratch in a process that is neutral to any system goals beyond simply generating *some* plausible explanation for the event being explained. However, applying that method in the rich domain of everyday events is impeded both by well-known problems of explanation construction cost and by another crucial problem: that the goodness of explanations depends not only on their validity but also on whether they provide the information that the explainer needs [Leake, 1991a; Leake, 1992; Ram and Leake, 1991].

The need for explanations to reflect explainer goals both complicates explanation selection and potentially increases the number of candidate hypotheses that must be generated. It is not sufficient for an explainer with overarching goals to generate *any* valid explanation; the explainer must generate a valid explanation that provides the information it needs. Although in restricted domains it may be possible to limit system rules in such a way as to assure that all explanations will be relevant to system goals, this is not possible in general. In a system with a rich knowledge base that builds everyday explanations, focusing explanation is a key problem.

This paper examines the significance of the case-based explanation process (e.g., [Kass and Leake, 1988; Leake, 1992; Schank, 1986; Schank and Leake, 1989]) and its theory of explanation evaluation (e.g., [Leake, 1991a; Leake, 1992]) as methods for focusing construction and selection of abductive hypotheses. In the case-based approach, which has been applied to the task of abductive understanding of anomalous events in news stories, explanation construction by chaining is replaced by analogical reasoning focused by three types of context: system beliefs and expectations, prior episodes that the system has explained, and the goals that currently motivate explanation. Case-based explanation builds new explanations by retrieving stored explanations for previous episodes and adapting them to fit current circumstances and needs. The case-based process models some aspects of human explanation [Read and Cesa, 1991], and, as we describe in the following sections, facilitates generation of plausible and useful explanations despite the problems of incomplete information and imperfect domain theories that mark everyday explanation.

Because detailed descriptions of the mechanisms involved have been published previously, this paper will not discuss specifics of how those mechanisms make case-based explanation construction practical. Instead, its goal is to delineate the relationship of this abductive reasoning process to other models of abductive understanding.

We begin with a brief overview of the case-based explanation model. We then identify five central issues in abductive understanding of everyday events, sketch how those issues are addressed in traditional and case-

*I would like to thank Ashwin Ram for helpful discussions of these issues and the IJCAI reviewers for their useful comments.

based explanation models, and discuss motivations for using the case-based approach to facilitate explanation in complex and imperfectly understood domains.

2 Overview of case-based explanation construction

The case-based model of abductive understanding was originally investigated in SWALE, a story understanding system that uses case-based reasoning to generate abductive explanations of anomalous events [Kass and Leake, 1988; Schank, 1986; Schank and Leake, 1989]. Specific aspects of the model were further refined in SWALE's descendents ABE [Kass, 1990] and ACCEPTER [Leake, 1992]. These systems use case-based reasoning to focus explanation of real-world events. Rather than assuming perfect knowledge, their explanation effort applies limited background knowledge to the incomplete information provided in simple news stories. The goal of their processing is to generate explanations that are plausible and that provide adequate information for particular types of overarching goals. Examples processed by ACCEPTER include the explosion of the space shuttle Challenger, the accidental shootdown of an Iranian airliner by the American warship Vincennes, and the namesake example of SWALE: the story of a star racehorse named Swale who was at the peak of his career when he was found dead in his stall.

In the Swale example, a vast range of possible causes could be hypothesized to explain the death—given the lack of information in initial news reports, almost any cause of death was a potential candidate. However, people confronted with the Swale story appeared to have little difficulty generating plausible hypotheses. As they generated these hypotheses, they often attributed the hypotheses to being reminded of prior episodes. For example, one person was reminded of the death of the runner Jim Fixx, who died when the exertion of recreational jogging overtaxed a hereditary heart defect. The explanation for Fixx's death does not apply directly to Swale—Swale was unlikely to do recreational jogging—but minor adaptation of the explanation, substituting horse racing for jogging, produces the plausible explanation that the stress of running in a race overtaxed a hereditary heart defect. This example and similar informally-collected accounts helped to suggest the case-based model of explanation used by the SWALE system, in which new explanations are built by retrieving and adapting explanations of previous episodes. Later psychological experiments support the psychological validity of the reminding-based explanation process and the tendency of people to favor explanations that are based on prior explanations of similar episodes [Read and Cesa, 1991].

In somewhat more detail, the basic steps of SWALE's case-based explanation algorithm are as follows:

- **Problem characterization:** Describe the information that a good explanation must provide, using the same vocabulary used to index explanations stored in memory.
- **Explanation retrieval:** Use the problem characterization as an index for retrieving relevant explanations of prior episodes from memory.[1]
- **Explanation evaluation:** Evaluate the retrieved explanations' plausibility and usefulness. Generate problem characterizations for any problems.
- **Explanation adaptation:** If problems were found, use the evaluator's problem characterization to select *adaptation strategies* [Kass, 1990] for modifying the explanation to repair the problems. Apply the strategies and return to the explanation evaluation phase to evaluate the new explanation.

In order to establish this algorithm as a viable alternative to standard abductive methods, the effectiveness of each phase of the process must be substantiated. That substantiation is beyond the scope of this paper but has been provided elsewhere: problem characterization and retrieval issues are addressed in [Leake, 1991b; Leake, 1992], explanation evaluation issues are addressed in [Leake, 1991a; Leake, 1992], and adaptation issues are addressed in [Kass, 1990].

3 Perspective on the case-based model

To delineate the ramifications of the case-based model and its relationship to other methods for abductive understanding, our central focus is to examine and put into perspective the tenets of the case-based model concerning five key issues: (1) how to build explanations, (2) the nature of explanations, (3) the plausibility evaluation process, (4) the role of anomalies in focusing an understander's explanation, and (5) the influence of overarching goals on explanation.

3.1 Building explanations

Explanation construction methods: In abductive understanding systems, standard theorem-proving chaining techniques are generally the mechanism for generating candidate explanations (e.g., [Hobbs *et al.*, 1990; Kautz and Allen, 1986]). A well-known problem for these methods is the cost of explanation construction, due to the combinatorial explosion of alternatives to consider. Methods have been proposed to control chaining cost (e.g., [Charniak, 1986; Hobbs *et al.*, 1990]), but despite the benefits of these methods, the difficulty of efficiently generating explanations remains acute in rich domains. Case-based explanation addresses this problem by adapting prior explanations to new situations, to avoid the cost of chaining from scratch.

The case-based approach to re-using explanations contrasts with another method that learns to facilitate explanation, explanation-based schema acquisition (e.g., [Mooney, 1990]). When confronted with novel situations, explanation-based schema acquisition builds

[1] SWALE, ACCEPTER and ABE all start with libraries of stored explanations. In general, a case-based explanation system's initial explanation library could be provided by external sources (e.g., by being told explanations or reading about explained episodes) or built up by traditional chaining methods, and the library is augmented by storing new explanations that the case-based explanation process generates.

new explanations by chaining and then immediately uses explanation-based generalization of the explanatory chain to form a general schema for future use. Instead of generalizing new explanations when they are first generated, case-based explanation construction stores variablized versions of the *specific explanations.* When those explanations are applied to new situations it adapts them to fit the new situation, using adaptation strategies that include not only generalization but also deletion, addition and replacement of components of the explanation [Kass, 1990]. Consequently, case-based explanation can apply prior explanations to a wider range of circumstances that explanation-based generalization. Although a stored explanation can potentially be applied to an enormous range of new situations, processing effort is controlled because adaptation is only done to the extent needed to explain a specific new situation.

As an example of the contrast between the case-based approach and explanation-based schema acquisition, we return to the Swale example. When the SWALE system explains Swale's death, one of the explanations it retrieves concerns the death of another young superstar, Janis Joplin. Joplin was driven to recreational drug use by the stress of being a star and the availability of recreational drugs, and she died from an accidental drug overdose. An explanation-based schema acquisition program that had previously explained Joplin's death would have generalized the explanation for Joplin's death at the time it was built, to form a general schema such as "stress and access to drugs can lead to death from accidental overdose." However, such a generalization does not apply to Swale. Consequently, the explanation-based schema acquisition system could not apply it and would chain together a new explanation without guidance from the Joplin explanation.

In SWALE, the applicability of a new explanation is not bound by a precomputed generalization. Instead, after retrieving the specific explanation of Joplin's death it attempts to decide how to adapt it in light of the particulars of the episode—Swale's death—to which it will be applied. During adaptation, SWALE abandons the parts of the explanation that do not apply and retains the kernel of the explanation that is potentially applicable: the hypothesis that a drug overdose caused the death. That hypothesis is unsupported, so adaptation takes drug overdose as a starting point (considerably narrowing the field of options to consider) and seeks additional support. Because its knowledge includes that racehorses are sometimes given performance-enhancing drugs, it generates the explanation that Swale might have died from an accidental overdose of performance-enhancing drugs. Thus the case-based approach uses experience to suggest alternatives even in situations that are not straightforwardly subsumed by generalizations of prior explanations, allowing more flexible reuse of the results of prior explanation construction. This process depends on having effective strategies for guiding adaptation, and a library of such strategies is proposed in [Kass, 1990].

We note that during the adaptation process, case-based explanation uses incremental evaluation of hypotheses to decide how to proceed further. This approach differs from models of abduction in which all explanation construction is assumed to precede any evaluation, and is in the same spirit as recent research on choosing which explanations to pursue according to validity estimates (e.g., [deKleer and Williams, 1989; Ng and Mooney, 1990]. However, the case-based model uses on-going evaluation not only to choose which explanation to pursue, but, because the evaluation suggests the type of information to seek to repair a faulty explanation, to provide very specific guidance of how to proceed when augmenting or modifying candidate explanations.

Ramifications for quality of explanations in imperfectly understood domains: Explanation-based generalization is guaranteed to produce correct generalizations even from a single episode, provided it starts with a perfect domain theory. Unfortunately, domain theories of the unrestricted everyday world are unavoidably imperfect. The case-based method generalizes and adapts explanations only to the extent required by new examples, and verifies the reasonableness of its results by evaluating those explanations in the current situation. This reduces the danger of faulty generalization.

In addition, the case-based approach helps to choose between competing explanations licensed by an inconsistent domain theory. By favoring explanations supported by specific similar experiences, case-based explanation takes advantage of regularities in the world—similar events are explained in similar ways—even if those regularities are not fully captured by the explainer's domain theory. Previous research on overcoming imperfect theory problems focuses on methods for repairing an imperfect domain theory when problems arise (e.g., [Dietterich and Flann, 1988; Rajamoney, 1988]); our focus is how to generate reasonable explanations despite possible imperfections in the domain theory.

3.2 The nature of explanations

The need of everyday explainers to use imperfect information is reflected in how case-based explanation views the nature of explanations. In most AI views of explanation, explanations are treated as deductive proofs. Abductive reasoning systems build their proofs by nondeductive methods, and additional assumptions may be required for those proofs to apply. However, their view is fundamentally the same in that if the abductive assumptions were shown to be true the resulting explanation would be considered a deductive proof.

The case-based approach, however, explicitly treats explanations as plausible reasoning chains that may be imperfect. In the case-based model, explanations are represented as *explanation patterns* (XPs) [Schank, 1986]. XPs trace a chain of reasoning that accounts for why antecedents of the XP provide support for belief in the consequent of the XP, but the reasoning chain is not considered to prove that the consequent must necessarily hold. (In this respect XPs are in the same spirit as Pearl's plausible inference networks [Pearl, 1988].) Consequently, decisions about the plausibility of an XP in a given situation depend both on its abductive assumptions (as in standard models) and on how well the internal derivation of the belief-support chain of the XP applies in the current situation [Leake, 1992].

3.3 Plausibility evaluation

In case-based explanation, the first criterion for selecting likely explanations is experience in similar situations: Explanations of new situations are considered most plausible if they have applied in similar prior situations. However, additional plausibility evaluation is needed to verify adapted parts of the retrieved explanation and to determine the plausibility of the explanation's assumptions and reasoning chain in the new circumstances.

In abductive reasoning systems, the dominant method for judging plausibility of explanations is to favor explanations that are in some sense structurally "minimal" (e.g. [Charniak, 1986; Kautz and Allen, 1986]). These methods stress factors such as the *number* of abductive hypotheses rather than their content. Instead, the case-based model relies primarily on the content of the assumptions and explanatory chain. Its approach to judging plausibility of assumptions is in the spirit of systems using probabilistic criteria (e.g., [Charniak and Goldman, 1991]), but the situations it attempts to explain are situations for which probability information is unlikely to be available. Consequently, the case-based approach evaluates individual parts of an explanation by similarity: It compares assumptions to standard patterns and considers the assumptions likely if they match those patterns. Similarity-based methods are not guaranteed to parallel correct probabilities and can sometimes lead to errors, but they are heuristics that people appear to use to estimate likelihoods when probabilities are unavailable [Kahneman *et al.*, 1982]. A full description of our model's plausibility evaluation can be found in [Leake, 1992].

3.4 The role of anomalies in focusing explanation

Standard abductive understanding systems take a neutral view of the events they explain: given an event to explain, they either accept any chain accounting for the event or always seek explanations focusing on a fixed aspect of the event (e.g., always trying to explain in terms of the goals that an action satisfies [Mooney, 1990]). In such systems, decisions of which explanation to select are based entirely on plausibility of the candidates (e.g., [Charniak, 1986; Charniak and Goldman, 1991; Hobbs *et al.*, 1990; Kautz and Allen, 1986; Ng and Mooney, 1990]).

However, plausibility alone is not sufficient to distinguish between candidate explanations—many different valid explanations can be generated for any event. One possible response is to attempt to pursue explanations from *all* perspectives [Ajjanagadde, 1991], but this aggravates the already expensive explanation task. Instead, case-based explanation takes the view that considerations of information needs should determine the aspects of an event to explain.

For example, consider a few of the many possible explanations for the statement "John used a blowtorch to break into an automatic teller machine (ATM).":

- John needed money to pay back a loan shark for gambling debts.[2]
- John believes that robberies of ATMs are more likely to succeed than bank robberies.
- Mark, a gang member who was originally expected to do the break-in, was sick, forcing John to replace him at the last minute.
- The bank's security camera had been removed for repair.

These examples show that plausibility considerations alone are inadequate to select a good explanation: All four explanations might simultaneously be valid, but their goodness depends on explainer motivation for explaining. For example, if the explanation effort was motivated by being surprised that John performed the robbery instead of Mark, the explanation that Mark was sick would be relevant but the explanation that the security camera was broken would not. If the explainer was surprised that the break-in succeeded despite bank precautions, the explanation that the security camera was missing would be relevant and Mark's illness would be irrelevant. Explanation construction processes that fail to reflect the reasons for explaining are doomed to impartially generate irrelevant candidates as well as relevant ones, exacerbating the problem of explanation construction cost and resulting in explanations that are irrelevant to the explainer's information needs.

In the case-based model, explanation retrieval is focused according to the explainer's information needs to resolve the anomalies that prompted explanation. ACCEPTER directs explanation search by retrieving explanations indexed as relevant to the anomalies prompting explanation [Leake, 1991b; Leake, 1992], and includes a processing phase that explicitly evaluates the relevance of candidate explanations to the anomaly being explained [Leake, 1992].

Although standard abductive understanding systems do not attempt to select the aspects of an event on which to focus, at first glance it appears that their basic backwards chaining mechanism would be sufficient to allow explanation to be focused on resolving anomalies. For example, the query to be explained by a backwards chaining system could be the particular aspect of the event that was anomalous, such as John's decision or the failure of John to be seen during the break-in. However, this provides only part of the needed focus for actually resolving the anomaly.

For example, suppose the anomalous aspect of the break-in is that John was not seen, and suppose that the explainer tries to explain it by chaining from the query

[2] We state the antecedents alone as a shorthand for the reasoning chain that makes up the complete explanation. In this example, the entire explanation might include the fact that loan sharks place their victims under duress to pay their debts; the inference that being placed under duress prompts a high-priority goal to obtain money; that robbery is a plan for obtaining money; that a step in robbery is forcibly gaining access to the money to be robbed, etc.

"John was not seen during the break-in." Many different explanations can still be generated for this query, and their relevance would depend on the explainer's previous expectations. For example, the explanation "the lights were broken" is relevant if the understander thought that John would be seen because of a reasoning chain including that the ATM was well-lighted; it is not relevant if the understander already knew that the ATM was dark but also knew that the guard had night-vision goggles that would enable him to see John regardless of the darkness. Thus relevance of an explanation depends not only on the aspect of the event being explained, which is easily captured by selection of the goal from which to do backwards chaining, but also on whether the derivation of that explanation provides information showing flaws in the understander's prior reasoning such as erroneous or overlooked beliefs.

In case-based explanation construction, it is possible to focus on candidate explanations that not only explain the right aspect of the event but also relate that aspect to likely flaws in prior system beliefs: stored explanations can be indexed by the *anomalies* they resolve (rather than only the surprising aspects of events), and only candidate explanations relevant to those anomalies need be considered. In [Leake, 1991b; Leake, 1992] we describe an indexing vocabulary for organizing a memory of explanations to facilitate this focused explanation retrieval.

3.5 The role of overarching goals

The previous section shows that explanation must reflect system information needs in order to generate explanations that are useful for resolving anomalies. In addition to being driven by anomalies, however, everyday explanation is shaped by overarching goals. For example, a bank security officer may wish to explain the break-in in order to make future robberies less likely to succeed. Achieving that goal requires an explanation of the breakdown in security. Alternatively, the robber's parents may wish to understand the break-in in order to better understand the robber's character, which requires an explanation of his motives. "John needed money to pay back a loan shark" provides information useful for understanding why he committed a robbery—that he was under duress—while "The bank's security camera had been removed for repairs" does not. However, the second explanation is a better explanation for protecting the ATM in the future.

Psychological evidence shows that people favor different explanations for an event if they have different goals. For example, subjects attempting to absolve themselves of blame will favor different explanations from those without that goal [Snyder *et al.*, 1983]. In general, in any multi-task system the only way to assure useful explanations is to explicitly evaluate their goodness according to current system goals [Leake, 1991a; Leake, 1992].

Although goal-based explanation selection is important regardless of the means used for explanation construction, the importance of goals beyond facilitating routine understanding has received little attention in research on other models of abductive understanding. However, the influence of overarching goals has begun to be investigated in research on abductive diagnosis (e.g., for controlling large-scale diagnosis [Freitag and Friedrich, 1991] and for integrating diagnosis and response [Rymon *et al.*, 1991]). In addition, much research in explanation-based learning (EBL) has addressed the question of what constitutes a useful explanation.[3] However, although EBL has been applied to a wide range of tasks (such as object recognition, problem solving, and search control), within these tasks explanations have been used for a single purpose: forming rules for concept recognition. (See [Keller, 1988] for a discussion of how diverse EBL systems can be placed within the concept recognition framework.) Consequently, all these systems reflect the concept recognition task by making two basic assumptions about the form of explanations. First, they require explanations to be complete proofs showing sufficient conditions for concept membership. Second, because the types of rules used within the derivations are irrelevant to concept recognition—only the antecedents are important to that task—they treat all explanations with the same antecedents as equivalent.

Neither of those assumptions applies to the everyday explanation of events. First, complete explanations may not be necessary; partial explanations are sufficient for some tasks. For example, an explanation for preventing an undesirable event needs only to identify a single necessary condition for the event that is preventable in the future. (E.g., a driver who knows that his car sometimes fails to start when it has been parked in the cold can prevent the problem on cold days by putting it in the garage, even if he does not know the other factors relevant to whether it will start.) In fact, everyday explanations are *necessarily* partial explanations; it is impossible to provide a complete account of the factors that are sufficient for an event to occur. In everyday explanation, it is vital for the explainer to be able to make a principled decision about which partial explanations to accept and to benefit from the information in those explanations even if more complete explanations are unavailable.

Second, the goodness of everyday explanations often depends on their internal structure and the types of rules that they use. For example, even if an explanation shows that a disease can be predicted with absolute certainty, based on a set of environmental factors, that explanation will be worthless for developing a vaccine for the disease unless the explanation shows *how* those environmental factors cause the disease.

In [Leake, 1991a; Leake, 1992] we discuss ten disparate sets of requirements for good explanations that arise from different uses for explanations. A subset of these requirements has been implemented in ACCEPTER's explanation evaluation process. By basing evaluation on these criteria that dynamically reflect explainer goals, the system can direct explanation construction towards satisfying current needs and can decide when to accept partial explanations.

[3] Although that research focuses on deductive explanation, its usefulness criteria are equally applicable within an abductive framework.

4 Conclusions

Despite the difficulties of abductive explanation in complex and imperfectly understood domains, human explainers have comparatively little difficulty controlling the search for candidate explanations. Case-based explanation is an explanation method modeled on human explanation that focuses the construction and selection of abductive hypotheses. It differs from standard models of explanation construction in building explanations by adapting specific explanations for prior experiences; by treating explanations as plausible reasoning, rather than as proofs; and by dynamically focusing explanation construction and evaluation to reflect goal-based explainer information needs. By combining reasoning based on specific experiences with focusing based on a theory of information needs and how they can be satisfied, case-based explanation provides guidance for explanation construction and facilitates generation of good explanations in domains that are complex and imperfectly understood.

References

[Ajjanagadde, 1991] V. Ajjanagadde. Incorporating background knowledge and structured explananda in abductive reasoning: A framework. Wilhelm Schickard-Institute für Informatik WSI 91-6, Universität Tübingen, 1991.

[Charniak and Goldman, 1991] E. Charniak and R. Goldman. A probabilistic model of plan recognition. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 160–165, Anaheim, CA, July 1991. AAAI.

[Charniak, 1986] E. Charniak. A neat theory of marker passing. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, pages 584–588, Philadelphia, PA, August 1986. AAAI.

[deKleer and Williams, 1989] J. deKleer and B. Williams. Diagnosis with behavioral modes. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 1324–1330, Detroit, MI, August 1989. IJCAI.

[Dietterich and Flann, 1988] T. Dietterich and N. Flann. An inductive approach to solving the imperfect theory problem. In *Proceedings of the 1988 AAAI Spring Symposium on Explanation-based Learning*, Stanford, CA., 1988. AAAI.

[Freitag and Friedrich, 1991] H. Freitag and G. Friedrich. Goal-driven structural focusing in model-based diagnosis. In *Working Notes of the Second International Workshop on Principles of Diagnosis*, Milano, Italy, 1991.

[Hobbs *et al.*, 1990] J. Hobbs, M. Stickel, D. Appelt, and P. Martin. Interpretation as abduction. Technical Report 499, SRI International, 1990.

[Kahneman *et al.*, 1982] D. Kahneman, P. Slovic, and A. Tversky. *Judgement under uncertainty: Heuristics and biases*. Cambridge University Press, Cambridge, 1982.

[Kass and Leake, 1988] A. Kass and D. Leake. Case-based reasoning applied to constructing explanations. In J. Kolodner, editor, *Proceedings of the Case-Based Reasoning Workshop*, pages 190–208, Palo Alto, 1988. DARPA, Morgan Kaufmann, Inc.

[Kass, 1990] A. Kass. *Developing Creative Hypotheses by Adapting Explanations*. PhD thesis, Yale University, 1990. Northwestern University Institute for the Learning Sciences, Technical Report 6.

[Kautz and Allen, 1986] H. Kautz and J. Allen. Generalized plan recognition. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, pages 32–37, Philadelphia, PA, August 1986. AAAI.

[Keller, 1988] R. Keller. Defining operationality for explanation-based learning. *Artificial Intelligence*, 35(2):227–241, 1988.

[Leake, 1991a] D. Leake. Goal-based explanation evaluation. *Cognitive Science*, 15(4):509–545, 1991.

[Leake, 1991b] D. Leake. An indexing vocabulary for case-based explanation. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 10–15, Anaheim, CA, July 1991. AAAI.

[Leake, 1992] D. Leake. *Evaluating Explanations: A Content Theory*. Lawrence Erlbaum, Hillsdale, NJ, 1992.

[Mooney, 1990] R. Mooney. *A General Explanation-based Learning Mechanism and its Application to Narrative Understanding*. Morgan Kaufmann Publishers, Inc., San Mateo, 1990.

[Ng and Mooney, 1990] H. Ng and R. Mooney. On the role of coherence in abductive explanation. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pages 337–342, Boston, MA, July 1990. AAAI.

[Pearl, 1988] J. Pearl. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann Publishers, Inc., San Mateo, 1988.

[Rajamoney, 1988] S. Rajamoney. Experimentation-based theory revision. In *Proceedings of the 1988 AAAI Spring Symposium on Explanation-based Learning*. AAAI, 1988.

[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*, pages 867–871, Chicago, IL, August 1991. Lawrence Erlbaum.

[Read and Cesa, 1991] S. Read and I. Cesa. This reminds me of the time when ...: Expectation failures in reminding and explanation. *Journal of Experimental Social Psychology*, 27:1–25, 1991.

[Rymon *et al.*, 1991] R. Rymon, B. Webber, and J. Clarke. Towards goal-directed diagnosis. In *Working Notes of the Second International Workshop on Principles of Diagnosis*, Milano, Italy, 1991.

[Schank and Leake, 1989] R.C. Schank and D. Leake. Creativity and learning in a case-based explainer. *Artificial Intelligence*, 40(1-3):353–385, 1989.

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

[Snyder *et al.*, 1983] C. Snyder, R. Higgens, and R. Stucky. *Excuses: Masquerades in Search of Grace*. Wiley, New York, 1983.