
Proceedings

**High Integrity Software
Conference**

October 15-16, 1997
Albuquerque, New Mexico

Sponsored by
The High Integrity Software Program
and
The IEEE Computer Society



Logic-Based and Heuristic Models for Reasoning with Uncertain or Incomplete Information

by

George F. Luger
University of New Mexico
Albuquerque NM 87131
luger@cs.unm.edu

&

Carl R. Stern
Vista Control Systems
Los Alamos NM 87544
stern@cs.unm.edu

All traditional logic habitually assumes that precise symbols are being employed. It is therefore not applicable to this terrestrial life but only to an imagined celestial existence.

—BERTRAND RUSSELL

It is the mark of an instructed mind to rest satisfied with that degree of precision which the nature of the subject admits, and not to seek exactness where only an approximation of the truth is possible.

—ARISTOTLE

1. Introduction

Human experts often feel in practicing their domains of expertise that they follow the model of reasoning used in formal logic: from correct premises, sound inference rules produce new, guaranteed correct conclusions. On reflection however, we realize there are many situations that will not fit this approach; that is, we must draw useful conclusions from poorly formed and uncertain evidence using unsound inference rules.

Drawing useful conclusions from incomplete and imprecise data with unsound reasoning is not an impossible task; we do it very successfully in almost every aspect of our daily life. We deliver correct medical diagnoses and recommend treatment from ambiguous symptoms; we analyze problems with our cars or stereos; we comprehend language statements that are often ambiguous or incomplete; and we successfully navigate the stock and money markets.

To demonstrate the problem of reasoning under uncertainty, consider a simple rule for diagnosing problems with an automobile:

```
if
  the engine does not turn over, and
  the lights do not come on
then
  the problem is battery or cables.
```

On the surface, this rule looks like a normal predicate relation to be used in sound inferencing (modus ponens). However, it is not; it is heuristic in nature. It could be possible, though very unlikely, that the battery and cables are fine but that the car simply has a bad starter motor and burned-out headlights. Failure of the engine to turn over and the lights to come on does not necessarily imply that the battery and cables are bad. It is interesting that the converse of the rule is true:

```
if
  the problem is battery or cables
then
  the engine does not turn over, and
  the lights do not come on.
```

Barring the supernatural, with a dead battery, neither the lights nor the starter will work!

Our expert rule offers an example of abductive reasoning. More formally, abduction states that from $P \Rightarrow Q$ and Q , under the right conditions, it is possible to infer P . Abduction is an unsound rule of inference, meaning that the

conclusion is not necessarily true for every interpretation in which the premises are true. The conditions under which abductive inferences are warranted requires considerable analysis (Stern 1996). In a knowledge-based system, we often attach a certainty factor to the rule to measure our confidence in its conclusion. For example, the rule, $P \text{ P } Q$ (.9), expresses the belief "If you believe P to be true, then you believe Q will happen 90% of the time." Thus, heuristic rules can express an explicit policy for belief.

Another constraining problem for expert reasoning is that useful results must be drawn from data sets with missing, incomplete, or incorrect information. We may also use certainty factors to reflect our belief in the quality of the data, for example, the lights do come on (.2) can indicate that the headlights do come on, but are weak and barely visible. Finally, beliefs and imperfect data are propagated through rule sets. This paper shows how heuristic rules can be combined to extend beliefs.

Although abduction is unsound, it is often essential to solving problems. The "logically correct" version of the battery rule is not very useful in diagnosing car troubles since its premise, determining whether or not the battery is bad, is our goal and its conclusions are the observable symptoms with which we must work. The rule must be used in an abductive fashion, as are rules in most diagnostic situations. Faults or diseases cause (imply) symptoms, not the other way around; but diagnosis must work from symptoms back to their causes.

In this paper, we discuss two of many possible ways of managing abductive inference and uncertainty, especially as it is required for knowledge-intensive problem solving. First, in Section 2, we consider the formalisms of logic to see how they might be extended to capture the constraints of abductive inference. Next, in Section 3, we consider causal networks, a heuristic methodology for reasoning on multiple levels: first on the data level where information about a situation is gathered and examined. On a second level, we reason about how pieces of information are part of sets of symptoms or causal patterns. Finally, on the highest level, we are able to organize potential explanations of discovered data sets. Causal networks have been developed in the domain of medicine, and our example is taken from the literature of that area. They may be applied, however, to any sufficiently rich area of diagnosis.

2 Set Cover and Clause-based Abduction

As noted in the introduction, in abductive reasoning, we have rules of the form $P \Rightarrow Q$, along with a reasonable belief in Q . We wish then to make a case for the truth of predicate P . Abductive reasoning is not sound, but what is often called reasoning to the best explanation for the presence of the data Q . In this section, we look more closely at the generation of explanations in domains of abductive inference.

The set cover approach defines an abductive explanation as a covering of the set of actual observations by a binary relation expressing causal associations. Reggia et al. (1983) set cover approach assumes that causality can be expressed through a simple causal relation R where R is a subset of $\{\text{Causes} \times \text{Observations}\}$. Given a set of observations $S2$, Reggia's algorithm searches for a minimal set covers, i.e., sets of causes that comprise a minimal cover of $S2$ using the causal relation R . The weakness of this approach is that it reduces explanation to a simple list of causes. In situations where there are interrelated or interacting causes or where an understanding of the structure or sequencing of causal interactions is required, the set cover model is inadequate.

Clause-based approaches to abduction on the other hand, rest on a more sophisticated notion of explanation. Levesque (1989) defines an abductive explanation of some previously unexplained set of observations O as a minimal set of hypotheses H consistent with an agent's background knowledge K that entails O . More formally:

$\text{abduce}(K, O) = H$, if and only if:

1. K does not entail O
2. $H \cup K$ entails O
3. $H \cup K$ is consistent, and
4. No subset of H has properties 1, 2, and 3.

Note that in general many sets of hypotheses may exist; that is, there may be many potential abductive sets of explanations for a given set of observations O .

The clause-based definition of abductive explanation suggests a corresponding mechanism for explanation discovery in the context of a knowledge-based system. If the explanatory hypotheses must entail the observations O , then the way to construct a complete explanation is to reason backwards from O . One approach would be to start from the conjunctive components of O and reason back from consequents to antecedents.

This “backchaining” approach also seems natural because the conditionals which support the backchaining can readily be thought of as causal laws, thus capturing the pivotal role which causal knowledge plays in the construction of explanations. The model is also convenient because it fits nicely to something with which the AI community already has experience: backchaining and computational models for deduction.

There are also clever ways of finding the complete set of abductive explanations. Assumption-based truth-maintenance systems ATMS (deKleer 1986), contain an algorithm for computing minimal support sets, the set of (non-axiom) propositions that logically entail a given proposition in a theory. To find all possible abductive explanations for a set of observations, we merely take the Cartesian product over the support sets, pruning as necessary inconsistent conjunctions of hypotheses.

As simple, precise, and convenient as the clause-based account of abduction is, there are two related shortcomings: high computational complexity and semantic weakness. Selman and Levesque (1990) found the complexity of abduction tasks similar to that involved in computing support sets for an ATMS. The standard proof that the ATMS problem is NP-hard depends on the existence of problem instances with an exponential number of solutions. Selman and Levesque avoid the number of potential solutions complexity issue by asking whether finding a smaller set of solutions is also NP-hard. Given a Horn clause knowledge base (Luger and Stubblefield 1998, chapter 12), Selman and Levesque produce an algorithm that finds a single explanation in order $O(kn)$ where k indicates the number of propositional variables and n the number of occurrences of literals. However, when restrictions are placed on the kinds of explanations sought, the problem again becomes NP-hard, even for Horn clauses.

One interesting result from the Selman and Levesque (1990) analysis is the fact that adding certain kinds of goals or restrictions to the abduction task actually makes computation significantly harder. From the naive viewpoint of the human problem solver, this added complexity is surprising. Human problem solvers assume that the addition of further constraints to the search for relevant explanations makes the task easier. The reason the abduction task is harder in the clause-based model is that it only contributes additional clauses to be processed in the problem solving, not additional structure to the activity of problem solving.

Explanation discovery in the clause-based model is characterized as the task of finding a set of hypotheses with certain logical properties. These properties, including consistency with the background knowledge and entailment of what is to be explained, are meant to capture the necessary conditions of explanations: the minimal conditions which a set of explanatory hypotheses must satisfy in order to count as an abductive explanation. Proponents of this approach believe that by adding additional constraints, the approach can be extended to provide a characterization of good or reasonable explanations.

One simple strategy for selecting good explanations is to define a set of fact clauses that are abducible, that is, from which candidate hypotheses must be chosen. This clause set allows search to be restricted in advance to those factors that can potentially play a causal role in the chosen domain. Another strategy is to add selection criteria for evaluating and choosing between explanations. Various selection criteria have been proposed, including decision procedures based on prior probabilities and Bayesian style conditional probabilities (Pearl 1988). Other approaches use simplicity and coherence (Ng and Mooney 1990) as preferential criteria.

Both simplicity and coherence can be seen as applications of Occam’s razor (Luger and Stubblefield 1998). Simplicity and coherence criteria are particularly appealing what is to be explained is not a simple proposition but rather a set of propositions. Ng and Mooney (1990) have argued that a coherence metric is superior to a simplicity metric for choosing explanations in the analysis of natural language text. They define coherence as a property of a proof graph where explanations with more connections between any pair of observations and fewer disjoint partitions are more coherent. The coherence criterion is based on the heuristic assumption that what we are asked to explain is a single event or action with multiple aspects. The justification for a coherence metric in natural language understanding is based on Gricean felicity conditions, that is, the speaker’s obligation to be coherent and pertinent (Grice 1975). It is not difficult to extend their argument to a variety of other situations. For example in diagnosis, the observations which comprise the initial set of things to be explained are brought together because they are believed to be related to the same underlying fault or failure mechanism.

Another mechanism for explanation selection, *cost-based abduction*, is also interesting because it takes into account both properties of the hypothesis set as well as properties of the proof procedure. Cost-based abduction places a cost on potential hypotheses as well as a cost on rules. The total cost of the explanation is computed on the basis of the total cost of the hypotheses plus the cost of the rules used to abduce the hypotheses. Competing hypothesis sets are then compared according to cost. One natural semantic that can be attached to this scheme is the probabilistic one (Charniak and Shimony 1990). Higher costs for hypotheses represent less likely events; higher costs for rules represent less probable causal mechanisms. Cost-based metrics can be combined with least-cost search algorithms, such as best-first search, considerably reducing the computational complexity of the task.

In Section 3, we consider architectures for reasoning that supports the manipulation of network-based descriptions of the world, a relaxation of the strictures of logic, but an explicit attempt to generate meaningful explanations.

3. Causal Networks

Causal models depict relationships as links between nodes in a graph or a network of nodes. These models are used quite extensively in a number of areas of reasoning, including diagnosis in medicine, the analysis of faults in electronic circuits, and story understanding. The approach in these applications is straightforward: map observations onto a network of nodes and then link the network nodes in a causally coherent pattern.

One of the first efforts to build an explicit model of causal relationships was for the diagnosis of various forms of the eye disease glaucoma. This program, CASNET (Weiss et al. 1977), was a kind of semantic network that represented a dynamic process occurring over time as a causal relationship among states. This network also related the nodes of the causal process to external manifestations; the observations, the evidence, and in this case medical classifications, that is, to diagnostic categories. More precisely, this representation has three connected levels: first, a level of pathophysiological states, second, a level of observations, and finally, a level of disease categories.

At the core of the model is the network of pathophysiological states connected by causal links. The notion of a "causal link" is interpreted loosely and not intended for exact Bayesian correlational analysis. The links connecting states are "weighted" with numerical confidence measures from 1, rarely causes, to 5, almost always causes, the pathophysiological states.

A complete causal pathway from a start node to a terminal node represents a complete disease process, while pathways that end in non-terminal nodes represent partial or incomplete evolution of a disease process. Confirmation of a state is derived either from associated observations, where the links between observations and pathophysiological states also have weights, or indirectly through the causal link to another state for which there is some evidence. Activation of the network proceeds by a weight-propagation algorithm (Weiss et al. 1977). The network state then drives an investigation based on the selection of tests suggested by the linked states. The final diagnosis is reflected by the classification of the paths found through the causal network.

The CASNET model handles causality somewhat superficially, representing causal processes as linear associations between states. Furthermore, the topological structure of the network is built by the system designer before the analysis begins, and modeling the physiological process involves little more than weight propagation and node activation through this pre-assembled net. A more sophisticated use of causal representations for diagnosis appeared shortly after CASNET. ABEL (Patil et al. 1981) reasons about Acid Base and Electrolyte imbalances in patients. ABEL's architecture is based on the observation that clinicians consider a case at several levels of detail that are eventually integrated into a higher level categorical understanding of the disease process, affording a detailed interpretation of the data collected.

Given patient data, ABEL develops a patient-specific model consisting of an interpretation of the data in a hierarchical causal network. ABEL uses initial patient data to generate a patient specific model. It then uses the same procedures to suggest clinical measurements and patient specific model revision. Unlike CASNET, where the causal network is constructed when the program was designed, ABEL dynamically instantiates its general medical knowledge in response to the data. ABEL's network has three levels, reflecting the different levels of detail at which human diagnosticians typically reason. These levels, as can be seen in Figure 1, are the clinical, the pathophysiological, and an intermediate transfer level.

Construction of the causal network is accomplished through five operators: aggregation, elaboration, decomposition, summation, and projection. Aggregation summarizes the description of the causal network at a given level into the next higher aggregate level. There are two types of aggregation, focal aggregation, which summarizes a node and its immediate causally connected neighborhood by a single node, and causal aggregation, which summarizes a chain of cause-effect relations by a single cause and effect relation. Elaboration is the inverse of aggregation, serving to expand the causal relationships represented at a given level into a more detailed set of

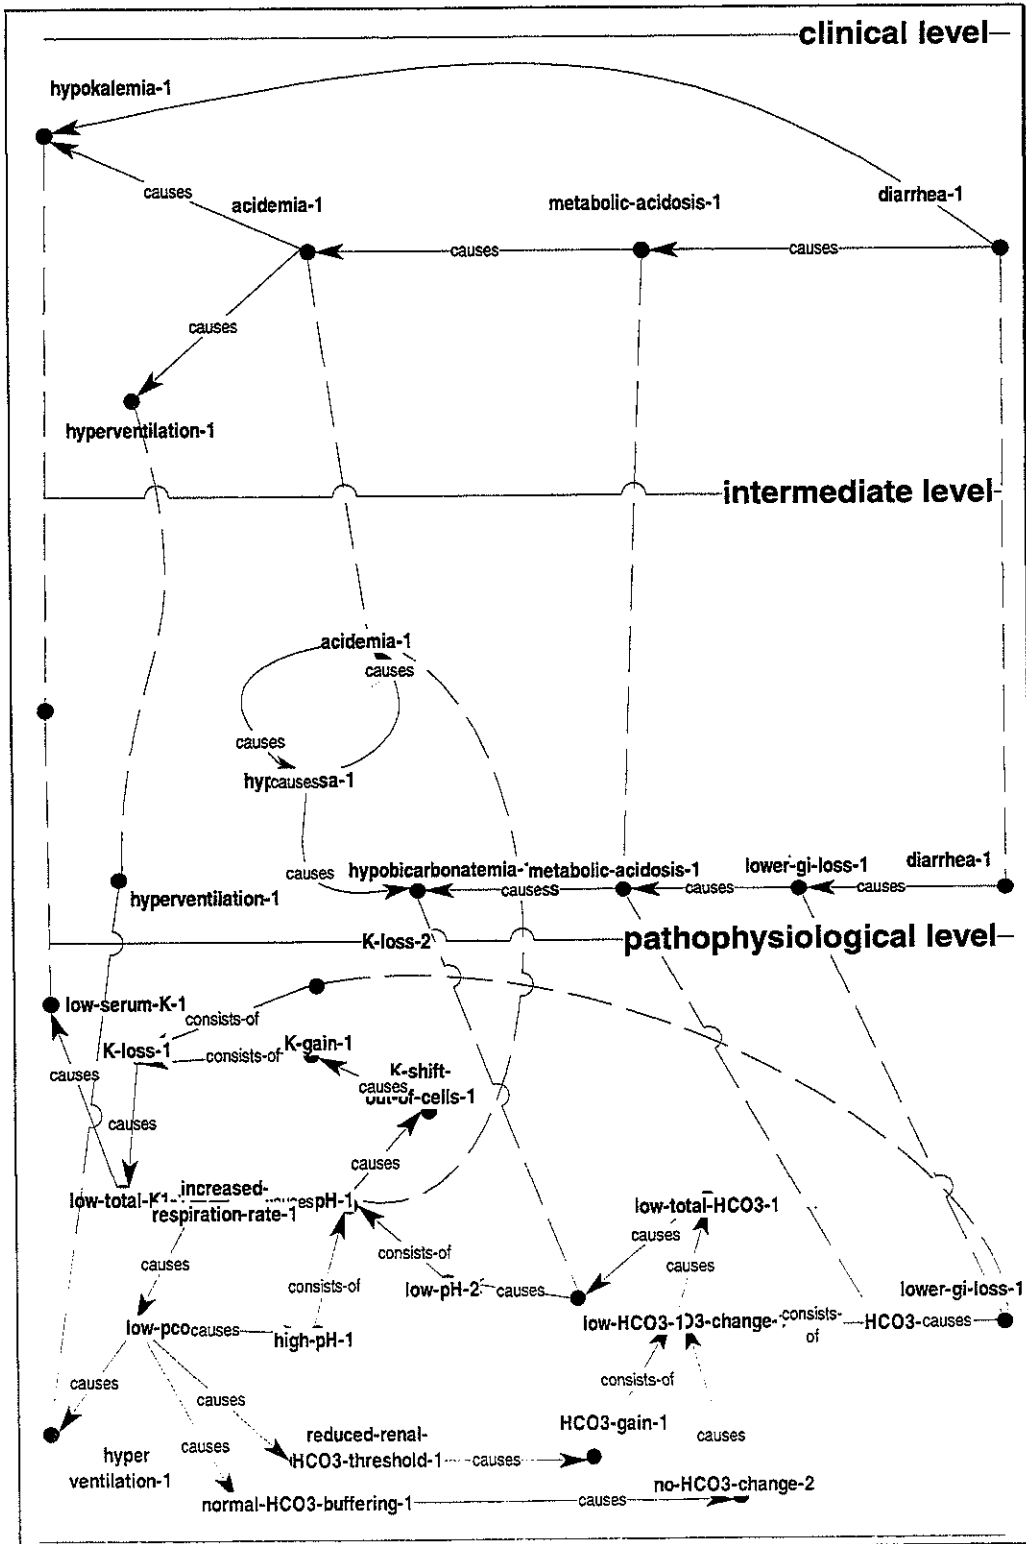


Figure 1. A Causal Network in ABEL.

relations at a lower level. There are two types of elaboration, focal and causal, the duals of the aggregation operators just mentioned.

The decomposition and summation operators relate components at the same level of detail, constraining a causally connected region of the network and enforcing the consistency of the summation of quantities distributed over causal links. Projection, perhaps the most interesting operator, is similar to elaboration in that it serves to expand a region of the network. It is essentially an abductive operator in that it is used to extrapolate the hypothetical causal relationships needed to account for otherwise unexplainable states or quantities in the network. Projection can serve to generate expectations and motivates the collection of diagnostic data.

A causal link in ABEL is a mapping relation that takes attributes of cause–instance pairs into attributes of an effect–instance. Causal links themselves have contextual attributes which, when they differ from default information, can induce functional changes in the mapping relation. This mapping relation supports the numerical summation and decomposition operators, which allow the system to reason about quantitative information such as electrolyte levels and pH. ABEL represented the state-of-the-art in clinical reasoning for its time and still remains unsurpassed in its hierarchical integration of causal reasoning across multiple levels of detail.

The causal network addresses the problem of abduction by explicitly developing this hierarchy relating problem data to causal mechanisms. Experts in a problem domain develop these hierarchies through careful analysis of their own diagnostic skills. Then, in addressing a new situation for analysis, the data available instantiate the appropriate parts of this network. Problem solving happens as the system propagates these constraints through the network. This is, of course, a highly interactive system as partial solutions can recommend more data acquisition in the context of attempting to establish a particular explanation of the problem.

Another approach to causal reasoning in well-understood contexts, scripts, is described in Luger and Stubblefield (1998). In scripts there are no certainty factors, simply a set of data structures that are intended to represent reasoning within well-understood situations, such as going to a restaurant or attending a child's birthday party.

4 Conclusions

We have discussed two relatively formal methods for accomplishing the abductive task. These are the set cover approach and clause-based abductive chaining. There are a number of further approaches available based on extensions of first order logic. These extensions include truth maintenance systems, circumscriptive logics, and other forms of nonmonotonic reasoning.

A second general class of approaches is based on network representations. In network approaches, sound inference rules are replaced with carefully crafted reasoning structures that capture the the domain expert's knowledge of causal relationships. Other heuristic approaches include the work of Stern in the analysis of failures of discrete component semiconductors (Stern 1996; Stern and Luger 1992, 1996). Finally, Bayesian belief networks support a well founded probabilistic approach to diagnostic problem solving.

Heuristic approaches to uncertain reasoning, i.e., those used in expert systems, should also not be overlooked. These involve straightforward algebraic techniques to capture and propagate uncertain information through imprecise rules. Especially important here are the Stanford Certainty Factor Algebra, the Dempster-Shafer calculus, and in more limited contexts, fuzzy systems. Examples and further references for all the approaches mentioned in this concluding section may be found in Luger and Stubblefield (1998).

5. References

- Charniak, E. and Shimony, S. 1990. Probabilistic Semantics for Cost Based Abduction. Proceedings of the Eighth National Conference on Artificial Intelligence, pp. 106–111. Menlo Park, CA: AAAI Press/ MIT Press.
- Clancy, W. J. and Shortliffe, E. H., eds. 1984. Readings in Medical Artificial Intelligence: the First Decade. Reading, MA: Addison Wesley.
- deKleer, J. 1986. An assumption based truth maintenance system, *Artificial Intelligence*, 28.
- Grice, H. P. 1975. Logic and conversation. In Cole and Morgan (1975).
- Levesque, H. J. 1989. A knowledge level account of abduction. Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, pp. 1051–1067. San Mateo, CA: Morgan Kaufman.

- Luger, G.F. and Stubblefield, W.A. 1998. *Artificial Intelligence: Structures and strategies for complex problem solving*, London: Addison-Wesley.
- Ng, H. T. and Mooney, R. J. 1990. On the Role of Coherence in Abductive Explanation. *Proceedings of the Eighth National Conference on Artificial Intelligence*, pp. 347–342. Menlo Park, CA: AAAI Press/ MIT Press.
- Patil, R., Szolovits, P., and Schwartz, W. 1981. Causal understanding of patient illness in medical diagnosis. *Proceedings of the International Joint Conference on Artificial Intelligence*: Palo Alto: Morgan Kaufmann.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Los Altos CA: Morgan Kaufmann.
- Reggia, J., Nau, D. S., and Wang, P. Y. 1983. Diagnostic expert systems based on a set covering model. *International Journal of Man–Machine Studies*, 19(5):437–460.
- Selman, B. and Levesque, H. J. 1990. Abductive and Default Reasoning: A Computational Core. *Proceedings of the Eighth National Conference on Artificial Intelligence*, pp. 343–348. Menlo Park, CA: AAAI Press/MIT Press.
- Stern, C. R. 1996. *An Architecture for Diagnosis Employing Schema-based Abduction*. PhD dissertation, Department of Computer Science, University of New Mexico.
- Stern, C. R. and Luger, G. F. 1992. A Model for Abductive Problem Solving Based on Explanation Templates and Lazy Evaluation. *International Journal of Expert Systems*, 5(3):249–265.
- Stern, C. R. and Luger, G. F. 1997. *Abduction and Abstraction in Diagnosis: A Schema-Based Account*. *Situated Cognition: Expertise in Context*. Ford et. al. ed. Cambridge, MA: MIT Press.
- Weiss, S. M., Kulikowski, C. A., Amarel, S., and Safir, A. 1977. A model-based method for computer-aided medical decision-making. *Artificial Intelligence*, 11(1–2):145–172.