Focusing on the most important explanations: Decision-theoretic Horn abduction

Paul O'Rorke
Department of Information and Computer Science
University of California, Irvine, CA 92717–3425
ororke@ics.uci.edu

Abstract

This paper describes a new method, called Decision-Theoretic Horn Abduction (DTHA), for generating and focusing on the most important explanations. A procedure is given that can be used iteratively to generate a sequence of explanations from the most to the least important. The new method considers both the likelihood and utility of partial explanations and is applicable to a wide range of tasks. This paper shows how it applies to an important engineering design task, namely Failure Modes and Effects Analysis (FMEA). A concrete example illustrates the advantages of the general approach in the context of FMEA.

Introduction

Abduction, the process of finding and evaluating explanations, is important in a number of areas of AI, including diagnosis and natural language understanding. One of the difficulties associated with abduction is that there are far too many explanations and it is difficult to focus on the best ones. The definition of "best" and of methods for comparing and evaluating explanations is also difficult. Many of the most advanced abduction methods address these problems by focusing on and preferring the most likely explanations (taking a Bayesian probabilistic approach to gauging likelihood). Poole (1992, 1993) describes a general approach called *Probabilistic Horn Abduction* (PHA) and shows how it can be applied to tasks such as diagnosis.

Probabilistic approaches represent an advance extending and improving previous methods. However, further improvement is needed because *sometimes the most likely explanations are not the most important ones.* For an example, consider a diagnostic situation involving symptoms that usually indicate a benign condition. But assume that sometimes these symptoms indicate a malignant and frequently fatal disorder. It may well be desirable to focus attention first on the diagnosis that corresponds to the potentially deadly problem even if it is far less likely.

The method described in this paper — Decision-Theoretic Horn Abduction (DTHA) — focuses on the most important explanations, not just the most likely ones. Importance is measured using decision-theory, which extends probability theory by combining numerical measures of likelihood or probability with measures of value or utility. The DTHA method extends Poole's (1992) Probabilistic Horn Abduction (PHA) procedure for finding maximally likely explanations of individual observations. DTHA considers multiple observations (a finite number of "outcomes" in the terminology of decision-theory) with differing importance (as indicated by given numerical "utility" scores). DTHA computes the importance of alternative explanations by multiplying the utility of the outcomes by the products of the prior probabilities of the assumptions underlying the corresponding explanations. Extreme values (whether maximal or minimal) are considered to be more important. DTHA focuses on the most important explanation: it will work on unlikely explanations if the outcomes are sufficiently valuable and it will pursue explanations of less valuable outcomes if they are sufficiently likely.

Decision-Theoretic Horn Abduction

A version of Decision-Theoretic Horn Abduction based closely on Poole's (1992) PHA procedure is given in table 1. The inputs are: a finite list of "assumables" $a_1, \ldots, a_m$ and corresponding probabilities $p_1, \ldots, p_m$; a finite list of inconsistent assumptions $nogood(a_i, a_j)$, where $i, j \in \{1, \ldots, m\}$; a Probabilistic Horn Abduction theory; and a finite set of outcomes $o_1, \ldots, o_n$ and corresponding utilities $u_1, \ldots, u_n$. The output is (one of) the most important explanation-outcome pair(s).

The notation used in table 1 is to be understood as follows. For each value of i from one to n, $\mathcal{D}_i$ contains done and $\mathcal{P}_i$ contains partial explanations of outcome o_i. Partial explanations have the form $\langle A, p, C \rangle$ where A contains the assumptions made so far, p is the probability of the partial explanation and C is a conjunction of conditions that remain to be explained. The probability p of the partial explanation is the prior probability of the set of assumptions A. Note that it is important to distinguish the probabilities associated with sets of assumptions and partial explanations

Table 1: A Version of Decision-Theoretic Horn Abduction (DTHA)

Initialization: for $i = 1, \ldots, n$ set $\mathcal{D}_i := \emptyset$ and set $\mathcal{P}_i := \{\langle \emptyset, 1.0, o_i \rangle\}$

Procedure: Do while $\exists i$ such that $1 \leq i \leq n \wedge \mathcal{P}_i \neq \emptyset$

1. let $i \in \{1, \ldots, n\}$ and the corresponding $\langle A, \hat{p}_i, C \rangle \in \mathcal{P}_i$ be such that

$$\hat{p}_i \times u_i = \max_{j=1,\ldots,n} \hat{p}_j \times u_j$$

2. set $\mathcal{P}_i := \mathcal{P}_i \setminus \{\langle A, \hat{p}_i, C \rangle\}$
3. if $C = true$

(a) then if $ok(A)$ then set $\mathcal{D}_i := \mathcal{D}_i \cup \{A\}$, output A, and halt

(b) else let $C = a \wedge R$

i. for each rule $h \leftarrow B$ such that $mgu(a, h)$ exists, set $\sigma := mgu(a, h)$ and set $\mathcal{P}_i := \mathcal{P}_i \cup \{\langle A, \hat{p}_i, B \wedge R \rangle \sigma\}$

ii. if $\exists j \in \{1, \ldots, m\}$ such that $a = a_j$ and $ok(A \cup \{a_j\})$ then set $\mathcal{P}_i := \mathcal{P}_i \cup \{\langle A \cup \{a_j\}, \hat{p}_i \times p_j, R \rangle\}$

from the probabilities associated with individual assumptions, although they are related. The connection is that $p = \prod_{a_i \in A} p_i$ under the independence assumptions associated with PHA. The meaning of the remaining notation is as follows. For $i = 1, \ldots, n$ let $\hat{p}_i = \max_{\langle A,p,C \rangle \in \mathcal{P}_i} p$. At the beginning of execution of the procedure, each $\hat{p}_i$ represents the probability of one of the most likely partial explanations of the outcome o_i. This explanation, denoted $\langle A, \hat{p}_i, C \rangle$, is at least as likely as any other explanation of the i-th outcome in $\mathcal{P}_i$.

The initialization step of the procedure sets all the groups of *done* explanations to the empty set $\emptyset$. The sets of partial explanations $\mathcal{P}_i$ are initialized to singletons containing the explanations $\langle \emptyset, 1.0, o_i \rangle$. This is because the initial goal is to explain the outcome, the initial estimate bounding the probability of the goal from above is one, and no assumptions have been made yet in pursuit of this goal. The body of the procedure is as follows.

Step 1 selects the most important partial explanation to work on (if there is a non-empty set of partial explanations). The importance of a partial explanation $\langle A, p, C \rangle$ of outcome o_i is defined here as $p \times u_i$. This measure of importance is desirable because it is larger for more likely explanations and for outcomes with larger utilities. Step 1 finds a particular value of i and a partial explanation for outcome o_i with maximal importance over *all* known partial explanations — including explanations of other outcomes. Keep in mind that the importance and probability associated with a partial explanation are an upper bound: additional assumptions are often needed to complete the explanation and these assumptions reduce the probability of the original partial explanation (see Step 3(b)ii). For this reason, the "most important" partial explanation at one stage may produce explanations that are less important later on. Step 1 and the following steps ensure that DTHA "loses interest" in such explanations once they become less important than some alternative candidate.

The remaining steps execute a round of Probabilistic Horn Abduction focused on $\langle A, \hat{p}_i, C \rangle$ and o_i. Step 2 deletes the partial explanation that is about to be processed. If the explanation is complete, step 3a checks its assumptions to see whether they are acceptable. If so, it records the explanation, communicates it to the user, and halts processing (at least for now). The acceptability test *ok* has two components: 1) a check *new* that makes certain that the same explanation of the same outcome has not been seen before and 2) a limited consistency check *consistent* that ensures that no pair of assumptions in the explanation is an instance of a *nogood*. (To be specific, $ok(A, i) \equiv new(A, i) \wedge consistent(A)$ where $new(A, i) \equiv [\nexists D \in \mathcal{D}_i [D \subseteq A]]$ and $consistent(A) \equiv \forall \{a_j, a_k\} \subseteq A [\forall x, y [nogood(x, y) \rightarrow \nexists \theta [\{a_j, a_k\} = \{x, y\}\theta]]]$.) If the partial explanation is not complete, there is more work to be done. At least one condition (a) remains to be proven, so the following steps focus on it. Backward chaining occurs in step 3(b)i. If there is a rule (with head h and body B) that concludes the condition that is desired, that condition is deleted and the conditions of the rule are added in its place to the remaining conditions R. Assumptions are made in step 3(b)ii. If the condition to be explained is assumable, it is deleted from the conditions to be explained and added to the assumptions supporting the explanation. The probability of the partial explanation is reduced by multiplying it by the prior probability of the new assumption.

The backward chaining and assumption steps are mutually exclusive under the assumptions made in Probabilistic Horn Abduction. Once one or the other step is executed, the procedure is repeated. Then, the selection step (1) may shift attention elsewhere.

The procedure given in table 1 halts when it arrives at a single explanation of a single outcome. However, if the main body — **procedure** — is called again, it will deliver the next most important explanation, and so on, until the partial explanations are exhausted for all the outcomes.

A variant of this procedure suitable for finding the most costly, most likely potential outcomes can be had by representing costs as negative utilities and taking min instead of max in step 1. This version will be discussed below. Equivalently, the absolute value can be taken before taking the max or the negative utilities can all be made positive prior to input.

The version of DTHA described in table 1 is suitable for abductive planning (Elkan, 1990; Krivičić & Bratko, 1993). In this application, the goal is to find a set of assumptions and a chain of inferences spec-

ifying a course of action leading to an outcome that maximizes expected utility. It is also possible to apply a variant of DTHA to consider potentially harmful outcomes. The goal then is ultimately to minimize distress. This is done by identifying the most likely and most harmful outcomes so that something can be done to avert or avoid them. Failure Modes and Effects Analysis (FMEA) is a good example of this type of application. FMEA is used to illustrate the method in the following section.

A Method for Failure Modes and Effects Analysis

This section presents a novel abductive approach to Failure Modes and Effects Analysis (FMEA). This form of reliability analysis and risk assessment is often required by governments of their contractors and in turn it is often required by these companies of their subcontractors. FMEA is extensively used in the aerospace and automotive industries worldwide. The stated purposes of FMEA are to determine the consequences of all possible failure modes of all components of a system, to assess risks associated with failures, and to recommend actions intended to reduce risks. (Henley & Kumamoto, 1991)

In FMEA, pairs of failure modes and outcomes are prioritized using so-called "risk priority numbers (RPNs)." These numbers take into account three things: the undesirability of the effects caused by failures, their likelihood, and their detectability. This paper ignores the issue of detectability and considers only how the likelihoods and costs of failures can be used to prioritize the automatic construction of FMEAs.

The key ideas are: 1) to associate FMEAs with explanations of how undesirable effects can be caused by failures and 2) to order the generation of FMEAs so that the most likely and most undesirable outcomes are considered first. The first idea is implemented by specifying a model and/or theory of the system undergoing FMEA, including the following: 1) normal and failure modes and their associated prior probabilities 2) outcomes and their utilities; and 3) rules governing causal connections between the failure modes and outcomes. The second idea is implemented by using DTHA and by considering a partial explanation P_1 to be more important than a partial explanation P_2 if the probability of P_1 times the cost of the associated outcome is larger than the probability of P_2 times the cost of its outcome.

A Model of an Automotive System

An automotive example is given in this section as an illustration of the ideas and as a demonstration of how to implement FMEA as a special case of DTHA. The example is loosely based on real events.

The first event involved engine fires in General Motors cars. The fires were caused by a fuel hose coming loose and leaking. This prompted a recall. A similar event reported in the Los Angeles Times (Nauss, 1993) involved engine fires in GM's new Saturn line caused by a short circuit in the generator leading to electrical overloads in a wiring harness. All Saturn sedans, coupes, and station wagons built between 1991 and 1993 were recalled. Automotive industry analysts estimated the direct cost of the recall as \$8–\$20 million. This example indicates that the costs that can be incurred when potential faults are not anticipated can be substantial.

Another event involved solder joints in the circuit connecting an electronic controller to a sensor. Most modern cars have electronic controllers that regulate ignition, fuel injection, and so on. Sensors placed in strategic parts of the engine provide the controller with the information needed to optimize the performance of the engine. (Schultz, Lees, & Heyn, 1990)

In pre-sales field testing by another manufacturer of an engine with an electronic controller and sensors, it was discovered that a wire connected to a sensor's housing could come loose due to faulty solder joints. Without the information provided by the sensor, the electronic controller cannot regulate the engine properly. This can result in an unacceptable increase in emissions (to a level above the maximum allowed by government regulations). The problem was caught before the cars were distributed to the public, so a costly recall was avoided. Computerized FMEA will enable us to anticipate similar problems prior to field testing during the final stages of design, thus avoiding the manufacture of trouble-prone systems.

The following example shows how DTHA can be applied to FMEA and illustrates the behavior of DTHA in the context of FMEA. Recall that this approach requires a model including normal and failure modes, probabilities, outcomes and their utilities, plus causal connections between failure modes and outcomes. The model provided in the example is sketched in figure 1. There are three subsystems modeled as chambers connected together by pipes. Fuel enters the fuel pump, then goes through the delivery pipe to the injection system. Next, fuel is injected into the combustion chamber through a nozzle. From the combustion chamber, the (burned) fuel goes through an exhaust pipe to the exhaust manifold. The model also has an electrical circuit. Two wires connect an electronic controller to a sensor (which might be in the injection system, the fuel pump, or elsewhere in the car, e.g., the tachometer). Connections are modeled explicitly as components: pipes are connected to chambers by pipe joints and wires are connected to other components by solder joints. The structural part of the model comprises the components and connections just described.

The model also provides information about normal and failure modes and how often they occur. All the basic components are either in a normal or a broken state.[1] Most of them are considered to be highly reli-

[1]Failure rates and states are assigned to basic but not

Figure 1: Sketch of an Automotive System

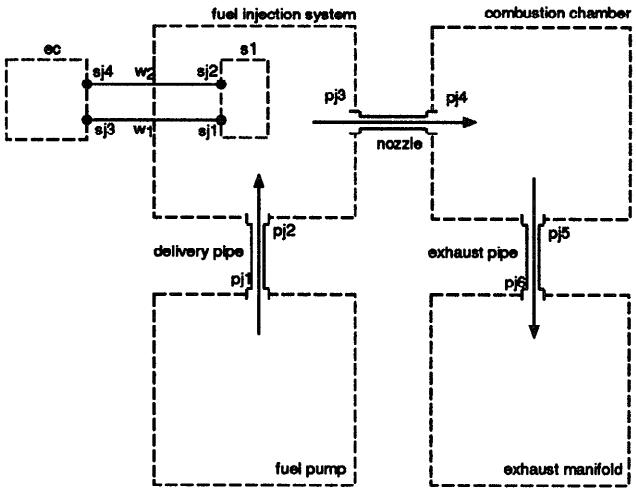

able: the prior probability of the broken state is only 0.0001 so the probability of normality is 0.9999.[2] The connections are less reliable: the prior probability of breakage for pipe joints and solder joints is 0.01. Sensors are the least reliable components: the prior probability that a sensor is broken is 0.05.

The outcomes of concern in this case are 1) there might be a fire in the engine compartment, 2) the exhaust emissions might be too high, and 3) the engine might run inefficiently. These outcomes are assigned costs of 10^6, 10^5, and 10^4 respectively.

The immediate causes of the outcomes are specified by rules. Two rules specify two independent causes of fires in the engine compartment due to fuel leaks in a pipe and a pipe joint near hot parts of the engine. Another rule states that the exhaust will be dirty if the electronic controller is uninformed. The final rule says that the engine will be less efficient if the electronic controller is uninformed.

Another set of rules specifies different aspects of the behaviors and functions of the components in the system. These rules provide the remaining connections linking the outcomes to their possible causes in terms of failure modes of the components. Two rules specify normal and failure modes of pipes and pipe joints. Normally, pipes connecting chambers cause the propagation of their contents. (To avoid cycles, flows are considered to be unidirectional.) When pipes or joints are broken, they leak. A fact states that the fuel pump supplies the fuel injection system with fuel. Additional rules specify that the electronic controller will be uninformed if a sensor is broken, or if the circuit connecting the electronic controller and the sensor is broken. A rule states that a circuit is broken if there is a component in the circuit that is broken. The rules for failure modes of solder joints say that when a connection between two wires is broken, the voltage on the wires goes to zero.

FMEA for the Automotive System

This section describes the behavior of the Decision-Theoretic Horn Abduction algorithm when it is invoked repeatedly given the model of the previous section as input. The following text summarizes a trace produced by an implementation of DTHA in PROLOG. The three outcomes are labelled in descending order on their costs as F, H, and L for engine fire, high emissions, and low efficiency respectively. Let these labels be variables that stand for the weighted cost of the most likely cause of the corresponding outcome. The variables are initialized with the outcome's costs since they might be unconditionally true. The labels are intended to be mnemonic and their alphabetic order reflects their initial numerical order: $F > H > L$. As the analysis proceeds, the variables will be updated and their order will change.

$\boxed{F = 10^6 > H = 10^5 > L = 10^4}$ — First, an attempt is made to explain the "engine fire" outcome without making any assumptions. The attempt fails but two partial explanations involving assumptions are added to the set of partial explanations for this outcome: one corresponds to a broken pipe and the other to a broken pipe joint. The pipe joint is considered to be relatively unreliable ($p_{pj} = 10^{-2}$) while the pipe is considered to be one of the more reliable ($p_p = 10^{-4}$) components. So the leading possibility is that the pipe joint will leak and cause an engine fire. This has a weighted cost of $10^{-2} \times 10^6 = 10^4$.

$\boxed{H = 10^5 > F = 10^4 = L}$ — Now "high emissions" has a higher weighted cost than the most likely possible cause of engine fires so it is pursued. Hypotheses about a faulty sensor ($p_s = 5 \times 10^{-2}$) and solder joints ($p_{sj} = 10^{-2}$) and wires ($p_w = 10^{-4}$) are added to the set of partial explanations for high emissions. The most likely hypothesis is that a sensor will be faulty. The corresponding weighted cost is 5×10^3.

$\boxed{F = 10^4 = L > H = 5 \times 10^3}$ — The focus returns to engine fires.[3] A possible cause, that the second pipe joint will break, is found and printed out. The next most likely explanations (involving wires) have weighted costs of 10^2.

$\boxed{L = 10^4 > H = 5 \times 10^3 > F = 10^2}$ — Next, low efficiency becomes the most important outcome. The same conditions that can contribute to high emissions can contribute to low efficiency so they are added to the set of partial explanations for low efficiency as well.

to complex components.

[2]The states of components are considered to be mutually exclusive and exhaustive.

[3]Low efficiency could have been chosen at this point since it has the same weighted cost.

Table 2: DTHA-FMEA on the Automotive Example

Causes & Consequences	Priorities
F	$> H > L$
H	$> F = L$
pipe joint 2$\rightarrow F$	$= L > H$
L	$> H > F$
sensor 1$\rightarrow H$	$> L > F$
solder joints 1,2,3, and 4$\rightarrow H$	$> L > F$
sensor 1$\rightarrow L$	$> F > H$
solder joints 1,2,3, and 4$\rightarrow L$	$= F > H$
delivery pipe$\rightarrow F$	$> H > L$
wires 1 and 2$\rightarrow H$	$> L$
wires 1 and 2$\rightarrow L$	

Again, the most likely is sensor failure ($P = .05$). So the weighted cost associated with the most likely cause of low efficiency is now 5×10^2.

$\boxed{H = 5 \times 10^3 > L = 5 \times 10^2 > F = 10^2}$ — Now the most important outcome is high emissions and the sensor is the most likely possible cause of this problem. This fact is reported to the user.

This process will continue as long as the outputs are sufficiently important to the user. In this example, a complete list of cause-consequence pairs is possible and is shown in table 2. The behavior of DTHA on the example is also summarized in table 2. The cause-consequence pairs are shown in order of generation in the first column. The consequences together with the second column shows the priorities of the system at each step. Failure Modes and Effects Analyses were done for the top priority outcome at each step resulting in the corresponding cause-consequence pairs.

Although the sensor is five times more likely to break than the pipe joint, the joint is chosen for the first FMEA because the outcome engine fire is ten times more costly than high emissions which is the most costly consequence of the sensor failing. Next, the possible causes of failure of the sensor circuit are enumerated. All of these are considered before the remaining cause of the most costly outcome because that cause (failure of the delivery pipe) is so unlikely. Finally, the most reliable components whose failure could lead to the less important outcomes are considered in turn. Ignoring the initial consideration of each outcome, the pattern in the example was F then H then L followed by a return to F then H then L.

Relation to Previous Work

Work on abduction in AI dates back to Pople (1973). There are many different approaches to explanation construction and evaluation, including case-based (Leake, 1992) and connectionist (Thagard, 1989) methods that address many of the same issues. The abduction method described in the present paper is a descendent of logic-based methods for generating explanations embodied in Theorist (Poole, Goebel, & Aleliunas, 1987) and Tacitus (Hobbs, Stickel, Martin, & Edwards, 1988). Recently, logical and symbolic approaches to abduction have been extended by adding probability and the most probable explanations have been considered to be the best ones. For example, Peng and Reggia (1990) adapt a probabilistic approach to diagnosis viewing it as a special case of abduction. Charniak and Goldman (1993) take a probabilistic approach to plan recognition, a special case of abduction that occurs in natural language processing. Poole (1992, 1993) describes a general integration of logical and probabilistic approaches to explanatory reasoning called *Probabilistic Horn Abduction* (PHA). A major weakness of these methods is that they do not take utilities or values into consideration when they evaluate and search for explanations. It is important to do so in many practical situations, for example in abductive planning and in considering failures that might cause undesirable outcomes (as in FMEA).

The abuctive approach to FMEA differs from existing AI approaches to FMEA described in (Hunt, Price, & Lee, 1993; Price, Hunt, Lee, & Ormsby, 1992) — largely due to the difference between postdiction and prediction. The abductive approach is an example of postdiction since it infers possible causes or reasons for given effects. Previous approaches use prediction to infer possible effects from given causes. Some use a simple forward-chaining simulator for prediction. More sophisticated qualitative and model-based reasoning is used for prediction in the approaches cited above. The main advantage of predictive (especially model-based) approaches to FMEA is that they can generate consequences that were not anticipated in advance. The main advantage of the DTHA approach is that it automatically focuses on the most important FMEAs first, although all FMEAs can be generated if required. Another advantage is that it works for multiple faults.

Limitations; Future Work

Although it appears to be adequate for significant practical tasks such as FMEA, the method described here is limited due due to the assumptions employed. For example, the model of outcomes and utilities is extremely simple: it is assumed that outcomes and utilities are given by the user. This is reasonable in the context of FMEA, since designers typically have a finite and small number of outcomes they are concerned about and they usually know the utilities of these undesirable effects of component failures. But in more complex situations, methods for acquiring and calculating utilities from users and from more basic information will be needed. Work on these issues in the new field of Decision-Theoretic Planning seems relevant.

Conclusion

This paper provides a new method that generates explanations focusing on the most important ones. Im-

portance is measured taking into account utilities or values in addition to probabilities so the method is called Decision-Theoretic Horn Abduction (DTHA). The addition of utilities is an important improvement over existing probabilistic approaches because in many situations the most likely explanations are not the best ones to focus on or generate. For example, in diagnosis, the most dangerous disease that explains the symptoms should often be considered before more common but less dangerous disorders. In abductive planning, in determining the best explanations of how to achieve a goal one should often take into consideration the value of the goal relative to alternatives and the costs of the actions involved in the plans in addition to the likelihood of success.

The paper provides an example illustrating the general DTHA method in the context of a task, Failure Modes Effects Analysis (FMEA), that involves utilities in addition to probabilities. In this context, the explanations correspond to assumptions about various components being in failure or normal modes and these cause outcomes that correspond to costly effects that the designers wish to avoid. The example demonstrates that the method is capable of keeping priorities straight in deciding which explanation-outcome pairs to generate. The method shifts the focus of attention: starting on one outcome, moving to another, and returning to an earlier focus. When the more costly or valuable outcome is at least as probable, or not too improbable, it is pursued. On the other hand, the method focuses attention on the most important explanations and outcomes even when this requires switching attention from more to less costly or from more to less probable outcomes.

Acknowledgments

The author is grateful to the reviewers; to David Poole for conversations and for Prolog code implementing PHA; and to Chris Price, Andy Ormsby, and David Pugh for conversations and for a demo of their FMEA system FLAME II. The author was a consultant with FAW (the Research Institute for Applied Knowledge Processing) during the research and the development of the prototype described here. Rüdiger Wirth, FAW FMEA group leader, provided assistance and support.

References

Charniak, E., & Goldman, R. P. (1993). A Bayesian model of plan recognition. *Artificial Intelligence, 64*, 53–79.

Elkan, C. (1990). Incremental, approximate planning. *The Eighth National Conference on Artificial Intelligence* (pp. 145–150). San Mateo, CA: AAAI Press/The MIT Press.

Henley, E. J., & Kumamoto, H. (1991). *Probabilistic risk assessment: Reliability engineering, design, and analysis.* Piscataway, NJ: IEEE Press.

Hobbs, J. R., Stickel, M., Martin, P., & Edwards, D. (1988). Interpretation as abduction. *Proceedings of the Twenty Sixth Annual Meeting of the Association for Computational Linguistics* (pp. 95–103). Buffalo, NY: The Association for Computational Linguistics.

Hunt, J. E., Price, C. J., & Lee, M. H. (1993). Automating the FMEA process. *Intelligent Systems Engineering, 2*(2), 119–132.

Krivičić, K., & Bratko, I. (1993). Abductive planning with qualitative models of dynamic systems. In N. Piera-Carreté, & M. G. Singh (Eds.), *Qualitative Reasoning and Decision Technologies* (pp. 389–395). Barcelona: CIMNE.

Leake, D. B. (1992). *Evaluating explanations: A Content theory.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Nauss, D. W. (1993). Nearly all Saturns recalled over potential engine fires. *Los Angeles Times (Orange County Edition).* Los Angeles, CA: A1 and A19.

Peng, Y., & Reggia, J. A. (1990). *Abductive inference methods for diagnostic problem solving.* New York: Springer-Verlag.

Poole, D. (1992). Logic programming, abduction and probability. *Proceedings of the International Conference on Fifth Generation Computer Systems (FGCS-92)* (pp. 530–538). Tokyo.

Poole, D. (1993). Probabilistic Horn abduction and Bayesian networks. *Artificial Intelligence, 64*, 81–129.

Poole, D. L., Goebel, R., & Aleliunas, R. (1987). Theorist: A logical reasoning system for defaults and diagnosis. In N. Cercone, & G. McCalla (Eds.), *The Knowledge Frontier: Essays in the Representation of Knowledge.* New York: Springer-Verlag.

Pople, H. E. (1973). On the mechanization of abductive logic. *Proceedings of the Third International Joint Conference on Artificial Intelligence* (pp. 147–152). San Mateo, CA: Morgan Kaufmann.

Price, C. J., Hunt, J. E., Lee, M. H., & Ormsby, A. R. T. (1992). A model-based approach to the automation of failure mode effects analysis for design. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering, 206*, 206–291.

Schultz, M., Lees, A. W., & Heyn, E. V. (Eds.). (1990). *What's wrong with my car? A guide to troubleshooting common mechanical and performance problems.* Mount Vernon, New York: Consumers Union.

Thagard, P. (1989). Explanatory coherence. *The Behavioral and Brain Sciences, 12*(3), 435–502.